

Special Issue Reprint

Recent Developments in Machine Design, Automation and Robotics, Second Edition

Edited by
Raul D. S. G. Campilho

mdpi.com/journal/machines

**Recent Developments in Machine
Design, Automation and Robotics,
Second Edition**

Recent Developments in Machine Design, Automation and Robotics, Second Edition

Guest Editor

Raul D. S. G. Campilho



Basel • Beijing • Wuhan • Barcelona • Belgrade • Novi Sad • Cluj • Manchester

Guest Editor

Raul D. S. G. Campilho
Department of Mechanical
Engineering
Polytechnic of Porto
Porto
Portugal

Editorial Office

MDPI AG
Grosspeteranlage 5
4052 Basel, Switzerland

This is a reprint of the Special Issue, published open access by the journal *Machines* (ISSN 2075-1702), freely accessible at: https://www.mdpi.com/journal/machines/special_issues/FA70DKQMSB.

For citation purposes, cite each article independently as indicated on the article page online and as indicated below:

Lastname, A.A.; Lastname, B.B. Article Title. <i>Journal Name</i> Year , Volume Number, Page Range.
--

ISBN 978-3-7258-8325-7 (Hbk)

ISBN 978-3-7258-8326-4 (PDF)

<https://doi.org/10.3390/books978-3-7258-8326-4>

© 2026 by the authors. Articles in this reprint are Open Access and distributed under the Creative Commons Attribution (CC BY) license. The reprint as a whole is distributed by MDPI under the terms and conditions of the Creative Commons Attribution-NonCommercial-NoDerivs (CC BY-NC-ND) license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>).

Contents

About the Editor	ix
Preface	xi
Raul D. S. G. Campilho Recent Developments in Machine Design, Automation and Robotics, Second Edition Reprinted from: <i>Machines</i> 2025 , <i>14</i> , 596, https://doi.org/10.3390/machines14060596	
	1
Yasser M. Alginahi, Omar Sabri and Wael Said Reinforcement Learning for Industrial Automation: A Comprehensive Review of Adaptive Control and Decision-Making in Smart Factories Reprinted from: <i>Machines</i> 2025 , <i>13</i> , 1140, https://doi.org/10.3390/machines13121140	
	7
Pierpaolo Dini, Damiano Nardi and Sergio Saponara Overview on Predictive Maintenance Techniques for Turbomachinery Reprinted from: <i>Machines</i> 2026 , <i>14</i> , 396, https://doi.org/10.3390/machines14040396	
	43
Syed Hassan Imam, Saqib Jamshed Rind, Saba Javed and Mohsin Jamil New Trends and Challenges in Electric and Hybrid Electric Vehicles: Powertrain Configurations, Traction Motors and Drive Control Techniques Reprinted from: <i>Machines</i> 2026 , <i>14</i> , 489, https://doi.org/10.3390/machines14050489	
	89
Johannes Muru and Anton Rassölkin A Scoping Review of Energy Consumption in Industrial Robotics Reprinted from: <i>Machines</i> 2025 , <i>13</i> , 542, https://doi.org/10.3390/machines13070542	
	127
Pedro Garcia Batista, André Costa Vieira and Pedro Dinis Gaspar Robotic Prostheses and Neuromuscular Interfaces: A Review of Design and Technological Trends Reprinted from: <i>Machines</i> 2025 , <i>13</i> , 804, https://doi.org/10.3390/machines13090804	
	157
Turac I. Ozcelik, Enrico Masi, Seyyed Masoud Kargar, Chiara Scagliarini, Adyan Fatima, Rocco Vertechy and Giovanni Berselli Recent Developments in Electrodehesion Grippers for Automated Fruit Grasping Reprinted from: <i>Machines</i> 2025 , <i>13</i> , 1128, https://doi.org/10.3390/machines13121128	
	191
Filippo Crescenzo and Nicola Bianchi Design, Analysis, and Prototyping of a Formula SAE Application Flux-Switching Permanent Magnet Motor Reprinted from: <i>Machines</i> 2026 , <i>14</i> , 332, https://doi.org/10.3390/machines14030332	
	210
Jinshi Du, Yu Zhang, Miguel Martínez García and Adrian Spencer A Novel Dynamic Surge Modeling Framework for Gas Turbines: Integration of Compressor Variable Geometry Reprinted from: <i>Machines</i> 2026 , <i>14</i> , 327, https://doi.org/10.3390/machines14030327	
	233
Eduardo Muñoz-Palomeque, Segundo Esteban and Matilde Santos Mechanical Design and Experimental Study of a Small-Scale Wind Turbine Model Reprinted from: <i>Machines</i> 2025 , <i>13</i> , 929, https://doi.org/10.3390/machines13100929	
	257
Shibin Yao, Xiaoyuan Li, Jian Zhou and Manoj Khandelwal Graded Evaluation and Optimal Scheme Selection of Mine Rock Diggability Based on the Multidimensional Cloud Model Reprinted from: <i>Machines</i> 2025 , <i>13</i> , 1019, https://doi.org/10.3390/machines13111019	
	278

Despoina Almpiani and George-Christopher Vosniakos A Study of Tau-Robot Configuration for Friction Stir Welding Reprinted from: <i>Machines</i> 2026 , <i>14</i> , 289, https://doi.org/10.3390/machines14030289	301
Mario Peñacoba-Yagüe and Jesús Enrique Sierra-García Generalized Design Methodology for Dual-Arm Robotic Platforms: From Conceptualization to Experimental Validation Within the MANiBOT Framework Reprinted from: <i>Machines</i> 2026 , <i>14</i> , 74, https://doi.org/10.3390/machines14010074	321
Ameur Gargouri, Mohamed Karray, Bechir Zalila and Mohamed Ksantini Design and Development of an Autonomous Mobile Robot for Unstructured Indoor Environments Reprinted from: <i>Machines</i> 2025 , <i>13</i> , 1044, https://doi.org/10.3390/machines13111044	369
Hongjoo Jin, Chanhyuk Moon, Taegyun Kim and Taewon Seo Belt Sanding Robot for Large Convex Surfaces Featuring SEA Arms and an Active Re-Tensioner with PI Force Control Reprinted from: <i>Machines</i> 2025 , <i>13</i> , 1012, https://doi.org/10.3390/machines13111012	396
Feifei Qin, Zexuan Liu, Yuanjie Xian, Binrui Wang, Qiaoye Zhang and Ye-Hwa Chen Design and Experimental Validation of an Adaptive Robust Control Algorithm for a PAM-Driven Biomimetic Leg Joint System Reprinted from: <i>Machines</i> 2026 , <i>14</i> , 84, https://doi.org/10.3390/machines14010084	414
André Moreira, Marco Pinto, Miguel Fernandes, João Costa, Jorge Fidalgo and Alessandro Fantoni Myoelectric Controlled Bionic Robotic Hand for Voluntary Finger Motion Driven by Neuromuscular Intent Reprinted from: <i>Machines</i> 2026 , <i>14</i> , 355, https://doi.org/10.3390/machines14030355	437
Connor Ruybalid, Christian Salisbury and Duke M. Bulanon Robotic Harvesting of Apples Using ROS2 Reprinted from: <i>Machines</i> 2026 , <i>14</i> , 433, https://doi.org/10.3390/machines14040433	456
Eva M. Rubio, Marian Sáenz-Nuño, Marta M. Marín and David Gómez Design and Industrial Integration of Automated Coordinate Measuring Machines for Automotive Production Reprinted from: <i>Machines</i> 2026 , <i>14</i> , 449, https://doi.org/10.3390/machines14040449	469
Mauricio Arturo Moreno-Gerena, Luis Manuel Navas-Gracia and Juan Gonzalo Ardila-Marín Design, Calibration, and Troubleshooting of a Modular Low-Cost 3D Printer Based on Open-Source Technologies Reprinted from: <i>Machines</i> 2026 , <i>14</i> , 261, https://doi.org/10.3390/machines14030261	493
Ray Tang, Zhijie Lyu and Anurag Purwar Kinematic Synthesis of Planar Leg Mechanisms Through Large-Scale Dataset Generation, Geometric Filtering, and Optimization Reprinted from: <i>Machines</i> 2026 , <i>14</i> , 253, https://doi.org/10.3390/machines14030253	519
Tibor Horak, Roman Ruzarovsky, Roman Zelník, Martin Csekei and Ján Šido Real-World Experimental Evaluation of DDoS and DRDoS Attacks on Industrial IoT Communication in an Automated Cyber-Physical Production Line Reprinted from: <i>Machines</i> 2026 , <i>14</i> , 258, https://doi.org/10.3390/machines14030258	543
Claus Pahl, Hamid R. Barzegar and Nabil El Ioini Quality Management for AI-Generated Self-Adaptive Resource Controllers Reprinted from: <i>Machines</i> 2026 , <i>14</i> , 25, https://doi.org/10.3390/machines14010025	581

Umar Farid, Bilal Khan, Sahibzada Muhammad Ali and Zahid Ullah A Digital Twin Model for UAV Control to Lift Irregular-Shaped Payloads Using Robust Model Predictive Control Reprinted from: <i>Machines</i> 2025 , <i>13</i> , 1069, https://doi.org/10.3390/machines13111069	607
Hamed Nozari and Agnieszka Szmelter-Jarosz A Predictive Maintenance Approach for Composting Plants Based on ERP and Digital Twin Integration Reprinted from: <i>Machines</i> 2025 , <i>13</i> , 1123, https://doi.org/10.3390/machines13121123	635
Ioannis Mavrelis, Avraam Chatzopoulos and Paraskevi Zacharia Hybrid Deep Learning for Predictive Maintenance in Industrial Machinery Using LSTM and MLP Models Reprinted from: <i>Machines</i> 2026 , <i>14</i> , 191, https://doi.org/10.3390/machines14020191	659
Faisal Saleem, Muhammad Umar and Jong-Myon Kim An Optimized Few-Shot Learning Framework for Fault Diagnosis in Milling Machines Reprinted from: <i>Machines</i> 2025 , <i>13</i> , 1010, https://doi.org/10.3390/machines13111010	682
Minseop Shin, Junyoung Seo, In-Bae Lee and Sojung Kim Defective Photovoltaic Module Detection Using EfficientNet-B0 in the Machine Vision Environment Reprinted from: <i>Machines</i> 2026 , <i>14</i> , 232, https://doi.org/10.3390/machines14020232	708
Diego Rigolli, Lorenzo Pompei, Massimo Manfredini, Massimiliano Vignoli, Vincenzo La Battaglia and Alessandro Giorgetti The Development of a Predictive Maintenance System for Gearboxes Through a Statistical Diagnostic Analysis of Lubricating Oil and Artificial Intelligence Reprinted from: <i>Machines</i> 2025 , <i>13</i> , 693, https://doi.org/10.3390/machines13080693	727

About the Editor

Raul D. S. G. Campilho

Raul D. S. G. Campilho is a professor at CIDEM, ISEP-School of Engineering, Polytechnic of Porto, and a collaborator at INEGI-Institute of Science and Innovation in Mechanical and Industrial Engineering, Portugal. His main research interests include advanced joining technologies, composite materials, manufacturing systems, and structural design for aerospace and mechanical applications. He has coordinated and participated in numerous national and international research projects focusing on structural performance, hybrid bonding techniques, and digital manufacturing. Raul has published extensively in peer-reviewed journals and is involved in the scientific community as editor in chief, associate editor and reviewer for scientific journals. He is also a supervisor of MSc and PhD students and regularly collaborates with industry on applied research and innovation initiatives.

Preface

This Reprint brings together a collection of high-quality contributions that reflect the current state of research and development in machine design, industrial automation, and robotics. Its scope includes fundamental methodologies, advanced mechanical systems, intelligent control approaches, and emerging digital technologies, with a strong focus on their integration within modern engineering practice.

The motivation behind this Reprint is to provide a coherent and comprehensive perspective on the rapid evolution of automated and robotic systems, driven by the increasing demand for flexibility, efficiency, and reliability in industrial environments. The selected contributions illustrate how classical mechanical design is being enhanced by artificial intelligence, data driven techniques, and cyber physical integration, to enable higher performance and autonomy.

This Reprint is addressed to researchers, engineers and practitioners working in the fields of mechanical engineering, automation, and robotics, as well as to those interested in the adoption of advanced technologies in industrial contexts. It aims to support knowledge dissemination, encourage interdisciplinary collaboration and highlight the challenges and opportunities that define the future of intelligent engineering systems.

Raul D. S. G. Campilho

Guest Editor

Editorial

Recent Developments in Machine Design, Automation and Robotics, Second Edition

Raul D. S. G. Campilho ^{1,2}

¹ CIDEM, ISEP, Polytechnic of Porto, R. Dr. António Bernardino de Almeida, 431, 4200-072 Porto, Portugal; raulcampilho@gmail.com; Tel.: +351-939-526-892

² INEGI, Pólo FEUP, Rua Dr. Roberto Frias, 400, 4200-465 Porto, Portugal

Machine design, automation, and robotics are central pillars of modern manufacturing and engineering practice. In a globalized and competitive environment, industries are obliged to implement highly efficient, flexible, and intelligent production systems that can respond quickly to shifting market demands [1]. In this context, the emergence of smart manufacturing principles, often associated with Industry 4.0, has redefined how machines are conceived, integrated, and operated within production ecosystems [2]. The incorporation of real-time monitoring, predictive maintenance, and autonomous decision-making capabilities has significantly enhanced system reliability and productivity [3]. Simultaneously, robotics has evolved from rigid and pre-programmed systems to adaptive and collaborative platforms capable of interacting safely with humans and operating in unstructured environments [4]. Moreover, the increasing focus on sustainability and energy efficiency has introduced new constraints and opportunities in machine design, encouraging the development of optimized systems with reduced environmental impact [5].

Machine design is a fundamental discipline to ensure the performance, reliability, and adaptability of industrial systems. Modern machine design is no longer limited to conventional mechanical considerations but increasingly integrates multidisciplinary approaches by combining structural mechanics, materials engineering, control systems, and digital design tools [6]. In this regard, advanced computational methods and optimization techniques enable engineers to develop highly efficient and lightweight components that are able to meet specific operational requirements (e.g., [7]). Furthermore, additive manufacturing technologies have introduced flexibility in design, allowing for complex geometries and material distributions that were previously unattainable. Open-source and modular design approaches also contribute to accelerating innovation cycles and reducing development costs [8]. In parallel, the demand for precision and quality has triggered the development of advanced measurement and calibration systems, which ensure that machine performance meets stringent operational standards [9].

Automation has suffered a significant transformation in recent years by shifting from rigid, rule-based systems to adaptive and intelligent programs. This change is mainly possible by integrating advanced control strategies, data analytics, and machine learning techniques [10]. As a result, systems can respond dynamically to variations in operating conditions [11]. Modern automation systems rely on predictive and adaptive control algorithms, which enhance process stability and efficiency and reduce production stoppages [12]. An example of this trend is the incorporation of hybrid deep learning models for predictive maintenance, which make possible the early fault detection and improved decision-making in complex industrial environments [13]. On the other hand, the joint use of automation and cyber-physical systems has improved communication between

physical assets and digital platforms [14]. The implementation of digital twins further provides virtual replicas of systems for real-time simulation and control purposes [15].

Robotics has experienced major advancements by transitioning from structured and repetitive task execution to autonomous, intelligent, and context-aware systems [16,17]. This evolution is supported by breakthroughs in sensing technologies, artificial intelligence, and real-time decision-making algorithms, which enable robots to operate effectively in dynamic and unstructured environments. A significant development lies in integrating learning-based approaches, particularly reinforcement learning and vision-based systems, which allow robots to acquire skills through interaction and perception [18]. The design of robotic platforms has become increasingly sophisticated since these incorporate advanced kinematic synthesis, modular architectures, and dual-arm coordination strategies to enhance dexterity and task versatility [19]. Human-centered robotics has also gained prominence, with significant progress in prosthetics and bioinspired systems that interface directly with neuromuscular signals, allowing intuitive and precise control [20]. Additionally, emerging actuation and gripping technologies, such as electroadhesion-based systems, are expanding the range of objects and materials that robots can handle efficiently [21].

The industrial sectors where research efforts are most concentrated and where technology transfer has been most significant from recent literature are as follows:

- **Automotive and Mobility Industry:** Innovation driver due to its high production volumes and quality requirements. Automation is extensively used in assembly, inspection, and logistics, while robotics enables precision operations such as welding, painting, and component handling [22,23];
- **Aerospace Industry:** Characterized by low production rates but extremely high precision and safety standards. Advanced machine design and automation are used for composite manufacturing, inspection, and assembly. Robotics is applied in automated drilling, fastening, and friction stir welding [24];
- **Agriculture and Agri-Food Systems:** Automation and robotics applied to address labor shortages and improve productivity. Research focuses on autonomous vehicles, robotic harvesting, and precision farming technologies [25,26]. Vision-based systems and AI-driven decision-making enable robots to operate outdoors [18];
- **Energy and Power Generation:** Wind, solar, and conventional energy systems, where automation supports monitoring, maintenance, and optimization of complex infrastructures [27]. Robotics is used for inspection and maintenance in hazardous or hard-to-access environments [28];
- **Manufacturing and Smart Factories:** Advanced manufacturing environments represent a central domain for automation and robotics research, particularly within the framework of Industry 4.0 [29,30]. The focus is on cyber-physical systems, digital twins, and intelligent control architectures that enable real-time optimization and predictive maintenance, with the integration of AI-driven models for fault diagnosis [31].

The main research lines that currently drive innovation in machine design, automation, and robotics are collected in different topics:

- **Intelligent Control and Adaptive Systems:** Development of advanced control architectures such as model predictive control, robust control strategies, and AI-enhanced controllers that integrate learning capabilities [32,33]. Recent contributions demonstrate how adaptive and robust control algorithms can significantly improve system performance and stability in complex mechatronic systems [34];
- **Predictive Maintenance and Condition Monitoring:** Techniques combining sensor data, digital twins, and hybrid AI models enable more accurate diagnostics and reduced downtime, as illustrated by recent developments in intelligent maintenance systems for industrial machinery [35];

- Additive Manufacturing and Advanced Machine Design: Research explores how additive manufacturing can be integrated into machine design to enable lightweight structures. Open-source platforms and modular architectures further enhance flexibility and accessibility, supporting low-cost equipment development [8];
- Cyber-Physical Systems and Digital Twin Integration: These technologies enable real-time monitoring, simulation, and optimization of industrial processes, supporting improved decision-making and lifecycle management. Their application in complex systems highlights their potential to improve operational efficiency [15];
- Artificial Intelligence and Machine Learning: AI-driven methodologies are increasingly applied in automation and robotics [36,37]. Research includes deep learning for fault detection, reinforcement learning for decision-making, and intelligent vision systems for quality inspection [38,39].

A discussion on recent advances in machine design, automation, and robotics also requires acknowledging the limitations that constrain their broader implementation:

- Limited Generalization and Robustness of AI Models: Many machine learning models, particularly deep learning approaches, struggle to generalize beyond the specific datasets on which they are trained. Variations in operating conditions, noise, and unforeseen scenarios can lead to performance degradation [40];
- Integration Complexity in Cyber-Physical Systems: Due to the heterogeneity of hardware, software, and communication protocols. Achieving perfect interoperability between legacy equipment and digital infrastructures requires customization, which increases both cost and deployment time [41];
- High Initial Investment and Scalability Constraints: Advanced automation and robotic systems typically involve substantial upfront costs, including equipment, integration, and workforce training [42,43]. Scaling pilot solutions to full industrial operation can be complex [44];
- Safety and Certification Challenges: As systems become more autonomous and collaborative, ensuring safety and compliance with regulatory standards becomes increasingly difficult [45,46]. Certification processes for AI-driven or adaptive systems are still evolving, which creates bottlenecks in industrial implementation [47];
- Energy Consumption and Sustainability Trade-Offs: Automation may lead to increased energy consumption due to high-performance computing, continuous sensing, and complex robotic systems [48]. Thus, there is a need for energy-aware design and optimization strategies in automated systems [49].

A forward-looking perspective on machine design, automation, and robotics reveals several promising research opportunities to shape industrial systems:

- Human-Centered and Collaborative Systems: Future research will increasingly focus on human-machine collaboration, toward true synergy between operators and automated systems: intuitive interfaces, wearable sensing technologies, and adaptive robotic behavior that responds to human intent in real time [50,51];
- Autonomous and Self-Optimizing Industrial Systems: The development of fully autonomous systems capable of self-configuration, self-diagnosis, and self-optimization represents a major opportunity [52]; to adapt continuously to changing environments and operational demands;
- Advanced Materials and Multifunctional Machine Design: Emerging materials, including smart materials, metamaterials, and self-healing composites, open new possibilities for machine design [53]. These materials enable structures with embedded sensing, actuation, and adaptive properties, leading to machines that are lighter, more efficient, and capable of responding dynamically to external stimuli [54];

- Sustainable and Energy-Aware Engineering Systems: Emphasis on energy-efficient machine design, circular economy principles, and low-impact manufacturing processes [55]. Optimization techniques that simultaneously consider performance, cost, and environmental impact in automated production systems [56];
- Digitalization, Data Spaces, and Interoperability Standards: The expansion of digital infrastructures, including digital twins, industrial data spaces, and standardized communication protocols, presents significant opportunities to improve system interoperability and reproducibility [57]. Creation of unified frameworks that optimize data exchange across platforms and organizations [58].

In summary, the contributions gathered in this Special Issue reflect the current research in machine design, automation, and robotics, highlighting technological maturity and emerging innovation topics. The presented works demonstrate how advanced design methodologies, intelligent control, and data-driven strategies are improving modern industrial systems. At the same time, these reveal existing challenges and motivate continued research toward more robust, adaptive, and sustainable solutions. It is expected that this Special Issue will stimulate further developments and promote interdisciplinary collaboration in both academia and industry.

Funding: This research received no external funding.

Conflicts of Interest: The author declares no conflict of interest.

References

1. Kostal, P.; Mudrikova, A.; Michal, D. Possibilities of intelligent flexible manufacturing systems. *IOP Conf. Ser. Mater. Sci. Eng.* **2019**, *659*, 012035. [CrossRef]
2. Mourtzis, D. Industry 4.0 and smart manufacturing. In *Manufacturing from Industry 4.0 to Industry 5.0*; Elsevier: Amsterdam, The Netherlands, 2024; pp. 13–61.
3. Nozari, H.; Szmelter-Jarosz, A. A predictive maintenance approach for composting plants based on ERP and digital twin integration. *Machines* **2025**, *13*, 1123. [CrossRef]
4. Gargouri, A.; Karray, M.; Zalila, B.; Ksantini, M. Design and development of an autonomous mobile robot for unstructured indoor environments. *Machines* **2025**, *13*, 1044. [CrossRef]
5. Muru, J.; Rassölkin, A. A scoping review of energy consumption in industrial robotics. *Machines* **2025**, *13*, 542. [CrossRef]
6. Fossdal, F.; Heldal, R.; Peek, N. Interactive digital fabrication machine control directly within a CAD environment. In Proceedings of the 6th Annual ACM Symposium on Computational Fabrication, New York, NY, USA, 28–29 October 2021; pp. 1–15.
7. Crescenzo, F.; Bianchi, N. Design, Analysis, and Prototyping of a Formula SAE Application Flux-Switching Permanent Magnet Motor. *Machines* **2026**, *14*, 332. [CrossRef]
8. Moreno-Gerena, M.A.; Navas-Gracia, L.M.; Ardila-Marín, J.G. Design, Calibration, and Troubleshooting of a Modular Low-Cost 3D Printer Based on Open-Source Technologies. *Machines* **2026**, *14*, 261. [CrossRef]
9. Rubio, E.M.; Sáenz-Nuño, M.; Marín, M.M.; Gómez, D. Design and Industrial Integration of Automated Coordinate Measuring Machines for Automotive Production. *Machines* **2026**, *14*, 449. [CrossRef]
10. Ricci, L. Control system design for industrial automation in the fourth industrial revolution. *Am. J. Instrum. Control Eng.* **2025**, *6*, 19–29. [CrossRef]
11. Baek, S. System integration for predictive process adjustment and cloud computing-based real-time condition monitoring of vibration sensor signals in automated storage and retrieval systems. *Int. J. Adv. Manuf. Technol.* **2021**, *113*, 955–966. [CrossRef]
12. Fernández, J.; Sánchez, M. Development of an adaptive control algorithm for autonomous robotic systems in manufacturing. *Int. J. Contemp. Res. Lit. Work* **2021**, *2*, 50–58.
13. Mavrelis, I.; Chatzopoulos, A.; Zacharia, P. Hybrid Deep Learning for Predictive Maintenance in Industrial Machinery Using LSTM and MLP Models. *Machines* **2026**, *14*, 191. [CrossRef]
14. Singh, K.K.; Nayyar, A.; Tanwar, S.; Abouhawwash, M. *Emergence of Cyber Physical System and IoT in Smart Automation and Robotics*; Springer: Berlin/Heidelberg, Germany, 2021.
15. Farid, U.; Khan, B.; Ali, S.M.; Ullah, Z. A digital twin model for UAV control to lift irregular-shaped payloads using robust model predictive control. *Machines* **2025**, *13*, 1069. [CrossRef]

16. Galeas, J.; Tudela, A.; Ballesteros, J.; Ayala, I.; Romero-Garcés, A.; Bandera, J.P. Context aware cognitive architecture for adaptive socially assistive robots. In Proceedings of the International Conference on Ubiquitous Computing and Ambient Intelligence, Prato/Florence, Italy, 26–28 November 2025; pp. 220–232.
17. Anjum, K.; Sadhu, V.; Pompili, D. Contextbots: Real-time context-aware inference on aerial robots. In Proceedings of the 2024 20th International Conference on Distributed Computing in Smart Systems and the Internet of Things (DCOSS-IoT), Abu Dhabi, United Arab Emirates, 29 April–1 May 2024; pp. 219–223.
18. Ruybalid, C.; Salisbury, C.; Bulanon, D.M. Robotic Harvesting of Apples Using ROS2. *Machines* **2026**, *14*, 433. [CrossRef]
19. Peñacoba-Yagüe, M.; Sierra-García, J.E. Generalized Design Methodology for Dual-Arm Robotic Platforms: From Conceptualization to Experimental Validation Within the MANiBOT Framework. *Machines* **2026**, *14*, 74. [CrossRef]
20. Moreira, A.; Pinto, M.; Fernandes, M.; Costa, J.; Fidalgo, J.; Fantoni, A. Myoelectric Controlled Bionic Robotic Hand for Voluntary Finger Motion Driven by Neuromuscular Intent. *Machines* **2026**, *14*, 355. [CrossRef]
21. Ozelik, T.I.; Masi, E.; Kargar, S.M.; Scagliarini, C.; Fatima, A.; Vertechy, R.; Berselli, G. Recent developments in electroadhesion grippers for automated fruit grasping. *Machines* **2025**, *13*, 1128. [CrossRef]
22. Babel, W. Automation of the automotive industry in fast motion: Robotics. In *Industry 4.0, China 2025, IoT: The Hype Around the World of Automation*; Springer: Berlin/Heidelberg, Germany, 2022; pp. 341–353.
23. Winkelhake, U. Vision of a Digitized Automotive Industry 2030. In *The Digital Transformation of the Automotive Industry: Catalysts, Roadmap, Practice*; Springer: Berlin/Heidelberg, Germany, 2025; pp. 95–167.
24. Almpani, D.; Vosniakos, G.-C. A Study of Tau-Robot Configuration for Friction Stir Welding. *Machines* **2026**, *14*, 289. [CrossRef]
25. Shamshiri, R.R. Electrical tractors for autonomous farming. In *Mobile Robots for Digital Farming*; CRC Press: Boca Raton, FL, USA, 2024; pp. 89–106.
26. Bhise, S.; Deshmukh, S.S.; Chougule, A.; Dambalkar, S.; Desai, H. AgriBot: A hexapod-based autonomous robot for precision agriculture and sustainable farming. In Proceedings of the 2025 2nd International Conference on Integration of Computational Intelligent System (ICICIS), Pune, India, 9–12 September 2025; pp. 1–5.
27. Yang, X.; Zou, T.; Xu, Y.; Chen, F.; Luo, H.; Li, R. Research dynamics and applications of tracking technology and devices in solar energy utilization system. *Sustain. Energy Technol. Assess.* **2025**, *75*, 104256. [CrossRef]
28. Muñoz-Palomeque, E.; Esteban, S.; Santos, M. Mechanical design and experimental study of a small-scale wind turbine model. *Machines* **2025**, *13*, 929. [CrossRef]
29. Adams, G. The impact of advanced manufacturing technologies on job market evolution. *Am. J. Ind. Prod. Eng.* **2024**, *5*, 6–10. [CrossRef]
30. Asif, S.; Bueno, M.; Ferreira, P.; Anandan, P.; Zhang, Z.; Yao, Y.; Ragunathan, G.; Tinkler, L.; Sotoodeh-Bahraini, M.; Lohse, N. Rapid and automated configuration of robot manufacturing cells. *Robot. Comput. Integr. Manuf.* **2025**, *92*, 102862. [CrossRef]
31. Saleem, F.; Umar, M.; Kim, J.-M. An optimized few-shot learning framework for fault diagnosis in milling machines. *Machines* **2025**, *13*, 1010. [CrossRef]
32. Nayak, R.R. AI-Enhanced Edge Computing for Intelligent Robotics and IoT Systems. *J. Comput. Anal. Appl.* **2026**, *35*, 449.
33. Deka, B.; Dash, S. Advanced control methods for biomedical systems: Applications in robotics and medical diagnosis. In *Medical Robotics and Intelligent Healthcare Technologies*; Elsevier: Amsterdam, The Netherlands, 2026; pp. 47–77.
34. Qin, F.; Liu, Z.; Xian, Y.; Wang, B.; Zhang, Q.; Chen, Y.-H. Design and Experimental Validation of an Adaptive Robust Control Algorithm for a PAM-Driven Biomimetic Leg Joint System. *Machines* **2026**, *14*, 84. [CrossRef]
35. Rigolli, D.; Pompei, L.; Manfredini, M.; Vignoli, M.; La Battaglia, V.; Giorgetti, A. The development of a predictive maintenance system for gearboxes through a statistical diagnostic analysis of lubricating oil and artificial intelligence. *Machines* **2025**, *13*, 693. [CrossRef]
36. Khan, A.; Memon, J. Artificial intelligence and machine learning in robotics. *Adv. Artif. Intell.* **2023**, 63–81.
37. Kilari, S.D. The role of artificial intelligence in the manufacturing of agricultural machinery. *IJSAT-Int. J. Sci. Technol.* **2025**, *16*, 1–4.
38. Wang, S.; Tao, J.; Jiang, Q.; Chen, W.; Liu, C. Manipulator joint fault localization for intelligent flexible manufacturing based on reinforcement learning and robot dynamics. *Robot. Comput. Integr. Manuf.* **2024**, *86*, 102684. [CrossRef]
39. Rufai, S.; Haq, I.U.; Shah, H.; Fayaz, M.A. Neural network approaches for intelligent decision-making in automation. In *Handbook of Intelligent Automation Systems Using Computer Vision and Artificial Intelligence*; Wiley-Scrivener: Austin, TX, USA, 2025; pp. 247–275.
40. Xie, A.; Lee, L.; Xiao, T.; Finn, C. Decomposing the generalization gap in imitation learning for visual robotic manipulation. In Proceedings of the 2024 IEEE International Conference on Robotics and Automation (ICRA), Yokohama, Japan, 13–17 May 2024; pp. 3153–3160.
41. Zampetti, F.; Nardone, V.; Di Penta, M. Problems and solutions in applying continuous integration and delivery to 20 open-source cyber-physical systems. In Proceedings of the 19th International Conference on Mining Software Repositories, Pittsburgh, PA, USA, 23–24 May 2022; pp. 646–657.

42. Rajapakshe, S.; Hettiarachchi, R. Design and development of a research oriented low cost robotics platform with a novel dynamic global path planning approach. In Proceedings of the 2022 8th International Conference on Control, Automation and Robotics (ICCAR), Xiamen, China, 8–10 April 2022; pp. 71–76.
43. Naik, S. System-Level Automation Strategies for Improving Positional Accuracy in Robotic Welding for High-Reliability Manufacturing. *Int. J. Sci. Res. Eng. Dev.* **2026**, *9*, 1952–1962.
44. Pandey, G.; Banarse, A.R.; Jayaram, V.; Ganeeb, K.K.; Gupta, P.; Krishnappa, M.S. Smart CRP using pega robotics: Enhancing customer relationship platforms with robotics process automation. In Proceedings of the 2024 International Conference on Intelligent Cybernetics Technology & Applications (ICICyTA), Ubud, Indonesia, 17–19 December 2024; pp. 18–23.
45. Chakraborty, S.; Zhu, Q.; Hu, J. Tutorial: Design automation for ML-enabled cyber-physical systems: From verification to synthesis. In Proceedings of the International Conference on Embedded Software, Taipei, Taiwan, 28 September–3 October 2025; pp. 1–2.
46. Kavallieratou, E. Towards ethical foundation models in robotics: Challenges and proposals. In Proceedings of the International Conference on Robot Ethics and Standards, Porto, Portugal, 3–4 July 2025; pp. 15–21.
47. Gudivaka, R.L. AI-driven optimization in robotic process automation: Implementing neural networks for real-time imperfection prediction. *Int. J. Bus. Gen. Manag. (IJBGM)* **2023**, *12*, 35–46.
48. Whig, P.; Kouser, S.; Bhatia, A.B.; Nadikattu, R.R.; Alkali, Y.J. IoT industrial application using wireless technologies. In *Wireless Communication Technologies*; CRC Press: Boca Raton, FL, USA, 2024; pp. 101–121.
49. Kumar, J.; Chaudhary, N.; Bhandwal, M.; Patil, P.; Babu, G.P.; Baig, Z.T. Optimization of energy-efficient wireless sensor networks for industrial process automation. In Proceedings of the 2025 14th International Conference on System Modeling & Advancement in Research Trends (SMART), Moradabad, India, 14–15 November 2025; pp. 166–173.
50. Cardoso, A.; Colim, A.; Bicho, E.; Braga, A.C.; Arezes, P. A novel human-centered methodology for assessing manual-to-collaborative safe conversion of workstations. *Saf. Sci.* **2025**, *181*, 106685. [CrossRef]
51. Gao, J.; Wu, K.; Li, W.; Glogowski, P.; Kuhlenkötter, B. Human-centered safe trajectory planning for human–robot collaborative assembly of electric vehicle battery copper busbar. *CIRP J. Manuf.* **2026**, *68*, 84–102. [CrossRef]
52. Radha, R.; Achary, R.; Paul, P.M. Intelligent autonomic computing: Achieving self-optimization through autonomous learning. In Proceedings of the 2025 3rd International Conference on Inventive Computing and Informatics (ICICI), Bengaluru, India, 4–6 June 2025; pp. 1463–1468.
53. Slaby, V.; Bajer, J.; Marcian, P.; Hrstka, M.; Hadas, Z. Novel metamaterial platform with piezoelectric sensors for self-sensing mechanical support. *Int. J. Mech. Sci.* **2026**, *314*, 111387. [CrossRef]
54. Wang, W.-J.; Yang, H.; Zhang, W.-M.; Shang, N.; Ma, L. Experimental study on the impact resistance of fill-enhanced mechanical metamaterials. *Int. J. Mech. Sci.* **2025**, *285*, 109799. [CrossRef]
55. Rengan, P.; Rajagopalan, A.; Kumar, R.; Nagarajan, K. Design and performance analysis of an energy-efficient single-phase capacitor run induction motor for washing machine applications. In *2024 IEEE 1st International Conference on Green Industrial Electronics and Sustainable Technologies (GIEST)*; IEEE: New York, NY, USA, 2024; pp. 1–6.
56. Yuan, G.; Liu, X.; Zhu, C.; Wang, C.; Zhu, M.; Sun, Y. Multi-objective coupling optimization of electrical cable intelligent production line driven by digital twin. *Robot. Comput. Integr. Manuf.* **2024**, *86*, 102682. [CrossRef]
57. Sreedharan, S.; Ramachandran, M.; Ramesh, D. Harnessing digital twins and industrial-IoT for cutting-edge mining automation: A methodological and technology assessment prototype. *Comput. Ind. Eng.* **2025**, *201*, 110871. [CrossRef]
58. Chaudhari, S.S.; Bhole, K.S.; Rane, S.B. Industrial automation and data processing techniques in IoT-based digital twin design for thermal equipment: A case study. *J. Inst. Eng. Ser. C* **2025**, *106*, 553–569. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Review

Reinforcement Learning for Industrial Automation: A Comprehensive Review of Adaptive Control and Decision-Making in Smart Factories

Yasser M. Alginahi ¹, Omar Sabri ^{2,*} and Wael Said ^{3,4}

¹ Department of Computer Science, Adrian College, Adrian, MI 49221, USA; yalginahi@adrian.edu

² Zekelman School of Business & IT, St. Clair College, Windsor, ON N9A 6S4, Canada

³ Computer Science Department, Faculty of Computers and Informatics, Zagazig University, Zagazig 44511, Egypt; wael.mohamed@zu.edu.eg

⁴ Computer Science Department, College of Computer Science and Engineering, Taibah University, Madinah 42353, Saudi Arabia

* Correspondence: osabri@stclaircollege.ca

Abstract: The accelerating integration of Artificial Intelligence (AI) in Industrial Automation has established Reinforcement Learning (RL) as a transformative paradigm for adaptive control, intelligent optimization, and autonomous decision-making in smart factories. Despite the growing literature, existing reviews often emphasize algorithmic performance or domain-specific applications, neglecting broader links between methodological evolution, technological maturity, and industrial readiness. To address this gap, this study presents a bibliometric review mapping the development of RL and Deep Reinforcement Learning (DRL) research in Industrial Automation and robotics. Following the PRISMA 2020 protocol to guide the data collection procedures and inclusion criteria, 672 peer-reviewed journal articles published between 2017 and 2026 were retrieved from Scopus, ensuring high-quality, interdisciplinary coverage. Quantitative bibliometric analyses were conducted in R using Bibliometrix and Biblioshiny, including co-authorship, co-citation, keyword co-occurrence, and thematic network analyses, to reveal collaboration patterns, influential works, and emerging research trends. Results indicate that 42% of studies employed DRL, 27% focused on Multi-Agent RL (MARL), and 31% relied on classical RL, with applications concentrated in robotic control (33%), process optimization (28%), and predictive maintenance (19%). However, only 22% of the studies reported real-world or pilot implementations, highlighting persistent challenges in scalability, safety validation, interpretability, and deployment readiness. By integrating a review with bibliometric mapping, this study provides a comprehensive taxonomy and a strategic roadmap linking theoretical RL research with practical industrial applications. This roadmap is structured across four critical dimensions: (1) Algorithmic Development (e.g., safe, explainable, and data-efficient RL), (2) Integration Technologies (e.g., digital twins and IoT), (3) Validation Maturity (from simulation to real-world pilots), and (4) Human-Centricity (addressing trust, collaboration, and workforce transition). These insights can guide researchers, engineers, and policymakers in developing scalable, safe, and human-centric RL solutions, prioritizing research directions, and informing the implementation of Industry 5.0-aligned intelligent automation systems emphasizing transparency, sustainability, and operational resilience.

Keywords: reinforcement learning; deep reinforcement learning; industrial automation; robotics; multi-agent systems

1. Introduction

Reinforcement Learning (RL), a subfield of Machine Learning (ML), has emerged as a transformative methodology for developing adaptive and autonomous decision-making

systems within complex industrial environments. Unlike traditional supervised or unsupervised learning techniques that depend on static, pre-labeled datasets, RL enables an agent to interact dynamically with its environment, learning optimal strategies through trial and error to maximize long-term rewards [1]. This unique capability allows RL systems to operate effectively under conditions of uncertainty, nonlinearity, and dynamic variability, i.e., features that define modern industrial systems [2]. In the era of Industry 4.0 and the transition toward Industry 5.0, RL has become a cornerstone technology in the development of intelligent manufacturing systems, offering the potential to enhance productivity, flexibility, and resilience through autonomous control and real-time optimization [3,4].

While Industry 4.0 established the infrastructure for data-driven automation, the emerging paradigm of Industry 5.0 complements this by emphasizing human-centricity, sustainability, and operational resilience [5,6]. This shift places new demands on RL systems, moving beyond pure performance optimization. For RL to be truly industrially ready, it must now also address interpretability to foster human trust, ensure safe collaboration in human–robot teams, and optimize for energy efficiency and long-term system resilience. Consequently, challenges such as scalability and safety are not merely technical bottlenecks but fundamental barriers to achieving this human–AI collaborative vision.

Industrial Automation is undergoing a profound transformation driven by the integration of cyber-physical systems (CPS), the Internet of Things (IoT), cloud computing, and AI. In this context, RL serves as a vital bridge between perception and control, enabling machines to make context-aware decisions, adapt to changing environments, and optimize operational efficiency. Smart factories, relying on distributed sensor networks and interconnected robotic systems, increasingly adopt RL-based frameworks to support adaptive process control, energy management, quality inspection, and predictive maintenance. The ability of RL to learn from continuous interaction makes RL particularly valuable in applications where explicit modeling is infeasible or cost-prohibitive due to complexity and stochastic system dynamics.

Classical RL algorithms, such as Q-Learning and SARSA (State Action Reward State Action refers to the sequence of events used to update the Q-value), represent the foundation of this paradigm and remain relevant for industrial tasks characterized by discrete state and action spaces. Their simplicity, interpretability, and low computational overhead make them well suited for process control, material transport, and rule-based decision systems [7,8]. However, classical approaches struggle with scalability in high-dimensional environments where the state-action space becomes exponentially large. To overcome these limitations, DRL integrates Neural Networks (NNs) with RL principles, enabling systems to approximate complex value functions and learn directly from raw, high-dimensional sensory data such as images, trajectories, or sensor streams [6,9,10].

Notable DRL algorithms, such as the Deep Q-Network (DQN), Proximal Policy Optimization (PPO), and Soft Actor–Critic (SAC), have demonstrated superior performance in complex control environments. These algorithms enable autonomous robots to execute tasks like high-speed pick-and-place operations, robotic assembly, and adaptive motion planning, often outperforming conventional Proportional–Integral–Derivative (PID) controller or rule-based controllers [11–14]. Their robustness, stability, and capacity for continuous learning make them suitable for process optimization in manufacturing lines, energy systems, and logistics. In particular, PPO and SAC have been widely applied for energy-efficient process control and dynamic resource allocation, where continuous action spaces and environmental uncertainty demand adaptability and robustness.

MARL, which constitutes 27% of the reviewed studies, extends these principles to collaborative industrial systems. While enabling decentralized control in production lines and warehouses, MARL introduces specific challenges distinct from single-agent settings. A primary issue is environment non-stationarity, where the environment becomes unpredictable from any single agent’s perspective as all other agents are simultaneously learning, complicating convergence. Secondly, the credit assignment problem, determining

each agent's contribution to a shared global reward, is non-trivial and critical for effective coordination. Finally, designing scalable communication architectures that prevent bottlenecks is essential for large-scale deployments [15–17]. Promising solutions emerging in the literature include Centralized Training with Decentralized Execution (CTDE) paradigms and attention mechanisms to manage complex multi-agent interactions.

Beyond these mainstream approaches, recent developments have introduced specialized frameworks such as hierarchical and modular RL, safe RL, and digital twin-integrated RL, which further expand the operational scope of RL in Industrial Automation [18,19]. Hierarchical RL decomposes complex manufacturing tasks into manageable subtasks, thereby improving sample efficiency and interpretability. Modular RL architectures facilitate knowledge reuse across tasks, accelerating learning in environments where similar subtasks recur. Safe RL incorporates formal safety constraints into the learning process, ensuring that exploratory actions do not cause instability or physical damage, an essential requirement for high-risk applications such as chemical processing or autonomous machining. Meanwhile, digital twin integration enables real-time simulation and virtual experimentation, allowing policies to be tested and refined in digital replicas of industrial systems before physical deployment.

Aligning with the goals of Industry 5.0, a growing body of research focuses on Human-Centric RL, which prioritizes the integration of the human operator into the learning and operational loop. A primary challenge is the 'black-box' nature of DRL policies, which undermines operator trust and compliance [20]. To address this, Explainable Reinforcement Learning (XRL) methods are being developed to provide post hoc explanations or create intrinsically interpretable models, making the agent's decision-making process transparent to human supervisors [17].

Furthermore, Human-in-the-Loop RL (HITL-RL) frameworks allow real-time human feedback to guide the learning process, improving sample efficiency and aligning outcomes with human intuition and ergonomic constraints [17]. This is particularly critical in applications like collaborative assembly, where the robot must adapt not only to the task but also to the human partner's behavior, pace, and physical comfort, thereby preventing strain and injury. These approaches mark a significant shift from RL systems that operate autonomously to those that operate collaboratively, fostering the trust and transparency required for widespread adoption in human-populated industrial environments.

Despite these promising advancements, significant barriers continue to hinder the widespread adoption of RL in industrial settings. Sample inefficiency remains one of the most critical challenges: many RL algorithms require millions of interactions to achieve convergence, which is impractical for physical industrial systems due to time, cost, and wear considerations. Safety and reliability are also paramount; as untested policies may cause accidents or equipment damage. Interpretability and transparency further limit adoption, as most DRL models act as 'black boxes,' offering limited insight into decision-making logic, a major concern for compliance and regulatory approval [8,21]. Moreover, the integration of RL into legacy infrastructures poses additional obstacles, as many factories operate with rigid control architectures and outdated communication protocols. Achieving seamless integration requires standardized interfaces, hybrid architectures, and interoperable software frameworks that can connect RL controllers to existing industrial systems.

While numerous studies have explored specific RL algorithms or applications within manufacturing and robotics, there remains a lack of comprehensive synthesis that connects algorithmic evolution with industrial readiness, integration technologies, and deployment challenges. Existing reviews often focus narrowly on algorithmic performance or domain-specific use cases without providing a holistic perspective that integrates methodological trends, industrial applicability, and future opportunities [5,9,22,23]. Accordingly, this study is guided by the following overarching research question:

“How has RL evolved as an enabling technology for Industrial Automation, and what strategies, technologies, and research trends are shaping its real-world deployment in smart factories?”

This question guides the analysis toward understanding both the theoretical and applied dimensions of RL in Industrial Automation, emphasizing methodological advancements, integration challenges, and emerging paradigms that bridge research and practice. Building upon this motivation, the objectives of this study are fourfold:

1. To review RL research in Industrial Automation, encompassing its algorithms, frameworks, and application domains.
2. To identify and categorize the major challenges hindering industrial adoption, including issues of safety, data efficiency, interpretability, and system integration.
3. To synthesize emerging strategies and frameworks such as digital twin-assisted RL, hierarchical and modular architectures, safe RL, and transfer learning techniques that address these limitations.
4. To highlight future research opportunities and industrial directions for RL within the broader context of Industry 4.0 and the transition toward Industry 5.0, focusing on human-centered, resilient, and sustainable automation systems [5,24,25].

The scope of this research encompasses peer-reviewed publications from major scientific databases that address RL and its applications within Industrial Automation contexts. These include process control, robotic manipulation, supply chain optimization, energy management, and predictive maintenance. The review excludes purely theoretical or non-industrial studies to maintain focus on practical relevance and deployment feasibility. Furthermore, the review examines both classical and advanced RL frameworks ranging from value-based and policy-gradient methods to hybrid, hierarchical, and multi-agent architectures. This broad scope ensures that the analysis captures not only algorithmic developments but also their integration with enabling technologies such as digital twins, IoT systems, and edge computing platforms. Taken together, this study contributes to the growing body of literature on intelligent industrial systems by offering an integrated perspective that connects RL theory, empirical application, and technological trend analysis. By emphasizing both methodological evolution and practical implications, it aims to inform future research and guide stakeholders—engineers, system designers, and policymakers—toward more efficient, interpretable, and human-centered AI-driven automation solutions. Ultimately, RL represents not only a computational paradigm but also a transformative force shaping the future of smart manufacturing, enabling industries to achieve higher levels of autonomy, adaptability, and operational excellence.

This evolution occurs alongside the strategic transition from Industry 4.0 toward Industry 5.0, which supplements a pure technology focus with core principles of human-centricity, sustainability, and operational resilience. This paradigm shift redefines the requirements for RL in industrial settings. Beyond raw performance, RL systems must now be interpretable to build operator trust, safe for seamless human–robot collaboration, and efficient in their use of energy and materials. Therefore, challenges such as scalability, safety, and interpretability are not merely technical hurdles but fundamental barriers to realizing the human–AI collaborative vision of future smart factories.

The remainder of this paper is structured as follows. The next section presents the literature review (Section 2), outlining the evolution of RL methodologies and their applications across diverse industrial domains. The methodology is explained in Section 3. Although the study is not a systematic review, the PRISMA framework is applied to structure the data collection and inclusion steps, ensuring clarity, transparency, and replicability. This section also details the modified Joanna Briggs Institute (JBI) methodological quality assessment, which was used to evaluate the rigor, transparency, and risk of bias of the included studies.

Next, Section 4 presents the results and discussion, focusing on key findings and emerging patterns derived from the analysis, with an emphasis on technological trends, industrial readiness, and strategic frameworks. Following this, Section 5 highlights future

research directions and potential opportunities for advancing RL adoption in Industrial Automation, concentrating on scalability, safety, interpretability, and human–machine collaboration. Finally, the paper concludes by summarizing the main findings, implications, and contributions, underscoring the importance of RL as a cornerstone of intelligent industrial systems.

2. Literature Review

RL has emerged as a transformative paradigm in Industrial Automation, enabling systems to learn optimal control and decision policies through interactions with dynamic and often stochastic environments. Unlike conventional control strategies that depend on pre-defined rules or mathematical models, RL facilitates adaptive and intelligent decision-making, allowing systems to respond effectively to uncertainty, variability, and evolving production conditions.

As highlighted by Eckardt et al. [26], RL’s capability to model complex, stochastic systems make it particularly suitable for modern smart factories. Likewise, Farooq and Iqbal [2] emphasize that RL promotes self-optimization and operational efficiency in manufacturing processes characterized by high complexity and interdependent tasks. In this context, RL supports the broader vision of Industry 4.0, where intelligent, data-driven autonomy forms the cornerstone of next-generation production systems.

The literature demonstrates that RL has been applied successfully across a wide range of industrial domains, including adaptive process control, robotics, and supply chain management. In adaptive process control, RL techniques have proven effective in managing nonlinear and uncertain systems such as press hardening and chemical production. Studies by Hou et al. [12]; Paranjape et al. [27]; and Rajasekhar et al. [14] show that both offline and model-based RL approaches can learn control policies from historical or simulated data, thereby improving process efficiency, product quality, and robustness. By continuously adjusting control parameters based on environmental feedback, RL achieves levels of precision and adaptability that traditional controllers, such as PID or Model Predictive Control (MPC) often cannot sustain under dynamic operating conditions.

In industrial robotics, RL has significantly advanced autonomous task learning and adaptation. Kober et al. [10] highlight RL’s ability to enable robots to perform complex operations such as assembly, material handling, and pick-and-place tasks. The integration of DRL architectures allows robots to process high-dimensional sensory inputs and develop sophisticated perception–action strategies. Transitioning from pre-programmed automation to RL-driven control has introduced more resilient and scalable robotic systems capable of error recovery and self-improvement without extensive human intervention, ultimately reducing programming effort and enhancing flexibility and efficiency [10].

A critical factor in deploying robotic systems is the economic and performance trade-off between different hardware configurations, which can be evaluated effectively through simulation prior to physical deployment. As demonstrated by Peta et al. [28] in a precision assembly case study, a direct comparison between single-arm and dual-arm collaborative robots revealed a significant performance trade-off: the dual-arm robot completed the assembly task 33% faster (14 s versus 21 s), while the single-arm variant was more energy-efficient. Crucially, an economic analysis showed that the higher initial cost of the dual-arm system was offset by its speed, achieving a return on investment within eight months. This study underscores the value of simulation tools for conducting multi-criteria evaluations, highlighting that optimal robot selection depends on a balance of speed, energy consumption, and long-term economic viability rather than a single performance metric.

Expanding to supply chain applications, RL has also been increasingly adopted in supply chain management and logistics. According to Estes et al. [29] and Rolf et al. [22], RL algorithms enhance inventory control, production scheduling, and logistics routing by dynamically adapting to fluctuations in real-time demand and supply. Unlike traditional optimization approaches that rely on static assumptions, RL introduces continuous learning

mechanisms that improve adaptability and responsiveness. This adaptability leads to better resource utilization, reduced operational costs, and higher service levels. Furthermore, RL's strength in handling multi-objective optimization problems makes it particularly relevant for smart factories, where trade-offs among efficiency, sustainability, and cost are common.

Beyond large-scale industrial and supply chain settings, RL has also shown significant promise in optimizing resource-constrained IoT networks, which form a critical data-gathering layer for smart factories. For instance, Khan et al. [30] proposed a Cognitive IoT-based Dynamic Power Management (CIoT-DPM) scheme to enhance energy efficiency in consumer electronics by mitigating energy waste from redundant data; their model achieved a 20–25% improvement in learning strategy and reduced computational complexity by 79.33%. Building on this focus on network efficiency, a subsequent study by Khan et al. [31] applied RL to cognitive IoT sensors through an Adaptive Scheduling Control Reinforcement Learning mechanism (ASC-RL). This complementary approach used a multi-parameter state-changing policy to dynamically manage sensor activity, which consequently boosted energy efficiency and reliability by 35% and increased the transmission success rate by 6.25%. Together, these studies underscore RL's significant potential in simultaneously addressing energy conservation and network performance challenges in the intelligent sensor networks that underpin industrial IoT systems.

2.1. Algorithmic Approaches and Comparative Analysis

To provide a consolidated overview, Table 1 summarizes key RL models, their industrial applications, advantages, and challenges as reported in the literature. Approaches such as Q-learning, DQN, PPO, and SAC have been widely explored, alongside specialized frameworks such as MARL, Offline RL, and Digital Twin-Integrated RL [1,2,8]. While each model offers unique benefits, ranging from scalability and robustness to safe exploration, common limitations include sample inefficiency, computational intensity, and the need for extensive tuning.

A cross-analysis of the studies in Table 1 reveals several key insights. First, there is a clear progression from discrete, low-dimensional algorithms such as Q-learning toward more advanced continuous-control and policy-gradient methods like PPO and SAC, reflecting the growing complexity of industrial applications. Second, integration-oriented approaches, particularly MARL, hierarchical RL, and Digital Twin-Integrated RL, highlight a shift toward decentralized, simulation-driven, and safety-aware frameworks suited to smart factories. Comparative studies suggest that PPO and SAC outperform DQN in continuous control tasks [10,27], particularly when environmental dynamics are highly stochastic, whereas DQN remains competitive for discrete action spaces with lower computational overhead. Finally, the persistence of challenges such as scalability, data scarcity, and hyperparameter sensitivity indicates that no single RL paradigm yet provides a universally practical solution for Industrial Automation. These comparative observations reinforce the need for hybrid, modular, and XRL approaches that can meet the operational demands of Industry 4.0 environments.

Table 1. Overview of RL Models in Industrial Automation.

RL Model/Approach	Industrial Application	Key Advantages	Key Challenges	Representative Studies	Real-World Deployment Status
Q-Learning	Adaptive process control, simple robotic tasks	Easy to implement; effective in discrete action spaces	Limited scalability; poor performance in high-dimensional continuous spaces	[2,8]	Mostly simulation and lab-scale deployment; limited industrial adoption
Collaborative Robotics (Single- vs. Dual-Arm RL-Assisted Control)	Precision assembly, dual-arm and single-arm robot evaluation, production station design	Improves assembly speed and path efficiency; dual-arm faster (20%); supports energy/performance trade-off analysis; simulation-driven optimization	Energy-speed trade-offs; requires high-fidelity simulation; profitability varies by task complexity	[28]	Simulation-based evaluation; applicable to industrial workstation design and robot deployment planning
Deep Q-Networks (DQN)	Robotic automation, assembly, pick-and-place	Handles high-dimensional sensory inputs; scalable to complex tasks	Requires large datasets; computationally intensive	[9,10,32]	Some pilot-scale deployments; emerging industrial adoption in robotics
Proximal Policy Optimization (PPO)	Supply chain management, production scheduling	Stable learning; effective for continuous action spaces; robust to hyperparameter changes	Sensitive to environment variability; may require extensive tuning	[11,33,34]	Primarily simulation; limited real-world application in dynamic supply chains
Soft Actor-Critic (SAC)	Energy-efficient process control, continuous robotics tasks	Sample-efficient; handles continuous actions; robust exploration	Complex implementation; sensitive to hyperparameter tuning	[35–37]	Mostly lab-scale; experimental deployments in process control and robotics
MARL	Multi-robot coordination, collaborative production	Decentralized decision-making; scalable for multi-agent environments	Communication overhead; coordination complexity; convergence issues	[1,15,38]	Early-stage pilot tests; limited industrial adoption
Offline/Batch RL	Processes with limited real-time data	Learns from historical datasets; safer for industrial deployment	Poor adaptability to unobserved states; performance depends on dataset quality	[2,39,40]	Some real-world trials in chemical and manufacturing processes
Hierarchical/Modular RL	Complex industrial systems with multiple sub-tasks	Enables task decomposition; improves scalability and transfer learning	Requires careful task structuring; may increase computational complexity	[41,42]	Mostly experimental and simulation-based; few industrial pilots
Constrained/Safe RL	Safety-critical industrial processes	Incorporates safety constraints; reduces risk of operational failures	Can limit exploration and performance; complex constraint design	[43–45]	Limited industrial deployment in high-risk manufacturing
Digital Twin-Integrated RL	Real-time optimization and predictive maintenance	Allows virtual simulation before deployment; reduces risk; enhances safety	Requires accurate digital twin models; high computational demand	[46,47]	Pilot projects in manufacturing, energy, and logistics; increasing adoption
IoT-Optimized RL	Energy management, sensor networks, consumer electronics	Reduces network congestion, improves energy efficiency, handles redundant data patterns	Limited to specific IoT architectures, requires sensor state management	[30,31]	Early industrial adoption in consumer electronics; emerging in industrial IoT

While Q-Learning and DQN have been validated in laboratory and small-scale industrial settings, approaches such as hierarchical RL and digital twin-integrated RL are predominantly explored in simulation, with limited deployment in operational factories. This highlights the gap between theoretical development and real-world applicability, emphasizing the need for scalable and safe RL implementations.

Despite advances in robotic automation, applications of RL in industrial energy management, adaptive supply chain optimization, and other complex, high-dimensional industrial processes remain underexplored, highlighting clear avenues for future research. The literature consistently demonstrates that RL has the potential to transform Industrial Automation by enabling adaptive control, intelligent decision-making, and continuous self-optimization. However, fully realizing this potential requires addressing several critical barriers, including data inefficiency, operational safety, model interpretability, and seamless integration with existing industrial infrastructure.

These challenges are particularly pronounced in real-world deployments. Sample inefficiency remains a core limitation, as most RL algorithms demand extensive interaction data to converge on optimal policies, making training in high-cost or high-risk industrial environments resource-intensive and time-consuming [8]. In addition, the black-box nature of DRL models constrains interpretability, which is essential for operator trust, regulatory compliance, and accountability in safety-critical settings [48]. Finally, the integration of RL into legacy systems poses significant technical hurdles, necessitating the development of middleware solutions, modular architectures, and hybrid deployment strategies to enable gradual adoption without disrupting established control frameworks [49]. Addressing these limitations has motivated emerging strategies such as digital twin integration, modular and hierarchical RL, and transfer learning, which will be explored in the subsequent section on Future Directions and Research Opportunities.

These observations underscore the need for targeted strategies that combine algorithmic innovation with system-level solutions. Approaches such as digital twin integration, hierarchical and modular RL, and safe RL frameworks offer promising pathways to overcome these barriers. The subsequent sections as well as the Future Directions and Research Opportunities section build upon these insights, outlining strategies for bridging the gap between theoretical RL advancements and practical, large-scale industrial implementation.

2.2. Research Dimensions and Synthesis

Building on the algorithmic overview in Table 1, Table 2 presents a comprehensive synthesis of key research dimensions in RL for Industrial Automation. It organizes findings from the reviewed literature by categorizing RL applications, core algorithms, integration technologies, and emerging strategies in smart manufacturing. The table illustrates the evolution of RL from classical, model-free approaches such as Q-Learning to advanced frameworks including DQN, PPO, SAC, and MARL. It also highlights the growing integration of RL with enabling technologies such as digital twins, CPS, and IoT, which collectively enhance autonomy, real-time optimization, and system resilience. In addition, the Key Trends & Insights column emphasizes the most significant developments, emerging approaches, and practical considerations, providing a clear view of the maturity, applicability, and future directions of RL in smart industrial systems.

Table 2. Summary of Key Dimensions in RL for Industrial Automation.

Aspect	Description/Coverage	Representative Studies	Key Insights/Contributions	Identified Gaps & Future Research Needs	Key Trends & Insights
Research Focus	Application of RL for adaptive control and intelligent decision-making in smart manufacturing environments	[2,6,9]	RL enables autonomous optimization of processes, robotics, and supply chain systems through continuous learning	Limited real-world deployment and integration with legacy automation systems	Shift from classical RL to multi-agent and hierarchical RL for complex systems
Core RL Algorithms	Q-Learning, DQN, PPO, SAC, MARL, Offline RL, Hierarchical RL, Safe RL, Digital Twin–Integrated RL	[1,6,9–11,13,14,27,29]	Progressive evolution from discrete control to continuous, multi-agent, and simulation-driven control	Sample inefficiency, hyperparameter tuning, and limited interpretability	Emergence of safe, explainable, and hybrid RL methods
Industrial Application Areas	Process control, robotic automation, logistics and supply chain, energy management, predictive maintenance	[9,24,25,28–31]	RL improves precision, adaptability, and energy efficiency across domains	Underexplored applications in energy management and large-scale coordination	Growing focus on sustainable, energy-efficient, and smart manufacturing
Integration Technologies	CPS, IoT, Digital Twins, Edge AI, Cloud-based RL frameworks	[3,4,16,30,40,41]	Enhances simulation-based optimization, safety, and real-time decision-making	High computational cost and need for real-time synchronization between physical and virtual layers	Digital twins and CPS integration facilitate sim-to-real transfer and pilot deployments
Advantages of RL in Automation	Adaptive decision-making, real-time optimization, autonomous control, self-learning, reduced programming effort	[2,8,25,30,31]	Significant improvement in operational flexibility and fault tolerance	Scalability and explainability remain key constraints	Supports modular, flexible, and self-optimizing production processes
Major Challenges	Sample inefficiency, safety and robustness, interpretability, integration with legacy systems	[6,18,34]	Recognized barriers to industrial deployment	Need for hybrid RL methods combining model-based, safe, and explainable learning	Safety, trust, and interpretability are critical bottlenecks
Emerging Strategies	Digital Twin–Integrated RL, Hierarchical and Modular RL, Safe and XRL, Transfer Learning	[16,17,35,36,40]	Improve scalability, safety, and transferability	Lack of standardized evaluation frameworks and benchmarks	Hybrid, modular, and transfer learning approaches are accelerating real-world adoption
Validation and Deployment	Predominantly simulation or pilot-scale; limited industrial field tests	[9,14,33,41]	Promising pilot results in robotics and predictive maintenance	Require large-scale case studies and longitudinal validation	Simulation remains primary testing environment; limited full-scale deployment
Ethical and Regulatory Aspects	Safety compliance, human–AI collaboration, decision transparency, accountability	[18,37,42,43]	Growing attention to AI ethics and trust	Need for policy-aligned frameworks for safe AI integration	Explainable and safe RL is essential for human–AI collaboration and regulatory compliance
Future Research Directions	Hybrid RL frameworks, digital twin co-simulation, data-efficient learning, explainable and safe RL for Industry 5.0	[19–21,30,31,40,43]	Pathway toward self-optimizing, human-centered, and sustainable smart factories	Bridging theoretical RL advances with industrial-scale implementation	Standardized benchmarks, data-efficient methods, and human-centered designs are needed for widespread adoption

As summarized in Table 2, RL in Industrial Automation has evolved from early single-agent, simulation-driven approaches to multi-agent and hierarchical frameworks capable of handling complex, distributed manufacturing systems. A major trend is the integration of RL with digital twins, CPS, and IoT infrastructures, enabling safe, efficient, and real-time decision-making while facilitating simulation-to-real transfer. Emerging research emphasizes safe, explainable, and hybrid RL methods to address challenges related to trust, interpretability, and regulatory compliance. Although most studies remain at the simulation or pilot scale, there is growing attention to sustainable manufacturing, energy efficiency, and modular system design. Collectively, these insights highlight a trajectory toward human-centered, data-efficient, and standardized RL frameworks that bridge theoretical advances with industrial-scale deployment, paving the way for Industry 5.0.

To provide a comprehensive overview of RL deployment in Industrial Automation, this research developed Table 3, which maps RL categories to representative algorithms, typical applications, integration technologies, advantages, limitations, deployment status, and emerging trends. The table summarizes key insights from the literature and bibliometric analyses, showing how classical RL continues to serve as a foundational benchmark, while DRL, multi-agent, modular, and safe/explainable frameworks are increasingly applied in pilot-scale or experimental industrial settings. This synthesis highlights that selecting an RL approach requires careful consideration of task complexity, operational constraints, agent coordination, and deployment objectives. By consolidating these elements, Table 3 serves as a practical guide for researchers and industrial practitioners seeking to implement scalable, safe, and efficient RL strategies in smart manufacturing environments.

Collectively, the literature indicates that RL has matured beyond algorithmic innovation toward industrial-ready, human-centered, and data-efficient frameworks. While foundational models remain critical for benchmarking, modern research emphasizes safe, explainable, modular, and hybrid approaches integrated with digital twins, CPS, and IoT. This evolution sets the stage for Industry 5.0, where adaptive, scalable, and ethically compliant RL frameworks bridge the gap between theoretical advances and real-world industrial deployment.

2.3. Positioning RL Against Classical Control and Optimization Methods

While this review focuses on RL's role in Industrial Automation, it is crucial to contextualize its value proposition against established control paradigms. Traditional methods like PID control and MPC excel in well-defined deterministic environments but face limitations in complex, stochastic settings. Similarly, optimization techniques like genetic algorithms and particle swarm optimization offer global search capabilities but lack RL's sequential decision-making and adaptive learning features.

The primary advantage of RL lies in its model-free adaptability and capacity for handling uncertainty and nonlinearity [1,2]. Unlike MPC that requires accurate system models or PID controllers needing precise tuning for specific operating points, RL agents learn optimal policies through environmental interaction, enabling them to operate effectively under conditions of dynamic variability and partial observability. This capability is particularly valuable in complex manufacturing environments where explicit modeling is infeasible due to system complexity and stochastic dynamics [2,7].

However, RL exhibits significant disadvantages compared to classical approaches. Sample inefficiency remains a critical limitation, as most RL algorithms require extensive interaction data, often millions of steps, to achieve convergence, making training impractical in high-cost physical systems [7,21]. This contrasts sharply with PID or MPC, which can be deployed rapidly with minimal data. Furthermore, safety and reliability concerns persist, as untested RL policies may cause equipment damage or operational failures, whereas classical controllers offer proven stability and safety records [43,44]. The black-box nature

of DRL models also limits interpretability, a major concern for regulatory compliance and operator trust compared to the transparent logic of traditional control systems [20,48].

Table 3. Research Dimensions, Applications, and Emerging Trends in RL for Industrial Automation.

RL Category	Representative Algorithms	Typical Applications	Integration Technologies	Advantages	Limitations	Deployment Status	Key Trends & Insights
Classical RL	Q-Learning	Discrete control tasks, simple process automation	None/basic simulation	Simple, low computational cost	Poor scalability, limited to low-dimensional problems	Simulation/pilot-scale	Still foundational for benchmarking; used for small-scale, low-risk tasks
DRL	DQN, PPO, SAC	High-dimensional robotic tasks, pick-and-place, assembly, process control, adaptive IoT sensor scheduling, collaborative assembly analysis	Digital Twins, IoT, CPS	Handles high-dimensional inputs, continuous actions; robust and stable	Sample inefficiency, black-box nature	Pilot-scale/industrial experiments	Core of DRL-driven automation; increasing use in multi-stage industrial systems
MARL	MADDPG, QMIX	Collaborative robotics, decentralized production lines, dual-arm cooperative assembly	CPS, IoT, networked environments dual-arm cooperative assembly	Supports coordination among multiple agents; scalable for distributed systems	Coordination complexity, communication overhead	Pilot-scale/research testbeds	Essential for multi-agent and distributed industrial applications
Hierarchical & Modular RL	Options framework, Feudal RL	Multi-stage production, task decomposition	Digital Twins, IoT	Knowledge transfer, improved learning efficiency, interpretable	Complex design, integration overhead	Pilot-scale	Emerging for scalability, modularity, and energy efficiency
Safe & XRL	Constrained RL, XRL	Safety-critical processes, predictive maintenance	Digital Twins, CPS	Ensures operational compliance, builds trust, improves reliability	Slower learning, limited generalization	Pilot-scale/small deployment	Key for regulatory compliance and human-centered design
Hybrid RL	Model-based + model-free approaches	Adaptive control, optimization under uncertainty, IoT power Management	Digital Twins, CPS, IoT	Combines sample efficiency with robustness; reduces sim-to-real gap	Design complexity, computational cost	Simulation/early industrial deployment	Trend toward bridging theory and practical deployment; supports Industry 5.0 objectives

As summarized in Table 4, RL is not a universal replacement but rather a complementary approach. Classical methods remain preferable for well-understood, deterministic processes requiring guaranteed stability and rapid deployment. RL excels in complex, dynamic environments where adaptability and autonomous optimization provide significant value, particularly when safe simulation or digital twin environments are available for training [46,47]. Future hybrid architectures that combine RL's adaptability with the stability guarantees of classical control represent a promising direction for Industrial Automation [19,45].

The literature review highlighted the growing interest in applying RL and DRL within Industrial Automation and robotics, while also revealing the need for a structured and quantitative assessment of research developments in this area. To address this gap, the present study employs a bibliometric review methodology to identify, screen, and analyze the relevant literature in a systematic manner. The following section presents the methodology used in this study.

Table 4. Comparative Analysis: RL vs. Established Industrial Control Methods.

Aspect	Classical Control (PID/MPC)	Optimization Algorithms	Reinforcement Learning
Model Requirements	Requires accurate system models (MPC) or transfer functions	Model-free or model-based variants	Model-free; learns from interaction
Adaptability	Limited; requires retuning for new conditions	Good for static optimization	High adaptability to dynamic environments
Handling Uncertainty	Limited performance under high uncertainty	Variable performance	Excels in stochastic, nonlinear environments
Sample Efficiency	High-minimal data for tuning	Medium-high	Very low-millions of interactions needed
Safety & Reliability	Proven stability guarantees	Depends on implementation	Safety concerns during exploration
Interpretability	Transparent and interpretable	Medium interpretability	Black-box nature limits trust
Computational Demand	Low to medium	Medium to high	High computational intensity
Integration Complexity	Mature integration with legacy systems	Moderate	Challenging integration with existing infrastructure
Real-world Deployment	Extensive industrial deployment	Widespread use	Limited to pilot-scale (22% of studies)
Best Suited For	Well-defined, deterministic processes	Static optimization problems	Complex, dynamic environments with simulation capability

3. Methodology

To ensure both comprehensive coverage and methodological rigor, this review adopted a two-stage analytical process. First, the PRISMA 2020 protocol was adapted and applied to systematically identify, screen, and select relevant studies, ensuring a transparent, reproducible, and unbiased construction of the final dataset [50]. Second, a modified JBI critical appraisal checklist was applied to evaluate the methodological quality and risk of bias of the included studies, ensuring that the subsequent synthesis and interpretations were grounded in robust methodological evidence [51].

Complementing these systematic review procedures, the study also employed a bibliometric analysis using the open source R package software (bibliometrix (version 4.5.1) and its biblioshiny web interface) to quantitatively examine the RL and DRL literature within Industrial Automation and robotics. Bibliometric techniques provide insights into publication trends, influential authors and institutions, collaboration networks, and thematic evolution, thereby enriching and contextualizing the qualitative findings from the systematic review [52–54]. By integrating bibliometric mapping with qualitative synthesis, this review offers a comprehensive and multidimensional understanding of research developments, knowledge gaps, and emerging trajectories in the field.

3.1. Data Sources and Search Strategy

The bibliographic data were retrieved primarily from Scopus, one of the most comprehensive databases for engineering and computer science research [54]. Scopus was selected because it provides extensive coverage of peer-reviewed journals and offers detailed citation information, ensuring both the breadth and quality of the retrieved literature. Furthermore, a significant portion of the high-quality journals indexed by Scopus are also covered by other prestigious indexes like ISI Web of Science, reinforcing its standing as a source for established, high-impact literature. Compared with other databases such as Web of Science or IEEE Xplore, Scopus encompasses a wider interdisciplinary scope, which is particularly valuable for identifying RL research that intersects computer science, engineering, and industrial applications. To capture relevant literature, a structured search

strategy was implemented using controlled keywords related to RL applications in Industrial Automation, robotics, and human–robot collaboration. The query was designed to balance precision and comprehensiveness by combining the core terms “reinforcement learning” and “deep reinforcement learning” with application-specific keywords such as “robotics,” “autonomous systems,” “manufacturing automation,” “path planning,” “assembly,” and “human–robot collaboration.” This approach ensured the inclusion of studies that not only focused on the theoretical development of RL but also its applied dimensions in industrial and collaborative robotic contexts [55,56]. The advanced Scopus query employed in this study was:

```
((TITLE-ABS-KEY("reinforcement learning") OR TITLE-ABS-KEY("deep reinforcement learning")) AND (TITLE-ABS-KEY("robotics") OR TITLE-ABS-KEY("autonomous systems") OR TITLE-ABS-KEY("manufacturing automation")) AND (TITLE-ABS-KEY("path planning") OR TITLE-ABS-KEY("assembly") OR TITLE-ABS-KEY("human-robot collaboration")) AND PUBYEAR > 2016 AND PUBYEAR < 2026 AND (LIMIT-TO(SUBJAREA,"COMP") OR LIMIT-TO(SUBJAREA,"ENGI") OR LIMIT-TO(SUBJAREA,"MATH") OR LIMIT-TO(SUBJAREA,"MATE") OR LIMIT-TO(SUBJAREA,"DECI")) AND (LIMIT-TO(DOCTYPE,"ar") OR LIMIT-TO(DOCTYPE,"re")) AND (LIMIT-TO(SRCTYPE,"j")) AND (LIMIT-TO(LANGUAGE,"English"))
```

The search was limited to English-language journal articles and reviews published between 2017 and 2026 within subject areas such as computer science, engineering, mathematics, materials science, and decision sciences. These parameters were chosen to ensure that the dataset focused on high-quality, peer-reviewed, and recent research outputs, capturing the most up-to-date developments in RL applications. The decision to exclude non-journal sources, such as conference proceedings or preprints, was made to maintain methodological rigor and ensure reliability of findings [57]. This search strategy enhanced the relevance, validity, and reproducibility of the review by systematically narrowing the scope to studies that reflect the state-of-the-art in RL-driven Industrial Automation and robotics while excluding less pertinent or non-scholarly materials.

3.2. Inclusion and Exclusion Criteria

The following criteria guided the selection process:

- Inclusion criteria: (i) peer-reviewed journal articles and reviews; (ii) published between 2017 and 2026; (iii) explicitly addressing RL/DRL in robotics, autonomous systems, or manufacturing automation; (iv) containing at least one of the target keywords in title, abstract, or keywords;
- Exclusion criteria: (i) conference proceedings, book chapters, editorials, or non-peer-reviewed sources; (ii) papers outside the scope of Industrial Automation or robotics; (iii) duplicates across databases.

3.3. Selection Process

The foundational step of this review drew on key principles from systematic literature review (SLR) methodology to ensure a transparent and well-structured dataset, as recommended in evidence-based research practices [58,59]. The initial Scopus query retrieved 1226 records. The selection process followed an adapted PRISMA framework, commonly used to enhance transparency and reproducibility in review studies [50,60], without constituting a full systematic review. The PRISMA-based flow diagram (Figure 1) outlines the identification, screening, and eligibility stages, ensuring that each inclusion and exclusion step is clearly documented (see Supplementary File and Table A1 in Appendix A).




Figure 1. PRISMA Screening Flow Diagram.

After the initial search, 412 records (books, editors, conferences, and proceedings) were excluded, leaving 814 documents for screening. A second round of screening, based on titles, abstracts, and keywords, eliminated 65 additional records. To maintain focus and accessibility, 77 non-English articles were removed, resulting in 672 final publications retained for bibliometric and qualitative analysis.

The application of PRISMA enhances this study by:

1. Ensuring methodological rigor—all decisions for inclusion/exclusion are systematic rather than arbitrary [61];
2. Enhancing transparency—readers can trace the exact filtering path from 1226 initial results to 672 final articles [50];
3. Improving reproducibility—future researchers can replicate or extend the review following the same protocol [62].

3.4. Data Extraction and Preparation

Bibliographic metadata, including titles, abstracts, authors, affiliations, keywords, citations, and references, were exported from Scopus in CSV and plain text/BibTeX formats. The datasets were merged, and duplicates were removed primarily via DOI-based matching, with fuzzy matching on titles, first author names, and year applied when DOIs were missing. Author names, institutional affiliations, and keywords were standardized to ensure consistency across the dataset.

Given the Scopus download limit of 200 records per batch, the full dataset of 672 records was obtained in four separate batches, merged into a single BibTeX file, and imported into R Biblioshiny for bibliometric analysis.

3.5. Bibliometric Analysis Tools

To systematically evaluate the literature on RL in Industrial Automation, bibliometric methods were applied [58]. Data analysis was conducted using the R package Bibliometrix (version 4.5.1) and its web interface Biblioshiny, which provided descriptive statistics

(e.g., annual scientific production, leading journals, authors, and countries), citation analysis (e.g., most cited documents, h-index), and network analyses (e.g., co-authorship, co-citation, bibliographic coupling). Several studies have successfully applied Bibliometrix in related domains, such as del Real Torres et al. [5], Samal [63], and Ouhadi et al. [64], demonstrating its reliability and versatility.

The bibliometric analysis was conducted for comprehensive data analysis and visualization. High-resolution bibliographic maps were generated directly within this environment to visualize keyword co-occurrence networks and thematic clusters, which were instrumental in delineating the conceptual structure of the research field. To ensure the robustness of the findings, multiple keyword occurrence thresholds were tested during the analysis.

Using bibliometric analysis in this research is especially useful for three reasons. First, RL in Industrial Automation is highly interdisciplinary, spanning computer science, control engineering, and operations research; bibliometric mapping helps reveal overarching patterns that narrative reviews might overlook. Second, it provides an objective, evidence-based way to track research communities, influential works, and emerging paradigms such as safe RL, digital twin integration, and multi-agent systems. Third, by combining quantitative mapping with qualitative synthesis, this approach improves transparency and reproducibility, ensuring insights are grounded in systematically processed data.

The bibliometric results were combined with qualitative content analysis of highly cited and thematically central papers to provide interpretive depth. Clusters derived from co-occurrence and co-citation networks guided the structuring of the literature review into thematic areas, including RL for motion planning, RL for assembly, and RL for human–robot collaboration. Research gaps and opportunities were identified by examining low-density areas in thematic maps and emerging keywords, providing a comprehensive understanding of trends, challenges, and future directions in RL applications for Industrial Automation.

3.6. Methodological Quality Assessment

To complement the bibliometric analysis and ensure the analytical rigor of our review, a methodological quality assessment was performed on the final corpus of publications. We employed a modified version of JBI critical appraisal checklist [51]. Each study was evaluated based on three key criteria:

1. Clarity of Research Aims: Were the objectives and research questions explicitly defined?
2. Methodological Rigor: Was the RL methodology (algorithm, state/action space, reward function) described with sufficient detail for replication?
3. Transparency of Results: Were the outcomes, including limitations and deployment challenges, clearly reported?

Each criterion was scored as ‘1’ if fully met, ‘0.5’ if partially met, and ‘0’ if not met or not reported. Refer to Table A1, Appendix A, where: ✓ = Criterion fully met; △ = Criterion partially met; X = Criterion not met. The studies achieving a total score of 3 were classified as ‘high quality.’ The overall quality index (Q) of the corpus was calculated as the percentage of high-quality studies using the following equation:

$$Q(\%) = (N_{\text{high quality}} / N_{\text{total}}) \times 100$$

where $N_{\text{high quality}}$ is the number of studies scoring 3, and N_{total} is the total number of included studies.

The results of this assessment are synthesized in Section 4, providing critical context for the bibliometric findings and ensuring that our review is grounded in a qualitatively evaluated evidence base.

4. Results and Discussion

The methodological quality assessment of the 90 included studies revealed a high overall standard, with 75.6% (68/90) of studies meeting all five quality criteria. This can be observed from Table A1, Appendix A, where none of the studies scored below 0.5 in all 5 criteria. This provides a robust and reliable foundation for the subsequent bibliometric and thematic analyses conducted in this review.

RL has emerged as a pivotal tool in Industrial Automation, enabling adaptive decision-making across robotics, process control, supply chain management, and predictive maintenance. Its defining characteristic—learning optimal policies through interaction with dynamic environments—allows industrial systems to respond effectively to uncertainty and evolving operational conditions [2,23].

Comparative analyses of RL models reveal distinct patterns of industrial applicability. Classical RL, particularly Q-Learning, remains widely used for discrete control tasks due to its simplicity and low computational requirements [6]. However, in high-dimensional or continuous environments, its scalability and performance are limited.

DRL models, such as DQN, PPO, and SAC, extend classical RL by incorporating NNs, enabling robots to process high-dimensional sensory data for tasks like pick-and-place operations, assembly, and process optimization. These algorithms demonstrate robustness, stability, and suitability for complex, continuous action spaces [8,33].

MARL supports decentralized decision-making for collaborative robotics and distributed production lines, addressing coordination and scalability challenges [17]. Hierarchical and modular architectures improve task decomposition and knowledge transfer, enhancing learning efficiency in complex multi-stage processes [64]. Safe and XRL frameworks ensure operational reliability in safety-critical contexts, while digital twin integration allows virtual experimentation to reduce deployment risks and enable real-time optimization [8,33,36].

Tables 2 and 3 provide complementary insights into RL evolution and deployment. Table 2 tracks the trajectory from early single-agent, simulation-driven methods to hybrid, human-centered frameworks, illustrating the historical and conceptual development of RL in Industrial Automation. Table 3, developed by the researchers, maps RL categories to algorithms, typical applications, integration technologies, advantages, limitations, deployment status, and emerging trends. This table serves as a practical guide for industrial adoption, showing how classical RL remains foundational for benchmarking, while DRL, multi-agent, modular, safe, and hybrid frameworks are increasingly deployed in pilot-scale or experimental industrial settings. Together, these tables highlight that algorithm selection must align with task complexity, operational constraints, and deployment objectives. Together, Tables 2 and 3 highlight the evolution from single-agent to human-centric, multi-agent frameworks, mapping the trajectory of RL research toward the complex, collaborative environments envisioned by Industry 5.0.

A comparative assessment of RL approaches highlights a clear clustering based on application suitability, complexity, and deployment readiness:

- Discrete vs. Continuous Control: Q-Learning excels in low-dimensional, discrete processes, whereas DQN, PPO, and SAC are preferable for high-dimensional, continuous tasks due to their NN-based approximations [8];
- Single-Agent vs. Multi-Agent Systems: Traditional RL and DQN are suitable for single-agent tasks; MARL is essential when multiple agents must coordinate in a shared environment, such as collaborative robotics or multi-stage production lines [17];
- Safety-Critical Environments: Safe RL and digital twin-integrated approaches are indispensable in contexts where operational failures could have severe financial or safety consequences, offering a balance between exploration and constrained policy execution [36];
- Transfer and Modular Learning: Hierarchical and modular RL architectures outperform classical RL in multi-task or variable production settings by enabling knowledge transfer and task decomposition, thus improving scalability and efficiency [65,66].

Overall, DRL algorithms dominate robotic automation research due to their capacity to handle high-dimensional inputs, while model-based, offline, or modular approaches are favored in process control and predictive maintenance for safety and sample efficiency [2,27]. Algorithm selection must align with task complexity, agent coordination, and operational safety. Modular and hierarchical architectures facilitate knowledge transfer across multi-stage processes, and safe RL frameworks, often integrated with digital twins, ensure reliable operation in critical contexts.

A deeper analysis of the literature reveals the specific dimensions of key challenges like scalability and interpretability. The scalability challenge manifests not just in computational cost, but in the “curse of dimensionality” for high-dimensional state-action spaces and the complexity of multi-agent coordination. In response, emerging trends focus on hierarchical RL to decompose problems, model-based RL to improve sample efficiency, and modular architectures that enable knowledge transfer across tasks, thereby reducing the learning burden for new scenarios.

This focus on scalable coordination is powerfully exemplified in the MARL sector, which constitutes 27% of the reviewed studies. A deeper analysis of this sector reveals distinct thematic trends and a focus on specific industrial problems. The bibliometric data and literature synthesis indicate that a significant portion of MARL research is dedicated to robotic swarm coordination in warehouse logistics and dynamic task allocation in flexible assembly lines. The primary “issue” dominating this subfield, as identified in the literature, remains the non-stationarity problem and the need for robust credit assignment mechanisms. In response, a clear “trend” has emerged: the CTDE paradigm is the predominant architectural framework, cited in over 60% of the MARL studies reviewed, as it effectively mitigates the non-stationarity problem during training while allowing for scalable decentralized operation. This is often combined with attention mechanisms to manage communication in large-scale agent systems, pointing to a consistent research direction focused on scalability and decentralized coordination.

Similarly, the interpretability challenge, a major barrier to trust and regulatory approval, stems from the inherent “black-box” nature of deep NNs used in DRL. The bibliometric analysis and literature synthesis confirm a rising trend in XRL, which aims to make agent decisions transparent. This includes methods for post hoc explanation of policy decisions and the development of intrinsically interpretable models that prioritize simplicity and clarity, ensuring that human operators can understand, trust, and effectively collaborate with RL systems in safety-critical industrial settings.

Research Trends and Insights from Bibliometric Analyses

Emerging strategies address core challenges in industrial RL. Digital twin-integrated RL enables safe virtual experimentation and pilot deployment, while transfer learning accelerates policy adaptation across similar tasks, enhancing sample efficiency [67]. Hierarchical and modular architectures improve knowledge reuse, task decomposition, and interpretability [68]. Safe and XRL frameworks ensure compliance with operational and regulatory constraints [20,69,70].

A synthesis of the literature highlights gaps in real-world validation, standardization protocols, and ethical frameworks. Many studies remain confined to simulation or laboratory settings, emphasizing the need for pilot deployments, cross-industry benchmarking, and regulatory alignment. Future research should focus on developing scalable, interpretable, and transferable RL frameworks that balance operational efficiency, safety, and adaptability in complex industrial contexts.

Bibliometric analyses provide further insights. The co-occurrence network in Figure 2 reveals three distinct yet interconnected thematic clusters: a methodological core (green), robotics applications (blue), and multi-agent systems (red) reflecting RL’s evolution in Industry 4.0. This structure visually confirms that algorithmic research is tightly coupled with applied domains like industrial robotics and collaborative systems.




Figure 2. Co-occurrence Network.

Figure 3, illustrating the temporal evolution of article production in RL for Industrial Automation by country, strongly validates the urgent need for this review by demonstrating the field's explosive growth and its pronounced geographic stratification.

These findings align with [71,72], which highlights China as the unequivocal Tier 1 global leader in RL research. The data clearly establishes China’s dominant position, with an article output projected to reach approximately 700 by 2026, underscoring the country’s central role in driving methodological innovations and applied developments in the field [73].

This dramatic, near-exponential acceleration after 2019 suggests an enormous, centrally driven national investment, which is consistent with external reports indicating China’s surge in industrial robotics adoption and its increasing robot density (e.g., reaching 470 robots per 10,000 employees in 2023, surpassing Germany and Japan). This high level of physical automation provides the ideal real-world platform and data-rich environment necessary for developing and testing complex RL policies in smart factories. The dominance of China in application-focused research, as seen in Figure 3, suggests that real-world deployment scales with the availability of physical industrial infrastructure, a key factor in bridging the sim-to-real gap.




Figure 3. Countries' Production over Time.

Conversely, the USA (Tier 2), while a robust second contributor (projected around 140–150 articles), shows a less aggressive, more linear growth pattern. This structure implies that while the theoretical and core algorithmic advances may still be heavily influenced by Western research (consistent with North America's lead in the broader AI market), the scale and volume of practical industrial application research is being overwhelmingly driven by Asia-Pacific economies, particularly China. This disparity underscores a central challenge of RL deployment: the sim-to-real gap and the struggle for sample efficiency; the countries with the highest deployment of physical systems (China, Republic of Korea, which has the world's highest robot density) are logically generating the most application-focused research [71,74].

Therefore, the comprehensive synthesis proposed in this review must prioritize examining the methodologies emanating from these high-volume research regions, particularly focusing on how they address safety, robustness, and integration challenges within highly automated, large-scale industrial systems. The structure of the research landscape shown in Figure 3 dictates that the global future of industrial RL will be shaped by the success of these massive, national-level automation initiatives.

Figures 2 and 3 visually reinforce the findings from Tables 2 and 3. The co-occurrence network, Figure 2, reveals three interconnected thematic clusters: the methodological core (green) encompassing RL, DRL, motion planning, and path planning; robotics applications (blue) with industrial robots and manipulators; and multi-agent/autonomous systems (red) emphasizing UAVs and distributed applications. The temporal evolution of publications by country (Figure 3) highlights China as the global leader, with projected output of ~700 articles by 2026, followed by the USA as a Tier 2 contributor (~140–150 articles). These patterns indicate that practical, application-focused research is largely concentrated in Asia-Pacific regions, reflecting their extensive automation infrastructures and data-rich environments.

The quantitative trends revealed by the bibliometric analysis are substantiated by concrete case studies that demonstrate RL's practical impact across different layers of

the Industrial Automation stack. For instance, at the physical robotics layer, the work of Peta et al. [28] provides a critical performance and economic validation, showing how RL-driven systems enable sophisticated trade-off analyses, such as a 33% speed increase with a dual-arm robot versus the energy efficiency of a single-arm system, directly informing high-value deployment decisions. Simultaneously, at the foundational data acquisition layer, studies by Khan et al. [30,31] demonstrate RL's role in optimizing the underlying IoT infrastructure, achieving metrics like a 35% boost in energy efficiency and a 6.25% higher transmission success rate for sensor networks. These case-based results confirm that the heavy research focus on “robotics,” “motion planning,” and “deep learning” identified in the bibliometric maps (Figure 2) is yielding tangible advancements. They illustrate a dual focus: on high-level, physically embodied intelligence for complex tasks, and on low-level, resource-aware intelligence for sustaining the pervasive sensing required in smart factories.

The thematic flow in Figure 4 quantitatively substantiates the strategic branching of research, showing a strong transition from core RL theory towards high-level decision-making and robot programming, confirming the field's shift toward autonomous intelligence. The internal flow from RL (2017–2024) to itself (RL—2025–2026) is the most frequent connection, accounting for 428 occurrences, with a high Weighted Inclusion Index of 0.59. These findings align with the results reported in [71,74].




Figure 4. Thematic Research Flow (2017–2026).

This indicates a persistent, high-volume effort dedicated to refining the core theory, particularly concerning “reinforcement learnings” and “deep learning,” which are essential for addressing the sample inefficiency challenges noted in the review. Crucially, the strongest practical transitions are seen in two primary application domains: the flow into decision-making (90 occurrences) and robot programming (159 occurrences), both exhibiting high Weighted Inclusion Indices of 0.71. This metric confirms that the existing RL knowledge is highly relevant and directly transferable to these high-level robotic tasks (keywords: “robotics, machine learning, decision-making, intelligent robots”), solidifying the review’s argument that RL is enabling the shift from human-coded control to autonomous, adaptive intelligence. A parallel specialization is confirmed in the motion planning research stream, which shows its most substantial flow toward autonomous vehicles (62 occurrences) with the highest Weighted Inclusion Index of 0.70. This quantitative split confirms a strategic focus: while RL refines its own algorithms, the established control expertise (motion planning) is being systematically applied to mobile platforms (keywords: “autonomous vehicles, autonomous driving”), highlighting the industry’s strategic investment in flexible, logistics-based automation.

Table 5 provides a critical mapping of the global distribution of academic efforts and resources, revealing that the ten largest research areas account for over 55% of the entire represented knowledge space. This concentration indicates that the current research agenda is predominantly centered on DRL as the core technology for achieving physical autonomy in industrial robotics, consistent with the findings of Rajamallaiiah et al. [75] and Gautam [76], who both emphasized the pivotal role of DRL in advancing control, resilience, and autonomy within power and energy systems. The primary research investment is clearly concentrated in the RL Core, with RL (594 occurrences, 14%) and its variants (358 occurrences, 8%) dominating the landscape. This dominance underscores the community's focus on policy optimization and foundational RL theory. The research effort cascades into areas concerned with managing complexity and physical movement, with motion planning (339 occurrences, 8%) and DL (233 occurrences, 5%) reflecting the strategic shift toward DRL architectures capable of handling continuous state and action spaces. This trajectory effectively reduces human-intensive efforts in robot programming (219 occurrences, 5%) by replacing static control laws with adaptive, self-learned policies for tasks such as path planning (179 occurrences, 4%). The remaining high-frequency terms—learning algorithms, robotics, learning systems, and intelligent robots—confirm that the collective research focus is on developing adaptable, intelligent robotic entities, where the human role transitions from programmer to system architect and safety supervisor, ensuring the successful orchestration of fully autonomous systems.

Table 5. Tree Map Analysis.

NO.	Terms	Frequency	%
1	reinforcement learning	594	14.00613063
2	reinforcement learnings	358	8.441405329
3	motion planning	339	7.993397784
4	deep learning	233	5.493987267
5	robot programming	219	5.163876444
6	deep reinforcement learning	205	4.833765621
7	path planning	179	4.220702664
8	learning algorithms	119	2.805941995
9	robotics	105	2.475831172
10	learning systems	95	2.240037727
...			
47	collisions avoidance	28	0.660221646
48	manipulators	27	0.636642301
49	performance	27	0.636642301
50	reinforcement learning (RL)	27	0.636642301

Figure 5, which tracks the cumulative frequency of key research terms from 2017 to 2026, precisely delineates the algorithmic and application focus driving the explosive publication growth observed in this domain. The data underscores that the primary scientific commitment in RL for Industrial Automation is not merely theoretical, but centers on the deep integration of advanced RL techniques with core robotic operations.

According to Li et al. [77], RL has rapidly emerged as a key approach for smart transportation, enabling autonomous decision-making in complex, dynamic systems, with bibliometric and literature analyses highlighting its development, leading topics, application domains, and future research directions over the past decade. Similarly, in the current study on smart factories, the term RL emerges as the overwhelmingly foundational keyword, reaching nearly 600 cumulative occurrences by 2026. This parallel dominance confirms that RL serves as an indispensable methodological framework for adaptive decision-making, whether in transportation systems or Industrial Automation, highlighting its central role across diverse complex, dynamic domains.




Figure 5. Cumulative Frequency of Key Research Terminology Over Time (2017–2026).

Immediately succeeding RL is DRL, which reaches approximately 350 occurrences, demonstrating that researchers are fundamentally concerned with applying DL architectures (which follows closely) to handle the complexity and high dimensionality of real-world industrial state-action spaces. The high prominence of DRL validates the review’s emphasis on algorithms like DQN, PPO, and SAC, which are essential for moving beyond simple discrete control tasks and into continuous, nonlinear processes.

Crucially, the next tier of highly frequent terms defines the critical application sector: robotics and its functional components, motion planning and path planning. The accelerating co-occurrence of ‘robotics’ with ‘motion planning’ and ‘deep learning’ in Figure 5 demonstrates that the research surge is channeled toward enabling autonomous, complex manipulation, directly connecting abstract algorithms to tangible industrial hardware.

In contrast, the slower-growing foundational terms like Learning Systems and Learning Algorithms suggest that the community’s attention has matured beyond discussing generic ML principles. The current focus is a specialized, applied one, targeting the specific computational mechanisms DRL needed to optimize the physical operations (robotics/planning) of smart factory infrastructure. This convergence of DRL and high-level robotic control establishes the core research front that this review is obligated to synthesize.

The chronological analysis in Table 6 further delineates the evolution of RL research toward autonomy and its human implications. The field initially focused on low-volume, foundational work, exemplified by Markov Decision Processes (Frequency 14; Median 2020). Research then rapidly pivoted to scale this complexity, with DL in robotics and Automation (Frequency 8; Median 2020) reflecting early efforts to reduce reliance on human-defined features and expert engineering, transferring the burden of optimization to learning algorithms.

Table 6. Research Terminology Trends and Chronology.

Term	Frequency	Year (Q1)	Year (Median)	Year (Q3)
reinforcement learning	593	2022	2024	2025
reinforcement learnings	358	2023	2024	2025
motion planning	339	2022	2024	2025
learning algorithms	119	2022	2023	2024
learning systems	95	2021	2023	2024
robot learning	87	2024	2025	2025
machine learning	72	2020	2023	2024
robots	56	2020	2021	2022
robotics	105	2021	2022	2024
adversarial machine learning	35	2024	2025	2025
markov processes	30	2020	2022	2024
agricultural robots	29	2020	2021	2023
reinforcement learning method	26	2020	2022	2025
multipurpose robots	26	2024	2025	2025
reinforcement learning approach	17	2021	2021	2023
markov decision processes	14	2019	2020	2024
deep learning in robotics and automation	8	2019	2020	2020
path planning problems	8	2019	2020	2022
robotic manipulators	5	2018	2018	2024

The most recent, high-volume research phase (Median 2024–2025) targets the most complex, human-like abilities in physical systems, including Robot Learning (Median 2025) and motion planning (Median 2025). This trajectory confirms a transition beyond simple automation toward adaptive autonomy, where robots adjust to dynamic environments with minimal or no human intervention.

Crucially, the emergence of newer, lower-frequency terms such as Adversarial ML (Frequency 35; Median 2025) underscores the final, critical step in the human–machine collaboration equation. As noted by Khurram et al. [78], the adoption of autonomous RL in high-risk manufacturing environments introduces significant concerns regarding safety, robustness, and interpretability—factors that inherently influence human trust and regulatory approval. Similarly, the growing emphasis on adversarial ML, as highlighted by Hagendorff [79] and Ilahi et al. [80], demonstrates that researchers are actively confronting these ethical and safety constraints. They recognize that for RL to gain widespread acceptance and fully integrate into industrial operations, the opaque and complex policies learned by these algorithms must be demonstrably safe, interpretable, and reliable enough to operate alongside, or even replace, human workers.

The three-field Sankey diagram (Figure 6) visually encapsulates the global research ecosystem, reinforcing the dominance of China and clarifying the disciplinary channels through which academic output flows into thematic applications. China emerges as the single largest contributor, with a massive volume of research flowing into the core methodologies of reinforcement learning and reinforcement learnings, underscoring its broad, national commitment to both fundamental theory and general application and this is confirmed by [3,81,82]. The flow from key sources, such as IEEE Transactions on Automation Science and Engineering and IEEE Robotics and Automation Letters, confirms that top-tier disciplinary journals are the primary conduits for this research. Critically, regardless of the country of origin (China, USA, India, etc.), the vast majority of research effort converges on a handful of destination keywords: reinforcement learning, DRL, and motion planning [83–85]. This convergence validates that the global community’s focus remains fixed on the essential pillars of adaptive control in smart factories—namely, using advanced RL to enable highly flexible and intelligent physical movements.




Figure 6. Three-Field plot: Source, Country, and Thematic Research Flow.

Figure 7 illustrates the distribution of the top 10 global flow frequencies, making clear a stark and asymmetric pattern of connectivity that is heavily skewed toward Australia as a destination market for high-volume, directional interactions, primarily driven by major global economic powers. The three highest positive “Frequency” values—China → Australia (≈ 334.48), United Kingdom → Australia (≈ 334.33), and USA → Australia (≈ 334.15)—are virtually identical and overwhelmingly dominate the dataset. This clustering suggests an intrinsic, nonlinear relationship (likely log-transformed or normalized trade, migration, or investment flows), where positive directional inflows from these three global giants converge toward a maximum. Collectively, they reflect Australia’s role as a unique global hub: raw material exports to China, historical/social migration from the UK, and investment and technology transfer from the USA.




Figure 7. Global High-Volume Flows and Connectivity.

Beyond these dominant triads, Figure 7 also highlights enduring regional and geopolitical linkages. The next tier of flows—USA → Korea (≈ 55.45), USA → Japan (≈ 54.91), and USA → Hong Kong (≈ 54.89)—underscores the United States’ persistent commercial and strategic presence in East Asia. Meanwhile, the India → UAE (≈ 54.30) pathway captures a different high-volume dynamic: labor migration and remittance flows, a hallmark of the South Asia–Gulf economic model. Taken together, the composition of these flows confirms the table’s representation of global connectivity as a blend of (i) traditional economic trade and investment links (Western/Asian powers → developed markets) and (ii) critical labor corridors (South Asia → GCC).

Crucially, Figure 7 also reinforces the observation that the most intense global relationships are overwhelmingly unidirectional, as the absence of significant negative flows in the top 10 indicates that high-frequency exchanges tend to move in one direction, reflecting asymmetries in global trade, migration, and capital distribution.

Figure 8 illustrates journal publication trends for RL in Industrial Automation, revealing a strong academic focus concentrated in key technical and applied science outlets, dominated by the Institute of Electrical and Electronics Engineers (IEEE), as confirmed by recent studies [63,77,80,85]. The field’s epicenter lies at the intersection of practical control and intelligent decision-making, with IEEE Robotics and Automation Letters leading with 55 publications, highlighting robot control, manipulation, and autonomous operation as the most active areas of RL research. The high volume in IEEE Access (48 documents) emphasizes the importance of rapid, open-access dissemination across engineering domains. Contributions from Sensors (33 documents) and IEEE Internet of Things Journal (15 documents) underscore the critical role of high-quality, real-time sensing and IoT infrastructure in defining the “state” for RL agents in smart factories. Overall, these publication trends confirm that successful industrial RL research integrates algorithmic innovation, control theory, sensing technology, and intelligent automation into a coherent, system-level framework.




Figure 8. Journal Trends in Industrial RL Research.

In summary, bibliometric analyses reveal that RL research in Industrial Automation is increasingly focused on DRL applications in robotics, motion planning, and autonomous systems, with China leading in publication volume and applied experimentation, as supported by documented industrial testbeds and nationwide digital twin initiatives [86,87]. Thematic clusters indicate a clear progression from foundational algorithmic development to high-impact industrial applications, while emerging areas such as safe RL, hierarchical architectures, and digital twin integration are gaining traction. These insights underscore the growing maturity and specialization of the field, while also highlighting persisting gaps in real-world validation, standardization, and safety-aware deployment, which are addressed in the future challenges and research opportunities section.

5. Future Directions and Research Opportunities

Although RL has demonstrated considerable promise in enabling adaptive control and intelligent decision-making in Industrial Automation, several critical challenges and underexplored areas continue to limit its large-scale industrial adoption. Recent research has begun to address these limitations through emerging trends that combine simulation, transfer learning, modular design, and ethical frameworks for autonomous operation.

5.1. Technological Advancements

5.1.1. Digital Twin Integration

The convergence of RL with digital twin technology is a major trend reshaping Industrial Automation. Digital twins provide high-fidelity virtual replicas of physical assets, allowing RL agents to train, validate, and optimize control policies in simulated environments before deployment. This approach enhances safety, efficiency, and scalability while reducing the risks associated with experimentation on real systems. Studies by Abo-Khalil et al. [88] and Xia et al. [3] demonstrate that digital twin-integrated RL systems enable real-time monitoring, predictive maintenance, and adaptive optimization in manufacturing and energy applications, thus bridging the gap between simulation and real-world deployment.

5.1.2. Transfer Learning for Industrial Adaptability

Another promising direction is the use of transfer learning to improve learning efficiency and generalization. RL models often struggle when deployed in new or dynamically changing environments, requiring retraining from scratch. Transfer learning mitigates this challenge by allowing RL agents to leverage knowledge from previously learned tasks. As shown by Hua et al. [89] transferring representations, policies, or value functions across tasks significantly reduces the data and time required for training, thereby accelerating deployment and improving scalability in complex industrial systems.

5.1.3. Modular and Hierarchical RL Architectures

To enhance adaptability across heterogeneous industrial tasks, modular RL frameworks are gaining traction. Raziei [90] introduces a Hyper-Actor Soft Actor-Critic (HASAC) framework that combines task modularization and transfer learning, allowing agents to adapt to varying conditions while maintaining stable performance. Similarly, hierarchical RL architectures facilitate task decomposition, enabling scalable and interpretable learning across multi-stage production processes. Such modularity is particularly valuable in factories with multiple interdependent subsystems requiring coordinated decision-making.

5.2. Implementation Considerations

5.2.1. Standardization and Interoperability

Despite progress in algorithmic development, the absence of standardized protocols for integrating RL within industrial systems remains a substantial barrier. As highlighted by Farooq and Iqbal [2] and Semeraro et al. [55], a lack of common frameworks complicates interoperability across different platforms and manufacturing environments. Establishing

industry-wide standards and best practices would facilitate consistent deployment, enhance safety verification, and promote trust among industrial stakeholders.

5.2.2. Ethical and Regulatory Considerations

The deployment of autonomous RL systems in industrial environments raises significant ethical and regulatory concerns that extend beyond technical performance metrics. As identified in our bibliometric analysis, the growing emphasis on human–AI collaboration and Industry 5.0 principles necessitates addressing fundamental ethical challenges that impact trust, safety, and societal acceptance.

A primary concern is the black-box nature of DRL models, which constrains interpretability and complicates accountability in safety-critical industrial settings [48]. This opacity creates challenges for regulatory compliance and operator trust, particularly in applications where human workers collaborate directly with RL-driven systems. The lack of transparent decision-making processes in most DRL architectures undermines the ability to conduct meaningful safety audits or explain operational decisions to regulatory bodies.

Safety and liability concerns represent another critical dimension. As noted by Rahanu et al. [43] and Martinho et al. [44], the absence of formal verification methods for RL policies creates uncertainty about system behavior in edge cases or unexpected scenarios. When RL systems cause operational failures or safety incidents, determining liability becomes complex—spanning algorithm developers, system integrators, and end-users. This ambiguity is particularly problematic in high-risk manufacturing environments where equipment damage or worker safety may be compromised during the exploration phase of RL training.

The integration of RL into human–robot collaborative environments introduces additional ethical considerations related to workforce impact and algorithmic bias. RL policies trained on operational data may inadvertently inherit and amplify existing biases in production systems, potentially leading to unfair resource allocation or discriminatory operational patterns. Furthermore, the transition toward autonomous RL systems raises questions about workforce displacement and the ethical responsibility of organizations to retrain and redeploy human workers.

Our analysis reveals that current regulatory frameworks are inadequately prepared to address these challenges. Existing industrial safety standards and machinery directives lack specific provisions for self-learning autonomous systems, creating a regulatory gap that hinders widespread adoption. The development of comprehensive guidelines for responsible AI deployment in industrial settings requires collaboration between researchers, industry stakeholders, and regulatory bodies to establish standards for testing, validation, and continuous monitoring of RL systems.

Emerging approaches such as Safe RL and XRL offer promising pathways to address these concerns. Safe RL frameworks incorporate formal safety constraints directly into the learning process, reducing risks during both training and deployment [43,44]. XRL methods enhance transparency by providing post hoc explanations of policy decisions or developing intrinsically interpretable models, thereby building operator trust and facilitating regulatory approval [20,48].

Future regulatory frameworks must balance innovation with protection, establishing clear requirements for:

1. Transparency and explainability in RL decision-making processes
2. Safety verification protocols for RL systems before and during deployment
3. Liability allocation frameworks for autonomous system failures
4. Continuous monitoring and update procedures for evolving RL policies
5. Human oversight mechanisms ensuring meaningful human control in critical decisions.

Addressing these ethical and regulatory considerations is essential for building stakeholder trust and enabling the responsible integration of RL technologies into industrial ecosystems aligned with Industry 5.0's human-centric vision.

5.2.3. Real-World Validation and Scalability

Finally, the literature reveals a notable gap between simulation-based studies and real-world implementation. Most existing research validates RL algorithms within controlled or synthetic environments, which often fail to capture the full complexity of industrial operations [6,8]. To advance the field, large-scale pilot deployments and cross-sector collaborations are needed to assess RL performance under real-world conditions, including variable process dynamics, human-machine interactions, and hardware constraints.

Future research should focus on bridging the gap between theoretical development and industrial implementation through standardization, safety assurance, and practical validation. As manufacturing moves toward Industry 5.0, advancing human-machine collaboration becomes increasingly important, shifting the focus from full automation to human-centric, resilient, and sustainable production systems. Ensuring that autonomous technologies complement human expertise will be essential for creating adaptable, trustworthy, and ergonomically aligned work environments.

Digital twin integration, transfer learning, and modular architectures represent key technological enablers, while ethical frameworks and interoperability standards will be critical for responsible and sustainable deployment. Strengthening human-AI teaming, particularly in decision-making, supervisory control, and cognitive assistance, will further support the transition toward Industry 5.0's vision of synergistic collaboration between humans and intelligent systems.

Addressing these opportunities will accelerate progress from the autonomous, data-driven factories of Industry 4.0 toward the human-centric, value-driven, and resilient manufacturing ecosystems envisioned under the Industry 5.0 paradigm.

6. Conclusions

This review has systematically examined the applications, challenges, and emerging trends of RL in Industrial Automation, covering domains such as adaptive process control, robotics, supply chain management, and predictive maintenance. RL distinguishes itself from traditional control and optimization methods by enabling systems to learn optimal policies through continuous interaction with dynamic environments, thereby facilitating adaptive, intelligent, and autonomous industrial operations. The literature reveals that classical RL methods like Q-Learning remain effective for low-dimensional, discrete tasks, while DRL algorithms, including DQN, PPO, and SAC, are particularly suited for high-dimensional, continuous environments. Multi-Agent RL and hierarchical or modular architectures further support complex, decentralized, and multi-stage industrial operations, highlighting RL's versatility across diverse industrial contexts.

Emerging approaches such as safe RL, digital twin integration, and transfer learning are increasingly addressing critical challenges associated with RL deployment, including operational safety, sample inefficiency, and adaptability across heterogeneous tasks. Bibliometric analyses indicate that China currently leads in application-focused RL research, reflecting a strong nexus between Industrial Automation infrastructure and real-world experimentation, while the USA, Europe, and other regions contribute predominantly in foundational algorithmic and theoretical development. Despite these advances, several persistent challenges hinder large-scale adoption, including sample inefficiency, limited interpretability, safety and robustness concerns, and the complexity of integrating RL systems with legacy industrial infrastructures. These limitations underscore the need for modular, explainable, and rigorously validated RL frameworks that can operate safely alongside human workers.

To systematically guide future efforts and bridge the identified gaps, this review proposes a strategic roadmap structured around four interconnected dimensions: (1) Algorithmic Readiness, focusing on the development of safe, explainable, and data-efficient RL methods; (2) Technological Integration, leveraging digital twins and IoT for seamless simulation-to-reality transfer and deployment; (3) Validation Maturity, establishing standardized benchmarks and promoting rigorous pilot-scale testing in real-world settings; and (4) Human-Centricity, ensuring operational transparency, safe human–robot collaboration, and supportive policies for workforce transition. This roadmap provides a clear framework for stakeholders to prioritize research and development, accelerating the path toward industrially viable RL systems.

Looking forward, the convergence of digital twin technology, transfer learning, and modular RL architectures represents a promising pathway toward scalable, safe, and interpretable autonomous industrial systems. At the same time, the rise of Industry 5.0 places renewed emphasis on human–machine collaboration, where RL-driven automation should augment rather than replace human expertise. Ensuring seamless cooperation between human operators and intelligent systems will require new interaction models, intuitive interfaces, and training mechanisms that keep humans in meaningful supervisory roles.

Equally important is the establishment of standardized protocols, ethical guidelines, and regulatory frameworks to support responsible deployment and build trust among industrial stakeholders. Bridging the gap between simulation-based research and real-world implementation remains a critical priority, as rigorous validation under realistic industrial conditions will ultimately determine RL’s practical value.

In summary, RL has the potential to reshape modern Industrial Automation by enabling adaptive, resilient, and intelligent operations. Achieving this vision demands coordinated progress in algorithmic innovation, human–machine teamwork, system-level integration, safety assurance, and ethical governance, paving the way for fully autonomous, human-centered, and sustainable smart factories aligned with both Industry 4.0 and Industry 5.0 paradigms.

Supplementary Materials: The following supporting information can be downloaded at: <https://www.mdpi.com/article/10.3390/machines13121140/s1>. This supplementary table provides a comprehensive summary of 88 studies examining reinforcement learning (RL) applications in Industrial Automation, spanning sectors from manufacturing and robotics to energy systems, supply chain management, and ethical AI considerations across multiple countries and methodologies.

Author Contributions: Conceptualization, Y.M.A. and O.S.; methodology, O.S.; software (R v 4.5.1), O.S.; validation, Y.M.A., O.S. and W.S.; formal analysis, W.S.; investigation, W.S.; resources, Y.M.A.; data curation, W.S.; writing—original draft preparation, O.S.; writing—review and editing, Y.M.A. and W.S.; visualization, O.S.; supervision, Y.M.A.; project administration, W.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The original contributions presented in this study are included in the article and its Supplementary Materials. Further inquiries can be directed to the corresponding author.

Acknowledgments: During the preparation of this manuscript, the authors used ChatGPT v5.1 and DeepSeek (latest model) to assist with rephrasing and improving the clarity of some text in the Introduction and Literature Review, and to review and correct language, grammar, and formatting throughout the manuscript. Google Gemini Flash 2.5 was used to enhance the resolution of figures. All AI-assisted outputs were reviewed, edited, and validated by the authors, who take full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial Intelligence
ASC-RL	Adaptive Scheduling Control using Reinforcement Learning
CIoT-DPM	Cognitive IoT-based Dynamic Power Management
CPS	cyber-physical systems
CTDE	Centralized Training with Decentralized Execution
DL	Deep Learning
DQN	Deep Q-Network
DRL	Deep Reinforcement Learning
HASAC	Hyper-Actor Soft Actor–Critic
HITL-RL	Human-in-the-Loop Reinforcement Learning
IEEE	Institute of Electrical and Electronics Engineers
IoT	Internet of Things
JBI	Joanna Briggs Institute
ML	Machine Learning
MARL	Multi-Agent Reinforcement Learning
MPC	Model Predictive Control
NNs	Neural Networks
PID	Proportional-Integral-Derivative
PPO	Proximal Policy Optimization
PRISMA	Preferred Reporting Items for Systematic Reviews and Meta-Analyses
RL	Reinforcement Learning
SAC	Soft Actor–Critic
XRL	Explainable Reinforcement learning

Appendix A

Table A1. Methodological quality assessment of included studies using a modified JBI checklist.

No.	Study	Clear Research Aims	Detailed Methodology Description	Comprehensive Results Reporting	Appropriate Data Analysis	Relevant Conclusions	Overall Quality Score
[1]	De Asis et al. (2018)	✓	✓	✓	✓	✓	5/5
[2]	Farooq & Iqbal (2024)	✓	✓	✓	△	✓	4/5
[3]	Xia et al. (2021)	✓	✓	✓	✓	✓	5/5
[4]	Ryalat et al. (2023)	✓	△	✓	✓	✓	4/5
[5]	del Real Torres et al. (2022)	✓	✓	✓	✓	✓	5/5
[6]	Langås et al. (2025)	✓	✓	✓	✓	✓	5/5
[7]	Moos et al. (2022)	✓	✓	✓	✓	✓	5/5
[8]	Arulkumaran et al. (2017)	✓	✓	△	✓	✓	4/5
[9]	Han et al. (2023)	✓	✓	✓	✓	✓	5/5
[10]	Kober et al. (2013)	✓	✓	✓	✓	✓	5/5
[11]	Schulman et al. (2017)	✓	✓	✓	✓	✓	5/5
[12]	Hou et al. (2020)	✓	✓	△	✓	✓	4/5
[13]	Sathya et al. (2024)	✓	△	✓	✓	✓	4/5

Table A1. Cont.

No.	Study	Clear Research Aims	Detailed Methodology Description	Comprehensive Results Reporting	Appropriate Data Analysis	Relevant Conclusions	Overall Quality Score
[14]	Rajasekhar et al. (2025)	✓	✓	✓	✓	✓	5/5
[15]	Hady et al. (2025)	✓	✓	✓	✓	✓	5/5
[16]	Jia & Pei (2025)	✓	✓	✓	✓	✓	5/5
[17]	Abbas et al. (2025)	✓	✓	✓	✓	✓	5/5
[18]	Schwung et al. (2018)	✓	△	✓	△	✓	3/5
[19]	Chang et al. (2023)	✓	✓	✓	✓	✓	5/5
[20]	Emami et al. (2024)	✓	△	△	△	✓	3/5
[21]	Hernandez-Leal et al. (2019)	✓	✓	✓	✓	✓	5/5
[22]	Rolf et al. (2023)	✓	✓	✓	✓	✓	5/5
[23]	Martins et al. (2025)	✓	✓	✓	✓	✓	5/5
[24]	Xu & Peng (2025)	✓	△	△	△	✓	3/5
[25]	Chen et al. (2021)	✓	✓	✓	✓	✓	5/5
[26]	Eckardt et al. (2021)	✓	✓	✓	△	✓	4/5
[27]	Paranjape et al. (2025)	✓	✓	✓	✓	✓	5/5
[28]	Peta et al. (2025)	✓	✓	✓	✓	✓	5/5
[29]	Esteso et al. (2023)	✓	✓	✓	✓	✓	5/5
[30]	Khan et al. (2025)—CIoT-DPM	✓	✓	✓	✓	✓	5/5
[31]	Khan et al. (2025)	✓	✓	✓	✓	✓	5/5
[32]	Mayer et al. (2021)	✓	✓	✓	✓	✓	5/5
[33]	Zhang et al. (2022)	✓	✓	✓	△	✓	4/5
[34]	Phan & Ngo (2023)	✓	✓	✓	✓	✓	5/5
[35]	Mohammadi et al. (2025)	✓	✓	✓	✓	✓	5/5
[36]	Terven (2025)	✓	△	△	△	✓	3/5
[37]	Sola et al. (2022)	✓	✓	✓	✓	✓	5/5
[38]	Jin et al. (2025)	✓	△	△	△	✓	3/5
[39]	Li (2022)	✓	△	△	△	✓	3/5
[40]	Yuan et al. (2024)	✓	✓	✓	✓	✓	5/5
[41]	Wang et al. (2025)	✓	✓	✓	✓	✓	5/5
[42]	Chen et al. (2024)	✓	△	△	△	✓	3/5
[43]	Rahanu et al. (2021)	✓	△	✓	✓	✓	4/5
[44]	Martinho et al. (2021)	✓	✓	✓	✓	✓	5/5
[45]	Leikas et al. (2019)	✓	△	✓	✓	✓	4/5
[46]	Hashmi et al. (2024)	✓	✓	✓	✓	✓	5/5
[47]	Clark et al. (2025)	✓	✓	✓	✓	✓	5/5
[48]	Rădulescu et al. (2020)	✓	✓	✓	✓	✓	5/5
[49]	Kallenberg et al. (2025)	✓	✓	✓	✓	✓	5/5
[50]	Page et al. (2021)	✓	✓	✓	✓	✓	5/5
[51]	Moola et al. (2015)	✓	✓	✓	✓	✓	5/5
[52]	Yu et al. (2020)	✓	✓	✓	✓	✓	5/5
[53]	Dervis (2019)	✓	✓	✓	✓	✓	5/5
[54]	Singh et al. (2021)	✓	✓	✓	✓	✓	5/5
[55]	Semeraro et al. (2023)	✓	✓	✓	✓	✓	5/5
[56]	Keshvarparast et al. (2024)	✓	✓	✓	✓	✓	5/5
[57]	Sizo et al. (2025)	✓	✓	✓	△	✓	4/5

Table A1. Cont.

No.	Study	Clear Research Aims	Detailed Methodology Description	Comprehensive Results Reporting	Appropriate Data Analysis	Relevant Conclusions	Overall Quality Score
[58]	Marzi et al. (2025)	✓	✓	✓	✓	✓	5/5
[59]	Mancin et al. (2024)	✓	✓	✓	✓	✓	5/5
[60]	Nezameslami et al. (2025)	✓	✓	✓	✓	✓	5/5
[61]	Sarkis-Onofre et al. (2021)	✓	✓	✓	✓	✓	5/5
[62]	Moher et al. (2009)	✓	✓	✓	✓	✓	5/5
[63]	Samal (2025)	✓	✓	✓	✓	✓	5/5
[64]	Ouhadi et al. (2024)	✓	✓	✓	✓	✓	5/5
[65]	Aromataris & Pearson (2014)	✓	✓	✓	✓	✓	5/5
[66]	Ahn & Kang (2018)	✓	✓	✓	✓	✓	5/5
[67]	Liu et al. (2022)	✓	✓	✓	✓	✓	5/5
[68]	Khdoudi et al. (2024)	✓	✓	✓	✓	✓	5/5
[69]	Wells & Bednarz (2021)	✓	✓	✓	✓	✓	5/5
[70]	Su et al. (2025)	✓	✓	✓	✓	✓	5/5
[71]	Liu (2025)	✓	✓	✓	✓	✓	5/5
[72]	Liengpunsakul (2021)	✓	△	✓	✓	✓	4/5
[73]	Jung & Cho (2025)	✓	✓	✓	✓	✓	5/5
[74]	Cheng & Zeng (2023)	✓	△	✓	✓	✓	4/5
[75]	Rajamallaiah et al. (2025)	✓	✓	✓	✓	✓	5/5
[76]	Gautam (2023)	✓	✓	✓	✓	✓	5/5
[77]	Li et al. (2023)	✓	✓	✓	✓	✓	5/5
[78]	Khurram et al. (2025)	✓	✓	✓	✓	✓	5/5
[79]	Hagendorff (2021)	✓	△	✓	✓	✓	4/5
[80]	Ilahi et al. (2021)	✓	✓	✓	✓	✓	5/5
[81]	Zhang et al. (2022)	✓	✓	✓	✓	✓	5/5
[82]	Sun et al. (2021)	✓	✓	✓	✓	✓	5/5
[83]	Bouali et al. (2025)	✓	✓	✓	✓	✓	5/5
[84]	Noriega et al. (2025)	✓	✓	✓	✓	✓	5/5
[85]	Xiao et al. (2022)	✓	✓	✓	✓	✓	5/5
[86]	Chen et al. (2020)	✓	✓	✓	✓	✓	5/5
[87]	Wei et al. (2023)	✓	✓	✓	✓	✓	5/5
[88]	Abo-Khalil (2023)	✓	✓	✓	✓	✓	5/5
[89]	Hua et al. (2021)	✓	✓	✓	✓	✓	5/5
[90]	Raziei & Moghaddam (2021)	✓	✓	✓	△	✓	4/5

Scoring: ✓ = Criterion fully met; △ = Criterion partially met; X = Criterion not met; Overall Quality Score: Sum of fully met criteria (5 = highest quality). Note. This table uses a modified version of the JBI Critical Appraisal Checklist with five simplified methodological criteria adapted for this study.

Overall Summary (References 1–90):

- 75.6% of all studies met all quality criteria (68/90 studies)
- 16.7% met 4 out of 5 criteria (15/90 studies)
- 7.8% scored 3 or below (7/90 studies)
- The overall high-quality scores validate the robustness of the literature corpus selected for this systematic review.

References

1. Asis, K.D.; Hernandez-Garcia, J.F.; Holland, G.Z.; Sutton, R.S. Multi-step reinforcement learning: A unifying algorithm. In Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence and Thirtieth Innovative Applications of Artificial Intelligence Conference and Eighth AAAI Symposium on Educational Advances in Artificial Intelligence, New Orleans, LA, USA, 2–7 February 2018; p. 354.
2. Farooq, A.; Iqbal, K. A Survey of Reinforcement Learning for Optimization in Automation. In Proceedings of the 2024 IEEE 20th International Conference on Automation Science and Engineering (CASE), Bari, Italy, 28 August–1 September 2024; pp. 2487–2494. [CrossRef]
3. Xia, K.; Sacco, C.; Kirkpatrick, M.; Saidy, C.; Nguyen, L.; Kircaliali, A.; Harik, R. A digital twin to train deep reinforcement learning agent for smart manufacturing plants: Environment, interfaces and intelligence. *J. Manuf. Syst.* **2021**, *58*, 210–230. [CrossRef]
4. Ryalat, M.; ElMoaqet, H.; AlFaouri, M. Design of a Smart Factory Based on Cyber-Physical Systems and Internet of Things towards Industry 4.0. *Appl. Sci.* **2023**, *13*, 2156. [CrossRef]
5. del Real Torres, A.; Andreiana, D.S.; Ojeda Roldán, Á.; Hernández Bustos, A.; Acevedo Galicia, L.E. A Review of Deep Reinforcement Learning Approaches for Smart Manufacturing in Industry 4.0 and 5.0 Framework. *Appl. Sci.* **2022**, *12*, 12377. [CrossRef]
6. Langás, E.F.; Zafar, M.H.; Sanfilippo, F. Exploring the synergy of human-robot teaming, digital twins, and machine learning in Industry 5.0: A step towards sustainable manufacturing. *J. Intell. Manuf.* **2025**, *89*, 102769. [CrossRef]
7. Moos, J.; Hansel, K.; Abdulsamad, H.; Stark, S.; Clever, D.; Peters, J. Robust Reinforcement Learning: A Review of Foundations and Recent Advances. *Mach. Learn. Knowl. Extr.* **2022**, *4*, 276–315. [CrossRef]
8. Arulkumaran, K.; Deisenroth, M.P.; Brundage, M.; Bharath, A.A. Deep Reinforcement Learning: A Brief Survey. *IEEE Signal Process. Mag.* **2017**, *34*, 26–38. [CrossRef]
9. Han, D.; Mulyana, B.; Stankovic, V.; Cheng, S. A Survey on Deep Reinforcement Learning Algorithms for Robotic Manipulation. *Sensors* **2023**, *23*, 3762. [CrossRef]
10. Kober, J.; Bagnell, J.A.; Peters, J. Reinforcement learning in robotics: A survey. *Int. J. Robot. Res.* **2013**, *32*, 1238–1274. [CrossRef]
11. Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; Klimov, O. Proximal policy optimization algorithms. *arXiv* **2017**, arXiv:1707.06347. [CrossRef]
12. Hou, Z.; Zhang, K.; Wan, Y.; Li, D.; Fu, C.; Yu, H. Off-policy maximum entropy reinforcement learning: Soft actor-critic with advantage weighted mixture policy (sac-awmp). *arXiv* **2020**, arXiv:2002.02829. [CrossRef]
13. Sathya, D.; Saravanan, G.; Thangamani, R. Reinforcement Learning for Adaptive Mechatronics Systems. In *Computational Intelligent Techniques in Mechatronics*; Wiley: Hoboken, NJ, USA, 2024; pp. 135–184.
14. Rajasekhar, N.; Radhakrishnan, T.K.; Samsudeen, N. Exploring reinforcement learning in process control: A comprehensive survey. *Int. J. Syst. Sci.* **2025**, *56*, 3528–3557. [CrossRef]
15. Hady, M.A.; Hu, S.; Pratama, M.; Cao, Z.; Kowalczyk, R. Multi-agent reinforcement learning for resources allocation optimization: A survey. *Artif. Intell. Rev.* **2025**, *58*, 354. [CrossRef]
16. Jia, L.; Pei, Y. Recent Advances in Multi-Agent Reinforcement Learning for Intelligent Automation and Control of Water Environment Systems. *Machines* **2025**, *13*, 503. [CrossRef]
17. Abbas, A.N.; Amazu, C.W.; Mietkiewicz, J.; Briwa, H.; Perez, A.A.; Baldissone, G.; Demichela, M.; Chasparis, G.C.; Kelleher, J.D.; Leva, M.C. Analyzing Operator States and the Impact of AI-Enhanced Decision Support in Control Rooms: A Human-in-the-Loop Specialized Reinforcement Learning Framework for Intervention Strategies. *Int. J. Hum. Comput. Interact.* **2025**, *41*, 7218–7252. [CrossRef]
18. Schwung, D.; Reimann, J.N.; Schwung, A.; Ding, S.X. Self Learning in Flexible Manufacturing Units: A Reinforcement Learning Approach. In Proceedings of the 2018 International Conference on Intelligent Systems (IS), Madeira, Portugal, 25–27 September 2018; pp. 31–38. [CrossRef]
19. Chang, X.; Jia, X.; Fu, S.; Hu, H.; Liu, K. Digital twin and deep reinforcement learning enabled real-time scheduling for complex product flexible shop-floor. *Proc. Inst. Mech. Eng. Part B J. Eng. Manuf.* **2022**, *237*, 1254–1268. [CrossRef]
20. Emami, Y.; Almeida, L.; Li, K.; Ni, W.; Han, Z. Human-in-the-loop machine learning for safe and ethical autonomous vehicles: Principles, challenges, and opportunities. *arXiv* **2024**, arXiv:2408.12548. [CrossRef]
21. Hernandez-Leal, P.; Kartal, B.; Taylor, M.E. A survey and critique of multiagent deep reinforcement learning. *Auton. Agents Multi-Agent Syst.* **2019**, *33*, 750–797. [CrossRef]
22. Rolf, B.; Jackson, I.; Müller, M.; Lang, S.; Reggelen, T.; Ivanov, D. A review on reinforcement learning algorithms and applications in supply chain management. *Int. J. Prod. Res.* **2023**, *61*, 7151–7179. [CrossRef]
23. Martins, M.S.E.; Sousa, J.M.C.; Vieira, S. A Systematic Review on Reinforcement Learning for Industrial Combinatorial Optimization Problems. *Appl. Sci.* **2025**, *15*, 1211. [CrossRef]
24. Xu, R.; Peng, J. A Comprehensive Survey of Deep Research: Systems, Methodologies, and Applications. *arXiv* **2025**, arXiv:2506.12594. [CrossRef]
25. Chen, W.; Qiu, X.; Cai, T.; Dai, H.N.; Zheng, Z.; Zhang, Y. Deep Reinforcement Learning for Internet of Things: A Comprehensive Survey. *IEEE Commun. Surv. Tutor.* **2021**, *23*, 1659–1692. [CrossRef]

26. Eckardt, J.N.; Wendt, K.; Bornhäuser, M.; Middeke, J.M. Reinforcement Learning for Precision Oncology. *Cancers* **2021**, *13*, 4624. [CrossRef] [PubMed]
27. Paranjape, A.; Quader, N.; Uhlmann, L.; Berkels, B.; Wolfschläger, D.; Schmitt, R.H.; Bergs, T. Reinforcement Learning Agent for Multi-Objective Online Process Parameter Optimization of Manufacturing Processes. *Appl. Sci.* **2025**, *15*, 7279. [CrossRef]
28. Peta, K.; Wiśniewski, M.; Kotarski, M.; Ciszak, O. Comparison of Single-Arm and Dual-Arm Collaborative Robots in Precision Assembly. *Appl. Sci.* **2025**, *15*, 2976. [CrossRef]
29. Estes, A.; Peidro, D.; Mula, J.; Díaz-Madroñero, M. Reinforcement learning applied to production planning and control. *Int. J. Prod. Res.* **2023**, *61*, 5772–5789. [CrossRef]
30. Khan, M.N.; Ullah, I.; Bashir, A.K.; Al-Khasawneh, M.A.; Arishi, A.; Alghamdi, N.S.; Lee, S. Reinforcement Learning-Based Dynamic Power Management for Energy Optimization in IoT-Enabled Consumer Electronics. *IEEE Trans. Consum. Electron.* **2025**, *71*, 8103–8114. [CrossRef]
31. Khan, M.N.; Lee, S.; Shah, M. Adaptive Scheduling in Cognitive IoT Sensors for Optimizing Network Performance Using Reinforcement Learning. *Appl. Sci.* **2025**, *15*, 5573. [CrossRef]
32. Phan, L.A.D.; Ngo, H.Q.T. Systematic Review of Smart Robotic Manufacturing in the Context of Industry 4.0. In *Context-Aware Systems and Applications, Proceedings of the 12th EAI International Conference, ICCASA 2023, Ho Chi Minh City, Vietnam, 26–27 October 2023*; Vinh, P.C., Tung, N.T., Eds.; Springer Nature: Cham, Switzerland, 2024; pp. 19–42. [CrossRef]
33. Mayer, S.; Classen, T.; Endisch, C. Modular production control using deep reinforcement learning: Proximal policy optimization. *J. Intell. Manuf.* **2021**, *32*, 2335–2351. [CrossRef]
34. Zhang, M.; Lu, Y.; Hu, Y.; Amaitik, N.; Xu, Y. Dynamic Scheduling Method for Job-Shop Manufacturing Systems by Deep Reinforcement Learning with Proximal Policy Optimization. *Sustainability* **2022**, *14*, 5177. [CrossRef]
35. Mohammadi, E.; Ortiz-Arroyo, D.; Hansen, A.A.; Stokholm-Bjerregaard, M.; Gros, S.; Anand, A.S.; Durdevic, P. Application of Soft Actor-Critic algorithms in optimizing wastewater treatment with time delays integration. *Expert Syst. Appl.* **2025**, *277*, 127180. [CrossRef]
36. Terven, J. Deep Reinforcement Learning: A Chronological Overview and Methods. *AI* **2025**, *6*, 46. [CrossRef]
37. Sola, Y.; Le Chenadec, G.; Clement, B. Simultaneous Control and Guidance of an AUV Based on Soft Actor–Critic. *Sensors* **2022**, *22*, 6072. [CrossRef]
38. Jin, W.; Du, H.; Zhao, B.; Tian, X.; Shi, B.; Yang, G. A comprehensive survey on multi-agent cooperative decision-making: Scenarios, approaches, challenges and perspectives. *arXiv* **2025**, arXiv:2503.13415. [CrossRef]
39. Li, Y. Reinforcement learning in practice: Opportunities and challenges. *arXiv* **2022**, arXiv:2202.11296. [CrossRef]
40. Yuan, Z.; Zhang, Z.; Li, X.; Cui, Y.; Li, M.; Ban, X. Controlling Partially Observed Industrial System Based on Offline Reinforcement Learning—A Case Study of Paste Thickener. *IEEE Trans. Ind. Inform.* **2025**, *21*, 49–59. [CrossRef]
41. Wang, T.; Ruan, Z.; Wang, Y.; Chen, C. Control strategy of robotic manipulator based on multi-task reinforcement learning. *Complex Intell. Syst.* **2025**, *11*, 175. [CrossRef]
42. Chen, J.; He, J.; Chen, F.; Lv, Z.; Tang, J.; Li, W.; Liu, Z.; Yang, H.H.; Han, G. Towards General Industrial Intelligence: A Survey of Continual Large Models in Industrial IoT. *arXiv* **2024**, arXiv:2409.01207. [CrossRef]
43. Rahanu, H.; Georgiadou, E.; Siakas, K.; Ross, M.; Berki, E. Ethical Issues Invoked by Industry 4.0. In *Systems, Software and Services Process Improvement, Proceedings of the 28th European Conference, EuroSPI 2021, Krems, Austria, 1–3 September 2021*; Yilmaz, M., Clarke, P., Messnarz, R., Reiner, M., Eds.; Springer International Publishing: Cham, Switzerland, 2021; pp. 589–606. [CrossRef]
44. Martinho, A.; Herber, N.; Kroesen, M.; Chorus, C. Ethical issues in focus by the autonomous vehicles industry. *Transp. Rev.* **2021**, *41*, 556–577. [CrossRef]
45. Leikas, J.; Koivisto, R.; Gotcheva, N. Ethical Framework for Designing Autonomous Intelligent Systems. *J. Open Innov. Technol. Mark. Complex.* **2019**, *5*, 18. [CrossRef]
46. Hashmi, R.; Liu, H.; Yavari, A. Digital Twins for Enhancing Efficiency and Assuring Safety in Renewable Energy Systems: A Systematic Literature Review. *Energies* **2024**, *17*, 2456. [CrossRef]
47. Clark, L.; Garcia, B.; Harris, S. Digital Twin-Driven Operational Management Framework for Real-Time Decision-Making in Smart Factories. *J. Innov. Gov. Bus. Pract.* **2025**, *1*, 59–87.
48. Rădulescu, R.; Mannion, P.; Roijers, D.M.; Nowé, A. Multi-objective multi-agent decision making: A utility-based analysis and survey. *Auton. Agents Multi-Agent Syst.* **2019**, *34*, 10. [CrossRef]
49. Kallenberg, M.; Baja, H.; Ilić, M.; Tomčić, A.; Tošić, M.; Athanasiadis, I. Interoperable agricultural digital twins with reinforcement learning intelligence. *Smart Agric. Technol.* **2025**, *12*, 101412. [CrossRef]
50. Page, M.J.; McKenzie, J.E.; Bossuyt, P.M.; Boutron, I.; Hoffmann, T.C.; Mulrow, C.D.; Shamseer, L.; Tetzlaff, J.M.; Akl, E.A.; Brennan, S.E.; et al. The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ* **2021**, *372*, n71. [CrossRef] [PubMed]
51. Moola, S.; Munn, Z.; Sears, K.; Sfetcu, R.; Currie, M.; Lisy, K.; Tufanaru, C.; Qureshi, R.; Mattis, P.; Mu, P. Conducting systematic reviews of association (etiology): The Joanna Briggs Institute’s approach. *JBI Evid. Implement.* **2015**, *13*, 163–169. [CrossRef]
52. Yu, D.; Xu, Z.; Wang, X. Bibliometric analysis of support vector machines research trend: A case study in China. *Int. J. Mach. Learn. Cybern.* **2020**, *11*, 715–728. [CrossRef]
53. Derviş, H. Bibliometric analysis using bibliometrix an R package. *J. Scientometr. Res.* **2019**, *8*, 156–160. [CrossRef]

54. Singh, V.K.; Singh, P.; Karmakar, M.; Leta, J.; Mayr, P. The journal coverage of Web of Science, Scopus and Dimensions: A comparative analysis. *Scientometrics* **2021**, *126*, 5113–5142. [CrossRef]
55. Semeraro, F.; Griffiths, A.; Cangelosi, A. Human–robot collaboration and machine learning: A systematic review of recent research. *Robot. Comput.-Integr. Manuf.* **2023**, *79*, 102432. [CrossRef]
56. Keshvarparast, A.; Battini, D.; Battaia, O.; Pirayesh, A. Collaborative robots in manufacturing and assembly systems: Literature review and future research agenda. *J. Intell. Manuf.* **2024**, *35*, 2065–2118. [CrossRef]
57. Sizo, A.; Lino, A.; Rocha, Á.; Reis, L.P. Defining quality in peer review reports: A scoping review. *Knowl. Inf. Syst.* **2025**, *67*, 6413–6460. [CrossRef]
58. Marzi, G.; Balzano, M.; Caputo, A.; Pellegrini, M.M. Guidelines for Bibliometric-Systematic Literature Reviews: 10 steps to combine analysis, synthesis and theory development. *Int. J. Manag. Rev.* **2025**, *27*, 81–103. [CrossRef]
59. Mancin, S.; Sguanci, M.; Andreoli, D.; Soekeland, F.; Anastasi, G.; Piredda, M.; De Marinis, M.G. Systematic review of clinical practice guidelines and systematic reviews: A method for conducting comprehensive analysis. *MethodsX* **2024**, *12*, 102532. [CrossRef] [PubMed]
60. Nezameslami, R.; Nezameslami, A.; Mehdikhani, B.; Mosavi-Jarrahi, A.; Shahbazi, A.; Rahmani, A.; Masoudi, A.; Yeganegi, M.; Akhondzardaini, R.; Bahrami, M.; et al. Adapting PRISMA Guidelines to Enhance Reporting Quality in Genetic Association Studies: A Framework Proposal. *Asian Pac. J. Cancer Prev.* **2025**, *26*, 1641–1651. [CrossRef] [PubMed]
61. Sarkis-Onofre, R.; Catalá-López, F.; Aromataris, E.; Lockwood, C. How to properly use the PRISMA Statement. *Syst. Rev.* **2021**, *10*, 117. [CrossRef]
62. Moher, D.; Liberati, A.; Tetzlaff, J.; Altman, D.G.; The, P.G. Preferred Reporting Items for Systematic Reviews and Meta-Analyses: The PRISMA Statement. *PLoS Med.* **2009**, *6*, e1000097. [CrossRef]
63. Samal, U. Evolution of machine learning and deep learning in intelligent manufacturing: A bibliometric study. *Int. J. Syst. Assur. Eng. Manag.* **2025**, *16*, 3134–3150. [CrossRef]
64. Ouhadi, A.; Yahouni, Z.; Di Mascolo, M. Integrating machine learning and operations research methods for scheduling problems: A bibliometric analysis and literature review. *IFAC-PapersOnLine* **2024**, *58*, 946–951. [CrossRef]
65. Aromataris, E.; Pearson, A. The Systematic Review: An Overview. *Am. J. Nurs.* **2014**, *114*, 53–58. [CrossRef]
66. Ahn, E.; Kang, H. Introduction to systematic review and meta-analysis. *Korean J. Anesth.* **2018**, *71*, 103–112. [CrossRef]
67. Liu, Y.; Xu, H.; Liu, D.; Wang, L. A digital twin-based sim-to-real transfer for deep reinforcement learning-enabled industrial robot grasping. *Robot. Comput.-Integr. Manuf.* **2022**, *78*, 102365. [CrossRef]
68. Khoudoudi, A.; Masrouf, T.; El Hassani, I.; El Mazgualdi, C. A Deep-Reinforcement-Learning-Based Digital Twin for Manufacturing Process Optimization. *Systems* **2024**, *12*, 38. [CrossRef]
69. Wells, L.; Bednarz, T. Explainable AI and Reinforcement Learning—A Systematic Review of Current Approaches and Trends. *Front. Artif. Intell.* **2021**, *4*, 550030. [CrossRef] [PubMed]
70. Su, T.; Wu, T.; Zhao, J.; Scaglione, A.; Xie, L. A Review of Safe Reinforcement Learning Methods for Modern Power Systems. *Proc. IEEE* **2025**, *113*, 213–255. [CrossRef]
71. Liu, Y. Bridging research and policy in China’s energy sector: A semantic and reinforcement learning framework. *Energy Strategy Rev.* **2025**, *59*, 101770. [CrossRef]
72. Liengpunsakul, S. Artificial Intelligence and Sustainable Development in China. *Chin. Econ.* **2021**, *54*, 235–248. [CrossRef]
73. Jung, W.; Cho, K. Deep Learning and NLP-Based Trend Analysis in Actuators and Power Electronics. *Actuators* **2025**, *14*, 379. [CrossRef]
74. Cheng, J.; Zeng, J. Shaping AI’s Future? China in Global AI Governance. *J. Contemp. China* **2023**, *32*, 794–810. [CrossRef]
75. Rajamallaiiah, A.; Naresh, S.V.K.; Raghuvamsi, Y.; Manmadharao, S.; Bingi, K.; Anand, R.; Guerrero, J.M. Deep Reinforcement Learning for Power Converter Control: A Comprehensive Review of Applications and Challenges. *IEEE Open J. Power Electron.* **2025**, *6*, 1769–1802. [CrossRef]
76. Gautam, M. Deep Reinforcement Learning for Resilient Power and Energy Systems: Progress, Prospects, and Future Avenues. *Electricity* **2023**, *4*, 336–380. [CrossRef]
77. Li, C.; Bai, L.; Yao, L.; Waller, S.T.; Liu, W. A bibliometric analysis and review on reinforcement learning for transportation applications. *Transp. B Transp. Dyn.* **2023**, *11*, 2179461. [CrossRef]
78. Khurram, M.; Zhang, C.; Muhammad, S.; Kishnani, H.; An, K.; Abeywardena, K.; Chadha, U.; Behdinan, K. Artificial Intelligence in Manufacturing Industry Worker Safety: A New Paradigm for Hazard Prevention and Mitigation. *Processes* **2025**, *13*, 1312. [CrossRef]
79. Hagendorff, T. Forbidden knowledge in machine learning reflections on the limits of research and publication. *AI Soc.* **2021**, *36*, 767–781. [CrossRef]
80. Ilahi, I.; Usama, M.; Qadir, J.; Janjua, M.U.; Al-Fuqaha, A.; Hoang, D.T.; Niyato, D. Challenges and Countermeasures for Adversarial Attacks on Deep Reinforcement Learning. *IEEE Trans. Artif. Intell.* **2022**, *3*, 90–109. [CrossRef]
81. Zhang, W.; Gu, X.; Tang, L.; Yin, Y.; Liu, D.; Zhang, Y. Application of machine learning, deep learning and optimization algorithms in geoengineering and geoscience: Comprehensive review and future challenge. *Gondwana Res.* **2022**, *109*, 1–17. [CrossRef]
82. Sun, H.; Zhang, W.; Yu, R.; Zhang, Y. Motion Planning for Mobile Robots—Focusing on Deep Reinforcement Learning: A Systematic Review. *IEEE Access* **2021**, *9*, 69061–69081. [CrossRef]

83. Bouali, M.; Sebaa, A.; Farhi, N. A Comprehensive Review on Reinforcement Learning Methods for Autonomous Lane Changing. *Int. J. Intell. Transp. Syst. Res.* **2025**. [CrossRef]
84. Noriega, R.; Pourrahimian, Y.; Askari-Nasab, H. Deep Reinforcement Learning based real-time open-pit mining truck dispatching system. *Comput. Oper. Res.* **2025**, *173*, 106815. [CrossRef]
85. Xiao, X.; Liu, B.; Warnell, G.; Stone, P. Motion planning and control for mobile robot navigation using machine learning: A survey. *Auton. Robot.* **2022**, *46*, 569–597. [CrossRef]
86. Chen, J.; Lin, C.; Peng, D.; Ge, H. Fault Diagnosis of Rotating Machinery: A Review and Bibliometric Analysis. *IEEE Access* **2020**, *8*, 224985–225003. [CrossRef]
87. Wei, Z.; Liu, H.; Tao, X.; Pan, K.; Huang, R.; Ji, W.; Wang, J. Insights into the Application of Machine Learning in Industrial Risk Assessment: A Bibliometric Mapping Analysis. *Sustainability* **2023**, *15*, 6965. [CrossRef]
88. Abo-Khalil, A.G. Digital twin real-time hybrid simulation platform for power system stability. *Case Stud. Therm. Eng.* **2023**, *49*, 103237. [CrossRef]
89. Hua, J.; Zeng, L.; Li, G.; Ju, Z. Learning for a Robot: Deep Reinforcement Learning, Imitation Learning, Transfer Learning. *Sensors* **2021**, *21*, 1278. [CrossRef]
90. Raziei, Z.; Moghaddam, M. Enabling adaptable Industry 4.0 automation with a modular deep reinforcement learning framework. *IFAC-PapersOnLine* **2021**, *54*, 546–551. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Overview on Predictive Maintenance Techniques for Turbomachinery

Pierpaolo Dini *, Damiano Nardi and Sergio Saponara

Department of Information Engineering, University of Pisa, Via G. Caruso 15, 55122 Pisa, Italy; damiano.nardi@ing.unipi.it (D.N.); sergio.saponara@unipi.it (S.S.)

* Correspondence: pierpaolo.dini@unipi.it

Abstract

Within the Industry 5.0 paradigm, the management of critical assets requires advanced digital architectures capable of ensuring resilience and operational sustainability. The present systematic review analyzes the state of the art in predictive maintenance (PdM) technologies for turbines and turbomachinery, providing a technical examination of anomaly and fault detection frameworks, extended to remaining useful life (RUL) estimation and root cause analysis (RCA). The work addresses inherent sectoral challenges, ranging from the processing of high-dimensional multivariate time series (MTS) from Supervisory Control and Data Acquisition (SCADA) systems to labeled data scarcity and signal non-stationarity in real-world environments. Both purely data-driven frameworks and hybrid physics-informed models, such as Physics-Informed Neural Networks (PINNs), are critically evaluated against performance indicators. A significant contribution of this study lies in the classification of methodologies based on their readiness for real-time inference, emphasizing the role of Explainable AI (XAI) in providing transparent insights to domain experts, who remain central to decision-making processes. The primary objective of this review is to offer an analytical overview of progress to date against current technological gaps, tracing a clear trajectory for future developments. In this regard, the adoption of Generative AI and Large Language Models (LLMs) is identified as a fundamental step toward evolving into interactive, human-centric decision support systems.

Keywords: predictive maintenance; turbomachinery; anomaly detection; deep learning; physics-informed neural networks; remaining useful life; root cause analysis; SCADA data; Industry 5.0

1. Introduction

1.1. Motivations

In recent years, Predictive Maintenance (PdM) has increasingly taken hold in the manufacturing landscape, particularly within the turbomachinery sector. This field has evolved from manual vibration monitoring toward advanced digital architectures based on artificial intelligence [1,2]. Such technological progression currently aligns with Industry 5.0 principles, which incorporate human-centricity, sustainability, and resilience to foster adaptive production processes and minimize environmental impact [3].

PdM has emerged as a superior solution to address the primary challenges posed by equipment failures, which can lead to significant financial losses, operational disruptions, and damage to a company's reputation [4]. By utilizing real-time data and advanced analytics to monitor equipment conditions continuously, such a predictive approach enables the

early detection of potential issues, allowing for timely and precise interventions. The proactive nature of PdM minimizes unplanned downtime, extends asset lifespans, and reduces wasteful maintenance spending [5].

The implementation of PdM in modern power generation facilities and industrial plants has demonstrated substantial operational benefits. For instance, the authors in [6] demonstrate the advantages of transitioning from condition-based maintenance (CBM) to a fully predictive maintenance strategy, showing significant reductions in operational downtime (70–75%) and maintenance costs (20–50%) for wind turbines. In other industrial contexts, Romanssini et al. [7] demonstrate that implementing PdM on rotating machinery yields reductions in operational downtime up to 50%, decreased maintenance costs (8–12% compared to conventional preventive maintenance), extended asset lifespan, improved operator safety, and enhanced overall plant reliability.

However, applying these strategies to turbomachinery presents unique challenges. The implementation of a PdM strategy begins with the analysis of the anomalies and failures that characterize the system. In turbomachinery, the main issues include metallurgical degradations such as creep, High Cycle Fatigue (HCF), and Low Cycle Fatigue (LCF), alongside corrosion and erosion [8]. An equally major category consists of rotordynamic anomalies, which primarily manifest as mass unbalance, shaft misalignment, and rubbing phenomena (involuntary rotor–stator interaction) [2,9]. Additionally, mechanical degradation of transmission components and fouling phenomena [10] are frequently observed, potentially accelerating long-term failures. In turbomachinery, the onset of such failures is often tracked by monitoring the condition of bearings, including Rolling (RB), Journal (JB), and Thrust (TB) types, in terms of vibrations and extreme temperatures [11–13]. Being wear components, they are assessed during both scheduled inspections and CBM strategies [14].

The evolution toward decarbonization imposes new operational challenges: gas turbines must adapt to low-carbon fuels (hydrogen) to back up renewables [15,16], while wind turbines face variable environmental conditions [17]. Such a shift to flexible and cyclic operation accelerates degradation mechanisms, reducing overall efficiency [16,18,19]. These evolving operational requirements highlight the need for advanced monitoring and prediction capabilities integrated within a Digital Twin (DT) [1,20,21]. Building on these foundations, the DT integrates AI, Machine Learning, and advanced analytics to enable learning from multi-source data, adaptive behavior under changing operating conditions, and support for predictive decision-making [20].

1.2. Background

The success of PdM stems from its unique ability to balance cost efficiency with operational reliability, enabled by the convergence of fundamental technologies: (1) IoT (Internet of Things): Sensors enable massive equipment data collection; (2) Big Data: Processes raw IoT data into actionable insights; (3) Deep Learning: Powers fault diagnosis and Remaining Useful Life (RUL) prediction; (4) Deep Reinforcement Learning (DRL): Optimizes maintenance decisions in dynamic environments; (5) Hardware: GPUs/TPUs accelerate complex model training [22,23]. Despite persistent implementation challenges for PdM across industries, it remains the most cost-effective approach for minimizing downtime and failure-related expenses compared to reactive maintenance (RM) and preventive maintenance (PM) [24].

From a data analysis perspective, signals collected via SCADA systems constitute complex Multivariate Time Series (MTS) [25]. One of the major challenges in this domain lies in the scarcity of labeled data [26,27]. In fact, many faults are extremely rare by design, especially in high-safety applications [28]. This is compounded by the difficulty of identifying the exact moment of the initial anomaly and by variations in operating

conditions (domain shifts), which lead to false alarms if the model is not robust [29]; finally, obtaining reliable ground truth is expensive and often impossible [30].

To overcome these obstacles, the PdM pipeline has shifted toward Deep Learning processes and hybrid approaches, which are necessary to manage long-term temporal dependencies and enable unsupervised learning. In this scenario, Anomaly Detection assumes a priority role over Fault Detection: while Fault Detection aims to isolate a specific fault requiring prior knowledge, Anomaly Detection identifies deviations from nominal behavior without needing labels, ensuring the timeliness essential for preventing catastrophic failures. The study by Sepe et al. [9] exemplifies this hybrid approach, accurately predicting creep damage in Baker Hughes turbomachinery with a 28.5% error margin relative to the failure interval, enabling the estimation of Remaining Useful Life (RUL) specific to each turbine.

Although advanced commercial frameworks for lifecycle monitoring exist (e.g., GE Vernova's SmartSignal [31], Siemens's ATOM [32], and EPRI's Digital Twin [33]), focusing on real-time diagnostics and predictive maintenance, the analysis in this review will cover both purely data-driven techniques (AI/ML) and hybrid approaches like Physics-Informed Neural Networks (PINNs) and Physics-Informed Machine Learning (PIML) [18]. The latter will be examined in the context of accelerating computational simulations (CFD, FEM), estimating parameters using a priori physical knowledge, and integrating mixed models without strict physical constraints [16,34].

1.3. Contributions

Existing surveys either focus purely on data-driven methods for anomaly detection or on prognostics or are limited to specific turbine types. A holistic, comparative review that integrates anomaly detection, RUL, and RCA across turbine types and critically evaluates the role of physics-informed AI is currently missing. Therefore, the objective of this work is not to propose a unique PdM framework but rather to provide a systematic state-of-the-art review of the fundamental methodologies that constitute its backbone: Anomaly Detection, Fault Detection, RUL estimation, and Root Cause Analysis (RCA). The main contribution of this survey is threefold. First, it provides a structured taxonomy and a comparative in-depth analysis of modern architectures for Anomaly and Fault Detection, categorizing them into two main families: purely Deep Learning methods and hybrid-physics (data-driven and physics-informed) approaches. Second, it separately and in-depth analyzes recent advancements in the other two pillars of PdM: RUL estimation and RCA, applying this same cross-cutting dichotomy. Finally, the work discusses open challenges and future perspectives, identifying research gaps and outlining promising directions, such as the synergistic integration of Large Language Models (LLM) and Human-in-the-Loop (HITL) [35] approaches to bridge the semantic gap between AI model outputs and operational decisions, in full adherence to Industry 5.0 paradigms [36].

To provide a clearer overview of the study's structure, a visual organizer of this review is presented in Figure 1.

The remainder of this paper is organized as follows: Section 2 presents the review methodology. Section 3 introduces a taxonomy of Deep Learning architectures for anomaly detection. Section 4 evaluates anomaly and fault detection methods (forecasting, reconstruction, contrastive, and physics-informed). Section 5 covers RUL estimation techniques. Section 6 reviews RCA approaches. Section 7 provides a comparative analysis. Section 8 discusses challenges and concludes.




Figure 1. Structural roadmap and organizer of the review paper.

2. Review Methodology

The typical PdM pipeline consists of several layers. One of the core stages, performed immediately after data acquisition and pre-processing, is anomaly detection [1,20]. The evolution of anomaly detection has largely moved beyond traditional machine learning techniques, increasingly shifting toward deep learning processes. The transition is primarily driven by distinct advantages, such as the ability to effectively handle long-term temporal dependencies (e.g., via LSTM, GRU, or Transformer architectures), capture non-linear relationships between variables, enable label-free learning (unsupervised or self-supervised approaches), and ensure robustness against noise and non-stationarity [37]. From a practical perspective, these aspects are vital. Indeed, data collected via acquisition units and sensors distributed across turbomachinery allow for the reception of signals sampled through SCADA systems. Mathematically, these signals constitute evolving time series. Since they often require processing via online inference, it is essential to adopt data-driven models optimized for real-time execution, ensuring that computational requirements are compatible with resource-constrained platforms. The analysis of these signals results in a Multivariate Time Series (MTS), representing heterogeneous physical quantities within turbomachinery monitoring [25]. During MTS analysis, anomalies can manifest in various forms as the data not only evolve temporally but are also interconnected by linear or non-linear relationships between variables. Consequently, a distinction is generally made between two main families of anomalies: intra-metric anomalies, relating to the temporal evolution within a single variable, and inter-metric anomalies, relating to inconsistencies caused by the violation of expected correlations between different metrics [38,39].

In rotating systems characterized by stationary dynamics, anomalies usually do not involve point-wise anomalies (where the anomalous behavior of a single time point may appear as a spike or glitch) [38]. Instead, typical anomalies often involve trend anomalies, where the series suddenly deviates from a long-term trend, or cyclic anomalies, such as the occurrence of a seasonal peak with altered amplitude, frequency, or periodicity, as shown in Figure 2. Furthermore, Figure 3 illustrates the system architecture, which begins with raw

data acquisition and subsequent pre-processing for cleaning and meaningful feature extraction. The resulting stage yields fundamental health indicators, such as time series, function transforms, and monitoring parameters, which feed into the diagnostic (Anomaly and Fault Detection) and prognostic modules for RUL estimation. In this context, RCA provides indispensable decision support for the operator, translating identified deviations into diagnostic information regarding the triggering causes, necessary for precisely planning maintenance interventions on the physical system. The following section addresses condition monitoring and fault detection techniques, with a specific focus on Anomaly Detection applied to turbines and turbomachinery in general, providing an analysis and review of Deep Learning methodologies, demonstrating their effectiveness in this field, and introducing a Deep Learning scheme for time series anomaly detection in the turbomachinery sector.




Figure 2. Visual classification of anomaly families in turbomachinery time series: (a) Point-wise anomalies; (b) Cyclic anomalies; (c) Trend anomalies [38].




Figure 3. Schematic representation of the PdM paradigm. The process comprises Diagnostics: encompassing Anomaly Detection for deviation identification and Fault Detection for cause isolation and Prognostics: which estimates the Remaining Useful Life (RUL) to anticipate functional failure.

Trends in Deep Learning for Turbomachinery Monitoring

Deep Learning is increasingly gaining ground in modern industrial monitoring applications. In fact, relying solely on physical modeling of the system under analysis is becoming increasingly difficult, particularly for systems that are complex in terms of both mathematics and physics. With the availability of a wide range of signals from various acquisition systems, including SCADA logs and raw sensor data, and enabled by modern architectures such as High-Performance Computing (HPC), enabling real-time processing, the analysis of these data streams becomes the fundamental input for system health monitoring. Consequently, in recent years, Anomaly Detection and Fault Analysis have

shifted increasingly from statistical regression techniques (such as AR, ARIMA, ARX, NLARX, or SARIMA) and Machine Learning approaches towards Deep Learning [38,40]. While statistical and regression-based methods prove effective in analyzing single-variable temporal trends or Univariate Time Series (UTS), they struggle to capture the complex correlations among hundreds of distinct sensors simultaneously. Similarly, algorithms such as One-Class Support Vector Machine (OCSVM) and Support Vector Data Description (SVDD) represent valid alternatives; however, traditional Machine Learning methods generally face inherent scalability and performance limitations when handling large, complex multivariate datasets or MTS [37,38,40]. Furthermore, as system complexity increases, these approaches encounter significant challenges driven by the scarcity of labeled anomaly data, as previously discussed [37,38].

In this regard, Zamanzadeh et al. [37] address the issue of data scarcity, noting that generic Machine Learning methods require balanced datasets that rarely exist in the real world. The advantage highlighted here is that Deep Learning enables unsupervised or self-supervised approaches (such as Representation-based methods) that learn the intrinsic characteristics of normal data without the need for fault labels, thereby overcoming the main bottleneck of classical statistical methods, which require a priori definitions. Parallel to this, there is a technical limitation regarding traditional One-class classifiers (such as OCSVM): designed for fixed-dimension data, they struggle to capture temporal dependencies [41–43]. In practice, they treat each data point as an isolated or static event, losing the fundamental information residing in the temporal sequence, which recurrent architectures (e.g., RNNs and LSTMs) or Transformers handle natively [44,45]. Finally, the difficulty in managing unpredictability persists; indeed, statistical regression methods based on anomaly forecasting fail when the system changes rapidly or is intrinsically chaotic. For instance, given a volatile time series, a statistical method will consistently generate a high prediction error, confusing normal behavior with an anomaly. Reconstruction-based Deep Learning methods, conversely, circumvent this issue by analyzing the internal structure of the current data rather than relying solely on its future projection [37,46]. The effectiveness of Deep Learning has been confirmed not only by its high computational performance but also by its increasing prevalence in the scientific literature dedicated to health monitoring, specifically in Anomaly and Fault Detection. In recent years, it has seen pervasive adoption, increasingly replacing classical Machine Learning approaches as well as purely physics-based models or visual inspections, a trend clearly observable within the turbomachinery sector [47]. To support this evidence, a systematic analysis of scientific publications over the last decade (2015–2025) was conducted. The literature acquisition process was executed through an automated pipeline utilizing the public APIs of Semantic Scholar, arXiv, Crossref, Elsevier (Scopus), IEEE Xplore, and OpenAlex. To mitigate potential publication bias and ensure a balanced overview, the search strategy explicitly included pre-print servers (e.g., arXiv) and international conference proceedings. This approach allows for capturing a broader spectrum of results, including early-stage theoretical advancements and industrial case studies that might otherwise be omitted from journal-only syntheses, which are often characterized by a higher frequency of positive outcomes. To ensure a systematic search, 16 targeted queries were used (e.g., “turbomachinery anomaly detection”, “instability detection in combustion systems”), strictly excluding overviews, reviews, and surveys to isolate original contributions only. The search queries concerned not only the detection phase but the entire spectrum of health-monitoring, diagnostics, and prognostics.

The identified works were cataloged into three methodological macro-categories:

- Deep Learning (DL): layered neural architectures capable of autonomously extracting feature hierarchies from complex data, including generative models, recurrent architectures, and hybrid solutions that integrate different deep learning paradigms.

- Machine Learning (ML): supervised and unsupervised learning algorithms based on mathematical models that do not exploit deep structures, focusing on data regression and partitioning aimed at identifying logical patterns or optimal decisions.
- Alternative Methods (AM): including physics-based models, deterministic models, autoregressive models, practical-visual analyses, and direct measurements.

The selection and classification were governed by a rigorous multi-stage filtering process:

- Inclusion Criteria: (i) original research implementing DL, ML, or AM architectures; (ii) specific focus on turbomachinery assets (gas, steam, wind, or hydro turbines and compressors); (iii) English-language full-text availability.
- Exclusion Criteria: (i) non-industrial applications (e.g., medical or clinical studies); (ii) conference abstracts lacking detailed mathematical formulations.

After a deduplication stage based on normalized title matching, 411 unique articles were selected. The assignment to the DL or ML categories was determined by a keyword-based scoring system: architectures featuring multi-layered structures (e.g., CNN, LSTM, PINNs) were assigned to DL, while shallow models (e.g., SVM, Random Forest) were categorized as ML. The results obtained from this analysis are illustrated in Figure 4. The analysis shows that 2020 marked a significant surge in the use of Deep Learning techniques compared to classical Machine Learning. It is also evident that, as of 2017, Machine Learning methods in general, including Deep Learning, have surpassed traditional physics-based approaches. The latter are often computationally too slow, as seen with FEM or CFD simulations, or limited in precision for effective real-time monitoring. This trend aligns with the observations of C. Tsallis et al. [48], who provide a systematic literature review confirming the expansion of AI within the broader PdM sector. Similarly, M. de Castro-Cros et al. [47] offer a detailed overview of machine learning applied specifically to gas turbine condition monitoring. While these studies provide valuable qualitative summaries, our review introduces a distinct methodological advancement. Unlike prior works that categorize literature by algorithm family or component type, we propose a dual-axis taxonomy that characterizes methods by both their functional objectives as pure data-driven approaches and their physics-informed integration paradigms. Moreover, we implement a quantitative benchmarking framework utilizing synthetic indices to evaluate industrial trade-offs, providing an analytical depth not explored in previous surveys.




Figure 4. Bibliometric analysis of the reviewed literature. **(Left)** Historical evolution of publication trends, showing a clear shift towards Deep Learning (DL) for Anomaly Detection (AD), Health Monitoring (HM), and Prognostics (PR) tasks, surpassing traditional Machine Learning (ML) and Alternative Methods (AM) in recent years. **(Right)** Distribution of thematic areas within the PHM framework and breakdown of the bibliographic sources.

The analysis results shown in Figure 4 represent the only works exclusively focused on Anomaly Detection while also considering its integration within a broader framework, such as Prognostics and Health Management (PHM), thereby addressing topics like health monitoring and system prognostics. Furthermore, many studies share common methodologies; this not only leads to attribution of multiple techniques to a single work but also demonstrates that certain scientific branches continue to persist and evolve, albeit to a lesser extent than others. Finally, the research procedure and data extraction stages for the literature review are shown in Figure 5.




Figure 5. Systematic workflow of the bibliographic search and selection process.

3. Taxonomy of Deep Learning Architectures for Anomaly Detection

The analysis in this section focuses primarily on Deep Learning architectures, their structural frameworks, and their applications. The inputs for these recent models have transitioned from individual time steps to sliding windows. These windows aggregate data over time to capture underlying temporal patterns and inter-variable relationships. In real-world scenarios, data often exhibit complex structures such as noise, seasonality, and trends; furthermore, anomalies typically manifest as persistent patterns over a duration rather than isolated points [38], as illustrated in Figure 2, especially in sophisticated systems such as turbines. Such a transition toward Deep Learning for processing complex datasets reflects a broader trend in safety-critical infrastructures. For instance, similar architectures are systematically employed in rail track monitoring to detect defects and ensure operational reliability [49]. To establish a consistent taxonomy across all PdM tasks, including Anomaly Detection, RUL estimation, and RCA, this review categorizes methodologies into two overarching families: the Purely Data-Driven approach and the Hybrid Physics-Informed paradigm. The former relies exclusively on statistical data patterns. Conversely, the latter acts as a cross-cutting layer where physical domain knowledge is seamlessly integrated into deep learning architectures, such as forecasting or reconstruction models, to constrain predictions within physically plausible bounds. In the context of data streaming, a clear taxonomy can be established based on the temporal execution of training and inference phases. Generally, three primary operational frameworks are identified: Online Training and Online Inference, where the model continuously updates its parameters and detects anomalies as data flows; Offline Training and Online Inference, the most common approach in industrial applications, where a validated model is

deployed to monitor live streams; and Offline Training and Offline Inference, typically used for post-process diagnostic analysis [50]. In the turbomachinery sector, the framework based on offline training and online inference is the most suitable for addressing specific operational challenges. The Online inference is essential for the timely analysis of sensor signals such as vibrations, temperatures, and pressures, enabling the detection of incipient faults in real-time. Simultaneously, offline training is preferred over online methods to ensure system stability and prevent continuous learning from erroneously interpreting slow degradation as normal behavior. This approach also overcomes the computational limitations of industrial hardware, allowing for the deployment of strictly validated models and ensuring the reliability required in PdM applications.

Another fundamental aspect concerns the distinction between step-by-step and end-to-end architectures [37]. In step-by-step solutions, the model initially generates an intermediate output, such as a reconstruction error or an anomaly score, which requires a subsequent statistical stage to establish the actual presence of a fault. In this context, adopting techniques such as Nonparametric Dynamic Thresholding (NDT) or Peaks-Over-Threshold (POT) is decisive for defining adaptive thresholds capable of converting raw data into decision labels [37,50]. Given that high safety requirements characterize the turbomachinery sector, the step-by-step structure appears preferable to the end-to-end counterpart, which integrates the entire process into a single neural network. Indeed, the availability of an intermediate score guarantees greater physical interpretability within Anomaly or Faults Detection. Similarly, the use of algorithms like POT or NDT allows for better stochastic management of the nature of industrial signals, offering superior control over system sensitivity compared to the ‘black-box’ nature of end-to-end models [37,38]. Such an approach can be observed in [51], where the authors implement NDT for turbine monitoring, achieving a drastic reduction in false alarms due to the threshold’s ability to autonomously adapt to the machine’s operational variations. Similar results are found in the study by [52], where the use of the POT technique allows for the statistical management of extreme events in turbomachinery signals, ensuring timely fault detection without constant supervision. On the other hand, Ref. [53] proposes an end-to-end framework that, while outperforming classical models in accuracy and class-imbalance management, sacrifices granular explainability for direct, integrated classification.

Shifting to the frequency domain offers substantial advantages in turbomachinery diagnostics, enabling the extraction of information latent in the time domain. Spectral analysis facilitates the identification of cyclic patterns (Figure 2). Since rotating systems exhibit periodic behavior, the use of Fast Fourier Transform (FFT) allows for isolating dominant frequencies to detect anomalies or faults related to vibrations and mechanical imbalances that would appear only as noise in the raw signal, using these as inputs for the network [54]. Furthermore, signal transformation achieves a reduction in complexity by filtering random fluctuations and allowing the network to focus on the spectral signature of the machinery. It improves precision and reduces false alarms, particularly in architectures featuring filtering blocks such as Wavelet-based models [55] or FEDformer, which employs a Frequency Enhanced Block to manage global dependencies [56]. The frequency domain also enables a multidimensional representation. Models such as TimesNet [57] exploit dominant periodicities to transform 1D time series into 2D representations, jointly capturing temporal evolution and signal rhythms. Moreover, FFT-based formulations improve computational efficiency. In this context, Autoformer [58] replaces classical self-attention with a period-based auto-correlation mechanism, efficiently identifying long-term dependencies and trends. This facilitates the detection of anomalies and supports real-time monitoring in wind turbine applications. A summary scheme of the discussed topics is shown in Figure 6, where, as opposed to the Temporal module, which captures time-dependent patterns, the

Spatial module learns the normality of the system through inter-variable dependencies, enabling the detection of inter-metric anomalies that are essential for identifying complex fault conditions in multivariate turbine datasets.




Figure 6. Proposed learning scheme for Time Series Anomaly Detection related to the turbomachinery sector.

4. Deep Learning Families for Anomaly Detection in Turbine Operations

Within the context of turbine monitoring, Deep Learning architectures have evolved to handle the high dimensionality and non-linear nature of sensor data. According to recent literature, anomaly detection approaches can be broadly divided into two overarching families: Purely Data-Driven and Hybrid Physics-Informed paradigms. The first family includes methods that rely exclusively on historical data patterns, primarily categorized into Forecasting, Reconstruction, and Contrastive methods [38]. Conversely, the second family, such as Physics-Informed Neural Networks (PINNs), integrates physical principles into the learning process [16], embedding domain knowledge directly into the model to form hybrid strategies that enhance generalizability and ensure that the detected anomalies remain physically consistent with the turbine's operating logic. Consequently, this section analyzes the specific methodologies reported in the literature for turbomachinery, moving from purely data-driven models to hybrid physics-based solutions. Furthermore, dedicated summary tables are provided to systematize the reviewed contributions. To facilitate a comparative technical evaluation and practical implementation, these tables extend beyond basic categorization (i.e., learning paradigm, turbomachinery system, and year) by introducing three additional classification criteria:

- **Input Data Types (Input):** detailing the nature of the ingested signals. These are categorized as Multivariate Time Series (MTS) for standard SCADA logs, Spatio-Temporal MTS (ST-MTS) for graph-based topologies, Vibration signals (VIB) for high-frequency monitoring, Images/Spectrograms (IMG/SPEC), or Physics (e.g., thermodynamic residuals, PDE constraints, or calculated physical parameters).
- **Validation Metrics (Metrics):** specifying the mathematical criteria used to assess model performance, such as RMSE, MAE, MAPE and R^2 for forecasting errors, or F1-Score, Accuracy, and AUC for classification capabilities.
- **Computational Complexity (Comp.):** providing a qualitative assessment based on the architectural depth and hardware requirements. This metric is classified as Low (L) for shallow neural networks or single-layered architectures (e.g., simple CNNs or single-

layer LSTMs) ensuring immediate inference on standard edge devices; Medium (M) for deep architectures or attention-based mechanisms (e.g., Deep LSTMs, Transformers, or CNN-LSTM hybrids) requiring GPU acceleration for efficient training; and High (H) for computationally intensive frameworks, such as Physics-Informed Neural Networks (PINNs) or Generative Adversarial Networks (GANs), which involve complex iterative mathematical operations or massive parameter spaces.

4.1. Purely Data-Driven Family

The Purely Data-Driven family for Anomaly Detection relies exclusively on historical sensor data to model the nominal behavior of turbomachinery without explicit physical rules. This category is structured into three primary methods, each utilizing specific architectures to process complex multivariate time series:

- **Forecasting:** Forecasting strategies utilize historical sensor streams to model the temporal evolution of turbine parameters, enabling the prediction of future operational states. In this framework, anomalies are identified by quantifying the residual between the model's estimation and the actual physical measurement, where a significant deviation signals a departure from nominal behavior. The reference architectures employed within this predictive paradigm primarily include models such as CNN, RNN, LSTM, GNN, and Transformers [37,59].
- **Reconstruction:** Reconstruction-based methods focus on modeling the distribution of nominal behavior by mapping input signals into a compressed latent space. Unlike predictive approaches, these models aim to synthesize the original input, using the reconstruction error as a metric for Anomaly Detection. In the context of gas turbines, Ref. [47] highlights that these unsupervised architectures are particularly effective for identifying hidden patterns in multidimensional time series. In fact, by accessing the full temporal context related to the use of sliding windows, these models can effectively reconstruct complex operational scenarios [37]. Within this framework, nominal behavior is captured by encoding training subsequences into low-dimensional latent spaces. Under the assumption that anomalies are rare, such patterns are typically difficult for the model to reconstruct accurately. Therefore, detection is initiated when the reconstruction error, defined as the difference between the actual and synthesized observations, is high or when the reconstruction probability remains below a specific threshold. While this approach may introduce a minor delay, it is often favored in the condition monitoring of turbines and their related plants because it provides the high accuracy needed to reduce maintenance expenses and prevent critical equipment failures [47]. The specific implementations related to this method rely on various frameworks, including AE, VAE, GAN, Transformer, and Diffusion models.
- **Contrastive:** Contrastive learning paradigms for Anomaly Detection define a self-supervised approach centered on the agreement between different views of the same data [37]. These methodologies typically employ a twin-branch architecture (often referred to as a dual-tower model) where multivariate time series are processed through distinct encoding networks to produce high-dimensional representation vectors [38]. The model identifies consistent features of nominal behavior by analyzing the signal from multiple perspectives, such as diverse temporal scales or data augmentations [59]. To ensure synchronization across complex sequences, a fusion module incorporates alignment techniques like upsampling to map these representations to each timestep [50]. The training objective is governed by a contrastive loss function that quantifies the similarity between branches, marking instances of low correlation as potential anomalies [38]. This framework establishes a robust baseline without requiring manual labels, making it ideal for the unsupervised monitoring of industrial

assets. The techniques included in these works primarily include Transformer-based, MLP-based, and LLM-based models.

The specific applications, architectural nuances, and performance trade-offs of these methodologies will be detailed in the following sub-subsections of this part dedicated to purely data-driven methods. To provide a clearer visual summary of these purely data-driven architectures, a conceptual representation is illustrated in Figure 7.




Figure 7. Overview of the purely data-driven families for Anomaly Detection: (a) forecasting, (b) reconstruction, and (c) contrastive learning [38].

4.1.1. Forecasting Methods

A relevant trade-off for managing complex time series in turbomachinery involves the integration of advanced statistical models with learning techniques. Along these lines, Goyal [60] proposed a hybrid approach for gas turbines, combining Generalized Additive Models (GAM) with ARFIMA and GARCH processes to account for residual volatility and non-stationary data. This concept extends to sensor virtualization, where predictive models validate physical instruments or replace missing ones. The research projects by Gori et al. [61] and Shetty et al. [62] developed Continual Learning and Granger causality strategies for Anomaly Detection in gas turbines. Using LSTM networks, these models develop virtual sensors to detect sensor drift or failure. In this context, Farahani [63] addresses the challenge of high-dimensional data in gas turbines by coupling an Autoencoder for feature extraction with a deep LSTM for forecasting. This hybrid strategy effectively reduces noise and dimensionality before performing time-series prediction, managing operational shifts from a forecasting perspective.

To improve sensitivity toward critical components such as main bearings, specialized architectures have been recently proposed [64,65]. While B. Wang et al. [64] introduced a Feature Alignment LSTM (FA-LSTM) to account for the component's thermal inertia by synchronizing previous temperature states with current variables, Chen et al. [65] developed the Norm-linear-ConvNeXt-TCN framework to isolate degradation signals with higher computational efficiency. Utilizing Temporal Convolutional Networks (TCN), this latter approach successfully separates thermal degradation patterns from standard operational fluctuations.

Similarly, in the hydropower sector, B. Liu et al. [66] integrated CNNs and LSTMs with Generative Adversarial Networks (GANs) to balance datasets lacking sufficient failure samples, enhancing the identification of anomalies in generation units under variable loads. Purely temporal models, however, often struggle to capture the dependencies between interconnected subsystems. Graph Neural Networks (GNNs) address this by representing the turbine as a network of nodes. Related to this, Pincirolì et al. [67] utilized a GAT-GRU

framework to weight sensor correlations, ensuring the system remains sensitive to fault patterns that automatic control systems might otherwise obscure. A significant contribution in this field is the TempGNN model by G. Jiang [68], which focuses on system-level monitoring through a decoupling stage designed to remove the influence of time-varying conditions, identifying overheating in the main bearing, generator components, and converters with high precision. Further extending this spatial analysis, Y. Zhan et al. [69] implemented deep spatio-temporal modeling for anomaly detection in small hydropower stations, successfully capturing the co-evolution of state variables. In the wind sector, Daenens et al. [70] utilized GNNs to model correlations across offshore wind farms, where interactions such as the wake effect significantly influence regional power forecasting.

Finally, the forecasting literature highlights two distinct philosophies in graph construction: Y. Zheng et al. [71] proposed a hierarchical GNN for wind turbine pitch systems, integrating autocorrelation mechanisms to detect blade jamming or encoder errors through a data-driven adjacency matrix. In contrast, X. Jin et al. [72] emphasized the use of prior physical knowledge to define the graph topology. Their GSTN framework maps a fault propagation chain across sensors, enhancing the explainability of the diagnostic results. Addressing regional-scale management across wind turbine clusters, J. Wang et al. [73] implemented an improved DBSCAN anomaly detection as a preprocessing step to clean SCADA data from maintenance or grid curtailment artifacts. This ensures that the subsequent Graph Convolutional Recurrent Network (GCRN) receives reliable inputs for power forecasting, maintaining the integrity of the dispatching process.

Pros and Cons. In summary, forecasting architectures provide a robust framework for real-time monitoring and early degradation detection, particularly when integrated with spatial awareness or physics-informed modules. Their primary advantage lies in the direct explainability of the prediction residual as a physical indicator of system health. Nevertheless, these methods remain sensitive to signal noise and non-stationary operational shifts [60,61], which can induce false alarms if the baseline is not meticulously refined against artifacts such as outliers or concept drift [74]. As evidenced by the latest scientific literature, the frontier of this paradigm lies in the trade-off between the flexibility of deep learning and the structural reliability provided by physical constraints [72] and robust data pre-processing [73]. The works related to forecasting are reported in Table 1.

Table 1. Summary of the illustrated works on forecasting for turbomachinery Anomaly and Fault Detection.

Work	Learning Paradigms	System	Year	Input	Metrics	Comp.
[60]	GAM, ARFIMA, and GARCH	Gas Turbine	2021	MTS	RMSE, MAE	L
[61]	RNN and Continual Learning	Turbomachinery	2022	MTS	MSE, RMSE	M
[62]	RNN-LSTM and Granger Causality	Turbomachinery	2023	MTS	MSE, RMSE	M
[63]	Hybrid Autoencoder and Deep LSTM	Gas Turbine	2021	MTS	MSE, RMSE	M
[64]	FA-LSTM (Feature Alignment)	Wind Turbine	2025	MTS	RMSE, MAE, R^2	M
[65]	Norm-linear, ConvNeXt, and TCN	Wind Turbine	2024	MTS	F1-score	M
[66]	CNN-LSTM-GAN	Hydropower	2025	MTS	Acc, F1-score	H
[67]	GAT and GRU (Graph Attention)	Wind Turbine	2025	ST-MTS	F1-score, AUC	M
[68]	TempGNN (Temperature-based GNN)	Wind Turbine	2022	ST-MTS	RMSE, MAE	M
[69]	Deep Spatio-Temporal Modeling	Hydropower	2025	ST-MTS	F1-score	M
[70]	Spatio-Temporal GNN (STGNN)	Wind Turbine	2025	ST-MTS	RMSE, MAE, MAPE	M
[71]	Hierarchical ST-Autocorrelation GNN	Wind Turbine	2024	ST-MTS	F1-score, RMSE	M
[72]	GSTN (Graph Sensor Transformer)	Wind Turbine	2024	ST-MTS	Acc, F1-score	M
[73]	Improved DBSCAN, GBM, and AGCRNN	Wind Turbine	2025	ST-MTS	RMSE, MAE, R^2	M

Note: PIL: Physics-Informed Loss; PIA: Physics-Informed Architecture; PDDG: Physics-Driven Data Generation; MTS: Multivariate Time Series; ST-MTS: Spatio-Temporal MTS; VIB: Vibration; IMG/SPEC: Image/Spectrogram; RMSE: Root Mean Square Error; MAE: Mean Absolute Error; Acc: Accuracy; F1: F1-score.

4.1.2. Reconstruction Methods

The foundations of this approach lie in deep autoencoders and their robust variants, designed to operate in industrial contexts where data is often unlabeled. W. Xie et al. [75] use a Sparse Autoencoder (SAE), a reconstruction network that introduces a sparsity constraint in the latent space to force the model to learn only the most significant features of vibrations in aeroderivative gas turbines. By integrating Transfer Learning techniques, the authors successfully transfer knowledge between different engines, solving data scarcity and achieving a diagnostic accuracy of 92%.

The vast expansion of wind farms, particularly in China, generates massive amounts of SCADA data often affected by noise and disturbances. C. Zhang [76] attributes these issues to electromagnetic interference and outliers during startup or shutdown phases, while J. Chen [77] highlights the impact of communication failures and packet loss during standstill periods. A widely adopted solution for managing these corrupted signals is based on the Stacked Denoising Autoencoder (SDAE). This architecture forces the network to reconstruct the original input from a version corrupted by artificial noise, pushing the model to ignore random components and learn only the most robust latent features of the signal. C. Zhang [76] integrates a Local Outlier Factor algorithm to clean the dataset before training and fuses the SDAE with LSTM cells and an XGBoost classifier to create a complex diagnostic pipeline. In contrast, J. Chen [77] proposes an approach based on the purity of reconstruction, where temporal dependency is handled through an integrated sliding window within the SDAE. In this architecture, training is performed across multiple noise levels to refine the network's ability to distinguish salient features, making the reconstruction error a highly sensitive anomaly indicator. Within this family of methods, Yan and Yu [78] apply the SDAE to gas turbines to extract hierarchical features from exhaust gas temperature profiles, overcoming the limitations of traditional combustor monitoring. To address the issue of training datasets polluted by unreported faults, consisting of unlabeled data that might contain anomalies mistaken for normality, S. Fu [79] introduces the Re-optimized Deep Auto-encoder (R-DAE) for gas turbines. This architecture solves the inefficiency of classic models, which, if trained on unfiltered data, tend to reconstruct faults well, making their identification difficult. The R-DAE implements a selection mechanism that purifies the training set and combines latent features with reconstruction residuals to feed an Isolation Forest, defining a more accurate decision boundary. Within the same Gas Turbine environment, G. Lee [26] optimizes monitoring efficiency via a Convolutional Auto-encoder (CAE); this specific reconstruction architecture utilizes convolutional filters to capture structural patterns while drastically reducing the number of network parameters compared to dense models, ensuring low computational costs for long-term malfunction detection. Following this technological path, D. Liu et al. [80] proposed the CSiamese framework, which integrates a Convolutional Auto-encoder with a Siamese network to measure reconstruction similarity. This method addresses the critical problem of data imbalance in gas turbines, where the extreme rarity of abnormal samples hinders the training of standard models. By optimizing parameters exclusively on normal samples, CSiamese enhances anomaly detection performance specifically for such imbalanced datasets.

Temporal analysis represents another fundamental challenge that reconstruction methods address. Jalil Pour [81] focuses on sensor modeling for gas turbines using LSTM-Autoencoder networks to identify early signals of internal degradation through the drift of thermodynamic parameters. In this context, Fahmi et al. [82] demonstrate the superiority of Temporal Convolutional Network-Autoencoder (TCN-AE) networks over LSTM in terms of speed and accuracy for gas turbines. A significant contribution to the development of digital twins is offered by Y. Ma and others [83], who developed a Performance Digital Twin based on LSTM-Autoencoder networks for aviation Gas Turbines. This re-

construction architecture does not merely model nominal behavior, but actively quantifies operational uncertainty, making detection reliable even during complex operational transients. To extend the applicability of reconstruction-based models to entire Wind Turbine fleets, Roelofs et al. [84] demonstrate that an Autoencoder pre-trained on a source turbine can be successfully adapted to target units with just a few months of historical data, thanks to Transfer Learning, effectively monitoring new plants where historical data is still insufficient.

Research in recent years has moved towards modeling spatial correlations via graphs and multi-scale analysis, placing itself among the most advanced reconstruction architectures. A critical problem for wind turbines is blade icing, which remains a significant concern for these due to its negative impact on aerodynamics and structural integrity. Wang et al. [85] address this issue via the beta-VGATAE framework, employing Graph Attention Networks (GAT) to extract the spatial structure of sensor interactions. Although ice identification often concerns diagnostics, the authors highlight how this model provides a fundamental anomaly detection tool, identifying the dynamic discrepancy between a nominal and an iced blade. Similarly, Duan et al. [86] introduce the TCAD model, based on Transformer and ResNet, to associate global and local dependencies in multivariate sequences. Following this direction, the work by Y. Zhan [69], previously discussed in the context of forecasting, exploits its hybrid structure to bridge the gap between predictive and reconstruction-based anomaly detection, integrating next-value prediction with spatial reconstruction via Graph Attention Networks. Through this approach, the model identifies anomalies by highlighting the discrepancy between predicted dynamics and global signal coherence. Luo [87] completes this landscape with the METG model, which combines graph networks and memory modules to prevent rare anomalous patterns from being reconstructed too faithfully, thereby increasing the system sensitivity in distinguishing critical deviations from nominal behavior.

The adoption of probabilistic generative models has further refined the detection of incipient faults under heavy noise. Fan et al. [88] propose the TCVAE-GAN model for wind turbine bearings, combining the stability of the VAE with the discriminative strength of GAN. For Steam Turbines, Xu and Zhang [89] present the ELSTMVAE (Enhanced Long Short-Term Memory Variational Autoencoder) architecture. This model integrates an LSTM architecture to handle time series with a Variational Autoencoder (VAE), utilizing a Gaussian Mixture Model (GMM) in the latent space; this integration allows modeling complex distributions and eliminating inherent training anomalies, ensuring precise classification. Alper [90] confirms the effectiveness of these approaches for gearboxes, while Yang et al. [91] fuse VAE and Neural ODE (Ordinary Differential Equation) to predict the dynamic response of Wind Turbine blades. This model identifies progressive degradation, such as erosion or cracks caused by mechanical fatigue and cyclic aerodynamic loads, by analyzing the deviation between reconstructed and measured dynamics. In the maritime sector, Dabaja [92] uses GAN-based architectures identify faults in marine engine air systems weeks in advance.

The evolution of these methodologies leads to the integration of reconstruction with specialized detection paradigms. For instance, Peng [93] introduces the DUA-SVDD (Deep Unsupervised Adaptive Support Vector Data Description) model, which defines a geometric hypersphere to enclose normal operational data. The network simultaneously minimizes the reconstruction error and the hypersphere volume, providing adaptive alarm thresholds for wind turbines. To conclude, Liang et al. [94] introduce ExpertAP, a framework applied to steam turbines that addresses the limitations of single-unit monitoring. By learning from historical patterns across multiple units, similar to human experts, the model mitigates

the scarcity of anomaly labels through multi-unit pretraining and knowledge transfer, significantly improving the identification of potential failures.

Pros and Cons. In summary, reconstruction methods for anomaly or fault detection in turbomachinery strike a balance between extraction capability and architectural complexity. While AE architectures, extensively applied to gas turbines and wind farms [76,77], allow for the automatic extraction of salient parameters, they often show limits in the physical interpretation of the reconstruction error. To address this, VAE models provide a solid probabilistic basis for health-monitoring systems [85], although the assumption of Gaussian distributions may limit their precision in capturing highly non-linear operational shifts. Higher fidelity in modeling such complex distributions is ensured by GANs [92], despite their potential for training instability. At the same time, diffusion models [78] prove effective against SCADA signal noise, yet they require specific adjustments to handle the rapid temporal dynamics of turbomachinery. This evolution currently trends toward hybrid and multi-unit frameworks [69,94], which aim to overcome individual constraints by integrating diverse paradigms to achieve more reliable and interpretable diagnostics. The studies focusing on reconstruction-based methods are summarized in Table 2.

Table 2. Summary of recent literature on reconstruction-based methods for turbomachinery Anomaly and Fault Detection.

Work	Learning Paradigms	System	Year	Input	Metrics	Comp.
[75]	Sparse Autoencoder (SAE) with Transfer Learning	Gas Turbine	2022	VIB	F1-score, Acc	M
[76]	LSTM-SDAE and XGBoost	Wind Turbine	2022	MTS	F1-score, AUC	M
[77]	Stacked Denoising Autoencoders (SDAE)	Wind Turbine	2020	MTS	F1-score, RMSE	M
[78]	SDAE for Hierarchical Features	Gas Turbine	2019	MTS	F1-score, Acc	M
[79]	Re-optimized Deep Auto-encoder (R-DAE)	Gas Turbine	2021	MTS	F1-score, AUC	M
[26]	Convolutional Auto-encoder (CAE)	Gas Turbine	2020	MTS	F1-score, MSE	M
[80]	CSiamese (CAE + Siamese Network)	Gas Turbine	2023	MTS	F1-score, Acc	M
[81]	LSTM-Autoencoder (Sensor Modeling)	Gas Turbine	2022	MTS	MSE, MAE	M
[82]	TCN-Autoencoder (TCN-AE)	Gas Turbine	2024	MTS	RMSE, R^2	M
[83]	VAE-based Performance Digital Twin	Gas Turbine	2023	MTS	RMSE, MAE	M
[84]	Transfer Learning with Autoencoders	Wind Turbine	2024	MTS	F1-score, AUC	M
[85]	Beta-VGATAE (Graph Attention Networks)	Wind Turbine	2024	ST-MTS	F1-score, Acc	M
[86]	TCAD (Transformer and ResNet)	Wind Turbine	2022	MTS	F1-score	M
[69]	Deep Spatio-Temporal Modeling	Hydropower	2025	ST-MTS	F1-score	M
[87]	METG (Memory-enhanced Transformer)	Wind Turbine	2024	ST-MTS	F1-score	M
[88]	TCVAE-GAN and Correlation Enhancement	Bearing	2023	VIB	F1-score, Acc	H
[89]	ELSTMVAE-DAF-GMM (Enhanced VAE)	Steam Turbine	2025	MTS	F1-score, AUC	M
[90]	AE, VAE, and Deviation Networks	Gearbox	2023	VIB	F1-score, AUC	M
[91]	VAE and Neural ODE Integration	Wind Turbine	2024	VIB	F1-score, RMSE	H
[92]	GAN-based Unsupervised Detection	Marine Engine	2024	MTS	F1-score	H
[93]	DUA-SVDD (Deep Unsupervised SVDD)	Wind Turbine	2025	MTS	F1-score, AUC	M
[94]	Semi-supervised Sequence Reconstruction	Steam Turbine	2025	MTS	F1-score	M

Note: PIL: Physics-Informed Loss; PIA: Physics-Informed Architecture; PDDG: Physics-Driven Data Generation; MTS: Multivariate Time Series; ST-MTS: Spatio-Temporal MTS; VIB: Vibration; IMG/SPEC: Image/Spectrogram; RMSE: Root Mean Square Error; MAE: Mean Absolute Error; Acc: Accuracy; F1: F1-score.

4.1.3. Contrastive Methods

The evolution from reconstruction-based detection toward learning discriminative latent representations has redefined the monitoring of rotating machinery. As discussed in the previous subsection, D. Liu et al. [80] introduced the CSiamese framework, which integrates a CAE with a Siamese network to measure reconstruction similarity. This method, while sharing roots with previously analyzed reconstruction models, employs a triangle loss to optimize parameters on imbalanced datasets, overcoming the limitations of standard models in gas turbines. A recurring problem in industrial contexts like turbines is the difficulty of obtaining a database covering all possible fault types. In this regard, J. Chen [95]

addresses the scarcity of labels: the model does not search for a specific fault but learns to master the machine's normal behavior. He proposes the TCN framework, an unsupervised approach that utilizes neural transformations and temporal contrastive learning to capture complex patterns without requiring negative samples, thereby preventing representation collapse. This network, which the author applies to steam turbines, acts as a sort of short-to-medium-term historical memory. Thanks to dilated convolutions, it can perceive a very wide temporal window. If the signal sequence (e.g., pressure rising while temperature falls anomalously) does not correspond to what was learned during training, the model generates an alert. The peculiarity of Chen's approach is that it does not require fault examples for training; instead, it compares the current signal with a transformed version of itself (a self-supervised approach) to verify the consistency of the latent representation. Managing the intrinsic challenges of SCADA data, such as imbalance and environmental noise, is a central theme of recent research. Sun et al. [96] developed the Matching Contrastive Learning (MCL) method, which introduces fixed reference points (bollards) in the latent space to guide data distribution toward stable centers, ensuring separability even with few fault examples. In another work, the same authors [97] adopt non-parametric regression to purify signals in scenarios of high meteorological variability. Since the relationship between environmental variables and turbine response is non-linear, this method models nominal behavior to calculate the discrepancy between expected and measured values. By isolating this clean signal, contrastive learning defines sharper boundaries between healthy and broken blades, regardless of aerodynamic loads or fluctuations induced by weather conditions, proving effective for the anomaly detection of blades. For managing rare or new faults, X. Liu et al. [98] utilize the Meta-Analogical Momentum Contrast Learning (MACL) method, based on a momentum encoder and one-shot learning logic, allowing for the identification of anomalies with a single example. Similarly, the problem of blade icing is addressed, as ice subtly alters the blade's mass and aerodynamic profile. Researchers like Z. Wang [99] apply the Unsupervised and Imbalanced Semi-Supervised Contrastive Learning (UISSCL) framework, which optimizes recognition in semi-supervised contexts. Using contrastive learning, this model learns to extract fundamental signal features (vibration, power, and wind speed) during training; then, the few available labeled data are introduced. The network uses these anchors to map the rest of the dataset. Finally, the framework optimizes the position of data in the latent space so that rare icing events are not suppressed or ignored by the massive amount of normal data. Changing strategy but remaining within the contrastive learning family, Guo et al. [100] construct a K-Nearest Neighbor Graph (KNNG) to map sensor interconnections and then use Graph Contrastive Learning (GCL) to train the network to recognize faults based on the graph's structure. The fusion of contrastive-based deep learning methods and physical knowledge represents a further methodological support. In this view, Qiao et al. [101] utilize prior knowledge by integrating fault characteristic frequencies into a model termed Prior Knowledge Embedding Contrastive Attention Learning Network (PKECALN). It consists of a backbone based on a 1D Deep Convolutional Neural Network (1D-DCNN) integrated with a custom Sequential Attention Module (SAM) to extract multi-scale time-frequency features from vibration signals and integrate prior knowledge (fault frequencies) to improve diagnosis with small samples. Miyamoto [102] and C. Wang [103] use contrastive learning to resolve visual confusion in anomaly detection. Specifically, Miyamoto proposes the Adaptive Activated Anomaly Detection (AAAD) model, which isolates inspection targets from variable backgrounds to find potential faults with few labeled images. Meanwhile, ref. [103] introduces the Spatial Contrast and Semantic Difference Perception Network (SSPN) for precise damage segmentation on aero-engine blades, achieving extreme accuracy at damage boundaries and preventing the loss of spatial details during image processing. Qin [104]

proposes the Time-series and Image Pre-trained Encoder (TIPE) model based on ResNet, which aligns pressure signals and combustion images to detect thermo-acoustic instabilities through a multi-modal contrastive learning framework within gas turbines. Finally, the current research frontier shifts toward the integration of descriptive semantics offered by LLM. For decision support, Y. Li et al. [105] developed LLM-YOLOMS, based on YOLO (You Only Look Once) [106], where visual detection results performed at Multi-Scale (MS) are converted into textual attributes and interpreted by an LLM to provide maintenance recommendations. Tang et al. [107] introduce an LLM-based interpreter capable of translating high-frequency SCADA data into natural language health assessments, overcoming traditional diagnostic rigidity, where it is nearly impossible for a human operator to identify anomalies without intermediate tools and where classical anomaly detection systems often provide only a binary alarm (0 or 1) without further explanation.

Pros and Cons. Contrastive learning offers significant advantages for rotating machinery monitoring, particularly in managing extreme class imbalance through stable latent centers [96] and learning from unlabeled data to capture nominal behavior [95]. Its ability to integrate prior knowledge [101] and align multimodal signals [104] improves diagnostic robustness under low-label conditions. However, these methods often require meticulous design of negative samples to prevent representation collapse, a limitation effectively addressed by hybrid architectures like CSiamese [80], which utilize reconstruction objectives and triangle loss to stabilize training. Furthermore, pure contrastive models can be sensitive to environmental noise without specific pre-processing [97]. Conversely, LLM-based frameworks [105,107] provide superior semantic interpretability for maintenance technicians and few-shot adaptability, transforming structured detection into actionable maintenance advice. Despite these strengths, the direct application of LLMs to time series data remains a challenge, requiring specific prompting strategies to be effective. Moreover, as highlighted in current surveys [38], LLMs risk generating hallucinations when indexing or explaining anomaly points and may struggle with highly complex, context-dependent temporal anomalies. Consequently, adhering to Industry 5.0 paradigms, the integration of a Human-in-the-Loop (HITL) approach remains essential to validate AI inferences against domain expertise, ensuring that deep learning outputs act as decision support rather than autonomous replacements [35]. The works related to contrastive learning are reported in Table 3.

Table 3. Summary of the literature on contrastive learning for turbomachinery Anomaly and Fault Detection.

Work	Learning Paradigms	System	Year	Input	Metrics	Comp.
[80]	CAE and Siamese Network	Gas Turbine	2023	MTS	Acc, F1-score	M
[91]	DCdetector (Dual Attention Contrastive)	Wind Turbine	2023	MTS	F1-score	M
[108]	1D-CNN and Supervised Contrastive	Steam Turbine	2024	VIB	Acc, F1-score	M
[99]	UISSCL (Semi-supervised Contrastive)	Wind Turbine	2023	MTS	F1-score, AUC	M
[101]	PKECALN (Contrastive Attention)	Bearings	2024	VIB	Acc, F1-score	M
[104]	TIPE (Multimodal Contrastive Learning)	Combustors	2023	IMG/SPEC + MTS	F1-score	M
[103]	SSPN (Spatial and Semantic Contrast)	Aeroengine Blade	2025	IMG/SPEC	F1-score	M
[105]	LLM and YOLOMS Semantic Interpretation	Wind Turbine	2025	IMG/SPEC	Acc, F1-score	H
[95]	CNT (Graph-based Contrastive Learning)	Steam Turbines	2025	ST-MTS	F1-score, AUC	M
[96]	Matching Contrastive Learning	Wind Turbine	2023	MTS	F1-score	M
[102]	AAAD (Contrastive and Segmentation)	Wind Turbine	2025	IMG/SPEC	AUC	M
[98]	One-Shot (Meta-Analogical Momentum)	Wind Turbine	2022	MTS	Acc, F1-score	M
[100]	Graph Contrastive Learning (GCL)	Wind Turbine	2025	ST-MTS	Acc, F1-score	M
[97]	Environment-adapted Contrastive	Wind Turbine	2023	MTS	F1-score	M
[107]	LLM-based SCADA Interpreter	Wind Turbine	2025	MTS	Acc, F1-score	H

Note: PIL: Physics-Informed Loss; PIA: Physics-Informed Architecture; PDDG: Physics-Driven Data Generation; MTS: Multivariate Time Series; ST-MTS: Spatio-Temporal MTS; VIB: Vibration; IMG/SPEC: Image/Spectrogram; RMSE: Root Mean Square Error; MAE: Mean Absolute Error; Acc: Accuracy; F1: F1-score.

4.2. Hybrid Physics-Informed Family

Representing the second family for Anomaly and Fault Detection, hybrid methodologies integrate domain knowledge. A primary approach involves sequential coupling, where physics-based models are used to generate reference features that are subsequently analyzed through deep learning architectures [109]. Alternatively, Physics-Informed Neural Networks (PINNs) embed partial differential equations (PDEs) directly into the loss function. These frameworks address forward problems by acting as efficient surrogates for traditional numerical methods (such as FEM or CFD), enabling real-time estimation of unmeasurable critical quantities like equivalent stress, fatigue, or entire velocity and pressure fields [110]. Simultaneously, in inverse problems, PINNs infer unknown parameters or source terms (such as heat fluxes) from observed sensor data [111]. For diagnostics, this enables Anomaly Detection based on physical inconsistencies and direct comparisons between model outputs and sensor readings or nominal references, facilitating direct fault quantification and providing a basis for complex turbine health management [16]. Although various distinct approaches fall under this Hybrid Physics-Informed family, they will be analyzed collectively within this single subsection.

Analysis of Hybrid Physics-Informed Architectures

Methods integrating neural networks with physical laws make it possible to overcome the limitations of each individual approach, creating robust and reliable systems. These approaches help data-driven techniques cope with data scarcity, overfitting, and scenarios in which physical-domain computations are costly, accelerating CFD simulations by up to 70% and enabling online monitoring and real-time optimization [112,113]. In particular, they allow known physical information to be exploited even when operational data are noisy or incomplete, ensuring predictions that remain consistent with domain laws.

The first research line integrates physical equations directly into the optimization process through PINNs. Lai et al. [114] propose a multi-stage framework for nuclear turbines: the model decomposes the problem into sequential phases by solving the Navier–Stokes equations, monitoring flow rate and flow coefficient (C_v), and detecting obstructions, internal leakage, and stiction phenomena with respect to reference benchmarks. In the wind energy sector, Pei et al. [115] develop a hierarchical architecture: a data-cleaning module removes anomalies from SCADA signals, a GRA-LSTM network evaluates the turbine state, and a BiLSTM-PINN predicts power constrained by the fundamental theoretical kinetic energy equation ($P \propto v^3$). While real-world turbines typically operate at a lower exponent (approximately 2–2.5) due to mechanical-electrical losses and the Betz limit, the architecture utilizes such physical grounding to correct inconsistent deviations. Mittal et al. [116] address faulty sensors using clustering-based imputation, enabling the PINN to regularize the power curve while preserving physical coherence even in the presence of uncertain data. Gijon et al. [117] combine PINNs with Evidential Deep Learning on real SCADA data, directly parameterizing probabilistic distributions on the outputs (power, torque, and power coefficient). Physical constraints make it possible to distinguish epistemic and aleatoric uncertainty, respecting fundamental aerodynamic relationships and reducing uncertainty in sparsely sampled operating regions. For mechanical components, Zhang et al. [118] propose CNN-IPINN: the CNN extracts features from gearbox vibration signals, while the PINN solves the gear mesh stiffness equations. This dual approach ensures that fault diagnosis is consistent with mechanical dynamics, going beyond simple statistical pattern recognition.

A second research line concerns embedding physical coherence into generative models and architectures. Zideh and Solanki [119] develop a Physics-Informed Convolutional Autoencoder (PIConvAE) to protect power grids with high penetration of Distributed Energy Resources (DER). Unlike traditional autoencoders, which minimize only the statistical

reconstruction error, this model integrates the residuals of nodal power balance equations into the loss. In this way, an anomaly is detected not only when data deviate from learned patterns, but also when physical constraints are violated, enabling the detection of stealth cyber-attacks such as False Data Injection, which are statistically camouflaged but physically inconsistent. In wind turbine monitoring, Wu et al. [120] develop a Physics-Informed Patch Transformer that primarily functions as a forecasting model while integrating physics directly into its architecture. This distinction serves to differentiate between autonomous feature learning and hybrid models where physical consistency is structurally enforced. The model constrains the attention mechanism through physical coherence over temporal segments. The approach stabilizes turbine behavior modeling, anticipates faults up to 5 days before the critical event, reduces post-repair false alarms, and allows maintenance personnel to accurately locate the fault source, optimizing maintenance interventions. In the context of geared transmissions, Yue et al. [121] propose the Physics-Informed Attention-aided Multiple Autoencoder (PI-AMAE) for zero-shot scenarios, where fault histories are unavailable. The framework incorporates a lumped-parameter dynamic model that generates synthetic fault data, referred to as signal twins, allowing the system to recognize anomalies never observed in operation, thanks to a prior physical signature. In parallel, Duan et al. [122] address noisy and stochastic data in wind turbines using a PyGRU (Physics-informed Gated Recurrent Unit), integrating physical constraints to model uncertainty. The PyGRU reduces epistemic uncertainty on normal data and increases sensitivity on anomalous data, improving detection and reducing false alarms caused by fixed thresholds.

Physics also facilitates the generation of synthetic data to train networks under limited data conditions. Alblawi [109] simulates gas turbine degradation to train Multi Feed-forward Artificial Neural Networks (MFANNs) and isolate environmental effects, while Schröder et al. [123] address wind load monitoring scarcity through Transfer Learning, using aeroelastic simulations for physical pre-training before refinement with real-world data. Such Virtual Sensing strategies enable accurate real-time fatigue load estimation using only SCADA data. Similarly, Huber et al. [124] demonstrate the effectiveness of physics-informed data augmentation by generating physics-based fault trajectories and mixing them with nominal data to train classifiers capable of recognizing rare faults. In the context of extreme class imbalance in rotating machinery, Li et al. [125] adopt a generative strategy based on physical models to create signal twins and train contrastive networks. In this case, the authors do not use standard PINN architectures: physics does not enter as a constraint in the cost function, but is employed upstream through a lumped-parameter dynamic model, generating synthetic replicas of fault vibrations that serve as training data. Perez-Sanjines et al. [126] and Jamil et al. [127] develop fault detection methodologies for wind turbines, arguing that raw data are often too noisy for direct deep learning. Instead, they use Cyclic Spectral Coherence, a signal processing technique based on rotating component kinematics to extract cyclostationary signatures hidden in noise. These physically meaningful 2D maps are used as inputs to deep classifiers (CNNs) and enable reliable detection of incipient faults in gearboxes across entire fleets. Freeman et al. [128] address rotor imbalance detection in marine turbines, a condition often masked by hydrodynamic instability. The authors integrate turbulence intensity, computed from the flow, into a 1D CNN. The combined use of TI as both an explicit input feature and as a regulator within a physics-informed loss function allows the model to separate environmental fluctuations from real structural anomalies, improving the detection of hidden faults. The integration of physical features with advanced optimization algorithms is investigated by Lee et al. [129,130] and Chen et al. [131]. Lee utilizes Particle Swarm Optimization (PSO) to calibrate CatBoost classifiers on imbalanced wind SCADA data, while Chen maximizes the discrepancy between real data and physically simulated fault states to optimize LSTM

sensitivity to damage symptoms. Finally, Khan et al. [132] apply a physics-guided Bayesian neural network to wind turbine sensor networks, exploiting physical correlations among operational variables to detect sensor faults with a probabilistic measure of confidence.

Pros and Cons. As evidenced by the review of the aforementioned works, the integration of physical knowledge into deep learning models can be traced back to three integration paradigms, each with specific advantages and limitations. The first paradigm, Physics-Informed Loss (PIL), typical of PINNs, utilizes physics as an external constraint: the network maintains a standard architecture but is penalized when its predictions violate fundamental domain laws [115–117]. Key advantages include high efficiency with limited data and a mesh-free nature, which simplifies the modeling of complex geometries without requiring discretization grids [114]. Furthermore, the ability to solve both forward and inverse problems within a single framework represents a unique strength for identifying unknown parameters [111]. However, PINNs face significant optimization challenges: stiff and multi-scale loss landscapes make convergence difficult, requiring extremely precise weight tuning. They also exhibit marked sensitivity to noise in inverse problems, where signal degradation can lead to biased parameter estimations. The second paradigm, Physics-Informed Architecture (PIA), ensures greater stability by structurally incorporating physical knowledge directly into the network’s building blocks, as implemented in the Patch Transformer [120], the PI-AMAE framework [121], or the PyGRU cell [122]. This method ensures that every intermediate model state adheres to physical constraints, improving structural robustness. Finally, the third paradigm, Physics-Driven Data Generation (PDDG), utilizes physics to create ideal synthetic datasets to train the network [109,123,125]. While this strategy radically solves class imbalance and the scarcity of real-world samples, its effectiveness depends entirely on the fidelity of the original physical model: any upstream modeling error is inevitably inherited by the deep learning system. Table 4 provides an overview of studies on physics-informed methods; in this case, instead of the publication date in the last column, the integration paradigms (PIL, PIA, and PDDG) are reported.

Table 4. Summary of recent literature on hybrid physics-informed methods for turbomachinery Anomaly and Fault Detection.

Work	Learning Paradigms	System	Integration	Year	Input	Metrics	Comp.
[114]	Multistage PINN	Nuclear Valve	PIL	2024	Physics + MTS	F1-score, RMSE	H
[115]	BiLSTM-PINN Architecture	Wind Turbine	PIL	2026	Physics + MTS	RMSE, MAE	H
[116]	Error Detection PINN	Wind Turbine	PIL	2024	Physics + MTS	RMSE, MAE	H
[117]	PINN with Uncertainty Quantification	Wind Turbine	PIL	2026	Physics + MTS	RMSE, MAE	H
[118]	CNN-IPINN (CNN + PINN)	Gearbox	PIL	2025	Physics + VIB	F1-score, Acc	H
[119]	PI-Convolutional Autoencoder	Power Grid	PIL	2025	Physics + MTS	F1-score, RMSE	H
[128]	Turbulence Intensity Infusion PINN	Marine Turbine	PIL	2022	Physics + MTS	F1-score, Acc	H
[129]	Deep CNN and AEPSo-Catboost	Wind Turbine	PIL	2025	Physics + MTS	F1-score, Acc	H
[130]	Deep CNN and AEPSo-XGBoost	Wind Turbine	PIL	2024	Physics + MTS	F1-score, Acc	H
[132]	Physics-Guided Bayesian NN	Wind Turbine	PIL	2025	Physics + MTS	F1-score, RMSE	H
[127]	PI-Multivariate Deep Learning	WT Drivetrain	PIL	2023	Physics + MTS	F1-score, Acc	H
[120]	Physics-informed Patch Transformer	Wind Turbine	PIA	2026	Physics + MTS	F1-score, Acc	H
[121]	Attention-aided Multiple AE (PI-AMAE)	Gearbox	PIA	2025	Physics + VIB	F1-score, Acc	H
[122]	Probabilistic Physics-informed AE	WT Electromech.	PIA	2025	Physics + MTS	F1-score, AUC	H
[126]	PI-Cyclic Spectral Coherence DL	WT Gearbox	PIA	2023	Physics + VIB	F1-score, Acc	H
[109]	Thermodynamic Model and MFANN	Gas Turbine	PDDG	2020	Physics + MTS	RMSE, MSE	M
[123]	PI-ML with Transfer Learning	Wind Turbine	PDDG	2022	Physics + MTS	RMSE, MAE	M
[124]	ML with Physics-informed Augmentation	Gas Turbine	PDDG	2023	Physics + MTS	F1-score, Acc	M
[125]	Physical-Knowledge Contrastive Learning	Rotating Mach.	PDDG	2025	Physics + VIB	F1-score, Acc	M
[131]	Physics-Informed LSTM (Hyperparameters)	Gearbox	PDDG	2022	Physics + VIB	F1-score, RMSE	M

Note: PIL: Physics-Informed Loss; PIA: Physics-Informed Architecture; PDDG: Physics-Driven Data Generation; MTS: Multivariate Time Series; ST-MTS: Spatio-Temporal MTS; VIB: Vibration; IMG/SPEC: Image/Spectrogram; RMSE: Root Mean Square Error; MAE: Mean Absolute Error; Acc: Accuracy; F1: F1-score; AUC: Area Under the Curve; MSE: Mean Square Error.

Synthesizing the Anomaly Detection landscape analyzed throughout this section, while the reviewed architectures have reached a high level of development, their ability to generalize across different turbine fleets and maintain accuracy under severe sensor noise represents a significant frontier, further examined in “Section Research Gaps and Future Perspectives”.

5. Prognostics and RUL Estimation in Modern Turbine Systems

In the past decade, predictive maintenance for turbine systems has progressively transitioned from classical machine learning to Deep Learning architectures [133,134]. This transformation is driven by the superior capability of deep models to handle high-dimensional monitoring data and autonomously extract degradation patterns, effectively bypassing manual feature engineering [135,136]. However, RUL estimation remains largely dependent on supervised paradigms requiring extensive datasets with complete run-to-failure histories [134]. To align with the overarching taxonomy established in the previous section for Anomaly Detection, the methodologies for RUL prediction can be strictly classified into two primary families: Pure Deep Learning for RUL and Physics-Guided RUL. To address the scarcity of comprehensive labeled data, recent research indicates that integrating physical principles with neural networks as a cross-cutting layer can overcome the limitations of purely data-driven models [16,137]. Despite the success of standard Deep Learning, the future of prognosis also lies in hybrid methodologies [36]. Reflecting this core methodological split, the following discussion analyzes recent developments by explicitly contrasting these two architectural domains.

5.1. Pure Deep Learning for RUL

Research within the purely data-driven domain focuses extensively on feature extraction from complex data streams and overcoming the dependence on heavily labeled datasets through semi-supervised paradigms or deep architecture optimizations. J. Sun [138] proposes a DT for the performance of aerospace gas turbines using a semi-supervised method. Instead of a complete physical model, they use Deep Learning to build a Performance Digital Twin (PDT) capable of extracting multi-dimensional health indicators. Similarly, Y. Wang et al. [139] introduce the FCDAE-CNN-LSTM model for turbofan engines, a semi-supervised system that employs autoencoders for signal denoising, reducing the error induced by environmental noise and the scarcity of diagnostic labels. Focusing strictly on deep feature extraction, Elsherif et al. [140] apply the supervised CAELSTM network to turbofan engines, monitoring multi-sensor telemetry (temperatures, pressures, and rotational speeds) from the C-MAPSS datasets. Despite the model’s complexity, the approach remains strictly data-tied, using a convolutional autoencoder for noise reduction and attention mechanisms to identify the most relevant temporal degradation sequences. In the context of wind turbine gearbox bearings, Ke He et al. [141] focus on high-frequency vibration signal analysis by introducing the Self-Calibration Temporal Convolutional Network (SC-TCN). This supervised model is purely Deep Learning, but introduces a significant mathematical innovation: an architectural self-calibration module designed to preserve local information from vibration sensors, which is typically lost in the dilated convolutions of standard TCNs. This allows for capturing micro-variations in the vibratory signal symptomatic of incipient wear, overcoming the gaps of classical temporal networks. The flexibility of these systems is further explored by Jiaze Li [142], who applies Multimodal Transfer Learning to aircraft turbines to transfer RUL knowledge across different operating regimes.

Finally, the expansion of prognosis to diverse physical signals completes the data-driven landscape. Atsafack et al. [143] use IoT data and supervised CNN-LSTM networks for the RUL of hydraulic turbines, monitoring electrical and hydraulic parameters in real-

time. In thermoelectric power plants, Dayang Li et al. [142] forecast the degradation trend of steam turbines by correlating online oil monitoring data (viscosity, humidity, and wear particles) with the output power through LSTM networks. This connects to the work of Lei Sun et al. [144], which confirms the effectiveness of supervised RNNs and CNNs for power prediction in steam turbines. Closing the review, Rengasamy [145] proposes the use of asymmetric loss functions in deep models for aero-engines, optimizing supervised prognosis to avoid RUL underestimations that could compromise operational safety.

5.2. Physics-Guided RUL

To bridge the interpretability and generalization limits of the purely statistical models discussed above, prognostic literature highlights a growing interest in the integration of physical models and deep neural architectures for Remaining Useful Life (RUL) estimation. In this context, Ang Li [146] presents the Phy-DeepLSTM framework applied to marine steam turbines. This represents a supervised, hybrid physics-informed approach to system-wide Digital Twin modeling; the model integrates thermodynamic constraints into the DeepLSTM network to ensure physical consistency during operational transients, thereby improving accuracy over traditional black-box methods. Similarly, Yucesan and Viana [110] introduce a supervised PINN for the RUL estimation of wind turbine main bearings. Their model is fully hybrid, designed to fuse physical layers (modeling fatigue based on torsional and axial loads) with data-driven layers tasked with capturing complex stochastic phenomena, such as the chemical degradation of lubricating grease.

The management of thermomechanical stress characterizes the work of Zhang et al. [147], who develop a PINN to predict the creep-fatigue life of components operating at elevated temperatures (gas and steam turbines). This supervised model integrates physical knowledge through both dedicated feature engineering and a constrained loss function that allows for effective operation even with reduced datasets. Within the wind energy sector, Zheng Wang et al. [148] address the RUL of gearbox bearings using a supervised hybrid PI-LSTM network. The integration of prior knowledge regarding mechanical wear mechanisms guides the network's learning even in the presence of limited samples. Similarly, Huber et al. [124] extend this philosophy to gas turbines, utilizing hybrid models where physics supports data augmentation to generate synthetic failure scenarios, making the supervised prognosis significantly more robust against the scarcity of real data. Expanding on the concept of physics-driven data generation, Lu et al. [137] propose a Digital Twin-driven architecture for gearbox RUL estimation. By employing a high-fidelity virtual model to simulate degradation trajectories, their framework effectively compensates for the lack of real-world run-to-failure data. Additionally, the authors introduce a water-wave information transmission network to process multivariate time series, which captures long-term dependencies more stably than conventional recurrent networks. An overview of both purely data-driven and physics-guided research on RUL prediction methods is presented in Table 5.

Concluding the review of prognostic methodologies, while physics-guided RUL models demonstrate high reliability, managing the inherent uncertainties of aging systems and optimizing these models for constrained edge hardware remain primary challenges, as detailed in “Section Research Gaps and Future Perspectives”.

Table 5. Summary of recent literature on RUL prediction and performance prognostics for turbomachinery.

Learning Paradigms	System	Family	Year	Input	Metrics	Comp.
[138] Performance Digital Twin (PDT)	Gas Turbine	Pure DL	2023	MTS + DP	RMSE, MAE	M
[139] FCDAE-CNN-LSTM (Semi-supervised)	Turbofan Engine	Pure DL	2023	MTS	RMSE, MAE	M
[140] CAELSTM with Attention	Turbofan Engine	Pure DL	2025	MTS	RMSE, MAE	M
[141] Self-Calibration TCN (SCTCN)	Wind Turbine	Pure DL	2022	VIB	RMSE, MAE	M
[149] Multimodal Transfer Learning	Turbine Engine	Pure DL	2025	MTS	RMSE, MAE	M
[143] CNN-LSTM and IoT Data	Hydraulic Turbine	Pure DL	2025	MTS	RMSE, MAE, R^2	M
[142] LSTM and Oil Monitoring	Power Plant Turbine	Pure DL	2024	MTS	RMSE, MAE, MAPE	M
[144] RNN and CNN Power Prediction	Steam Turbine	Pure DL	2021	MTS	RMSE, MAE, MAPE	M
[145] DL and Asymmetric Loss Functions	Gas Turbine	Pure DL	2020	MTS	RMSE, MSE	M
[146] Physics-informed DeepLSTM	Steam Turbine	Phys-Guided	2025	Phys + MTS	RMSE, MAE	H
[110] Physics-informed Neural Network	Wind Turbine	Phys-Guided	2021	Phys + MTS	RMSE, MAE	H
[147] PINN (Feature Engineering + Loss)	Steam Turbine	Phys-Guided	2021	Phys + MTS	RMSE, R^2	H
[148] PI-LSTM and Prior Knowledge	Wind Turbine	Phys-Guided	2022	Phys + MTS	RMSE, MAE	H
[124] ML and Physics-informed Augmentation	Gas Turbine	Phys-Guided	2023	Phys + MTS	RMSE, MAE	M
[137] Digital Twin and WITRAN	Gearbox	Phys-Guided	2025	Phys + VIB	Acc, RMSE, MAE	H

Note: PIL: Physics-Informed Loss; PIA: Physics-Informed Architecture; PDDG: Physics-Driven Data Generation; MTS: Multivariate Time Series; ST-MTS: Spatio-Temporal MTS; VIB: Vibration; Phys: Physics-based inputs; DP: Design Parameters; RMSE: Root Mean Square Error; MAE: Mean Absolute Error; MAPE: Mean Absolute Percentage Error; MSE: Mean Square Error; R^2 : Coefficient of Determination; Acc: Accuracy.

6. Root Cause Analysis in Modern Turbomachinery Diagnostics

Starting from the classics of Root Cause Analysis (RCA), such as the Ishikawa Diagram and FMEA, in the Industry 5.0 era, the analysis of root causes is evolving towards a modern and highly automated approach [150]. The high operational complexity of modern turbomachinery renders purely manual diagnosis obsolete, necessitating a shift towards Data-Driven methodologies, as widely discussed in the literature [151]. In the current industrial landscape, established frameworks rely on probabilistic and causal reasoning. Bayesian Networks (BN) assume a central role: utilizing joint probability and a priori knowledge, they allow for precise Root-Cause Isolation by tracing fault propagation across subsystems. Alongside these, Fuzzy Cognitive Maps (FCM) are widely employed to model causal relationships under conditions of uncertainty, providing a logical layer that is easily interpretable by operators. To handle the complexity of raw sensor streams, these reasoning tools are increasingly integrated with Deep Learning architectures. In this context, CNNs are adopted to extract specific fault signatures (feature extraction) rather than merely for classification, while LSTM networks and Attention mechanisms have become the de facto standard for capturing long-term temporal dependencies. However, as highlighted by Pietsch et al. [35], the massive adoption of these algorithms encounters the limit of sparse semantics: although Deep Learning is powerful in detecting numerical correlations, it often fails to explain the physical reasons. To bridge this gap, the most advanced research is moving towards GNN to preserve the machine topology and LLM to process unstructured knowledge such as maintenance logs. This shift paves the way for a Hybrid Intelligence, where computational power is validated by the presence of the expert (human-in-the-loop). This synergy perfectly embodies the Industry 5.0 paradigm: a human-centric approach where technology supports the operator, enhancing irreplaceable skills such as complex decision-making and advanced problem-solving. Accordingly, this section analyzes recent works related to turbomachinery RCA, focusing exclusively on purely data-driven models.

Frontiers of Root Cause Analysis in Turbomachinery

As prognostics evolve, research has increasingly targeted the automation of RCA to shift diagnostics from reactive to proactive by integrating probabilistic models, graphs, and artificial intelligence. In this field, Gugaliya [152] proposes an approach based on Multi-PCA for gas turbines. This method enhances traditional PCA capabilities by managing system non-linearities through the clustering of operational data; the added value lies in

isolating faults via the identification of statistical deviations in individual sensors, acting as a link between classical statistics and machine learning.

To enhance explainability, C. Zhang et al. [108] introduce an AIOps framework for wind turbines using LSTM-AVAGMM models and the SHAP technique. The application of SHAP addresses the opacity of deep learning by quantifying the contribution of each sensor to the onset of the anomaly, providing a transparent diagnosis to the operator.

The most recent research frontier explores the use of graph-based architectures to respect the physical topology of the machinery. L. Yuan et al. [153] present the GAT-BN model, which combines Graph Attention Networks and Bayesian Networks for aircraft engines. This architecture allows for the autonomous learning of node criticality, modeling fault propagation along the physical structure of the engine. Building on this topological perspective, Mourya and Tyagi [154] propose a system based on GNN and Dynamic Bayesian Networks for offshore turbomachinery, introducing the concept of the Maintenance Causality Loop to bridge the gap between siloed signals and prioritized maintenance action.

Real-time diagnostic capabilities are further expanded by Ofoedu et al. [155], who employ LSTM networks to capture temporal causality in sequential data and Bayesian Networks to model probabilistic dependencies in offshore process failures. Their hybrid framework effectively integrates distributed sensor streams to isolate root causes with high accuracy; specifically, the combined use of deep learning for time series and probabilistic inference allows for managing the intrinsic uncertainty of live data, drastically reducing false alarms compared to threshold-based systems. Finally, generative AI is transforming the management of unstructured knowledge. Vitale [156] explores the use of LLMs and information retrieval for gas turbines and compressors, translating natural language descriptions into structured diagnostic insights. Complementing this approach, R. Shi et al. [157] propose an intelligent RCA method for steam turbines based on Knowledge Graphs and LLMs. Their system constructs a domain-specific graph of turbine components (e.g., blades, bearings) and employs the reasoning capabilities of the language model to identify root causes of quality accidents and operational failures. Studies addressing RCA methods are listed in Table 6.

Table 6. Summary of recent literature on RCA and advanced diagnostics for turbomachinery.

Work	Learning Paradigms	System	Year	Input	Metrics	Comp.
[152]	Multi-PCA and HDBSCAN Clustering	Gas Turbine	2020	MTS	F1-score, Acc	L
[156]	Generative AI and LLMs	Turbomachinery	2024	MTS	F1-score, Acc	H
[157]	Knowledge Graph and LLM	Steam Turbine	2025	Physics + MTS	F1-score, Acc	H
[153]	Improved Bayesian Network and GAT	Aircraft Engine	2026	ST-MTS	F1-score, Acc	M
[154]	AI-Enabled Operational Intelligence	Offshore O&G	2025	MTS + VIB	F1-score, Acc	M
[108]	LSTM-AVAGMM and SHAP-based RCA	Wind Turbine	2024	MTS	F1-score, AUC	M
[155]	LSTM and Bayesian Networks	Offshore Processes	2022	MTS	F1-score, Acc	M

Note: PIL: Physics-Informed Loss; PIA: Physics-Informed Architecture; PDDG: Physics-Driven Data Generation; MTS: Multivariate Time Series; ST-MTS: Spatio-Temporal MTS; VIB: Vibration; IMG/SPEC: Image/Spectrogram; Phys: Physics-based inputs; RMSE: Root Mean Square Error; MAE: Mean Absolute Error; Acc: Accuracy; F1: F1-score; AUC: Area Under the Curve.

Reflecting on the diagnostic strategies for root cause identification, while the emergence of XAI and Large Language Models provides a semantic bridge for human operators, the formal integration of physical consistency and the mitigation of technical hallucinations in critical operations remain the primary challenges, as detailed in “Section Research Gaps and Future Perspectives”.

7. Comparative Analysis

To evaluate architectures objectively, five synthetic indices are introduced to reflect fundamental aspects of industrial monitoring:

- **Performance Index (P-Index):** Quantifies diagnostic reliability by prioritizing the F1-score or, in its absence, Accuracy, AUC, and prediction error or normalized prediction error.
- **Data Efficiency Index (DE-Index):** Evaluates the volume of historical data required to make the system operational, rewarding models capable of converging with minimal training sets compared to architectures necessitating extensive data baselines.
- **Inference Time Index (IT-Index):** Assesses responsiveness and suitability for deployment on limited hardware, distinguishing between fast parallelizable architectures and dynamic models burdened by the recalculation of spatial relationships.
- **Interpretability Index (I-Index):** Measures the ability to support the identification of fault causes, distinguishing between black-box models and transparent architectures that allow for the explicit visualization of the sensors driving the alarm.
- **Generalization Capability Index (GC-Index):** Quantifies the readiness of the model to be extended to a fleet without the need for retraining. In turbomachinery, generalization is hindered by sensor tolerances, the environment, and the degradation state of systems.

7.1. Methodology for Synthetic Index Definition

To ensure reproducibility and objectivity, the assigned scores derive from a meta-analytic synthesis algorithm. This approach allows for mapping the technical characteristics and numerical results of the surveyed works onto a common metric.

7.1.1. Methodological Derivation of the P-Index

The performance index is determined according to a rigorous protocol based on the following eight rules:

1. **F1-score Priority:** This constitutes the primary metric for monitoring. Given the rarity of faults in the sector, the F1-score is the most reliable indicator for evaluating the actual anomaly detection capability on imbalanced classes.
2. **Accuracy/AUC Hierarchy:** In the absence of the F1-score, Accuracy or AUC is adopted as a proxy indicator. Operating already on a 0–1 scale, these values are integrated directly into the calculation.
3. **Inverse Normalization of Errors:** Utilized in the presence of reconstruction or prognostic errors when F1-score, Accuracy, or AUC are unavailable. To compare such error metrics with performance indicators, Root Mean Square Error (RMSE), Mean Absolute Error (MAE), and Mean Squared Error (MSE) values are transformed using inverse linear normalization:

$$P_i = 1 - \frac{M_i - M_{min}}{M_{max} - M_{min}}$$

where P_i represents the index result for the i -th work and M_i is the reported error metric value. The parameters M_{min} and M_{max} are the minimum and maximum values, respectively, extracted within the same methodological family or related to the same benchmark, especially when such data results are isolated (i.e., present in only 1 or 2 works).

4. **MAPE Management:** In the absence of RMSE and MAE, the Mean Absolute Percentage Error (MAPE) is treated analogously to other error indicators, being normalized and inverted to preserve statistical information in forecasting works.

5. Heterogeneous Metrics Management: If a sub-family presents mixed metrics, individual normalization is performed for each work reporting prediction errors, making them comparable with native 0–1 scale values (F1-score, Accuracy, AUC).
6. Maximum Value Rule (SOTA): For each study, the result of the best-performing configuration is extracted, thus mapping the maximum potential of the proposed architecture.
7. Single Case Management (External Benchmark): If an error metric appears in a very small number of works (1 or 2) within a sub-family, hindering a meaningful internal normalization, the value is compared against a consolidated external benchmark or the known SOTA for that system.
8. Pooled Average Calculation: The final P-Index of the sub-family is the global arithmetic mean of all individual P_i values, ensuring that each study contributes equally to the final statistical weight.

Certain methodological boundaries apply to the proposed framework. Scoring depends on author-reported metrics of varying rigor, while arithmetic aggregation weights all studies equally regardless of research scale. Additionally, diverse datasets and protocols hinder direct cross-study comparisons. The resulting indices serve as relative benchmarks for identifying architectural trends rather than absolute performance measurements.

7.1.2. Scoring Criteria for Qualitative Indices (DE, IT, I, GC)

The indices related to architectural and operational properties are determined through an ordinal mapping of the declared technical characteristics. The final score of each sub-family is obtained by calculating the average of the values assigned to individual works according to the following scoring structure Table 7:

Table 7. Ordinal scoring rubric for qualitative synthetic indices based on architectural and operational features.

Value	DE-Index	IT-Index	I-Index	GC-Index
1.0	Physics-Informed or Few-Shot paradigms. Requires minimal training sets due to prior physical knowledge.	Low (L) complexity. Shallow models or parallelizable TCNs ready for real-time edge computing.	Physics-explicit architectures or structural transparency models allowing the visualization of sensors driving the alarm.	Fleet-wide or Zero-Shot systems. Validation on machines different from those used in training.
0.7	Semi-supervised or Hybrid paradigms. Utilizes unlabeled data or partial physical regularization.	Medium (M) complexity. Deep LSTM or Transformers. Requires GPU acceleration but is manageable in industrial contexts.	Models with Attention Weights or GNNs. Visualization of critical sensors (structural XAI).	Validation on real-world data, multi-regime, and with sensor noise, limited to a specific variant.
0.3	Pure DL paradigms. Requires massive historical databases and full seasonal cycles for convergence.	High (H) complexity. Iterative PINNs or massive GANs. High computational load or cloud latency.	Black-Box models (Standard LSTM/CNN). They signal anomalies without providing evidence regarding the root causes.	Validation on a single system or exclusively on synthetic data (numerical simulations).

In this framework, the DE-Index assignments primarily reflect a qualitative synthesis of the training requirements inherent to each paradigm. Higher scores (1.0) are typically associated with Physics-Informed or Few-Shot models that utilize prior knowledge to reach convergence with limited samples, whereas lower scores (0.3) generally characterize Purely Data-Driven architectures that often necessitate extensive historical baselines to effectively model system dynamics. We acknowledge that this meta-analytic synthesis, while based on rigorous criteria, is influenced by author-reported data and the specific experimental configurations of the original works. The methodology, however, allows for

overcoming purely qualitative assessments, providing a quantitative guide for the adoption of AI solutions in the turbomachinery sector.

Relying on these evaluative metrics, the following subsections systematically assess and benchmark the methodological architectures across the three primary PdM tasks: Anomaly and Fault Detection, RUL estimation, and RCA.

7.1.3. Illustrative Example: GC-Index Calculation for Graph-Based Methods

To address the requirement for methodological transparency, Table 8 illustrates the scoring process for the Generalization Capability Index (GC-Index) within the “Graph-based Reconstruction” sub-family. This example demonstrates how qualitative features extracted from the literature are mapped to ordinal scores (0.3, 0.7, 1.0) to derive the final aggregate value.

As shown, the resulting index of 0.6 is not an arbitrary assignment but the result of a reproducible aggregation of the architectural characteristics and validation setups declared by the authors. Similar procedures were applied to all papers for the calculation of P, DE, IT, and I indices, ensuring a standardized benchmarking across different methodological families.

Table 8. Detailed scoring example for the GC-Index in Graph-based Reconstruction methods.

Work [Ref]	Validation Framework	Score	Rationale (from Table 7)
Wang [85]	Validated on synthetic datasets with limited noise scenarios.	0.3	Synthetic data/Single regime
Zhan [69]	Tested on standard benchmark datasets without cross-domain validation.	0.3	Single system/Lab-scale
Luo [87]	Validated on multi-sensor industrial data with varying operating conditions.	0.7	Real-world data/Multi-regime
Liang [94]	Validated across different turbine fleets with structural prior knowledge.	1.0	Fleet-wide/Zero-Shot
Aggregate	Arithmetic Mean: $(0.3 + 0.3 + 0.7 + 1.0)/4$	0.575	≈ 0.6 (Final GC-Index)

7.2. Anomaly and Fault Detection: Methodological Benchmarking

This subsection evaluates Anomaly and Fault Detection methods, explicitly contrasting the Purely Data-Driven family against the Hybrid Physics-Informed one. Subsequently, physics-guided methods are structured and analyzed according to their specific integration paradigm. In conclusion, a Global Comparison aggregates the average results to objectively benchmark the overall operational impact between these two families concerning Anomaly Detection.

7.2.1. Comparative Analysis for the Purely Data-Driven Family

Purely data-driven methods (forecasting, reconstruction, and contrastive) are analyzed here by dividing them into specific sub-families; this approach allows for the grouping of works that employ similar Deep Learning architectures, thereby evaluating their varying performances according to the indices defined previously in this section.


Forecasting-Based Architectures

The examination of the works focused on forecasting allows for a clear mapping of the technological evolution in turbomachinery monitoring, which is shifting from simple temporal analysis toward models capable of mapping the physical layout and relationships between different sensors. Within this landscape, four main architectural sub-families have been identified, as summarized in Table 9.

Table 9. Summary of architectural sub-families for forecasting-based methods.

Sub-Family	Key Features & Applications	References
GNN	Represents the dominant and most innovative trend, capable of mapping spatial and topological relationships between sensors.	[67–73]
RNN/LSTM	The established standard for handling complex sequential and temporal dependencies.	[61–64]
CNN/TCN	Dedicated to computational efficiency and highly suitable for edge monitoring applications.	[65,66]
Statistical Hybrid	Employs classical models (e.g., GAM, ARFIMA) for operational contexts characterized by high stationarity.	[60]

Once the individual indices were calculated for each work in the forecasting group, the values were aggregated by family using the arithmetic mean to generate the Radar Plots presented in Figure 8. This quantitative synthesis not only highlights the strengths of each paradigm but also provides a decision-making guide for selecting the most suitable architecture based on specific industrial constraints. Despite GNNs being the most prominent sub-family in recent literature, the metrics highlight a significant trade-off: while these architectures excel in interpretability and spatial accuracy, they demand considerably higher computational resources and data volumes. In contrast, LSTM and TCN-based models remain highly efficient solutions, offering a superior balance between execution speed and industrial readiness, albeit at the expense of some diagnostic transparency.

**Figure 8.** Comparative analysis of forecasting architectural sub-families.

Reconstruction-Based Architectures

Adopting the same evaluation protocol used previously, we examined the 22 contributions belonging to the Reconstruction Methods macro-family. The literature analysis allowed for the identification of four distinct architectural sub-families, as summarized in Table 10.

The average performance of these sub-families, calculated according to the five introduced indices, is visualized in the Radar Plot in Figure 9. The data reveal that the Recurrent AE sub-family emerges as the most balanced solution, guaranteeing maximum diagnostic reliability with moderate computational costs. However, when the primary constraint is execution speed, such as in deployment on edge systems or high-frequency vibration analysis, the Convolutional sub-family clearly dominates in terms of temporal efficiency, with only minimal sacrifice in precision. A notable result concerns the Graph-based sub-family: although computationally more demanding, it excels not only in interpretability but also in generalization capability. This result is driven by recent contributions such as ExpertAP [94], demonstrating how the graph structure facilitates knowledge transfer (Transfer Learning) between different units of a fleet, overcoming the limits of traditional models trained on a single unit.




Figure 9. Comparative analysis of reconstruction architectural sub-families.

Table 10. Summary of architectural sub-families for reconstruction-based methods.

Sub-Family	Key Features & Applications	References
RNN-AEs	Represents the most numerous and consolidated group, effectively constituting the industrial standard for the reconstruction of complex time series.	[76,77,81,83,84,89,91,93]
Convolutional (CNN/TCN)	Focused on computational efficiency through the use of spatial or temporal filters for fast feature extraction.	[26,75,78–80,82,86]
Graph-based	The most recent cluster, specifically oriented toward capturing topological dependencies and interactions between sensors.	[69,85,87,94]
Generative (GAN/VAE)	Offers probabilistic approaches to deeply model the statistical distribution of nominal data.	[88,90,92]

Contrastive-Based Architectures

The analysis of contributions related to Contrastive Learning highlights how this paradigm is redefining the limits of generalization in industrial anomaly detection. Unlike reconstructive methods, which focus on point-wise error, these approaches learn a metric space where distinct operating conditions are geometrically separated. Within this landscape, four main architectural sub-families have been identified, as summarized in Table 11.

Table 11. Summary of architectural sub-families for contrastive-based methods.

Sub-Family	Key Features & Applications	References
Metric & Matching	Directly optimizes the distance between classes using Siamese architectures or Triplet losses for high-accuracy applications.	[80,96,101,108]
Graph & Temporal	Extends the contrastive metric space to complex data structures, capturing spatio-temporal dependencies.	[91,95,99,100]
Multimodal & Visual	Represents an emerging trend that integrates visual data or semantic descriptions through the use of LLMs.	[102,103,105,107]
Transfer & Few-Shot	Employs pre-training strategies to operate effectively in scenarios characterized by extreme data scarcity (One-shot learning).	[97,98,104]

The performance profiles, illustrated in the Radar Plot in Figure 10, reveal significant operational trade-offs across these groups. The Metric & Matching sub-family proves to be the highest performing in terms of inference speed and pure accuracy, making it ideal for standard applications with stringent latency requirements. However, the innovation lies elsewhere: the Transfer & Few-Shot sub-family dominates in data efficiency, demonstrating superior capabilities in adapting to new machines with very few examples through One-shot learning strategies. In parallel, the integration of LLMs within the Multimodal sub-family has pushed interpretability and generalization to unprecedented levels, providing natural language explanations for detected anomalies, albeit at the cost of high computational demand, which currently limits real-time responsiveness.

7.2.2. Comparative Analysis for the Hybrid Physics-Informed Family

In this subsection, the methodologies belonging to the Hybrid Physics-Informed family are structured and divided into specific sub-families based on their integration paradigms. This approach allows for a targeted analysis of their performance variations and operational trade-offs across different structural approaches. The integration of domain knowledge into deep learning models represents the frontier for reliable prognostics in safety-critical systems. Analyzing the 20 selected works, we identified three distinct integration strategies, as summarized in Table 12.




Figure 10. Comparative analysis of contrastive architectural sub-families.

Table 12. Summary of architectural sub-families for hybrid physics-informed methods.

Sub-Family	Key Features & Applications	References
PIL	Embeds physical laws (e.g., fluid dynamics or energy conservation) as regularization terms in the loss function.	[114–119,127–130,132]
PIA	Embeds physical constraints directly into the neural network layers, such as using specialized attention mechanisms or signal processing layers.	[120–122,126]
PDDG	Utilizes thermodynamic simulations or physical models to generate synthetic faulty data, addressing the chronic lack of run-to-failure records.	[109,123–125,131]

The quantitative assessment in Figure 11 reveals clear operational distinctions. The PIA sub-family achieves the highest Data Efficiency Index, as the structural constraints significantly reduce the volume of training data required to reach convergence. Conversely, the PDDG approach excels in inference speed, as the physics is offloaded to the training phase, leaving a lightweight standard model for deployment. The PIL methods offer the best balance for diagnostic transparency, as the adherence to physical laws provides an intrinsic validation of the model's predictions, although this often comes at the cost of increased training complexity.




Figure 11. Comparative analysis of physics-informed integration strategies.

7.2.3. Global Comparison: Purely Data-Driven vs. Hybrid Physics-Informed

To characterize the global operational behavior of the reviewed architectures for Anomaly and Fault Detection, performance metrics were aggregated and averaged to directly contrast the Purely Data-Driven family against the Hybrid Physics-Informed family. This overarching analysis highlights the macroscopic trade-offs between empirical learning and domain-knowledge integration, guiding industrial model selection based on the defined performance indices. The quantitative results demonstrate that Hybrid Physics-Informed architectures excel in overall reliability (P-Index) and data efficiency (DE-Index). The mathematical embedding of physical constraints acts as a powerful regularizer, effectively compensating for the scarcity of labeled run-to-failure samples. Furthermore, these models achieve a higher Interpretability Index (I-Index), as the adherence to physical laws inherently provides a validation of the network's internal logic. Conversely, while the Purely Data-Driven family offers a highly versatile baseline, its aggregated scores highlight limitations in Generalization (GC-Index) across varying operating conditions. Notably, the Physics-Informed family also exhibits superior average Inference Time scores (IT-Index). This counterintuitive result is explained by the fact that strategies like PDDG and PIL offload the heavy computational burden of physical simulations entirely to the offline training phase, leaving a lightweight and highly responsive model for real-time deployment, whereas the Data-Driven average is penalized by computationally intensive dynamic graph models. The trends emerging from this methodological dichotomy are summarized graphically in the radar diagram in Figure 12.




Figure 12. Global comparative analysis of anomaly detection families: Purely Data-Driven vs. Hybrid Physics-Informed. Note: The high Inference Time Index (IT-Index) of the hybrid family stems from offloading physical computations to the training phase, as explained in Section 7.2.3.

7.3. Operational Profiling of RUL Models: Pure Deep Learning vs. Physics-Guided

Following the overarching taxonomic split established in Section 5, the performance evaluation of RUL models is structured by explicitly contrasting the Pure Deep Learning family against the Physics-Guided family. This binary classification distills the evolution of prognostic methodologies, moving from purely statistical algorithms toward hybrid architectures capable of embedding component degradation mechanisms. The main characteristics and references for these two domains are summarized in Table 13.

Table 13. Summary of architectural families for RUL prediction methods.

Family	Key Features	References
Pure Deep Learning	Employs purely statistical and data-driven algorithms to model degradation mechanisms, ensuring fast real-time predictions without heavy thermodynamic constraints.	[138–145,149]
Physics-Guided	Integrates physical domain knowledge (e.g., fatigue laws, digital twins) acting as a robust mathematical regularizer to ensure diagnostic reliability and data efficiency.	[110,124,137,146–148]

The quantitative assessment of these two domains, visualized in the Radar Plot in Figure 13, reveals significant operational trade-offs. The Physics-Guided family overwhelmingly dominates in the DE-Index and the I-Index. The integration of physical domain knowledge, such as fatigue and creep laws or high-fidelity digital twin simulations, acts as a robust mathematical regularizer; this structural constraint allows the network to converge effectively even with extremely limited historical run-to-failure samples, particularly when synthetic degradation trajectories are utilized to augment the dataset, and offers an intrinsically transparent logic that justifies the remaining life estimation.

Conversely, the Pure DL family, while exhibiting lower structural transparency, presents a higher IT-Index. The advantage stems from optimized architectures (CNNs/LSTMs) that bypass physical constraints during inference, while hybrid models offload this complexity to the training phase. However, recent hybrid frameworks, including dedicated acceleration modules, are beginning to close this computational gap. Regarding absolute diagnostic reliability (P-Index), the physics-informed approaches maintain a slight superiority, confirming that constraining a neural network within physically plausible bounds prevents inconsistent extrapolations, especially as the component approaches the end of its life cycle.




Figure 13. Comparative radar plot illustrating the operational trade-offs between Pure Deep Learning and Physics-Guided models for RUL estimation.

7.4. Root Cause Analysis: Methodological Benchmarking

The automation of fault diagnosis in turbomachinery requires not merely detecting anomalies but precisely isolating their origin. Given the statistical and data-driven nature of the selected RCA methodologies, the evaluation framework defined in the previous section was maintained. Analyzing the selected works, we identified four distinct methodological sub-families, as summarized in Table 14.

Table 14. Summary of methodological sub-families for Root Cause Analysis.

Sub-Family	Key Features	References
Causal & Graph	Utilizes tools like Bayesian Networks not only to localize anomalies but also to reconstruct the directional propagation of faults, mapping causal influences between signals.	[153–155]
Knowledge-Augmented	Employs Generative AI, LLMs, and Knowledge Graphs to reason over unstructured data, delivering high interpretability.	[156,157]
Explainer (XAI)	Focuses on making “black-box” models transparent through feature attribution techniques.	[108]
Statistical/Multivariate	Relies on dimensionality reduction to isolate pure mathematical deviations with low computational cost.	[152]

The quantitative assessment in Figure 14 reveals clear operational trade-offs. It is essential to note that for the calculation of the P-Index, priority was given to Accuracy over F1-score in this context. While the F1-score is the standard for anomaly detection due to class imbalance, RCA studies typically operate on balanced subsets of known fault categories. In such scenarios, Accuracy provides a more direct measure of the model’s ability to unequivocally identify the faulty component among competing failure modes, effectively guiding maintenance decisions and preventing erroneous interventions. Under this lens, the Causal & Graph and Explainer methods exhibit the most robust performance. However, the Knowledge-Augmented sub-family sets the benchmark for the I-Index: the use of LLMs allows for delivering diagnoses in natural language, effectively bridging the semantic gap between sensor data and human understanding. Conversely, the Statistical/Multivariate approach excels in inference speed (IT-Index), owing to the lower computational cost of linear algebra compared to deep neural networks.




Figure 14. Comparative analysis of root cause analysis strategies.

7.5. Industrial Deployment: Technical Trade-Offs and Decision Guidelines

Building on the previous methodological analysis, this section synthesizes the core technical trade-offs and provides practical guidelines for model selection based on industrial infrastructure. The objective is to facilitate deployment by balancing diagnostic accuracy with real-world operational constraints.

7.5.1. Global Synthesis of Technical Trade-Offs

The aggregated metrics reveal three fundamental tensions that govern architectural selection in turbomachinery monitoring:

- **Reliability vs. Training Complexity (P-Index vs. Complexity):** Hybrid Physics-Informed models achieve superior reliability and data efficiency by embedding physical laws as mathematical constraints [16]. While their high complexity (Class H) requires significant offline training effort, they provide an important industrial advantage: they maintain a high Inference Time Index (IT-Index) during deployment by offloading physical simulations to the training phase.
- **Interpretability vs. Deployment Versatility (I-Index vs. Purely Data-Driven):** Graph-based architectures (GNNs) represent the benchmark for interpretability by mapping the physical sensor topology. However, they require a detailed structural map of the asset. Conversely, Purely Data-Driven models (e.g., RNN-AEs) are more versatile and quicker to deploy across different systems, though they operate as “black boxes” with lower diagnostic transparency.

- Data Efficiency vs. Generalization (DE-Index vs. GC-Index): Contrastive Learning frameworks excel in the GC-Index, making them ideal for managing diverse fleets. However, in scenarios with a total absence of historical faulty data, Physics-Informed solutions and performance Digital Twins are unparalleled (High DE-Index), as physical constraints serve as a surrogate for missing empirical information [16,83].

7.5.2. Guidelines for Industrial Implementation

Model selection must align with specific hardware capabilities and operational environments, focusing on the balance between edge and cloud processing:

- Scenario A: Real-time Edge Monitoring: For deployments integrated into local SCADA or edge devices with limited hardware, Convolutional (CNN/TCN) or lightweight RNN architectures are recommended [158]. Their balance between execution speed and resource demand minimizes network latency and optimizes streaming data management in decentralized nodes [159].
- Scenario B: High-Fidelity Diagnostics for Critical Assets: In safety-critical contexts (e.g., steam turbines or nuclear valves), Hybrid Physics-Informed (PINN) models and Digital Twins with uncertainty quantification are the primary choice [16,83]. Their physical consistency meets non-negotiable safety requirements; however, while the offline training and high-fidelity synchronization of such models require scalable cloud infrastructures to handle the underlying mathematical complexity, their real-time inference remains lightweight and suitable for responsive monitoring [158].
- Scenario C: Large-scale Fleet Management: For managing distributed assets, such as offshore wind farms, cloud-based remote monitoring platforms are the most efficient [160]. These architectures facilitate Transfer Learning strategies, reducing the need for unit-specific retraining and minimizing long-term O&M costs [16,160].

The practical adoption of these architectures depends on their alignment with international standards, such as ISO 13374 for data processing and ISO 13379-1 for diagnostic interpretation [161,162]. Specifically, the Interpretability Index (I-Index) serves as a metric for diagnostic traceability, reflecting the transparency requirements mandated by these norms. Architectures with high I-Index scores facilitate the verification of fault root causes, ensuring that AI-driven insights remain auditable and compliant with industrial safety protocols. Mapping the evaluation framework to normative compliance reinforces the transition from experimental black-box models to certified monitoring solutions.

8. Conclusions

This work presents a systematic and thorough review of the state of the art in PdM applied to the critical turbomachinery sector. Through the analysis of recent literature, primarily concentrated over the last five or six years, the core pillars composing the modern PdM pipeline have been examined: from diagnostics, focused on anomaly and fault detection, to prognostics for RUL and the complex discipline of RCA. Various frameworks used in the presented works were observed, providing not only a general overview but also analyzing the positive and negative aspects of using these methods. As a technical and application-oriented review, the analysis focused not only on architectural theory but also on the management of dataframes and real-world data flows [20,35,47]. In the field of diagnostics, the approach chosen to classify scientific contributions follows a clear tripartition of forecasting, reconstruction, and contrastive learning strategies [38]. The quantitative benchmarking through synthetic indices revealed that architectural choice is strictly driven by industrial trade-offs: GNN-based models excel in the Interpretability Index by mapping sensor topologies, while RNN and TCN architectures remain the most viable solutions for real-time edge deployment due to their superior Inference Time Index [67,71].

These methods are increasingly supplemented by the integration of physics-informed components, including specific architectures like PINNs [109,121,126]. This study highlights that Physics Integration, particularly the PIL and PIA approaches, acts as a primary enabler for the Data Efficiency Index, allowing models to converge with significantly reduced historical datasets and overcoming the traditional requirement for full seasonal operational cycles [110,114,127]. Regarding time-series forecasting, advanced frameworks utilized hybrid approaches to capture transient turbine dynamics [67,72], while signal reconstruction-based architectures, such as those employing spatial and temporal autoencoders [85,89], proved essential for identifying subtle anomalies in noisy environments. Regarding RUL estimation, the integration of physics and Deep Learning emerged as a winning strategy. Works based on physics-informed Deep-LSTM architectures [146] and hybrid models for modeling mechanical fatigue in bearings [110] demonstrated how domain knowledge can stabilize predictions over long time horizons. In parallel, RCA has undergone a topological and semantic revolution. The quantitative assessment identifies Causal & Graph approaches, such as improved Bayesian networks [153], as essential for reconstructing the directional propagation of faults. Furthermore, the pioneering use of LLMs [156,157] has set a new benchmark for the Interpretability Index, introducing automated reasoning capabilities that effectively bridge the semantic gap between raw sensor data and human understanding. Overall, this work confirms that PdM for turbomachinery is now a mature digital ecosystem, where the fusion of data and physics guarantees higher levels of reliability, safety, and sustainability in line with Industry 5.0 requirements.

However, to ensure a transparent and objective perspective, it is essential to acknowledge the limitations of this review itself. First, a potential publication bias may exist in the surveyed literature, as studies reporting positive or highly accurate results are inherently more likely to be published. Second, the search scope was restricted to English-language publications, which may exclude relevant theoretical advancements or industrial applications documented in other languages. Lastly, the deliberate focus on post-2015 deep learning and hybrid AI trends means that some long-validated traditional methods, which remain the standard practice in certain industrial sectors, might not be fully represented in this analysis.

Research Gaps and Future Perspectives

Despite the progress described, several technical bottlenecks define future research directions:

- **Generalizability and Domain Shift:** Purely data-driven models often struggle to adapt across different engine variants or operational sites, requiring costly unit-specific retraining to handle varying sensor noise and operating envelopes [16,38].
- **Edge Scalability:** Transitioning from single-unit success to fleet-wide monitoring exposes hardware constraints in edge nodes, demanding more robust model compression techniques and secure, low-latency data pipelines [159,160].
- **Hybrid Integration Challenges:** Balancing physical consistency with empirical data in a joint loss function remains mathematically complex, particularly for aging assets where degradation and fluctuating environmental conditions make nominal physical laws increasingly imprecise [16,83].

The benchmarking analysis conducted in Section 7 provides a direct empirical grounding for these directions. The low GC-Index recorded for purely data-driven architectures quantitatively confirms the severity of the domain shift problem, motivating the development of future models that explicitly incorporate fleet-level transfer mechanisms, such as Few-Shot contrastive pre-training validated across heterogeneous turbine populations. The DE-Index advantage consistently observed for Physics-Informed approaches points

toward the development of lightweight PIA variants deployable on constrained edge hardware, directly addressing the scalability bottleneck identified above. Finally, the IT-Index penalty associated with dynamic graph models highlights the need for architectural compression strategies that preserve topological expressiveness without sacrificing real-time responsiveness. This progression is consistent with the hybrid AI maturity roadmap outlined by [16]. It confirms that the directions identified here represent validated next steps rather than speculative research avenues. Addressing these gaps is essential to move beyond isolated numerical RUL estimations toward operationally integrated monitoring systems. To bridge the semantic gap that limits the practical impact of even high-performing models, a promising direction lies in the evolution of LLMs as engines for contrastive learning [95,105], creating a semantic bridge between AI outputs and operational decisions. Integrating LLMs with explainability techniques [107] enables a transition toward prescriptive maintenance through a three-stage architecture: a predictive core identifies anomalies, while a semantic translator based on Retrieval-Augmented Generation (RAG) cross-references findings with technical manuals [156]. The objective is a conversational interface explaining not only when a component will fail but also why, suggesting immediate actions.

However, the deployment of LLMs in safety-critical turbomachinery environments is constrained by the risk of technical hallucinations, where models may generate ungrounded or factually false diagnostic prescriptions [163]. To ensure operational safety, future research must implement proactive mitigation strategies, such as utilizing physics-informed frameworks to evaluate LLM-generated scenarios against established physical constraints [164]. Furthermore, managing the computational trade-offs remains a fundamental challenge; while large-scale models offer superior reasoning, the energy sector and edge environments require efficient architectures to balance inference latency with the high energy demands of these models [165].

Alongside these challenges, a significant gap exists regarding standardization and interoperability. The lack of open standards currently hinders the large-scale adoption of AI models across industrial fleets [16]. Closely tied to this is the alignment between scholarly methodologies and the normative frameworks governing real industrial deployments. Standards such as ISO 13374 [166] and ISO 13379-1 [167] impose strict constraints on model validation and transparency that purely data-driven approaches do not inherently satisfy. Additionally, reconciling academic research with industrial ‘gray literature’, such as technical white papers and OEM best practices, represents a pivotal direction. Aligning research outputs with normative compliance ensures that data-driven solutions meet the operational requirements of the turbomachinery sector. Ultimately, the role of the human operator (Human-in-the-loop) remains the core element of this evolution. In line with Industry 5.0 paradigms, artificial intelligence is not intended to replace the expert but to enhance their decision-making capabilities. Regarding the implementation of HITL, two main paths emerge: LLM interpreters and XAI-enhanced RCA. While LLMs offer the most immediate solution for improving human–machine interaction through natural language [156], XAI-enhanced methods provide the technical transparency required for industrial safety [108]. We argue that an integrated approach, combining the interpretability of RCA with the accessibility of LLMs, represents the most viable trajectory for Industry 5.0 applications. The challenge for the future will therefore be the creation of digital mentoring systems capable of preserving the knowledge of senior experts and democratizing access to complex diagnoses [35].

Author Contributions: Conceptualization, P.D., D.N. and S.S.; methodology, P.D., D.N. and S.S.; investigation, P.D. and D.N.; resources, S.S.; writing—original draft preparation, P.D. and D.N.; writing—review and editing, P.D., D.N. and S.S.; supervision, S.S.; project administration, P.D. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Zhu, T.; Ran, Y.; Zhou, X.; Wen, Y. A survey of predictive maintenance: Systems, purposes and approaches. *arXiv* **2019**, arXiv:1912.07383.
2. Pusey, H.C. Turbomachinery condition monitoring and failure prognosis. *Sound Vib.* **2007**, *41*, 10–15.
3. Hein-Pensel, F.; Winkler, H.; Brückner, A.; Wölke, M.; Jabs, I.; Mayan, I.J.; Kirschenbaum, A.; Friedrich, J.; Zinke-Wehlmann, C. Maturity assessment for Industry 5.0: A review of existing maturity models. *J. Manuf. Syst.* **2023**, *66*, 200–210. [CrossRef]
4. Jezzini, A.; Ayache, M.; Elkhansa, L.; Makki, B.; Zein, M. Effects of predictive maintenance (PdM), Proactive maintenance (PoM) & Preventive maintenance (PM) on minimizing the faults in medical instruments. In *Proceedings of the 2013 2nd International Conference on Advances in Biomedical Engineering*; IEEE: New York, NY, USA, 2013; pp. 53–56.
5. Achouch, M.; Dimitrova, M.; Ziane, K.; Sattarpanah Karganroudi, S.; Dhouib, R.; Ibrahim, H.; Adda, M. On predictive maintenance in industry 4.0: Overview, models, and challenges. *Appl. Sci.* **2022**, *12*, 8081. [CrossRef]
6. Tazi, N.; Châtelet, E.; Bouzidi, Y. Using a hybrid cost-FMEA analysis for wind turbine reliability analysis. *Energies* **2017**, *10*, 276. [CrossRef]
7. Romanssini, M.; de Aguirre, P.C.C.; Compassi-Severo, L.; Girardi, A.G. A review on vibration monitoring techniques for predictive maintenance of rotating machinery. *Eng* **2023**, *4*, 1797–1817. [CrossRef]
8. Dowson, D. Metallurgical Failure Analysis of Steam Turbine, Compressor, and Hot Gas Expander Components. In *Proceedings of the Asia Turbomachinery & Pump Symposium*, Singapore, 12–15 March 2018.
9. Sepe, M.; Graziano, A.; Badora, M.; Di Stazio, A.; Bellani, L.; Compare, M.; Zio, E. A physics-informed machine learning framework for predictive maintenance applied to turbomachinery assets. *J. Glob. Power Propuls. Soc.* **2021**, *2021*, 1–15. [CrossRef]
10. Suman, A.; Morini, M.; Aldi, N.; Casari, N.; Pinelli, M.; Spina, P.R. A compressor fouling review based on an historical survey of asme turbo expo papers. *J. Turbomach.* **2017**, *139*, 041005. [CrossRef]
11. Shin, D.; Yang, J.; Tong, X.; Suh, J.; Palazzolo, A. A review of journal bearing thermal effects on rotordynamic response. *J. Tribol.* **2021**, *143*, 031803. [CrossRef]
12. Decker, T.; Jacobs, G.; Raddatz, M.; Röder, J.; Betscher, J.; Arneth, P. Detection of particle contamination and lubrication outage in journal bearings in wind turbine gearboxes using surface acoustic wave measurements and machine learning. *Forsch. Ingenieurwesen* **2025**, *89*, 17. [CrossRef]
13. Malla, C.; Panigrahi, I. Review of condition monitoring of rolling element bearing using vibration analysis and other techniques. *J. Vib. Eng. Technol.* **2019**, *7*, 407–414. [CrossRef]
14. Ding, A.; Ren, X.; Li, X.; Gu, C. Friction power analysis and improvement for a tilting-pad journal bearing considering air entrainment. *Appl. Therm. Eng.* **2018**, *145*, 763–771. [CrossRef]
15. Farhat, H.; Salvini, C. Novel gas turbine challenges to support the clean energy transition. *Energies* **2022**, *15*, 5474. [CrossRef]
16. Farhat, H.; Altarawneh, A. Physics-Informed Machine Learning for Intelligent Gas Turbine Digital Twins: A Review. *Energies* **2025**, *18*, 5523. [CrossRef]
17. Castellani, F.; Eltayesh, A.; Becchetti, M.; Segalini, A. Aerodynamic analysis of a wind-turbine rotor affected by pitch unbalance. *Energies* **2021**, *14*, 745. [CrossRef]
18. Liu, Y.; Jiang, X.; Ge, X.; Wei, M. A physics informed machine learning approach for performance degradation monitoring of gas turbine. In *Proceedings of the PHM Society Asia-Pacific Conference*, Tokyo, Japan, 11–14 September 2023; Volume 4.
19. Malik, T.H.; Bak, C. Full-scale wind turbine performance assessment using the turbine performance integral (TPI) method: A study of aerodynamic degradation and operational influences. *Wind Energy Sci.* **2024**, *9*, 2017–2037. [CrossRef]
20. Mołęda, M.; Małysiak-Mrozek, B.; Ding, W.; Sunderam, V.; Mrozek, D. From corrective to predictive maintenance—A review of maintenance approaches for the power industry. *Sensors* **2023**, *23*, 5970. [CrossRef]
21. Zhang, M.; Sui, F.; Liu, A.; Tao, F.; Nee, A. Digital twin driven smart product design framework. In *Digital Twin Driven Smart Design*; Elsevier: Amsterdam, The Netherlands, 2020; pp. 3–32.

22. Nunes, P.; Santos, J.; Rocha, E. Challenges in predictive maintenance—A review. *CIRP J. Manuf. Sci. Technol.* **2023**, *40*, 53–67. [CrossRef]
23. Zhang, W.; Yang, D.; Wang, H. Data-Driven Methods for Predictive Maintenance of Industrial Equipment: A Survey. *IEEE Syst. J.* **2019**, *13*, 2213–2227. [CrossRef]
24. Sanchez-Londono, D.; Barbieri, G.; Fumagalli, L. Smart retrofitting in maintenance: A systematic literature review. *J. Intell. Manuf.* **2023**, *34*, 1–19. [CrossRef]
25. Miele, E.S.; Bonacina, F.; Corsini, A. Deep anomaly detection in horizontal axis wind turbines using graph convolutional autoencoders for multivariate time series. *Energy AI* **2022**, *8*, 100145. [CrossRef]
26. Lee, G.; Jung, M.; Song, M.; Choo, J. Unsupervised anomaly detection of the gas turbine operation via convolutional auto-encoder. In *Proceedings of the 2020 IEEE International Conference on Prognostics and Health Management (ICPHM)*; IEEE: New York, NY, USA, 2020; pp. 1–6.
27. Zhang, C.; Yang, T. Anomaly detection for wind turbines using long short-term memory-based variational autoencoder wasserstein generation adversarial network under semi-supervised training. *Energies* **2023**, *16*, 7008. [CrossRef]
28. Michau, G.; Fink, O. Unsupervised transfer learning for anomaly detection: Application to complementary operating condition transfer. *Knowl. Based Syst.* **2021**, *216*, 106816. [CrossRef]
29. Rombach, K. Fault Diagnostics Under Label and Data Scarcity. Ph.D. Thesis, ETH Zurich, Zurich, Switzerland, 25 July 2023.
30. Algan, G.; Ulusoy, I. Image classification with deep learning in the presence of noisy labels: A survey. *Knowl. Based Syst.* **2021**, *215*, 106771. [CrossRef]
31. Cook, J. GE Vernova Talks AGP XPAND's Impact, Supply-Chain Resiliency. *Turbomachinery International*, 8 May 2025.
32. Surendra, A. ATOM: Digital Twin of Siemens Gas Turbine Fleet Operations. 2023. Available online: <https://www.anylogic.com/resources/case-studies/atom-digital-twin-of-siemens-gas-turbine-fleet-operations/> (accessed on 28 March 2026).
33. Lim, J.; Perullo, C.A.; Milton, J.; Whitacre, R.; Jackson, C.; Griffin, C.; Noble, D.; Boche, L.; Seachman, S.; Angello, L.; et al. The EPRI Gas Turbine Digital Twin—A Platform for Operator Focused Integrated Diagnostics and Performance Forecasting. In *Proceedings of the Turbo Expo: Power for Land, Sea, and Air*; American Society of Mechanical Engineers: New York, NY, USA, 2021; Volume 84966, p. V004T09A009.
34. Abdelillah, F.M.; Nora, H.; Samir, O.; Sidi-Mohammed, B. Hybrid Data-Driven and Knowledge-Based Predictive Maintenance Framework in the Context of Industry 4.0. In *Proceedings of the Model and Data Engineering*; Mosbah, M., Kechadi, T., Bellatreche, L., Gargouri, F., Eds.; Springer: Cham, Switzerland, 2024; pp. 319–337.
35. Pietsch, D.; Matthes, M.; Wieland, U.; Ihlenfeldt, S.; Munkelt, T. Root cause analysis in industrial manufacturing: A scoping review of current research, challenges and the promises of AI-driven approaches. *J. Manuf. Mater. Process.* **2024**, *8*, 277. [CrossRef]
36. Zhang, W.; Vatn, J.; Rasheed, A. A review of failure prognostics for predictive maintenance of offshore wind turbines. In *Proceedings of the Journal of Physics: Conference Series*; IOP Publishing: Bristol, UK, 2022; Volume 2362, p. 012043.
37. Zamanzadeh Darban, Z.; Webb, G.I.; Pan, S.; Aggarwal, C.; Salehi, M. Deep learning for time series anomaly detection: A survey. *ACM Comput. Surv.* **2024**, *57*, 1–42. [CrossRef]
38. Wang, F.; Jiang, Y.; Zhang, R.; Wei, A.; Xie, J.; Pang, X. A survey of deep anomaly detection in multivariate time series: Taxonomy, applications, and directions. *Sensors* **2025**, *25*, 190. [CrossRef] [PubMed]
39. Li, Z.; Zhao, Y.; Han, J.; Su, Y.; Jiao, R.; Wen, X.; Pei, D. Multivariate time series anomaly detection and interpretation using hierarchical inter-metric and temporal embedding. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*; ACM: New York, NY, USA, 2021; pp. 3220–3230.
40. Munir, M.; Chattha, M.A.; Dengel, A.; Ahmed, S. A comparative analysis of traditional and deep learning-based anomaly detection methods for streaming data. In *Proceedings of the 2019 18th IEEE International Conference on Machine Learning and Applications (ICMLA)*; IEEE: New York, NY, USA, 2019; pp. 561–566.
41. Darban, Z.Z.; Webb, G.I.; Pan, S.; Aggarwal, C.C.; Salehi, M. CARLA: Self-supervised contrastive representation learning for time series anomaly detection. *Pattern Recognit.* **2025**, *157*, 110874. [CrossRef]
42. Yue, Z.; Wang, Y.; Duan, J.; Yang, T.; Huang, C.; Tong, Y.; Xu, B. Ts2vec: Towards universal representation of time series. In *Proceedings of the AAAI Conference on Artificial Intelligence*; AAAI: Washington, DC, USA, 2022; Volume 36, pp. 8980–8987.
43. Zhang, X.; Zhao, Z.; Tsiligkaridis, T.; Zitnik, M. Self-supervised contrastive pre-training for time series via time-frequency consistency. *Adv. Neural Inf. Process. Syst.* **2022**, *35*, 3988–4003.
44. Ruff, L.; Vandermeulen, R.; Goernitz, N.; Deecke, L.; Siddiqui, S.A.; Binder, A.; Müller, E.; Kloft, M. Deep one-class classification. In *Proceedings of the International Conference on Machine Learning*; PMLR: Westminster, UK, 2018; pp. 4393–4402.
45. Shen, L.; Li, Z.; Kwok, J. Timeseries anomaly detection using temporal hierarchical one-class network. *Adv. Neural Inf. Process. Syst.* **2020**, *33*, 13016–13026.
46. Golestani, A.; Gras, R. Can we predict the unpredictable? *Sci. Rep.* **2014**, *4*, 6834. [CrossRef]
47. de Castro-Cros, M.; Velasco, M.; Angulo, C. Machine-learning-based condition assessment of gas turbines—A review. *Energies* **2021**, *14*, 8468. [CrossRef]

48. Tsallis, C.; Papageorgas, P.; Piromalis, D.; Munteanu, R.A. Application-wise review of Machine Learning-based predictive maintenance: Trends, challenges, and future directions. *Appl. Sci.* **2025**, *15*, 4898. [CrossRef]
49. Ji, A.; Woo, W.L.; Wong, E.W.L.; Quek, Y.T. Rail track condition monitoring: A review on deep learning approaches. *Intell. Robot.* **2021**, *1*, 151–175. [CrossRef]
50. Correia, L.; Goos, J.C.; Klein, P.; Bäck, T.; Kononova, A.V. Online model-based anomaly detection in multivariate time series: Taxonomy, survey, research challenges and future directions. *Eng. Appl. Artif. Intell.* **2024**, *138*, 109323. [CrossRef]
51. Kini, K.R.; Harrou, F.; Madakyaru, M.; Sun, Y. Enhancing wind turbine performance: Statistical detection of sensor faults based on improved dynamic independent component analysis. *Energies* **2023**, *16*, 5793. [CrossRef]
52. Zou, Z.; Chen, G.; Xie, L.; Wang, J.; Yang, Z. Thermo-Mechanical Fault Diagnosis for Marine Steam Turbines: A Hybrid DLinear–Transformer Anomaly Detection Framework. *J. Mar. Sci. Eng.* **2025**, *13*, 2050. [CrossRef]
53. Cheng, X.; Shi, F.; Liu, X.; Zhao, M.; Chen, S. A novel deep class-imbalanced semisupervised model for wind turbine blade icing detection. *IEEE Trans. Neural Netw. Learn. Syst.* **2021**, *33*, 2558–2570. [CrossRef] [PubMed]
54. Aravanis, T.; Papadopoulos, P.; Georgikos, D. Fourier Feature-Enhanced Neural Networks for Wind Turbine Power Modeling. *Electricity* **2025**, *6*, 70. [CrossRef]
55. Yao, K.; Wang, Y.; Fan, S.; Fu, J.; Wan, J.; Cao, Y. Improved and accurate fault diagnostic model for gas turbine based on 2D-wavelet transform and generative adversarial network. *Meas. Sci. Technol.* **2023**, *34*, 075104. [CrossRef]
56. Gou, P.; Zheng, D.; Tang, L.; Liu, M. A Transformer-Based Time Series Prediction Model with Contrastive Learning for Nuclear Power System State Monitoring. In *Proceedings of the 2024 Global Reliability and Prognostics and Health Management Conference (PHM-Beijing)*; IEEE: New York, NY, USA, 2024; pp. 1–6.
57. Qiao, W.; Liu, X.; Huang, J.; Wu, G. Fault diagnosis method based on the timesnet model under variable speed and strong noise. In *Proceedings of the 10th International Symposium on Test Automation & Instrumentation (ISTAI 2024)*; IET: Stevenage, UK, 2024; Volume 2024, pp. 110–117.
58. Kishore, B.; Kabilan, K.; Rahul, L.; Priyadarshan, G.P.; Vishnuthierth, E.; Satheesh, R.; Kolhe, M.L. Advancing short-term wind power forecasting by AI-driven models for improved accuracy. *Electr. Eng.* **2025**, *107*, 14795–14810. [CrossRef]
59. Schmidl, S.; Wenig, P.; Papenbrock, T. Anomaly Detection in Time Series: A Comprehensive Evaluation. *Proc. VLDB Endow.* **2022**, *15*, 1779–1797. [CrossRef]
60. Goyal, V. Anomaly Detection and Failure Prediction in Gas Turbines. Ph.D. Thesis, University of Central Florida, Orlando, FL, USA, December 2021.
61. Gori, V.; Veneri, G.; Ballarini, V. Continual Learning for anomaly detection on turbomachinery prototypes—A real application. In *Proceedings of the 2022 IEEE Congress on Evolutionary Computation (CEC)*; IEEE: New York, NY, USA, 2022; pp. 1–7.
62. Shetty, S.; Gori, V.; Bagni, G.; Veneri, G. Sensor Virtualization for Anomaly Detection of Turbo-Machinery Sensors—An Industrial Application. *Eng. Proc.* **2023**, *39*, 96.
63. Farahani, M. Anomaly Detection on Gas Turbine Time-Series’ Data Using Deep LSTM-Autoencoder. Master’s Thesis, Umeå University, Umeå, Sweden, 2021.
64. Wang, B.; Zhou, B.; Zhu, D.; Zou, M.; Rao, Z.; Luo, H. Feature Alignment LSTM for Detecting Anomalies in Wind Turbine Main Bearing Temperature Based on SCADA Data. *J. Eng. Res.* **2025**, *14*, 303–314. [CrossRef]
65. Chen, N.; Shao, C.; Wang, G.; Wang, Q.; Zhao, Z.; Liu, X. Anomaly detection of wind turbine based on norm-linear-ConvNeXt-TCN. *Meas. Sci. Technol.* **2024**, *35*, 076107. [CrossRef]
66. Liu, B.; Wang, X.; Zhang, Z.; Zhao, Z.; Wang, X.; Liu, T. Fault Prediction of Hydropower Station Based on CNN-LSTM-GAN with Biased Data. *Energies* **2025**, *18*, 3772. [CrossRef]
67. Pinciroli, L.; Baraldi, P.; Zio, E. Graph Neural Networks for Anomaly Detection in Wind Turbines. In *Proceedings of the 35th European Safety and Reliability & the 33rd Society for Risk Analysis Europe Conference (ESREL-SRA-E2025)*, Stavanger, Norway, 15–19 June 2025.
68. Jiang, G.; Li, W.; Fan, W.; He, Q.; Xie, P. TempGNN: A temperature-based graph neural network model for system-level monitoring of wind turbines with SCADA data. *IEEE Sens. J.* **2022**, *22*, 22894–22907. [CrossRef]
69. Zhan, Y.; Yang, B.; Ma, Y. Anomaly Detection for Small Hydropower Based on Deep Spatio-Temporal Modeling. In *Proceedings of the 2025 6th International Conference on Mechatronics Technology and Intelligent Manufacturing (ICMTIM)*; IEEE: New York, NY, USA, 2025; pp. 624–630.
70. Daenens, S.; Verstraeten, T.; Daems, P.J.; Nowé, A.; Helsen, J. Spatio-temporal graph neural networks for power prediction in offshore wind farms using SCADA data. *Wind Energy Sci.* **2025**, *10*, 1137–1152. [CrossRef]
71. Zheng, Y.; Wang, C.; Huang, C.; Li, K.; Yang, J.; Xie, N.; Liu, B.; Zhang, Y. Hierarchical spatial–temporal autocorrelation graph neural network for online wind turbine pitch system fault detection. *Neurocomputing* **2024**, *586*, 127574. [CrossRef]
72. Jin, X.; Lv, S.; Kong, Z.; Yang, H.; Zhang, Y.; Guo, Y.; Xu, Z. Graph spatio-temporal networks for condition monitoring of wind turbine. *IEEE Trans. Sustain. Energy* **2024**, *15*, 2276–2286. [CrossRef]

73. Wang, J.; Kou, M.; Li, R.; Qian, Y.; Li, Z. Ultra-short-term wind power forecasting jointly driven by anomaly detection, clustering and graph convolutional recurrent neural networks. *Adv. Eng. Inform.* **2025**, *65*, 103137. [CrossRef]
74. Blázquez-García, A.; Conde, A.; Mori, U.; Lozano, J.A. A review on outlier/anomaly detection in time series data. *ACM Comput. Surv.* **2021**, *54*, 1–33. [CrossRef]
75. Xie, J.; Schmidt, K.; Budeanu, N.; Letendre, V.; Zhao, Y.F. Combining feature learning and transfer learning in balancing anomaly detection for gas turbine engine vibration analysis. In *Proceedings of the International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*; American Society of Mechanical Engineers: New York, NY, USA, 2022; Volume 86212, p. V002T02A082.
76. Zhang, C.; Hu, D.; Yang, T. Anomaly detection and diagnosis for wind turbines using long short-term memory-based stacked denoising autoencoders and XGBoost. *Reliab. Eng. Syst. Saf.* **2022**, *222*, 108445. [CrossRef]
77. Chen, J.; Li, J.; Chen, W.; Wang, Y.; Jiang, T. Anomaly detection for wind turbines based on the reconstruction of condition parameters using stacked denoising autoencoders. *Renew. Energy* **2020**, *147*, 1469–1480. [CrossRef]
78. Yan, W.; Yu, L. On accurate and reliable anomaly detection for gas turbine combustors: A deep learning approach. *arXiv* **2019**, arXiv:1908.09238. [CrossRef]
79. Fu, S.; Zhong, S.; Lin, L.; Zhao, M. A re-optimized deep auto-encoder for gas turbine unsupervised anomaly detection. *Eng. Appl. Artif. Intell.* **2021**, *101*, 104199. [CrossRef]
80. Liu, D.; Zhong, S.; Lin, L.; Zhao, M.; Fu, X.; Liu, X. CSiamese: A novel semi-supervised anomaly detection framework for gas turbines via reconstruction similarity. *Neural Comput. Appl.* **2023**, *35*, 16403–16427. [CrossRef]
81. Jalil Pour, Z. Sensor Modelling for Anomaly Detection in Time Series Data. Master's Thesis, Linköping University, Linköping, Sweden, 1 June 2022.
82. Fahmi, A.T.W.K.; Reza Kashyzadeh, K.; Ghorbani, S. Advancements in Gas Turbine Fault Detection: A Machine Learning Approach Based on the Temporal Convolutional Network–Autoencoder Model. *Appl. Sci.* **2024**, *14*, 4551. [CrossRef]
83. Ma, Y.; Zhu, X.; Lu, J.; Yang, P.; Sun, J. Construction of data-driven performance digital twin for a real-world gas turbine anomaly detection considering uncertainty. *Sensors* **2023**, *23*, 6660. [CrossRef]
84. Roelofs, C.M.; Gück, C.; Faulstich, S. Transfer learning applications for autoencoder-based anomaly detection in wind turbines. *Energy AI* **2024**, *17*, 100373. [CrossRef]
85. Wang, L.; He, Y.; Shao, K.; Xing, Z.; Zhou, Y. An unsupervised approach to wind turbine blade icing detection based on beta variational graph attention autoencoder. *IEEE Trans. Instrum. Meas.* **2023**, *73*, 2500912. [CrossRef]
86. Duan, Y.; Xiang, M.; Zhou, B.; Fu, D.; Liu, H. TCAD: Unsupervised anomaly detection based on global local representation differences. *IEEE Access* **2022**, *10*, 114683–114693. [CrossRef]
87. Luo, Q.; Dong, J. Time Series Anomaly Detection Model Based on Memory-enhanced Transformer and Graph Network Joint Training. In *Proceedings of the 2024 4th International Conference on Computational Modeling, Simulation and Data Analysis*; ACM: New York, NY, USA, 2024; pp. 357–363.
88. Fan, Z.; Wang, Y.; Meng, L.; Zhang, G.; Qin, Y.; Tang, B. Unsupervised anomaly detection method for bearing based on VAE-GAN and time-series data correlation enhancement (June 2023). *IEEE Sens. J.* **2023**, *23*, 29345–29356. [CrossRef]
89. Xu, W.; Zhang, P. Steam turbine anomaly detection: An unsupervised learning approach using enhanced long short-term memory variational autoencoder. *Appl. Therm. Eng.* **2025**, *278*, 127138. [CrossRef]
90. Alper, O.C.; Doğan, H.; Öztürk, H. Gear pitting fault detection: Leveraging anomaly detection methods. In *Proceedings of the 2023 14th International Conference on Electrical and Electronics Engineering (ELECO)*; IEEE: New York, NY, USA, 2023; pp. 1–5.
91. Yang, Z.; Xu, M.; Wang, S.; Li, J.; Peng, Z.; Jin, F.; Yang, Y. Detection of wind turbine blade abnormalities through a deep learning model integrating VAE and neural ODE. *Ocean Eng.* **2024**, *302*, 117689. [CrossRef]
92. Dabaja, H.; Youssef, A.; Noura, H.; Ouladsine, M. Fault Detection in Scavenge Air System of a Marine Dual-Fuel Engine using GAN-based Unsupervised Anomaly Detection. In *Proceedings of the 6th International Conference on Control and Fault-Tolerant Systems (SysTol'25)*, Ayia Napa, Cyprus, 6–8 October 2025.
93. Peng, D.; Desmet, W.; Gryllias, K. Reconstruction-based deep unsupervised adaptive threshold support vector data description for wind turbine anomaly detection. *Reliab. Eng. Syst. Saf.* **2025**, *260*, 110995. [CrossRef]
94. Liang, Y.; Zheng, X.; Pang, D.; Zhu, W.; Liu, M.; Xue, S.; Cao, H. ExpertAP: Leveraging Multi-unit Operational Patterns for Advanced Turbine Anomaly Prediction. In *Proceedings of the International Conference on Intelligent Robotics and Applications*; Springer: Berlin/Heidelberg, Germany, 2024; pp. 344–359.
95. Chen, J. Deep Learning Frameworks for Anomaly Detection in Time Series and Graphs with Limited Labels. Ph.D. Thesis, University of Waterloo, Waterloo, ON, Canada, 29 April 2025.
96. Sun, S.; Hu, W.; Liu, Y.; Wang, T.; Chu, F. Matching contrastive learning: An effective and intelligent method for wind turbine fault diagnosis with imbalanced SCADA data. *Expert Syst. Appl.* **2023**, *223*, 119891. [CrossRef]
97. Sun, S.; Li, Q.; Hu, W.; Liang, Z.; Wang, T.; Chu, F. Wind turbine blade breakage detection based on environment-adapted contrastive learning. *Renew. Energy* **2023**, *219*, 119487. [CrossRef]

98. Liu, X.; Guo, H.; Liu, Y. One-shot fault diagnosis of wind turbines based on meta-analogical momentum contrast learning. *Energies* **2022**, *15*, 3133. [CrossRef]
99. Wang, Z.; Qin, B.; Sun, H.; Zhang, J.; Butala, M.D.; Demartino, C.; Peng, P.; Wang, H. An imbalanced semi-supervised wind turbine blade icing detection method based on contrastive learning. *Renew. Energy* **2023**, *212*, 251–262. [CrossRef]
100. Guo, J.; Liu, C.; Liu, S.; Liu, W. Graph contrastive learning for semi-supervised wind turbine fault diagnosis with few labeled SCADA data. *Measurement* **2025**, *245*, 116531. [CrossRef]
101. Qiao, W.; Liu, X.; Huang, J.; Wu, G. A Prior Knowledge Embedding Contrastive Attention Learning Network for Variable Working Conditions Bearing Fault Diagnosis With Small Samples. *IEEE Sens. J.* **2024**, *24*, 39967–39980. [CrossRef]
102. Miyamoto, K.; Lee, T.Y.; Minezawa, A. AAAD: Adaptive Activated Anomaly Detection on Varied Backgrounds. In *Proceedings of the 2025 IEEE International Conference on Advanced Visual and Signal-Based Systems (AVSS)*; IEEE: New York, NY, USA, 2025; pp. 1–6.
103. Wang, C.; Chen, H.; Wang, Y.; Zhao, S.; Liu, K. Spatial contrast and semantic difference perception network for aeroengine blade damage segmentation. *IEEE Trans. Aerosp. Electron. Syst.* **2024**, *61*, 2040–2056. [CrossRef]
104. Qin, Z.; Wang, X.; Han, X.; Lin, Y.; Zhou, Y. Pre-trained combustion model and transfer learning in thermoacoustic instability. *Phys. Fluids* **2023**, *35*, 037117. [CrossRef]
105. Li, Y.; Wang, Y.; Li, M.; Li, X.; Feng, J. LLM-YOLOMS: Large Language Model-based Semantic Interpretation and Fault Diagnosis for Wind Turbine Components. *arXiv* **2025**, arXiv:2511.10394.
106. Redmon, J.; Divvala, S.; Girshick, R.; Farhadi, A. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*; IEEE: New York, NY, USA, 2016; pp. 779–788.
107. Tang, Y.; Fu, J.; Cui, Q.; Dai, B.; Xiong, J. Wind Turbine SCADA Data Interpreter Based on Large Language Models. In *Proceedings of the Electrical Artificial Intelligence Conference*; Springer: Berlin/Heidelberg, Germany, 2024; pp. 530–541.
108. Zhang, C.; Hu, D.; Yang, T. Research of artificial intelligence operations for wind turbines considering anomaly detection, root cause analysis, and incremental training. *Reliab. Eng. Syst. Saf.* **2024**, *241*, 109634. [CrossRef]
109. Alblawi, A. Fault diagnosis of an industrial gas turbine based on the thermodynamic model coupled with a multi feedforward artificial neural networks. *Energy Rep.* **2020**, *6*, 1083–1096. [CrossRef]
110. Yucesan, Y.A.; Viana, F.A. A physics-informed neural network for wind turbine main bearing fatigue. *Int. J. Progn. Health Manag.* **2020**, *11*, 003. [CrossRef]
111. Gao, H.; Zahr, M.J.; Wang, J.X. Physics-informed graph neural Galerkin networks: A unified framework for solving PDE-governed forward and inverse problems. *Comput. Methods Appl. Mech. Eng.* **2022**, *390*, 114502. [CrossRef]
112. Zerrougui, I.; Li, Z.; Hissel, D. Physics-Informed Neural Network for modeling and predicting temperature fluctuations in proton exchange membrane electrolysis. *Energy AI* **2025**, *20*, 100474. [CrossRef]
113. Sahibzada, S.; Malik, F.S.; Nasir, S.; Lodhi, S.K. AI-Augmented Turbulence and Aerodynamic Modelling: Accelerating High-Fidelity CFD Simulations with Physics-informed Neural Networks. *Int. J. Innov. Res. Comput. Sci. Technol.* **2025**, *13*, 91–97. [CrossRef]
114. Lai, C.; Ahmed, I.; Zio, E.; Li, W.; Zhang, Y.; Yao, W.; Chen, J. A multistage physics-informed neural network for fault detection in regulating valves of nuclear power plants. *Energies* **2024**, *17*, 2647. [CrossRef]
115. Pei, F.; Zhang, B.; Long, S.; Yuan, M. BiLSTM-PINN-based Wind Turbine Power Prediction Architecture and Anomaly Data Identification. *Renew. Energy* **2026**, *261*, 125182. [CrossRef]
116. Mittal, S.; Narasinh, V.; Kulkarni, N.; Minz, R.L.; Chakravorty, N.; Mital, P. *Robust Power Prediction of Wind Turbine Using Error Detection, Clustering-Based Imputation and Physics-Informed Learning*; IARIA: Wilmington, DE, USA, 2024.
117. Gijón, A.; Pujana-Goitia, A.; Perea, E.; Molina-Solana, M.; Gómez-Romero, J. Prediction of wind turbines power with physics-informed neural networks and evidential uncertainty quantification. *Eng. Appl. Artif. Intell.* **2026**, *164*, 113331. [CrossRef]
118. Zhang, Z.; Wan, C.; Li, L.; Pang, Y.; Xiong, G.; Zhang, C.; Fan, X.; Wang, Y. CNN-IPINN: A method study on applying physics-informed neural networks to gear fault diagnosis. *Meas. Sci. Technol.* **2025**, *36*, 106101. [CrossRef]
119. Zideh, M.J.; Solanki, S.K. Multivariate physics-informed convolutional autoencoder for anomaly detection in power distribution systems with widespread deployment of distributed energy resources. *Sustain. Energy Grids Netw.* **2025**, *44*, 102022. [CrossRef]
120. Wu, R.; Luo, D.; Han, T. A Physics-Informed Patch Transformer with Interpretable Reasoning for Early Anomaly Warning in Wind Turbines. *Meas. Sci. Technol.* **2025**, *37*, 036205. [CrossRef]
121. Yue, K.; Li, Z.; Yao, X.; Du, W.; Yu, W.; Huang, W.; Wang, L. Physics-informed attention-aided multiple autoencoder for gear anomaly detection under zero-shot condition. *Measurement* **2025**, *257*, 118725. [CrossRef]
122. Duan, H.; Zhao, K.; Li, C.; Zhang, W.; Xu, Z. Reliable anomaly intelligent detection for electromechanical system of wind turbines based on probabilistic physics-informed autoencoders. In *Proceedings of the IET Conference Proceedings CP949*; IET: Stevenage, UK, 2025; Volume 35, pp. 1841–1847.
123. Schröder, L.; Dimitrov, N.K.; Verelst, D.R.; Sørensen, J.A. Using transfer learning to build physics-informed machine learning models for improved wind farm monitoring. *Energies* **2022**, *15*, 558. [CrossRef]

124. Huber, L.G.; Palmé, T.; Chao, M.A. Physics-informed machine learning for predictive maintenance: Applied use-cases. In *Proceedings of the 2023 10th IEEE Swiss Conference on Data Science (SDS)*; IEEE: New York, NY, USA, 2023; pp. 66–72.
125. Li, G.; Liu, Q.; Chen, Z.; Cheng, Y.; Wei, M.; Wu, D. Intelligent fault diagnosis of rotating machinery driven by physical information and contrastive learning under extreme sample imbalance conditions. *Meas. Sci. Technol.* **2025**, *36*, 106132. [CrossRef]
126. Perez-Sanjines, F.; Peeters, C.; Verstraeten, T.; Antoni, J.; Nowé, A.; Helsen, J. Fleet-based early fault detection of wind turbine gearboxes using physics-informed deep learning based on cyclic spectral coherence. *Mech. Syst. Signal Process.* **2023**, *185*, 109760. [CrossRef]
127. Jamil, F.; Peeters, C.; Verstraeten, T.; Helsen, J. Wind turbine drivetrain fault detection using physics-informed multivariate deep learning. In *Proceedings of the Surveillance, Vibrations, Shock and Noise*, Toulouse, France, 10–13 July 2023.
128. Freeman, B.; Tang, Y.; Huang, Y.; VanZwieten, J. Physics-informed turbulence intensity infusion: A new hybrid approach for marine current turbine rotor blade fault detection. *Ocean Eng.* **2022**, *254*, 111299. [CrossRef]
129. Lee, C.Y.; Maceren, E.D.C.; Huang, C.H. An Integrated Physics-Informed Deep CNN and Adaptive Elite-Based PSO-Catboost for Wind Energy Systems Fault Classification. *IET Renew. Power Gener.* **2026**, *20*, e70175. [CrossRef]
130. Lee, C.Y.; Maceren, E.D.C. Physics-informed anomaly and fault detection for wind energy systems using deep CNN and adaptive elite PSO-XGBoost. *IET Gener. Transm. Distrib.* **2025**, *19*, e13289. [CrossRef]
131. Chen, Y.; Rao, M.; Feng, K.; Zuo, M.J. Physics-Informed LSTM hyperparameters selection for gearbox fault detection. *Mech. Syst. Signal Process.* **2022**, *171*, 108907. [CrossRef]
132. Khan, M.A.; Rahman, A.; Mahmud, F.U.; Bishnu, K.K.; Nabil, H.R.; Mridha, M.; Hossen, M.J. A Physics-Guided Bayesian Neural Network for Sensor Fault Detection in Wind Turbines. *IEEE Open J. Comput. Soc.* **2025**, *6*, 931–942. [CrossRef]
133. Badihi, H.; Zhang, Y.; Jiang, B.; Pillay, P.; Rakheja, S. A comprehensive review on signal-based and model-based condition monitoring of wind turbines: Fault diagnosis and lifetime prognosis. *Proc. IEEE* **2022**, *110*, 754–806. [CrossRef]
134. Sun, T.; Yu, G.; Gao, M.; Zhao, L.; Bai, C.; Yang, W. Fault diagnosis methods based on machine learning and its applications for wind turbines: A review. *IEEE Access* **2021**, *9*, 147481–147511. [CrossRef]
135. Shah, S.S.; Daoliang, T.; Kumar, S.C. RUL forecasting for wind turbine predictive maintenance based on deep learning. *Heliyon* **2024**, *10*, e39268. [CrossRef]
136. Pandit, R.K.; Astolfi, D.; Durazo Cardenas, I. A review of predictive techniques used to support decision making for maintenance operations of wind turbines. *Energies* **2023**, *16*, 1654. [CrossRef]
137. Lu, Q.; Huang, X.; Wu, G.; Shen, X.; Zhu, D. Digital twin-driven water-wave information transmission and recurrent acceleration network for remaining useful life prediction of gear box. *Eng. Res. Express* **2025**, *7*, 025202. [CrossRef]
138. Sun, J.; Yan, Z.; Han, Y.; Zhu, X.; Yang, C. Deep learning framework for gas turbine performance digital twin and degradation prognostics from airline operator perspective. *Reliab. Eng. Syst. Saf.* **2023**, *238*, 109404. [CrossRef]
139. Wang, Y.; Wang, Y. A denoising semi-supervised deep learning model for remaining useful life prediction of turbofan engine degradation. *Appl. Intell.* **2023**, *53*, 22682–22699. [CrossRef]
140. Elsherif, S.M.; Hafiz, B.; Makhoul, M.; Farouk, O. A deep learning-based prognostic approach for predicting turbofan engine degradation and remaining useful life. *Sci. Rep.* **2025**, *15*, 26251. [CrossRef]
141. He, K.; Su, Z.; Tian, X.; Yu, H.; Luo, M. RUL prediction of wind turbine gearbox bearings based on self-calibration temporal convolutional network. *IEEE Trans. Instrum. Meas.* **2022**, *71*, 3501912. [CrossRef]
142. Li, D.; Zhou, F.; Gao, Y.; Yang, K.; Gao, H. Power plant turbine power trend prediction based on continuous prediction and online oil monitoring data of deep learning. *Tribol. Int.* **2024**, *191*, 109083. [CrossRef]
143. Atsafack, B.M.; Kabiri, C.; Rushingabigwi, G. Predictive maintenance for hydraulic turbine unit: A comparative deep learning approach using internet of things data in real-time. *IEEE Access* **2025**, *13*, 158340–158352. [CrossRef]
144. Sun, L.; Liu, T.; Xie, Y.; Zhang, D.; Xia, X. Real-time power prediction approach for turbine using deep learning techniques. *Energy* **2021**, *233*, 121130. [CrossRef]
145. Rengasamy, D.; Rothwell, B.; Figueredo, G.P. Asymmetric loss functions for deep learning early predictions of remaining useful life in aerospace gas turbine engines. In *Proceedings of the 2020 International Joint Conference on Neural Networks (IJCNN)*; IEEE: New York, NY, USA, 2020; pp. 1–7.
146. Li, A.; Guan, F.; Hu, H.; Liu, Z.; Xiao, B.; Huang, W.; Yang, Y. Physics-informed DEEPLSTM for digital twin modeling of steam turbine full-operating condition. *Appl. Therm. Eng.* **2025**, *278*, 127439. [CrossRef]
147. Zhang, X.C.; Gong, J.G.; Xuan, F.Z. A physics-informed neural network for creep-fatigue life prediction of components at elevated temperatures. *Eng. Fract. Mech.* **2021**, *258*, 108130. [CrossRef]
148. Wang, Z.; Gao, P.; Chu, X. Remaining useful life prediction of wind turbine gearbox bearings with limited samples based on prior knowledge and PI-LSTM. *Sustainability* **2022**, *14*, 12094. [CrossRef]
149. Li, J.; Yang, Z. Remaining Useful Life Prediction of Turbine Engines Using Multimodal Transfer Learning. *Machines* **2025**, *13*, 789. [CrossRef]

150. Wolniak, R.; Gajdzik, B.; Grebski, W. The usage of Root Cause Analysis (RCA) in Industry 4.0 conditions. In *Scientific Papers of Silesian University of Technology; Organization and Management Series*; Silesian University of Technology: Gliwice, Poland, 2023; pp. 223–235.
151. Papageorgiou, K.; Theodosiou, T.; Rapti, A.; Papageorgiou, E.I.; Dimitriou, N.; Tzovaras, D.; Margetis, G. A systematic review on machine learning methods for root cause analysis towards zero-defect manufacturing. *Front. Manuf. Technol.* **2022**, *2*, 972712. [CrossRef]
152. Gugaliya, J.K.; Vij, R.K.; Ramaswamy, S.; K, L.S.M. Multi-PCA Driven Approach for Fault Detection and Root Cause Analysis of Process Equipment. In *Proceedings of the AAAI Spring Symposium: Combining Machine Learning with Knowledge Engineering*, Palo Alto, CA, USA, 23–25 March 2020; Volume 1.
153. Yuan, L.; Han, G.; Dong, P. Improved bayesian network with graph attention and prior algorithm for aircraft engine fault root cause analysis. *Sci. Rep.* **2026**, *16*, 5924. [CrossRef]
154. Mourya, M.; Tyagi, T. Predict, Prevent, Decide: AI-Enabled Root Cause Analysis and Operational Intelligence for Offshore Performance. In *Proceedings of the Abu Dhabi International Petroleum Exhibition and Conference*; SPE: Richardson, TX, USA, 2025; p. 229267.
155. Ofoedu, A.T.; Ozor, J.E.; Sofoluwe, O.; Jambol, D.D. A Root Cause Analytics Model for Diagnosing Offshore Process Failures Using Live Operational Data. *Int. Sci. Ref. Res. J.* **2022**, *5*, 226–244.
156. Vitale, M.; Youssef, A.; Mishra, P.; Shetty, S.; Sharma, M.; Vanzo, G.; Veneri, G.; Vieri, L.; Bettini, A. Harnessing Generative AI for Interactive System Failure Diagnostics: A User-Centric Approach to Streamlined Problem Solving and Maintenance. In *Proceedings of the Abu Dhabi International Petroleum Exhibition and Conference*; SPE: Richardson, TX, USA, 2024; p. 222033.
157. Shi, R.; Liu, W.; Guo, C.; He, Y. Intelligent Root Cause Analysis Approach for Quality Accidents Based on Knowledge Graph and Large Language Model. In *Proceedings of the 2025 16th International Conference on Reliability, Maintainability and Safety (ICRMS)*; IEEE: New York, NY, USA, 2025; pp. 1–6.
158. Ferrari, P.; Rinaldi, S.; Sisinni, E.; Colombo, F.; Ghelfi, F.; Maffei, D.; Malara, M. Performance evaluation of full-cloud and edge-cloud architectures for Industrial IoT anomaly detection based on deep learning. In *Proceedings of the 2019 II Workshop on Metrology for Industry 4.0 and IoT (MetroInd4.0 & IoT)*; IEEE: New York, NY, USA, 2019; pp. 420–425.
159. Yang, Q.; Duan, L.; Song, W.; Zhang, S. A service-based cloud-edge fusion approach for abnormality detection of power generation equipment. *IEEE Access* **2024**, *12*, 51556–51569. [CrossRef]
160. Oliver, S.; Engber, M.; Mich, R.; Bohra, D. Cloud Based Implementation of a Gas Turbine Remote Monitoring System. In *Proceedings of the GPPS Forum 18*, Montreal, QC, Canada, 7–9 May 2018.
161. Hurtado Carreon, A.; Veldhuis, S.C. Practical Application of Condition-Based Monitoring (CBM) Technologies in the Modern Manufacturing Industry: A Review. *Processes* **2025**, *13*, 4084. [CrossRef]
162. Kurkin, A.; Kabaldin, Y.; Zhelonkin, M.; Mancero, S.; Anosov, M.; Shatagin, D. Multiparametric Vibration Diagnostics of Machine Tools within a Digital Twin Framework Using Machine Learning. *Appl. Sci.* **2026**, *16*, 982. [CrossRef]
163. Tikka, P.; Karjalainen, J.; Alesani, A.; Goriachev, V. Large Language Model hallucination mitigation in three industrial use cases. *IEEE Access* **2026**, *14*, 25564–25576. [CrossRef]
164. Jia, M.; Cheng, Q.; Tao, C.; Hu, Y.; Hong, Q.; Cheng, W.; Liu, Z. A physics-informed train on synthetic and test on real method for evaluating large language model-generated safety-critical traffic scenarios. *Comput.-Aided Civ. Infrastruct. Eng.* **2025**, *40*, 5153–5169. [CrossRef]
165. Mirshekali, H.; Shadi, M.R.; Ladani, F.G.; Shaker, H.R. A Review of Large Language Models for Energy Systems: Applications, Challenges, and Future Prospects. *IEEE Access* **2025**, *13*, 163162–163188. [CrossRef]
166. ISO 13374-1; Condition Monitoring and Diagnostics of Machines—Data Processing, Communication and Presentation—Part 1: General Guidelines. Technical Report; International Organization for Standardization: Geneva, Switzerland, 2003.
167. ISO 13379-1; Condition Monitoring and Diagnostics of Machine Systems—Data Interpretation and Diagnostics Techniques—Part 1: General Guidelines. Technical Report; International Organization for Standardization: Geneva, Switzerland, 2025.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Review

New Trends and Challenges in Electric and Hybrid Electric Vehicles: Powertrain Configurations, Traction Motors and Drive Control Techniques

Syed Hassan Imam ^{1,2,*}, Saqib Jamshed Rind ³, Saba Javed ² and Mohsin Jamil ^{4,*}

¹ Department of Electrical Engineering Technology, Hamdard University, Karachi 74600, Pakistan

² Department of Electronic Engineering, NED University of Engineering & Technology, Karachi 75270, Pakistan

³ Department of Automotive and Marine Engineering, NED University of Engineering & Technology, Karachi 75270, Pakistan

⁴ Department of Engineering, Brock University, St. Catharines, ON L2S 3A1, Canada

* Correspondence: mjamil@brocku.ca; Tel.: +709-864-2751

Abstract

The requirement of sustainable mobility and a clean environment has accelerated the development and adoption of electric vehicles (EVs) and hybrid electric vehicles (HEVs) as an alternative, practical and promising solution against conventional vehicles globally. Such alternative energy vehicles not only provide a critical solution to mitigate fossil fuel dependency and reduce greenhouse gas emissions, but also contribute to producing an energy-efficient transportation system. However, the operational performance, efficiency, and cost-effectiveness of EVs and HEVs are hugely dependent on their powertrain architectures, selection of traction motors and associated control techniques. This paper systematically compares major hybrid architectures: series, parallel, and series–parallel, plug-in, as well as battery and fuel cell electric vehicle platforms, highlighting trade-offs in component sizing, cost, and system integration complexity. The paper critically analyses traction motor technologies with respect to torque–speed characteristics, efficiency behavior, material constraints, and power density. A detailed comparative assessment of traction motor technologies is presented. Furthermore, classical and advanced motor control strategies, including field-oriented control (FOC), direct torque control (DTC), model predictive control (MPC) and AI-enhanced control frameworks, are evaluated with respect to transient performance, robustness, computational requirements, and scalability. The review identifies key technological milestones, emerging next-generation drive technologies, existing limitations, and unresolved research challenges. Finally, critical research gaps and future development pathways are articulated to support the advancement of high-efficiency, reliable, and cost-effective EV/HEV powertrain systems.

Keywords: electric vehicles; hybrid electric vehicles; motor control techniques; traction motors; battery

1. Introduction

The burning of fossil fuels in order to meet the energy requirements is a major source of greenhouse gas (GHG) emissions. GHG caused significant changes in the climate globally, which could be seen clearly in the form of glaciers melting and global warming [1,2]. As it is already estimated that global temperatures will rise by 2 degrees Celsius by 2050, the International Energy Agency (IEA) has developed major goals to reduce the average

global temperature rise [3,4]. GHG emissions are expected to more than triple by 2050 and transportation accounts for about 25% of GHG emissions. Effective measures are essential to mitigate emissions from the transportation sector. Therefore, alternative fuels are being explored that can provide a clean energy source, minimize GHG emissions, and offer a cost-effective solution with high performance efficiency [5,6].

To address geopolitical concerns about fossil fuel availability, energy security, global climate change, and rising petroleum prices, transport electrification is regarded as the most appropriate solution. Furthermore, multiple sources of creating electrical energy, like renewable energy, provide for electrical energy diversification. The transportation industry is one of the largest sources of GHG and carbon dioxide emissions, as the energy conversion of internal combustion engines (ICEs) is between 12 and 30% [7]. EVs are the ideal solution as they have a high conversion rate and are more attractive for users over ICE [8].

Global EV demand is steadily expanding. Over 6 million battery electric vehicles (BEVs) and plug-in hybrid electric vehicles (PHEVs) were dispatched at a growth rate of 40% globally. In May 2023, China's new EV passenger car sales grew by 60%, while EV sales in the United States increased by 48% and Europe-10 EV sales increased by 33%. Global EV sales prediction remains unchanged in April 2023, at 13.9 million units, with BEVs accounting for 72% of the total. The overall car industry recovery appears to be solid, with +23% year on year growth in the top-15 EV markets [9]. BYD is the leading brand in the global EV market in 2025. As seen in the year-by-year comparison, it took the most commanding sales position in the first half of 2025 with an increase of 33.3% [10]. HEV sales will reach thirty million globally by 2030, and 100 million by 2050 [11]. In Europe, 34% improvement in the EV market was observed in Q3 of 2025 [8]. Plug-in showed better sales progress compared to BEV in Q3 by a rise of 10% compared to Q2 of 2025. Over 25% of the new vehicles sold globally are EVs and the EV race is expanding to 39 countries where new EV manufacturers are contributing to 10% [12]. After a record-breaking Q3, a sharp decline in Q4 was observed in the US [13]. The year-to-year sales graph comparison of the USA is demonstrated in Figure 1.




Figure 1. Sales graph of EV and HEV (2025) [14].

Figure 2 illustrates EV sales growth comparison between 2024 and 2025 in Europe, the USA and China. This growth is continuously increasing, with full and mild HEVs growing by 14%, PHEVs by 5%, BEVs by 28% and overall EVs by 18%. However, shares of EV registrations in Europe, the USA and China for 2025 are demonstrated in Figure 3 [11,15].




Figure 2. Electric vehicle sales growth (2024 and 2025) [13].




Figure 3. Shares of electric vehicle registrations in 2025 [13].

Electrified transportation is a multidisciplinary field that includes the contributions and knowledge of different areas, as demonstrated in Figure 4. Traction motors play an important role in electric vehicle technology by improving performance and fuel efficiency [16,17]. The performance of the traction drive motor can either improve or degrade the performance of EVs. The environmental benefits of electric powertrains are seen and marketed as the true motivator for transportation electrification [18]. Manufacturers are constantly striving to create more efficient EVs [19]. EVs employed a battery bank to hold the charge and motors to power the vehicle, whereas ICE vehicles employed fuel and engines [20,21]. HEVs have high performance and efficiency, making them the most popular. Their large operating range and mature refueling infrastructure make them an alternative for the next generation [22]. There are various architectures of HEV propulsion systems that are already in place and providing future direction [23,24]. Most vehicles are charged while parked; the transition to EVs will place additional strain on the grid, particularly during peak hours [25]. To address this issue, an energy management system must be designed and executed. Various techniques have been presented; however, they are not physically implemented [26,27].




Figure 4. Electric vehicles and associated fields.

EVs employ a variety of motors; however, controlling an AC machine is a classic challenge [28]. Induction motors (IM) remain a smart alternative for EVs, even though the permanent magnet synchronous motor (PMSM) has been regarded as the dominant choice in recent years. Switch reluctance motors (SRMs) are also considered and employed in order to obtain a more affordable solution. Furthermore, it needed to be further improved for efficient performance. Although brushless direct current (BLDC) motors are highly efficient, they are preferred for lightweight vehicles. DC motors are the most basic motors, yet they are not appealing to manufacturers for EV applications these days. Double-stator and double-rotor motors are also under consideration for EV applications.

Motor speed control has a considerable impact on EV applications. Scalar control is a very basic approach, although DTC and FOC are superior speed controlling strategies. Intelligent controllers and MPC offer advanced speed control options. As technology advances in all dimensions, EVs are making strides towards smart cities through artificial intelligence and intelligent driving applications [29,30].

This paper is organized into six sections. Section 2 presents the classification of electric and hybrid vehicles, focusing on various powertrain architectures, their merits, limitations, and commercial utilization. Section 3 discusses the classification of traction motors, highlighting their applications and current challenges. Sections 4 and 5 explore high-performance traction motor control techniques with an emphasis on artificial intelligence integration. Section 6 addresses the present challenges, proposes research recommendations for potential solutions, and provides the overall conclusion of the paper, respectively.

2. Classification of Vehicles

Vehicles are mainly categorized into internal combustion engine (ICE) based automobiles and EVs, either fully or partially electric. HEVs include full hybrids and plug-in hybrids, while all-electric vehicles (AEVs) consist of battery electric vehicles (BEVs) and fuel cell electric vehicles (FCEVs) as demonstrated in Figure 5 [31]. ICEVs suffer from inefficiency and harmful emissions, prompting a shift toward battery-powered EVs to enhance efficiency and reduce air pollution for a healthy and clean environment. EVs and HEVs offer several advantages such as higher efficiency, lower emissions, reduced noise, and conservation of petroleum resources for alternative applications [32,33].




Figure 5. Classification of vehicles.

2.1. Hybrid Electric Vehicle

In HEV design, the primary focus is to regulate power flow between the electric source and the load to reduce power losses. Both the traction motor and ICE are utilized for the vehicle propulsion, achieving performance comparable to ICE vehicles as shown in Figure 6 [34,35]. Enhance dynamic performance in HEVs, which relies on careful energy configuration and hybrid power in powertrain systems [36,37]. On the basis of electrification level, HEVs are classified in micro-hybrid EV (mHEV), mild-hybrid EV (MHEV), full hybrid EV (HEV), plug-in hybrid EV (PHEV), and extended range EV (EREV) [32,38]. Full hybrid configurations are further characterized as series, parallel, and series-parallel [39,40]. HEVs are preferred due to their fuel efficiency, ultra-low emissions and regenerative braking advantages [38]. However, a hybrid powertrain increases vehicle weight, manufacturing cost and maintenance complexity. To overcome these issues, hub motors are introduced, which are positioned directly at the wheel hub, eliminating conventional axles and gearsets. The electric axle (E-Axle) drive integrates the propulsion motor, inverter and transmission in a single unit, optimizing energy conversion, torque distribution and overall power efficiency [41]. The hybridization factor (HF) expresses the relationship between the maximum power of the electric motor (PEM) and the power of the ICE motor (PICE). HF for EV is 1, whereas HEV performance could be optimized at (HF = 0.3–0.5). After that, improving the propulsion system capability did not increase HEV performance [32].

$$HF = PEM / (PEM + PICE) = PEM / PHEV \quad (1)$$




Figure 6. HEV architecture components.

2.1.1. Series HEV Systems

Series HEVs (SHEVs) are the simplest HEVs, consisting only of a battery and engine-generator set and propulsion via a traction motor, as demonstrated in Figure 7 [42]. ICE charges the battery bank while consuming minimal fuel. The drive train is employed in hybrid, engine, electric, charging, or regenerative braking mode. EREV is a prime example of SHEVs. SHEVs are appropriate for short-distance driving and stop-and-go traffic [43]. The advanced architecture provides optimal power control between motor and generator with a simple drive train structure without gears and clutches [44]. However, dual power conversion, huge traction motor size and large battery banks are the concerns. SHEVs operational characteristics, present challenges and system applications are demonstrated in Table 1.




Figure 7. Series HEV configuration.

Table 1. Characteristics of series HEV configuration [9,14,27,32].

Advantages	Disadvantages	System Application
The traction drive line is optimized and efficient. Small and efficient engine. Zero-emission operation is possible. The installation flexibility of machines. No multi-step transmission.	Energy loss due to multiple energy conversions. Large traction motor drive. Increased cost because of the additional generator.	BMW i3 Fisker Karma Tesla Buses Mercedes Citaro bus Tesla ultra-light rail

2.1.2. Parallel HEV Systems

Parallel HEV configurations are propulsion systems that connect both the ICE and the battery in parallel with the traction motor, as demonstrated in Figure 8. However, a switch allows both systems to operate concurrently. The battery runs the vehicle at moderate speeds, while the ICE runs the vehicle at high speeds. Therefore, it provides a compact system with superior dynamic performance. The recent advancement in compact integration of motor transmission coupling structure provides torque blending, regenerative braking while reducing drivetrain complexity [44]. However, difficult control, complex space packaging, and high system cost are the concerns. Parallel HEVs' operational characteristics, present challenges and system applications are demonstrated in Table 2.




Figure 8. Parallel HEV configuration.

Table 2. Characteristics of parallel HEV configuration [9,13,27].

Advantages	Disadvantages	System Application
Increased efficiency as a result of less power conversion. No additional generator. High speed due to two power sources. Zero emission is possible. One machine is required for hybrid operation. Less space is required for additional installation.	Complex space packaging. Expensive system. Complex control.	Ford Escape Hybrid SUV Honda insight Mercedes-Benz 400 Honda Civic Hybrid. Honda Accord. Lexus Hybrid SUV.

2.1.3. Series–Parallel (SP-HEV)

Series–parallel HEVs provide the capabilities of both series and parallel HEVs with the help of a planetary gearbox, as shown in Figure 9 [45]. At low speeds, electric motors generate the most torque, whereas at high speeds, ICE takes the lead. This setup improves the vehicle's power capability, low emissions, and fuel efficiency. An advanced hybrid drivetrain provides fuel efficiency and flexible power flow [44]. However, control complexity, high system cost, and complex drive train layout are the concerns [46,47]. SP-HEVs' operational characteristics, present challenges and system applications are summarized in Table 3. In addition, several power-split modes are addressed in [48].




Figure 9. Series–parallel configuration.

Table 3. Characteristics of SP-HEV configuration [39,49,50].

Advantages	Disadvantages	System Application
Vehicle power capability is improved. Torque is transmitted directly to the drive wheel. Better fuel efficiency. Zero emission is possible. Multi-step conversion is not required. It can be switched b/w series and parallel.	Two motors are required Cost increases System complexity	Toyota Prius. Toyota Camry Hybrid Hyundai Ioniq

Comparison of different operational modes, which are summarized in Table 4.

Table 4. Comparison of operational modes in series–parallel HEV [40,42,47,49,50].

Modes	ICE Dominating	EM Dominating
Starting phase	Electric motor and battery, then ICE	Electric motor and battery, then ICE
Full throttle acceleration	Both active, but ICE provides dominant power	Both active, but EM provides dominant power
Normal drive	ICE	EM
Braking/Deceleration	EM acts as a generator to recharge batteries	EM acts as a generator to recharge batteries
Charging during driving	ICE + generator charge batteries	ICE + generator charge batteries
Vehicle at rest	ICE drives the generator to charge batteries	ICE drives the generator to charge batteries

2.1.4. Plug-in HEV

Plug-in HEVs employ an external source to charge the battery. When the battery is depleted, the engine will resume functioning [51]. If the battery is charged, it has good fuel efficiency and low emissions, but as the battery discharges, the fuel consumption and CO₂ emissions increase. PHEVs are gaining popularity among consumers as their sales growth is rising by 40% in 2023, by 16% in 2024 and 20% in 2025 [10,52]. PHEVs are the electric drive train as the primary source of propulsion and energy. It is critical to pick the appropriate power train to drive plug-in vehicles to improve efficiency and mileage while reducing emissions [53]. Plug-in drawbacks are relatively large in size, which means the fuel efficiency is not good for highway travel or when the battery is not fully charged. However, Lithium-ion batteries are now extensively employed in Plug-in EVs and other types of EVs to boost performance [4,54]. Plug-in-HEVs operational characteristics, present challenges and system applications are summarized in Table 5.

Table 5. Characteristics of the plug-in HEV system [9,47,55].

Advantages	Disadvantage	System Application
Low emission Low fuel consumption Extended electric drive range	Dependent on the grid High initial cost Complex structure	Hyundai Ioniq Plug-in Volvo XC90 T8 Mini Countryman PHEV Porsche Cayenne e-Hybrid Volkswagen Phideon

2.2. Pure Electric Vehicle

The vehicle propulsion mechanism in pure electric vehicles is powered by electrical energy. It may also use renewable charging sources such as solar cells, wind energy, or other alternatives. This helps reduce dependence on the grid power supply [56–58]. Manufacturers are employing regenerative braking technologies during braking and as well as when traveling freely under the force of gravity [59]. The benefits include zero emissions, high mileage, very low noise pollution, compact size, and excellent efficiency. Development of unmanned, self-driving, and intelligent electric vehicles with one or more power train driving motors contributes more in the future [60,61]. However, the concerns are low load-bearing capacity, slow speed when compared to HEV and ICE vehicles, complex structure, and high cost owing to a huge battery bank.

BYD's sales more than doubled to 641,000 units, propelling it to the top of the global sales rankings. Among BEVs, Tesla continues to lead by a significant margin in 2024. However, it is the only brand whose sales dropped to 13.1% in 2025. FCEV and ICEV sales are down by −9% and −16% in 2025, respectively, from 2024. The comparison of each arrangement is presented in Table 6. Plug-ins (BEV and PHEV) accounted for an increase of 34.5% in 2025 H1. BEVs made up 10.4% of the market, while PHEVs made up 4.4%. Norway had the greatest PEV share in new light vehicle sales year to date (80.6%), followed by Hong Kong (68.5%), Iceland (54.9%), Sweden (52.3%), and Finland (46.8%) [9,62].

Table 6. Comparison of the global passenger EV market shares [62,63].

Global EV Market Share Up to Q3 2025							
Brands	Q1 2024	Q2 2024	Q3 2024	Q4 2024	Q1 2025	Q2 2025	Q3 2025
BYD Auto	15%	17%	16%	16%	15%	18%	16%
Tesla	20%	17%	17%	14%	12%	12%	13%
Geely Holdings	8%	8%	9%	9%	11%	11%	10%
Others	57%	58%	59%	61%	61%	60%	61%

In 2024, BEV continues its selling trend by 12% YoY in Q4 and BYD's sales overtake Tesla by 13% YoY. In China, the growth rate is reaching 56% in Q2, while America and Germany are following the increasing trend. The best-selling car in the world ranking for the year 2024 is BYD Song, led by 467,750 units but BYD Seagull leads in 2025, Tesla Model Y securing +3.6% market leadership in 2024, despite a drop in 2025, while Geely Holdings is in third place [9,10]. BYD and Tesla were still dominant at the end of 2025 [12].

2.2.1. Battery Electric Vehicle

A battery electric vehicle has three major components: an electric motor, a controller, and a rechargeable battery bank or super capacitor (SC), as demonstrated in Figure 10 [64]. The propulsion system is powered entirely by electrical energy. The electric motor is the primary source of EV movement and torque [65]. The recent development provides a compact module by integrating motor, inverter and transmission in an electric drive unit to enhance the efficiency. The adoption of a multi-level inverter mitigates switching losses and improves further performance [66]. The transition of operating voltage from 400 to 800 reduces current and minimizes power losses while enabling faster response and enhanced power density [67]. Due to the enormous battery bank, size and cost increases. It requires external charging, whereas HEVs have internal charging capabilities [68]. Furthermore, the charging infrastructure is not yet mature enough all over the world, despite rapid advancement in developing countries.




Figure 10. Basic BEV configuration.

Level 1 charging provides slow charging; however level 2 and level 3 charging stations are also advancing with fast charging topologies [69,70] and a fully charged battery provides approximately 200 km [71]. However, BEVs operational characteristics, present challenges and system applications are summarized in Table 7.

Table 7. Characteristics of BEV [9,13,62,72].

Advantages	Disadvantages	System Application
Zero emission. Good fuel efficiency. Reasonably high-power density.	Bigger in size and high cost. Battery dependency. Immature charging infrastructure.	Audi e-tron BMW i3, i7 Cadillac Lyriq Chevrolet Bolt EV Ford Mustang Mach-e

2.2.2. Fuel Cell Electric Vehicle

Fuel cell electric vehicles run on hydrogen as a fuel. Compressed hydrogen is stored in a fuel stack, which employs oxygen from the air to generate electrical energy, water and heat as a byproduct. Figure 11 shows the basic block diagram of FCEV. Platinum is employed as a catalyst to shorten the reaction time [73]. FCEV has high power density and fast refueling, which makes it suitable for long-range applications [74,75]. However, the conversion rate of hydrogen is about 40–60% as well as reliance on hydrogen infrastructure limits practical application [76]. Similarly, the complexity of the system, greater starting costs, fuel costs, reliability, and cyclic life are the challenges. FCEV has substantially lower specific power and lower starting performance; therefore, it is not sufficient by itself [77]. It requires batteries or a supercapacitor in hybridization, which increases cost and size. Therefore, A hybrid optimization technique is proposed for dynamic load sharing to tackle this issue for FCEV [78,79]. However, emerging electrified platforms accommodate multiple energy sources within a dedicated propulsion architecture to improve system efficiency [80]. Another method to deal with non-linearity and power losses based on fractional sliding mode control (SMC) design is proposed for FCEV [77].




Figure 11. FCEV components configuration.

As pure hydrogen is not available for refueling, the following are the commonly used compounds in FCEV. Direct methanol fuel cell (DMFC), alkaline fuel cell (AFC), phosphoric acid fuel cell (PAFC), solid oxide fuel cell (SOFC), and other types of compounds are commonly utilized in FCEV. A comparison of FC compound parameters is demonstrated in Table 8.

Table 8. Comparison of FC compound parameters [81].

Compound	Fuel Type	Power Density (W/m ³)	Temperature (°C)	Efficiency (%)
DMFC	Methanol	1500–3500	<100	30–40
PAFC	Hydrogen	800–2500	150–220	40–55
AFC	Hydrogen	1000–3000	50–150	40–60
PEMFC	Hydrogen	3500–6500	50–100	45–60

PEMFC is widely utilized in FCEV. It has a small size, a good start at low temperatures, a simple structure and maintenance-free operation, a reasonably high-power density, and the capacity to run under harsh conditions [82]. However, FCEVs operational characteristics, present challenges and system applications are summarized in Table 9.

Table 9. Characteristics of FCEV [83].

Advantages	Disadvantages	System Application
Low emission. Good fuel efficiency. Hydrogen as fuel	Additional chemicals are required. Low starting performance. Lack of infrastructure to produce pure hydrogen.	Toyota Mirai BMW i8 Audi H-tron quattro Mercedes-Benz GLC F-cell

A detailed comparison of ICE, BEV, HEV, PHEV and FCEV is demonstrated in Table 10.

Table 10. Comparison of ICEV, BEV, HEV, PHEV and FCEV [14,15,83–85].

Characteristics	ICEV	BEV	HEV	PHEV	FCEV
Propulsion System	IC engine	Electric Motor	Electric Motor and IC engine	Electric Motor and IC engine	Electric Motor
Energy Storage	Fuel tank Gas cylinder	Battery Super Capacitor Flywheel	Fuel tank, Battery Super Capacitor Flywheel	Fuel tank, Battery Super Capacitor Flywheel	Fuel cell Super Capacitor Battery
Energy Source Infrastructure	Fuel Station	Electric grid charging station	Electrical charging station Fuel pump	Electrical charging station (optional) Fuel pump	Hydrogen PEMFC, AFC, DMFC, SOFC

Table 10. Cont.

Characteristics	ICEV	BEV	HEV	PHEV	FCEV
Advantages	Fully matured High performance Reliable Availability of refueling infrastructure	Zero emission. High efficiency Low noise Independent of gasoline Commercialized	High performance Reliable Low emission Two sources of energy Higher fuel economy Durability Commercialized	Low emission Low fuel consumption Extended Electric drive range Smooth operation Capability of V2G or G2V	Ultralow emission. High energy efficiency Independent of gasoline Reliable
Disadvantages	Harmful emission Poor fuel economy Less efficient Comparatively bulk	High cost Poor dynamic response Long charging time Immature charging infrastructure Limited driving range	More complex due to two energy sources Bulky High cost Size and weight increased	Dependent on the grid High initial cost Complex structure Charging station infrastructure Battery technology Size and weight increased due to batteries and ICE	Pure hydrogen availability Slow dynamic response High cost Immature charging infrastructure Sophisticated electronic controller Conditioning, refilling and storage of hydrogen
Comparison	Still most dominant	Attracting users	Trending	Most trending EV	Decreasing trend
2021 vs. 2022	Sales declined by 16%	BEVs grew by 75%	Full Hybrids grew by 14%	PHEVs grew by 37%.	FCEV declined by 9%
2023 vs. 2022	ICE also increased +15%	BEV increased by 49%	HEVs increased 33%	PHEVs increased by 55%	FCEV decreased by 30% in Q1 2023
2024 vs. 2023	Decreased by 9%	BEV increased by 9% globally	Increased by up to 30%	About 25% increases	FCEV degrowth by 36.4%
2025 vs. 2024	Decline by 18.7%	Increased by 17%	Increased by 9.8%	Increased by 20%	Declined by 27.2%

2.3. Batteries in Electric Vehicles

The battery system serves as the core energy source in pure EVs, fundamentally determining driving range, power capability, efficiency, and operational safety. Unlike ICE-based systems, EV propulsion relies entirely on electrochemical energy storage, making the battery pack a critical subsystem rather than a passive component. It typically consists of interconnected cells arranged in modules, integrated with a battery management system (BMS), thermal management unit, and protection circuitry. The coordinated operation of these elements ensures reliable energy delivery under diverse driving conditions.

2.3.1. Role of Batteries in EV Powertrain

Within the EV powertrain, the battery pack supplies DC power to the inverter, which drives the traction motor. It must support both continuous energy demand during cruising and transient high-power requirements during acceleration. Additionally, the battery enables regenerative braking by absorbing recovered energy, thereby improving overall system efficiency [86].

The BMS performs essential monitoring and control functions, including estimation of state-of-charge (SOC) and state-of-health (SOH), voltage balancing across cells, and temperature supervision. Accurate estimation of these parameters is necessary to prevent overcharging, deep discharging, and thermal instability, which can significantly degrade battery performance and lifespan [87].

From a structural perspective, battery placement usually occurs under the vehicle floor to enhance stability by lowering the center of gravity. However, this configuration imposes constraints on packaging, thermal dissipation, and crash safety design.

2.3.2. Battery Chemistries for EV Applications

Lithium-ion (Li-ion) batteries currently dominate EV applications due to their high energy density, long cycle life, and relatively low self-discharge rate. Among various Li-ion chemistries, the most widely adopted include:

- **Lithium Nickel Manganese Cobalt Oxide (NMC):** Offers a balanced combination of energy density, power capability, and durability, making it suitable for a wide range of EVs.
- **Lithium Iron Phosphate (LFP):** Characterized by superior thermal stability, enhanced safety, and longer cycle life, though with comparatively lower energy density.
- **Lithium Nickel Cobalt Aluminum Oxide (NCA):** Provides high specific energy and is commonly used in long-range EVs, albeit with stricter thermal management requirements.

Recent advancements in Li-ion technology have focused on improving electrode materials, electrolyte formulations, and cell design to enhance performance metrics and safety. Beyond conventional Li-ion systems, next-generation battery technologies such as solid-state batteries, lithium-sulfur, and sodium-ion batteries are being actively investigated. Solid-state batteries, in particular, offer the potential for higher energy density and improved safety by eliminating flammable liquid electrolytes [87].

2.3.3. Key Performance Metrics for EV Batteries

The suitability of a battery for EV applications is determined by several critical performance parameters:

- **Energy Density (Wh/kg):** Directly influences driving range and vehicle weight.
- **Power Density (W/kg):** Determines acceleration capability and regenerative braking performance.
- **Life Cycle:** Reflects durability and long-term cost-effectiveness.
- **Charging Rate Capability:** Defines the feasibility of fast charging without excessive degradation.
- **Thermal Stability:** Ensures safe operation across varying environmental conditions.
- **Cost (\$/kWh):** A key factor affecting commercial viability and large-scale adoption.

These parameters are inherently interdependent, requiring careful trade-off analysis during battery selection and system design as shown in Table 11. For example, high-energy-density chemistries often exhibit reduced thermal stability, necessitating advanced cooling strategies and control mechanisms.

Table 11. Comparison of batteries [88–90].

Parameters	Li-Ion	LFP	NMC	Solid State Battery
Energy Density (Wh/kg)	150–250	120–200	200–300	400–500 (target)
Power Density (W/kg)	1800	1000–1500	1500–4000	2000–5000
Charging Speed	Moderate–fast	Moderate	Fast	Very fast (future potential)
Life Cycle	500–1000	1500–3000	500–1250	1000–3000 (expected)
Thermal Stability	Moderate	Very high	Moderate–low	Very high
Key Materials	Graphite + Li-metal oxides	Iron, phosphate	Nickel, cobalt, manganese	Lithium metal + solid electrolyte
Cost (\$/kWh)	Medium	Low	High	Very high (currently)
EV Application	General EVs	Buses, entry EVs	Long-range EVs	Future EVs

2.3.4. Challenges in EV Battery Technologies

Despite significant technological progress, several challenges continue to hinder the widespread adoption and optimization of EV batteries:

1. **Performance Degradation:** Battery aging mechanisms, including solid electrolyte interphase (SEI) growth, lithium plating, and electrode degradation, lead to capacity fade and increased internal resistance over time.
2. **Thermal Management:** Maintaining uniform temperature distribution is critical for performance and safety. Non-uniform thermal profiles can accelerate degradation and increase the likelihood of thermal runaway.
3. **Safety Risks:** Lithium-ion batteries are susceptible to thermal runaway under mechanical, electrical, or thermal abuse conditions, necessitating robust monitoring and fault mitigation strategies.
4. **Fast Charging Limitations:** High charging rates can induce excessive heat generation and accelerate degradation, requiring optimized charging protocols.
5. **Material Constraints:** The reliance on critical raw materials such as lithium, cobalt, and nickel introduces concerns related to cost, availability, and environmental impact.
6. **End-of-Life Management:** Efficient recycling and second-life utilization remain essential for improving sustainability and reducing environmental burden [88].

2.3.5. Emerging Trends and Research Directions

Ongoing research efforts are focused on addressing the aforementioned challenges through innovations in materials, system design, and control strategies. Artificial intelligence-based BMS approaches have demonstrated improved accuracy in SOC and SOH estimation, enabling predictive maintenance and adaptive control [91].

Furthermore, advancements in battery thermal management systems, including liquid cooling and phase-change materials, are enhancing safety and performance under high-load conditions. The development of solid-state batteries represents a promising direction for achieving higher energy density and improved intrinsic safety, although challenges related to scalability and manufacturing persist.

Another emerging approach involves hybrid energy storage systems that combine batteries with supercapacitors to handle peak power demands while reducing stress on the battery. Additionally, second-life applications of EV batteries in stationary energy storage systems are gaining attention as a means to extend battery utility and improve lifecycle sustainability.

3. Classification of Traction Motors

The role of the traction motor is crucial for the advancement of the electrified transportation system. Electric motors are widely employed not only in domestic and industrial applications, but also play a vital role in the development of modern EVs and HEVs. Figures 12 and 13 present various types of AC and DC traction motors [92]. However, Figure 14 presents cross-sectional views of key traction motor technologies: IM, PMSM, and SRM. Advances in motor design, materials, electronics, and control techniques continue to enhance efficiency and reliability across various applications [72,93].

The choice of an electric propulsion system depends on factors such as vehicle constraints, driver demands, and energy sources [31,80]. Various traction motors are available today; however, the selection of any particular motor for an electric propulsion system depends on factors such as simplicity, durability, lightweight design, compact size, cost-effectiveness, flexible control, fault tolerance, low noise, high efficiency, and reliable performance across diverse driving conditions [79,80]. Additionally, market acceptance, influ-

enced by the availability and cost of compatible power converters, plays a very important role in the motor selection.




Figure 12. Classification of AC traction motors.




Figure 13. Classification of DC traction motors.




Figure 14. Traction motor technologies: IM, PMSM and SRM.

3.1. DC Motor

The technical and economic advantages of DC motors historically made them the preferred choice for most EV and HEV manufacturers initially [94]. Their widespread adoption was driven by simple speed control, technological maturity, lower manufacturing costs, and superior speed regulation compared to other electromechanical conversion devices [95]. However, the use of commutators and brushes in DC motors reduces reliability, rendering them unsuitable for maintenance-free and high-speed applications like EVs and HEVs. Additionally, issues such as electromagnetic interference, bulky construction, low

efficiency, and low specific power density further limit their suitability for EV and HEV applications [96]. Compact design and high efficiency are crucial for the development of an electrified transportation system; therefore, BLDC motors are under consideration.

3.2. Induction Motor

IM is one of the most suited and mature electric propulsion systems due to its high dynamic performance. IMs can extend their constant power range up to 4–5 times the base speed, meeting traction requirements. However, high-speed operation is constrained by pullout torque limitations [95]. Figure 15 shows a typical speed–torque characteristic of an IM and its variations with rotor speed.




Figure 15. A typical speed–torque characteristics curve of IM.

IM propulsion systems have various advantages, including minimal noise and ripples, mature and efficient power electronic drivers, reliability, and simplicity of design. Variable frequency and voltage control can produce high starting torque and low starting current, while high efficiency in IMs can be achieved through minimum slip control [96,97]. Pulse width modulation (PWM) is typically employed to control the speed of the motors. However, it generates unwanted common-mode voltage, reducing motor life [98]. IM has achieved efficiency of 90%; however, thermal losses, larger size, and heavy weight are concerns that reduce the manufacturer's attraction [41]. PMSMs challenge IM dominance with 95% efficiency, higher power density, and lower rotor losses.

To enhance IM drive efficiency, several control techniques tailored for HEV applications have been proposed [78]. An optimization technique is proposed for deep regulation of magnetic flux to improve voltage drop and efficiency [79].

The development of better control techniques eroded IM's cost advantage. EV vane makers and Tesla had already deployed PMSM in their new models. IM's speed control is broader and much easier to achieve than PMs. Due to the absence of PM's high energy density, IM has a low torque performance. IMs operate with less than 3% slip, minimizing rotor core losses and high efficiency [88]. However, maximum efficiency is achieved when iron and copper losses are balanced [89]. Nevertheless, core and resistance losses are the significant contributors to IM efficiency when compared to stray losses [98,99]. The following Table 12 provides specifics about each loss.

Table 12. Characteristics of IM [99–101].

Types of Losses	Proposed Method	Adopted Technology		Advantage
Copper loss	Increase conductivity	Replacement of an aluminum conductor with a copper conductor		Copper loss reduced Efficiency improved
Iron losses	Improve model calculation	3D finite element analysis		Calculation efficiency improved
		A combination of vector control and the Fe loss model		Independent control of torque and speed Fe loss reduced
		Develop a controller on the basis of a stationary coordinate system		Fe loss reduced Less calculation Simplified model
	Fe loss model-based control technology	Combination of DTC and Fe loss model		Fe loss reduced Torque response is faster
		Combination of sliding mode control and Fe loss model		Robustness improved
	Search method	Iterative flux for input power reduction		Fe loss reduced Motor parameters are not sensitive
	Combine search and model	Initially, use the Fe loss model for the calculation of the magnetic flux approximate value, then go for the optimal value		Fe loss reduced Comparatively high accuracy Calculation speed increased
Structure losses	Employed Method	Efficiency (%)	Torque Density (Nm/L)	Advantage
Prove the structure design	Increase axial length	88	Not given	Efficiency improved Implementation is easy
	Skew rotor structure	87.3	33.17	Efficiency improved
		86.4	36.17	Torque performance improved
	Multi-objective optimization	89 86.5–87.7	22.4–39.1	To maximize efficiency, other performances are taken into account

3.3. PMSM

Since the last decade, PMSM has been the most demanding motor in the EV sector. Permanent magnets could be surface-mounted on the rotor or interior-mounted, with the magnet inside the rotor [102,103]. Advantages of PMSM include excellent efficiency, high starting torque, compact size, smooth control, decreased noise, low power consumption, good dynamic performance, and dependability [104,105]. Furthermore, low rotor losses and high-power density make them the most suitable motor [106–108]. The smaller size enables greater dependability, effective heat dispersion into the atmosphere, and maintenance-free operation. However, the motor cost is high due to rare-earth permanent magnets and necessitates the use of a motor control drive [109]. Sensored and sensorless control strategies are employed to achieve further improvement. A hybrid smooth switching, speed controlling strategy from low to high speed is discussed [110,111]. The impact of the air gap of PMSM is studied for shaft torque [110]. A fault-tolerant issue is studied by increasing the power electronic converter rating due to high-speed control and management of back emf.

As rare-earth magnets are costly due to their non-renewable nature and low productivity, alternative magnets or rare-earth-free designs are under consideration. However, the lower power density of rare-earth-free magnets remains a challenge. Ongoing studies aim to improve the power density and efficiency to make them more viable for advanced applications. Ferrite PM, which has a very low material cost, is one of the least expensive alternatives. A roadmap has been proposed to improve the efficiency and power density

with reduced cost [112]. The following is a comparison of two rare-earth magnets (Samarium cobalt and Neodymium) and two rare-earth-free magnets (Alnico and Ferrites), as demonstrated in Table 13 [113].

Table 13. Comparison of magnetic materials [114,115].

Property (Symbol, Unit)	Alnico	Ferrites (Ceramic)	Samarium Cobalt	Neodymium
Coercive Force (H_c , kA/m)	37–143	0.23–0.41	480–840	760–1030
Density (d , g/cm ³)	6.8–7.3	4.9	8.4	7.4
Electric resistivity (r , Ω cm)	$(50–75) \times 10^{-6}$	10^{-6}	$(53–86) \times 10^{-6}$	160×10^{-6}
Max. Service Temperature (T_{max} , °C)	450–550	800	300–350	150
Max. energy Product ($(BH)_{max}$, kJ/m ³)	10.7–71.6	8.35–31.8	130–240	220–336
Price, USD/k	58	7.1	100	75
Remanence (BR, T)	0.7–1.28	0.23–0.1	0.83–1.16	1.00–1.41

3.4. BLDC Motors

BLDC motors are electronically commutated; they do not require brushes, slip rings, or a physical commutator [116]. The windings are typically concentrated in nature and permanent magnets are placed on the rotor. A multi-step switching sequence is employed to generate a trapezoidal waveform. The back emf is employed to adjust the speed of the motor. Hall sensors are employed to regulate the speed and position of the rotor. BLDC motors have a very high efficiency, making them the best choice for electric bikes and rickshaws [117]. Furthermore, advantages include compact size, high efficiency, high starting torque, low power consumption, and reduced thermal losses [118]. However, the disadvantages of BLDC motors are the cost of neodymium magnets, magnetic field weakening in the in-runner motors and soft starting due to the high starting torque. An online monitoring algorithm is proposed to improve the reliability of the BLDC motor [119].

3.5. Switched Reluctance Motor (SRM)

SRM is described as an electronically commutated variable reluctance motor that employs reluctance torque to function. The stator has a single winding that is concentrated in nature, with no PMs in the rotor. Typically, stator poles have parallel sides, which means the edges are straight and equidistant. SRM operates on the variable reluctance concept, which means that the rotor will always follow the path with the least reluctance. The excitation sequence of the stator poles ensures that the SRM moves continuously. The number of poles on the rotor is mostly responsible for SRM's outstanding performance. Furthermore, steady torque and power are impaired at high speeds. As a result of its simplified structure and low cost, the SRM is a highly suitable motor for electric vehicles. Additional benefits include high starting torque, good efficiency, a wide speed range, and a high torque to inertia ratio [120]. However, complex excitation circuitry, high torque ripples, iron core losses, noise, and increased production costs are the concerns [72,121,122]. Newly proposed C-core hybrid SRM topologies and optimized magnet placement strategies enhance torque density while keeping structural robustness and low cost. Further studies are being conducted on the reduction of high torque ripples to strongly utilize these motors in electrified vehicles [123,124]. Comparison of torque ripple reduction strategies in SRM is demonstrated in Table 14.

Table 14. Comparison of torque ripple reduction strategies in SRM [122].

Method	Employed Techniques	Advantages	Disadvantage
Increase the number of rotor poles	Increased number of poles of the rotor	Average torque increased Torque ripple reduced Copper loss reduced	Complicated rotor configuration. Rotor material increased
Pole shape	Modification in rotor shapes, pole arc, pole shoe and air gap	Torque ripple reduced Efficiency increased High-speed performance	Complicated optimization for offline calculation
Modulation in current and angle	Turn-on and -off angles are optimized	Efficiency improved Torque–speed capability enhanced Torque ripple reduced	Limited current and torque control Large memory required to store current profiles
ATC and DTC	The hysteresis controller regulates the torque with online estimation	Torque ripple reduced Direct control of the instantaneous torque	Machine parameters prior knowledge required
Torque sharing function (TSF) method	TSF profile definition Implement hysteresis control with current reference from torque reference	Controlled torque easily Torque waveforms determined Smooth torque over a broad speed range	i-T- θ characteristics needed offline designed torque waveform
Feedback Loop (FBL) control	Transformation of a non-linear model into a linear model	Torque ripple reduced FBL is free from the non-linear term Provide the required decoupling of the current	Complex linearization algorithm. No-adaptability of uncertain parameters changes
Iterative learning control (ILC)	Current compensation to the phase current reference is added for current tracking	System parameter identification not needed Perfect current tracking on different conditions is achieved	Degraded performance to transient Complex learning control law Iteration cycle restricted
Intelligent Control	Offline and Online optimization of phase current	Strong self- learning Torque ripple reduced Adaptive ability Independent of machine parameters	Complex computational algorithm

3.6. Advanced Motor for EVs and HEVs

To boost torque density, a new breed of electric machine, known as a double-stator (DS) machine, is proposed [122]. The innovation of the double-rotor electric machine is developed from research on HEV power split devices [122]. Existing power split systems use mechanical gearboxes as power transmission devices, which are inefficient, noisy, and vibrational. Therefore, double-rotor electric machines can efficiently handle the problem since they employ magnetic force to convey torque and power, eliminating the need for physical touch. However, there is still a long way to go before the deployment of double-rotor and DS electric machines commercially. Their mechanical construction is complex and their power factor is low when compared to single-rotor electric machines. In some ways, these new topologies have a compact and complicated structure that is suited for EV powertrains as shown in Figure 16 [125]. DS IM provides more robustness and reliability with SMC [126]. Design analysis of an oil pump motor with PM is discussed to achieve maximum efficiency by using a different rotor [127]. Multiple criteria of different motors are discussed and analyzed [128].




Figure 16. Advance motors [125].

Table 15 compares all motor key characteristics in order to provide guidelines for choosing an appropriate motor for EV applications. PMSM is the most demanding motor, followed by IM and SRM, whereas DC motors are employed for small applications and are not chosen for EVs. Different traction motors employed in commercially available EVs and HEVs are shown in Table 16.

Table 15. Comparison of different motors for EV applications [43,73,98,128].

Characteristics	IM	PMSM	SRM	DC-M
Manufacturability	5	3	4	3
Controllability	5	4	3	5
Cost	5	3	4	3.5
Robustness	5	4	4.5	3.5
Reliability	5	4	4.5	3
Lifetime	5	4	4.5	3.5
Torque ripple/noise	4.5	4	3	3.5
Technical maturity	4.5	4	3.5	5
Efficiency	4	5	4.5	3
Size and weight	4	4.5	4	3
Power Density	4	5	3.5	3
Overload capability	4	4.5	4	3
Speed range	4	5	5	2.5
Torque density	3.5	5	4	3
Potential for improvement	3	4.5	5	2.5
Trend	2	5	3	1
Total Score	67.5	68.5	64	51

Table 16. Motor employed in EV or HEV [113,129–133].

EV Model	Motor Employed	Power kW	Year
Rivian R2	PMSM	220	2026
Toyota Land Cruiser Se	PMSM	400	2026
Volkswagen ID. GTI	PMSM	210	2026
Lucid Gravity	PMSM	600	2025
BYD Seal	PMSM	150	2025
Tesla Model Y long range RWD	IPM-SynRM	250	2024
BYD Seagull	PMSM	55	2024
Tesla Model 3 Standard Range	PMSM	239	2023
Tesla Model Y long range RWD	IPM-SynRM	250	2024
BYD Seagull	PMSM	55	2024
Tesla Model 3 Standard Range	PMSM	239	2023

Table 16. Cont.

EV Model	Motor Employed	Power kW	Year
Renault Megane E-Tech	PMSM	96	2022
Porsche Taycan ST	PMSM	300	2022
Mercedes-Benz EQS	PMSM	265	2022
Kia Niro EV	IM	150	2022
Tesla Model Y SR	IM	150	2022
Tesla Model 3 long range	IPM-SynRM	258	2021
BYD Tang	PMSM	380	2021
Hyundai Ioniq 5	PMSM	125	2021
Tata Nexon	PM	85	2020
Volkswagen id.3	PM	100	2020
Kia e-Soul	PM	100	2020
Mercedes EQC	IM + IM	300	2019
Smart EQ Fortwo	PM	60	2019
BYD S2	PM	70	2019

4. Control Techniques of Traction Motor

There are various techniques for controlling motor drive and all have their particular significance when dealing with motor control with respect to specific applications. Some are simple but less accurate (e.g., scalar control), while others are sophisticated but produce accurate findings (DTC, FOC, MPC, etc.). As a result, traction motor control is critical while heading towards smart applications [134–136]. Both high dynamic response and accurate speed control are the key characteristics required from the traction drive operation. Advanced motor controllers are responsible for providing not only good steady state, but also better transient response. In general, key characteristics required from the traction motor drive include the capability of high torque at low speed for starting and climbing, high power for cruising, fast torque response, and high-power density. Moreover, it should maintain high efficiency across wide speed ranges, while being compact, lightweight, cost-effective, reliable, and robust under varying vehicle operating conditions. Figure 17 illustrates the typical speed–torque characteristic curve essential for EV and HEV applications.




Figure 17. Typical speed–torque characteristic curve for EV and HEV [43].

current and q-axis current, where the optimal operation often enforces zero d-axis current for surface-mounted PMSMs. This leads to high efficiency and fast dynamic response. However, in interior PMSMs, reluctance torque introduces an additional degree of freedom, requiring advanced control strategies such as maximum torque per ampere (MTPA) and field-weakening operation [142].

In BLDC motors, the application of conventional VC is less straightforward due to the trapezoidal back-EMF profile. Instead, quasi-vector control or modified FOC approaches are employed to approximate sinusoidal current excitation, thereby reducing torque ripple and acoustic noise. The main challenge lies in accurate commutation and current shaping, especially at low speeds where position estimation becomes less reliable [142]. Recent improvements include hybrid control schemes that combine FOC with predictive current control to enhance dynamic performance.

For SRMs, the absence of permanent magnets and the highly nonlinear inductance profile complicate the direct application of classical VC. Nevertheless, extended vector control frameworks have been proposed by defining pseudo d-q-axes or utilizing flux-linkage-based transformations. These approaches aim to achieve smoother torque production and reduced ripple. However, strong magnetic saturation and phase coupling effects limit the accuracy of such models [142]. Consequently, adaptive and nonlinear control techniques are increasingly integrated with VC to address these limitations.

Despite its advantages, VC faces several challenges across all machine types, including sensitivity to parameter variations, reliance on accurate rotor position information, and increased computational complexity. These issues become more pronounced in EV applications where wide speed ranges and varying load conditions are common.

There are several improved FOC controllers available to provide the required operational performance for EVs. EV acceleration can be improved by adjusting the maximum torque and ampere ratio with FOC. Field weakening design and FOC can be employed to increase the full speed operation of EVs. It will increase the motor's speed range. Precision in output torque and efficiency are dependent on accurate rotor flux estimation, which is accomplished via FOC [143,144]. The FOC scheme associated issues and some suggested solutions are summarized in Table 17.

Table 17. FOC issues' proposed solution and outcomes [130,132,133,145].

Issues	Proposed Techniques/Method	Advantages	Disadvantages
Assessment of suitability	IFOC for asynchronous traction motor drives for EVs	Suitability of IFOC is verified.	No additional improvement suggested.
Degradation of speed performance due to thermal effects	LPV controller–observer is employed	Low current and supply voltage are required. Noise rejection improved. Better tracking performance. Robust against inverter nonlinearity and control mechanism.	More complicated than the conventional scheme.
Motor parameter variation impact on the performance.	MRAS estimator based on back EMF is proposed	Independent of stator inductance and resistance. Improved speed performance. Tracking performance, energy consumption and recovery.	The results presented consider only 50% change in the stator parameter.
Comparison of conventional PI and Fuzzy logic controller.	Use of the ECE-15 drive cycle	Improved power factor and efficiency. Reduced voltage distortion.	Improvement in IFOC is not considered.
Analysis of drivetrain efficiency.	FOC PMSM. Three-phase inverters, battery connected with a DC-DC converter.		The energy optimization technique is not considered a variation in motor flux.

4.3. Direct Torque Control

DTC needs a switching table to determine the current state of the inverter in order to generate the correct output torque [146]. It has a PI speed controller and two hysteresis controllers. It employs two different controllers, achieving a better dynamic response. A three-level comparator serves as a torque hysteresis controller, while a two-level comparator serves as a flux hysteresis controller with a different bandwidth [89,147].

The flux hysteresis bandwidth influences motor current distortion, while the torque hysteresis bandwidth influences switching losses and switching frequency. As a result, it is critical to select the hysteresis bandwidth with care. DTC allows for precise and rapid control of IM torque and flux without the use of sophisticated controlling algorithms [148]. The use of DTC for IM traction drive control is demonstrated in Figure 19.




Figure 19. DTC for IM traction drive control.

DTC includes switching vector modulation (SVM), switching table-based, and direct self-control. Its qualities of not being parameter dependent make it particularly significant for electric car applications, with resilience strengthened as it is not dependent on parameter variation and a smaller number of controllers. DTC is appropriate for high-speed operation and enables frequent starting/stopping and acceleration. DTC allows for the implementation of strong flux weakening control as well as dynamic motor operating. DTC has a slow start-up response, with high current and torque ripples [94,149]. However, low speed flux, torque estimate, torque control, torque ripples, and high current all have an impact on the DTC electric drive system. As a result, several variations or combinations are required for improved performance. DTC is implemented using a fuzzy logic controller (FLC) and a PI controller, and it is compared to a sliding mode control (SMC) scheme as demonstrated in Table 18 [150]. Ben Salem et al. proposed an inventive adaptive fractional order sliding mode (FO-SM) control method specifically designed for DTC-SVM applied to IM drives, which addressed the issues of parameter variation in load and temperature with better control and reduced torque ripples [151].

In PMSMs, DTC has been adapted to exploit the fast torque response while addressing the inherent torque and flux ripple associated with hysteresis-based switching. Unlike induction motors, PMSMs exhibit a strong coupling between flux and torque due to the presence of permanent magnets. This necessitates refined switching table designs and the incorporation of SVM to achieve constant switching frequency and reduced ripple [142]. Recent advancements include model predictive DTC (MPDTC), which enhances control accuracy by predicting system behavior over a finite horizon [142].

For BLDC motors, DTC implementation is comparatively simpler due to the discrete nature of back-EMF and commutation. However, conventional DTC strategies may lead

to significant torque ripple because of the non-sinusoidal characteristics. Modified DTC schemes incorporating duty ratio control and advanced switching strategies have been proposed to mitigate these effects [142]. Additionally, the integration of intelligent controllers has improved robustness under parameter uncertainties.

In the case of SRMs, DTC is particularly attractive due to the direct relationship between phase inductance and torque production. Unlike PMSMs and BLDC motors, SRMs do not require rotor flux estimation, simplifying control design. However, the highly nonlinear magnetic characteristics and significant torque ripple present major challenges. Advanced DTC approaches, including torque-sharing functions and predictive control, have been developed to address these issues. These methods aim to distribute torque production smoothly across phases, thereby reducing vibration and acoustic noise [142].

Although DTC provides fast transient response and reduced dependence on machine parameters, it suffers from variable switching frequency, high torque ripple, and difficulties in low-speed operation. These limitations necessitate the use of enhanced modulation techniques and intelligent control algorithms.

Table 18. DTC issues' proposed solution and outcomes [131,144,145,152].

Issues	Proposed Method	Advantages	Disadvantages
Minimization of issues in conventional DTC	Utilization of a multi-level neural network instead of a conventional switching table	Torque and current ripples were reduced significantly	Efficiency improvement is not investigated. A simulation study on the common drive cycle has not been done. A very small-scale model based on simulation is employed but without using recognized drive cycles.
Efficiency improvement of DTC	Variable flux reference selection	Efficiency improved	
Control method for a four-in-wheel EV drive	An additional electronic differential system provides speed references to in-wheel motors. Modified torque hysteresis controller.	Handling improved. Good dynamic and steady-state speed tracking.	High torque ripple is observed
Reduction in torque ripple in CDTC	The MRAS estimator is employed for rotor speed and stator resistance.	Torque error, torque ripple, current THD and flux error improved.	Optimal bandwidth is employed, which may be difficult to achieve in large-scale systems for EVs.
Improved integral time-weighted absolute error (ITAE) and speed response	Fractional order PI controller.	ITEA is reduced as compared to the PI controller. Lower energy requirement. Improved efficiency.	Issues associated with CDTC have still not improved.
Efficiency improvement in DTC	Stator current minimization technique. DC link voltage is employed.	Excessive computational resources are not required. No convergence issues. Not sensitive to parameter variation.	Maximum torque/ampere is allowed during operation but below base speed.

The applicability of VC and DTC across different traction motors highlights a fundamental trade-off between control precision and implementation complexity. VC offers superior steady-state accuracy and low torque ripple, making it suitable for applications requiring smooth operation, such as PMSM-based EV drives. In contrast, DTC provides

faster dynamic response and simpler structure but requires additional modifications to achieve comparable performance.

Recent research trends indicate a convergence of these techniques through hybrid control strategies. For instance, predictive control frameworks are being integrated with both VC and DTC to improve dynamic performance while maintaining robustness. Similarly, AI-based approaches are increasingly employed for parameter estimation, fault diagnosis, and adaptive control across all motor types [142].

Another important direction involves the development of sensor-less control methods to eliminate mechanical sensors, thereby reducing cost and improving system reliability. This is particularly relevant for EV applications where harsh operating conditions can affect sensor performance.

5. Artificial Intelligence Controllers

Outstanding operational performance of the present artificial intelligence controllers (AICs) has piqued the interest of many researchers in recent years [144]. The primary goals of ICs are to obtain high dynamic performance while reducing the processing time and increasing the robustness against the machine parameter dependency [145,146]. Motor speed and torque are regulated using optimized parameter settings for PD, PI, and PID controllers and model predictive control (MPC) enables fault-tolerant operation of the system. A hierarchical control method integrating MPC with a torque vectoring algorithm has been proposed to reduce torque requirements and achieve significantly improved results [153]. Furthermore, advanced controllers are necessary to achieve increased accuracy and precision control. artificial neural network controller (ANNC), adaptive neural fuzzy inference system controller (ANFIS Controller), and fuzzy logic controller (FLC) are some of the advanced controllers [147,148].

5.1. MPC Controller

In a recent study, a major shift was observed towards the MPC controller to deal with the EV control system and real-time parameter improvement. ANFIS-based MPC is proposed to reduce current and torque ripples [154]. A hierarchical predictive optimal control is proposed to control IPMSM drives to improve the torque control with improved energy efficiency of 37% [153]. Hybrid MPC is employed to deal with stability control for a distributed drive vehicle on real-time conditions of different road surfaces at high speeds [155]. Another improved drive system for five-phase PMSM is proposed based on MPC. This article proposed an emerging speed tracking technique to deal with the faults in open-phase motor [156].

5.2. ANN Controller

To increase efficiency and fast processing, conventional technologies are progressing towards intelligent systems. An intelligent controller is necessary to achieve higher precision and eliminate losses in the conventional controllers. Figure 20 shows the Basic 3-inputs and 1-output structure of an ANN. Motor parameter dependency, torque ripples and efficiency are the issues in motor controllers. An adaptive neural network-based torque controller is designed for the inner current loop to mitigate torque ripple and improve propulsion motor efficiency, which improves EV performance [157,158]. An ANN-based model predictive controller is used for optimal torque performance [153]. Another ANN-based model with MPC is employed for better speed estimation [156]. ANFIS is compared to FLC; the results show that ANFIS has greater dynamic performance [159,160].




Figure 20. ANN architecture.

ANN is utilized as an intelligent controller to estimate motor parameters and detect motor faults [161,162]. A more robust controller is proposed in [163], in which PMSM will benefit from two innovative deep neural network (DNN)-based vector controllers. The numerical burden and complex computation are the issues to be addressed in this type of controller.

5.3. ANFIS Controller

The ANFIS controller is a hybrid of artificial neural networks and fuzzy logic. It is designed to take advantage of the unique qualities of both networks. It is an intelligent controller employed to estimate motor parameters without the use of a mathematical model. The ANFIS model prediction current control method is employed to reduce torque and current ripples in PMSM [159]. ANFIS has the problem of requiring a big data set for data training and learning, which takes a long time and requires complicated computing processing. As compared to a traditional controller, advanced ANFIS and FLC produced efficient and better HEV performance [162–164]. An ANFIS-based fractional order PID controller exhibits robustness against external disturbances and provides potential EV speed regulation control [165]. PI-ANFIS and particle swarm optimization (PSO)-PI-ANFIS algorithms were contrasted, and PSO-PI-ANFIS forces the motor speed to closely follow its reference. Furthermore, it models with simpler structures, which are more reliable than PI-ANFIS models [166]. The advantages of these approaches are their high accuracy and extremely fast computational speed for prediction and adjustment of the classical PI controller parameters; however, they have a complex structure [167].

5.4. FLC

Fuzzy Logic Controllers (FLCs) are intelligent controllers that use a fuzzy logic network. FLC operates in four basic steps. Fuzzification comes first, followed by the rule base, inference engine, and defuzzification. The initial stage is to determine the number of inputs, the position of fuzzy MFs, and the outputs. As stated in the equation, the inputs are error and error change [168].

$$e(t) = x_n^* - x_n \quad (2)$$

$$de(t) = e(t) - e(t - 1) \quad (3)$$

By analyzing each input in a group, the fuzzification design characterizes the right linguistic value with inputs. This is known as a one-of-a-kind MF label. The number of MFs employed determines the precision of the FLC. Triangular or trapezoidal MFs are employed to depict the output. The architecture of an inference engine indicates the control rules and language concepts of decision-making. It is usually divided into two categories,

Mamdani and Takagi-Sugeno. Due to its basic structure and easy design, the Mamdani approach is widely employed. As a fuzzy rule, the if-then clause is employed. The basic architecture of FLC is demonstrated in Figure 21.




Figure 21. FLC architecture.

The final stage is defuzzification, which generates the output values as crisp values. This method regulates the crispness of output MFs and makes minor adjustments. FLC has found application in the operation of HEVs as it is simple, effective, and offers the option of real-time supervisory control based on a deterministic rule-based approach [145]. When the FLC's speed tracking performance was compared to that of the PI controller, the FLC performed better at constant speed [169]. FLC is employed for speed improvement; therefore, the speed of the motor is improved, and the error is reduced [170]. FLC can enhance starting current amplitude and conserve more power [171]. Improved adaptive FLC-based techniques are used for traction drive and energy management systems [172]. Moreover, the Neuro-fuzzy controller is more advanced and is based on a nonlinear and adaptive method. As a result of its precision, extensive supervisory control, and low error rate, this adaptable system can be employed in even more complex real-time circumstances [173]. The comparison table of AI controllers is shown in Table 19.

Table 19. Comparison of AI controllers [166,174–178].

Aspects	MPC	ANN	ANFIS	FLC
Concept	Model-based optimization using a prediction horizon	Data-driven learning to approximate nonlinear system behavior	A hybrid of fuzzy logic + neural networks for adaptive rule-based learning Combines interpretability + learning	Rule-based control using linguistic variables and expert knowledge
Advantages	Handles multiple constraints Optimal control performance Fast dynamic response Suitable for multi-motor coordination	High adaptability Strong nonlinear approximation Fault tolerance Self-learning capability	Effective for nonlinear systems Adaptive parameter tuning More robust than pure FLC	Simple implementation No mathematical model required Robust to uncertainties Low computational cost
Limitations	High computational complexity Requires an accurate model Difficult real-time implementation at high speed Torque control in PMSM drives	Needs large training data Lacks interpretability Generalization issues No guaranteed stability	Rule explosion problem High computational cost Complex training Scalability issues	Heuristic design No optimality guarantee Limited adaptability Performance depends on rule tuning
Applications	Multi-motor coordination Stability and energy optimization	Fault diagnosis and tolerant control Parameter estimation Adaptive torque control	Adaptive speed control Hybrid EV control systems Nonlinear drive control	Low-cost EV controllers Basic speed and torque control Real-time embedded systems

6. Present Challenges and Recommendations

In spite of substantial research and development work, EVs and HEVs still face numerous technical and large-scale commercial adoption challenges that require more research, innovation, and technological refinement for optimized performance and mature technology.

6.1. EV and HEV Powertrain Architectures

6.1.1. Present Challenges

- High System production cost, weight optimization and complexity due to multi-source systems, such as engine, battery, power converters and generator integration, are still persistent issues.
- Limited energy density of current batteries results in restricted driving range and frequent recharging requirements.
- Limited battery capacity, cost-effective energy storage and charging infrastructure limitations remain key challenges for pure EVs.
- Energy management and packaging constraints in compact vehicle designs.
- Limited energy storage capability and inefficient energy management under highly dynamic driving conditions.
- Reliability and durability concerns in commercial fleets with high utilization rates.

6.1.2. Recommendations

- Development of an optimized lightweight design of modular and scalable architectures approach, which integrates renewable energy sources and is easier to implement both on light-duty and heavy-duty vehicles.
- Integration of advanced solid-state batteries, supercapacitors, and hybrid energy storage systems for improved driving range and lifecycle.
- Application of predictive energy management strategies using artificial intelligence.
- Enhanced energy management strategies using predictive and data-driven models can further optimize fuel economy and reduce emissions in hybrid power-split configurations.
- Research on power electronics with wide-bandgap semiconductors (SiC, GaN) to improve efficiency and reduce system losses.

6.2. Traction Motors

6.2.1. Present Challenges

- Dependence on rare-earth permanent magnets for PMSMs increases the overall system cost and supply chain vulnerability.
- IMs suffer from relatively high losses at low load and limited efficiency at high speeds.
- BLDC motors and SRMs offer potential alternatives but require further refinement in noise reduction, efficiency, and control mechanism complexity.
- Thermal limitations reduce lifespan and reliability in high-power applications.
- Packaging and weight constraints limit the adoption of large machines in compact EV design.

6.2.2. Recommendations

- Advancement and development of rare-earth-free designs such as SRMs and advanced induction motor topologies.
- Exploration of axial flux and transverse flux machines offering higher torque density and compact design.
- Improved cooling strategies, such as direct liquid cooling and advanced thermal interface materials, for better heat dissipation.

- Design of fault-tolerant motor architectures to ensure reliability in safety-critical applications.
- Optimization of multi-motor propulsion systems (e.g., in-wheel motors) for distributed drive and enhanced control flexibility.
- Integration of novel materials, additive manufacturing, and fault-tolerant designs is expected to play a key role

6.3. Motor Control Techniques

6.3.1. Present Challenges

- Conventional methods such as vector control and DTC are sensitive to parameter variations during operation and disturbances.
- Both aforementioned methods are employed for high-performance applications such as in EVs and HEVs, but suffer under dynamic and uncertain driving conditions, torque ripple and high efficiency across wide operating ranges.
- High computational demands in real-time control for nonlinear and multi-variable systems are required.
- Face challenges in ensuring robustness under variable driving conditions, limited robustness in handling motor faults, inverter faults, and grid interaction scenarios.
- Cybersecurity risks in intelligent and connected control systems are becoming increasingly important for safety and reliability in autonomous and connected EVs.

6.3.2. Recommendations

- Incorporation of artificial intelligence, machine learning, and adaptive control for predictive torque control, less computational burden and efficiency optimization.
- Development of self-tuning algorithms that can adjust controller parameters automatically under changing load and environmental conditions.
- Advanced sensor-less control techniques are also important to reduce hardware costs and improve reliability.
- Integration of fault-diagnosis and fault-tolerant control techniques for enhanced safety.
- Research into distributed and cooperative control strategies for multi-motor EV architectures.
- Enabling motor controllers to support vehicle-to-grid (V2G) operations, including bidirectional power flow and grid stabilization.

These challenges in EV and HEV architectures, traction motors, and motor control require interdisciplinary collaboration across materials science, power electronics, control theory, artificial intelligence and system integration. Future EV and HEV systems will likely rely on high-efficiency, rare-earth-free traction motors, modular powertrain designs, and intelligent controllers capable of adaptive learning. Together, these developments will accelerate the transition toward sustainable, reliable, and cost-effective electrified transportation.

7. Conclusions

This paper has presented a structured and critical assessment of electric and hybrid electric vehicle powertrain architectures, traction motor technologies, and advanced control methodologies. A structured comparative framework was employed to evaluate series, parallel, and series-parallel architectures with respect to efficiency, cost-performance balance, structural complexity, and commercial feasibility. The review further analyzed dominant traction motor technologies with emphasis on torque-speed capability, efficiency mapping and material considerations, identifying both technological breakthroughs and unresolved limitations. Although PMSM, IM, and related topologies demonstrate high performance, challenges associated with rare-earth utilization and scalability remain open

research concerns. Advanced control strategies, including vector-based methods, direct torque approaches, model predictive control and emerging AI-enabled frameworks, were comparatively evaluated in terms of dynamic response, robustness, computational demand, and scalability. Key research gaps were explicitly identified in areas such as rare-earth material dependency, high-speed efficiency optimization, integrated energy management, and real-time intelligent control implementation. The findings underscore the need for multidisciplinary innovation in next-generation motor design, power electronics integration, and adaptive control architectures to achieve reliable, high-efficiency, and economically sustainable electrified mobility solutions.

Author Contributions: Conceptualization, investigation, formal analysis and writing—original draft preparation, S.H.I. and S.J.R.; formal analysis and data curation, S.J.; supervision and writing—review and editing, M.J. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: No new data were created or analyzed in this study.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Nordelöf, A.; Grunditz, E.; Lundmark, S.; Tillman, A.M.; Alatalo, M.; Thiringer, T. Life cycle assessment of permanent magnet electric traction motors. *Transp. Res. D Transp. Environ.* **2019**, *67*, 263–274. [CrossRef]
2. Kumar, R.R.; Alok, K. Adoption of electric vehicle: A literature review and prospects for sustainability. *J. Clean. Prod.* **2020**, *253*, 119911. [CrossRef]
3. Available online: <https://iea.blob.core.windows.net/assets/ed5f4484-f556-4110-8c5c-4ede8bcb637/GlobalEVO Outlook2021.pdf> (accessed on 22 July 2022).
4. Yong, J.Y.; Ramachandaramurthy, V.K.; Tan, K.M.; Mithulanathan, N. A review on the state-of-the-art technologies of electric vehicle, its impacts and prospects. *Renew. Sustain. Energy Rev.* **2015**, *49*, 365–385. [CrossRef]
5. Richardson, D.B. Electric vehicles and the electric grid: A review of modeling approaches, Impacts, and renewable energy integration. *Renew. Sustain. Energy Rev.* **2013**, *19*, 247–254. [CrossRef]
6. Patel. Fuel Economy Model 2022. 2022.
7. Maroti, P.K.; Padmanaban, S.; Bhaskar, M.S.; Ramachandaramurthy, V.K.; Blaabjerg, F. The state-of-the-art of power electronics converters configurations in electric vehicle technologies. *Power Electron. Devices Compon.* **2022**, *1*, 100001. [CrossRef]
8. Mei, J.; Zuo, Y.; Lee, C.H.T.; Kirtley, J.L. Modeling and Optimizing Method for Axial Flux Induction Motor of Electric Vehicles. *IEEE Trans. Veh. Technol.* **2020**, *69*, 12822–12831. [CrossRef]
9. EV-Volumes-The Electric Vehicle World Sales Database. Available online: <https://www.ev-volumes.com/country/total-world-plug-in-vehicle-volumes> (accessed on 5 July 2023).
10. Propfe, B.; Redelbach, M.; Santini, D.J.; Friedrich, H. Cost analysis of Plug-in Hybrid Electric Vehicles including Maintenance & Repair Costs and Resale Values. *World Electr. Veh. J.* **2012**, *5*, 886–895. [CrossRef]
11. Electric Vehicle Sales Review Q2 2025 Foresight to Drive the Industry. 2025. Available online: <https://www.pwc.com/hu/hu/kiadvanyok/assets/pdf/strategyand-electric-vehicle-sales-review-q2-2025.pdf> (accessed on 12 October 2025).
12. Graham, E. The EV Leapfrog-How Emerging Markets Are Driving a Global EV Boom. Ember. Available online: <https://ember-energy.org/latest-insights/the-ev-leapfrog-how-emerging-markets-are-driving-a-global-ev-boom/> (accessed on 24 March 2026).
13. Wimmer, H.; Neuhausen, J. “Electric Vehicle Sales Review Q4-2025 | PwC and Strategy&,” January 2026. Available online: <https://www.strategyand.pwc.com/de/en/industries/automotive/electric-vehicle-sales-review-q4-2025.html> (accessed on 11 February 2026).
14. ‘Despite Q4 Collapse, 2025 EV Sales Decline Only 2% Versus 2024; Policy Shifts, New Product Set Stage for Next Chapter-Cox Automotive Inc.’ Cox Automotive Inc. 13 January 2026. Available online: <https://www.coxautoinc.com/insights-hub/q4-2025-ev-sales-report-commentary/> (accessed on 24 February 2026).
15. Singh, K.V.; Bansal, H.O.; Singh, D. *A Comprehensive Review on Hybrid Electric Vehicles: Architectures and Components*; Springer: Berlin/Heidelberg, Germany, 2019. [CrossRef]
16. Prasanthi, A.; Shareef, H.; Errouissi, R.; Asna, M.; Mohamed, A. Energy Conversion and Management: X Hybridization of battery and ultracapacitor for electric vehicle application with dynamic energy management and non-linear state feedback controller. *Energy Convers. Manag. X* **2022**, *15*, 100266. [CrossRef]

17. Baskar, S.; Vijayan, V.; Premkumar, I.J.I.; Arunkumar, D.; Thamaran, D. Design and material characteristics of hybrid electric vehicle. *Mater. Today Proc.* **2020**, *37*, 351–353. [CrossRef]
18. Mahmoud, M.; Garnett, R.; Ferguson, M.; Kanaroglou, P. Electric buses: A review of alternative powertrains. *Renew. Sustain. Energy Rev.* **2016**, *62*, 673–684. [CrossRef]
19. Wang, N.; Tang, L.; Pan, H. A global comparison and assessment of incentive policy on electric vehicle promotion. *Sustain. Cities Soc.* **2019**, *44*, 597–603. [CrossRef]
20. El, M.; Benbouzid, H.; Diallo, D.; Zeraoulia, M. Advanced Fault-Tolerant Control of Induction-Motor Drives for EV / HEV Traction Applications: From Conventional to Modern and Intelligent Control Techniques. *IEEE Trans. Veh. Technol.* **2007**, *56*, 519–528. [CrossRef]
21. Cha, K.S.; Kim, D.M.; Jung, Y.H.; Lim, M.S. Wound field synchronous motor with hybrid circuit for neighborhood electric vehicle traction improving fuel economy. *Appl. Energy* **2020**, *263*, 114618. [CrossRef]
22. Ehsani, M.; Gao, Y.; Miller, J.M. Hybrid electric vehicles: Architecture and motor drives. *Proc. IEEE* **2007**, *95*, 719–728. [CrossRef]
23. Hannan, M.A.; Azidin, F.A.; Mohamed, A. Hybrid electric vehicles and their challenges: A review. *Renew. Sustain. Energy Rev.* **2014**, *29*, 135–150. [CrossRef]
24. Salman, W.; Zhang, X.; Li, H.; Wu, X.; Li, N.; Azam, A.; Zhang, Z. A novel energy regenerative shock absorber for in-wheel motors in electric vehicles. *Mech. Syst. Signal Process.* **2022**, *181*, 109488. [CrossRef]
25. Tie, S.F.; Wei, C. A review of energy sources and energy management system in electric vehicles. *Renew. Sustain. Energy Rev.* **2013**, *20*, 82–102. [CrossRef]
26. Cuma, M.U.; Koroglu, T. A comprehensive review on estimation strategies used in hybrid and battery electric vehicles. *Renew. Sustain. Energy Rev.* **2015**, *42*, 517–531. [CrossRef]
27. Sabri, M.F.M.; Danapalasingam, K.A.; Rahmat, M.F. A review on hybrid electric vehicles architecture and energy management strategies. *Renew. Sustain. Energy Rev.* **2016**, *53*, 1433–1442. [CrossRef]
28. Wang, J.; Wang, F.; Wang, G. Based Robust Finite Control Set Predictive Current Control for Induction Motor Systems with Time-Varying Disturbances. *IEEE Trans. Ind. Inform.* **2018**, *14*, 4159–4168. [CrossRef]
29. Podder, A.K.; Chakraborty, O.; Member, S. Control Strategies of Different Hybrid Energy Storage Systems for Electric Vehicles Applications. *IEEE Access* **2021**, *9*, 51865–51895. [CrossRef]
30. Lv, Z.; Member, S.; Qiao, L.; Member, G.S.; Cai, K.; Wang, Q. Big Data Analysis Technology for Electric Vehicle Networks in Smart Cities. *IEEE Trans. Intell. Transp. Syst.* **2021**, *22*, 1807–1816. [CrossRef]
31. Akhtar, M.A.; Saha, S. Positive Current Reference Generation based Current Control Technique for BLDC Motor Drives Applications. In *Proceedings of the 2019 5th International Conference on Advanced Computing and Communication Systems, ICACCS 2019*; IEEE: Piscataway, NJ, USA, 2019. [CrossRef]
32. Barroso, D.G.; Yang, Y.; Member, S.; Emadi, A. Electrified Automotive Propulsion Systems: State-of-the-Art Review. *IEEE Trans. Transp. Electrification* **2022**, *8*, 2898–2914. [CrossRef]
33. Gao, B.; Liang, Q.; Xiang, Y.; Guo, L.; Chen, H. Gear ratio optimization and shift control of 2-speed I-AMT in electric vehicle. *Mech. Syst. Signal Process.* **2015**, *50–51*, 615–631. [CrossRef]
34. Emadi, A.; Lee, Y.J.; Rajashekara, K. Power electronics and motor drives in electric, hybrid electric, and plug-in hybrid electric vehicles. *IEEE Trans. Ind. Electron.* **2008**, *55*, 2237–2245. [CrossRef]
35. Altun, Y.E.; Kutlar, O.A. Energy Management Systems' Modeling and Optimization in Hybrid Electric Vehicles. *Energies* **2024**, *17*, 1696. [CrossRef]
36. Chen, S.Y.; Wu, C.H.; Hung, Y.H.; Chung, C.T. Optimal strategies of energy management integrated with transmission control for a hybrid electric vehicle using dynamic particle swarm optimization. *Energy* **2018**, *160*, 154–170. [CrossRef]
37. Hu, X.; Member, S.; Han, J.; Tang, X.; Lin, X. Powertrain Design and Control in Electrified Vehicles: A Critical Review. *IEEE Transp. Electrification* **2021**, *7*, 1990–2009. [CrossRef]
38. Munir, M.F.; Ahmad, I.; Siffat, S.A.; Qureshi, M.A.; Armghan, H.; Ali, N. Non-linear control for electric power stage of fuel cell vehicles. *ISA Trans.* **2020**, *102*, 117–134. [CrossRef]
39. Saponara, S.; Lee, C.H.T.; Wang, N.X.; Kirtley, J.L. Electric Drives and Power Chargers: Recent Solutions to Improve Performance and Energy Efficiency for Hybrid and Fully Electric Vehicles. *IEEE Veh. Technol. Mag.* **2020**, *15*, 73–83. [CrossRef]
40. Williamson, S.S.; Emadi, A.; Rajashekara, K. Comprehensive Efficiency Modeling of Electric Traction Motor Drives for Hybrid Electric Vehicle Propulsion Applications. *IEEE Trans. Veh. Technol.* **2007**, *56*, 1561–1572. [CrossRef]
41. Ezugwu, E.O.; Bhattacharya, I.; Ayomide, A.I.; Dhason, M.V.A.; Soyoye, B.D.; Banik, T. Powertrain in Battery Electric Vehicles (BEVs): Comprehensive Review of Current Technologies and Future Trends Among Automakers. *World Electr. Veh. J.* **2025**, *16*, 573. [CrossRef]
42. Kim, D.M.; Lee, S.G.; Kim, D.K.; Park, M.R.; Lim, M.S. Sizing and optimization process of hybrid electric propulsion system for heavy-duty vehicle based on Gaussian process modeling considering traction motor characteristics. *Renew. Sustain. Energy Rev.* **2022**, *161*, 112286. [CrossRef]

43. Rind, S.J.; Ren, Y.; Hu, Y.; Wang, J.; Jiang, L. Configurations and Control of Traction Motors for Electric Vehicles: A Review. *Chin. J. Electr. Eng.* **2017**, *3*, 1–17. [CrossRef]
44. Mohammad, S.F.; Bakhsh, F.I.; Ibrahim, M.; Mumtaz, N.; Hameed, S. Detailed modelling and performance analysis of power flow topology in a hybrid electric vehicle having series-parallel architecture. *Renew. Energy Focus* **2024**, *49*, 100579. [CrossRef]
45. Qin, Y.; Tang, X.; Jia, T.; Duan, Z.; Zhang, J.; Li, Y.; Zheng, L. Noise and vibration suppression in hybrid electric vehicles: State of the art and challenges. *Renew. Sustain. Energy Rev.* **2020**, *124*, 109782. [CrossRef]
46. Moutafidis, I. *Architecture and Impacts of Electric Vehicles Architecture and Impacts of Electric Vehicles*; International Hellenic University: Thessaloniki, Greece, 2011.
47. Ieee, L.F.; Ieee, S.M.; Ieee, M. State of the Art and Trends in Electric and Hybrid Electric Vehicles. *Proc. IEEE* **2021**, *109*, 967–984. [CrossRef]
48. López, I.; Ibarra, E.; Matallana, A.; Andreu, J.; Kortabarria, I. Next generation electric drives for HEV/EV propulsion systems: Technology, trends and challenges. *Renew. Sustain. Energy Rev.* **2019**, *114*, 109336. [CrossRef]
49. El-refaie, A.; Osama, M. High Specific Power Electrical Machines: A System Perspective. *CES Trans. Electr. Mach. Syst.* **2019**, *3*, 88–93. [CrossRef]
50. Types of Hybrid Electric Vehicles (HEV)–X-Engineer.org. Available online: <https://x-engineer.org/types-hybrid-electric-vehicles-hev/> (accessed on 5 December 2025).
51. Vehicles, P.P.H.E.; Liu, T.; Hu, X.; Member, S.; Hu, W.; Member, S. A Heuristic Planning Reinforcement Learning-Based Energy Management for Power-Split Plug-in Hybrid Electric Vehicles. *IEEE Trans. Ind. Inform.* **2019**, *15*, 6436–6445. [CrossRef]
52. Wirasingha, S.G.; Emadi, A. Classification and review of control strategies for plug-in hybrid electric vehicles. *IEEE Trans. Veh. Technol.* **2011**, *60*, 111–122. [CrossRef]
53. Abdelkareem, M.A.; Maghrabie, H.M.; Abo-Khalil, A.G.; Adhari, O.H.K.; Sayed, E.T.; Radwan, A.; Rezk, H.; Jouhara, H.; Olabi, A. Thermal management systems based on heat pipes for batteries in EVs/HEVs. *J. Energy Storage* **2022**, *51*, 104384. [CrossRef]
54. Wirasingha, S.G.; Gremban, R.; Emadi, A.; Member, S. Source-to-Wheel (STW) Analysis of Plug-in Hybrid Electric Vehicles. *IEEE Transection Smart Grid* **2012**, *3*, 316–331. [CrossRef]
55. Al-Sahlawi, A.A.K.; Ayob, S.M.; Tan, C.W.; Ridha, H.M.; Hachim, D.M. Optimal Design of Grid-Connected Hybrid Renewable Energy System Considering Electric Vehicle Station Using Improved Multi-Objective Optimization: Techno-Economic Perspectives. *Sustainability* **2024**, *16*, 2491. [CrossRef]
56. Malozyomov, B.V.; Martyushev, N.V.; Kukartsev, V.V.; Konyukhov, V.Y.; Oparina, T.A.; Sevryugina, N.S.; Gozbenko, V.E.; Kondratiev, V.V. Determination of the Performance Characteristics of a Traction Battery in an Electric Vehicle. *World Electr. Veh. J.* **2024**, *15*, 64. [CrossRef]
57. Mwasilu, F.; Justo, J.J.; Kim, E.K.; Do, T.D.; Jung, J.W. Electric vehicles and smart grid interaction: A review on vehicle to grid and renewable energy sources integration. *Renew. Sustain. Energy Rev.* **2014**, *34*, 501–516. [CrossRef]
58. Manzetti, S.; Mariasiu, F. Electric vehicle battery technologies: From present state to future systems. *Renew. Sustain. Energy Rev.* **2015**, *51*, 1004–1012. [CrossRef]
59. He, H.; Han, M.; Liu, W.; Cao, J.; Shi, M.; Zhou, N. MPC-based longitudinal control strategy considering energy consumption for a dual-motor electric vehicle. *Energy* **2022**, *253*, 124004. [CrossRef]
60. Zhang, Y.; Ai, Z.; Chen, J. Energy-Saving Optimization and Control of Autonomous Electric Vehicles with Considering Multiconstraints. *IEEE Transection Cybern.* **2021**, *52*, 10869–10881. [CrossRef]
61. Types of Electric Vehicles: BEVs, PHEVs, HEVs-What's the Difference? Available online: <https://www.evgo.com/ev-drivers/types-of-evs/#bevs> (accessed on 22 October 2024).
62. 'Global Electric Vehicle Market Share Quarterly'. 8 December 2025. Available online: <https://counterpointresearch.com/en/insights/global-electric-vehicle-market-share-quarterly> (accessed on 11 February 2026).
63. Fathabadi, H. Novel fuel cell/battery/supercapacitor hybrid power source for fuel cell hybrid electric vehicles. *Energy* **2018**, *143*, 467–477. [CrossRef]
64. Liu, L.; Kong, F.; Liu, X.; Peng, Y.; Wang, Q. A review on electric vehicles interacting with renewable energy in smart grid. *Renew. Sustain. Energy Rev.* **2015**, *51*, 648–661. [CrossRef]
65. Sachs, C.; Neuburger, M. A data-based review of battery electric vehicle and traction inverter trends. In *Proceedings of the IECON 2025–51st Annual Conference of the IEEE Industrial Electronics Society*; IEEE: Piscataway, NJ, USA, 2025; pp. 1–8.
66. Mousavi, M.H.; Moradi, H.; Rouzbehi, K. Transition of DC Link Voltage from 400V to 800V in Electric Vehicles: Performance, Trade-Offs, and Technical Insights. In *Proceedings of the 2025 10th International Conference on Technology and Energy Management (ICTEM)*; IEEE: Piscataway, NJ, USA, 2025; pp. 1–7. [CrossRef]
67. Manoj, V.; Pilla, R.; Sura, S.R. A Comprehensive Analysis of Power Converter Topologies and Control Methods for Extremely Fast Charging of Electric Vehicles. *J. Phys. Conf. Ser.* **2023**, *2570*, 012017. [CrossRef]
68. Safayatullah, M.; Elrais, M.T.; Ghosh, S.; Rezaei, R.; Batarseh, I. A Comprehensive Review of Power Converter Topologies and Control Methods for Electric Vehicle Fast Charging Applications. *IEEE Access* **2022**, *10*, 40753–40793. [CrossRef]

69. Li, Y.; Yang, J.; Song, J. Design principles and energy system scale analysis technologies of new lithium-ion and aluminum-ion batteries for sustainable energy electric vehicles. *Renew. Sustain. Energy Rev.* **2017**, *71*, 645–651. [CrossRef]
70. Poullikkas, A. Sustainable options for electric vehicle technologies. *Renew. Sustain. Energy Rev.* **2015**, *41*, 1277–1287. [CrossRef]
71. Fathabadi, H. Combining a proton exchange membrane fuel cell (PEMFC) stack with a Li-ion battery to supply the power needs of a hybrid electric vehicle. *Renew. Energy* **2019**, *130*, 714–724. [CrossRef]
72. Kumar, L.; Jain, S. Electric propulsion system for electric vehicular technology: A review. *Renew. Sustain. Energy Rev.* **2014**, *29*, 924–940. [CrossRef]
73. Pan, G.; Bai, Y.; Song, H.; Qu, Y.; Wang, Y.; Wang, X. Hydrogen Fuel Cell Power System—Development Perspectives for Hybrid Topologies. *Energies* **2023**, *16*, 2680. [CrossRef]
74. Hames, Y.; Kaya, K.; Baltacioglu, E.; Turksoy, A. Analysis of the control strategies for fuel saving in the hydrogen fuel cell vehicles. *Int. J. Hydrog. Energy* **2018**, *43*, 10810–10821. [CrossRef]
75. Lukic, S.M.; Member, S. Topological Overview of Hybrid Electric and Fuel Cell Vehicular Power System Architectures and Configurations. *IEEE Trans. Veh. Technol.* **2005**, *54*, 763–770. [CrossRef]
76. Ihm, J.; Amghar, B.; Chun, S.; Park, H. Optimum Design of an Electric Vehicle Charging Station Using a Renewable Power Generation System in South Korea. *Sustainability* **2023**, *15*, 9931. [CrossRef]
77. Qayyum, N.; Khan, L.; Wahab, M.; Mumtaz, S.; Ali, N.; Khan, B.S. Fractional-Order Swarming Intelligence Heuristics for Nonlinear Sliding-Mode Control System Design in Fuel Cell Hybrid Electric Vehicles. *World Electr. Veh. J.* **2025**, *16*, 351. [CrossRef]
78. Cha, K.S.; Jung, Y.H.; Park, S.H.; Park, M.R. Optimal Design Considering AC Copper Loss of Traction Motor Applied HSFF Coil for Improving Electric Bus Fuel Economy. *Mathematics* **2025**, *13*, 1509. [CrossRef]
79. Zhang, Y.; Huang, X.; Liu, J.; Zhuge, C. The potential uptake and climate impacts of Hydrogen-Fuel-Cell vehicles in Beijing. *Transp. Res. D Transp. Environ.* **2026**, *154*, 105253. [CrossRef]
80. Chan, C.C. *The State of the Art of Electric and Hybrid Vehicles*; IEEE: Piscataway, NJ, USA, 2002; Volume 90, pp. 245–246.
81. Armenta-Déu, C. Improving Sustainability in Urban and Road Transportation: Dual Battery Block and Fuel Cell Hybrid Power System for Electric Vehicles. *Sustainability* **2024**, *16*, 2110. [CrossRef]
82. Lee, C.H.T.; Member, S.; Hua, W.E.I.; Member, S. A Critical Review of Emerging Technologies for Electric and Hybrid Vehicles. *IEEE Open J. Veh. Technol.* **2022**, *2*, 471–485. [CrossRef]
83. Yuvaraj, T.; Suresh, T.D.; Christy, A.A.; Babu, T.S.; Nastasi, B. Modelling and Allocation of Hydrogen-Fuel-Cell-Based Distributed Generation to Mitigate Electric Vehicle Charging Station Impact and Reliability Analysis on Electrical Distribution Systems. *Energies* **2023**, *16*, 6869. [CrossRef]
84. Enescu, F.M.; Birleanu, F.G.; Raboaca, M.S.; Raceanu, M.; Bizon, N.; Thounthong, P. Electric Vehicle Charging Station Based on Photovoltaic Energy with or without the Support of a Fuel Cell–Electrolyzer Unit. *Energies* **2023**, *16*, 762. [CrossRef]
85. Bashir, H.; Yaqoob, A.; Jawaid, I.; Khalid, W.; Javed, M.Y.; Sultan, W. A Review of Battery Management System and Modern State Estimation Approaches in Lithiumion Batteries for Electric Vehicle. In *Proceedings of the 2022 5th International Conference on Energy Conservation and Efficiency, ICECE 2022-Proceedings*; Institute of Electrical and Electronics Engineers Inc.: Piscataway, NJ, USA, 2022. [CrossRef]
86. Dannier, A.; Brando, G.; Ribera, M.; Spina, I. Li-Ion Batteries for Electric Vehicle Applications: An Overview of Accurate State of Charge/State of Health Estimation Methods. *Energies* **2025**, *18*, 786. [CrossRef]
87. Suganya, R.; Joseph, L.M.I.L.; Kollem, S. Understanding lithium-ion battery management systems in electric vehicles: Environmental and health impacts, comparative study, and future trends: A review. *Results Eng.* **2024**, *24*, 103047. [CrossRef]
88. Ali, M.U.; Zafar, A.; Nengroo, S.H.; Hussain, S.; Alvi, M.J.; Kim, H.J. Towards a smarter battery management system for electric vehicle applications: A critical review of lithium-ion battery state of charge estimation. *Energies* **2019**, *12*, 446. [CrossRef]
89. Aruchamy, K.; Ramasundaram, S.; Divya, S.; Chandran, M.; Yun, K.; Oh, T.H. Gel Polymer Electrolytes: Advancing Solid-State Batteries for High-Performance Applications. *Gels* **2023**, *9*, 585. [CrossRef]
90. Sang, V.T.D.; Duong, Q.H.; Zhou, L.; Arranz, C.F.A. Electric Vehicle Battery Technologies and Capacity Prediction: A Comprehensive Literature Review of Trends and Influencing Factors. *Batteries* **2024**, *10*, 451. [CrossRef]
91. Wang, Z.; Zhou, J.; Rizzoni, G. A review of architectures and control strategies of dual-motor coupling powertrain systems for battery electric vehicles. *Renew. Sustain. Energy Rev.* **2022**, *162*, 112455. [CrossRef]
92. Zhu, Z.Q.; Howe, D. Electrical machines and drives for electric, hybrid, and fuel cell vehicles. *Proc. IEEE* **2007**, *95*, 746–765. [CrossRef]
93. de Santiago, J.; Bernhoff, H.; Ekerghard, B.; Eriksson, S.; Ferhatovic, S.; Waters, R.; Leijon, M. Electrical motor drivelines in commercial all-electric vehicles: A review. *IEEE Trans. Veh. Technol.* **2012**, *61*, 475–484. [CrossRef]
94. Saidur, R.; Mekhilef, S.; Ali, M.B.; Safari, A.; Mohammed, H.A. Applications of variable speed drive (VSD) in electrical motors energy savings. *Renew. Sustain. Energy Rev.* **2012**, *16*, 543–550. [CrossRef]
95. Wang, H.; Yang, Y.; Chen, D.; Ge, X.; Li, S.; Zuo, Y. Speed-Sensorless Control of Induction Motors with an Open-Loop Synchronization Method. *IEEE J. Emerg. Sel. Top. Power Electron.* **2022**, *10*, 1963–1977. [CrossRef]

96. Mir, T.N.; Member, S.; Singh, B.; Bhat, A.H. FS-MPC-Based Speed Sensorless Control of Matrix Converter Fed Induction Motor Drive with. *IEEE Trans. Ind. Electron.* **2021**, *68*, 9185–9195. [CrossRef]
97. Zeraoulia, M.; Benbouzid, M.E.H.; Diallo, D. Electric motor drive selection issues for HEV propulsion systems: A comparative study. *IEEE Trans. Veh. Technol.* **2006**, *55*, 1756–1764. [CrossRef]
98. Salem, A.; Narimani, M. A Review on Multiphase Drives for Automotive Traction Applications. *IEEE Trans. Transp. Electrification* **2019**, *5*, 1329–1348. [CrossRef]
99. Ding, X.; Wang, Z.; Zhang, L.; Wang, C. Longitudinal Vehicle Speed Estimation for Four-Wheel-Independently-Actuated Electric Vehicles Based on Multi-Sensor Fusion. *IEEE Trans. Veh. Technol.* **2020**, *69*, 12797–12806. [CrossRef]
100. Odhano, S.A.; Pescetto, P.; Awan, H.A.A.; Hinkkanen, M.; Pellegrino, G.; Bojoi, R. Parameter Identification and Self-Commissioning in AC Motor Drives: A Technology Status Review. *IEEE Trans. Power Electron.* **2019**, *34*, 3603–3614. [CrossRef]
101. Kodkin, V.; Anikin, A.; Baldenkov, A. Optimization of Traction Electric Drive with Frequency Control. *World Electr. Veh. J.* **2025**, *16*, 139. [CrossRef]
102. Kano, Y.; Matsui, N. Rotor Geometry Design of Saliency-Based Sensorless Controlled Distributed-Winding IPMSM for Hybrid Electric Vehicles. *IEEE Trans. Ind. Appl.* **2018**, *54*, 2336–2348. [CrossRef]
103. Pellegrino, G.; Vagati, A.; Boazzo, B.; Guglielmi, P. Comparison of Induction and PM Synchronous Motor Drives for EV Application Including Design Examples. *IEEE Trans. Ind. Appl.* **2012**, *48*, 2322–2332. [CrossRef]
104. Sun, X.; Member, S.; Zhang, Y.; Tian, X.; Cao, J. Speed Sensorless Control for IPMSMs Using a Modified MRAS with Gray Wolf Optimization Algorithm. *IEEE Trans. Transp. Electrification* **2022**, *8*, 1326–1337. [CrossRef]
105. Ding, X.; Guo, H.; Xiong, R.; Chen, F.; Zhang, D.; Gerada, C. A new strategy of efficiency enhancement for traction systems in electric vehicles. *Appl. Energy* **2017**, *205*, 880–891. [CrossRef]
106. Widmer, J.D.; Martin, R.; Kimiabeigi, M. Electric vehicle traction motors without rare earth magnets. *Sustain. Mater. Technol.* **2015**, *3*, 7–13. [CrossRef]
107. Sawma, J.; Seferian, V.; Khatounian, F.; Monmasson, E.; Ghosn, R. An improved anti-rollback control algorithm for gearless traction motor in elevator applications. *Mechatronics* **2021**, *79*, 102659. [CrossRef]
108. Kommuri, S.K.; Defoort, M.; Karimi, H.R.; Veluvolu, K.C. A Robust Observer-Based Sensor Fault-Tolerant Control for PMSM in Electric Vehicles. *IEEE Trans. Ind. Electron.* **2016**, *63*, 7671–7681. [CrossRef]
109. Popsi, N.R.S.; Anik, A.; Verma, R.; Viana, C.; Iyer, K.L.V.; Kar, N.C. Influence of Electric Motor Manufacturing Tolerances on End-of-Line Testing: A Review. *Multidiscip. Digit. Publ. Inst.* **2024**, *17*, 1913. [CrossRef]
110. Sun, X.; Cao, J.; Lei, G.; Guo, Y.; Zhu, J. Speed Sensorless Control for Permanent Magnet Synchronous Motors Based on. *IEEE Trans. Ind. Electron.* **2020**, *67*, 6089–6100. [CrossRef]
111. Usman, A.; Saxena, A. Technical Roadmaps of Electric Motor Technology for Next Generation Electric Vehicles. *Multidiscip. Digit. Publ. Inst.* **2025**, *13*, 156. [CrossRef]
112. Jahns, T. *Getting Rare-Earth Magnets Out of EV Traction Machines*; IEEE: Piscataway, NJ, USA, 2017; pp. 6–18. [CrossRef]
113. Wang, Z.; Ching, T.Z.E.W.; Member, S.; Huang, S.; Wang, H.; Xu, T. Challenges Faced by Electric Vehicle Motors and Their Solutions. *IEEE Access* **2021**, *9*, 5228–5249. [CrossRef]
114. Sarlioglu, B.; Morris, C.T.; Han, D.; Li, S. Driving Toward Accessibility: A Review of Technological Improvements for Electric Machines, Power Electronics, and Batteries for Electric and Hybrid Vehicles. *IEEE Trans. Ind. Appl.* **2017**, *23*, 14–25. [CrossRef]
115. Tadeus, D.Y.; Winarno, H.; Sasmoko, P. Variable time delay switching to enable stepper control of brushless traction motor. *Mater. Today Proc.* **2022**, *63*, S255–S261. [CrossRef]
116. Konijeti, M.S.N.K.; Bharathi, M. Extraction of maximum power from solar with BLDC motor driven electric vehicles based HHO algorithm. *Adv. Eng. Softw.* **2022**, *170*, 103137. [CrossRef]
117. Tseng, C.Y.; Yu, C.H. Advanced shifting control of synchronizer mechanisms for clutchless automatic manual transmission in an electric vehicle. *Mech. Mach. Theory* **2015**, *84*, 37–56. [CrossRef]
118. Qian, Z.; Gong, S.; Xiao, S.; Lin, Z.; Li, X. Online Monitoring Method for Capacitor Lifetime in Brushless DC Motor Drive Systems with DC-Link Series Switch. *World Electr. Veh. J.* **2025**, *16*, 330. [CrossRef]
119. Zhou, D. Four-Quadrant Position Sensorless Operation of Switched Reluctance Machine for Electric Vehicles over a Wide Speed Range. *IEEE Trans. Transp. Electrification* **2021**, *7*, 2835–2847. [CrossRef]
120. Riba, J.R.; López-Torres, C.; Romeral, L.; Garcia, A. Rare-earth-free propulsion motors for electric vehicles: A technology review. *Renew. Sustain. Energy Rev.* **2016**, *57*, 367–379. [CrossRef]
121. Gan, C.; Wu, J. A Review on Machine Topologies and Control Techniques for Low-Noise Switched Reluctance Motors in Electric Vehicle Applications. *IEEE Access* **2018**, *6*, 31430–31443. [CrossRef]
122. Xiao, D.; Filho, S.R.; Fang, G.; Ye, J.; Member, S.; Emadi, A. Position-Sensorless Control of Switched Reluctance Motor Drives: A Review. *IEEE Trans. Transp. Electrification* **2022**, *8*, 1209–1227. [CrossRef]
123. Davarpanah, G.; Mohammadi, S. Connected C-core hybrid SRMs for EV applications. In *Proceedings of the 2025 IEEE International Electric Machines & Drives Conference (IEMDC)*; IEEE: Piscataway, NJ, USA, 2025; pp. 501–505.

124. Liu, C.; Chau, K.T.; Lee, C.H.T.; Song, Z. A Critical Review of Advanced Electric Machines and Control Strategies for Electric Vehicles. *Proc. IEEE* **2021**, *109*, 1004–1028. [CrossRef]
125. Benbouya, B.; Cheghib, H.; Chrenko, D.; Delgado, M.T.; Hamoudi, Y.; Rodriguez, J.; Abdelrahman, M. Sliding Mode Control of an Electric Vehicle Driven by a New Powertrain Technology Based on a Dual-Star Induction Machine. *World Electr. Veh. J.* **2024**, *15*, 155. [CrossRef]
126. Liu, H.C. Design, Analysis, and Comparison of Electric Vehicle Electric Oil Pump Motor Rotors Using Ferrite Magnet. *World Electr. Veh. J.* **2025**, *16*, 50. [CrossRef]
127. Kinoti, E.; Mosetlhe, T.C.; Yusuff, A.A. Multi-Criteria Analysis of Electric Vehicle Motor Technologies: A Review. *World Electr. Veh. J.* **2024**, *15*, 541. [CrossRef]
128. Shao, L.; Karci, A.E.H.; Tavernini, D.; Sornioti, A.; Cheng, M. Design Approaches and Control Strategies for Energy-Efficient Electric Machines for Electric Vehicles—A Review. *IEEE Access* **2020**, *8*, 116900–116913. [CrossRef]
129. Ramesh, P.; Lenin, N.C. High Power Density Electrical Machines for Electric Vehicles—Comprehensive Review Based on Material Technology. *IEEE Trans. Magn.* **2019**, *55*, 0900121. [CrossRef]
130. Wang, W.; Chen, X.; Wang, J.; Member, S. Motor / Generator Applications in Electrified Vehicle Chassis—A Survey. *IEEE Trans. Transp. Electr.* **2019**, *5*, 584–601. [CrossRef]
131. Loukas. Different Types of Electric Motors Used in EVs. Available online: https://www.arenaev.com/different_types_of_electric_motors_used_in_evs-news-214.php (accessed on 5 December 2025).
132. EV Database. 2024. Available online: <https://ev-database.org/compare/newest-upcoming-electric-vehicle#sort:path~type~order=.id~number~desc> (accessed on 10 October 2025).
133. Amin, M.; Member, S.; Aziz, G.A.A.; Durkin, J.; Member, S.; Mohammed, O.A. A Hardware-in-the-Loop Realization of Speed Sensorless Control of PMA-SynRM with Performances Enhancement. *IEEE Trans. Ind. Appl.* **2019**, *55*, 5331–5342. [CrossRef]
134. Poorfakhraei, A.; Narimani, M.; Emadi, A. A Review of Modulation and Control Techniques for Multilevel Inverters in Traction Applications. *IEEE Access* **2021**, *9*, 24187–24204. [CrossRef]
135. Haque, T.S.; Rahman, H.; Islam, R.; Razzak, A.; Badal, F.R.; Ahamed, H.; Moyeen, S.I.; Das, S.K.; Ali, F.; Tasneem, Z.; et al. A Review on Driving Control Issues for Smart Electric Vehicles. *IEEE Access* **2021**, *9*, 135440–135472. [CrossRef]
136. Hannan, M.A.; Ali, J.A.; Mohamed, A.; Hussain, A. Optimization techniques to enhance the performance of induction motor drives: A review. *Renew. Sustain. Energy Rev.* **2018**, *81*, 1611–1626. [CrossRef]
137. IUzhno-Ural'skii Gosudarstvennyi Universitet and Institute of Electrical and Electronics Engineers. *Proceedings of the 2017 International Conference on Industrial Engineering, Applications and Manufacturing (ICIEAM), Chelyabinsk, Russia, 16–19 May 2017*; IEEE: Piscataway, NJ, USA, 2017.
138. Taa, S.; Mokhtari, B. Proposal of an Optimal Control of an Electric Vehicle by Combined FOC and DTC Techniques. *Electrica* **2024**, *24*, 660–669. [CrossRef]
139. Zhang, X. Sensorless Induction Motor Drive Using Indirect Vector Controller and Sliding-Mode Observer for Electric Vehicles. *IEEE Trans. Veh. Technol.* **2013**, *62*, 3010–3018. [CrossRef]
140. Loop, P. Speed Sensorless Control of SPMSM Drives for EVs with a Binary Search Algorithm-Based. *IEEE Trans. Veh. Technol.* **2020**, *69*, 4968–4978. [CrossRef]
141. Oubelaid, A.; Taib, N.; Nikolovski, S.; Alharbi, T.E.A.; Rekioua, T.; Flah, A.; Ghoneim, S.S.M. Intelligent Speed Control and Performance Investigation of a Vector Controlled Electric Vehicle Considering Driving Cycles. *Electronics* **2022**, *11*, 1925. [CrossRef]
142. Mohanraj, D.; Gopalakrishnan, J. Critical Aspects of Electric Motor Drive Controllers and Mitigation of Torque Ripple—Review. *IEEE Access* **2022**, *10*, 73635–73674. [CrossRef]
143. Ali, S.M.N.; Hossain, M.J.; Wang, D.; Lu, K.; Rasmussen, P.O.; Sharma, V.; Kashif, M. Robust Sensorless Control Against Thermally Degraded Speed Performance in an IM Drive Based Electric Vehicle. *IEEE Trans. Energy Convers.* **2020**, *35*, 896–907. [CrossRef]
144. Mapelli, F.L.; Tarsitano, D.; Cheli, F. MRAS rotor resistance estimators for EV vector controlled induction motor traction drive: Analysis and experimental results. *Electr. Power Syst. Res.* **2017**, *146*, 298–307. [CrossRef]
145. De Klerk, M.L.; Saha, A.K. A Comprehensive Review of Advanced Traction Motor Control Techniques Suitable for Electric Vehicle Applications. *IEEE Access* **2021**, *9*, 125080–125108. [CrossRef]
146. Tazerart, F.; Mokrani, Z.; Rekioua, D.; Rekioua, T. ScienceDirect Direct torque control implementation with losses minimization of induction motor for electric vehicle applications with high operating life of the battery. *Int. J. Hydrogen Energy* **2015**, *40*, 13827–13838. [CrossRef]
147. Bose, B.K. *Modern Power Electronics and AC Drives*; Phi Learning PP: New Delhi, India, 2012.
148. Lee, J.; Kim, J.; Park, B. Fast Anti-Slip Traction Control for Electric Vehicles Based on Direct Torque Control with Load Torque Observer of Traction Motor. In *Transportation Electrification Conference & Expo*; IEEE: Piscataway, NJ, USA, 2021; pp. 321–326. [CrossRef]
149. Sutikno, T.; Idris, N.R.N.; Jidin, A. A review of direct torque control of induction motors for sustainable reliability and energy efficient drives. *Renew. Sustain. Energy Rev.* **2014**, *32*, 548–558. [CrossRef]

150. Aktas, M.; Awaili, K.; Ehsani, M.; Arisoy, A. Direct torque control versus indirect field-oriented control of induction motors for electric vehicle applications. *Eng. Sci. Technol. Int. J.* **2020**, *23*, 1134–1143. [CrossRef]
151. Salem, F.B.; Almousa, M.T.; Derbel, N. Direct Torque Control with Space Vector Modulation (DTC-SVM) with Adaptive Fractional-Order Sliding Mode: A Path Towards Improved Electric Vehicle Propulsion. *World Electr. Veh. J.* **2024**, *15*, 563. [CrossRef]
152. Chinthakunta, U.R.; Prabhakar, K.K.; Singh, A.K.; Kumar, P. Direct torque control induction motor drive performance evaluation based on torque error status selection methods. *IET Electr. Syst. Transp.* **2019**, *9*, 113–127. [CrossRef]
153. Lekshmi, S.; Lal Priya, P.S. Hierarchical predictive optimal control for range extension of EV with ANN based torque control for IPMSM drives. *e-Prime-Adv. Electr. Eng. Electron. Energy* **2024**, *10*, 100772. [CrossRef]
154. Sangar, B.; Singh, M.; Sreejeth, M. An improved ANFIS model predictive current control approach for minimizing torque and current ripples in PMSM-driven electric vehicle. *Electr. Eng.* **2024**, *106*, 5897–5907. [CrossRef]
155. Wang, S.; Zheng, Z.; Zhao, X.; Li, Z.; Tian, J. Hybrid model predictive control-based integration of handling stability control for distributed drive electric vehicles. *J. Traffic Transp. Eng. (Engl. Ed.)* **2025**, *12*, 795–811. [CrossRef]
156. Hassan, A.M.; Metwally, M.E. Improving performance of electric vehicle drive system based a five-phase PMSM under fault using ANN and MPC. *Sci. Rep.* **2025**, *15*, 42943. [CrossRef]
157. Banda, G.; Kolli, S.G. An intelligent adaptive neural network controller for a direct torque controlled ecar propulsion system. *World Electr. Veh. J.* **2021**, *12*, 44. [CrossRef]
158. Saleeb, H.; Kassem, R.; Sayed, K. Artificial neural networks applied on induction motor drive for an electric vehicle propulsion system. *Electr. Eng.* **2022**, *104*, 1769–1780. [CrossRef]
159. Fadil, N.D.B.M.; Zaidi, A.F.A.; Leong, J.H.; Azalan, M.S.Z.; Azmi, S.A.; Wahab, S.P.A. Assessing Torque-Ripple Mitigation Strategies for BLDC Motors in Electric Vehicles. In *Proceedings of the 2025 9th International Conference on Man-Machine Systems (ICoMMS)*; IEEE: Piscataway, NJ, USA, 2025; pp. 216–221. [CrossRef]
160. Joshi, G.; Pius, A.P. ANFIS controller for vector control of three phase induction motor. *Indones. J. Electr. Eng. Comput. Sci.* **2020**, *19*, 1177–1185. [CrossRef]
161. Hannan, M.A.; Ali, J.A.; Ker, P.J.; Mohamed, A.; Lipu, M.S.H.; Hussain, A. Switching techniques and intelligent controllers for induction motor drive: Issues and recommendations. *IEEE Access* **2018**, *6*, 47489–47510. [CrossRef]
162. Araria, R.; Negadi, K.; Marignetti, F. Design and Analysis of the Speed and Torque Control of IM with DTC Based ANN Strategy for Electric Vehicle Application. *Tec. Ital.-Ital. J. Eng. Sci.* **2019**, *63*, 181–188. [CrossRef]
163. Fatemimoghdam, A.; Yan, Y.; Iyer, L.V.; Kar, N.C. Permanent Magnet Synchronous Motor Drive Using Deep-Neural-Network-Based Vector Control for Electric Vehicle Applications. In *Proceedings of the 2022 International Conference on Electrical Machines, ICEM 2022*; Institute of Electrical and Electronics Engineers Inc.: Piscataway, NJ, USA, 2022; pp. 2358–2364. [CrossRef]
164. Suhail, M.; Akhtar, I.; Kirmani, S.; Jameel, M. Development of Progressive Fuzzy Logic and ANFIS Control for Energy Management of Plug-In Hybrid Electric Vehicle. *IEEE Access* **2021**, *9*, 62219–62231. [CrossRef]
165. Islam, A.; Singh, J.A.I.G.; Member, S.; Jahan, I.; Lipu, M.S.H.; Jamal, T. Modeling and Performance Evaluation of ANFIS Controller-Based Bidirectional Power Management Scheme in Plug-In Electric Vehicles Integrated with Electric Grid. *IEEE Access* **2021**, *9*, 166762–166780. [CrossRef]
166. George, M.A.; Kamat, D.V.; Kurian, C.P. Electric vehicle speed tracking control using an ANFIS-based fractional order PID controller. *J. King Saud Univ.-Eng. Sci.* **2022**, *36*, 256–264. [CrossRef]
167. Subbarao, M.; Dasari, K.; Duvvuri, S.S.; Prasad, K.R.K.V.; Narendra, B.K.; Krishna, V.B.M. Design, control and performance comparison of PI and ANFIS controllers for BLDC motor driven electric vehicles. *Meas. Sens.* **2024**, *31*, 101001. [CrossRef]
168. Intidam, A.; El Fadil, H.; Housny, H.; El Idrissi, Z.; Lassioui, A.; Nady, S.; Laafou, A.J. Development and Experimental Implementation of Optimized PI-ANFIS Controller for Speed Control of a Brushless DC Motor in Fuel Cell Electric Vehicles. *Energies* **2023**, *16*, 4395. [CrossRef]
169. Yin, H.; Zhou, W.; Li, M.; Ma, C.; Zhao, C. An adaptive fuzzy logic-based energy management strategy on battery/ultracapacitor hybrid electric vehicles. *IEEE Trans. Transp. Electr.* **2016**, *2*, 300–311. [CrossRef]
170. Paulo, S.; Aktaş, T.Ö.M. Research on Control Strategy and Energy Consumption for Electric Vehicles. In *Proceedings of the 11th IFAC Workshop on Intelligent Manufacturing Systems*; The International Federation of Automatic Control: Sao Paulo, Brazil, 2013; pp. 444–449. [CrossRef]
171. Shenbagalakshmi, R.; Mittal, S.K.; Subramaniyan, J.; Vengatesan, V.; Manikandan, D.; Ramaswamy, K. Adaptive speed control of BLDC motors for enhanced electric vehicle performance using fuzzy logic. *Sci. Rep.* **2025**, *15*, 12579. [CrossRef]
172. Kassem, R.; Sayed, K.; Kassem, A.; Mostafa, R. Power optimisation scheme of induction motor using FLC for electric vehicle. *IET Electr. Syst. Transp.* **2020**, *10*, 275–284. [CrossRef]
173. Guler, N.; Ismail, Z.M.; Hazem, Z.B.; Naik, N. Adaptive Fuzzy Logic Controller-Based Intelligent Energy Management System Scheme for Hybrid Electric Vehicles. *IEEE Access* **2024**, *12*, 173441–173454. [CrossRef]

174. Tian, H.; Wang, X.; Lu, Z.; Huang, Y.; Tian, G. Adaptive Fuzzy Logic Energy Management Strategy Based on Reasonable SOC Reference Curve for Online Control of Plug-in Hybrid Electric City Bus. In *IEEE Transactions on Intelligent Transportation Systems*; Institute of Electrical and Electronics Engineers Inc.: Piscataway, NJ, USA, 2018; pp. 1607–1617. [CrossRef]
175. Hou, S.; Chen, H.; Liu, X.; Cui, J.; Zhao, J.; Gao, J. Hierarchical model predictive control for energy management and lifespan protection in fuel cell electric vehicles. *Energy* **2025**, *319*, 134968. [CrossRef]
176. Guo, Z.; Chen, H.; Xu, F.; Kong, X.; Lin, J. Learning-based Model Predictive Control for Four-Wheel Drive Electric Vehicle Stability under Environmental Disturbance. *IFAC-PapersOnLine* **2024**, *58*, 421–426. [CrossRef]
177. Kasri, A.; Ouari, K.; Belkhier, Y.; Bajaj, M.; Zaitsev, I. Optimizing electric vehicle powertrains peak performance with robust predictive direct torque control of induction motors: A practical approach and experimental validation. *Sci. Rep.* **2024**, *14*, 14977. [CrossRef]
178. Li, Y.; Zhao, C.; Zhou, Y.; Qin, Y. Model predictive torque control of PMSM based on data drive. *Energy Rep.* **2020**, *6*, 1370–1376. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Review

A Scoping Review of Energy Consumption in Industrial Robotics

Johannes Muru * and Anton Rassõlkin *

Department of Electrical Power Engineering and Mechatronics, Tallinn University of Technology,
19086 Tallinn, Estonia

* Correspondence: johannes.muru@taltech.ee (J.M.); anton.rassolkin@taltech.ee (A.R.)

Abstract

The increasing adoption of industrial robots has significantly advanced manufacturing efficiency and flexibility. However, this expansion introduces new energy consumption challenges, especially as electricity has become the dominant energy source in automated systems. As the industrial sector faces rising energy costs and ambitious sustainability goals, understanding and minimizing the energy consumption of robotic systems is imperative. This review presents a structured analysis of energy consumption in industrial robots, linking mechanical design, actuation systems, and control strategies to their energetic effects. We first discuss different industrial robot types and their kinematic configurations, identifying how structural characteristics influence energy use. The article then categorizes energy consumption optimization strategies into software-based and hardware-based approaches. A comparative SWOT analysis highlights the strengths and limitations of each approach. The review also explores emerging trends such as DC microgrid integration. The future directions underline the need for standardized energy assessment frameworks and the development of hybrid optimization strategies that combine the reviewed approaches, suitable for being applied in real-world industrial robot applications. This work provides a comprehensive foundation for establishing best practices in energy consumption optimization for industrial robots.

Keywords: automation; robotics; industrial robotics; energy consumption; energy optimization

1. Introduction

The adoption of Industry 4.0 technologies has driven a rapid increase in the use of industrial robots (IRs) over the past decade. According to the International Federation of Robotics (IFR), more than 4 million IRs were installed globally by 2024—a 10% increase from 2023 and nearly three times the amount in 2014 [1]. This surge reflects the growing demand for automation to improve productivity, quality, and operational flexibility across manufacturing sectors.

ISO 8373:2021 defines an IR as “an automatically controlled, reprogrammable multi-purpose manipulator, programmable in three or more axes, which can be either fixed in place or fixed to a mobile platform for use in automation applications in an industrial environment” [2]. This definition encompasses not only the manipulator but also its actuators, the robot controller, including all associated hardware, software communication interfaces, and the teach pendant. The term reprogrammable indicates that the robot’s motion instructions can be modified through software without any physical changes to the system, while multipurpose underscores the robot’s versatility across a range of applications.

Energy efficiency has become a central priority for industrial manufacturing, particularly in light of rising energy costs and climate goals. The European Union (EU) has

committed to reducing overall energy consumption by at least 11.7% by 2030 [3]. In 2022, the industrial sector was the third-largest energy consumer in the EU, accounting for 25.1% of final energy use, with electricity comprising the largest share at 33% [4]. Improving energy performance in industrial operations is therefore crucial not only for meeting sustainability targets but also for reducing manufacturing costs and maintaining competitiveness.

Recognizing this challenge, the IFR has identified sustainability and energy efficiency as key trends shaping the future of robotics [5]. When properly integrated, robotic systems can enhance production efficiency and resource utilization [6]. Robots often execute tasks with greater efficiency and consistency, which enables increased production throughput, ultimately contributing to improved energy efficiency [7,8]. Additionally, as noted in [9], fully automated facilities can implement “lights-out” manufacturing, where operations proceed without human intervention. This approach further reduces energy usage by eliminating non-essential energy consumers such as lighting and other support systems. Recent studies demonstrate that the application of IRs plays a significant role in optimizing energy consumption within the entire manufacturing company [10]. This impact underscores the importance of integrating robots in the evolving industrial landscape and modern energy transition strategies.

However, the widespread deployment of IRs may increase overall energy demand. IRs are complex electromechanical systems that consume significant energy during operation. They require power to drive actuators, control electronics, auxiliary components, and end effectors. Numerous studies have explored energy consumption optimization (ECO) in industrial robotics from various angles, including dynamic modeling, trajectory planning, control strategy optimization, and hardware-level improvements, as stated in recent topical reviews [11–14]. Relatively few works have demonstrated how their findings are implemented in real-world IR applications within manufacturing settings like in [15]. While ECO in IRs has been studied from various perspectives, including various control algorithms, trajectory planning, and hardware enhancements, relatively little attention has gone towards how the kinematic structure of IRs, together with its controllable aspects, influences energy usage. This constitutes a crucial gap in the literature, especially considering the diverse kinematic configurations in use today. This review aims to bridge that gap.

A literature search was performed to conduct this scoping review using the databases Google Scholar, IEEE Xplore, and Scopus. The search strategy involved combinations of keywords such as “industrial robot”, “energy optimization”, “energy consumption”, “energy-efficient control”, “energy-efficient trajectory”, and “kinematic structure”. The search included peer-reviewed journal articles, conference papers, and preprints, published in the English language. Studies were considered eligible if they addressed energy consumption or energy optimization in the context of IRs, with a focus on either control strategies, mechanical design, or system-level considerations. Papers that focused solely on mobile robots, legged robots, or service robots were excluded. The selection followed the PRISMA 2020 guidelines [16], specifically the extension for scoping review (PRISMA-ScR) [17]. Titles and abstracts were screened for relevance, followed by full-text review. Articles lacking technical depth or applicability to IRs were removed.

The remainder of this paper is organized as follows: Section 2 presents an overview of IR types and kinematic structures, linking mechanical design with energetic characteristics. Section 3 presents a comprehensive review of ECO strategies, encompassing both software-based approaches, such as parameter tuning, task scheduling, and trajectory planning, and hardware-based methods, including structural design, energy recovery systems, and the application of functional redundancy. Additionally, a SWOT analysis is provided

to compare the various software- and hardware-based approaches. This is followed by a discussion and suggestions for future research directions in Section 4.

2. Background: Characteristics of Industrial Robots

2.1. Definition and Types of Industrial Robots

According to IFR, IRs can be categorized based on their mechanical structure into the following types [18]:

- Articulated robot;
- SCARA robot;
- Cartesian robot;
- Parallel/Delta robot;
- Cylindrical robot;
- Polar robot.

Articulated robots are among the most widely deployed robot types in industrial settings, primarily due to their high versatility [19]. This is largely enabled by their kinematic structure, which allows for functional redundancy, and a wide range of motion relative to their physical size. Articulated robots are built from several interconnected links, which are connected through revolute joints (Figure 1). Each joint provides the robot one degree of freedom (DOF). The most common configurations of articulated robots are 4-DOF to 6-DOF. Articulated robots are adaptable to various automation tasks that require complex motions. However, complex configuration means more complex programming, as well as higher costs and more frequent maintenance for the robot.




Figure 1. (a) ABB IRB 1600 industrial robot [20], (b) kinematic diagram of articulated robot.

A SCARA (Selective Compliance Assembly Robot Arm) robot usually consists of three revolute and one prismatic joint (Figure 2). SCARA robots are mainly designed for

tasks requiring high-speed, precise, and repeatable movements within a horizontal plane. The term “selective compliance” comes from the characteristic compliance of the joints in the X-Y direction, which makes the robot best suited for assembly applications; yet they are extensively used in material handling and pick-and-place tasks as well. Thanks to their compact size, simple design, and cost effectiveness, SCARA robots excel in environments with limited space and high throughput demands.




Figure 2. (a) KUKA SCARA KR 6 R500 Z200-2 [21], (b) kinematic diagram of SCARA robot.

A Cartesian robot, also known as a linear robot or gantry robot, operates on three prismatic joints perpendicular to each other, forming a rectangular coordinate system (Figure 3). This simple design provides high precision and repeatability, making Cartesian robots ideal for applications such as CNC machining and 3D printing, as well as simple pick-and-place tasks. Their structure allows for scalability, enabling larger workspaces and higher payload capacities compared to other robot types. Cartesian robots are widely used in manufacturing, logistics, and laboratory automation due to their straightforward programming and efficiency in tasks requiring linear movements.




Figure 3. (a) Festo YXCR gantry robot [22], (b) kinematic diagram of Cartesian robot.

A Delta robot is a high-speed parallel robot with a lightweight structure. Most commonly, the robot comprises three arms with a parallelogram structure connected to revolute joints at the base allowing three DOFs (Figure 4). Some configurations also use prismatic joints instead of the revolute joints in the form of linear axes to drive the arms [23].

The structural characteristics that make Delta robots effective in industrial applications include symmetry, distributed actuation, and rigid constraint-based motion. These principles can be more generally described through the framework of parallel mechanism synthesis [24,25], which allows for the systematic exploration of alternative limb configura-

tions and constraint arrangements that achieve similar or expanded motion capabilities. The Delta robot is typically implemented using a 3-RRR configuration (Figure 4b). However, other topologies, such as the 3-PRR configuration shown in Figure 4d, belong to the same family of parallel mechanisms.

These unique designs keep the orientation of the end effector unchanged, and the structure provides high stability, high-speed operation, and low inertia. In most cases, the end plate is equipped with an additional revolute joint to allow rotational movement for the end effector. The robot excels in performing precise and rapid movements and is often used in tasks requiring quick pick-and-place operations, sorting, and packaging in industries such as food processing, pharmaceuticals, and electronics.




Figure 4. (a) Igus Delta LCA-DR1000-04 robot [26], (b) kinematic diagram of delta robot with rotational links, (c) Festo EXPT-45-E1 robot [27], (d) kinematic diagram of Delta robot with linear links.

Cylindrical and polar robots are generally regarded as obsolete in contemporary industrial applications due to their limited usage. A modern alternative to cylindrical robots is the 4-axis articulated robot, which is widely employed in industry today, particularly for palletizing tasks [28]. Polar robots share a similar workspace configuration with higher DOF articulated robots, offering comparable reach and motion capabilities, although with reduced flexibility and precision.

The kinematic design of an IR plays an important role in its energy consumption. Robots with serial kinematics, such as articulated robots, often experience cumulative inertial effects along their kinematic chain. Each additional joint must support the mass and movement of all subsequent joints, which leads to higher torque requirements, especially

in the base joints. This increases the demand for more powerful motors and leads to higher energy consumption during dynamic operations. Strategies such as robot positioning to minimize gravitational torque are applicable here to minimize the consumed energy, especially during standstill [29,30].

Another notable characteristic of articulated manipulators is kinematic redundancy, which provides additional degrees of freedom beyond those minimally required for a given task. This redundancy enhances the robot's flexibility, as the system has multiple possible configurations to reach the same end-effector pose. However, this also introduces challenges in energy-efficient control, as not all configurations are energetically equivalent. Choosing a posture that minimizes joint torques becomes important for reducing energy consumption, while intelligent path-planning algorithms can exploit the redundancy to optimize not only for time and collision avoidance but also for energy [31,32]. Hence, the presence of redundancy affects the control system as well as the joint and link design that can efficiently execute a range of configurations.

In parallel robot architectures, the payload is supported directly by all actuators, which can typically be positioned on or near the robot base. This configuration allows the connecting links between the base and the mobile platform to be significantly lighter, resulting in a substantially higher payload-to-robot-mass ratio [33]. In terms of accuracy and repeatability, parallel robots offer clear advantages over serial counterparts. In serial manipulators, errors such as backlash, friction, and joint compliance accumulate along the kinematic chain, with greater influence from joints located closer to the base. Parallel robots avoid this cascading error effect and exhibit high structural rigidity, even with lightweight components. Additionally, due to their favorable payload-to-mass ratio and reduced joint coupling, parallel robots demonstrate superior dynamic performance compared to serial configurations [34]. Some recent studies have also investigated energy minimization strategies for Delta robots [35,36].

The researchers in [37] showed that parallel manipulators tend to be on average more energy efficient than serial manipulators with similar work capacity. The energetic characteristic of parallel robots is mostly affected by the actuator placement and load distribution. In many parallel configurations, the heavy actuators are fixed to the base, which significantly reduces the moving mass [38]. This, in turn, leads to lower inertia [39], meaning less energy is required for acceleration and deceleration. Furthermore, the symmetry and rigidity of the structure allow for efficient force transmission, and the closed-loop kinematics enable multiple limbs to share the load, reducing peak actuator demand [40].

Variations in kinematic structure directly affect component selection, especially regarding motors, gearboxes, and power electronics. For example, a robot designed for high-speed operations with short reach and low payload may utilize low-inertia motors and harmonic drives to achieve a desired efficiency [41,42]. On the other hand, a long-reach articulated robot operating under high payload demands requires high-torque motors and robust gearboxes, which may increase both energy consumption and system complexity [42]. These design decisions represent a balance between functional versatility and energy cost, emphasizing the need for an integrated co-design approach that merges mechanical architecture with energy-aware component selection. Motor sizing and gear ratio choices must accommodate worst-case torque demands, which in turn contribute substantially to the overall energy budget of the system.

In high-precision applications such as welding, assembly, or machining, energy consumption is often secondary to motion accuracy and repeatability. In these cases, trajectory smoothness and tool-path precision are prioritized, sometimes at the expense of optimal energy use. With optimization techniques such as trajectory planning or workstation re-

design, studies have shown that the possible energy savings are modest, typically in the range of 10% [43,44].

Applications involving high-speed, repetitive motions, and larger movements—such as sorting, palletizing, and packaging—present greater opportunities for energy consumption optimization. These tasks typically involve predictable motion patterns and shorter cycle times, making them ideal for various ECO strategies. Generally, reported savings in such use cases are substantially higher—up to 30% in palletizing [28], 45% in pick-and-place, and 25% in bin-picking scenarios [45], and in moderate-complexity operations such as drilling, energy reductions approaching 18% have been achieved when optimization strategies are applied [45].

Table 1 summarizes typical application areas in relation to different kinematic configurations of IRs. It is important to note that the listed use cases refer to the kinds of applications that can be realized with a given kinematic structure—not to the specific systems themselves. The ECO priority levels presented in the table offer a reasoned estimation of how critical or suitable energy optimization strategies are for each use case. However, these assessments are not yet universally validated. The literature in this domain remains limited, often relying on isolated case studies.

Table 1. Comparison of industrial robot configurations by degrees of freedom, typical application areas, and relative energy optimization potential. The energy optimization priority reflects how suitable or critical energy consumption strategies are for each robot type and task context.

DOF/Motors	Robot Type	Typical Use Cases	ECO Priority
3	Cartesian	CNC machining, 3D printing	Low
4	SCARA, Delta, Articulated	High-speed assembly, sorting, packaging, pick-and-place, palletizing	High
5	Articulated	Pick and place, palletizing	High
6	Articulated	Welding, painting, machining, assembly	Low - Medium
7	Articulated	Human-robot collaboration, flexible tasks	Medium

2.2. Fundamentals of Energy Consumption in Industrial Robots

IRs are complex systems composed of various interconnected components, each having a specific function in the robotic system. These components—such as motors, transmissions, controllers, sensors, and end effectors—work together to make the robot move with high accuracy and efficiency as is desired in a manufacturing environment. The total energy consumption of an IR is not determined by any single component but rather by the interaction between all of them; however, the drive system, which is responsible for robot motion, is the primary contributor to the overall energy consumption [46].

Traditionally, the actuators have been powered by three primary sources: electrical, pneumatic, and hydraulic systems [47,48]. Hydraulic actuators were preferred for IRs during the early stages of robotics development [49,50] due to their high power-to-weight ratio, robustness, and ability to handle large payloads with rapid dynamics and precise

positioning. These systems operate by converting pressurized fluid energy into mechanical force, typically using oil as the working fluid. Hydraulic robots are especially prevalent in heavy machinery and construction equipment, where high strength and quick response are essential. However, they tend to be complex, nonlinear, and require careful control design to manage their dynamics effectively [51].

Pneumatic actuators, sharing a similar working principle as hydraulic systems, are often favored in automation environments requiring clean and dry conditions, focusing primarily on the gripping mechanisms of IRs [52,53]. Pneumatic systems are generally lighter and simpler than hydraulic ones but face challenges in precise control due to the air compressibility and hysteresis effects inherent in the actuators [47]. However, the power delivery in pneumatic systems is continuous via compressors, and the systems in general tend to have low efficiencies [54]. Energy consumption in Europe of pneumatic systems accounts for around 10% of total industrial power consumption [55].

Electric actuation systems in IRs convert electrical energy directly into mechanical motion, typically through servo motors coupled with various transmission mechanisms. Electric actuation is by far the most common source used in IRs [47]. They offer advantages in the smoothness of motion, precise velocity and acceleration control, and energy efficiency. Electric actuators are widely used in applications requiring high precision and repeatability. Compared to fluid power systems, electric actuators generally have a smaller footprint and produce less noise, making them suitable for a broad range of industrial tasks [51].

In electrically powered robots, energy is supplied primarily through an AC grid connection, which is internally converted to a DC supply via power electronics. This DC bus powers the robot's servo drives, which in turn control brushless AC servo motors at each joint [56]. These motors are responsible for executing precise motions, with varying torque and velocity demands depending on the task dynamics and the robot's kinematic structure. In addition to the actuator system, electrical energy also supports the robot's control system, including the central processing unit, safety modules, sensors, human-machine interfaces (HMIs), and cooling mechanisms [56]. While the control-side consumption remains relatively constant, the majority of electrical energy is expended in driving the manipulator joints [46], where the efficiency is influenced by factors such as joint kinematics, mechanical transmission efficiency, motion profiles, and dynamic loads [29]. Therefore, analyzing and optimizing electrical power consumption at the manipulator level is crucial for developing energy-efficient robotic systems.

One of the emerging trends in industrial automation is the integration of DC microgrids with robotic systems, aimed at improving energy efficiency, system modularity, and renewable energy compatibility [57–59]. Unlike traditional AC-based distribution, DC microgrids offer several advantages in robotic applications, including reduced conversion losses, simplified power electronics, and better compatibility with onboard energy storage and regenerative braking systems [60]. IRs, particularly those operating in multi-robot cells or high-speed repetitive tasks, can generate significant amounts of regenerative energy during deceleration phases. When integrated into a shared DC microgrid, this energy can be redistributed to other robots or stored locally for later use, minimizing dissipation through braking resistors [60,61]. Moreover, DC microgrids facilitate the seamless integration of renewable sources [57], capacitive storage [60], and intelligent energy management systems [61]. These solutions shift toward a holistic energy management system in manufacturing environments, where IRs function not only as energy consumers but also as active participants in decentralized, intelligent energy networks.

Figure 5 inspired by [56] shows a schematic of the power distribution within an IR system. A standard 6-DOF articulated IR is actuated by a multi-drive servo system including six servo motors in each joint and the related power converters, composed of

a rectifier (AC/DC), DC-bus, and inverters (DC/AC), which are usually referred to as servo drives.




Figure 5. Power distribution from the grid to the subsystems of a 6-DOF industrial. Incoming AC power (P_{ac}) is rectified by an AC/DC converter to generate DC-bus power (P_{dc}), which supplies the servo drives (P_{sd}) and motor power (P_m). Load-independent consumption is divided between control system components (P_{cab})—including electronics, pendant, and cooling—and constant power draws from the rectifier ($P_{rec, const}$) and drives ($P_{sd, const}$). The DC-bus is stabilized by a capacitor (C_{dc}), while a brake resistor (R_{dc}) dissipates excess regenerative energy. The regenerated power P_{regen} is fed back to the DC-bus. Six servo drives (SD_1 – SD_6) independently control actuators (M_1 – M_6), each paired with a mechanical brake (br_1 – br_6). The mechanical power output (P_{mech}) is delivered through the motor shafts to move the robot joints P_{jnt} and payload P_{pyl} . Energy losses (*loss*) occur at each conversion stage and are annotated with dashed lines. The diagram highlights the distinction between the control system (red) and the mechanical manipulator (blue).

The overall power fed from the main grid P_{ac} is converted at the rectifier P_{dc} and divided between the power delivered to the servo drives P_{sd} , to the control cabinet P_{cab} and to the electromechanical brakes P_{br} . P_{sd} contributes to the motors P_m . Each conversion stage is associated with relevant electric power losses ($P_{rec, loss}$, $P_{sd, loss}$, $P_{m, loss}$). Mechanical power P_{mech} is delivered via motor shafts to the robot's mechanical structure to move the joints

(P_{jnt}) and the payload (P_{pyl}). Mechanical power loss $P_{mech,loss}$ also exists due to mechanical friction within the manipulator.

The term P_{cab} refers to the load-independent power required to operate the robot's control infrastructure, including the controller, teach pendant, cooling systems, and other auxiliary electronics. This also includes the constant power requirements of internal control components, such as the rectifier ($P_{rec,const}$) and the servo drives ($P_{sd,const}$). Losses ($P_{cab,loss}$) also exist here. The DC-bus capacitance (C_{dc}) functions as an energy buffer during regenerative braking events, storing the recuperated energy. However, due to the limited storage capacity of the capacitor, surplus energy that cannot be buffered is dissipated through the DC-bus balancing resistor (R_{dc}) to prevent over-voltage [60].

The relationship between the joint torques of an industrial manipulator and the changes in its joint positions can be characterized through inverse dynamics computations, which can be represented by the function

$$\tau = f(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}}), \quad (1)$$

where $\tau \in \mathbb{R}^n$ (n —number of joints) is the vector of joint torques, $\mathbf{q} \in \mathbb{R}^n$ is a vector of joint positions, $\dot{\mathbf{q}} \in \mathbb{R}^n$ is a vector of joint velocities, $\ddot{\mathbf{q}} \in \mathbb{R}^n$ is a vector of joint accelerations.

The dynamic behavior of serial manipulators with revolute joints in joint space is commonly described by a generalized nonlinear differential equation [62]

$$\tau = \mathbf{M}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} + \mathbf{G}(\mathbf{q}) + \mathbf{F}_f(\dot{\mathbf{q}}), \quad (2)$$

where $\mathbf{M}(\mathbf{q}) \in \mathbb{R}^{n \times n}$ is the generalized inertia matrix, $\mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}}$ denotes the Coriolis and centrifugal forces, $\mathbf{G}(\mathbf{q})$ is a vector of gravity terms, and $\mathbf{F}_f(\dot{\mathbf{q}})$ represents friction effects in the system. This equation captures how joint torques must compensate for inertia, Coriolis and centrifugal forces, gravity, and mechanical friction to follow a desired motion trajectory.

Parallel robots feature closed-loop kinematic chains where multiple limbs simultaneously support and actuate the end effector. This structural difference affects both the dynamic behavior and the control complexity of such systems. The closed-loop nature introduces kinematic constraints that must be respected at all times. These constraints relate the motion of different joints and ensure that the end effector follows a physically valid path. For parallel robots, the general dynamic model can be expressed as [63,64]

$$\tau + \mathbf{J}^T(q)\lambda = \mathbf{M}(q)\ddot{q} + \mathbf{C}(q, \dot{q})\dot{q} + \mathbf{G}(q) + \mathbf{F}_f(\dot{q}). \quad (3)$$

The constraint forces are incorporated into the dynamic model through the Jacobian, $\mathbf{J}^T(q)$, and the vector of Lagrange multipliers, λ , which are associated with the closed-loop kinematic constraints. Physically, the Lagrange multipliers represent the internal constraint forces required to maintain the closed-loop structure, and these forces are mapped to individual actuators via the Jacobian [65].

Cartesian robots are composed exclusively of orthogonal prismatic joints. Their dynamics are inherently simpler than those of articulated or parallel manipulators because of their linear and decoupled nature of motion. Coriolis or centrifugal effects are minimal or absent. The dynamic model thus reduces to a set of second-order linear differential equations, where actuator forces directly control axis accelerations. We can generalize the model as [66]

$$\mathbf{F} = \mathbf{M}\ddot{\mathbf{p}} + \mathbf{G}(\mathbf{p}) + \mathbf{F}_f(\dot{\mathbf{p}}), \quad (4)$$

where $\mathbf{p}$ is a vector of Cartesian position. The gravitational term mainly affects the vertical (z) axis. Cartesian structures are often more energy efficient due to simpler dynamics and lower internal motion losses.

The torque exerted by the robot joints is linked to instantaneous mechanical power

$$P(t) = \sum_{i=1}^n \tau_i(t) \cdot \dot{q}_i(t), \quad (5)$$

and the mechanical energy consumption

$$E = \int_{t_0}^{t_f} P(t) dt = \int_{t_0}^{t_f} \sum_{i=1}^n \tau_i(t) \cdot \dot{q}_i(t) dt. \quad (6)$$

2.3. Assessing the Energy Performance of an IR

ISO 9283:1998 [67] describes the methods for testing various performance characteristics for IRs, such as pose accuracy and pose repeatability. The performance is tested using a defined cubical volume within the robot's workspace (Figure 6). The test trajectories must be contained within that cube. There are certain constraints that must be applied. Firstly, the cube must be located in a portion of the workspace with the greatest anticipated use. Secondly, the cube must have the maximum possible volume with the edges parallel to the base coordinate system. One diagonal plane is then selected. The measurement plane, on which five measurement points are selected for testing, must be parallel to the selected diagonal plane.




Figure 6. The methodology shown in ISO:9283:1998 for robot performance characteristics evaluation: (a) example showing the plane $C_1 - C_2 - C_7 - C_8$ within the test cube with measurement points $P_1 - P_2 - P_3 - P_4 - P_5$, and (b) example showing the cube placed inside the robot workspace.

Some robot manufacturers have taken this procedure to assess the energy consumption of their robots and present it in the data sheet [68]. However, this solution does not represent an objective view on the EC, as the specific position of the cube is selected by the user and not all of the motion parameters are defined for the robot. The use of an ISO-cube to assess the energy consumption in an IR is shown in [46].

VDMA 24608:2013-11 [69] outlines a standardized methodology for energy measurements in 6-axis IRs, which is directly derived from ISO 9283:1998 [67]. In this approach, the reference cube is also defined to occupy the largest possible volume within the robot's working envelope, with its edges aligned parallel to the base coordinate frame, but the side lengths must be multiples of 200 mm. Trajectory points P_2 to P_5 are positioned inward from the cube's edges by 10% of the cube's diagonal length (Figure 6a), while P_1 is recorded as the center point, providing a fixed spatial reference for the test.

Static and dynamic measurements are then conducted as specified by the test procedure. During the static measurement, the robot is held at point P_1 , and energy is recorded

under conditions with the motor drives activated and deactivated. For dynamic measurements, the test cycle is defined as $30 \times (P2 - P3 - P4 - P5 - P2)$, executed at maximum speed under nominal load, which is defined by the maximum load at the wrist as per the manufacturer's data sheet [69].

While the documentation indicates that the type of interpolation (linear or joint motion) can be randomly selected for the motion of the IR [69], multiple studies have shown that the interpolation method can affect energy consumption under certain operating conditions [29,70].

The article [71] presents another standardized methodology for the comparison of 6-axis IRs, developed by Fraunhofer IWU in collaboration with German automotive automation initiative (AIDA). The aim here was to develop the methodology to compare IRs used particularly in the automotive industry.

The core of the method is an energy consumption evaluation based on a standardized reference trajectory. Additionally, IRs are categorized into six load-capacity classes, ranging from 0 to 250 kg in 50 kg increments, with a separate class for robots exceeding 250+ kg [71]. The reference trajectory was derived by statistically analyzing the operational parameters, such as point coordinates, movement types, trajectory lengths, velocities, accelerations, and path approximation values used in real robot programs in automotive body shops [72].

To ensure realistic testing conditions in terms of used payload, tool data were also analyzed to determine representative average values for each load-capacity class. Based on this, a configurable loading tool was developed to simulate standardized and comparable load conditions across reference trajectories for consistent energy evaluation [71,72].

The methodology includes three trajectory variants [71]:

- Standard—baseline cycle with fixed parameters defined by the statistical data used to create the reference trajectory.
- Performance—cycle with maximum speed and acceleration to analyze the maximal performance.
- Efficient—cycle with incrementally reduced velocity and acceleration to explore energy-cycle time trade-offs and identify optimal operating points.

The authors in [71] propose to compile the results into a data sheet, which includes detailed energy metrics of the IR, such as active, reactive, and apparent power, power factor, and energy per cycle, enabling transparent comparison between different robot models and manufacturers.

3. Comparative Review of ECO Techniques

Several previous reviews have categorized energy optimization strategies into hardware and software solutions, typically defining hardware as the selection of correct robot type, component replacements or additions, and software as encompassing trajectory and operation planning [11–13]. This review proposes a refined classification that focuses more specifically on the individual IR system level. The intention is to present a system-centric perspective that links the robot's physical structure and control strategy directly to its energy optimization potential. Our approach is schematically presented in Figure 7. Within software-based approaches, the strategies are divided into parameter optimization (e.g., tuning speed, acceleration, and jerk values along a predefined path) and trajectory optimization (e.g., shaping or modifying the path itself). While these two are closely inter-related and often employed together, they represent distinct levels of decision-making in energy-aware planning. A third category, task scheduling, focuses on the energy-efficient arrangement of tasks. Hardware-based optimization emphasizes the electrical and mechanical design of IRs, including structural design, the use of energy-efficient components, and, emerging in recent literature, the integration of energy regeneration units. The selected

strategies are discussed in more detail in the following sections with relevant references to the existing research.




Figure 7. Classification of energy optimization strategies.

3.1. Software

3.1.1. Parameter Optimization

Parameter optimization focuses on identifying the most energy-efficient operating parameters, most commonly velocity and acceleration, that minimize energy consumption in IRs without compromising task performance. These parameters are directly linked to motion profiles and actuator loads during operation. Some studies also consider external factors such as payload and drive lubricant temperature, which influence both mechanical performance under varying conditions.

Guerra-Zubiaga and Luong [73] focused on determining how linear speed, acceleration, payload, and temperature affect the IR energy consumption. The researchers used a Kawasaki ZZX130L IR in a simulated pick-and-place task. The results show that linear speed and acceleration contribute to 95% of energy consumption, while changes in payload and temperature have minor impacts. Additionally, higher speeds and accelerations are generally more energy efficient because of their effect on reducing cycle times.

Garcia et al. [29] investigated the impact of motion parameters on robot energy consumption also through simulations. Three experiments were conducted focusing on stand-still energy use, motion paths, and the effects of load, speed, and lubricant temperature. The findings highlighted, as expected, that minimizing dwell times and using energy-efficient positions (positions with low gravity torques) during standstill reduces idle energy consumption. Speeds between 50 and 100% of the maximum were identified to balance cycle time and energy use, with continuous motion paths preferred to avoid excessive deceleration–acceleration transitions. Payloads were most efficient when limited to 80% of the maximum, particularly at lower speeds.

Izagirre et al. [30] also investigated idle-time energy consumption, proposing a torque-based methodology to find the most optimal robot poses during standstill. By minimizing the torque applied to robot joints, this approach reduces both mechanical stress and energy consumption. Their experiments used a genetic algorithm to identify optimal poses, starting with tests on KUKA KR3 robot in a controlled laboratory environment. Validation was carried out on ABB IRB 6400R robot in a real assembly line, where the results showed a 31.37% reduction in average torque, highlighting its potential in practical applications. An analytical friction model in the third experiment showed that while speed remains the dominant factor, lower speeds reduced consumption in friction-model-based scenarios, i.e., load torque was not considered. High lubricant temperatures were found to lower friction-related consumption at high speeds. Applying these findings to an industrial welding robot demonstrated an energy reduction of 11–16% [30].

Torayev et al. [45] proposed a black-box model-based method that enables parameter optimization in real time without stopping production. Their algorithm starts with an initial guess of input parameters—velocity and acceleration—and then iteratively searches for more optimal parameters using various well-known optimization algorithms. When an optimal set of new parameters is found, the new parameters are set in the IR. They demonstrated the applicability of the solution on three applications—pick-and-place, bin picking, and drilling—and demonstrated significant energy savings compared to baseline parameters, which were defined as the minimum and maximum allowable parameter configurations for each task. Energy reductions were 44.59% in pick-and-place, 25% in bin picking, and 17.64% in drilling, where the focus was on the end-effector spindle speed and feed rate optimization of the IR [45].

Carabin and Scalera [74] presented an analytical model and experimental results on finding minimum-energy motion parameters by adjusting the acceleration time, deceleration time, and total motion duration. They considered two optimization scenarios: one with fixed acceleration and deceleration times and free total time, and another with fixed total time and free acceleration/deceleration times. Their experiments were conducted on two robotic systems: a linear axis of a Cartesian manipulator and a 1-DOF system composed of two coupled servomotors, either directly connected or coupled via a planetary gear. The Cartesian robot experiments demonstrated a potential of up to 33.6% reduction in energy consumption.

Pellicciari et al. [75] demonstrated a method for the electromechanical modeling of serial and parallel manipulators to compute energy-optimal trajectories based on predefined TCP position profiles, specifically focusing on pick-and-place operations. The approach optimized the task execution time of predefined joint trajectories, i.e., maintaining the same TCP position profiles while adjusting the speed of task execution. They conducted an experimental validation to highlight the existence of an energy consumption minimum in a pick-and-place application on ABB IRB6600. For the mathematical optimization problem, they implemented two simulation case studies, one involving a PUMA 560 serial manipulator and another involving a Delta robot. Both case studies confirmed the efficacy of the time-scaling method, showing reductions in energy consumption of up to 7.7% for the PUMA arm and 9.8% for the Delta robot.

The reviewed studies consistently emphasize velocity and acceleration as the most influential parameters in determining energy consumption in IRs. Optimal operating ranges typically fall between 50% and 80% of the robot's rated maximums. While secondary factors like payload and lubricant temperature have smaller overall effects, they may become critical under specific task or environmental conditions. Furthermore, several studies highlight the importance of reducing idle-time losses. These findings suggest strong potential for energy savings through parameter tuning—particularly in repetitive, short-cycle applications—without the need for major hardware modifications. However, practical implementation faces challenges such as limited access to low-level control parameters in commercial robot controllers, the need for parameter tuning across varying payload conditions, and ensuring that cycle time constraints are still met.

3.1.2. Task Scheduling

In multi-robot manufacturing environments, task scheduling plays a pivotal role in ensuring efficient system coordination. Traditionally, scheduling strategies have aimed to optimize productivity by sequencing robot operations to avoid collisions, reduce cycle time, and ensure timely task completion [76–78]. However, recent studies have demonstrated that these methods can be extended to support ECO without compromising throughput or violating operational constraints. A key energy-saving strategy is to coordinate robot

actions in ways that reduce idle times, balance workload across robots, and activate power-saving modes during non-productive periods.

The researchers in [15] introduced a mathematical model that incorporates robot speeds, positions, operational order, and power-saving modes. A mixed-integer linear programming formulation was applied; however, this was only suitable for instances with a small number of robots. For larger robot cells, the authors propose a hybrid heuristic algorithm, which uses parallelization for faster and more efficient computations. The proposed method was tested on a real IR cell containing six robots. The method achieved approximately a 20% reduction in energy consumption [15].

In [79] the researchers formulated a similar mixed-integer linear programming approach but integrated IR energy consumption models with production constraints. This approach allowed for the allocation of tasks and synchronization of multiple robots to avoid collisions and idle periods, thereby reducing unnecessary energy usage. The scheduling model incorporated variables such as operation start and stop times, cycle times, and energy consumption during active and idle states, enabling a holistic optimization of the robotic cell's energy profile [79].

In addition to task sequencing, trajectory time-scaling methods have been investigated for non-process-critical movements, such as transitions to and from home positions. Meike et al. [80] showed that adjusting the velocity of these movements and leveraging the earlier release of mechanical brakes during idle states can reduce energy consumption without affecting cycle time or process quality. These techniques are practical for deployment, as they require minimal changes to existing robot programs and no additional hardware modifications.

Task scheduling appears as a practical strategy for energy optimization in multi-robot production systems. By incorporating energy consumption models into scheduling algorithms, researchers have demonstrated energy reductions of up to 20% in real-world robotic cells [15]. Approaches combining scheduling with time-scaling and brake-release strategies further enhance efficiency without disrupting existing workflows [80]. Compared to other optimization strategies, task scheduling benefits from its system-level perspective, allowing coordinated adjustments across robots and tasks. However, challenges remain mainly regarding the scalability of the methods. Mixed-integer linear programming-based models offer precise optimization but face computational limitations as system complexity increases. Heuristic and learning-based methods show promise but require further validation and benchmarking. Additionally, expanding these methods to accommodate dynamic environments and unstructured tasks will be essential for broader industrial adoption.

3.1.3. Trajectory Optimization

Trajectory planning for IRs is a critical control strategy with significant potential for reducing energy consumption. Instead of following predefined paths, these approaches reformulate the motion of the IR as an optimization problem that considers energy efficiency alongside task constraints, dynamic feasibility, and limitations. By leveraging detailed kinematic and dynamic models, trajectory optimization enables IRs to reduce energy usage without compromising task performance. This subsection highlights the key contributions from recent research, showcasing the range of strategies and their reported effectiveness.

Liu et al. [28] introduced an energy-optimal trajectory planning method for palletizing robots, beginning with developing a rigid-flexible coupling dynamics model of the 4-DOF palletizing robot for torque prediction. The model was verified against a virtual model of the robot to ensure consistency. To identify the energy-efficient trajectory, the researchers formulated a mathematical optimization problem with boundary conditions specifying zero velocity and acceleration at both the start and end points—ensuring the robot begins

and finishes each motion from a standstill. In addition to trajectory optimization, the study also considers determining the optimal pick position to minimize energy usage. Experimental results demonstrated that the proposed trajectory optimization method reduces energy consumption by approximately 15%, and optimal pick point selection can further enhance efficiency, achieving up to a 30% reduction in total energy use for palletizing tasks [28]. Similarly, the work in [81] employed B-spline functions and a detailed robot model including robot dynamics, friction losses, servo drive losses, and energy exchange between axes through DC-bus coupling to optimize PTP trajectories, highlighting the importance of considering servo drive and inverter losses for realistic energy savings. The model was applied to KUKA KR 210-2 IR, where a reduction of up to 10% in energy consumption was achieved.

Otani et al. [82] proposed a novel method to find the energy-efficient trajectory under task deadline constraints. The method enables robots to adopt energy-efficient postures within the trajectory by low joint torque loads before proceeding to task completion within the specified cycle time. Unlike conventional approaches that focus on minimizing task duration or reducing acceleration and deceleration along predefined paths, their method allows the robot to adopt intermediate postures with low joint torque demands before proceeding to task completion within the given cycle time. By discretizing the end effector's position in both spatial and temporal dimensions, the search space for trajectory planning is expanded. This enables the optimization to identify energy-efficient configurations where the robot can momentarily remain stationary in low-torque postures, rather than strictly following the most direct or time-optimal route. The approach exploits the structural characteristics of the IR, such as link length and mass, to minimize torque requirements during the entire motion. Simulation results using models including the PUMA500 IR demonstrated a 10% reduction in total torque required in the new trajectory for executing the predefined task.

Garriz and Domingo [83] focused on maximizing manipulator performance while minimizing electrical energy consumption in an industrial sealing solution featuring a KUKA KR30-3 IR. In this application, trajectory optimization is critical for ensuring quality sealer application. The authors proposed a multiobjective optimization method based on the Kalman algorithm, which iteratively adjusts trajectories by balancing manipulability and energy consumption criteria. The cost function incorporates constraints such as joint limits or velocity restrictions. The method integrates the kinematic and dynamic model of the IR and applies an objective function with experimentally determined weights to guide optimization. Simulation results for three different trajectories demonstrated significant improvements, with electrical energy consumption reduced by up to 70.65%, and trajectory times shortened by up to 10.75%. The approach outperforms several previous methods in both manipulability and energy efficiency.

Some approaches consider energy recuperation and collision avoidance strategies within motion planning. Knöchelmann et al. [84] address the challenge of reducing energy consumption in industrial multi-robot cells through trajectory optimization while ensuring collision avoidance and exploiting energy sharing between robots over a common DC circuit. They implemented a two-level optimization approach, where the first level determines the optimal task sequence among the robots, and the second level refines individual robot trajectories to coordinate and balance energy flow across the shared DC-bus. The trajectory optimization incorporates a nonlinear collision detection constraint to prevent infeasible trajectories that could cause collisions in the packed robot cell. Simulation results on a three-robot system demonstrate that trajectory optimization without collision detection yields approximately 12.7% energy savings but results in collisions, whereas incorporating collision detection maintains around 9% energy savings.

Energy optimization can also be achieved along a fixed trajectory by adjusting motion timing. Li et al. [85] proposed a computationally efficient method that combines dynamic time-scaling with an energy characteristic parameter model to optimize energy use without altering the path itself. Their approach achieved 33.7% energy savings compared to the initial trajectory, while also reducing the computational cost associated with typical dynamic programming methods. This work highlights the potential of time-domain optimization techniques to achieve significant improvements in energy consumption within fixed motion paths.

Smart trajectory optimization offers a robust software-based solution for improving energy efficiency in IRs without hardware modifications. The reviewed methods generally frame the problem as multi-objective optimization, balancing energy use with additional goals such as manipulability [83], task timing [82], or collision avoidance [84]. These strategies leverage detailed physical and electrical models, yielding energy savings typically in the range of 10–30%, and in some cases even higher. However, computational complexity and the need for accurate system modeling can limit real-time deployment.

As discussed in [86], the computational speed of trajectory planning is a critical factor for practical deployment, particularly in dynamic or time-sensitive environments. One of the key challenges arises as the number of DOFs in the robot increases—the dimensionality of the optimization space grows, leading to significantly higher computational complexity [86]. This problem extends when additional constraints are introduced, such as joint limits, obstacle avoidance, or energy consumption bounds. These constraints increase the nonlinearity of the problem, often requiring more sophisticated solvers [86,87]. Furthermore, optimization outcomes are often task- and system-specific and not easily generalized.

Despite these challenges, trajectory optimization is a promising direction—especially when combined with energy-sharing mechanisms and adaptive scheduling.

Table 2 concludes an analysis of the main software-based ECO strategies' characteristics.

Table 2. SWOT analysis—software ECO strategies.

Strengths	Weaknesses
- Software based optimization can significantly reduce energy consumption without altering hardware. - Universally effective for wide variety of applications. - Online and real-time optimization techniques enable adaptation to varying operational conditions. - Can be applied at different levels – motion, task scheduling, and trajectory – offering multi-layered efficiency gains.	- Performance is highly dependent on accurate models; Artificial intelligence (AI) and machine learning (ML) based methods may require extensive training data. - Optimization strategies may vary significantly by task and robot type. - Implementation complexity increases when integrating real-time optimization within existing control systems. - May conflict with other goals such as cycle time or throughput if not well balanced.

Table 2. *Cont.*

Opportunities	Threats
- Integration with AI and ML allows adaptive, context-aware optimization. - Scalable to multi-robot cells through task scheduling and robot coordination. - Trajectory optimization can take advantage of electromechanical energy recovery (e.g., regenerative braking, DC bus coupling) highlighting the opportunity for hybrid strategies. - Combining real-time sensor data with optimization models could enable closed-loop energy-aware control.	- Over-reliance on complex algorithms may reduce interpretability and increase system debugging difficulty. - For large-scale robot systems, computational time or model complexity may become cumbersome. - Lack of standardization in optimization methods can hinder adoption across robot types and manufacturers. - Some algorithms do not scale efficiently. As the number of robots, tasks, or decision variables increases, the computational complexity can grow exponentially, making real-time or large-scale implementation impractical.

3.2. Hardware

3.2.1. Structural Design

One of the primary hardware-oriented strategies for enhancing energy efficiency in IRs involves optimizing the structural design to reduce the mass and inertia of moving components. Lighter moving elements directly translate to lower joint torque requirements, which in turn reduces actuator effort and overall energy consumption during operation. This is especially relevant for articulated robots, where each joint not only drives its own link but also bears the weight of mechanically linked actuators, as actuators are mounted directly at the joints. To address this, several studies propose relocating actuators closer to the robot base and transmitting motion through cable- or belt-driven systems [88–90]. In parallel, structural optimization through material selection, topology refinement, and dimensional adjustments has shown promising results in improving energy efficiency while maintaining stiffness and precision [91,92].

Aziz et al. [88] designed a lightweight 3-DOF planar industrial manipulator aimed at reducing the torque requirements and overall weight of the robotic arm by locating all motors at the base. This configuration transfers motor motion via belts, aluminum pulleys, and bearings to the links, resulting in lighter links, allowing the potential use of smaller motors. Simulation results showed a significant reduction in joint torques, which is up to 3.8 times lower compared to traditional manipulators, demonstrating the effectiveness of this design for more efficient operation in industrial automation. A similar design of a cable-driven robotic arm was presented in [89].

Nguyen et al. [90] introduced a low-cost, 6-DOF cable-driven robotic arm designed to achieve low inertia movement comparable in scale and payload capacity to the UR5 robotic arm. The design relocates the heavy wrist motors to the shoulder base, significantly reducing the wrist's weight, thereby decreasing torque demands on the elbow and shoulder motors. A hybrid actuation approach is employed, with direct-drive mechanisms at the shoulder and elbow joints to ensure precise positional repeatability, while the wrist's three

degrees of freedom are cable-driven through a sophisticated differential gear mechanism. Experimental evaluations demonstrated the arm's capability to manipulate a 5 kg payload with positional repeatability under 0.1 mm, alongside effective load handling, cable longevity, and precise manipulation, all achieved with a lightweight carbon fiber frame and cost-effective materials.

Jia and Sun [91] presented a comprehensive methodology for the structural optimization of IRs by integrating topological and dimensional parameter design. The proposed method aims to reduce the mass of robot components—particularly the lower arm—while maintaining structural stiffness and minimizing deformation. Their study models stiffness–mass relationships and employs a multi-objective genetic algorithm for optimization, reducing mass by 15% while achieving a deformation of less than 0.352 mm, confirming its viability for real-world application.

Another structural optimization strategy was described in [92], where the authors selected aluminum as the primary lightweight material for the robot structure. DC brushless motors were combined with harmonic drive gearboxes to enhance dynamic performance and achieve a compact design. The results demonstrated that their approach led to a 15.7% reduction in total mass and enabled the effective use of motors with up to 57% lower power ratings.

These studies highlight the significant impact that structural design has on energy consumption in IRs. Designs that relocate actuators to the base using cable- or belt-driven transmissions can dramatically reduce joint torques, and incorporating lightweight materials, such as aluminum or carbon fiber, and applying optimization algorithms to minimize link mass and deformation further enhance efficiency. Reductions in total system mass and the possibility of using less powerful motors have been demonstrated. These findings confirm that structural optimization, when approached holistically, can play a pivotal role in developing energy-efficient robots. However, retrofitting existing IRs with optimized structures is rarely feasible. Lightweight designs may introduce compliance or reduce vibration damping, necessitating additional suppression strategies [93]. Cable- or belt-driven solutions, while reducing inertia, also introduce additional maintenance and control complexity. Such solutions may also not meet the strict IP ratings required in harsh industrial environments.

3.2.2. Energy Recovery and Reuse

Energy recovery strategies aim to reclaim the electrical and kinetic energy that is typically lost during robotic operations, particularly during deceleration, braking, or idle phases [29]. In conventional IR systems, regenerative braking energy is briefly buffered in the internal capacitors of the DC-bus (Figure 5) [60]. Recent advancements aim to improve reusing or storing the recovered energy through mechanisms such as shared DC-bus architectures [60,94], external energy storage, and compliant or elastic recovery elements [36,95,96], demonstrating the feasibility of integrating both passive and active energy recovery systems in real-world applications.

Meike et al. [60] presented two complementary strategies to reduce energy consumption in IRs through capacitive energy buffering and robot-to-robot energy sharing. The first strategy involves the use of external capacitor banks connected to the DC-bus of a single robot controller. The results demonstrate energy savings of up to 20% during dynamic tasks. However, due to robots spending a comparatively small fraction of production time in motion, the return on investment for this hardware is relatively long; therefore, a more scalable approach involves connecting multiple robot controllers via a shared DC-bus, allowing braking energy from one robot to be reused by another robot. Simulations with 20 robots show that energy dissipation on brake resistors dropped from 24% to below 0.5%.

This method offers similar energy savings as capacitive buffering but without the need for expensive energy storage hardware.

Building on the energy buffering concept, the researchers in [94] explored the use of a shared storage capacitor connected to a common DC-bus among multiple IRs. Simulation results demonstrated that such a configuration could lead to energy savings of up to 30%, highlighting the potential of using buffer capacitors in multi-robot systems.

Another study by Meike et al. [61] introduced a smart power converter that builds upon the concept of a shared DC-bus across multiple robotic systems. The proposed system enables each robot's drive to store recuperated energy within an independent DC subgrid and to return the stored energy only when required. Experimental validation with two high-payload IRs demonstrated energy savings exceeding 20%.

Palomba et al. [95] proposed a method to add energy-recovering and energy-storing devices in the form of regenerative drives and compliant elements (springs that store elastic energy) that would function together to lower the energy consumption in mechanisms performing cyclic tasks. The compliant elements would reduce the necessary load from the actuators to drive the system. However, the compliant element needs to be designed optimally in terms of stiffness and preload to fit each system. The solution was tested on a Five-Bar Mechanism where an overall energy reduction of 70% was achieved, and on a SCARA robot, where the energy consumption was reduced by 60%. The approach works with an arbitrary motion trajectory and aims for universal applicability.

Scalera et al. [96] demonstrated the use of torsional springs to enhance energy efficiency in a parallel robot designed for pick-and-place tasks. The springs are mounted in parallel with the motors, enabling conversion between elastic and kinetic energy to reduce actuator torque demands. The results showed an energy consumption reduction of up to 67.8%.

Expanding here, the work in [36] demonstrated the use of variable stiffness springs in a parallel configuration with the motors. Similarly, one end of each spring is connected to the motor, while the other is linked to the robot joint; however, the stiffness of the springs is dynamically tunable with a dedicated motor, allowing the spring to adapt to different conditions during each phase of the robot's motion. Their aim was to target the braking phases of the robot's motion, where the energetic losses are usually the highest. The results show that energy consumption is minimized by up to 70% compared to the standard configuration without the springs.

The literature highlights multiple hardware-oriented strategies for recovering otherwise wasted energy in IRs. Capacitive buffering on the DC-bus of a single robot controller can yield some energy savings during motion-intensive operations, while multi-robot shared DC-bus systems can reduce brake resistor losses significantly, making them highly scalable for energy reuse across collaborative cells. Complementary to electrical strategies, mechanical solutions involving elastic or variable-stiffness elements have also shown good results. However, electrical energy recovery systems often involve additional hardware costs and complex system integration, and mechanical recovery methods must be carefully tuned to task-specific dynamics, requiring sophisticated control schemes and introducing additional maintenance complexity. Therefore, future work should aim to generalize these strategies across robot types and tasks. Ultimately, hybrid approaches combining electrical and mechanical recovery with intelligent software-based methods stand as a promising direction. From a deployment perspective, energy recovery systems often necessitate physical modifications to the robot's hardware and drive architecture as shown in [97]. Capacitive storage elements or shared DC-bus configurations may pose compliance challenges and require certification from system vendors to meet regulations.

3.2.3. Functional Redundancy

Functional redundancy refers to the intentional design or exploitation of additional DOFs beyond what is strictly necessary to achieve a task. This redundancy enables multiple kinematic solutions for the same end-effector position and orientation, allowing the optimization of secondary criteria such as joint torque, positioning [98], or energy consumption [31,32]. Beyond this, parallel redundancy introduces an additional layer of flexibility, where robot links are actuated by more than one actuator [99]. Though not presented as a standalone category in this review's ECO classification, functional and parallel redundancies offer a compelling example of a hybrid energy optimization strategy combining structural design with motion planning and control.

Boscariol et al. [31] explored how functional redundancy can lead to energy efficiency through smart motion planning. By building the dynamic and electrical model of the UR5 robot, they could simulate the energy consumption. Energy optimization can be achieved by adjusting the redundant degrees of freedom according to some task-specific constraints. Tasks with high inertial forces benefit the most from the proposed method, achieving up to 20.8% reduction in energy consumption [31].

Another work by the same author [32] explored the use of kinematic redundancy focusing on a SCARA robot integrated with an additional linear unit. The proposed method optimizes the trajectory by adjusting time parameters between via-points and motion profile parameters for the added axis, forming a numerical optimization problem. A detailed inverse dynamic and electromechanical model allows the precise estimation of energy drawn by each joint and the linear unit. Results show that introducing redundancy via an external axis enables both faster execution and energy savings—up to 33% [32].

While the earlier work in [32] focuses on hardware-extended redundancy by adding an external linear unit, the research in [31] shifts the focus to software-based functional redundancy, leveraging redundant degrees of freedom in joint configurations without modifying the robot structure. Both works use dynamic modeling and trajectory optimization but address redundancy from complementary angles—hardware augmentation versus internal joint-space exploitation.

Lee et al. [99] proposed an energy-saving strategy for parallel robotic manipulators by utilizing redundant parallel mechanical actuation. The concept is based on adding more actuators to the robot structure than functional DOFs, allowing the distribution of joint torques. This redundancy enables the optimization of actuator torques to minimize energy losses in three key areas: regenerative power dissipation, electric power loss, and frictional losses in reducers.

Using a 2-DOF parallel manipulator—structurally comparable to traditional palletizing robots with parallel linkage architecture—the researchers in [28] developed detailed kinematic and dynamic models and analyzed energy consumption in both non-redundant and redundant configurations. Their findings showed that redundant actuation reduced total energy consumption by 26.1%, mainly due to decreased frictional losses (26.7%) and regenerative power dissipation (75.3%) [99]. The optimized torque distribution also allows for the use of smaller motors and gearboxes, reducing the mechanical load and improving energy efficiency.

Functional redundancy presents the intersection of hardware and software strategies for energy optimization. In its simplest form, redundancy can be exploited entirely at the software level—particularly in articulated robots—by optimizing joint configurations without requiring any physical hardware modification. This makes it a highly accessible and low-cost solution for improving energy efficiency in existing systems. On the other hand, hardware-extended parallel redundancy, such as the integration of additional actuated axes or parallel actuators, can yield even greater energy savings by enabling more

flexible and distributed force generation. However, these approaches require modifications to the robot's physical structure, making them less feasible in commercial IRs. The need for custom mechanical design, additional components, and integration complexity limits their practical deployment in standardized manufacturing environments. Both approaches underscore the importance of co-optimizing robot structure and control strategy. Future work should explore the development of scalable frameworks that combine internal functional redundancy with adaptive motion planning and real-time optimization with AI-based methods.

Table 3 concludes an analysis of the main hardware-based ECO strategies' characteristics.

Table 3. SWOT analysis—hardware ECO strategies.

Strengths	Weaknesses
- Significant reduction in joint torques by relocating actuators to the base - Lightweight materials and optimized link design reduce inertia and enable use of smaller motors - Structural optimization through multi-objective algorithms yields stiffness-mass-efficient designs - Enables safer and more precise motion due to lower moving mass	- Increased mechanical complexity in transmission systems. - Maintenance and calibration of cable- or belt-driven systems may be more demanding. - Trade-offs between weight reduction and structural rigidity or load capacity. - Topology optimization requires high computational resources.
Opportunities	Threats
- Potential integration with energy recuperation systems - Applicable to a wide range of robots and payload classes - Functional redundancy can be leveraged for further energy optimization via smart motion planning	- Increased design complexity may limit applicability in cost-sensitive or compact environments. - Cable or belt fatigue and durability could limit long-term performance. - Dependence on highly specific mechanical configurations may limit modularity or scalability. - Structures with optimized topologies often feature open or exposed geometries, which can compromise their suitability for harsh industrial environments by reducing their ingress protection (IP) rating.

4. Discussion

Automation technologies are undergoing a rapid increase in adoption, following a similar trend observed in the electric vehicle (EV) sector, where sustainability and energy efficiency have become central priorities [100,101]. Just as EVs have urged the creation of standardized testing protocols and frameworks to support transparent performance evaluation, a similar shift is needed in IRs. However, in the field of IRs, the evaluation of energy performance remains fragmented. The reported energy savings in previous studies are often referenced against baseline or worst-case conditions (e.g., maximum speed and acceleration, and unoptimized trajectories), which limits the ability to make consistent

comparisons across robot types, tasks, or manufacturers. This variability complicates the objective assessment of energy optimization strategies and slows the establishment of best practices within the field. Currently, there is no universally accepted framework for evaluating or reporting the energy performance of IRs. As a result, researchers as well as integrators and end-users lack reliable tools for benchmarking and decision-making regarding energy-efficient robot deployment. To address this gap, future research should advocate for the development of the following:

- Standardized benchmarking procedures as an expansion of ISO 9283 [67] or VDMA [69];
- Development of a universal energy optimization tool that is usable with standardized robotic hardware across multiple manufacturers.

This standardization would simplify the comparisons of energy performance and promote transparency in the energy-related specifications provided by IR manufacturers. It would also form the foundation for developing industry-relevant ECO strategies that go beyond academic simulations and can be embedded into real-world manufacturing environments.

While the majority of energy optimization research in IRs has emerged from academic studies, often centered on dynamic modeling, simulation-based optimization, or laboratory-scale experiments, there is a growing need to transition these findings into practical, scalable solutions for industrial environments. Academic methods frequently emphasize efficiency gains under idealized conditions. This advocates for a shift toward implementation-focused research that considers integration into existing robotic systems, user accessibility, and maintainability.

An example of such an industry-oriented solution is the patented method by Björkman et al. [102], which introduces a software control method to dynamically switch between energy optimization mode and time optimization mode during robot program execution. The optimization can be started through an energy-optimization block within the program code, allowing energy-wise optimization over portions of the program. The energy is minimized under the constraint of a fixed path.

Building upon this, future research directions should prioritize the integration of hybrid energy optimization strategies and adaptive control algorithms that enable real-time decision-making. Current control-based ECO techniques often rely on static or predefined motion parameter sets, which may not be optimal under changing operational or environmental conditions. These approaches must not only respond dynamically to variations in the production environment, task load, cycle time, and robot condition but also integrate seamlessly into existing IR workflows.

In this context, the growing use of AI-based tools offers new opportunities to enhance the effectiveness of hybrid optimization. AI-based algorithms can analyze large volumes of sensory and operational data to continuously adjust motion parameters, predict energy demand, or identify inefficient behaviors [103]. When deployed within closed-loop IR systems, such intelligence can dynamically coordinate between control, scheduling, and actuation components. Prior research has explored data-driven techniques for motion planning [104,105], and even from an energy-efficiency point of view [106], as well as the optimization of motion parameters [107,108]. However, most of these studies have focused on offline or simulation-based applications. Applying AI methods in real-time optimization contexts where energy usage is monitored and minimized during robot operation remains an open research problem.

Recent developments indicate that AI-based methods are increasingly capable of bridging the gap between offline energy optimization and real-time application in industrial robotics. For instance, the use of long short-term memory (LSTM) deep neural networks for predicting energy consumption in IRs has been demonstrated in [109]. In some cases,

LSTM models have been integrated with model-based approaches to estimate energy use based on joint torques and velocities [110]. Dona' et al. [111] proposed a real-time trajectory generation method that produces near-optimal energy profiles with significantly reduced computation time. Other studies have addressed real-time optimization under unknown payload conditions, such as the method proposed in [112], which combines LSTM-based load estimation with particle swarm optimization (PSO) for online trajectory tuning. Their approach achieved more than 50% energy savings in simulated environments. Extending such strategies to real-world applications such as pick-and-place and palletizing holds strong potential for practical impact.

Real-time optimization would benefit from integrating power measurement directly at the manipulator–controller interface, rather than at the whole system input. Typical experimental validations assess power consumption at the full system level (including controllers, cooling systems, and auxiliary electronics), masking the relevant electromechanical losses specific to the manipulator.

Despite this progress, most studies remain limited to offline or simulation-based settings. The application of AI techniques in real-time energy optimization—where energy usage is monitored and minimized continuously during robot operation—remains largely unexplored. Progress in this area would benefit from integrating direct power measurements at the manipulator–controller interface as opposed to measuring system-wide energy consumption. Current experimental validations typically assess power usage at the full system level, which includes controllers, cooling units, and auxiliary electronics, thereby obscuring the electromechanical losses that are directly attributable to the manipulator.

By proper estimation of the consumed power at the manipulator interface, it is possible to isolate energy flow through the actuators, drive systems, and mechanical linkages, aligning the experimental data more closely with the assumptions and focus of dynamic modeling. This opens the door for training more accurate data-driven models, developing predictive analytics for joint-specific losses, and enabling closed-loop AI agents to learn and apply energy-aware behaviors in real time. In such a framework, the models could infer energy-optimal motion trajectories and parameters in dynamic industrial environments.

The application of high-resolution energy monitoring and intelligent control architecture could redefine energy optimization in IR, shifting the paradigm from pre-planned static simulation-based tuning to responsive, autonomous, and self-improving systems. Herein, the authors propose a standalone power measurement and energy optimization tool that can be integrated between the manipulator and the control system, effectively isolating the mechanical subsystem to enable precise evaluation of energy losses within the manipulator itself. Furthermore, connecting the tool to the IR control system would allow for continuous monitoring, dynamic evaluation, and real-time optimization of energy usage, specifically facilitating joint-level measurement. In this context, AI integration would be well-suited to enhance energy-aware trajectory planning, motion parameter tuning, and task scheduling.

5. Conclusions

As the usage of IRs continues to expand across manufacturing sectors, optimizing their energy consumption has become increasingly critical for achieving both economic efficiency and environmental sustainability. This review provided a detailed exploration of the controllable aspects of energy use in IRs, linking kinematic structures, component-level characteristics, and control strategies to energy performance. Through the classification of energy optimization approaches into software- and hardware-based strategies, the review has highlighted the different methods available for reducing energy consumption, from motion parameter tuning and trajectory planning to structural redesign, energy recov-

ery systems, and applying functional redundancy. A comparative SWOT analysis further clarified the trade-offs and complementarities between these strategies.

Software-based energy optimization methods offer a flexible and non-intrusive approach to improving energy efficiency in IRs. These strategies can be implemented without altering the IR hardware and are generally applicable across various robot types and applications. Their compatibility with AI and ML models enable advanced predictive and adaptive control, especially when integrated with real-time sensor feedback. Furthermore, software-based methods are usable in multi-robot environments, allowing coordination and optimization at the task level. However, the effectiveness of these techniques heavily relies on the accuracy of the dynamic and energy models of the IR. Optimization outcomes may vary depending on the complexity of the task, and there can be trade-offs between energy efficiency and other priorities such as cycle time or precision. Some algorithms may suffer from limited scalability and increased computational demands, potentially impacting system responsiveness and interpretability.

Hardware-based strategies offer substantial energy savings by directly reducing joint torques, moving masses, and inertial loads, thereby enhancing the baseline energy efficiency of IRs independent of control strategies. Structural optimizations can also improve motion precision and contribute to safer and more responsive robotic behavior due to reduced moving masses. The integration of energy recuperation mechanisms, such as regenerative braking or compliant elements, presents further opportunities for capturing and reusing otherwise wasted energy, especially in cyclic tasks. However, implementing hardware modifications typically requires physical redesign, retrofitting, or the adoption of non-standard components. This may be impractical in existing industrial settings due to cost, downtime, integration complexity, or environmental constraints. Their effectiveness also varies depending on the application domain of the IR, and they may introduce additional mechanical complexity or maintenance requirements.

Emerging trends such as the integration of DC microgrids and regenerative energy sharing point toward a future in which robots not only consume energy efficiently but also participate actively in intelligent energy ecosystems. Equally important is the role of modular power measurement systems that enable precise energy monitoring and continuous feedback, creating the foundation for real-time optimization.

Looking ahead, the convergence of these developments with AI and ML presents a powerful opportunity: the creation of hybrid, adaptive, and context-aware energy optimization frameworks that are capable of scaling to complex, real-world applications, with the control layer in particular. ML models can dynamically adjust motion parameters, select energy-efficient trajectories, or even predict system behavior under varying workloads. Reinforcement learning and neural network-based controllers may eventually enable robots to autonomously fine-tune their actions in response to changes in task complexity, load, or workspace configuration. However, the successful deployment of such systems will depend on robust real-time data acquisition, especially at the manipulator level.

Taking into account the growing trend in automatization, similar to that observed in the EV industry, future efforts should also address the open question of how to develop standardized energy certification schemes and publicly accessible consumption profiles for commercial IRs. These initiatives would promote transparency, support benchmarking, and drive manufacturers toward designing inherently more energy-efficient IR systems. By synthesizing current research, technological advances, and future directions, this review aims to support the development of best practices and inspire further innovation toward energy-aware IRs.

Author Contributions: Conceptualization, J.M. and A.R.; writing—draft preparation, J.M.; writing—review and editing, A.R.; visualization, J.M.; supervision, A.R. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Acknowledgments: The authors acknowledge the use of ChatGPT (OpenAI) to support English language editing and phrasing improvements during the preparation of this manuscript.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. IFR. Record of 4 Million Robots in Factories Worldwide. Available online: <https://ifr.org/ifr-press-releases/news/record-of-4-million-robots-working-in-factories-worldwide> (accessed on 1 November 2024).
2. ISO 8373; Robotics—Vocabulary. ISO (International Organization for Standardization): Geneva, Switzerland, 2021.
3. Parliament, E. Parliament Adopts New Rules to Boost Energy Savings | News | European Parliament. Available online: <https://www.europarl.europa.eu/news/en/press-room/20230707IPR02421/parliament-adopts-new-rules-to-boost-energy-savings> (accessed on 15 November 2024).
4. Eurostat. Final Energy Consumption in Industry—Detailed Statistics. Available online: https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Final_energy_consumption_in_industry_-_detailed_statistics (accessed on 15 November 2024).
5. IFR. TOP 5 Global Robotics Trends 2025. Available online: <https://ifr.org/ifr-press-releases/news/top-5-global-robotics-trends-2025> (accessed on 29 April 2025).
6. Barosz, P.; Gołda, G.; Kampa, A. Efficiency Analysis of Manufacturing Line with Industrial Robots and Human Operators. *Appl. Sci.* **2020**, *10*, 2862. [CrossRef]
7. Huang, G.; He, L.Y.; Lin, X. Robot adoption and energy performance: Evidence from Chinese industrial firms. *Energy Econ.* **2022**, *107*, 105837. [CrossRef]
8. Wang, E.Z.; Lee, C.C.; Li, Y. Assessing the impact of industrial robots on manufacturing energy intensity in 38 countries. *Energy Econ.* **2022**, *105*, 105748. [CrossRef]
9. Barnett, N.; Costenaro, D.; Rohmund, I. Direct and Indirect Impacts of Robots on Future Electricity Load. In Proceedings of the ACEEE Summer Study on Energy Efficiency in Industry, Denver, CO, USA, 15–18 August 2017.
10. Liu, Y.; Zhao, C.; Kong, Z.; Taghizadeh-Hesary, F. Pathways to stable economic policy: The role of industrial robotics and energy transition. *Renew. Energy* **2025**, *244*, 122751. [CrossRef]
11. Carabin, G.; Wehrle, E.; Vidoni, R. A Review on Energy-Saving Optimization Methods for Robotic and Automatic Systems. *Robotics* **2017**, *6*, 39. [CrossRef]
12. Sękal, A.; Błaszczyk, T.; Foit, K.; Kost, G. Selected Issues, Methods, and Trends in the Energy Consumption of Industrial Robots. *Energies* **2024**, *17*, 641. [CrossRef]
13. Vászárhelyi, J.; Salih, O.M.; Rostum, H.M.; Benotsname, R. An Overview of Energies Problems in Robotic Systems. *Energies* **2023**, *16*, 8060. [CrossRef]
14. Yao, M.; Shao, Z.; Zhao, Y. Review on Energy Consumption Optimization Methods of Typical Discrete Manufacturing Equipment. In Proceedings of the Intelligent Robotics and Applications, Yantai, China, 22–25 October 2021; Liu, X.J., Nie, Z., Yu, J., Xie, F., Song, R., Eds.; Springer: Cham, Switzerland, 2021; pp. 48–58. [CrossRef]
15. Bukata, L.; Sucha, P.; Hanzalek, Z.; Burget, P. Energy Optimization of Robotic Cells. *IEEE Trans. Ind. Inform.* **2017**, *13*, 92–102. [CrossRef]
16. Page, M.J.; McKenzie, J.E.; Bossuyt, P.M.; Boutron, I.; Hoffmann, T.C.; Mulrow, C.D.; Shamseer, L.; Tetzlaff, J.M.; Akl, E.A.; Brennan, S.E.; et al. The PRISMA 2020 Statement: An Updated Guideline for Reporting Systematic Reviews. *BMJ* **2021**, *372*, n71. [CrossRef]
17. Tricco, A.C.; Lillie, E.; Zarin, W.; O'Brien, K.K.; Colquhoun, H.; Levac, D.; Moher, D.; Peters, M.D.; Horsley, T.; Weeks, L.; et al. PRISMA Extension for Scoping Reviews (PRISMA-ScR): Checklist and Explanation. *Ann. Intern. Med.* **2018**, *169*, 467–473. [CrossRef]
18. IFR. Industrial Robots. Available online: <https://ifr.org/industrial-robots> (accessed on 10 April 2025).
19. Amiri, P.; Müller, M.; Southgate, M.; Theodoridis, T.; Wei, G.; Richards-Brown, M.; Holderbaum, W. A Statistical Analysis of Commercial Articulated Industrial Robots and Cobots. *J. Manuf. Mater. Process.* **2024**, *8*, 216. [CrossRef]
20. ABB. IRB 160. Available online: <https://new.abb.com/products/robotics/robots/articulated-robots/irb-1600> (accessed on 20 May 2025).
21. my.KUKA. KR 6 R500 Z200-2. Available online: <https://my.kuka.com/s/product/kr-6-r500-z2002/01t1i0000025DCIAA2> (accessed on 20 May 2025).

22. FESTO. Three-Dimensional Gantry YXCR. Available online: https://www.festo.com/ee/en/p/three-dimensional-gantry-id_YXCR/ (accessed on 20 May 2025).
23. Bouri, M.; Clavel, R. The Linear Delta: Developments and Applications. In Proceedings of the ISR 2010 (41st International Symposium on Robotics) and ROBOTIK 2010 (6th German Conference on Robotics), Munich, Germany, 7–9 June 2010; pp. 1–8.
24. Kong, X.; Gosselin, C.; Richard, P.L. Type Synthesis of Parallel Mechanisms With Multiple Operation Modes. *J. Mech. Des.* **2006**, *129*, 595–601. [CrossRef]
25. Jia, G.; Huang, H.; Wang, S.; Li, B. Type Synthesis of Plane-Symmetric Deployable Grasping Parallel Mechanisms Using Constraint Force Parallelogram Law. *Mech. Mach. Theory* **2021**, *161*, 104330. [CrossRef]
26. Igus. 3-Axis Delta Robot. Available online: <https://www.igus.com/product/22609?artNr=LCA-DR1000-03> (accessed on 20 May 2025).
27. FESTO. Parallel Kinematic System, Delta Robot EXPT. Available online: https://www.festo.com/us/en/p/parallel-kinematic-system-delta-robot-id_EXPT/ (accessed on 20 May 2025).
28. Liu, Y.; Liang, L.; Han, H.; Zhang, S. A Method of Energy-Optimal Trajectory Planning for Palletizing Robot. *Math. Probl. Eng.* **2017**, 5862457. [CrossRef]
29. Garcia, R.R.; Bittencourt, A.C.; Villani, E. Relevant factors for the energy consumption of industrial robots. *J. Braz. Soc. Mech. Sci. Eng.* **2018**, *40*, 464. [CrossRef]
30. Izagirre, U.; Arcin, G.; Andonegui, I.; Eciolaza, L.; Zurutuza, U. Torque-based methodology and experimental implementation for industrial robot standby pose optimization. *Int. J. Adv. Manuf. Technol.* **2020**, *111*, 2065–2072. [CrossRef]
31. Boscariol, P.; Caracciolo, R.; Richiedei, D.; Trevisani, A. Energy Optimization of Functionally Redundant Robots through Motion Design. *Appl. Sci.* **2020**, *10*, 3022. [CrossRef]
32. Boscariol, P.; Richiedei, D. Trajectory Design for Energy Savings in Redundant Robotic Cells. *Robotics* **2019**, *8*, 15. [CrossRef]
33. Ma, J.; Lian, B.; Wang, M.; Dong, G.; Li, Q.; Wu, J.; Yang, Y. Optimal design of a parallel assembling robot with large payload-to-mass ratio. *Robot. Comput.-Integr. Manuf.* **2023**, *80*, 102474. [CrossRef]
34. Khalil, W.; Dombre, E. Introduction to geometric and kinematic modeling of parallel robots. In *Modeling, Identification and Control of Robots*, 1st ed.; Elsevier Science & Technology: Amsterdam, The Netherlands, 2004; pp. 171–190.
35. Vodovozov, V.; Lehtla, M.; Raud, Z.; Semjonova, N.; Petlenkov, E. Managing Energy Consumption of Linear Delta Robots Using Neural Network Models. *Energies* **2024**, *17*, 4081. [CrossRef]
36. Balderas Hill, R.; Briot, S.; Chriette, A.; Martinet, P. Minimizing the Energy Consumption of a Delta Robot by Exploiting the Natural Dynamics. In *ROMANSY 23—Robot Design, Dynamics and Control*; Venture, G.; Solis, J.; Takeda, Y.; Konno, A., Eds.; Springer International Publishing: Cham, Switzerland, 2021; Volume 601, pp. 213–221; Series Title: CISM International Centre for Mechanical Sciences. [CrossRef]
37. Li, Y.; Bone, G. Are Parallel Manipulators More Energy Efficient? In Proceedings of the the 2001 IEEE International Symposium on Computational Intelligence in Robotics and Automation, Banff, AB, Canada, 29 July–1 August 2001; pp. 41–46; Cat. No.01EX515.
38. Kurilova-Palaisaitiene, J.; Permin, E.; Mannheim, T.; Buhse, K.; Lorenz, M.; Schmitt, R.; Corves, B.; Björkman, M. Industrial energy efficiency potentials: An assessment of three different robot concepts. *Int. J. Sustain. Eng.* **2017**, *10*, 185–196. [CrossRef]
39. Yigit, A.; Breton, D.; Zhou, Z.; Laliberte, T.; Gosselin, C. Kinematic Analysis and Design of a Novel (6+3)-DoF Parallel Robot with Fixed Actuators. *arXiv* **2023**, arXiv:2304.12499.
40. Fabris, G.; Scalera, L.; Gasparetto, A. Dynamic Modelling and Energy-Efficiency Optimization in a 3-DOF Parallel Robot. *Int. J. Adv. Manuf. Technol.* **2024**, *132*, 2677–2699. [CrossRef]
41. Morikawa, K.; Katsura, S. In-Link Actuators for Low-Inertia Robots. *IEEE/ASME Trans. Mechatron.* **2024**, *30*, 2122–2131. [CrossRef]
42. Pham, A.D.; Ahn, H.J. High Precision Reducers for Industrial Robots Driving 4th Industrial Revolution: State of Arts, Analysis, Design, Performance Evaluation and Perspective. *Int. J. Precis. Eng. Manuf.-Green Technol.* **2018**, *5*, 519–533. [CrossRef]
43. Peta, K.; Suszyński, M.; Wiśniewski, M.; Mitek, M. Analysis of Energy Consumption of Robotic Welding Stations. *Sustainability* **2024**, *16*, 2837. [CrossRef]
44. Zhou, J.; Jiang, P.; Huang, C.; Ping, K. An energy-saving optimization method for robot laser welding systems based on multi-objective metaheuristic algorithms. *J. Meas. Eng.* **2025**, *13*, 24349. [CrossRef]
45. Torayev, A.; Martínez-Arellano, G.; Chaplin, J.C.; Sanderson, D.; Ratchev, S. Online and Modular Energy Consumption Optimization of Industrial Robots. *IEEE Trans. Ind. Inform.* **2024**, *20*, 1198–1207. [CrossRef]
46. Uhlmann, E.; Reinkober, S.; Hollerbach, T. Energy Efficient Usage of Industrial Robots for Machining Processes. *Procedia CIRP* **2016**, *48*, 206–211. [CrossRef]
47. Almurib, H.A.F.; Al-Qrimli, H.F.; Kumar, N. A review of application industrial robotic design. In Proceedings of the 2011 Ninth International Conference on ICT and Knowledge Engineering, Bangkok, Thailand, 12–13 January 2012; pp. 105–112. [CrossRef]
48. Šitum, Z. Fluid power drives in robotic systems. In Proceedings of the International Conference Fluid Power 2017, Maribor, Slovenia, 14–15 September 2017; pp. 11–24. [CrossRef]

49. Habibi, S.; Richards, R.; Goldenberg, A. Hydraulic actuator analysis for industrial robot multivariable control. In Proceedings of the 1994 American Control Conference—ACC '94, Baltimore, MD, USA, 29 June–1 July 1994; Volume 1; pp. 1003–1007. [CrossRef]
50. Lischinsk, P.; Canudas-de Wit, C.; Morel, G. Friction compensation for an industrial hydraulic robot. *IEEE Control Syst.* **1999**, *19*, 25–32. [CrossRef]
51. Pustavrh, J.; Hočevár, M.; Podržaj, P.; Trajkovski, A.; Majdič, F. Comparison of hydraulic, pneumatic and electric linear actuation systems. *Sci. Rep.* **2023**, *13*, 20938. [CrossRef] [PubMed]
52. Mykhailishyn, R.; Savkiv, V.; Maruschak, P.; Xiao, J. A systematic review on pneumatic gripping devices for industrial robots. *Transport* **2022**, *37*, 201–231. [CrossRef]
53. Maeda, S.; Tsujiuchi, N.; Koizumi, T.; Sugiura, M.; Kojima, H. Development and Control of a Pneumatic Robot Arm for Industrial Fields. *Int. J. Adv. Robot. Syst.* **2012**, *9*, 59. [CrossRef]
54. Shi, Y.; Cai, M.; Xu, W.; Wang, Y. Methods to Evaluate and Measure Power of Pneumatic System and Their Applications. *Chin. J. Mech. Eng.* **2019**, *32*, 42. [CrossRef]
55. Bonfà, F.; Salvatori, S.; Benedetti, M.; Introna, V.; Ubertini, S. Monitoring compressed air systems energy performance in industrial production: Lesson learned from an explorative study in large and energy-intensive industrial firms. *Energy Procedia* **2017**, *143*, 396–403. [CrossRef]
56. Meike, D.; Pellicciari, M.; Berselli, G. Energy Efficient Use of Multirobot Production Lines in the Automotive Industry: Detailed System Modeling and Optimization. *IEEE Trans. Autom. Sci. Eng.* **2014**, *11*, 798–809. [CrossRef]
57. Grebers, R.; Gadaleta, M.; Paugurs, A.; Senfelds, A.; Avotins, A.; Pellicciari, M. Analysis of the Energy Consumption of a Novel DC Power Supplied Industrial Robot. *Procedia Manuf.* **2017**, *11*, 311–318. [CrossRef]
58. Gutwald, B.; Liu, H.; Barth, M.; Reichenstein, T.; Römer, P.; Franke, J. Conversion of Industrial Robots to DC Power Supply: Efficiency, Protection Concepts and Recuperation Potential. In Proceedings of the 2024 IEEE 3rd Industrial Electronics Society Annual On-Line Conference (ONCON), Virtual Event, 8–10 December 2024; pp. 1–6. [CrossRef]
59. Senfelds, A.; Avotins, A.; Ribickis, L.; Apse-Apsitis, P. DC Microgrid for Robotic Manufacturing—Field demonstration and laboratory experience. *RE&PQJ* **2024**, *17*, 194–197. [CrossRef]
60. Meike, D.; Ribickis, L. Recuperated energy savings potential and approaches in industrial robotics. In Proceedings of the 2011 IEEE International Conference on Automation Science and Engineering, Shanghai, China, 10–12 June 2011; pp. 299–303. [CrossRef]
61. Meike, D.; Senfelds, A.; Ribickis, L. Power converter for DC bus sharing to increase the energy efficiency in drive systems. In Proceedings of the IECON 2013—39th Annual Conference of the IEEE Industrial Electronics Society, Vienna, Austria, 10–13 November 2013; pp. 7199–7204. [CrossRef]
62. Siciliano, B.; Khatib, O. Dynamics. In *Springer Handbook of Robotics*, 2nd ed.; Springer Handbooks; Springer: Cham, Switzerland, 2016; pp. 149–195.
63. Park, S.B.; Kim, H.S.; Song, C.; Kim, K.J. Dynamics Modeling of a Delta-Type Parallel Robot. In Proceedings of the IEEE ISR 2013, Seoul, Republic of Korea, 24–26 October 2013; pp. 1–5.
64. Thanh, T.D.; Kotlarski, J.; Heimann, B.; Ortmaier, T. On the Inverse Dynamics Problem of General Parallel Robots. In Proceedings of the 2009 IEEE International Conference on Mechatronics, Malaga, Spain, 14–17 April 2009; pp. 1–6.
65. Lynch, K.M.; Park, F.C. *Modern Robotics*; Cambridge University Press: Cambridge, UK, 2017.
66. Tuli, T. Mathematical Modeling and Dynamic Simulation of Gantry Robot Using Bond Graph. In Proceedings of the Information and Communication Technology for Development for Africa, Bahir Dar, Ethiopia, 25–27 September 2017; Mekuria, F., Nigussie, E., Dargie, W., Edward, M., Tegegne, T., Eds.; Springer: Cham, Switzerland, 2018; pp. 228–237.
67. ISO 9283; Manipulating Industrial Robots—Performance Criteria and Related Test Methods. ISO (International Organization for Standardization): Geneva, Switzerland, 1998.
68. IRB 1600ID | ABB. Available online: <https://new.abb.com/products/robotics/robots/articulated-robots/irb-1660id> (accessed on 11 April 2025).
69. VDMA 24608; Energy performance for industrial robots. VDMA (Verband Deutscher Maschinen- und Anlagenbau e.V.): Frankfurt/Main, Germany, 2013. Available online: <https://www.dinmedia.de/en/technical-rule/vdma-24608/197844344> (accessed on 19 May 2025).
70. Gadaleta, M.; Berselli, G.; Pellicciari, M.; Grassia, F. Extensive experimental investigation for the optimization of the energy consumption of a high payload industrial robot with open research dataset. *Robot. Comput.-Integr. Manuf.* **2021**, *68*, 102046. [CrossRef]
71. Findeisen, M.; Todtermuschke, M.; Schaffrath, R.; Putz, M. A method for energetic comparison of 6-axis industrial robots and its further scope for resource-efficient plant design. *Procedia Manuf.* **2018**, *21*, 251–258. [CrossRef]
72. Todtermuschke, M.; Findeisen, M.; Bauer, A. Methodology for Creation a Reference Trajectory for Energetic Comparability of Industrial Robots in Body Shop. *Procedia CIRP* **2014**, *23*, 122–126. [CrossRef]

73. Guerra-Zubiaga, D.A.; Luong, K.Y. Energy consumption parameter analysis of industrial robots using design of experiment methodology. *Int. J. Sustain. Eng.* **2021**, *14*, 996–1005. [CrossRef]
74. Carabin, G.; Scalera, L. On the Trajectory Planning for Energy Efficiency in Industrial Robotic Systems †. *Robotics* **2020**, *9*, 89. [CrossRef]
75. Pellicciari, M.; Berselli, G.; Leali, F.; Vergnano, A. A method for reducing the energy consumption of pick-and-place industrial robots. *Mechatronics* **2013**, *23*, 326–334. [CrossRef]
76. Zacharia, P.T.; Aspragathos, N.A. Optimal robot task scheduling based on genetic algorithms. *Robot. Comput.-Integr. Manuf.* **2005**, *21*, 67–79. [CrossRef]
77. Pupa, A.; Van Dijk, W.; Brekelmans, C.; Secchi, C. A Resilient and Effective Task Scheduling Approach for Industrial Human-Robot Collaboration. *Sensors* **2022**, *22*, 4901. [CrossRef]
78. Tejer, M.; Szczepanski, R.; Tarczewski, T. Robust and efficient task scheduling for robotics applications with reinforcement learning. *Eng. Appl. Artif. Intell.* **2024**, *127*, 107300. [CrossRef]
79. Vergnano, A.; Thorstensson, C.; Lennartson, B.; Falkman, P.; Pellicciari, M.; Leali, F.; Biller, S. Modeling and Optimization of Energy Consumption in Cooperative Multi-Robot Systems. *IEEE Trans. Autom. Sci. Eng.* **2012**, *9*, 423–428. [CrossRef]
80. Meike, D.; Pellicciari, M.; Berselli, G.; Vergnano, A.; Ribickis, L. Increasing the energy efficiency of multi-robot production lines in the automotive industry. In Proceedings of the 2012 IEEE International Conference on Automation Science and Engineering (CASE), Seoul, Republic of Korea, 20–24 August 2012; pp. 700–705. [CrossRef]
81. Hansen, C.; Oltjen, J.; Meike, D.; Ortmaier, T. Enhanced approach for energy-efficient trajectory generation of industrial robots. In Proceedings of the 2012 IEEE International Conference on Automation Science and Engineering (CASE), Seoul, Republic of Korea, 20–24 August 2012; pp. 1–7. [CrossRef]
82. Otani, T.; Nakamura, M.; Kimura, K.; Takanishi, A. Energy Efficient Path and Trajectory Optimization of Manipulators With Task Deadline Constraints. *IEEE Access* **2023**, *11*, 107441–107450. [CrossRef]
83. Garriz, C.; Domingo, R. Trajectory Optimization in Terms of Energy and Performance of an Industrial Robot in the Manufacturing Industry. *Sensors* **2022**, *22*, 7538. [CrossRef]
84. Knöchelmann, E.; Steinke, D.; Greenyer, J.; Spindeldreier, S.; Ortmaier, T. Trajectory Optimization Methods for Robotic Cells Considering Energy Efficiency and Collisions. In *Annals of Scientific Society for Assembly, Handling and Industrial Robotics*; Springer: Berlin/Heidelberg, Germany, 2020; pp. 229–240. [CrossRef]
85. Li, X.; Lan, Y.; Jiang, P.; Cao, H.; Zhou, J. An Efficient Computation for Energy Optimization of Robot Trajectory. *IEEE Trans. Ind. Electron.* **2022**, *69*, 11436–11446. [CrossRef]
86. Liu, J.; Yap, H.; Khairuddin, A. Review on Motion Planning of Robotic Manipulator in Dynamic Environments. *J. Sens.* **2024**, *2024*, 5969512. [CrossRef]
87. Zhang, S.; Zanchettin, A.M.; Villa, R.; Dai, S. Real-Time Trajectory Planning Based on Joint-Decoupled Optimization in Human-Robot Interaction. *Mech. Mach. Theory* **2020**, *144*, 103664. [CrossRef]
88. Aziz, M.A.S.; Zhanibek, M.; Elsayed, A.S.A.; AbdulRazic, M.O.M.; Yahya, S.; Almurib, H.A.; Moghavvemi, M. Design and analysis of a proposed light weight three DOF planar industrial manipulator. In Proceedings of the 2016 IEEE Industry Applications Society Annual Meeting, Portland, OR, USA, 2–6 October 2016; pp. 1–7. [CrossRef]
89. Kim, Y.J. Anthropomorphic Low-Inertia High-Stiffness Manipulator for High-Speed Safe Interaction. *IEEE Trans. Robot.* **2017**, *33*, 1358–1374. [CrossRef]
90. Nguyen, V.P.; Chow, W.T.; Dhyhan, S.B.; Zhang, B.; Han, B.S.; Wong, H.Y.A. Low-Cost Cable-Driven Robot Arm with Low-Inertia Movement and Long-Term Cable Durability. *Robotics* **2024**, *13*, 128. [CrossRef]
91. Jia, J.; Sun, X. Structural Optimization Design of a Six-Degrees-of-Freedom Serial Robot with Integrated Topology and Dimensional Parameters. *Sensors* **2023**, *23*, 7183. [CrossRef]
92. Yin, H.; Huang, S.; He, M.; Li, J. An overall structure optimization for a light-weight robotic arm. In Proceedings of the 2016 IEEE 11th Conference on Industrial Electronics and Applications (ICIEA), Hefei, China, 5–7 June 2016; pp. 1765–1770. [CrossRef]
93. Thomsen, D.; Sørensen, R.; Balling, O.; Zhang, X. Vibration Control of Industrial Robot Arms by Multi-Mode Time-Varying Input Shaping. *Mech. Mach. Theory* **2021**, *155*, 104072. [CrossRef]
94. Rankis, I.; Meike, D.; Senfelds, A. Utilization of Regeneration Energy in Industrial Robots System. *Enerġetika Un Elektrotehnika* **2013**, *31*, 95–100.
95. Palomba, I.; Wehrle, E.; Carabin, G.; Vidoni, R. Minimization of the Energy Consumption in Industrial Robots through Regenerative Drives and Optimally Designed Compliant Elements. *Appl. Sci.* **2020**, *10*, 7475. [CrossRef]
96. Scalera, L.; Carabin, G.; Vidoni, R.; Wongratanaphisan, T. Energy efficiency in a 4-DOF parallel robot featuring compliant elements. *Int. J. Mech. Control* **2019**, *20*, 49–57.
97. Khalaf, P.; Richter, H. Trajectory Optimization of Robots with Regenerative Drive Systems: Numerical and Experimental Results. *arXiv* **2018**, arXiv:1804.00311. [CrossRef]

98. Erdős, G.; Kovács, A.; Váncza, J. Optimized joint motion planning for redundant industrial robots. *CIRP Ann.* **2016**, *65*, 451–454. [CrossRef]
99. Lee, G.; Sul, S.K.; Kim, J. Energy-saving method of parallel mechanism by redundant actuation. *Int. J. Precis. Eng. Manuf.-Green Technol.* **2015**, *2*, 345–351. [CrossRef]
100. Di Martino, A.; Miraftabzadeh, S.M.; Longo, M. Strategies for the Modelisation of Electric Vehicle Energy Consumption: A Review. *Energies* **2022**, *15*, 8115. [CrossRef]
101. Weiss, M.; Winbush, T.; Newman, A.; Helmers, E. Energy Consumption of Electric Vehicles in Europe. *Sustainability* **2024**, *16*, 7529. [CrossRef]
102. Björkman, M.; NORRLÖF, M. A Method for Reducing the Energy Consumption of an Industrial Robot and an Industrial Robot System. WO WO2011042293A1, 14 April 2011.
103. Biswas, P.; Rashid, A.; Biswas, A.; Nasim, M.A.A.; Chakraborty, S.; Gupta, K.D.; George, R. AI-driven approaches for optimizing power consumption: A comprehensive survey. *Discov. Artif. Intell.* **2024**, *4*, 116. [CrossRef]
104. Tamizi, M.G.; Yaghoubi, M.; Najjaran, H. A review of recent trend in motion planning of industrial robots. *Int. J. Intell. Robot. Appl.* **2023**, *7*, 253–274. [CrossRef]
105. Meyes, R.; Tercan, H.; Roggendorf, S.; Thiele, T.; Büscher, C.; Obdenbusch, M.; Brecher, C.; Jeschke, S.; Meisen, T. Motion Planning for Industrial Robots using Reinforcement Learning. *Procedia CIRP* **2017**, *63*, 107–112. [CrossRef]
106. Wang, X.; Cao, J.; Cao, Y.; Zou, F. Energy-efficient trajectory planning for a class of industrial robots using parallel deep reinforcement learning. *Nonlinear Dyn.* **2025**, *113*, 8491–8511. [CrossRef]
107. Yao, M.; Su, Y.; Shao, Z.; Huo, Y. Data-driven Energy Evaluation and Optimization Method for Industrial Robots. In Proceedings of the 2023 8th International Conference on Automation, Control and Robotics Engineering (CACRE), Hong Kong, China, 13–15 July 2023; pp. 287–293. [CrossRef]
108. Zhang, M.; Yan, J. A data-driven method for optimizing the energy consumption of industrial robots. *J. Clean. Prod.* **2021**, *285*, 124862. [CrossRef]
109. Jiang, P.; Wang, Z.; Li, X.; Wang, X.V.; Yang, B.; Zheng, J. Energy consumption prediction and optimization of industrial robots based on LSTM. *J. Manuf. Syst.* **2023**, *70*, 137–148. [CrossRef]
110. Jiang, P.; Zheng, J.; Wang, Z.; Qin, Y.; Li, X. Industrial Robot Energy Consumption Model Identification: A Coupling Model-Driven and Data-Driven Paradigm. *Expert Syst. Appl.* **2025**, *262*, 125604. [CrossRef]
111. Dona', D.; Franzese, G.; Della Santina, C.; Boscariol, P.; Lenzo, B. Real-Time Generation of Near-Minimum-Energy Trajectories via Constraint-Informed Residual Learning. *arXiv* **2025**, arXiv:2501.09450. [CrossRef]
112. Chang, Q.; Yuan, T.; Li, H.; Chen, Y.; Wang, X.; Gao, S.; Ren, H.; Zhao, X.; Wang, L. A Data-Driven Method for Predicting and Optimizing Industrial Robot Energy Consumption Under Unknown Load Conditions. *Actuators* **2024**, *13*, 516. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Robotic Prostheses and Neuromuscular Interfaces: A Review of Design and Technological Trends

Pedro Garcia Batista ^{1,2}, André Costa Vieira ^{1,3} and Pedro Dinis Gaspar ^{1,3,*}

¹ Department of Electromechanical Engineering, University of Beira Interior, Rua Marquês d'Ávila e Bolama, 6201-001 Covilhã, Portugal; pedro.g.batista@ubi.pt (P.G.B.); andre.costa.vieira@ubi.pt (A.C.V.)

² Cova da Beira University Hospital Center (ULS Cova da Beira), 6200-251 Covilhã, Portugal

³ C-MAST—Center for Mechanical and Aerospace Science and Technologies, 6201-001 Covilhã, Portugal

* Correspondence: dinis@ubi.pt

Abstract

Neuromuscular robotic prostheses have emerged as a critical convergence point between biomedical engineering, machine learning, and human–machine interfaces. This work provides a narrative state-of-the-art review regarding recent developments in robotic prosthetic technology, emphasizing sensor integration, actuator architectures, signal acquisition, and algorithmic strategies for intent decoding. Special focus is given to non-invasive biosignal modalities, particularly surface electromyography (sEMG), as well as invasive approaches involving direct neural interfacing. Recent developments in AI-driven signal processing, including deep learning and hybrid models for robust classification and regression of user intent, are also examined. Furthermore, the integration of real-time adaptive control systems with surgical techniques like Targeted Muscle Reinnervation (TMR) is evaluated for its role in enhancing proprioception and functional embodiment. Finally, this review highlights the growing importance of modular, open-source frameworks and additive manufacturing in accelerating prototyping and customization. Progress in this domain will depend on continued interdisciplinary research bridging artificial intelligence, neurophysiology, materials science, and real-time embedded systems to enable the next generation of intelligent prosthetic devices.

Keywords: neuromuscular prostheses; machine learning; artificial intelligence; sensor integration

1. Introduction

Limb loss profoundly impacts autonomy, employability, and mental health. This disability severely impairs human mobility and physical function, imposing a significant global health economic burden. Despite decades of research on robotic prosthetic hands (RPHs) leading to breakthroughs in robotics, a relevant paradox persists: most amputees still rely on technologies that have barely evolved in half a century. This apparent disparity is less surprising when it is considered the difficulty of replicating the human hand. The hand holds a large sensory representation in the brain, and its grasping function is the most sophisticated coordination task. Additionally, the limb, which possesses both the highest density of mechanoreceptors and the most Degrees of Freedom (DoF), enables humans to both perceive and manipulate [1]. Over the past two decades, engineering and clinical research have converged to create neuromuscular robotic prostheses—artificial limbs that sense the user's residual neuromuscular activity, translate it into multi-degree-of-freedom (DoF) motion, and increasingly return haptic or proprioceptive feedback.

1.1. Rationale and Scope of This Review

Previous surveys have treated myoelectric control, hybrid interfaces, or specific actuation technologies separately. The present review adopts an interdisciplinary perspective, integrating recent advances in sensor technologies (sEMG and nerve cuffs); signal processing; actuation and human–machine interfaces (targeted muscle reinnervation, osseointegrated gateways, brain–computer interfaces). By providing both foundational concepts and state-of-the-art insights, this review aims to present a comprehensive understanding of the current state of neuromuscular robotic prostheses and their potential for future development, thereby maximizing their real-world impact.

While several prior reviews have surveyed individual subdomains, for example, the mechanical design and kinematic performance of upper-limb prostheses, the electrophysiology and acquisition methods for myoelectric control, or machine-learning approaches to signal decoding, relatively few works synthesize these topics in a single review. This manuscript fills that gap by integrating hardware design, neuromuscular and implanted sensing modalities, control, and machine-learning architectures for real-time decoding and surgical/clinical adjuncts (e.g., TMR and osseointegration) that affect device integration.

The research question of this paper is: which hardware and algorithmic bottlenecks, across sensing, actuation, and control, currently limit neuromuscular robotic prostheses, and which integrated solutions show the most promise to mitigate those bottlenecks?

As a working hypothesis, it is proposed that the greatest near-term gains will come from co-designed solutions that trade off sensing modality stability, control complexity, and actuator energy density, rather than incremental improvements confined to a single subsystem.

1.2. Socio-Economic Relevance

Multifunctional robotic hands (e.g., i-Limb Quantum, bebionic, Michelangelo) have been increasingly adopted worldwide since their introduction. Nevertheless, factors such as cost, weight, battery life, and limited robustness continue to constrain their use, particularly in low-resource settings (Figure 1) [1]. Some clinical trials have shown that when myoelectric prostheses are coupled with pattern-recognition control, users report significant gains in satisfaction measures, with a clear preference for pattern recognition control over conventional direct control during home use [2]. The humanitarian and economic incentive for further innovation in this field is therefore clear.

1.3. Clarification of Terminology

In the context of this review, several terms are used to describe different yet interconnected aspects of assistive technologies for limb replacement. The term “robotic prosthetics” is a general term sometimes used to describe robotic technologies applied to prosthetic limbs, regardless of the control method. When associated with the term “neuromuscular,” it refers specifically to robotic prosthetic systems that are controlled via neuromuscular signals, typically captured through electromyography. These systems integrate sensors, actuators, and intelligent algorithms to interpret user intention from muscle activity. The expression “neuromuscular interfaces” designates the interface layer between the human neuromuscular system and the prosthesis, encompassing signal acquisition methods, signal processing algorithms, and the pathways used to relay control commands and feedback. Finally, “robotic prosthetic devices” refers to the physical electromechanical constructs of the prosthetic limbs themselves, including structural components, actuators, and embedded electronics, independent of the control strategy.

Although the term “neuromuscular prostheses” may suggest applications across all limb types, this review focuses specifically on upper limb prostheses. This is because

the biomechanical and functional complexity of the human hand necessitates a far more sophisticated engineering approach than that required for lower-limb prostheses. While lower-limb prostheses primarily address simple tasks (locomotion and weight-bearing), upper-limb devices must replicate fine motor control, multi-degree-of-freedom movements, and intuitive interaction. Consequently, advanced neuromuscular technologies are almost exclusively developed and implemented in upper-limb prostheses.




Figure 1. Examples of functional hand replacements include the following: (a) A customized cosmetic hand solution (ITOP, Italy) offering a natural appearance. (b) A body-powered prosthesis (Ottobock, Germany) commonly used for upper-limb amputees. (c) An sEMG-based pattern recognition system that classifies grasp types. (d) A system combining sEMG control with extracellular stimulation via an implanted FINE, providing sensory feedback for home use; the implant remained stable for over five years, with sensory feedback tested for up to 13 days. (e) Fully implanted myoelectric sensors that deliver robust, reliable signals irrespective of arm positioning, socket rotation or perspiration; (f) Six-degree-of-freedom prosthetic hand control (i-Limb Ultra, Össur, Iceland) utilizing threshold-based sEMG control for grasp switching. (g) Sensory feedback through intraneural TIMEs, which encodes information on object shape and stiffness. (h) Hand transplant remains a promising yet not routine procedure due to several surgical and clinical challenges. Abbreviations: DOF—degree of freedom; FINE—flat interface nerve electrode; sEMG—surface electromyography; TIME—transverse intrafascicular multichannel electrode (adapted with permission from [1]).

2. Literature-Search Strategy

This review on neuromuscular robotic prostheses was conducted as a narrative state-of-the-art analysis. The methodology comprised a literature search of relevant publications to elucidate the anatomical, physiological, and technological foundations underpinning advanced prostheses. The review primarily focused on the following:

- **Literature Search Strategy:** A detailed search was performed across scientific databases like PubMed, IEEE Xplore, Scopus, ScienceDirect, and SpringerLink using the following terms: “neuromuscular prostheses”, “robotic prostheses”, “surface electromyogra-

phy”, “myoelectric control”, “biomechanical design”, and “sensor integration”. The literature search encompassed publications from January 2000 to July 2025 to ensure an up-to-date perspective.

- *Selection Criteria:* Publications were selected based on their relevance to the current challenges and technological trends. Emphasis was placed on studies addressing the integration of sensor technologies, the underlying anatomophysiological principles, and innovative design approaches. Both experimental studies and review articles were considered, provided they contributed critical insights into the field.
- *Data Extraction and Synthesis:* Key data such as methodological approaches, technological innovations, and challenges in the design and control of neuromuscular prostheses were extracted from the selected articles. The extracted information was organized into thematic sections reflecting the major aspects of the state-of-the-art discussion, such as anatomical fundamentals, physiological mechanisms, signal processing techniques, and emerging actuator technologies.

This structured approach enabled a thorough examination of scientific literature, forming the basis for the subsequent discussion on current trends and future directions in the development of robotic prosthetic devices. Given the layered structure of the following section, a flowchart is provided in Figure 2 to summarize the hierarchical organization of the content of the manuscript.




Figure 2. Overview of the structural organization of Section 3.

3. Survey on Neuromuscular Robotic Prostheses

This section presents a detailed synthesis of the key findings derived from the state-of-the-art analysis of neuromuscular robotic prostheses. The results are organized into thematic subsections that capture the critical technological advancements, their interpretation, and the conclusions drawn from the literature review. In practice, neuromuscular prostheses are primarily developed for the upper limb, where precision, dexterity, and multifunctional control are essential. In contrast, lower-limb prostheses typically rely on simpler control schemes focused on gait stabilization.

3.1. Sensors and Actuators in Robotic Prostheses

3.1.1. Human Hand Biomechanics and Direct-Drive Architectures

Recreating the biomechanics of a human hand is a difficult challenge. Prosthetic hands rely on various actuators and actuation methods to mimic the movement and configurations of human hands. The various possible movements of the fingers allow for a range of configurations or patterns, enabling adaptation to the objects being grasped. Although the possible postures are infinite, some are regularly used in daily life activities. Hand configurations for grasping objects can be categorized as grips or pinches. The main pinches are referred to as digit-to-digit pinches, which involve two fingers, with the thumb opposing the index or any other finger and can be further divided into force pinches and precision pinches. There are also tridigital pinches used for writing. Grips, on the other hand, are a method of grasping objects, more robust compared to pinches, since they allow bigger contact areas and the use of all fingers. A detailed description of grasp patterns is given by [3].

Early prototypes sought to replicate the skeletal structure of the fingers, but only recently has it become feasible to effectively implement biomimetic actuation using muscle-like actuators. One major challenge is integrating actuators, sensors, and electronic components into a prosthesis that maintains the dimensions and weight of a natural hand. Significant systems integration and miniaturization are essential before such devices can be effectively employed by amputee patients. Concerning sensory feedback, one current objective is to simulate the mechanoreceptors found in human skin, each sensitive to different stimuli and responsible for distinct sensory functions.

Actuators are the components responsible for generating motion in fingers and applying force according to the kinematics of articulating parts. Often, numerous joints will be coupled to form a finger. Knowing the actuator position, the position of all joints in that finger can be determined. Therefore, actuator position is a parameter or a degree of freedom (DOF) that controls a set of configurations [4]. In this case, there is a direct drive, where each finger or even a joint has its own actuator, allowing independent movement and high dexterity. The most common configuration of prosthetic hands has six DOFs. One for each finger (flexion/extension), and the thumb usually has two DOFs, using two independent actuators, one to control flexion/extension and another to control circumduction. The circumduction rotation along the circumduction axis is required to alternate between a lateral grasp and a power or precision grasp [5].

3.1.2. Underactuation Principles and Implementations

The underactuation method uses a single actuator to control several independent degrees of freedom [6]. In this sense, the single actuator parameter cannot be used to describe the position of the joints since they are dependent on the contact state of each finger link with the object. These mechanisms use differential systems since they allow the other fingers to continue their movement when one of the fingers is stopped by the grasping object.

Underactuation is a widely adopted strategy to simplify the mechanical design of prosthetic devices while preserving functional dexterity. An underactuated system is characterized by having fewer actuation degrees (DOA—degree of actuation) than DOFs. In passive DOFs (non-actuated), the system exploits the capacity to adapt to the contacting surface, enabling a self-adjusting grip without controlling each joint individually. Such systems reduce the number of required motors, thereby simplifying complexity, weight, and cost [6]. Several commercial robotic prosthetic solutions employ underactuated mechanisms. These include the Michelangelo prosthesis (Ottobock, Germany), the i-Limb Ultra (Össur, Iceland), the bebionic (Ottobock), and the VINCENTevolution 3 (Vincent Systems, Germany) [1].

3.1.3. Soft Materials, 3D Printing, and High-Resolution Tactile Sensing

The most common actuator used in prostheses today is a direct current (DC) motor. These motors are small and lightweight and can be packaged in the hand or forearm. To reduce the speed and increase the limited torque from these devices, gearing, lead screws, and even harmonic drives may be used. Cables are also usually used to transmit force from actuators to fingers, enabling compact designs [7]. Among recent solutions are the soft pneumatic actuators (SPAs), which consist of rubbery structures composed of one or more interconnected and physically constrained chambers that, when pressurized pneumatically or hydraulically, perform flexural or torsional movements [8].

Currently, various research groups are developing innovative solutions such as 3D printing with biomaterials. These soft materials and flexible tactile sensors can replicate the mechanoreceptor responses that take place in human skin, offering enhanced sensory feedback. Some researchers have demonstrated a fully 3D-printed biomimetic skin where interdigitated capacitive electrodes and piezoelectric structures are co-printed within elastomeric layers (achieving both static and dynamic tactile sensing) and interfaced with a myoelectric band for closed-loop force feedback in prosthetic hands [9,10]. There are also examples of hybrid hands with a rigid endoskeleton and soft robotic joints equipped with independent neuromorphic tactile sensing layers. Tactile sensor designs can be categorized as capacitive, resistive, piezoresistive, piezoelectric, triboelectric, and optical. Flexible tactile sensors based on triboelectric, piezoelectric, and capacitive mechanisms using nanomaterials achieve good spatial resolutions and low detection thresholds for high-fidelity human-scale touch reproduction [11]. Closed-loop prosthetic systems integrate capacitive tactile sensing with feedback actuation for grip force regulation, enabling naturalistic sensory experiences through dynamic modulation of applied force, as exemplified by wearable devices such as Tacsac [12]. When combined with neuromorphic tactile sensors, in which spiking neural activity can be converted into targeted nerve stimulation, these systems may further enhance the naturalness of sensory experiences in the future.

Neuromorphic tactile sensing is a relatively recent technology that has demonstrated the ability of fine textures discrimination. For example, optical neuromorphic sensors encode tactile events as spike trains using spatial and temporal coding schemes, achieving robust classification across variations in contact force [13]. Machine learning algorithms can process these patterns to modulate grasp force and respond adaptively. Such approaches support more naturalistic and responsive tactile feedback loops in prosthetic control systems.

A major barrier is the lack of standardized protocols for evaluating sensitivity, linearity, hysteresis, and longevity. Unified benchmarks are needed to compare disparate designs objectively. Durability remains critical, with cyclic loading and environmental factors degrading sensor performance. Open-source platforms like BioIn-Tacto could democratize sensor development, enhance reproducibility, and accelerate innovation [14].

Bringing together the tactile perception capabilities outlined earlier with motor control is a crucial step in enabling their application to robotic functions such as grasping and object manipulation. By providing tactile feedback with both high sensitivity and rapid response, neuromorphic tactile sensors make it possible for robots to modulate grip force dynamically [13]. Most neuroprosthetic devices currently encode neural stimulation patterns as artificial pulse waveforms. While effective, such non-biomimetic signals often produce perceptual experiences that feel unnatural to the user. Neuromorphic tactile systems offer a significant shift in approach by delivering sensory feedback through naturally occurring spike trains, providing a more faithful, biomimetic representation of touch. This capability enables prosthetic hands equipped with neuromorphic sensors to reinstate neuromuscular reflexes and markedly enhance performance in fine-motor tasks. For example, integrating these sensors has been shown to restore reflexive control in amputees, improving precision grasping and delicate manipulation [13].

By using biomimetic signal encoding to stimulate either the user's peripheral nerves or the somatosensory cortex, neuromorphic systems have the potential to reproduce complex somatosensory experiences. These capabilities can also improve grasp stability and coordinated motor control [13]. Furthermore, innovations involving real-time neural decoding and adaptive feedback frameworks, powered by advanced deep learning models, will continue to drive and accelerate these advances in the years to come.

3.1.4. Soft Pneumatic Actuators

Recent investigations have shown that soft manipulators can achieve performance levels comparable to those of rigid systems in various applications. Research in this area addresses not only the materials involved but also the mechanical structure and integrated sensory capabilities of prosthetic devices, particularly for applications requiring delicate object manipulation. Flexible sensors and soft robotic components have contributed to safer and more effective human–prosthesis interaction, thereby improving usability and mitigating injury risks. However, further research is needed to evaluate whether soft architecture represents a viable solution for more widespread use.

The mechanical characterization of thermoplastic polyurethane, a flexible polymer widely used in soft robotics, reveals that it is suitable for the design and manufacturing of industrial and agricultural soft grippers [15]. The Fin Ray effect is a biomimetic structural principle in which a flexible structure with angled internal cross-beams bends toward an applied force, allowing it to wrap around and adapt to the shape of an object for stable grasping. A soft robotic gripper inspired by this effect demonstrates improved mechanical performance, adaptability, and reduced structural failures under testing conditions [16]. Further computational optimization of a Fin Ray effect soft robotic finger achieves a better balance between gripping force and flexibility [17]. In parallel, advances in low-dimensional carbon-based flexible sensors have enabled their integration into wearable networks for environmental and health monitoring, with porous carbonaceous materials providing enhanced sensitivity and durability for human–machine interfaces [18].

Emerging designs include prosthetic hands controlled by magnetically actuated muscles implanted in the forearm, providing intuitive and near-natural motion through thought-based commands. These designs represent a shift toward reducing external hardware dependence while improving user satisfaction and control fidelity. Innovative actuation systems, such as McKibben pneumatic artificial muscles (PAMs), offer the advantage of integrating structure with sensing, thereby enhancing safety and efficiency. McKibben PAMs have emerged as a promising alternative for actuation in prostheses, offering a unique blend of advantages such as lightweight design, flexibility, substantial force generation, and a remarkable similarity to the structure and function of biological muscles [19,20].

3.1.5. Pneumatic Artificial Muscles: Principles and Applications

PAMs are pneumatic actuators inspired by biological muscles. They consist of an elastomeric tube surrounded by a braided mesh of inextensible fibers (such as Kevlar or nylon) [19]. When the tube is inflated with compressed air, it expands radially and contracts axially, generating force and movement analogous to human muscle contraction. The force generated is directly proportional to the air pressure and the muscle's cross-sectional area, allowing for a wide range of actuation (from delicate movements to the generation of significant force). These features make PAMs ideal for a variety of applications in robotic prostheses. For upper limbs, they can be used to actuate fingers, hands, and arms, enabling grasping, object manipulation, and the performance of daily tasks with dexterity and natural movement. The integration of pressure sensors and tactile feedback can further enhance the precision and sensitivity of these prostheses. In lower limbs, PAMs can simulate the action of leg, ankle, and foot muscles, offering a more natural and efficient gait adaptable to various terrains and activities. Originally created for prosthetic limbs in the 1950s and 60s, the McKibben pneumatic artificial muscle has recently been commercialized and re-engineered for use in robotics and skeletal models. It is essentially a tube with a braided mesh shell. When air pressure is applied to the inside, the tube shortens and creates tension, acting like a spring. This design provides several key advantages: it is flexible, lightweight, and has a unique spring-like behavior with variable stiffness [20].

The use of robots for rehabilitation is also gaining attention, with the use of multi-joint robotic arms in physical therapy programs, as these are more effective than single-joint machines for recreating natural human movement. A key concern for such robots is safety, which is addressed by using PAMs that are inherently compliant due to their use of compressed air [21]. Their ability to generate force smoothly and adaptively makes them ideal for amplifying muscle strength and reducing fatigue. In pursuit of heightened human-like dexterity, researchers are employing a highly biomimetic design methodology to replicate the intricate manipulation capabilities of the human hand within robotic systems [22].

Pneumatic muscle systems are distinguished from alternative actuators, such as electric motors or hydraulic systems, by their exceptionally high power-to-weight and power-to-volume ratios. These characteristics make them highly suitable for the design of human-centric exoskeleton systems intended to provide strength and mobility assistance [23]. Despite these undeniable advantages, PAMs still face challenges for wide-spread application in robotic prostheses. Their inherently nonlinear dynamics require sophisticated control strategies, such as gain-scheduled or feedback-linearization controllers, to achieve accurate force and movement regulation [23]. Additionally, the materials used in PAMs may experience wear and fatigue, particularly in applications demanding high frequency and intensity of actuation. Therefore, advances in material science are expected to have a significant impact on their durability and performance.

Ongoing research into new materials and manufacturing techniques may enhance durability and reduce maintenance needs. Moreover, integrating PAMs with EMG or neural signals could revolutionize the market for robotic prostheses, making them more intuitive and responsive to user commands. Recent efforts in prosthetic technologies focus on systems with compliant actuation and flexible architectures, enabling devices to deliver more natural and precise movements while remaining adaptable and accessible to a broad range of users.

Current actuation strategies in neuromuscular prostheses encompass a wide range of technologies, from conventional DC motors with gear reductions to pneumatic artificial muscles and soft actuators. Compact DC motors are often reported to provide relatively high fingertip forces and fast response, but this generally comes at the cost of weight

and limited intrinsic compliance. Pneumatic artificial muscles, by contrast, tend to offer favorable specific power and biologically realistic compliance, though they are typically constrained by the need for external compressors and by reduced efficiency. Soft robotic actuators are frequently highlighted for their safety in human–robot interaction, but their force output and durability remain limited. Rather than enabling strict comparison, these examples illustrate the diversity of engineering trade-offs across actuation strategies.

A major hurdle for clinical use is the long-term reliability of these systems. Daily prosthesis uses subjects' actuators to extensive cyclic loading, yet durability under clinically realistic conditions remains underexplored. Pneumatic artificial muscles are often reported to fail through bladder fatigue, while soft actuators are prone to material degradation. In contrast, geared DC motor drives tend to demonstrate greater longevity, although typically at the expense of increased wear in transmission components. Overall, durability assessments remain inconsistent across the literature, in part due to the lack of standardized testing protocols. This gap highlights the need for systematic evaluation under clinically realistic conditions before many emerging actuators can be considered viable for long-term clinical deployment.

3.2. Categories of Neuromuscular Robotic Prosthesis

Prosthetic devices represent an essential solution for restoring lost functionality in amputee patients. One of their primary contributions is the significant improvement in the quality of life. The diversity of these devices allows them to be classified based on multiple criteria, including the desired functionality, the activation mechanism, and the control method used.

In terms of functionality, prosthetic devices are divided into two main categories: passive and active prostheses. Passive prostheses are primarily designed to restore the physical appearance of the limb without providing active functional capabilities [1]. They are valued mainly for their esthetic appeal, which helps the user achieve a more integrated and harmonious body image. The design of these devices focuses on visual appearance, comfort during use, and minimal weight. Their applicability spans all levels of amputation, as their purpose is restricted to enhancing visual appeal and psychological comfort.

Passive prostheses can be further divided into esthetic and functional types. Esthetic prostheses aim exclusively at the visual replacement of the missing limb, striving to achieve a high degree of realism. In contrast, functional prostheses are designed to perform specific tasks, adapting their appearance according to their intended function. Additionally, passive prostheses can be classified into two subcategories: static and adjustable. Static prostheses offer no movement, whereas adjustable prostheses incorporate mechanisms that allow adaptation to the user's specific needs [24].

Active prostheses, on the other hand, are defined by their ability to execute movements through internal forces, either by using mechanical energy generated by the user's body or by actuators powered by external energy sources. These medical devices are recognized for providing mobility to the amputated limb and are categorized based on their activation system. Body-powered activation operates via a mechanical interface—such as cables attached to the patient's shoulder—allowing specific movements to be translated into actions by the prosthesis. An example of this mechanism is the ability to open or close a prosthetic claw through arm extension, offering a simple and intuitive control interface [24].

For electrically activated prostheses, control methods may include external button control, myoelectric control, or hybrid systems [24]. Button-controlled prostheses require an external power source to drive the motors for articulated movements, allowing the user greater freedom by enabling individual movement control through buttons or pressure sensors. Sensor technology plays a crucial role in the accessibility and performance of

myoelectric control systems. Affordable EMG electrodes are crucial in resource-constrained environments [25]. Studies that focus on human–machine interfaces for upper-limb prostheses reveal a growing trend toward integrating various control strategies, including EMG-based, hybrid, and advanced signal processing approaches, to enhance user comfort, device precision, and overall usability [26].

Feedback related to command execution is crucial for prosthetic device users, as it facilitates interaction with the device. This feedback can take several forms, vibrational, tactile, electrical, or somatosensory, each providing a distinct method to simulate the sensation of touch or interaction with the environment. In vibrational feedback, sensors incorporated into the prosthesis extremities activate vibratory devices on the user’s skin upon contact with an object, providing an indirect perception of touch that significantly enhances functionality and control in myoelectric prostheses, where capturing electromyographic signals is critical. Alternatively, tactile feedback applies a slight pressure on the user’s skin instead of vibration, simulating a more direct physical sensation of touch [27].

For electrical feedback, low-intensity current pulses are used to simulate the sensation of touch when applied to the skin. Despite its potential to enhance prosthetic capabilities, electrical feedback may interfere with the acquisition of myoelectric signals. Finally, somatosensory feedback involves the direct application of electrical stimuli to the nervous system to replicate the natural sensation of force and proprioception derived from sensors. Although this approach is at the forefront of research, its implementation faces significant challenges, such as the need for precise neural mapping to effectively and selectively activate the appropriate areas [27].

Each feedback method possesses unique characteristics and is at varying stages of development and research. Nonetheless, a shared objective across all approaches is the enhancement of prosthetic functionality and user experience, fostering greater integration and more natural interaction with the surrounding environment.

3.3. *Electroneuromyography: From Physiology to Signal Processing*

Locomotion is a fundamental human characteristic, largely driven by muscles acting as biological motors that generate force upon contraction. Electromyography uses electrodes to evaluate the electrical activity of nerves and muscles [28]. Originally applied for diagnosing neuromuscular pathologies—from congenital myopathies to adult-onset disorders—EMG has been crucial since the early 20th-century cellular studies of muscle physiology. Two main approaches exist: classical neurological EMG (where electrically stimulated responses are measured under static conditions) and kinesiological EMG [28,29]. The kinesiological electromyography focuses on voluntary muscle activation during functional tasks, making it valuable in clinical assessment, rehabilitation, sports training, and in the development of robotic prostheses.

During muscle contraction, electrical signals are produced by the sequential activation of many muscle fibers. These signals are complex and chaotic due to the superposition of multiple action potentials and are captured using multi-channel amplifiers, data acquisition systems, and specialized processing software [28,29]. The amplitude and frequency of the EMG signal increase with muscle force, and the basic functional unit, the motor unit (comprising a motor neuron and its innervated fibers), ensures that all fibers contract synchronously. With a broad frequency range and amplitudes, EMG signals offer distinct advantages over other bio-signals such as ECG or EEG [28]. Key signal characteristics include amplitude (the potential difference between positive and negative peaks), phase, rise time, and overall duration. These features are essential for understanding motor unit action potentials (MUAPs), which represent the summation of individual fiber action potentials. Analyzing MUAPs provides insights into motor unit recruitment—both spatially

(adding more motor units) and temporally (increasing firing rates)—and helps diagnose neuropathic or myopathic conditions [28].

EMG signals can be obtained invasively or noninvasively (Table 1). Invasive techniques, such as needle electrodes, provide localized, high-resolution recordings but are more uncomfortable for patients. In contrast, sEMG uses electrodes placed on the skin to capture the aggregated activity of numerous motor units, making it especially useful for movement analysis and control. The quality of sEMG recordings depends critically on electrode placement and skin preparation. Guidelines like those from the SENIAM project help standardize electrode dimensions and inter-electrode distances to ensure reliable and comparable results. Proper skin cleaning (e.g., with alcohol) minimizes artifacts, and using multiple sensors can improve the decoding of movement intentions [30].

Table 1. Sensor technologies for neuromuscular prostheses.

Technology	Interface/Placement	Selectivity	Key Advantages	Key Limitations
sEMG	Ag/AgCl electrodes	Low–moderate (2–16 ch)	Non-invasive; quick fitting; low cost	Motion/sweat artifacts; crosstalk
High-density sEMG	Closely spaced electrodes on a flexible array	Moderate–high (up to 128 ch)	Spatial features; tolerant to small shifts	Larger data rate; alignment is still critical
Intramuscular EMG	IMES implants within residual muscles	High (1–32 ch per muscle)	Stable over time; immune to skin impedance	Surgical implantation; longevity unknown
Peripheral nerve cuffs	Multi-polar spiral or FINE cuffs around the nerve trunk	Moderate (6–16 contacts)	Bidirectional; chronic	Limited fascicle selectivity; fibrotic encapsulation
Intrafascicular arrays	Penetrating electrodes within fascicles	Very high (20–100 contacts)	Highest selectivity; natural sensory feedback	Highly invasive; fibrotic encapsulation
Osseointegration	Titanium percutaneous implant	Depends on embedded tech	Eliminates socket; rigid, infection-resistant feed	Requires two-stage surgery; skin seal upkeep

The core EMG signal is the result of electrical activity produced by skeletal muscles. This myoelectric data, stemming from motor neurons in the central nervous system, is interpreted through electrical readings [31]. Implantable sensors receive power and commands from an external telemetry controller that activates a coil integrated into a sleeve housing the prosthesis. Each implant functions as an independent differential amplifier connected to two electrodes, one at each end, to detect the muscle contraction signal. The primary drawback of these implantable myoelectric sensors (IMES) is their requirement for surgical implantation along with prolonged rehabilitation and training, which delays rapid results. Needle electrodes form another category, ideal for detecting changes in motor unit size and internal structure and for recording spontaneous fibrillation potentials. Although widely used, needle electrodes are perceived as more painful and yield less favorable results during forceful contractions. They are classified as monopolar, bipolar, or concentric. Monopolar electrodes tend to be cheaper and less painful, but they are electrically less stable and require a separate surface reference, unlike concentric designs [31–34].

Neural interfaces remain a long-standing challenge in achieving effective communication between prosthetic devices and the nervous system. Various approaches have been developed to both record and stimulate nerve fiber activity, with the primary differences being the degree of invasiveness and selectivity. A balance must be struck between increased invasiveness for higher selectivity and reduced invasiveness to minimize tissue damage and ensure chronic stability [35]. For viable restoration of motor and sensory functions, interfaces must provide both topographical and functional selectivity, which is aided by understanding the fascicular organization of peripheral nerves. Current invasive neural interfaces are divided into extraneural (e.g., cuff electrodes that wrap around nerves) and intraneural (intrafascicular) types. Cuff electrodes, while useful for activating or recording

from groups of nerve fibers, offer reduced selectivity since they engage superficial, myelinated fibers. Innovative cuff designs may improve this, but often require multiple cuffs to achieve adequate selectivity, thereby increasing implantation risks [35].

Implantable flexible nerve electrodes (IFNEs) have been applied in neurological disorders for recording and stimulating nerve signals. However, mechanical mismatches between rigid electrodes and soft tissues can lead to chronic inflammation and perineural fibrosis. Advances in electrode design have also led to improved selectivity for targeting specific axon populations. For example, innovative structures such as spiral or flat interface nerve electrodes (FINE) have been developed. FINEs use a non-conductive cuff containing multiple conductive elements to reshape the nerve under controlled pressure without compromising blood flow, thus enhancing stimulation, or recording selectivity. Designs integrating embedded multiplexers further reduce the number of wires needed, lowering the risk of neurovascular damage while maintaining a flexible, conformal interface with the nerve. Intrafascicular electrodes, such as the Longitudinal Intrafascicular Electrode (LIFE) and the Transverse Intrafascicular Multichannel Electrode (TIME), are implanted directly within nerve fascicles, offering higher signal-to-noise ratios, lower stimulation thresholds, and superior selectivity. While LIFEs are widely used, TIMEs—capable of interacting with multiple fascicles—have demonstrated improved target muscle selectivity with lower stimulation currents [35].

A recent innovation is the Neurotassel, an implantable neural probe developed in neuroscience to adapt effectively to brain tissue. It consists of a matrix of ultra-flexible microelectrodes with a high aspect ratio (i.e., very thin compared to their length). Notably, these microelectrodes self-organize into a fine fiber, causing minimal neuronal loss during chronic implantation and reducing inflammatory responses. Furthermore, the Neurotassel's scalable design allows its matrix to comprise multiple micro-filaments, forming a fine fiber that serves as a powerful and minimally invasive tool for recording neural activity [36].

In the broader technological context, flexible sensors have emerged as a significant innovation. By emulating human touch, these sensors enable intelligent systems to acquire detailed information on pressure, temperature, and texture. When integrated into wearable devices, they allow real-time monitoring of vital signs and physical activities, a crucial asset for healthcare, sports performance, and early disease detection. In robotics, flexible sensors enhance environmental awareness, enabling robots to handle delicate objects or navigate complex environments. Flexible sensors improve human-machine interfaces by providing intuitive tactile feedback that helps users control grip force and pressure precisely. In biomedical applications, they are used in implantable devices for continuous vital sign monitoring, although challenges remain regarding durability, sensitivity, and large-scale integration [35,36].

Non-invasive surface electrodes are typically made of conductive materials and are fixed to the skin. Gel electrodes are commonly used due to their lower impedance and noise, although they can dry out over time and may cause inflammatory reactions. In contrast, solid metal electrodes can be used over longer periods without signal degradation, though they are more expensive and may be more susceptible to movement artifacts. Electrodes may be passive (transmitting the signal for external amplification) or active (including built-in pre-amplifiers to reduce noise). Furthermore, EMG signals can be recorded in monopolar, bipolar, or multipolar configurations. Monopolar setups use a distant reference electrode and are suited for small muscles, while bipolar configurations capture differential signals from two nearby sites, reducing common noise. Accurate electrode placement requires anatomical knowledge since aspects such as amputation type or previous surgeries can affect the underlying anatomy and compromise pattern recognition [30,31].

3.4. Control Schemes in Myoelectric Prostheses

The ultimate goal in controlling robotic prostheses is to replicate the intuitive control of a natural hand, enabling uniform finger and wrist movements. Increasing degrees of freedom and achieving proportional control can enhance dexterity but may compromise the robustness of signal decoding [24]. One promising solution is to integrate robotic automation and shared control strategies, wherein user commands are combined with automated adjustments for grip force and grasp patterns [37–39].

Some studies highlight advancements in human-like robotic hands due to integration of tactile and force/torque sensors to enable stable grasping, dexterous manipulation and expressive gestures like playing musical instruments. Innovative designs range from adaptive actuation mechanisms and compliant low-cost servo hands to multimodal EMG-vision fusion for prosthetic control. These studies point out the urgent need for continued research and collaboration in developing advanced robotic and prosthetic hand technologies [40,41].

For transradial amputees, decoding peripheral nerve signals can enable highly accurate analysis of various grasp types. Advanced regenerative interfaces that reconnect severed nerves with muscle grafts act as bioamplifiers, enhancing both the specificity and long-term stability of the captured EMG signals. Some studies also research the promising benefits of new robust framework that integrates EMG and force myography (FMG) signals to enhance grasp type classification during dynamic arm position changes. Unlike EMG-only systems, the combined EMG-FMG approach demonstrates stable performance across varying window sizes, proving more robust than either signal alone. Overall, this framework has the potential to significantly enhance the robustness of natural grasp classification under dynamic conditions, offering a promising solution to bridge the gap between academic research and practical prosthetic applications [42].

Neural electrodes are an emerging technology that aims to mimic natural sensory processing at both cortical and peripheral levels. The current focus is on evolving from basic feedback to the induction of more complex, natural sensations through direct nerve stimulation [42–45]. The integration of additional sensors is increasingly important in bidirectional prosthetic interfaces. For example, the Mia hand developed by Controzzi incorporates force sensors at the fingertips to measure normal and tangential forces. Myoelectric control systems have significantly advanced, focusing on restoring motor function in amputees and individuals with motor impairments, though each control modality presents its advantages and limitations [24].

Real-life scenarios underline the importance of system robustness. For instance, the “hot coffee” problem illustrates that even a good laboratory success rate in grasp detection may not suffice when a slip could result in injury. Studies also indicate that a considerable proportion of both child and adult amputees abandon prostheses due to functional dissatisfaction. Control strategies for prostheses are generally classified into sequential and simultaneous approaches. Most upper-limb prostheses today use sequential control, which is further divided into pattern recognition (using classifiers to distinguish EMG signal patterns) and non-pattern recognition (predefined command sequences). While simultaneous control offers the potential for more natural multi-degree-of-freedom movement, it faces substantial challenges in integration and reliability [37,38].

A typical EMG control flow involves several key stages that enable effective signal processing and motor control (Table 2). The process begins with data segmentation, where raw EMG signals are divided into manageable segments for analysis, allowing relevant information to be isolated for further processing. Next, Finite State Machine (FSM) control is employed, transitioning between different states (such as opening or closing a prosthetic hand) once predefined signal thresholds are reached. This ensures accurate and context-appropriate responses to muscle activity [24].

Table 2. Control and machine-learning strategies for upper-limb prostheses.

Strategy	Input Signals	Latency (ms)	DOF Controlled	Strengths	Weaknesses
Threshold	1–2 sEMG channels	50–100	1 active DOF	Simple; robust	Poor dexterity
Finite-state machine	2–4 sEMG + triggers	80–120	1 (plus state switching)	Intuitive for fixed task sets	State cycling burdensome
Pattern-recognition	6–12 sEMG or HD-sEMG	100–150	Up to 6 discrete grips	Multi-grip	retraining
Proportional control	4–16 sEMG or intramuscular	50–100	2–3 continuous DOF	Natural speed/force scaling	crosstalk
Deep-learning	HD-sEMG, IMU or multi-modal	80–200	≥ 3 discrete + proportional	Learns features automatically	Compute cost
Hybrid (EMG + vision/EEG)	Multimodal (EMG + cameras, BCI)	120–250	Task-dependent	Context-aware; compensates EMG dropouts	Sensor fusion complexity

Following this, pattern recognition algorithms identify specific movement patterns within the segmented data, allowing the system to interpret user intent and facilitate intuitive control over assistive devices. The system then utilizes regression control, mapping signal amplitude continuously to motor output. This enables proportional control, where motor responses are directly correlated with the intensity of the user's muscle contractions. Additionally, postural control and ON/OFF strategies are applied using fixed signal thresholds to either maintain or initiate movement, ensuring that actions are triggered or halted as necessary. Finally, proportional control ensures that motor output is directly proportional to the detected signal amplitude, providing motor responses that are finely tuned to the user's muscle activity [43,44].

Pattern recognition methods analyze short-duration signal segments by extracting time-domain, frequency-domain, or time–frequency features. The classification process involves pre-processing (using techniques like PCA, LDA, or MDA), classification (via neural networks, Bayesian systems, or fuzzy logic), and post-processing to refine outputs. However, real-world variability, such as stress, electrode displacement, or changes in muscle morphology, can challenge the transfer of laboratory results to clinical practice [43].

Non-pattern recognition strategies are simpler to implement but offer limited functionality. Often, a hybrid approach that combines both methods yields the best results. For example, proportional myoelectric control directly correlates EMG amplitude with motor output, allowing for continuous and intuitive control, especially in systems requiring fine movements. In contrast, FSM control uses a limited number of predefined states and is effective for a fixed set of movements but lacks flexibility. Threshold-based control, which activates responses once a signal surpasses a set level, is also common but may not provide the smooth adjustment seen with proportional methods. Simultaneous control is regarded as the most natural approach, enabling the prosthesis to perform multiple functions at once. Algorithms based on regression can decode concurrent commands, such as differentiating between wrist rotation and hand opening, to enhance control flexibility. Hybrid control systems may also integrate additional sensor inputs, such as pressure sensors, accelerometers, EEG signals, or computer vision, to further improve accuracy and robustness [24,41].

Commercial systems often use a reduced number of electrodes and predominantly sequential control strategies, sometimes augmented with external control methods (Table 2). Examples include the i-limb Quantum, which uses gesture control, and the Bebionic system, which incorporates finger-position encoders to prevent object drops. Some manufacturers are exploring the simultaneous activation of two electrodes to switch between different grasp types, as seen with the Michelangelo hand by Ottobock, although this method can be less intuitive. Direct control via pattern recognition is employed in systems like the

Coapt and Ottobock Myo Plus, which achieve high classification rates with 2–12 sEMG electrodes [1,24,37,38].

Conventional approaches such as FSM and pattern recognition have been enhanced by machine learning techniques, particularly HD-sEMG paired with deep learning algorithms. Recent studies using sequential temporal regression and reinforcement learning have enabled natural control of wrist and finger movements, increasing system responsiveness and dexterity. Incorporating incremental learning strategies allows control systems to adapt to long-term variations in EMG signal patterns, supporting sustained, user-specific calibration over time. Hybrid control strategies that combine EMG, force, and motion sensors are increasingly being tested in real-time scenarios, enhancing system robustness and reducing command misclassification in dynamic conditions [43,44].

Advancements in prosthetic technology are continuously striving to create devices that are more intuitive, functional, and comfortable for the user. A key focus is on improving the human–machine interface, which allows for more natural and seamless control of the prosthetic limb. One promising approach involves the surgical creation of electro-neuromuscular constructs, where nerves and muscles are restructured to create new control signals that can be more easily interpreted by the prosthetic device. This method aims to provide a more direct and reliable link between the user’s brain and the artificial limb [45].

Researchers are exploring various sensor technologies to enhance the control of lower prosthetics. For instance, studies have investigated using brain signals, specifically those from electroencephalography (EEG) and functional near-infrared spectroscopy (fNIRS), to detect a user’s intent to move [46]. This approach allows for a direct neural connection to the prosthetic, bypassing the need for muscle signals entirely. In an equivalent way, other research has focused on the use of continuous neural control to restore a more biomimetic gait after amputation, enabling the prosthetic to move in a way that closely mimics natural human walking [47]. Furthermore, preserving the biomechanical relationship between opposing muscle groups, known as agonist–antagonist muscle strain, in the residual limb is crucial for maintaining motor control and a sense of limb perception [48].

Another critical area of development is the integration of advanced sensors directly into prosthetic devices and their interfaces. Flexible and stretchable sensors are becoming increasingly important for creating more comfortable and functional prosthetic skins [49]. The application of flexible sensors in human–machine interfaces is a rapidly growing field, with these sensors being used to detect various inputs such as bending, pressure, and tactile feedback [50–53]. For example, wearable sensors can be used to recognize a user’s intent for adaptive control of intelligent ankle–foot prosthetics, allowing the device to adjust its function in real-time based on the wearer’s movements [51]. The use of artificial intelligence (AI) is also enabling more intuitive control by processing complex nerve signals in real-time, translating a user’s intent into precise prosthetic movements [54].

3.5. Sensory Feedback Systems and Human–Machine Interfaces

Overall, decades of research in robotic prostheses have yielded both invasive and non-invasive solutions integrated with the human body. The efficacy of these interfaces is primarily measured by their selectivity, signal quality, and stability, with ongoing research focusing on improving electrode impedance and filtering methods.

There is still no consensus on the best non-invasive approaches for providing sensory feedback in prostheses. Transcutaneous electrical nerve stimulation (TENS) has been explored as a non-invasive option to elicit natural-like sensations in amputees, but implantable electrodes, especially intraneural interfaces, have shown markedly superior outcomes. These invasive methods not only improve the functionality (e.g., enhanced

texture and shape discrimination) but also reduce the cognitive load for the user, delivering a feedback experience that feels nearly natural [42].

The challenges in the motor and sensory domains are distinct. On the motor side, robust signal decoding (often utilizing deep learning techniques) is essential, whereas in the sensory domain, implantable electrodes are emerging as the most promising solution, even though advances have primarily focused on tactile feedback. Thermal perception, mediated by fast-conducting A δ fibers and slower C fibers, remains particularly challenging; current electrodes cannot effectively target these minute fibers. Future developments may involve higher-selectivity electrodes or sensory remapping techniques to simulate temperature feedback [40,47].

Recent advances in surgically embedded systems led to the development of the Agonist-Antagonist Myoneural Interface (AMI). This interface enables natural neural feedback and improved proprioception through physiological muscle reconnection in the residual limb, facilitating more adaptive gait and balance control in lower-limb amputees [40]. Flexible and stretchable sensors continue to play a leading role in human-machine interfaces, particularly for continuous monitoring and adaptive control.

Recent clinical trials have demonstrated that patients with below-knee amputations who undergo AMI surgery walk on average 41% faster and navigate stairs and obstacles with reduced cognitive load compared to conventional amputees [48]. The foundational AMI framework, covering conceptual design, surgical technique, and preclinical efficacy, has highlighted the improvements in pain mitigation, phantom-limb range of motion, and controllability through preserved fascicle dynamics and central proprioceptive pathways. By specifically preserving agonist-antagonist muscle relationships, the AMI maintains spindle-fiber and Golgi-tendon-organ afferents that underpin natural proprioception. Neuroimaging studies further confirm that AMI recipients exhibit proprioceptive response strengths in the primary sensorimotor cortex comparable to intact-limb controls when proprioceptors are passively stimulated [48].

Concurrently, materials innovations have produced stretchable, self-powered sensors that conform to residual-limb anatomies. Triboelectric nanogenerators can harvest biomechanical energy to power continuous monitoring without batteries, converting gait-induced motions into electrical signals with high efficiency [49–51]. Aerosol jet-printed silver strain gauges on low-melting-point plastic substrates enable custom geometries on non-planar surfaces, achieving sub-millimeter resolution and robust cycling performance [52]. Artificial intelligence-assisted mechanoluminescent films integrate deep-learning color-decoding pipelines directly on flexible substrates, transducing complex deformation patterns in real time for haptic feedback [50,51]. High-density, dry-electrode EMG arrays fabricated on flexible PCBs map muscle activity across broad areas with signal quality matching rigid grids, facilitating both clinical gait analyses and at-home training [52–54].

In summary, advanced control architectures fuse multimodal biosignals with machine learning to deliver intuitive and adaptive behavior. Continuous EMG controllers modulate ankle and knee actuators in real time, achieving smooth postural adjustments and load-transfer improvements substantiated in direct-control studies. Hybrid EEG-fNIRS brain-computer interfaces decode cortical engagement to adjust assistance levels dynamically, enhancing user experience and reducing mental workload. At the peripheral-nerve level, recurrent neural network decoders translate multichannel nerve recordings into movement commands with high accuracy and stable performance.

3.5.1. Targeted Muscle Reinnervation

Targeted muscle reinnervation (TMR) is an innovative neural interface procedure designed to improve control of upper-limb prostheses. In TMR, nerves that once con-

trolled the amputated limb are surgically transferred to remaining intact muscles. After reinnervation, these muscles generate electrical signals in response to the original motor commands, enabling a more intuitive and natural control of the prosthesis while enhancing sensory feedback [55,56]. Severed nerves can remain viable for regeneration even years after amputation. Ensuring complete denervation of the target muscle is crucial so that the resulting EMG signals originate solely from the transferred nerves, thereby avoiding mixed signals that could complicate prosthetic control. For optimal independent control of various movements, reinnervated muscle segments must produce distinct EMG signals; minimizing the subcutaneous fat layer can further enhance signal amplitude and differentiation [57–60].

This surgical procedure has rapidly evolved into a multifaceted neural-interface strategy that not only enhances intuitive myoelectric control of upper-limb prostheses but also provides durable relief from neuroma and phantom-limb pain. Originally developed to reroute residual peripheral nerves into expendable muscles, TMR creates new myoelectric signal sources that correspond directly to motor intentions, yielding markedly improved prosthetic dexterity and user satisfaction. Over the years, indications for TMR have broadened in cases of neuroma pain after primary amputation, with a stable pain reduction and functional gains postoperatively [61–63].

At its core, TMR entails coapting a transected mixed motor or sensory nerve into the motor branch of a denervated target muscle, harnessing the muscle's trophic milieu to guide axonal regrowth and reestablish afferent–efferent loops [60]. Ideal donor nerves include those that previously controlled wrist or finger movement; common recipients are segments of the pectoralis or biceps in transhumeral amputations and the brachi-radialis or flexor digitorum in more distal levels [60]. Precise dissection and intraoperative nerve stimulation ensure complete denervation of the target muscle, preventing signal contamination by residual native fibers and maximizing the specificity of the regenerated EMG signals.

When an amputation occurs, the protective sheath of a nerve is damaged, causing the axons inside to grow chaotically into the surrounding scar tissue. This disorganized growth frequently results in the development of painful neuromas. To address this, reinnervation techniques such as TMR and Regenerative Peripheral Nerve Interface (RPNI) are employed. Both TMR and RPNI redirect the severed nerve into a segment of healthy, denervated muscle. This muscle segment acts as a conduit, providing a more organized pathway for the nerve to regrow and find new targets (motor end plates) for reinnervation. By preventing the nerve from sprouting randomly into fibrotic tissue, these procedures lower the risk of neuroma formation and help reduce post-amputation pain. They are both equally effective in pain relief, though TMR may confer superior control signals when ample muscle targets exist. Combining TMR with RPNI (“TMRp-ni”) may further optimize neuroma prevention, but more studies are needed [64].

The spatially distinct activation patterns can be exploited by HD-sEMG for high-fidelity control. Unlike conventional systems, HD-sEMG sensors can be seamlessly integrated into clothing, creating a non-intrusive myoelectric interface that can be used throughout a user's daily routine. A recent study introduced a user-friendly HD-sEMG device designed to interface with an 8-degree-of-freedom mobile manipulator. The research evaluated the use of real-time myoelectric gesture recognition to enable precise control over a range of self-care and household tasks. The study demonstrated the potential of this wearable system for controlling mobile manipulators within a home environment, highlighting its efficacy in physically assistive roles [65].

Modern TMR-based controllers integrate EMG pattern recognition algorithms, such as support-vector machines and convolutional neural networks, to map complex activation patterns to movements in real time. Hybrid approaches that fuse EMG, inertial

measurement unit (IMU), and pressure sensor data achieve >95% classification accuracy across multiple tasks. Inspired by multimodal strategies used in fatigue detection systems, where EMG and IMU signals are fused via deep learning architectures for robust real-time classification, emerging bidirectional interfaces incorporate afferent stimulation to the reinnervated muscles, closing the loop for sensory feedback and further refining motor control [66].

Ideal TMR candidates are those with residual limb lengths permitting effective nerve transfer, typically mid-humeral or more proximal for upper-limb amputations, and sufficient muscle mass to host reinnervation sites. Each candidate is individually evaluated based on the quality of soft tissues, nerve function, and their commitment to the rehabilitation process. Younger patients often exhibit better outcomes due to the higher regenerative capacity of nerves. The surgical procedure is facilitated by flexible soft tissues and the absence of prior skin grafts, although previous surgeries do not automatically preclude TMR. Longer residual limbs allow for more extensive nerve transfers. Contraindications include proximal brachial plexus injuries or severe soft tissue scarring that precludes adequate nerve mobilization. Postoperative protocols begin with an immobilization period, followed by progressive physical rehabilitation training [55,56].

While generally safe, TMR carries surgical risks such as hematoma, infection, and transient sensory deficits in donor-nerve distribution [61]. Surgical risks include hemorrhage, infection, paresthesia, loss of sensitivity, exacerbation of phantom pain, and possibly the need for additional interventions. Postoperative protocols typically require an initial prosthesis-free period to ensure proper wound healing, and patients are advised that full reinnervation may take several months. Often, amputation is combined with neurectomies to manage neuroma formation; physical examination, supported by simple radiographs, is critical to assess residual bone length and soft tissue conditions of the patient. Phantom pain may transiently worsen during the reinnervation phase but typically subsides as new afferent pathways mature over the following months [57].

TMRpni addresses cases with limited local targets, ensuring no regenerating axon is left without a functional end organ. Additionally, fully implantable systems that combine intramuscular electrodes with external transmitters could, in the future, demonstrate stable signal quality and user acceptance [61].

Ongoing advances aim to miniaturize implantable electronics, integrate direct neural recording from mixed sensory–motor fascicles, and leverage AI for adaptive signal decoding. As surgical techniques refine and closed-loop sensory feedback systems mature, TMR stands poised to deliver truly biomimetic experiences, restoring not just function but naturalistic embodiment.

3.5.2. The Role of Osseointegration

Amputation treatments are often performed under emergency conditions, prioritizing patient survival over long-term esthetics or functionality. Consequently, many patients are left with a residual stump that is both functionally poor and painful, which can compromise subsequent outcomes. In recent years, significant investments have been made in innovative prostheses and control strategies, making the functional preparation of the stump crucial for advanced robotic prostheses. Traditional surgical techniques aim to preserve as much limb length as possible to form a cushion for standard prosthetic sockets. However, such approaches result in a passive stump that is not ideal for the dynamic control required by robotic devices [67–69].

Discomfort and poor socket fit are among the primary reasons for prosthetic device rejection. To mitigate these issues, the interface must provide a comfortable, stable, and individually tailored fit that protects the residual limb. Excessive and continuous pressure,

as well as localized pressure peaks, can lead to pain, edema, eczema, and even neuroma formation. These factors adversely affect both control and overall quality of life [68,69]. Different fixation systems exist for prostheses, which can be categorized as invasive or non-invasive. Non-invasive methods include the use of sleeves or harnesses, vacuum systems, and systems like the BOA Fit, which offer precise, adjustable fittings. However, non-invasive systems may not always provide optimal comfort or stability. Osseointegration represents a more invasive but transformative approach that involves implanting a titanium screw into the residual bone, allowing the prosthesis to be directly affixed. The procedure typically requires two surgical stages: the first to insert the implant and the second to attach the percutaneous abutment. Although osseointegration offers better esthetics and a greater sense of security, it carries risks such as infection and may necessitate multiple interventions [70].

The pioneering work of Ingvar Brånemark, who demonstrated stable titanium–bone integration in rabbit tibiae and later in dental implants, laid the foundation for modern osseointegration. Today, the technique is applied in hip arthroplasty, dental implants, craniofacial reconstruction, and the direct anchoring of prosthetic devices. This method not only eliminates the complications associated with conventional socket interfaces but also offers benefits such as reduced energy expenditure, improved range of motion, and indirect sensory feedback—referred to as osseoperception, which enhances the user’s ability to perceive vibratory sensations through bone [68–71]. The success of osseointegration depends on the bone’s capacity to incorporate a non-biological implant, initiating a process that starts with hematoma formation and progresses through lamellar bone formation and subsequent remodeling. Patient selection is critical; candidates must have adequate bone stock, adhere to rehabilitation protocols, and be in good overall health. Preoperative planning, often involving CT scans to assess bone quality and dimensions, is essential to minimize risks and ensure implant stability [68–72].

Several osseointegrated systems have been developed for amputees. For example, the OPRA (Osseointegrated Prostheses for the Rehabilitation of Amputees) system, introduced in 1998 in Sweden, uses an intramedullary implant with an external abutment for prosthesis attachment. A modified version, e-OPRA, incorporates a bidirectional communication gateway between implanted neuromuscular electrodes and the external device, thereby enabling both motor control and sensory feedback [69–72]. Recent advances include both intramedullary and extramedullary designs. Extramedullary devices (Figure 3), which fix the prosthesis to the external surface of the bone, may offer lower risks of bone resorption and infection, increased stability, and enhanced bone regeneration compared to their intramedullary counterparts [68,69].




Figure 3. Schematic diagram of the novel external cap implant and the internal columnar implant. (A) Novel design of the extramedullary cap-shaped implant; (B) current design of the intramedullary implanted columnar implant (adapted with permission from [71]).

Despite improvements in implant design, surgical techniques, and rehabilitation protocols, infection remains a major concern. Osseointegration surgery is highly demanding and should be performed only by experienced orthopedic surgeons, with rigorous patient selection based on clinical criteria. Studies indicate that osseointegration improves comfort, increases range of motion, and enhances mobility compared to conventional socket prostheses. A key benefit is osseoperception, the enhanced ability to discern vibratory sensations via bone conduction, which contributes to a more natural incorporation of the prosthesis into the user's body image. This integration supports not only better motor control but also improved psychosocial well-being [67–69].

3.5.3. Brain–Computer Interfaces

Brain–computer interfaces (BCIs) enable direct communication between the brain and external devices by translating neural activity into control commands. These systems are particularly valuable for individuals with neurological disorders, paralysis, or severe motor impairments, allowing them to operate assistive technologies or even participate in consumer applications through their brain signals. Most non-invasive BCIs acquire signals via electroencephalography (EEG), prized for its high temporal resolution, low cost, and safety. However, EEG is inherently nonstationary and prone to noise and artifacts, which demand rapid and precise extraction of relevant features for accurate decoding [73–75].

A foundational paradigm in EEG-based BCIs is motor imagery, where users imagine specific movements without physical execution. The standard processing pipeline comprises five stages: acquisition, pre-processing, feature extraction, feature selection, and classification. Pre-processing tackles noise and artifacts, while feature extraction methods, such as wavelet transforms, short-time Fourier transform (STFT), common spatial pattern (CSP), and filter bank CSP (FBCSP), capture time, frequency, and spatial characteristics. Extracted features are then fed into machine learning classifiers, including convolutional neural networks (CNNs) and sparse autoencoders (SAE), to distinguish between imagined actions [76].

Despite the effectiveness of this pipeline, conventional workflows rely heavily on manual preprocessing and handcrafted feature engineering, creating barriers to real-time and clinical applications. To address this, recent frameworks employ end-to-end CNN architectures that operate directly on raw EEG data. By learning hierarchical representations autonomously, these models classify multi-class motor imagery, such as hand, foot, or tongue movements, without explicit filtering, while integrating skip connections and adaptive pooling to improve resilience against inter- and intra-subject variability [76]. Such generalized models not only maintain competitive accuracy across diverse datasets but also reduce computational overhead, facilitating deployment on embedded or mobile platforms.

In neurorehabilitation, EEG-BCI systems harness motor imagery combined with feedback modalities to drive neural plasticity and restore motor function. A growing body of research highlights promising outcomes for upper extremity recovery, often employing robotic orthoses or virtual training scenarios that translate neural commands into assisted movements. The integration of virtual reality (VR) technologies enhances patient engagement and adherence; while non-immersive VR setups dominate clinical studies, fully immersive environments show potential for even greater therapeutic efficacy by creating more compelling sensorimotor experiences [76].

Lower-limb BCI applications are less prevalent but have demonstrated feasibility in tasks like foot movement control and wheelchair navigation. Innovative approaches incorporate haptic feedback to complete the sensory-motor loop, providing users with tactile cues that reinforce intended movements. To date, unified frameworks addressing

both upper and lower extremities remain rare, representing a significant opportunity for future development in holistic rehabilitation solutions [76].

Despite these advances, widespread clinical adoption of EEG-BCIs faces several hurdles: a lack of standardized protocols, extensive calibration and training requirements, and limited validation across heterogeneous patient populations. Overcoming these challenges will depend on advancements in signal-processing algorithms, the creation of more affordable and user-friendly hardware, and the integration of BCIs with complementary rehabilitation technologies. For instance, coupling BCIs with functional electrical stimulation can enhance motor recovery by directly activating muscles in synchrony with neural intent. Moreover, adaptive BCI systems that personalize therapy based on real-time neural feedback promise to optimize rehabilitation outcomes and reduce the burden of manual adjustment [76].

While EEG-based systems offer safety and convenience, they are limited in spatial resolution and signal fidelity. Invasive BCIs—using implanted microelectrode arrays and wireless systems—seek to overcome these constraints by providing high-resolution, stable recordings over extended periods. Biocompatible materials minimize immune responses and scar tissue formation, ensuring long-term device viability [75,76].

A leading example is Neuralink's N1 chip, designed to capture high-bandwidth neural data through 3072 analog pixels. Each pixel contains a low-noise amplifier, achieving a system gain to match the resolution of on-chip Analog-to-Digital Converters (ADCs). Data from the amplifiers feed into ADCs, are multiplexed, and then transmitted serially. The SoC incorporates sophisticated power management strategies, backpressure buffering to alleviate data congestion and event-based sampling to conserve resources. A JTAG interface based on Boundary Scan and a dedicated test access port (TAP) controller enables configuration, diagnostics, and on-chip testing without invasive probing [77]. This architecture exemplifies the balance between high-density neural recording, low power consumption and robust data handling required for next-generation BCI applications.

Together, these technological innovations, spanning non-invasive end-to-end decoding frameworks, immersive neurorehabilitation platforms and advanced implantable hardware, are driving a paradigm shift toward more accessible, scalable, and effective BCIs for both clinical and consumer markets.

3.5.4. Electromyographic Signal Acquisition and Processing

The use of EMG in robotic devices, especially in prostheses, has been a strong area of research, given that it is crucial for developing more natural and effective human-machine interfaces. In this type of device, the components' size and portability have implications for their operational effectiveness and efficiency. As previously mentioned, the data captured by electromyographic electrodes vary depending on the nature of the muscle contraction, the size of the analyzed muscle, and other technical or methodological variables [78].

One of the invasive electromyography techniques consists of inserting a needle electrode through the skin directly into the muscle. Although it can provide a localized and high-resolution description of the electrical activity of the muscle, it is a somewhat painful experience for the patient. The pick-up area of the needle electrode depends on its specific design and may include only one or two muscle fibers. The monopolar needle electrode requires a reference electrode placed away from the needle insertion point in an electrically neutral location (such as over a bone). The concentric needle electrode avoids the reference electrode by referencing the active surface of the electrode to the cannula of the needle. Both types of needle electrodes are used in clinical settings. By placing several needle electrodes in various locations, a more detailed description of the electrical activity of a larger muscle group can be obtained [78].

An additional development of the surface electrode technique is the linear electrode array, designed to provide a spatial description of myoelectric activity. The resulting multichannel EMG signal from several electrodes placed equidistantly along the direction of the muscle allows studying the generation and extinction of action potentials and estimating the conduction velocity. This method can be recorded at lower sampling rates than needle EMG, as the intervening tissue between the motor units and the surface electrode acts as a low-pass filter of the electrical signal. It also has most of its spectral power below 400–500 Hz, which implies that a sampling rate of 1 kHz or higher is necessary. For needle EMG, the sampling rate should be chosen so that different waveforms of MUAPs, which can contain frequencies up to 10 kHz, are accurately reproduced; therefore, a sampling rate of 50 kHz is often used [78].

In sEMG, the size of the electrode is determined by the conductive area and orientation relative to the muscle fibers. The dimensions of the electrode should remain within a range that preserves sensitivity to signal variations. Electrodes of varied sizes and shapes are commercially available, with circular and rectangular formats being the most common. The spacing between bipolar electrodes needs to be chosen according to the muscle being studied, since excessive separation can introduce instability. In smaller muscles, the distance between electrodes should not exceed one-quarter of the muscle fiber length, thus avoiding unstable readings caused by tendon effects. The material that constitutes the contact area with the skin must ensure good electrode–skin contact and consistent low impedance. Pre-gelled Ag/AgCl electrodes are generally the most recommended [30,31,78].

Several steps must be followed when using sEMG. The first should be skin preparation with good electrode–skin contact to ensure the best signal acquisition, minimizing electrical artifacts and interference. It is recommended to clean the skin with alcohol, which disinfects and evaporates quickly, preparing the skin for sensor placement. The use of multiple EMG sensors to decode movement intentions is also vital. Traditionally, sensors were placed between the muscle innervation areas and the tendinous zones, which are known to offer better quality signals. However, the complexity of the overlapping muscles in the human forearms obscures the median muscle lines on the skin surface, which makes it difficult to capture electromyographic signals [78].

Several studies demonstrate that small deviations of an electrode over the muscle innervation area can significantly alter the signal amplitude. Furthermore, the process of positioning the sensors requires careful customization, with significant inter-individual variability due to the anthropometric characteristics of each individual. Creating a functional myoelectric interface requires considerable adjustment effort, especially when increasing the number of sensors. The signal-to-noise ratio can be improved by positioning the electrodes as close as possible to the signal source. The reference electrode should be placed in a location where the risk of signal disturbances is minimal, preferably in electrically inactive tissues. Depending on the muscle under analysis, it is common to position the reference electrode on the wrist, elbow, chest, tibia, ankle, or the C7 vertebra [78–80].

The manipulation capabilities of the human hand are governed by the activation patterns of the numerous muscles of the forearm and hand. Some of these muscles are superficial and easily accessible, but most are in deeper compartments. For this reason, there are only a small number of superficial forearm muscles from which sEMG signals can be reliably detected. In addition, the relatively small size and proximity of many muscles are also a concern. Once acquired, EMG signals undergo several processing steps: amplification, filtering, and rectification, to remove noise and extract useful features. Usual challenges include movement artifacts, electrocardiographic interference (especially in trunk or neck recordings), and electromagnetic interference from nearby devices. Advanced

noise-cancellation and narrowband filtering techniques, like those used in ECG processing, are applied to preserve the EMG spectral content [78–81].

In terms of control, EMG signals from electrodes on the residual limb are analyzed to extract control commands. Simple controllers may use signal amplitude to differentiate functions (such as opening/closing a hand), while multifaceted prostheses combine inputs from multiple electrodes and employ advanced algorithms to capture transient signal patterns at the onset of rapid contractions. Real-time processing is critical, as response times must be short enough to prevent perceptible delays [78–82]. The amplitude of the sEMG signal is a key parameter for clinical and control applications and can be estimated using statistical techniques (e.g., standard deviation or dispersion estimates) over sliding time windows. Due to the stochastic nature of sEMG (resulting from random motor unit recruitment, variable firing rates, and external interference), advanced statistical tools and models (often based on linear time-invariant, LTI, filter concepts) are employed. Although these models, such as the filtered noise model, simplify the complex relationship between neural activity and the EMG signal, they face limitations due to nonlinear tissue properties and variable electrode–tissue interfaces [78–82].

3.5.5. General Steps in EMG Signal Processing

The acquisition and processing of electromyographic (EMG) signals for controlling a bionic hand involve several sequential stages that transform raw muscle signals into actionable commands. The process begins at the signal source, typically the forearm, where muscle contractions generate interpretable signals. Strategically placed electrodes capture these signals, which are then transmitted to a data acquisition device. Here, analog signals are converted into discrete digital values, a step that is crucial for accurate digital processing [83,84].

The digitized signals undergo pre-processing to remove noise and artifacts, ensuring that only high-quality, relevant data are analyzed. This includes segmenting the continuous signal into temporal windows, which is essential for subsequent feature extraction. In this phase, key metrics such as amplitude, frequency, and other parameters are extracted to differentiate between distinct types of muscular movements. These extracted features are then fed into classification algorithms that interpret the user's intended movement and convert it into a specific command for the prosthetic hand. The movement controller receives this command, activates the corresponding motors via the servo driver, and simultaneously sends performance and position feedback back into the system for real-time adjustments, ensuring a more intuitive and responsive control experience [83,84].

3.5.6. Artificial Intelligence in EMG Signal Inference

Recent advances in EMG pattern recognition, bolstered by deep learning techniques, promise significant improvements in the functionality and versatility of control systems. Deep learning models enable the discrimination of a wide array of movement intentions, supporting the control of prostheses with multiple degrees of freedom. However, deploying these advanced techniques in real-world environments remains challenging due to individual variability and the dynamic nature of everyday activities. Understanding artificial intelligence (AI) in biomedical applications begins with grasping the fundamentals of deep learning.

Discriminative models classify input data into predefined labels by learning discriminative features through adaptive nonlinear transformations and probabilistic predictions [85]. Classification is achieved using a softmax layer combined with one-hot label encoding. Common discriminative architectures include Multilayer Perceptrons (MLPs), Recurrent Neural Networks (RNNs), and Convolutional Neural Networks (CNNs), along

with their variants [85–87]. The softmax layer is an activation function used in the final layer of classification models. It converts output values (logits) into probabilities, ensuring each value lies between 0 and 1 and that their sum equals 1. During training, true labels are converted into one-hot vectors, where only the position corresponding to the true class is set to 1. This combination allows error calculation and network weight adjustments. Softmax ensures interpretable probabilities, while one-hot encoding enables direct numerical comparison between predictions and true labels, critical for gradient-based learning [85–87].

One-hot encoding represents categorical labels as binary vectors, where only one position is “1” (active) and others are “0”. For example, categories 0, 1, 2, and 3 are encoded as (1, 0, 0, 0), (0, 1, 0, 0), (0, 0, 1, 0), and (0, 0, 0, 1), respectively. This method avoids implicit ordinal relationships between categories. However, for large datasets, it increases dimensionality and storage inefficiency due to sparse vectors (e.g., 100 categories require 100-dimensional vectors) [85,86].

In contrast, representational models like autoencoders (AE), restricted Boltzmann machines (RBM) and deep belief networks (DBN) are used to autonomously extract the most salient features from the data without manual intervention. Although these models excel at capturing underlying data distributions, they often require a combination with discriminative techniques to perform classification tasks effectively. Generative models such as variational autoencoders (VAEs) and generative adversarial networks (GANs) can also synthesize new data samples, which is especially useful for augmenting limited training datasets [85,86].

Hybrid deep learning models integrate both representational and discriminative approaches to harness the strengths of each, achieving robust and flexible performance in EMG signal interpretation. Notably, recent research using 3D Convolutional Neural Networks (3DCNNs) has shown promise in extracting muscle synergy features by converting one-dimensional sEMG signals into two-dimensional anatomical representations and then compiling these into three-dimensional segments. This approach captures both spatial and temporal components of muscle activity, outperforming traditional CNNs and autoencoders in terms of accuracy and efficiency while also reducing training times [85,86].

Representational deep learning models autonomously learn pure, high-level features from raw data without manual feature engineering. Unlike traditional machine learning, these models identify underlying patterns through hierarchical representations. Examples include autoencoders (AEs), Restricted Boltzmann Machines (RBMs) and Deep Belief Networks (DBNs). AEs compress input data into latent representations and reconstruct them, while RBMs and DBNs learn probability distributions and hierarchical features. However, these models focus solely on feature extraction, not classification. Generative models learn the joint probability distribution of input data and labels, enabling synthetic data generation. In biosignal processing, they are used to augment training datasets. Key examples include VAE (which learns probabilistic latent spaces to generate data) and GAN (which uses a generator-discriminator framework to produce realistic synthetic samples). Unlike conventional AEs, VAEs prioritize generation over reconstruction [85,86].

Hybrid models integrate multiple algorithm types, such as combining representational models for feature extraction with discriminative models for classification. A common example pairs DBNs with softmax layers for decision-making. True hybrid systems explicitly merge distinct paradigms (e.g., GANs with CNNs for image processing or VAEs with RNNs for sequential data). These models leverage the strengths of different approaches to tackle complex challenges [85,86].

Advanced algorithms and high-quality open-source datasets are enhancing EMG pattern recognition, attracting more researchers to the field. Deep learning techniques now

outperform conventional methods in grasp classification and movement detection. However, the need for large training datasets remains a challenge. Studies using surface EMG (sEMG) to predict continuous wrist movements are particularly relevant for post-stroke rehabilitation, where intention prediction can personalize robotic orthosis assistance [85,86].

Recent studies propose 3DCNNs to extract muscle synergies from sEMG signals. These models transform 1D sEMG data into 2D anatomical maps, compiled into 3D segments capturing spatial-temporal dynamics. 3DCNNs outperform traditional CNNs and autoencoders in accuracy and efficiency while reducing training time. Limitations include validation gaps in complex movements and physically impaired populations. Future work aims to expand predictions to multi-degree-of-freedom motions and integrate biomechanical constraints [87].

This integration of cutting-edge algorithms (Figure 4) and open datasets is driving innovation in EMG-based rehabilitation technologies, paving the way for personalized and efficient therapeutic solutions.




Figure 4. Main categories of deep learning model classification. Adapted with permission from [85].

4. Challenges in Robotic Hand Prostheses

The excessive cost of robotic prostheses remains a significant barrier to broader adoption. With most advanced solutions priced between USD 10,000 and USD 20,000, many researchers advocate for more affordable alternatives, especially for developing countries. Open-source prostheses present a promising alternative by drastically reducing development and distribution costs. Although rigid architecture currently prevails, there is growing interest in developing more flexible systems that facilitate safe interactions. These more compliant architectures combine simple actuation with the performance of a highly adaptable hand.

The advancements in neuromuscular robotic prostheses have been remarkable, integrating EMG signal processing, machine learning algorithms, and BCIs to enhance prosthetic control and functionality. The studies reviewed demonstrate significant progress in decoding neural and muscular signals to achieve more intuitive and precise movements. However, despite these advancements, several challenges remain.

One of the key findings is the necessity of robust signal acquisition and processing methods to ensure real-time and accurate control. Many studies highlight the importance of minimizing signal noise and improving electrode placement strategies to enhance the reliability of EMG-based control systems. Additionally, the integration of AI has shown great potential in adapting responses to user-specific needs, but further optimization is required to refine learning algorithms and reduce computational latency. Another critical aspect discussed in recent research is the bidirectional communication between prostheses and the user. Sensory feedback mechanisms, such as electrical stimulation and haptic

interfaces, are being developed to restore a sense of touch and proprioception in amputees. While promising, these technologies still face hurdles in terms of biocompatibility, long-term stability, and user adaptation. Moreover, current designs of robotic prostheses are often constrained by power consumption and portability. The miniaturization of electronic components and the development of energy-efficient actuation systems are essential to improve the wearability and usability of these devices. Emerging energy harvesting techniques and low-power computing architectures could provide viable solutions for future prosthetic systems.

The recent advances in neuromuscular robotic prostheses illustrate a compelling integration of signal processing, sensor technology, and intelligent control systems. One of the major engineering challenges lies in the design of analog front ends and digital conversion circuits that can reliably amplify and convert low-amplitude bioelectrical signals, typically in the microvolt range, into robust digital data without introducing excessive noise. The selection of high-quality electrodes, proper amplifier design, and low-noise ADCs are critical to ensuring that even the faintest signals are preserved for further processing. Modern signal processing techniques have evolved to include sophisticated filtering, segmentation, and feature extraction algorithms that are implemented on embedded platforms. These methods not only remove artifacts from the raw signal but also extract meaningful parameters such as signal amplitude, frequency characteristics, and time-domain features. This is crucial for the subsequent steps where pattern recognition algorithms based on deep learning techniques interpret these features to infer user intention.

Control strategies have also evolved with the integration of both classic control systems and data-driven approaches that exploit the power of deep neural networks. For instance, CNNs are employed to decipher complex patterns from EMG data, enabling a more natural and multi-DOF control of prosthetic limbs. This shift toward incorporating AI has enabled prosthetic systems to continuously learn and adapt to the user's specific muscular and neural signatures, thereby enhancing both precision and responsiveness. Another exciting development is the integration of BCIs with traditional EMG-based systems. This hybrid approach combines complementary data streams from peripheral muscles and central nervous system activities, which can significantly improve system robustness and response accuracy. However, achieving seamless, bidirectional communication between these diverse sensors remains a challenging task, particularly in terms of synchronizing high-bandwidth neural data with lower-frequency muscle signals. Addressing issues such as latency, data fusion, and real-time processing remains paramount for achieving the objective of natural and intuitive control.

The hardware and embedded system design aspects are equally critical. Engineers must optimize the trade-offs between power consumption, processing speed, and overall device miniaturization. Advances in semiconductor technology and low-power design have enabled the development of highly integrated SoCs that aggregate the necessary analog and digital functionalities into a compact form factor. Wireless communication modules, efficient battery management, and the integration of energy harvesting techniques are emerging as key enablers for the next generation of wearable prostheses.

From a clinical perspective, there is an increasing emphasis on the standardization of testing protocols and validation methods to facilitate the transition from laboratory research to real-world applications. Interdisciplinary collaboration between engineers, neuroscientists, and clinicians is crucial to ensuring that the solutions developed meet the practical needs of users. Future research directions should focus on improving the robustness of signal processing techniques, enhancing sensory feedback integration, and developing lightweight, energy-efficient prosthetic models. Furthermore, the adoption of

hybrid control strategies may offer a more inclusive approach to prosthetic control, leading to a new generation of intelligent and adaptive robotic limbs.

5. Limitations of the Existing Works and Future Road-Map

After decades of research in the field of robotic prostheses, a paradoxical situation has emerged: although the development of new prostheses represents one of the most promising areas in robotics, most amputee patients still use technologies that have evolved little over the past decades. This paradox is even more striking when considering the colossal challenge of developing a prosthesis capable of replicating the sensory and motor functions of a human hand. The hand possesses one of the most complex motor and sensory representations in the human brain, making the execution of grasping tasks extremely complex.

Current challenges are vast and multifaceted, requiring a combination of advanced mechatronic solutions to develop highly sensorized robotic hand prostheses, along with novel approaches to establish robust and efficient interfaces with patients' nervous systems. There is presently a diverse range of options for restoring hand function, from custom cosmetic solutions to state-of-the-art prostheses featuring sensory feedback and advanced control. Each solution offers its benefits and drawbacks, with the ideal choice depending on the patient's individual needs and preferences. For instance, the custom cosmetic solution (Figure 1a) provides a natural and realistic appearance that enhances the patient's self-esteem and confidence. It can be personalized to match the skin tone, texture, and characteristics of the unaffected hand; however, it does not restore motor or sensory functions.

The body-powered prosthetic solution (Figure 1b) harnesses the patient's muscle strength to control the prosthesis, offering intuitive and direct control. It does not require an external power source and is both lightweight and portable, though it may demand more physical effort and offer more limited movement compared to myoelectric prostheses. In contrast, the surface electromyography (sEMG) pattern recognition system (Figure 1c) analyzes myoelectric signals from the patient's residual muscles to identify movement patterns. In this system, electromyographic patterns are translated into control commands for the prosthesis, enabling more complex and precise movements. Furthermore, with sEMG control combined with extracellular stimulation via an implanted FINE (flat interface nerve electrode) electrode (Figure 1d), there is a merging of myoelectric control and neural stimulation to provide sensory feedback to the prosthetic hand. The implanted FINE electrode stimulates sensory nerves by transmitting information about the external environment, thereby offering a more natural and intuitive sensory experience.

In fully implanted myoelectric sensor systems (Figure 1e), the need for external electrodes is eliminated, resulting in a more natural and discreet appearance. These sensors capture myoelectric signals directly from the muscles, thereby reducing noise and interference. However, they require surgical implantation and may have a limited lifespan. The prosthesis designated as i-Limb Ultra (Figure 1f) operates using sEMG control. Additionally, in systems where sensory feedback is provided by intraneural TIMEs (Transverse Intrafascicular Multichannel Electrode) electrodes (Figure 1g), direct stimulation of sensory nerves within the peripheral nerve delivers high-resolution feedback. These electrodes represent a recent technological advancement with the potential to revolutionize sensory feedback in prostheses [1].

Table 3 shows a qualitative evidence map regarding the technological approaches discussed throughout this review and the associated limitations.

Table 3. Technological landscape of neuromuscular prostheses.

Category	Typical Capabilities	Known Limitations	Representative Examples
Surface EMG sensors	Non-invasive, accessible, and well-established for prosthetic control	Signal degradation due to sweat, electrode displacement, variable skin impedance; reduced long-term stability	Ottobock myoelectric prostheses
High-Density EMG AR-Rays	Higher spatial resolution, potential for better DOF decoding	Increased computational load, complexity, and cost	HD-sEMG pattern recognition methods
Implanted myoelectric sensors	Stable long-term recordings and less affected by skin impedance	Requires surgery and risk of failure or encapsulation	IMES implants
Peripheral nerve interfaces	High selectivity, access to motor and sensory pathways	Surgical implantation, variable long-term stability	Utah and cuff electrodes
Osseointegration	Stable skeletal fixation and improved Comfort through direct load transfer	Risk of infection and percutaneous interface care needed	OPRA implant
Targeted muscle reinnervation	Redirects neural signals to residual muscles, improves EMG separation, and intuitive control.	Specialized surgery with a limited candidate pool	Nerve transfer surgeries performed on amputees with neuroma pain
Pneumatic artificial muscles	High power-to-weight, compliant motion, bio-inspired	Bulky compressors, low durability under cycles	McKibben muscles
Machine Learning control	Higher accuracy, adaptive to user variability, multi-DOF potential	Data hungry, computationally heavy	Deep learning-based EMG classifiers
Brain–Computer interfaces	Direct neural decoding, potential for high-DOF intuitive control	Low bandwidth (EEG), invasive surgery, calibration burden, and regulatory hurdles	Neuralink N1 SoC

Table 4 provides an evidence map of neuromuscular prosthesis technologies, covering the domains of sensing, control, actuation, feedback, and BCI. This table includes typical capabilities, known limitations, and readiness levels.

This analysis highlights the technological improvements and ongoing challenges in the design of neuromuscular robotic prostheses, with particular focus on signal acquisition technologies, control strategies, and actuator typologies. The inherent complexity of replicating human hand function, especially in terms of tactile sensitivity, proprioception, and multi-joint control, has driven the development of increasingly integrated solutions. The diversity of control approaches available today reflects the trade-offs required to simultaneously optimize accuracy, robustness, and natural movement. Thus, the progress in this field depends on the gathered-evolution of engineering, neuroscience, and medical technology, with direct implications for functional rehabilitation and the quality of life of upper-limb prosthesis users.

Table 4. Evidence map of neuromuscular prosthesis technologies.

Domain	Typical Capabilities	Limitations	Readiness (Low/High)	References
Sensing	sEMG provides non-invasive control; HD-sEMG enables spatial resolution; intramuscular and nerve cuffs give stable long-term signals	sEMG sensitive to sweat, motion artifacts, electrode shift; invasive sensors require surgery and present long-term biocompatibility concerns	High (sEMG) Low (invasive)	[28–36]
Control	Threshold, FSM, and pattern-recognition allow multi-grasp decoding; deep learning enables proportional and simultaneous multi-DOF control.	Variability across subjects; retraining needed; robustness in real-world limited; computational cost of deep models	Low (pattern recognition) High (threshold /FSM)	[37–45]
Actuation	DC motors are robust and compact; PAMs and soft actuators provide compliance	DC motors might be heavy and have low compliance; PAMs require compressors and present durability issues; soft actuators have low force and limited lifespan	High (DC motors) Low (soft actuators and PAMs)	[19–23]
Feedback	Vibrotactile, electrotactile, and intraneural interfaces restore partial sensation; AMI and osseointegration enhance proprioception.	Non-invasive feedback feels unnatural; invasive approaches have risks (fibrosis and infection); lack of standardized testing	Low (invasive electrodes) High (osseointegration)	[40,47,55,67]
BCI	EEG-based BCIs allow motor imagery decoding; invasive arrays enable high resolution; VR/rehab integration	EEG-based BCIs exhibit low signal-to-noise ratios, bitrates below 30–50 bits/minute, and high inter and intra-subject variability; high training demand; invasive BCIs carry surgical risks, and limited long-term data	High (EEG), Low (rehab-focused BCI)	[73–77]

6. Conclusions

Recent advancements in neuromuscular robotic prostheses reflect a transformative synergy between artificial intelligence, sensorimotor engineering, and neurotechnology. EMG-based control systems, brain–computer interfaces, and machine learning have significantly enhanced the precision and intuitiveness of robotic prostheses. This review has outlined the evolution of the main subsystems of these devices, highlighting the technical innovations driving performance gains in real-world applications.

Modern signal acquisition techniques combined with deep learning architectures can achieve good decoding accuracy. These methods exhibit resilience to signal drift, electrode displacement, and noise. Moreover, recent developments in hybrid control systems, fusing EMG and EEG inputs, offer enhanced robustness across dynamic conditions. Sensory feedback mechanisms, both invasive and noninvasive, have matured significantly in recent years, which simultaneously improves signal selectivity and enables bidirectional flow of information, facilitating artificial proprioception and tactile feedback.

The convergence of three pillars, intelligent control algorithms, surgically optimized neuromuscular interfaces, and osseointegrated anchoring, has shifted the clinical paradigm from passive assistance to active neurointegration. Osseointegration, beyond providing

mechanical fixation, is now recognized as an electromechanical interface supporting implanted electronics and transcutaneous energy transfer, crucial for embedded sensing and stimulation without compromising skin integrity. These systems pave the way for fully implantable, maintenance-free prosthetic architectures with long-term usability.

Despite these progresses, critical challenges persist. The long-term reliability of surface EMG remains hindered by motion artifacts, muscle fatigue, and variable skin-electrode contact. Implanted neural interfaces, while offering higher fidelity, are limited by biological encapsulation, signal degradation, and the risk of chronic immune responses. Power consumption also remains a significant barrier; integrating actuators, batteries, sensors, and processors into a compact, lightweight system without compromising function is a key design constraint. Additionally, many state-of-the-art control algorithms remain confined to controlled experimental settings, lacking robustness for deployment in unstructured environments.

Another major bottleneck is the absence of standardized validation frameworks. Discrepancies in signal processing pipelines, outcome metrics, and safety protocols hinder reproducibility and regulatory translation. Real-time bidirectional control also remains underdeveloped, as latency, even on the order of milliseconds, can disrupt the synchrony between user intent and execution. Energy-efficient hardware, adaptive actuators, and neuromorphic computing platforms are promising directions to address this constraint.

Future research must target these gaps through a systems-level approach. Embedding AI models within low-power edge devices enables responsive, wearable prosthetics with minimal energy overhead. Material innovation, such as bioactive electrode coatings and flexible substrates, may improve long-term biocompatibility. Meanwhile, the democratization of prosthetic development through open-source platforms and additive manufacturing is accelerating accessibility and customization. Equally vital is the refinement of user adaptation strategies and longitudinal validation studies. Rigorous clinical trials with amputee cohorts are essential to assess functional outcomes, cognitive load, and long-term neural plasticity. Cross-disciplinary collaboration among engineers, clinicians, neuroscientists, and designers is no longer optional but foundational to the field's advancement. Development teams must address not only the mechanical and electrical interfaces but also the physiological, psychological, and social dimensions of limb loss and restoration.

In summary, the integration of advanced biosignal acquisition, AI-driven control frameworks, and closed-loop sensory feedback is redefining the frontiers of neuromuscular prosthetics. Achieving intuitive control with real-time feedback is within reach, but will require persistent investigations in embedded systems, materials science, and neuroengineering. As the field moves from proof-of-concept to clinical translation, the emphasis must remain on human-centered design, reproducible methodologies, and sustainable deployment. Ultimately, the goal is to create prosthetic systems that do not merely restore movement but reinstate embodiment, devices that respond to the user's intent with naturalness and precision.

Author Contributions: Conceptualization, P.G.B., P.D.G. and A.C.V.; methodology, P.G.B., P.D.G. and A.C.V.; software, P.G.B.; validation, P.D.G. and A.C.V.; formal analysis, P.G.B., P.D.G. and A.C.V.; investigation, P.G.B., P.D.G. and A.C.V.; resources, P.D.G.; data curation, P.G.B.; writing—original draft preparation, P.G.B., P.D.G. and A.C.V.; writing—review and editing, P.D.G. and A.C.V.; supervision, P.D.G. and A.C.V.; funding acquisition, P.D.G. All authors have read and agreed to the published version of the manuscript.

Funding: The authors would like to express their gratitude to Fundação para a Ciência e Tecnologia (FCT) and C-MAST (Center for Mechanical and Aerospace Science and Technologies) for their support in the form of funding under the project UIDB/00151/2020.

Data Availability Statement: Not applicable.

Acknowledgments: The authors acknowledge the support provided by LITecS (Laboratory of Innovation and Technologies for Sustainability) (Available online: <https://lites.ubi.pt/en/> (accessed on 10 March 2025)).

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Mendez, V.; Iberite, F.; Shokur, S.; Micera, S. Current solutions and future trends for robotic prosthetic hands. *Annu. Rev. Control. Robot. Auton. Syst.* **2021**, *4*, 595–627. [CrossRef]
2. Hargrove, L.J.; Miller, L.A.; Turner, K.; Kuiken, T.A. Myoelectric pattern recognition outperforms direct control for transhumeral amputees with targeted muscle reinnervation: A randomized clinical trial. *Sci. Rep.* **2017**, *7*, 13840. [CrossRef] [PubMed]
3. Zheng, J.Z.; De La Rosa, S.; Dollar, A.M. An investigation of grasp type and frequency in daily household and machine shop tasks. In Proceedings of the 2011 IEEE International Conference on Robotics and Automation, Shanghai, China, 9–13 May 2011; pp. 4169–4175. [CrossRef]
4. Coert, J.H.; Van Dijke, G.A.H.; Hovius, S.E.R.; Snijders, C.J.; Meek, M.F. Quantifying thumb rotation during circumduction utilizing a video technique. *J. Orthop. Res.* **2003**, *21*, 1151–1155. [CrossRef] [PubMed]
5. Belter, J.T.; Dollar, A.M. Performance characteristics of anthropomorphic prosthetic hands. In Proceedings of the 2011 IEEE International Conference on Rehabilitation Robotics, Zurich, Switzerland, 29 June–1 July 2011. pp. 1–7. [CrossRef]
6. Laliberté, T.; Birglen, L.; Gosselin, C. Underactuation in Robotic Grasping Hands. *Jpn. J. Mach. Intell. Robot. Control* **2002**, *4*, 77–87.
7. Yoder, Z.; Kellaris, N.; Chase-Markopoulou, C.; Ricken, D.; Mitchell, S.K.; Emmett, M.B.; Weir, R.F.F.; Segil, J.; Keplinger, C. Design of a High-Speed Prosthetic Finger Driven by Peano-HASEL Actuators. *Front. Robot. AI* **2020**, *7*, 586216. [CrossRef]
8. Robertson, M.A.; Sadeghi, H.; Florez, J.M.; Paik, J. Soft Pneumatic Actuator Fascicles for High Force and Reliability. *Soft Robot.* **2017**, *4*, 23–32. [CrossRef] [PubMed]
9. Huang, F.; Sun, X.; Xu, Q.; Cheng, W.; Shi, Y.; Pan, L. Recent Developments and Applications of Tactile Sensors with Biomimetic Microstructures. *Biomimetics* **2025**, *10*, 147. [CrossRef]
10. Navaraj, W.T.; Nassar, H.; Dahiya, R. Prosthetic hand with biomimetic tactile sensing and force feedback. In Proceedings of the 2019 IEEE International Symposium on Circuits and Systems (ISCAS), Sapporo, Japan, 26–29 May 2019; pp. 1–4. [CrossRef]
11. Mahmoud, M.; Takele, N.A.; Derege, O.; Rifiki, N.; El Khalil, M.; Tiwari, V.; Zhong, J. Tactile Sensors: A Review. *Measurement* **2024**, *238*, 115332. [CrossRef]
12. Pirozzi, S. Tactile Sensors for Robotic Applications. *Sensors* **2020**, *20*, 7009. [CrossRef]
13. Liu, T.; Brayshaw, G.; Li, A.; Xu, X.; Ward-Cherrier, B. Neuromorphic touch for robotics—A review. *Neuromorphic Comput. Eng.* **2025**, *5*, 032001. [CrossRef]
14. Oliveira, T.E.A.; Fonseca, V.P. BioIn-Tacto: A compliant multi-modal tactile sensing module for robotic tasks. *HardwareX* **2023**, *16*, e00478. [CrossRef]
15. Lang, L.; Antunes, R.; Dutra, T.A.; Aguiar, M.L.d.; Pereira, N.; Gaspar, P.D. Mechanical Characterization and Computational Analysis of TPU 60A: Integrating Experimental Testing and Simulation for Performance Optimization. *Materials* **2025**, *18*, 240. [CrossRef]
16. Antunes, R.; Lang, L.; de Aguiar, M.L.; Dutra, T.A.; Gaspar, P.D. Design of Fin Ray Effect Soft Robotic Gripper for Improved Mechanical Performance and Adaptability: Numerical Simulations and Experimental Validation. In Proceedings of the 2024 20th IEEE/ASME International Conference on Mechatronic and Embedded Systems and Applications (MESA), Genova, Italy, 2–4 July 2024; pp. 1–6. [CrossRef]
17. Antunes, R.; Lang, L.; de Aguiar, M.L.; Dutra, T.A.; Gaspar, P.D. Enhancing the Performance of Fin Ray Effect Soft Robotic Finger via Computational Design and Simulation. In Proceedings of the 2024 IEEE International Conference on Autonomous Robot Systems and Competitions (ICARSC), Paredes de Coura, Portugal, 18–19 April 2024; pp. 189–194. [CrossRef]
18. Wang, F.; Liu, S.; Shu, L.; Tao, X.M. Low-dimensional carbon-based sensors and sensing network for wearable health and environmental monitoring. *Carbon* **2017**, *121*, 353–367. [CrossRef]
19. Klute, G.K.; Czerniecki, J.M.; Hannaford, B. McKibben artificial muscles: Pneumatic actuators with biomechanical intelligence. *IEEE/ASME Trans. Mechatron.* **1999**, *7*, 418–430. [CrossRef]
20. Chou, C.P.; Hannaford, B. Measurement and modeling of McKibben pneumatic artificial muscles. *IEEE Trans. Robot. Autom.* **1996**, *12*, 90–102. [CrossRef]
21. Noritsugu, T.; Tanaka, T. Application of Rubber Artificial Muscle Manipulator as a Rehabilitation Robot. *IEEE/ASME Trans. Mechatron.* **1997**, *2*, 259–267. [CrossRef]
22. Gong, D.; Yu, J. Design and Control of the McKibben Artificial Muscles Actuated Humanoid Manipulator. *IntechOpen* **2022**. [CrossRef]

23. Repperger, D.W.; Phillips, C.A.; Krier, M. Controller design involving gain scheduling for a large scale pneumatic muscle actuator. In Proceedings of the 1999 IEEE International Conference on Control Applications, Kohala Coast, HI, USA, 22–27 August 1999; pp. 285–290.
24. Geethanjali, P. Myoelectric control of prosthetic hands: State-of-the-art review. *Med. Devices Evid. Res.* **2016**, *9*, 247–255. [CrossRef]
25. Kamavuako, E.N.; Brown, M.; Bao, X.; Chihi, I.; Pitou, S.; Howard, M. Affordable Embroidered EMG Electrodes for Myoelectric Control of Prostheses: A Pilot Study. *Sensors* **2021**, *21*, 5245. [CrossRef] [PubMed]
26. Ribeiro, J.; Fernandes, H.; Lopes, C.; Magalhães, F.; Teixeira, S. Analysis of Man-Machine Interfaces in Upper-Limb Prosthesis: A Review. *Robotics* **2019**, *8*, 16. [CrossRef]
27. Jabban, L.; Dupan, S.; Zhang, D.; Ainsworth, B.; Nazarpour, K.; Metcalfe, B.W. Sensory Feedback for Upper-Limb Prostheses: Opportunities and Barriers. *IEEE Trans. Neural Syst. Rehabil.* **2022**, *30*, 738–747. [CrossRef]
28. Konrad, P. *The ABC of EMG: A Practical Introduction to Kinesiological Electromyography, Version 1.0*; Noraxon Inc.: Scottsdale, AZ, USA, 2005; pp. 1–30.
29. Naik, G. *Biomedical Signal Processing: Advances in Theory, Algorithms and Applications*; Springer: Singapore, 2019. [CrossRef]
30. Hermens, H.J.; Freriks, B.; Disselhorst-Klug, C.; Rau, G. Development of recommendations for SEMG sensors and sensor placement procedures. *J. Electromyogr. Kinesiol.* **2000**, *10*, 361–374. [CrossRef] [PubMed]
31. Gohel, V.; Mehendale, N. Review on electromyography signal acquisition and processing. *Biophys Rev.* **2020**, *12*, 1361–1367. [CrossRef]
32. Valentinuzzi, M.E.; Sörnmo, L.; Laguna, P.; Merletti, R.; Parker, P. Bioelectrical Signal Processing in Cardiac and Neurological Applications. *Biomed. Eng. Online* **2007**, *6*, 27. [CrossRef]
33. Pasquina, P.F.; Childress, D.S.; Armiger, R.S.; Vandersea, J.; Tsao, J.W. First-in-man demonstration of a fully implanted myoelectric sensors system to control an advanced electromechanical prosthetic hand. *J. Neurosci. Methods* **2015**, *244*, 85–93. [CrossRef] [PubMed]
34. Merrill, D.R.; Lockhart, J.; Troyk, P.R.; Weir, R.F.; Hankin, D.L. Development of an Implantable Myoelectric Sensor for Advanced Prosthesis Control. *Artif. Organs* **2011**, *35*, 249–252. [CrossRef]
35. Badia, J.; Boretius, T.; Andreu, D.; Azevedo-Coste, C.; Stieglitz, T.; Navarro, X. Comparative Analysis of Transverse Intrafascicular Multichannel, Longitudinal Intrafascicular and Multipolar Cuff Electrodes for the Selective Stimulation of Nerve Fascicles. *J. Neural Eng.* **2011**, *8*, 036023. [CrossRef]
36. Hong, J.W.; Yoon, C.; Jo, K.; Won, J.H.; Park, S. Recent advances in recording and modulation technologies for next-generation neural interfaces. *iScience* **2021**, *24*, 103550. [CrossRef]
37. Roche, A.D.; Rehbaum, H.; Farina, D.; Aszmann, O.C. Prosthetic Myoelectric Control Strategies: A Clinical Perspective. *Curr. Surg. Rep.* **2014**, *2*, 44. [CrossRef]
38. Calado, A.; Soares, F.; Matos, D. A Review on Commercially Available Anthropomorphic Myoelectric Prosthetic Hands, Pattern-Recognition-Based Microcontrollers and sEMG Sensors Used for Prosthetic Control. In Proceedings of the 2019 IEEE International Conference on Autonomous Robot Systems and Competitions (ICARSC), Paredes, Portugal, 24–26 April 2019; pp. 1–6. [CrossRef]
39. Moore, J.E., Jr.; Maitland, D.J. (Eds.) *Biomedical Technology and Devices*, 2nd ed.; CRC Press: Boca Raton, FL, USA, 2013. [CrossRef]
40. Secco, E.L.; Noh, Y. Editorial: Human-like Robotic Hands for Biomedical Applications and Beyond. *Front. Robot. AI* **2024**, *11*, 1414971. [CrossRef]
41. Madusanka, D.G.K.; Wijayasingha, L.N.S.; Gopura, R.A.R.C.; Amarasinghe, Y.W.R.; Mann, G.K.I. A Review on Hybrid Myoelectric Control Systems for Upper Limb Prosthesis. In Proceedings of the 2015 Moratuwa Engineering Research Conference (MERCon), Moratuwa, Sri Lanka, 7–8 April 2015; pp. 136–141.
42. Ke, A.; Huang, J.; Wang, J.; He, J. Improving the Robustness of Human-Machine Interactive Control for Myoelectric Prosthetic Hand during Arm Position Changing. *Front. Neurorobot.* **2022**, *16*, 853773. [CrossRef] [PubMed]
43. Domenico, D.D.; Siboni, M.H.; Fiore, G.; Pirozzi, S.; Dosen, S.; Farina, D. Long-Term Upper-Limb Prosthesis Myocontrol via High-Density sEMG and Incremental Learning. *IEEE Robot. Autom. Lett.* **2024**, *9*, 9938–9945. [CrossRef]
44. Gao, S.; Gong, J.; Chen, B.; Zhang, B.; Luo, F.; Yerabakan, M.O.; Pan, Y.; Hu, B. Use of Advanced Materials and Artificial Intelligence in Electromyography Signal Detection and Interpretation. *Adv. Intell. Syst.* **2022**, *4*, 2200063. [CrossRef]
45. Zbinden, J.; Artoni, F.; van der Kooij, H.; Millán, J.D.R.; Micera, S. Improved Control of a Prosthetic Limb by Surgically Creating Electro-Neuromuscular Constructs with Implanted Electrodes. *Sci. Transl. Med.* **2023**, *15*, eabq3665. [CrossRef] [PubMed]
46. AlQahtani, N.J.; Al-Naib, I.; Althobaiti, M. Recent Progress on Smart Lower Prosthetic Limbs: A Comprehensive Review on Using EEG and fNIRS Devices in Rehabilitation. *Front. Bioeng. Biotechnol.* **2024**, *12*, 1454262. [CrossRef]
47. Song, H.; Hsieh, T.H.; Yeon, S.H.; Shu, T.; Nawrot, M.; Landis, C.F. Continuous Neural Control of a Bionic Limb Restores Biomimetic Gait After Amputation. *Nat. Med.* **2024**, *30*, 2010–2019. [CrossRef]
48. Song, H.; Israel, E.A.; Gutierrez-Arango, S.; Teng, A.C.; Srinivasan, S.S.; Freed, L.E.; Herr, H.M. Agonist–Antagonist Muscle Strain in the Residual Limb Preserves Motor Control and Perception After Amputation. *Commun. Med.* **2022**, *2*, 97. [CrossRef]

49. Xu, B.; Zhong, L.; Zhang, G.; Liang, X.; Virtue, D.; Madan, R.; Bhattacharjee, T. CushSense: Soft, Stretchable and Comfortable Tactile-Sensing Skin for Physical Human–Robot Interaction. *arXiv* **2024**, arXiv:2405.03155. [CrossRef]
50. Zheng, L.Q.; Wang, J.; Wang, J.; Lei, R.; Jin, H.; Chai, W.; Ge, H.; Jin, D. Design and Application of Flexible Sensors in Human–Machine Interaction. *Adv. Polym. Technol.* **2024**, *2024*, 5535918. [CrossRef]
51. Kumar, V.; Pratihari, D.K. Wearable sensor-based intent recognition for adaptive control of intelligent ankle-foot prosthetics. *Meas. Sens.* **2025**, *39*, 101865. [CrossRef]
52. Luu, D.K.; Nguyen, A.T.; Jiang, M.; Drealan, M.W.; Xu, J.; Wu, T. Artificial Intelligence Enables Real-Time and Intuitive Control of Prostheses via Nerve Interface. *IEEE Trans. Biomed. Eng.* **2022**, *69*, 3051–3063. [CrossRef]
53. Liu, Y.; Zhou, X.; Wang, W.; Zhang, Z.; Li, J. Recent Advances in Flexible Bending Sensors and Their Applications. *Int. J. Smart Nano Mater.* **2024**, *15*, 697–729. [CrossRef]
54. Varghese, R.J.; Pizzi, M.; Kundu, A.; Grison, A.; Burdet, E.; Farina, D. Design, Fabrication and Evaluation of a Stretchable High-Density Electromyography Array. *Sensors* **2024**, *24*, 1810. [CrossRef]
55. Schultz Feuser, A.E.; Barlow, A.K. *Targeted Muscle Reinnervation: A Neural Interface for Artificial Limbs*, 1st ed.; CRC Press: Boca Raton, FL, USA, 2013. [CrossRef]
56. Le, E.L.H.; Iorio, M.L.; Greyson, M.A. Targeted Muscle Reinnervation in Upper Extremity Amputations. *Eur. J. Orthop. Surg. Traumatol.* **2024**, *34*, 3717–3725. [CrossRef]
57. Lauzon, J.C.; Boyd, K.U.; Dudek, N.L. What to Expect Following Targeted Muscle Reinnervation/Regenerative Peripheral Nerve Interface: Pain Outcomes in an Amputee Population. *J. Plast. Reconstr. Aesthetic Surg.* **2024**, *99*, 373–376. [CrossRef]
58. Dumanian, G.A.; Potter, B.K.; Mioton, L.M. Targeted Muscle Reinnervation Treats Neuroma and Phantom Pain in Major Limb Amputees: A Randomized Clinical Trial. *Ann. Surg.* **2019**, *270*, 238–246. [CrossRef] [PubMed]
59. Cheesborough, J.E.; Souza, J.M.; Dumanian, G.A.; Bueno, R.A. Targeted Muscle Reinnervation in the Initial Management of Traumatic Upper Extremity Amputation Injury. *Hand* **2014**, *9*, 253–257. [CrossRef] [PubMed]
60. Souza, J.M.; Cheesborough, J.E.; Ko, J.H. Targeted Muscle Reinnervation: A Novel Approach to Postamputation Neuroma Pain. *Clin. Orthop. Relat. Res.* **2014**, *472*, 2984–2990. [CrossRef]
61. Kuiken, T.A.; Barlow, A.K.; Hargrove, L.; Dumanian, G.A. Targeted Muscle Reinnervation for the Upper and Lower Extremity. *Tech. Orthop.* **2017**, *32*, 109–116. [CrossRef] [PubMed]
62. Roubaud, M.S.; Hassan, A.M.; Shin, A. Outcomes of Targeted Muscle Reinnervation and Regenerative Peripheral Nerve Interfaces for Chronic Pain Control in the Oncologic Amputee Population. *J. Am. Coll. Surg.* **2023**, *237*, 644–654. [CrossRef]
63. Hoyt, B.W.; Gibson, J.A.; Potter, B.K.; Souza, J.M. Practice Patterns and Pain Outcomes for Targeted Muscle Reinnervation: An Informed Approach to TMR Use in the Acute Amputation Setting. *J. Bone Joint Surg. Am.* **2021**, *103*, 681–687. [CrossRef]
64. Yuan, M.; Gallo, M.; Gallo, L. Targeted Muscle Reinnervation and Regenerative Peripheral Nerve Interfaces versus Standard Management in the Treatment of Limb Amputation: A Systematic Review and Meta-Analysis. *Plast. Surg.* **2024**, *32*, 253–264. [CrossRef]
65. Yang, J.; Shibata, K.; Weber, D.; Erickson, Z. High-density electromyography for effective gesture-based control of physically assistive mobile manipulators. *Npj Robot.* **2025**, *3*, 2. [CrossRef]
66. Hwang, S.; Kwon, N.; Lee, D.; Kim, J.; Yang, S.; Youn, I.; Moon, H.-J.; Sung, J.-K.; Han, S. A Multimodal Fatigue Detection System Using sEMG and IMU Signals with a Hybrid CNN-LSTM-Attention Model. *Sensors* **2025**, *25*, 3309. [CrossRef]
67. Ciuffreda, M. The Role of Osseointegration in Bionic Prostheses of the Upper Limb. Doctoral Dissertation, Università Campus Bio-Medico di Roma, Rome, Italy, 2022.
68. Hoellwarth, J.S.; Tetsworth, K.; Rozbruch, S.R.; Handal, M.B.; Coughlan, A.; Al Muderis, M. Osseointegration for Amputees: Current Implants, Techniques and Future Directions. *JBJS Rev.* **2020**, *8*, e0043. [CrossRef] [PubMed]
69. Brånemark, R.; Berlin, Ö.; Hagberg, K.; Bergh, P.; Gunterberg, B.; Rydevik, B. A Novel Osseointegrated Percutaneous Prosthetic System for the Treatment of Patients with Transfemoral Amputation. *Bone Jt. J.* **2014**, *96-B*, 106–113. [CrossRef] [PubMed]
70. Akhtar, M.A.; Low, C.; Tiemessen, C. Current Challenges and Future Prospects of Osseointegration Limb Reconstruction for Amputees. *SN Compr. Clin. Med.* **2024**, *6*, 4. [CrossRef]
71. Sun, Y.; Xu, J.; Lv, S.; Xu, Z.; Li, L.; Li, Y.; Li, Y. Extramedullary Osseointegration—A Novel Design of Percutaneous Osseointegration Prosthesis for Amputees. *Front. Bioeng. Biotechnol.* **2022**, *10*, 81128. [CrossRef]
72. Sierakowski, A.; Watts, C.; Thomas, K.; Elliot, D. Long-term outcomes of osseointegrated digital prostheses for proximal amputations. *J. Hand Surg. Eur. Vol.* **2011**, *36*, 116–125. [CrossRef]
73. Camargo-Vargas, D.; Callejas-Cuervo, M.; Mazzoleni, S. Brain–Computer Interfaces Systems for Upper and Lower Limb Rehabilitation: A Systematic Review. *Sensors* **2021**, *21*, 4312. [CrossRef]
74. Ari, E.; Taçgın, E. NF-EEG: A Generalized CNN Model for Multi-Class EEG Motor Imagery Classification without Signal Preprocessing for Brain–Computer Interfaces. *Biomed. Signal Process. Control* **2024**, *92*, 106081. [CrossRef]
75. Choi, S.H.; Yoon, H.; Baek, H.J.; Long, X. Biomedical Signal Processing and Health Monitoring Based on Sensors. *Sensors* **2025**, *25*, 641. [CrossRef] [PubMed]

76. Liu, Z.; Shore, J.; Wang, M.; Yuan, F.; Buss, A.; Zhao, X. A systematic review on hybrid EEG/fNIRS in brain-computer interface. *Biomed. Signal Process. Control* **2021**, *68*, 102595. [CrossRef]
77. Musk, E.; Neuralink. An Integrated Brain–Machine Interface Platform with Thousands of Channels. *J. Med. Internet Res.* **2019**, *21*, e16194. [CrossRef]
78. Merletti, R.; Farina, D. (Eds.) *Surface Electromyography: Physiology, Engineering and Applications*; John Wiley & Sons, Inc.: Hoboken, NJ, USA, 2016.
79. Cheon, S.; Man, S.; Lazea, G.; Marcu, C.; Raica, P.; Robotin, R. Single EMG Sensor-Driven Robotic Glove Control for Reliable Augmentation of Power Grasping. *IEEE Trans. Med. Robot. Bionics* **2021**, *3*, 179–189. [CrossRef]
80. Favieiro, G.W. Controle de uma Prótese Experimental do Segmento Mão-Braço por Sinais Mioelétricos e Redes Neurais Artificiais. Master’s Thesis, Universidade Federal do Rio Grande do Sul, Porto Alegre, Brasil, 2009.
81. Wei, Z.; Li, M.; Zhang, Z.Q.; Xie, S.Q. Continuous Prediction of Wrist Joint Kinematics Using Surface Electromyography from the Perspective of Muscle Anatomy and Muscle Synergy Feature Extraction. *IEEE J. Biomed. Health Inform.* **2025**, *1*, 43–55. [CrossRef]
82. Foroutannia, A.; Akbarzadeh, M.; Akbarzadeh, A. A deep learning strategy for EMG-based joint position prediction in hip exoskeleton assistive robots. *Biomed. Signal Process. Control* **2022**, *75*, 103557. [CrossRef]
83. Jaramillo-Yáñez, A.; Benalcázar, M.E.; Mena-Maldonado, E. Real-Time Hand Gesture Recognition Using Surface Electromyography and Machine Learning: A Systematic Literature Review. *Sensors* **2020**, *20*, 2467. [CrossRef]
84. Castro, A.; Pinheiro, W.C.; Rigolin, G. A Hybrid 3D-Printed Hand Prosthesis Prototype based on sEMG with Embedded Computer Vision. *Front. Neurorobot.* **2021**, *15*, 751282. [CrossRef]
85. Zhang, X.; Yao, L.; Wang, X.; Monaghan, J.; McAlpine, D.; Zhang, Y. A Survey on Deep Learning-based Non-Invasive Brain Signals: Recent Advances and New Frontiers. *J. Neural Eng.* **2019**, *18*, 031002. [CrossRef] [PubMed]
86. Choi, B.J. Removing Neural Signal Artifacts with Autoencoder-Targeted Adversarial Transformers (AT-AT). *arXiv* **2025**, arXiv:2502.05332. [CrossRef]
87. Chen, J.; Bi, S.; Zhang, G.; Cao, G. High-Density Surface EMG-Based Gesture Recognition Using a 3D Convolutional Neural Network. *Sensors* **2020**, *20*, 1201. [CrossRef] [PubMed]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Review

Recent Developments in Electroadhesion Grippers for Automated Fruit Grasping

Turac I. Ozcelik ^{1,2,*}, Enrico Masi ², Seyyed Masoud Kargar ³, Chiara Scagliarini ², Adyan Fatima ⁴, Rocco Vertechy ^{2,5} and Giovanni Berselli ^{3,6}

¹ Robotics and Intelligent Machines, University of Genova, Via Balbi, 5, 16126 Genova, Italy

² Department of Industrial Engineering, University of Bologna, Viale del Risorgimento, 40136 Bologna, Italy; enrico.masi10@unibo.it (E.M.); chiara.scagliarini4@unibo.it (C.S.); rocco.vertechy@unibo.it (R.V.)

³ Department of Mechanical, Energy, Management and Transportation Engineering (DIME), University of Genova, Via Balbi, 5, 16126 Genova, Italy; seyyed.masoud.kargar@edu.unige.it (S.M.K.); giovanni.berselli@unige.it (G.B.)

⁴ Department of Computer Science, Bioengineering, Robotics and Systems Engineering—DIBRIS, University of Genova, Via Balbi, 5, 16126 Genova, Italy; fatima.adyan@unibo.it

⁵ CNR-STIIMA, National Research Council of Italy—Institute of Intelligent Industrial Technologies and Systems for Advanced Manufacturing, Via Alfonso Corti 12, 20133 Milan, Italy

⁶ Advanced Robotics Department (ADV), Fondazione Istituto Italiano di Tecnologia (IIT), Via Morego, 30, 16163 Genova, Italy

* Correspondence: ilkayturac.ozcelik2@unibo.it

Abstract

As global food demand rises and agricultural labor shortages intensify, robotic automation has become essential for sustainable fruit grasping. Among emerging technologies, ElectroAdhesion (EA) grippers offer a promising alternative to traditional mechanical end-effectors, enabling gentle, low-pressure handling through electrostatically induced adhesion. This paper presents a methodical review of EA grippers applied to fruit grasping, focusing on their advantages, limitations, and key design considerations. A targeted literature search identified ten EA-based and hybrid EA gripping systems tested on fruit manipulation, though none has yet been tested in real-world environments such as fields or greenhouses. Despite a significant variability in experimental setups, materials, and grasp types, qualitative insights are drawn from our analysis demonstrating the potentialities of EA technologies. The EA grippers found in the targeted review are effective on diverse fruits, shapes, and surface textures; they can hold load capacities ranging from 10 g (~0.1 N) to 600 g (~6 N) and provide minimal compressive stress at high electrostatic shear forces. Along with custom EA grippers designed accordingly to specific use cases, field and greenhouse testing will be crucial for advancing the technology readiness level of EA grippers and unlocking their full potential in automated crop harvesting.

Keywords: robotic gripper; end-effector; electroadhesion; electroadhesion gripper; fruit; harvesting; grasping; picking; manipulation

1. Introduction

The accelerating rate of global population growth, coupled with profound socio-economic challenges, underscores the urgent need for sustainable food production in the coming decades. According to the United Nations, the world population is projected to reach 9.8 billion by 2050 [1]. Meanwhile, long-run assessments by the Food and Agriculture Organization (FAO) indicate that feeding the mid-century population will require raising

global food production by roughly 70% relative to 2005/07 levels [2]. However, agriculture continues to rely heavily on labor-intensive practices, particularly in the handling and harvesting of fruits and vegetables. These operations are time-consuming and costly, and they are increasingly undermined by labor shortages and climate-driven shifts in ripening patterns, which together contribute to postharvest losses and reduced efficiency [3].

In this context, research has focused on developing robotic manipulators capable of performing automated harvesting missions with the aim of substituting or complementing human labor, offering advantages such as increased throughput, reduced waste, and more consistent quality. Unlike human operators, robotic systems are fatigue-free, can function continuously, and are less prone to error [4]. Harvesting robots can be classified into two groups: fully integrated systems and subsystems used in harvesting missions, such as vision systems, control systems, and grippers.

Nevertheless, automated harvesting faces key challenges. Fruits and vegetables are delicate, and effective harvesting requires gentle interaction to avoid bruising, scarring, or other surface damage that reduces market value. Field conditions add complexity: unstructured environments, occlusions by foliage/branches, highly variable lighting, and a wide diversity of plant morphologies all complicate perception, planning, and manipulation [5–7]. With advances in perception and control, modern systems are increasingly capable of adapting to the variability in the size, shape, pose, and seasonal availability of agricultural products [5,6,8].

Early work in the 1980s–1990s established the vision of robotic fruit picking but was constrained by perception and actuation limits [9]. Over the past decade, rapid advances in computer vision and deep learning—combined with progress in soft and compliant manipulation—have revitalized the field and significantly improved detection, pose estimation, and grasp success rates [3,10], leading to more effective and innovative robotic system designs [11–15].

A critical aspect of harvesting robots involves the design and development of end-effectors that must grasp the fruit firmly and gently to avoid damage, thereby preserving quality [16]. Grippers directly interact with crops and thereby determining both the feasibility and efficiency of harvesting systems. Given the delicate nature of fruit, grippers must strike a fine balance between secure grasping and low surface stress. This has given rise to a spectrum of designs that can be usefully grouped into three broad categories: rigid [17,18], soft [19,20], and hybrid [21,22]. Rigid grippers offer precision and structural robustness, but risk damaging delicate tissues [23,24]. Soft grippers, typically using elastomeric and compliant bodies driven by different actuation technology, conform to fruit geometry and reduce contact pressure, though they may sacrifice precision or grasp stability [10,25]. Crucially, the adaptability and low-pressure nature of compliant grippers, a sub-category of soft grippers, offer a compliant design that makes them particularly suitable for grasping small or irregularly shaped produce, such as berries and clustered fruits, where traditional rigid systems might struggle [26–29].

Hybrid designs combine rigid structures with compliant interfaces or multi-modal actuation to trade off these advantages and limitations [10].

The bibliometric analysis (Figure 1) illustrates the research trend patterns in the field of agricultural robotic grippers and end-effectors (additional details of the bibliometric analysis are provided in Appendix A). In particular, the study examines the extent of attention devoted to different design features. Publication activity on agricultural grippers and harvesting robots has risen sharply since 2019 (Figure 1a). Designs utilizing soft and hybrid solutions are increasing rapidly, largely due to their safety in grasping delicate objects (Figure 1b). Designs span from single-finger (monolateral grippers such as suction devices) to multi-finger grippers; two- and three-finger grippers dominate, balancing

adaptability with control complexity (Figure 1c). With respect to actuation, while motor-driven systems account for the majority of designs, pneumatic actuators are also widely employed (Figure 1d). Moreover, most studies are tailored to specific crops—such as tomatoes, strawberries, cucumbers, apples, and grapes (Figure 1e). Even when the overall structures are rigid, soft materials are commonly integrated at the tips to minimize damage, highlighting the practical importance of hybrid designs (Figure 1f). Regarding sensory systems for protecting delicate agricultural produce, some studies employed vision-based systems (e.g., RGB cameras), others used tactile sensors at the tips (e.g., force or pressure sensors), and several adopted hybrid configurations combining both approaches, while a few have no sensory components altogether (Figure 1g).




Figure 1. Bibliometric analysis of agricultural grippers and end-effectors: (a) trend of publications from 2019 to 2025; (b) distribution of designs by gripper type (rigid, soft, and hybrid); (c) distribution of designs by finger count; (d) distribution of designs by actuation/transmission mechanisms; (e) frequency of target fruits and agricultural produce; (f) types of materials used for the tips of grippers/end-effectors; and (g) sensory system.

Among the different types of soft grippers, those with controlled adhesion offer unique and interesting features for gripping fragile objects such as fruit [30]. Adhesion refers to the attraction at the interface between two surfaces, generating shear stresses that are proportional to the normal pressure. At the same time, the normal force on the object's surface is much lower than that of conventionally actuated grippers, allowing for the manipulation of highly fragile objects. Adhesion is very effective for soft and deformable objects because the interface attraction automatically follows the object's deformations when the gripper structure is appropriately designed to be compliant. Among the controlled adhesion category, electroadhesion (EA) grippers represent a relatively underexplored yet promising direction for agriculture. EA grippers generate controllable adhesion by applying an electric field across patterned electrodes, inducing polarization in the target surface and producing attractive forces at low normal loads without any restriction on the target material nature. This enables gentle grasping with minimal mechanical squeezing

and offers three salient advantages for delicate produce: (i) gentle handling with low contact pressure; (ii) low power operation relative to pneumatic/vacuum systems; and (iii) design versatility for integration with both soft and rigid structures [31–33]. Despite these benefits, EA has seen a broader development in electronics/semiconductor handling (electrostatic chucks), textile manipulation, and micro-/meso-scale manipulation than in agriculture [34–36].

From the state-of-the-art analysis illustrated in Appendix A, no evidence was found for grippers that fully or partially integrate electroadhesion. This result demonstrates that the practical application of EA grippers in robotic systems for fruit harvesting is immature, and, to the authors' knowledge, no EA gripper has been tested on an autonomous robotic system for fruit picking in a real orchard.

However, some scientific works have shown how EA grippers can also be used for gripping fruits and vegetables, showing their ability to grasp fruits and vegetables in laboratory environmental conditions. In Section 3, the method used to identify such contributions will be presented, and the performance and the characteristic aspects of the proposed designs will be discussed together, along with the conditions and the target fruits considered for their testing.

The purpose of this paper is to demonstrate the current level of application of EA grippers in fruit and vegetable harvesting and to provide an overview of the advantages and disadvantages of using this technology and the key design aspects of such systems.

The remainder of this paper is organized as follows: Section 2 introduces EA grippers and explains their working principles. Section 3, as anticipated, shows the application of EA grippers for fruit grasping in detail. Section 4 provides a detailed discussion of the insights presented in the paper.

2. Electroadhesion Grippers

According to the classification proposed by G.K. Monkman [37], electroadhesive devices (EADs) are a type of astrictive grippers increasingly used for prehension tasks in research and industrial environments. They operate by exploiting an electrostatic field to adhere to both conducting and insulating materials. Historically, the term electroadhesion was coined in 1917 by A. Johnsen and K. Rahbek [38] to target the astrictive force that they experimentally observed at the contact interface between polished lithographic stone and metal surfaces when applying a potential difference between them. This phenomenon has been subsequently documented among different materials and was first investigated in an application study for aerospace purposes in 1968 by R.P. Krape [39]. Over the last two decades, EA technologies have gained renewed interest for their versatility and potential in robotics applications for a wide variety of industrial settings, ranging from logistics and automation [40,41] to the space environment [42,43].

2.1. Working Principle

The astrictive action of EA technologies relies on electrostatic binding forces generated by the electrostatic charge displacement field building up inside a dielectric medium, constituting the electro-active gripping layer of the EAD, upon electrical activation. Featuring capacitive architecture, electrostatic grippers exhibit either a monopolar or a bipolar configuration. Monopolar systems consist of a conductive plate, functioning as a single-pole electrode, and a dielectric layer, which separates the electrode from the object to be handled [44]. In this configuration, depicted in Figure 2(1), the target object is placed on another conductive, electrically grounded surface. When a potential difference in the kV range is supplied between the electrode plate and the conductive bench, opposing electrical charges accumulate on the dielectric layer of the gripper and the target object,

which is thus electrostatically attached. Relevant examples of EA technologies integrating monopolar grippers can be found in electrostatic chucks for silicon wafer handling [45] and the electrostatic gripping of fabric plies for automated garment assembly [44].




Figure 2. Sketch of the transversal section of a monopolar and an interdigitated bipolar gripper architecture, respectively, illustrating the working principle of electroadhesion grippers when grasping objects of different nature: (1) for conductive targets, opposing image charges appear at the dielectric–object interface; (2) for dielectric objects, dipole charges align with the applied electric field.

In the simplest configuration for a single-pole EA gripper, i.e., when the grasped object is an electrical conductor, the system can be approximated to a parallel-plate capacitor. A theoretical model is presented for this configuration for simplicity. Under the above-mentioned assumptions, the electrostatic retention pressure σ can be computed from the balance of energy stored in a parallel-plate capacitor, and expressed by the following:

$$\sigma = \frac{F}{A} = \frac{1}{2} \epsilon_0 \epsilon_r \left(\frac{V}{d} \right)^2 \quad (1)$$

where A is the interface area between the gripper surface and the object, F is the electrostatic force in the normal direction, V is the potential difference supplied to the single-pole electrode with respect to the electrical ground, d is the distance between the gripper electrode and the surface of the target object, and ϵ_0 and ϵ_r are the vacuum permittivity and the relative permittivity of the dielectric layer, respectively.

Bipolar electrostatic grippers, instead, present either a pair or a set of both positive and negative electrodes on the gripper's body. When a high voltage is applied between the alternating poles, the dielectric layer of the gripper is exposed to an electric field and polarized. In response to the dielectrics' polarization field, an induced charge distribution appears in the object to be grasped, and an attractive electrostatic force is generated between the two surfaces. For EA grippers exploiting Coulomb interactions, the charge distribution induced inside the contacting object differs according to the electrical nature of the target material: when the EAD is in contact with a conducting target, image charges are gathered at the conductor's surface, with opposing signs with respect to the polarization charges inside the gripper's dielectric layer. If a dielectric object is grasped, electric dipoles inside the material reorient themselves, with this process depending on the electrical, microstructural, and polarization properties of the material, and align with the electric field lines: macroscopically, regions with bond charges of opposing signs are displayed at

the dielectrics' interface, resulting in the buildup of the EA force. The electro-active layer thickness is typically in the range of a few tens of microns, and its electrical and mechanical properties have a significant influence on the magnitude of the electric field that the EAD can withstand, as well as the maximal adhesive force delivered. In general, it also acts as mechanical support for electrode manufacturing.

While monopolar EADs are associated with simpler design requirements and a reduced risk of dielectric breakdown and are useful for lifting conductive or semiconductive materials, they have proven to be inefficient or ineffective in gripping dielectric materials [33]. In bipolar devices, different arrangements of the alternating positive and negative electrodes have been reported as more effective in grasping selected objects, and, depending on the use case, geometric optimization has been demonstrated to be a valid strategy for achieving enhanced EA forces [46,47]. Some of the most popular bipolar EAD geometries feature concentric [48], spiral [49], and interdigitated or comb-shaped electrodes [50,51].

2.2. Advantages and Potential Applications

EA grippers are establishing themselves as fascinating alternatives to more traditional mechanical impactive techniques thanks to their capability of exerting continuous, controllable, and high shear forces on the target object without applying mechanical compressive stress. Recent research has demonstrated that electroadhesion can be easily integrated with other grasping technologies based on different physical mechanisms, such as gecko-inspired surfaces exploiting van der Waals interactions [43,52] and suction cups [53], for the realization of technologically advanced adhesive interfaces. As reported by J. Guo et al. [33], the main key benefits that make electrostatically enhanced grippers competitive grasping methods for a broad range of industrial settings are the following:

- Versatility and adaptability to both conducting and insulating materials, and porous and rough surfaces;
- Capability of operating in harsh conditions, such as dusty environments, at low pressures, and in space;
- High shear forces without applying mechanical compressive stresses;
- Lightweight, thin-film, and compliant functional layers;
- Low-cost constitutive materials, easy process via additive manufacturing and printed electronics methods, and simple electrical control components;
- Low energy consumption (<30 mW) despite a relatively high voltage, with typical currents in the μA range measured for EADs;
- Safe handling of fragile, delicate, and soft objects, without exerting any mechanical compressive pressure.

Due to their operational advantages, EA grippers fabricated with a soft functional layer (i.e., polymer-based materials such as polydimethylsiloxane (PDMS), silicone epoxy, acrylic) appear highly promising in the context of fruit and crop harvesting, alongside more traditional robotic end-effectors. End-effectors integrating soft EADs enable the harvesting and pick-and-place operations of delicate crops while lowering the risk of damage. Based on gripper mechanics, multiple strategies have been successfully adopted aimed at the optimization of the grasping action and the effectiveness of the release tasks [54,55]. The most relevant EA gripper solutions identified through a systematic literature review in the field of robotic fruit harvesting are discussed throughout the following section.

3. EA Grippers' Grasping Capability on Fruit

This section presents the results of the state-of-the-art research on EA grippers applied to fruits and their grasping capability. To identify relevant studies, the following keywords were used: (electroadhes OR electro-adhes OR electro adhes OR EAD) AND (grip*). The

search was conducted in the Scopus database, restricted to publications from 2019 to 2025 and limited to English-language articles. After applying the keywords, 53 papers were found. Each instance was manually inspected to verify whether the presented gripper successfully grasped the fruits, and only those that did were included in this study.

3.1. EA Gripper 1: A Variable Stiffness Electroadhesive Gripper Based on Low-Melting-Point Alloys

Researchers in [56] present an EA soft gripper with a variable stiffness. It features a three-fingered design, and its actuation is a combination of a pneumatic drive and electrical resistance for heating the material. There are two steps: first, pneumatic pressure is used to bend the fingers to have the same shape as the object. Second, a high voltage is applied to the low-melting-point alloys (LMPAs) to generate electrostatic adhesion to ensure a secure grasp. The gripper is bilateral, which means that it uses three fingers to grasp an object from opposing sides. The main structure of the gripper is composed of silicone (Ecoflex™ 00-30, Smooth-On Inc. Macungie, PA, USA) and LMPAs. The EA patch is made of LMPAs as electrodes, a silicone dielectric layer, and electric resistance wires. The effective electrode area is 111 mm by 20 mm, and electrode spacings are tested for 2 mm, 3 mm, and 4 mm. The dielectric layer thickness is tested with the dimensions of 0.4 mm, 0.6 mm, 0.8 mm, and 1 mm, with 0.6 mm being chosen for the final design for safety and performance. The EA patch is powered by a 3.2 kV, 3.8 kV, and 4.4 kV DC voltage for 2 mm, 3 mm, and 4 mm spacings, respectively. The results, tested on paper substrate, demonstrate that the maximum tangential EA force that the gripper can generate is 0.512 N and the maximum normal force is 0.10 N at an activation voltage of 4.4 kV. For the three-fingered gripper, 20 kPa to 60 kPa with 10 kPa increments of pressure is applied. The maximum grasping force is 10.74 N for envelope grasping and 6.78 N for fingertip grasping. The total weight of the gripper is 340.2 g. The tests were conducted under conditions with a temperature of 22.6 °C for the repeatability test and 27.7 °C for the grasping force test. Additionally, the humidity levels were 55% for repeatability and 60% for the grasping tests. For the tests, the researchers demonstrate the grasping capability of the gripper on target fruits such as apples (202.2 g weight and 60 kPa pressure) and oranges (129.67 g weight and 30 kPa pressure), with the three-fingered gripper mounted on a frame. They display images of tests with various objects where the gripper is in a closed position, providing grasping. This study has a technology readiness level (TRL) of 4, since it has experimental validation.

3.2. EA Gripper 2: Additive Manufacturing of Flexible 3D Surface Electrodes for Electrostatic Adhesion Control and Smart Robotic Gripping

This study [57] demonstrates a soft gripper that combines mechanical shear and EA forces. It features two flexible fingers on a bi-stable support that enables them to fold and grasp objects. The fingers are covered with a multitude of flexible composite electrodes arranged in a hair-like manner that conforms to the object and generates electroadhesion. The working principle is as follows: first, the flexible surface and the hair-like electrodes deform to increase the contact area, which generates the shear force; then, voltage is applied to the electrodes, creating an electrical field that increases the adhesive force. The main structure of the gripper is made of an elastomeric resin (Elastic Resin, Formlab Inc. Somerville, MA, USA). The paper considered the grasping of several objects but specifically mentioned fresh fruits and conducted trials using a tangerine. The flexible hair-like electrodes are made of a soft elastomer with conductive metal nanowires obtained from copper and silver.

The results provided in the paper demonstrated that the flexible hair-like electrodes provide up to 72% better gripping force when the voltage is applied (tested on PP 40 mm diameter cylinder object). Tests with the electrodes slanted at different angles showed

that the biggest increase in lateral adhesion with 3 kV voltage was 162% for the smallest angle of 15°, and the lowest was 20%, with the largest angle of 75°. The gripper with electrodes slanted by 30° recorded a grasp with a force up to 1 N when the voltage was on and 0.6 N when the voltage was off. Tangerine grasping was successfully accomplished with electrodes having a width of 1.5 mm, a length of 8 mm, a slanted angle of 30°, and an activation at 3 kV. This study has a technology readiness level (TRL) of 4, since it has experimental validation.

3.3. EA Gripper 3: Delicate Yet Strong: Characterizing the Electroadhesion Lifting Force with a Soft Gripper

The soft gripper in this research [58] uses dielectric elastomer actuators (DEAs) and electroadhesion, featuring two fingers. The actuation is based on the DEA's ability to bend in reaction to an applied voltage, and so the gripper is driven electrically. The gripper is designed to generate both bending motions as produced by the DEA fingers and electroadhesion at the fingertips. Both are controlled by a single monolithic structure using the same electrodes. The DEAs bend fingers to the object and conform the shape. A high voltage creates an electric field, and it leads to EA forces that grasp the object. The gripper's structure is a 280 µm thick silicone-based membrane, including two passive layers and a central dielectric membrane fabricated by blade casting Sylgard 184 and 186 polydimethylsiloxane (PDMS) by Dow Corning Inc. Freeland, MI, USA, respectively. The gripper has been presented with cherry tomato holding and a heavy test load. The EA patch is composed of compliant electrodes which consist of a PDMS + carbon mixture, and it is integrated into a multilayered structure. The area of electroadhesion is 1 cm². This study considered two electrode designs with a width of 0.5 mm and electrode spacings, respectively, of 2.1 mm and 0.5 mm. a 3.5 kV DC voltage. The gripper is powered by a 3.5 kV DC voltage. The results showed that the lifting force was 16 N, shown on a cylinder covered with paper-based tape, but was strongly dependent on the object's shape, and it can lift a 10 g cherry tomato. The total weight of the gripper is 1.5 g; thus, it can lift loads over 1000 times its own weight. Furthermore, it shows a high durability, surviving over 100 cycles at high loads. This gripper has a technology readiness level (TRL) of 4, since it has experimental validation.

3.4. EA Gripper 4: Design and Development of a Variable Structure Gripper with Electroadhesion

This study [55] proposes a soft gripper having three fingers and a combined pneumatic-electrical actuation. The fingers are made of a silicone rubber (HY-E630, Shenzhen Hongye Silicone, Shenzhen, China) and can conform to the object's shape once pneumatically activated. An adjustment to objects of different sizes and curvatures is made by the electric motor, which regulates the opening angle of the fingers. EA films with coplanar interdigitated copper electrodes embedded in an insulating material are placed at the bottom of the fingers to enhance grip and allow the gripper to hold flat objects upon high-voltage activation. The paper mentions that the gripper has potential for fruit picking in the agricultural industry and shows its grasping ability on tomatoes (80 mm diameter and 297 g). The soft finger has dimensions of 144 mm × 40 mm × 25 mm. Each air chamber is 4 mm wide and 6 mm thick. The maximum finger bending angle is 130.2° at a pressure of 60 kPa. The EA film thickness is 0.05 mm, and its effective area is 3410 mm². Shear force testing of the EA patch was performed with an activation voltage ranging from 0 to 4 kV, which provided a maximum of 2.58 N at 4 kV on an adsorption substrate without specifying the material base. By varying the opening angle of the fingers from 0° to 90°, the gripper enables them to grasp spherical objects with a diameter in the range 0–482 mm. The maximum load that the gripper can reach is 10.91 N. All the tests have been conducted with a relative humidity of 63%, a temperature of 25.3 °C, and an ambient pressure of 1020.5 hPa.

The technology has reached a TRL of 4, which means it has a working lab prototype. However, it requires further more development before it can be used for industrial trials.

3.5. EA Gripper 5: EA-SoGripper: Electroadhesion-Stiffening Self-Adaptive Soft Robotic Gripper

This paper [59] introduces a soft gripper with two compliant fingers that uses electroadhesion to achieve a variable stiffness and self-adaptability. Finger opening/closing is regulated by an electrical actuator that drives a ball screw transmission. Each compliant finger consists of a band that is tensioned by a spring, allowing it to conform to the object to be grasped with a given compliance. To regulate the compliance, each finger also comprises an EA clutch, which acts on the band to exclude the spring once the desired level of tension is achieved. Unlike the previous designs, in this gripper, electroadhesion is not used to directly retain the object. Nonetheless, this concept allows us to securely clamp objects of various sizes and shapes. It can support loads up to 400 g while adapting its stiffness by controlling the EA clutch voltage. The EA clutch consists of two electrodes separated by a dielectric layer. The dielectric layer is a polyimide (PI) film with a thickness of 0.025 mm and a relative permittivity of 3.4; the high-voltage electrode is made of a copper tape with dimensions 80 mm × 30 mm; the low-voltage electrode is made of an aluminized polyethylene terephthalate (PET) film with dimensions 250 mm × 30 mm and a thickness of 0.075 mm. The weight of the gripper is 376 g and the maximum weight that can be lifted is 400 g. The tunable grasping stiffness ranges from 0.06 to 1.25 N/mm. Finger closing/opening requires 1 s. EA clutch activation is achieved with an AC square wave with a fixed frequency of 20 Hz; an amplitude of 400 V was sufficient to consistently grasp various objects including grapes. This study has a technology readiness level (TRL) of 4, since it has experimental validation.

3.6. EA Gripper 6: Electroadhesion Suction Cups

The study in [53] investigated a monolateral gripper exploiting the combined effect of electroadhesion and vacuum suction. The gripper consists of a circular EA membrane made of silicone elastomers having an Elastosil 2030 (Wacker Chemie AG, Munich, Germany) PDMS bottom layer and a Sylgard 184 (Dow Corning Inc. Freeland, MI, USA) PDMS upper layer. It features embedded interdigitated electrodes that are screen-printed on the PDMS bottom layer using a silicone-based electrically conductive silver ink, with a central pillar for external force application. Upon electrical activation with a high voltage, the EA membrane zips onto the object surface, forming an airtight seal that generates a vacuum suction effect once the object is lifted by pulling the central pillar. In this way, the strengths of both vacuum suction and electroadhesion are used, providing a very compact solution. The EA membranes had a diameter in the range of 3.3–5.3 cm and electrodes with a spacing of 0.5 mm. The gripper was operated at 1–3 kV DC for normal operations and at 3 kV AC with a 5 Hz frequency occasionally for charge neutralization. This gripper successfully handled objects of around 5–20 cm in size and 1 kg in weight, including grapefruits, oranges (230 g), jars and plastics. It enabled reliable grasping even on the rough and slippery surfaces of fruits, where conventional EA pads struggle. The addition of suction increases the holding force significantly (up to 18 N on PMMA substrate of roughness $S_q = 1.9$ at a vertical speed of $s = 0.2 \text{ mm s}^{-1}$ and an applied DC voltage $V = 3 \text{ kV}$, with a membrane of 5.3 cm diameter) compared to EA alone ($\sim 0.4 \text{ N}$), especially for rough and moist surfaces with near-atmospheric-vacuum pressures ($\approx -101 \text{ kPa}$). This gripper exhibited a high reliability, with more than 11,000 pick-and-place cycles, rapid release ($< 0.15 \text{ s}$ with valve), and energy-efficient operation. Its development corresponds to a technology readiness level (TRL) of 4, demonstrating a validated laboratory prototype on a robotic arm, although it is not yet mature enough for industrial trials.

3.7. EA Gripper 7: Peeling in Electroadhesion Soft Grippers

A peeling-controlled EA gripper is introduced in [54]. It features a bilateral design with two EA films with a trapezoidal shape, which are mounted on lightweight PLA and PMMA structures that can be brought together or separated by an electric motor. Each EA film is made of Sylgard 184 silicone elastomer with embedded interdigitated electrodes realized in carbon-black and PDMS. It has a length of 45 mm and a thickness of 90 μm , features an active area of 20×30 mm, and is energized through a high-voltage bipolar AC signal with a frequency of 1 Hz and an amplitude in the range of 1–4 kV. The gripper can switch between maximum adhesion for secure grasping and rapid detachment for proper release by adjusting the peeling angle through the electric motor. Demonstrations on cherry tomatoes, limes, and mangoes show that this design combines a high holding force with a reliable release. For smaller angles ($\theta = 0^\circ$), adhesion forces reached over 1 N on a substrate without specifying the material base, corresponding to shear stresses of about 2–3 kPa. This enables the grasping of fruits ranging from a 10 g cherry tomato to a 600 g mango. At higher angles ($\theta > 30^\circ$), the effective adhesion dropped sharply, resulting in a rapid release (i.e., within 300 ms). The power consumption of EA films was less than 1 W. This technology can be positioned at an intermediate TRL (4), presenting a promising outcome for fruit harvesting and other soft robotic applications.

3.8. EA Gripper 8: Hybrid Soft Electrostatic Metamaterial Gripper for Multi-Surface, Multi-Object Adaptation

In [60], the authors propose a hybrid soft electrostatic metamaterial (SEM) gripper offering a localized compliance that enables the gripper to conform to surface irregularities and achieve a high directional adhesion (up to 65 times larger with respect to when in the off state). Thanks to this feature, the gripper can grasp objects of various sizes and shapes without active adjustments. In addition to increasing adhesion, this design also mitigates residual peeling force, thereby enabling a faster release. The SEM gripper is soft and is demonstrated primarily in a two-finger configuration. It consists of an EA film made of a urethane elastomer (Vytaflex 20, Smooth-on Inc., Macungie, PA, USA), having embedded interdigitated electrodes realized in carbon-black adhesive sheet (ARclad 8006, Adhesive Research, Glen Rock, PA, USA), with a PET backing (L480.385, Rausch Packing, Meisterschwanden, Switzerland) for structural support and with properly engineered cuts to create the localized compliance. The testing of SEM grippers featuring grids of $10 \text{ mm} \times 15 \text{ mm}$, interdigitated electrodes in either 4×4 or 2×4 arrays, and activated to generate an electric field in the range of 0–6 kV/mm resulted in adhesion stresses up to 55.6 kPa and pulling forces of up to 33 N on a 77 mm diameter pipe, enabling the lifting of payloads up to 17 times its own weight. Successful gripping was demonstrated for apples (191 g weight), coffee bags, tape rolls, pamphlets, and PVC pipes of various diameters. The SEM gripper is driven at a low speed and exhibits a high efficiency. Overall, the SEM gripper showed laboratory validation, but it is still not industrially deployed, placing it at a TRL in the range of 4–5.

3.9. EA Gripper 9: Electroadhesion Zipping with Soft Grippers on Curved Objects

The study in [50] uses the zipping effect of EA grippers to address the challenge of curved object grasping with simple devices, showing the capability of successfully catching a range of fruits (grapefruit, mango, lemon). The zipping EA gripper comprises two fingers, each coinciding with an EA film made of Sylgard 184 PDMS, with embedded interdigitated electrodes realized in carbon-black (Ketjenblack EC-300J by Nouryon Functional Chemicals B.V., Radnor, PA, USA) and PDMS mixture. Thanks to the zipping effect, the same electrical activation that generates adhesion also makes the fingers autonomously wrap around

curved objects, thereby increasing the contact area without requiring any additional means of actuation. The interdigitated electrodes feature a width of 0.5 mm, a pitch of 0.7 mm and an active area of 432 mm². The silicone elastomer dielectric layer has a thickness of 45 µm. Electrical activation is achieved with AC or DC voltages with an amplitude of up to 4 kV. Experimental testing has been performed on different cylindrical objects with radii in the range of 30–45 mm and with different coatings (e.g., paper, copper, PVDF-TrFE-CTFE, PLA), demonstrating a high adhesion force (larger than 1 kg per gram of gripper mass), fast response time (about 0.4 s for zipping and 0.3 s for unzipping), and high efficiency (a power consumption less than 1 W). A smooth performance over AC/DC cycles was shown in the reliability test. Also, this gripper stands at TRL 4, as it has not been tested in the field.

3.10. EA Gripper 10: Enhancing Compliant Gripper Performance: Exploiting Electroadhesion to Increase Lifting Force over Grasping Force

This work [61] demonstrates a soft gripper that uses highly compliant FinRay fingers with custom EA pads to improve lifting force and minimize object squeezing. The gripper features a bilateral configuration with two fingers moved by an electrical actuator inside the end-effector of a robot arm. FinRay fingers are made of a very soft urethane rubber, which enables them to conform to the object shape, with minimal contact pressures that are unable to provide retention. The holding force is indeed obtained upon EA pad activation. The EA pad comprises a PIT1N/210 polyimide (PI) film (Caplinq, Assendelft, The Netherlands) with 25 µm thickness as the main dielectric layer, silver interdigitated electrodes (Anapro DGP 40LT-15C, Cheongwon-gun, Republic of Korea) deposited by ink-jet printing, and a blade-casted backing encapsulation made of a silicone elastomer (Silpuran 6000/05, Wacker Chemie AG, Munich, Germany) 100 µm. A double-coated polyester silicone tape with a 94 µm thickness is used for bonding the EA pad to the FinRay finger. The EA pad electrodes have a width of 400 µm, a gap of 400 µm, an active area of 10 mm × 70 mm, and are powered by a DC voltage up to 4 kV. The experimental and simulation results prove that the gripper can reduce squeezing forces by 77% compared to a commercial FESTO FinRay with the same lifting capacity. The gripper also demonstrated a great capability in the grasping of lemons, mandarins, apples, and tomatoes. This study has a technology readiness level (TRL) of 4, as the gripper has not been tested in a real environment, such as a crop or an industrial field.

The comparison of the given grippers can be found in Table 1. They are investigated under 9 categories.

Table 1. Comparison of the grippers mentioned in Section 3.

Gripper	Prehension Mechanism	Fruit Type	Electrode Geometry	Dielectric Material	Active Area (cm ²)	Max. Voltage (kV), AC/DC	Overall Load Capacity (kg)	Response Time (s)	TRL
1	EA + Pneumatic	Apple, orange	Interdigitated electrode; 2, 3, and 4 mm spacing	Ecoflex™ 00-30 Silicone	22.2	4.4	0.2	-	4
2	EA + Gecko Adhesion	Tangerine	Electrodes with width = 1.5 mm, length = 8 mm, patterned on rectangular and zigzag-T 3D pillars (slanted 15–75°)	Elastomer Resin	Up to 5	3 (DC)	-	-	4
3	DEA + EA	Cherry tomato	Interdigitated, width = 0.5 and 2.1 mm and gap = 0.5 mm	Sylgard 186 PDMS	1	3.5 (DC)	1.63	<5	4
4	EA + Pneumatic	Tomato	Interdigitated electrodes	Silicone rubber (HY-E630)	34.1	4 (DC)	1.11	-	4

Table 1. Cont.

Gripper	Prehension Mechanism	Fruit Type	Electrode Geometry	Dielectric Material	Active Area (cm ²)	Max. Voltage (kV), AC/DC	Overall Load Capacity (kg)	Response Time (s)	TRL
5	EA + Pneumatic stiffening	Cherry tomato, apple, lemon, grape	Two electrodes arranged in an electrostatic clutch	PI film	24	1 (AC)	0.4	2	4
6	EA + Passive Vacuum	Grape, orange	Interdigitated circular electrodes, 0.5 mm spacing	Elastosil 2030 PDMS	22	3 (DC)	1	0.2 (pick)/0.15 (release)	4
7	EA	Cherry tomato, lime, mango	Interdigitated electrodes	Sylgard 184 PDMS	6	4 (AC)	0.6	<0.3 s (release)	4
8	Hybrid EA + Metamaterial Adhesion	Apple	Interdigitated electrodes, width = gap = 1 mm in square metamaterial pattern (4 × 4 grid, 10 mm × 15 mm)	Urethane upper layer PET bottom layer	6	4–6 (DC)	0.33	-	4
9	EA	Grape, mango, lemon	Interdigitated, width = 0.5 mm and gap = 0.7 mm	Sylgard 184 PDMS	4.32	4 (DC and AC)	-	0.7	4
10	FinRay + EA	Lemon, tangerine, apple, tomato	Interdigitated, width and gap = 0.4 mm	Caplinq PIT1N/210 PI film	7	4 (DC)	0.384	-	4

Due to copyright restrictions, the images corresponding to grippers 3, 4, and 5 could not be included in Figure 3.




Figure 3. Grippers from the reviewed works: (a) EA gripper 1, (b) EA gripper 2, (c) EA gripper 6, (d) EA gripper 7, (e) EA gripper 8, (f) EA gripper 9, (g) EA gripper 10.

4. Discussion

This review aims to assess the current state and potential of EA grippers for fruit and vegetable harvesting, focusing on their advantages, limitations, and key design considerations within the context of agricultural robotics. Although no evidence of testing EA grippers in real environments, such as fields or greenhouses, has been documented to date, a systematic review of the current state of the art has revealed several promising cases of robotic grippers based on the EA working principle. Ten potential gripping systems are analyzed in this work that successfully integrate electroadhesion and address fruit handling applications.

By comparing the gripping technologies considered in Section 3, it becomes apparent that there is a great dissimilarity between the experimental conditions occurring in the evaluation of the effectiveness of the different EA grippers in fruit grasping operations, since different setups, functional materials, dimensions, and surface textures of the targeted crops are involved. It is thus impossible to establish a quantitative and equal comparison between them and to identify optimal candidates in the context of fruit harvesting and picking. Yet some qualitative performance considerations can be drawn from the data and the design details reported for the considered gripping systems. This analysis provides a solid starting point for the development of EA grippers specifically designed for crop handling in all of its aspects, from picking operations to harvesting in greenhouses and in real orchards. The main design features and applicability conditions emerging from this methodical review are listed below:

- Fruit selection: apples, tangerines, oranges, limes, lemons, tomatoes, cherry tomatoes, mangoes, and grapefruits are considered and compared in this analysis.
- Load capacity range: from 10 g (~0.1 N) for cherry tomatoes to 600 g (~6 N) for mangoes.
- Demonstrated adaptability to a variety of target shapes: thanks to their modular structure, the hybrid EA grippers reviewed display an improved gripping capability of flat surfaces, convex and concave objects, and targets having more complex shapes.
- Demonstrated grasping ability for a variety of contact interfaces: EADs prove to be effective on rough surfaces too.
- Application of minimal compressive stress and high EA shear force on the fruits to be grasped.

The limited number of real-world evaluations of electroadhesive (EA) grippers in agricultural environments can be attributed to several practical constraints. First and foremost, the development of controlled adhesion gripping systems remains relatively limited compared to other grasping technologies. Consequently, the application of such systems to the specific context of fruit and vegetable harvesting further narrows the number of relevant use cases.

Additionally, the lack of detailed information on the shear and normal stress performance of electroadhesive devices (EADs) for individual fruit types, as well as the absence of studies focused on optimizing the EAD performance for specific target fruits, hinders further progress in this field. Essential design parameters for these gripping systems are therefore still missing. Moreover, several factors known to limit EAD performance, such as moisture and dust, have not been quantitatively assessed in the context of fruit and vegetable applications, restricting the development of grippers specifically engineered to mitigate these environmental limitations.

Another aspect insufficiently addressed in the literature, and consequently limiting EA gripper development, concerns the influence of soft or rigid backing structures supporting the EAD. These structures inherently affect the finger stiffness and, consequently, its ability to conform closely to the target object, a crucial condition for maximizing the electroadhesive effectiveness. In the context of fruit and vegetable handling, this aspect is fundamental not only for designing efficient gripper systems but also in relation to robot dynamics, as the chosen backing structure must ensure a secure grasp during harvesting and picking operations at high dynamics. Furthermore, some EA gripper designs entirely omit a backing structure, which severely limits their applicability in cultivated environments due to the inadequate grip stability of robots' dynamics, which has to be competitive with human harvesting productivity.

Finally, many existing EA grippers remain at an early research stage, primarily validated under controlled laboratory conditions to demonstrate feasibility before integration

into fully operational robotic systems. Collectively, these factors have contributed to the lack of real-world testing reported in the literature to date.

A more general review [62] on robotic end-effectors for fruit and vegetable picking highlighted the key role that automated crop harvesting solutions will continue to have in the development of agricultural robots to be used in greenhouse and field applications. According to them, the current and future tendencies constituting fundamental advancement steps of the next-generation robotic end-effectors for fruit harvesting are the following: (1) light weight and adaptability; (2) development of bionic manipulators; (3) integration with multi-sensing technologies; (4) intelligent control systems; (5) flexibility; and (6) modular and reconfigurable design. From this perspective, testing and introducing custom-designed EA and hybrid EA grippers in greenhouses and field environments would certainly address these topics and lead to competitive advantages apart from intelligent control system. In fact, by limiting the discussion to the sole EA-based end-effectors, the reviewed grippers that integrate EA technologies effectively enable the following:

- (1) Lightweight, adaptive end-effectors: EA interfaces can lift up to 33 N, with payloads up to 17 times their own weight [60]. The structural optimization of hybrid EA grippers leads to an improved conformability to the fruit and introduces mechanical stiffness, contributing to prehension force enhancement. Their electrical power source can be located externally or miniaturized, with a minimal contribution to the weight of the end-effector. Adaptability to a wide variety of shapes and surface textures allows for a reduced damage risk, higher rates of success for picking operations, and improved efficiency.
- (2) In the context of bionic manipulators, the soft FinRay gripper reviewed in Section 3.10 represents a promising solution capable of conforming to delicate objects, including crops, without damaging them, by integrating a soft urethane backing structure and a stable grip through electroadhesion.
- (3) Due to their capacitive architecture, EADs will soon integrate proprioceptive capabilities, which hold the potential for proximity, contact, and slipping detection, with no need for bulky hardware for additional sensors and vision systems [63].
- (4) The combination of flexible materials for grippers is the new frontier in robotic handling and can compensate for the occurrence of excessive stress, which is frequently observed in traditional mechanical systems instead. This would further reduce the risk of crop damage by compression. All the systems considered in this analysis feature a soft element, either using a thin-film EAD system connected to a mechanical actuator promoting contact and release, or a hybrid design with a soft backing structure.
- (5) Compared to traditional end-effectors that achieve grasping by exploiting fixed components such as mechanical or pneumatic-driven fingers, EAD's integration on any pre-existing structure does not represent a problem, and an enhanced performance is often achievable by simply adding a conformable backing layer. This drastically favors system modularity and reconfigurability and aligns well with the market demand.

As a final remark, future developments of EA grippers would greatly benefit from the definition of optimal characteristics tailored to specific prehension tasks, such as the EAD active area, driving signal, functional material, electrode design, and, for hybrid EA grippers, the backing structure. This can be achieved through an experimental campaign by measuring the electroadhesive shear stress that EADs can exert upon contact with different crop varieties. Another topic deserving further exploration is the investigation of the mechanical coupling between the EAD layer and the backing structure of the gripper, aimed at conforming to and maximizing the contact area with the fruit to be grasped, which is expected to lead to improved electroadhesion performances. Finally, testing EA

gripper prototypes both in real fields and in greenhouse environments would provide additional insights that can dramatically advance the development and TRL of the current technologies addressing automated fruit harvesting operations.

Author Contributions: Conceptualization, E.M., R.V. and G.B.; methodology, E.M., R.V. and G.B.; formal analysis, T.I.O., E.M., S.M.K., C.S., A.F., R.V. and G.B.; writing—original draft preparation, T.I.O., E.M., S.M.K., C.S. and A.F.; writing—review and editing, T.I.O., E.M., C.S., R.V. and G.B.; supervision, R.V. and G.B.; project administration, E.M., R.V. and G.B. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded under the National Recovery and Resilience Plan (NRRP), Mission 04 Component 2 Investment 1.5—NextGenerationEU, Call for tender n. 3277 dated 30 December 2021 Award Number: 0001052 dated 23 June 2022.

Data Availability Statement: No new data were created or analyzed in this study.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

EAD	Electroadhesion Device
TRL	Technology Readiness Level
EA	Electroadhesion

Appendix A

Appendix A.1

Search string of the initial search: (harvesting robot* OR harvesting robots*) AND (gripping* OR End effector* OR gripper* OR grippers*) AND (harvesting* OR picking* OR grasping* OR harvest* OR fruit harvesting* OR crop* OR grip* OR grasp*) AND (fruit* OR fruits* OR agricult* OR strawberries* OR apple* OR oranges* OR cucumbers* OR peppers* OR peaches* OR nectarines* OR tangerines* OR mandarins* OR clementines* OR plums* OR sloes* OR tomatoes* OR lychee* OR watermelon* OR blueberries* OR berries* OR kiwi*) AND (review*).

Appendix A.2

The objective of this trend analysis is to identify and evaluate research patterns in the field of agricultural robotic grippers and end-effectors. In particular, the study examines the extent of attention devoted to different design features and technological advancements of robotic grippers within agricultural applications, based on peer-reviewed academic sources. To guide the analysis, the following research questions were formulated:

- (1) Is the end-effector/gripper rigid, soft, or hybrid?
- (2) How many fingers does the gripper employ?
- (3) What is the actuation method (e.g., pneumatic, motor-driven, tendon-driven, etc.)?
- (4) What materials are used at the gripper tips?
- (5) Which fruits or agricultural products have been targeted for grasping or harvesting?
- (6) Which sensory systems are used to protect delicate agricultural produce?

These questions shaped both the development of the search string and the selection criteria for relevant studies. The dataset was refined using the PRISMA methodology, which involves the sequential stages of identification, screening, eligibility, and inclusion. This rigorous approach ensured that only relevant and high-quality studies were retained.

To qualify for inclusion, each publication had to present the design of at least one gripper or end-effector explicitly tested on fruits or other agricultural produce.

The classification of grippers proceeded in three stages. First, the structural type of the gripper was identified as rigid, soft, or hybrid. Second, the number of fingers and the type of actuation/transmission mechanism were recorded. Third, the tip material was categorized. Since the material data were highly diverse, they were consolidated into five major groups to enable effective visualization:

- (1) Elastomers/Rubbers: silicone rubber, thermoplastic polyurethane (TPU), polyurethane elastomer (PUE), thermoplastic elastomer (TPE), generic thermoplastic polymers (TP), and polydimethylsiloxane (PDMS);
- (2) Plastics/Polymers: acrylonitrile butadiene styrene (ABS), polylactic acid (PLA), polylactic polymer (PL), polyethylene terephthalate (PET), polyvinyl chloride (PVC), and generic polymers;
- (3) Foams and Sponges;
- (4) Hydrogels/Soft Materials: hydrogels and polymeric membranes;
- (5) Metals.

A similar categorization approach was applied to agricultural produce, grouping crops into eight broad commodity-based classes:

- (1) Temperate fruits: apple, pear, grape, peach, and kiwi.
- (2) Tropical/subtropical fruits: banana, mango, papaya, pineapple, avocado, litchi, and dragon fruit.
- (3) Citrus fruits: orange, lemon, tangerine, and citrus.
- (4) Vegetables: tomato, cherry tomato, tomato bunches, cucumber, eggplant, pepper, chili pepper, sweet pepper, radish, potato, paprika, asparagus, and broccoli.
- (5) Berries: strawberry, blueberry, blackberry, and raspberry.
- (6) Nuts/Seeds: chestnut burr, nut, seeds, safflower, sunflower, and hemp plant.
- (7) Industrial crops: cotton, cotton plant, cotton ball, and tobacco leaves.
- (8) Mushrooms/Fungi: mushroom, morel, and *Agaricus bisporus*.

The same rationale of categorization was applied to define the sensory system used to protect agricultural produce, grouping them into four broad categories:

- (1) Vision-Based: Designs that utilize vision systems such as RGB cameras.
- (2) Sensor at the tip: Designs that employ various types of sensors at the contact interface of the gripper or end-effector to protect the produce. These include photoelectric/photonic, pressure/force, triboelectric, piezoelectric, PVDF, magnetic/electromagnetic, capacitive, ultrasonic sensors, etc.
- (3) Hybrid: Designs that incorporate both vision-based systems and sensors at the tip.
- (4) None: Designs that do not employ any of the aforementioned sensory solutions for protection.

The initial search yielded 1033 publications, from which only 275 studies met the inclusion criteria after refinement. These studies formed the final dataset used for bibliometric and trend analysis.

References

1. United Nations Department of Economic and Social Affairs (UN DESA). World Population Projected to Reach 9.8 Billion in 2050, and 11.2 Billion in 2100. Available online: <https://www.un.org/development/desa/en/news/population/world-population-prospects-2017.html> (accessed on 19 September 2025).
2. Food and Agriculture Organization of the United Nations (FAO). Global Agriculture Towards 2050. 2009. Available online: https://www.fao.org/fileadmin/templates/wsfs/docs/Issues_papers/HLEF2050_Global_Agriculture.pdf (accessed on 19 September 2025).

3. Elfferich, J.F.; Dodou, D.; Della Santina, C. Soft Robotic Grippers for Crop Handling or Harvesting: A Review. *IEEE Access* **2022**, *10*, 75428–75443. [CrossRef]
4. Liu, Y.; Hou, J.; Li, C.; Wang, X. Intelligent Soft Robotic Grippers for Agricultural and Food Product Handling: A Brief Review with a Focus on Design and Control. *Adv. Intell. Syst.* **2023**, *5*, 2300233. [CrossRef]
5. Zeeshan, S.; Aized, T.; Riaz, F. In-Depth Evaluation of Automated Fruit Harvesting in Unstructured Environment for Improved Robot Design. *Machines* **2024**, *12*, 151. [CrossRef]
6. Huang, Y.; Xu, S.; Chen, H.; Li, G.; Dong, H.; Yu, J.; Zhang, X.; Chen, R. A review of visual perception technology for intelligent fruit harvesting robots. *Front. Plant Sci.* **2025**, *16*, 1646871. [CrossRef]
7. He, Z.; Liu, Z.; Zhou, Z.; Karkee, M.; Zhang, Q. Improving picking efficiency under occlusion: Design, development, and field evaluation of an innovative robotic strawberry harvester. *Comput. Electron. Agric.* **2025**, *237*, 110684. [CrossRef]
8. Wang, W.; Li, C.; Xi, Y.; Gu, J.; Zhang, X.; Zhou, M.; Peng, Y. Research Progress and Development Trend of Visual Detection Methods for Selective Fruit Harvesting Robots. *Agronomy* **2025**, *15*, 1926. [CrossRef]
9. Sarig, Y. Robotics of Fruit Harvesting: A State-of-art Review. *J. Agric. Eng. Res.* **1993**, *54*, 265–280. Available online: <https://www.sciencedirect.com/science/article/abs/pii/S0021863483710206> (accessed on 19 September 2025). [CrossRef]
10. Navas, E.; Fernández, R.; Sepúlveda, D.; Armada, M.; Gonzalez-De-santos, P. Soft grippers for automatic crop harvesting: A review. *Sensors* **2021**, *21*, 2689. [CrossRef]
11. Han, W.; Luo, J.; Wang, J.; Gu, Q.; Lin, L.; Gao, Y.; Chen, H.; Luo, K.; Zeng, Z.; He, J. Design of a Chili Pepper Harvesting Device for Hilly Chili Fields. *Agronomy* **2025**, *15*, 1118. [CrossRef]
12. Yu, Y.; Xie, H.; Zhang, K.; Wang, Y.; Li, Y.; Zhou, J.; Xu, L. Design, Development, Integration, and Field Evaluation of a Ridge-Planting Strawberry Harvesting Robot. *Agriculture* **2024**, *14*, 2126. [CrossRef]
13. Sytsma, M.; van Marrewijk, B.M.; Tielen, T.; Vroegop, A.; Ruizendaal, J. Design and evaluation of a robotic prototype for gerbera harvesting, performing actions at never-seen locations. *Comput. Electron. Agric.* **2025**, *228*, 109671. [CrossRef]
14. Xie, H.; Zhang, D.; Yang, L.; Cui, T.; He, X.; Zhang, K.; Zhang, Z. Development, Integration, and Field Evaluation of a Dual-Arm Ridge Cultivation Strawberry Autonomous Harvesting Robot. *J. Field Robot.* **2025**, *42*, 1783–1798. [CrossRef]
15. Choi, D.W.; Park, J.H.; Yoo, J.H.; Ko, K.E. AI-driven adaptive grasping and precise detaching robot for efficient citrus harvesting. *Comput. Electron. Agric.* **2025**, *232*, 110131. [CrossRef]
16. Zhang, B.; Xie, Y.; Zhou, J.; Wang, K.; Zhang, Z. State-of-the-art robotic grippers, grasping and control strategies, as well as their applications in agricultural robots: A review. *Comput. Electron. Agric.* **2020**, *177*, 105694. [CrossRef]
17. Liu, W.; Tan, Z.; Xie, W. A hand-like gripper embedded with flexible gel sensor for tomato harvesting: Soft contact and intelligent ripeness sensing. *J. Food Meas. Charact.* **2025**, *19*, 3150–3161. [CrossRef]
18. Qiu, Y.; Ye, Z.; Tan, X.; Dai, M.; Ge, S.; Zhao, X.; Kong, D.; Ruan, Y. Fruit grasping evaluation based on a humanoid underactuated manipulator and an adaptive grasping algorithm. *Smart Agric. Technol.* **2025**, *11*, 101007. [CrossRef]
19. Wang, Q.; Bai, K.; Zhang, L.; Sun, Z.; Jia, T.; Hu, D.; Li, Q.; Zhang, J.; Knoll, A.; Jiang, H.; et al. Towards reliable and damage-less robotic fragile fruit grasping: An enveloping gripper with multimodal strategy inspired by Asian elephant trunk. *Comput. Electron. Agric.* **2025**, *234*, 110198. [CrossRef]
20. Sun, J.; Sun, L.; Zhao, G.; Liu, J.; Chen, Z.; Jing, L.; Cao, X.; Zhang, H.; Tang, W.; Wang, J. Triboelectric force feedback-based fully actuated adaptive apple-picking gripper for optimized stability and non-destructive harvesting. *Comput. Electron. Agric.* **2025**, *237*, 110725. [CrossRef]
21. Ansari, S.; Gohil, M.K.; Maeda, Y.; Bhattacharya, B. Design and development of a new hybrid auxetic gripper system for adaptive precision tomato harvesting. *Int. J. Intell. Robot. Appl.* **2025**, 1–24. [CrossRef]
22. Zhu, H.; Qin, H.; Qiu, Z.; Chen, X.; Xue, J.; Gu, X.; Lu, M. A Compliant Active Roller Gripper with High Positional Offset Tolerance for Delicate Spherical Fruit Handling. *Agriculture* **2025**, *15*, 220. [CrossRef]
23. Ward-Cherrier, B.; Rojas, N.; Lepora, N.F. Model-Free Precise in-Hand Manipulation with a 3D-Printed Tactile Gripper. *IEEE Robot. Autom. Lett.* **2017**, *2*, 2056–2063. [CrossRef]
24. He, L.; Lu, Q.; Abad, S.A.; Rojas, N.; Nanayakkara, T. Soft Fingertips with Tactile Sensing and Active Deformation for Robust Grasping of Delicate Objects. *IEEE Robot. Autom. Lett.* **2020**, *5*, 2714–2721. [CrossRef]
25. Wang, Z.; Chathuranga, D.S.; Hirai, S. 3D printed soft gripper for automatic lunch box packing. In Proceedings of the 2016 IEEE International Conference on Robotics and Biomimetics (ROBIO), Qingdao, China, 3–7 December 2016; IEEE: Piscataway, NJ, USA, 2016; pp. 503–508. [CrossRef]
26. Shan, Y.; Ding, B.; Zhong, J.; Li, Y. Design and optimization of a decoupled serial constant force microgripper for force sensitive objects manipulation. *Robotica* **2023**, *41*, 2064–2078. [CrossRef]
27. Wu, Z.; Wang, Y.; Chen, M.; Ding, B. Design, analysis, and experimental investigations of an asymmetrical under-actuated micro-gripper. *J. Intell. Mater. Syst. Struct.* **2024**, *35*, 960–970. [CrossRef]
28. Ding, B.; Li, X.; Li, C.; Li, Y.; Chen, S.C. A survey on the mechanical design for piezo-actuated compliant micro-positioning stages. *Rev. Sci. Instrum.* **2023**, *94*, 101502. [CrossRef]

29. Kargar, S.M.; Berselli, G. TriCoM Gripper—Part I: Mechanical Design and Synthesis of a 3-Finger Compliant Monolithic Gripper. In Proceedings of the MESA 2024—20th International Conference on Mechatronic, Embedded Systems and Applications (MESA), Genova, Italy, 2–4 September 2024; Institute of Electrical and Electronics Engineers Inc.: Piscataway, NJ, USA, 2024. [CrossRef]
30. Shintake, J.; Cacucciolo, V.; Floreano, D.; Shea, H. Soft Robotic Grippers. *Adv. Mater.* **2018**, *30*, 1707035. [CrossRef]
31. Huang, P.; Xin, Y.; Lee, P.S. Soft electroadhesion systems for soft robotics. *npj Robot.* **2025**, *3*, 29. [CrossRef]
32. Rajagopalan, P.; Muthu, M.; Liu, Y.; Luo, J.; Wang, X.; Wan, C. Advancement of Electroadhesion Technology for Intelligent and Self-Reliant Robotic Applications. *Adv. Intell. Syst.* **2022**, *4*, 2200064. [CrossRef]
33. Guo, J.; Leng, J.; Rossiter, J. Electroadhesion Technologies for Robotics: A Comprehensive Review. *IEEE Trans. Robot.* **2020**, *36*, 313–327. [CrossRef]
34. Choi, K.; Kim, Y.C.; Sun, H.; Kim, S.-H.; Yoo, J.W.; Park, I.-K.; Lee, P.-C.; Choi, H.J.; Choi, H.R.; Kim, T.; et al. Quantitative Electrode Design Modeling of an Electroadhesive Lifting Device Based on the Localized Charge Distribution and Interfacial Polarization of Different Objects. *ACS Omega* **2019**, *4*, 7994–8000. [CrossRef]
35. Digumarti, K.M.; Cacucciolo, V.; Shea, H. Dexterous textile manipulation using electroadhesive fingers. In Proceedings of the IEEE International Conference on Intelligent Robots and Systems, Prague, Czech Republic, 27 September–1 October 2021; Institute of Electrical and Electronics Engineers Inc.: Piscataway, NJ, USA, 2021; pp. 6104–6109. [CrossRef]
36. An, Y.; He, B.; Ma, Z.; Guo, Y.; Yang, G.Z. Microassembly: A Review on Fundamentals, Applications and Recent Developments. *Engineering* **2025**, *48*, 323–346. [CrossRef]
37. Monkman, G.J. An analysis of astrictive prehension. *Int. J. Robot. Res.* **1997**, *16*, 1–10. [CrossRef]
38. Johnsen, A.; Rahbek, K. A physical phenomenon and its applications to telegraphy, telephony, etc. *Eng. Comput. Sci. Phys.* **1923**, *61*, 713–725. [CrossRef]
39. Krape, R.P. *Applications Study of Electroadhesive Devices*; NASA Contractor Report CR-1211; Chrysler Corp. Space Division: New Orleans, LA, USA, 1968. Available online: <https://ntrs.nasa.gov/citations/19680027258> (accessed on 25 April 2025).
40. Hwang, G.; Park, J.; Cortes, D.S.D.; Hyeon, K.; Kyung, K.U. Electroadhesion-Based High-Payload Soft Gripper with Mechanically Strengthened Structure. *IEEE Trans. Ind. Electron.* **2022**, *69*, 642–651. [CrossRef]
41. Guo, J.; Bamber, T.; Zhao, Y.; Chamberlain, M.; Justham, L.; Jackson, M. Toward Adaptive and Intelligent Electroadhesives for Robotic Material Handling. *IEEE Robot. Autom. Lett.* **2017**, *2*, 538–545. [CrossRef]
42. Mikela, R.; David, B. Geometry characterization of electroadhesion samples for spacecraft docking application. In Proceedings of the 2017 IEEE Aerospace Conference: Yellowstone Conference Center, Big Sky, MT, USA, 4–11 March 2017; IEEE: Piscataway, NJ, USA, 2017. [CrossRef]
43. Narayanan, S.; Barnhart, D.; Rogers, R.; Ruffatto, D.; Schaler, E.; Van Crey, N.; Dean, G.; Bhanji, A.; Bernstein, S.; Singh, A.; et al. REACCH-Reactive Electro-Adhesive Capture Cloth Mechanism to Enable Safe Grapple of Cooperative/Non-Cooperative Space Debris. In Proceedings of the AIAA Scitech 2020 Forum, Orlando, FL, USA, 6–10 January 2020.
44. Zhang, Z.; Chestney, J.A.; Sarhadi, M. Characterizing an electrostatic gripping device for the automated handling of non-rigid materials. *Proc. Inst. Mech. Eng. Part B J. Eng. Manuf.* **2001**, *215*, 21–36. [CrossRef]
45. Qin, S.; McTeer, A. Wafer dependence of Johnsen-Rahbek type electrostatic chuck for semiconductor processes. *J. Appl. Phys.* **2007**, *102*, 064901. [CrossRef]
46. Ruffatto, D.; Shah, J.; Spenko, M. Optimization and Experimental Validation of Electrostatic Adhesive Geometry. In Proceedings of the 2013 IEEE Aerospace Conference, Big Sky, MT, USA, 2–9 March 2013; IEEE: Piscataway, NJ, USA, 2013. [CrossRef]
47. Guo, J.; Bamber, T.; Chamberlain, M.; Justham, L.; Jackson, M. Optimization and experimental verification of coplanar interdigital electroadhesives. *J. Phys. D Appl. Phys.* **2016**, *49*, 415304. [CrossRef]
48. Yuan, K.L.; Chen, X.; Peng, J. Optimization of Electrostatic Adhesion Using Concentric Ring Electrodes. In *Intelligent Robotics and Applications, Proceedings of the 17th International Conference, ICIRA 2024, Xi'an, China, 31 July–2 August 2024*; Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics); Springer Science and Business Media Deutschland GmbH: Berlin/Heidelberg, Germany, 2025; pp. 360–371. [CrossRef]
49. Guo, J.; Xiang, C.; Zanini, P.; Rossiter, J. Magnetic Augmented Self-sensing Flexible Electroadhesive Grippers. *IEEE Robot. Autom. Lett.* **2019**, *4*, 2364–2369. [CrossRef]
50. Mastrangelo, M.; Caruso, F.; Carbone, G.; Cacucciolo, V. Electroadhesion zipping with soft grippers on curved objects. *Extrem. Mech. Lett.* **2023**, *61*, 101999. [CrossRef]
51. Shintake, J.; Rosset, S.; Schubert, B.; Floreano, D.; Shea, H. Versatile Soft Grippers with Intrinsic Electroadhesion Based on Multifunctional Polymer Actuators. *Adv. Mater.* **2016**, *28*, 231–238. [CrossRef]
52. Alizadehyazdi, V.; Bonthron, M.; Spenko, M. An Electrostatic/Gecko-Inspired Adhesives Soft Robotic Gripper. *IEEE Robot. Autom. Lett.* **2020**, *5*, 4679–4686. [CrossRef]
53. Caruso, F.; Shea, H.; Cacucciolo, V. Electroadhesion Suction Cups. *Adv. Mater.* **2025**, *37*, 2420231. [CrossRef]
54. Cacucciolo, V.; Shea, H.; Carbone, G. Peeling in electroadhesion soft grippers. *Extrem. Mech. Lett.* **2022**, *50*, 101529. [CrossRef]

55. An, S.; Xiang, C.; Ji, C.; Liu, S.; He, L.; Li, L.; Wang, Y. Design and development of a variable structure gripper with electroadhesion. *Smart Mater. Struct.* **2024**, *33*, ad3bf8. [CrossRef]
56. Xiang, C.; Li, W.; Guan, Y. A Variable Stiffness Electroadhesive Gripper Based on Low Melting Point Alloys. *Polymers* **2022**, *14*, 4469. [CrossRef]
57. Kim, D.G.; Je, H.; Hart, A.J.; Kim, S. Additive manufacturing of flexible 3D surface electrodes for electrostatic adhesion control and smart robotic gripping. *Friction* **2023**, *11*, 1974–1986. [CrossRef]
58. Cacucciolo, V.; Shintake, J.; Shea, H. Delicate yet Strong: Characterizing the electro-adhesion lifting force with a soft gripper. In Proceedings of the 2019 2nd IEEE International Conference on Soft Robotics (RoboSoft), Seoul, Republic of Korea, 14–18 April 2019; IEEE: Piscataway, NJ, USA, 2019.
59. Xiong, Q.; Li, D.; Li, X.; Liu, L.; Zhou, X.; Ju, H.; Zhu, Y.; Yeow, R.C.-H. EA-SoGripper: Electroadhesion-Stiffening Self-Adaptive Soft Robotic Gripper. *IEEE/ASME Trans. Mechatron.* **2025**, 1–11. [CrossRef]
60. Kanno, R.; Nguyen, P.H.; Pinski, J.; Howard, D.; Song, S.; Kovac, M. Hybrid Soft Electrostatic Metamaterial Gripper for Multi-surface, Multi-object Adaptation. In Proceedings of the 2024 IEEE 7th International Conference on Soft Robotics, RoboSoft 2024, San Diego, CA, USA, 14–17 April 2024; Institute of Electrical and Electronics Engineers Inc.: Piscataway, NJ, USA, 2024; pp. 851–857. [CrossRef]
61. Carloni, A.; Valori, M.; Bertolucci, F.; Agostini, L.; Berselli, G.; Fassi, I.; Tosatti, L.M.; Vertechy, R. Enhancing compliant gripper performance: Exploiting electro-adhesion to increase lifting force over grasping force. *Robot. Comput. Integr. Manuf.* **2025**, *91*, 102843. [CrossRef]
62. Han, C.; Lv, J.; Dong, C.; Li, J.; Luo, Y.; Wu, W.; Abdeen, M.A. Classification, Advanced Technologies, and Typical Applications of End-Effector for Fruit and Vegetable Picking Robots. *Agriculture* **2024**, *14*, 1310. [CrossRef]
63. Rizzello, G.; Naso, D.; York, A.; Seelecke, S. A Self-Sensing Approach for Dielectric Elastomer Actuators Based on Online Estimation Algorithms. *IEEE/ASME Trans. Mechatron.* **2017**, *22*, 728–738. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Article

Design, Analysis, and Prototyping of a Formula SAE Application Flux-Switching Permanent Magnet Motor

Filippo Crescenzo and Nicola Bianchi *

Department of Industrial Engineering, University of Padova, Via Gradenigo 6A, 35131 Padova, Italy;
filippo.crescenzo@studenti.unipd.it

* Correspondence: nicola.bianchi@unipd.it

Abstract

This paper investigates the design, analysis, and prototyping of a Flux-Switching Permanent Magnet (FSPM) motor for Formula SAE electric vehicle applications. The stringent competition requirements demand traction motors with high torque and power density, and reliable operation at elevated speeds. An analytical model based on air-gap permeance and magnetomotive force distributions is developed to provide an effective preliminary design tool and to reduce computational effort. The proposed model is used to define the main geometrical parameters of a 12-slot, 10-rotor-tooth FSPM machine, which is subsequently validated through finite element analysis. Analytical and numerical results are compared in terms of air-gap flux density, flux linkage, and torque capability, showing good agreement. Manufacturing-driven design choices, including optimized magnet slot geometry, laminated permanent magnets for eddy-current loss mitigation, and a mechanically robust lightweight rotor, are introduced to ensure high-speed operability and assembly reliability.

Keywords: Flux Switching Permanent Magnet Motor; Formula SAE electric; high-speed electric machines; analytical modeling; finite element analysis; manufacturing-oriented design

1. Introduction

The Formula SAE electric competition imposes strict constraints on traction systems, including a maximum battery voltage of 600 V and a maximum discharge power of 80 kW. Within these limits, motors must provide high torque and power density, low mass, and reliable operation at rotational speeds of up to 20 krpm.

High-speed permanent magnet synchronous machines (PMSMs) are commonly adopted. However, rotor-mounted permanent magnet topologies may face mechanical and thermal challenges due to centrifugal stresses and limited magnet cooling, which are particularly relevant in lightweight motors for motorsport applications. The presented motor design procedure addresses the following:

- Electromagnetic performance (torque, flux linkage, power capability).
- High-speed mechanical integrity.
- Mass reduction.
- Manufacturing feasibility.

Loss mechanisms and thermal behavior must also be considered to ensure full compliance and reliable operation under sustained load. However, detailed modeling of these aspects is beyond the scope of this work and will be addressed in future developments.

1.1. Objective of This Work

This paper investigates the suitability of Flux-Switching Permanent Magnet (FSPM) machines [1] for Formula SAE traction applications. Interest in these machines in this kind of application was taken due to their attractive torque density capability and a rotor structure that makes them compatible with high-speed operations.

More specifically, the objective is to design, analyze, and prototype an FSPM motor specifically tailored for Formula SAE electric vehicles, such as the University of Padova Formula SAE electric race car shown in Figure 1, using a linear analytical model based on air-gap permeance and magnetomotive force as a preliminary tool for the definition of the main geometric parameters of the motor.



Figure 1. University of Padova Formula SAE electric vehicle.

1.2. Contributions of This Paper

The main contributions can be summarized as follows:

- Demonstration of the analytical model as a reliable tool for preliminary FSPM motor sizing, with relative errors below 6% compared to linear FEA for air-gap flux density, no-load flux linkage, and electromagnetic torque.
- Electromagnetic and mechanical performance improvements, including laminated permanent magnets for eddy-current loss reduction and rotor mass reduction, validated through structural stress analysis at 20 krpm.
- Description of the adopted manufacturing procedures and practical considerations encountered during prototyping, leading to the realization of a 3.32 kg active-mass prototype.

2. Flux-Switching Permanent Magnet Motor

2.1. General Structure

In the most common PM synchronous machine configurations, permanent magnets are located on the rotor, exposing them to the risk of permanent demagnetization due to insufficient heat dissipation [2]. Moreover, high-speed applications often require the use of restraining sleeves to prevent the magnets from being torn away from the rotor due to centrifugal forces [3]. Both of these problems can be avoided by placing the magnets in the stator: magnets can be directly cooled and no restraining sleeves are required. Compared to other stator-PM machines, such as DSPM and FRPM machines, FSPM machines present higher torque and power densities and a more sinusoidal back-emf [4]. Figure 2 illustrates the basic structure of this kind of machine, which consists of [5]:

- U-shaped stator laminations, with non-overlapping concentrated windings.
- Circumferentially magnetized PMs with alternate polarity, sandwiched between stator laminations.
- Rotor made of iron with a mechanically robust structure, similar to that of switched reluctance motors.




Figure 2. Cross-section of a 12-slot, 10-rotor-tooth FSPM machine. The red arrows indicate the magnetization direction of each magnet.

2.2. Working Principle

Both the stator and the rotor present a salient structure; bipolar flux linkage variations are produced through rotor rotation. A complete electrical period corresponds to a mechanical rotation equal to the angle between two adjacent rotor teeth; consequently, the number of pole pairs is equal to the number of rotor teeth. Concerning the torque production mechanism, it has been extensively studied in [6,7], and also an alternative perspective is given in [8]. The doubly salient structure of this motor results in a distorted air-gap flux density, compared with IPM and SPM machines; this is caused by the modulation of the magnetic field produced both by PMs and the excitation currents in stator windings. As a result, 95% of the average torque is produced by several synchronized air-gap harmonics produced by PMs and excitation currents [4].

3. Analytical Model

In order to obtain a preliminary machine geometry definition, an analytical model has been employed, avoiding the long simulation times associated with FEM while still providing acceptable accuracy. The motor design obtained analytically has been validated through FEM and subsequently improved.

3.1. Limitations

It is important to highlight that the approach used does not include the nonlinear behavior of the iron as well as the tangential components of the air-gap flux density, making it suitable for early-stage design and preliminary dimensioning. The effect of saturation and of the tangential components is then evaluated through finite element analysis.

3.2. Requirements

In the design procedure, the specifications of the motors currently in use are taken as a reference, while assuming the same battery and inverter configuration as specified in Table 1.

Regarding the FSPM motor slot/pole configuration, in order to limit the excitation fundamental frequency at a given speed, configurations with a smaller number of slots and poles are preferable for high-speed applications. Only a limited number of slot/pole combinations are feasible. Although the minimum number of stator slots is six, as shown in [9], this configuration is not suitable for high-speed applications due to the resulting

unbalanced back-EMF waveform, which leads to increased harmonic losses, as well as to the presence of unbalanced radial forces acting on the rotor. The next feasible slot number is twelve, with a minimum of ten rotor teeth. This configuration is the most commonly adopted for this type of machine. Although it is not the optimal one in terms of torque density—which can be improved by using alternative stator teeth and magnet configurations [10,11]—the 12/10 two-tooth layout is selected in this work for the sake of simplicity. Additionally, alternative permanent magnet configurations for outer-rotor in-wheel applications have been investigated in [12].

Table 1. Current Powertrain specifications.

Parameter	Value	Unit
Machine type	3-phase/IPM	-
Max. battery voltage	565	V _{dc}
Rated voltage	350	V _{rms}
Rated current	41	A _{rms}
Peak current	105	A _{rms}
Rated power	12.3	kW
Peak power (1.24 s)	35	kW
Rated torque	9.8	N m
Peak torque (1.24 s)	21	N m
Rated speed	12,000	rpm
Max. speed	20,000	rpm

3.3. Model Description

Different approaches have been developed for the analytical modeling of FSPM machines, such as the nonlinear magnetic equivalent circuit model [13] or through the solution of field equations at the air gap by means of Fourier analysis [14]. The method used in this work, similarly to the one presented by B. Gaussens in [15], is based on air-gap permeance and MMF functions. Following the permeance-based analytical formulation adopted in [16], the total magnetic permeance distribution between stator and rotor can be expressed as

$$\Lambda(\vartheta_m^s, \vartheta_m^r, \Delta\vartheta_{r,s}) = \frac{\mu_0}{g_0 + g_s(\vartheta_m^s) + g_r(\vartheta_m^r, \Delta\vartheta_{r,s})} \quad (1)$$

where g_0 is the physical air gap between the stator and the rotor; $g_s(\vartheta_m^s)$ is the stator equivalent air-gap function, which can be decomposed in the winding slot and magnet slot contributions, $g_{s,w}(\vartheta_m^s)$ and $g_{s,m}(\vartheta_m^s)$, respectively; $g_r(\vartheta_m^r)$ is the rotor equivalent air-gap function; ϑ_m^s is the angular position in the stator reference frame; ϑ_m^r is the angular position in the rotor reference frame; and $\Delta\vartheta_{r,s}$ is the angular shift between the rotor reference frame and the stator reference frame. The air-gap flux density distribution can be expressed as

$$B(\vartheta_m^s, \vartheta_m^r, \Delta\vartheta_{r,s}) = \Lambda(\vartheta_m^s, \vartheta_m^r, \Delta\vartheta_{r,s})[MMF_m(\vartheta_m^s) + MMF_w(\vartheta_m^s)] \quad (2)$$

where $MMF_m(\vartheta_m^s)$ and $MMF_w(\vartheta_m^s)$ are the functions which define the magnetomotive force produced by the magnets and the winding currents

3.3.1. Stator Equivalent Air-Gap Function

The stator equivalent air-gap function is computed by considering only the effect of winding and magnet slots, thus assuming a smooth rotor. Firstly, the equivalent air-gap function $g_{s,w}(\vartheta_m^s)$, associated with winding slots, is computed, by assuming indefinitely deep slots and a simplified circular path for the leakage flux lines as in [15]. The equivalent path of the flux lines is obtained by means of two parallel paths, $\Gamma_1(\vartheta_m^s)$ and $\Gamma_2(\vartheta_m^s)$,

as shown in Figure 3, where the radius of these circular paths depends on the angular coordinate ϑ_m^s .

$$R_1(\vartheta_m^s) = \frac{D_{i,s}}{2} \vartheta_m^s \quad R_2(\vartheta_m^s) = \frac{D_{i,s}}{2} (\Delta\vartheta_{so} - \vartheta_m^s) \quad (3)$$

$$\Gamma_1(\vartheta_m^s) = \frac{\pi}{2} R_1(\vartheta_m^s) \quad \Gamma_2(\vartheta_m^s) = \frac{\pi}{2} R_2(\vartheta_m^s) \quad (4)$$

$$g_{s,w}(\vartheta_m^s) = \frac{\Gamma_1(\vartheta_m^s) \Gamma_2(\vartheta_m^s)}{\Gamma_1(\vartheta_m^s) + \Gamma_2(\vartheta_m^s)} \quad (5)$$

where $D_{i,s}$ is the stator inner diameter. Then, substituting (1) and (2) in (3), the expression becomes

$$g_{s,w}(\vartheta_m^s) = \frac{\pi}{4} D_{i,s} \left(-\frac{1}{\Delta\vartheta_{so}} \vartheta_m^{s^2} + \vartheta_m^s \right) \quad (6)$$




Figure 3. Flux paths within the stator winding slot.

Figure 4 provides a graphical representation of the equivalent air-gap function due to the winding slots, which can be expressed by means of the Fourier series expansion as

$$g_{s,w}(\vartheta_m^s) = a_0^{(s,w)} + \sum_{k=1}^{+\infty} a_k^{(s,w)} \cos(z\vartheta_m^s) + b_k^{(s,w)} \sin(z\vartheta_m^s) \quad (7)$$

where

$$z = k \frac{2\pi}{\tau_s} \quad (8)$$

$$a_0^{(s,w)} = \frac{\Delta\vartheta_{so}^2 \pi D_{i,s}}{24\tau_s} \quad (9)$$

$$a_k^{(s,w)} = \frac{\pi D_{i,s}}{2\tau_s} \frac{1}{z^2} \left[\frac{2}{\Delta\vartheta_{so}} \frac{\sin(z\Delta\vartheta_{so})}{z} - \cos(z\Delta\vartheta_{so}) - 1 \right] \quad (10)$$

$$b_k^{(s,w)} = \frac{\pi D_{i,s}}{2\tau_s} \frac{1}{z^2} \left[\frac{2}{\Delta\vartheta_{so}} \frac{1}{z} - \frac{2}{\Delta\vartheta_{so}} \frac{\cos(z\Delta\vartheta_{so})}{z} - \sin(z\Delta\vartheta_{so}) \right] \quad (11)$$

The equivalent air-gap function associated with magnet slots $g_{s,m}(\vartheta_m)$ is obtained with the same procedure, by substituting the angle corresponding to the winding slot opening, $\Delta\vartheta_{so}$, with the angle corresponding to the magnet slot opening, $\Delta\vartheta_{tm}$:

$$g_{s,m}(\vartheta_m^s) = a_0^{(s,m)} + \sum_{z=1}^{+\infty} a_k^{(s,m)} \cos(z\vartheta_m^s) + b_k^{(s,m)} \sin(z\vartheta_m^s) \quad (12)$$

where

$$z = k \frac{2\pi}{\tau_s} \quad (13)$$

$$a_0^{(s,m)} = \frac{\Delta\vartheta_{tm}^2 \pi D_{i,s}}{24\tau_s} \quad (14)$$

$$a_k^{(s,m)} = \frac{\pi D_{i,s}}{2\tau_s} \frac{1}{z^2} \left[\frac{2}{\Delta\vartheta_{tm}} \frac{\sin(z\Delta\vartheta_{tm})}{z} - \cos(z\Delta\vartheta_{tm}) - 1 \right] \quad (15)$$

$$b_k^{(s,m)} = \frac{\pi D_{i,s}}{2\tau_s} \frac{1}{z^2} \left[\frac{2}{\Delta\vartheta_{tm}} \frac{1}{z} - \frac{2}{\Delta\vartheta_{tm}} \frac{\cos(z\Delta\vartheta_{tm})}{z} - \sin(z\Delta\vartheta_{tm}) \right] \quad (16)$$

A phase shift is added to the two functions so that the magnet slots are centered on the stator teeth, and the origin of the system of reference is centered with stator teeth. Finally, the stator equivalent air-gap function is obtained as

$$g_s(\vartheta_m) = g_{s,w}(\vartheta_m^s) + g_{s,m}(\vartheta_m^s) \quad (17)$$




Figure 4. Stator equivalent air-gap function associated with winding slots.

3.3.2. Rotor Permanence Distributions

The rotor equivalent air-gap function is obtained by assuming a slotless stator, no magnet or winding slots; in this case, rotor slots correspond to the spacing between two teeth. The procedure is the same as the one for the equivalent air-gap function associated with winding slots, so it only requires the following changes:

- $\tau_s \rightarrow \tau_r$
- $\Delta\vartheta_{so} \rightarrow \Delta\vartheta_{ro}$

Then, a phase shift is introduced in order to align rotor teeth with the origin of the rotor reference frame. Finally, the additional phase shift $\Delta\vartheta_{r,s}$ can be used to rotate the rotor in the stator reference frame. A visualization of the flux paths within the rotor teeth is provided by Figure 5, meanwhile Figure 6 shows the resulting rotor equivalent air-gap function.




Figure 5. Flux paths within the rotor teeth.




Figure 6. Rotor equivalent air-gap function.

Finally, the stator and rotor air-gap distributions when a rotor tooth is aligned with the reference stator tooth are shown in Figure 7.




Figure 7. Equivalent air-gap functions: stator contribution (**top**) and rotor contribution (**bottom**).

3.3.3. Magnets MMF

In order to derive the Magnetomotive Force function, associated with the excitation provided by the magnets, the no-load maximum value of the air-gap flux density produced by the PMs is computed by applying Ampere's law and Gauss's law, as well as the magnet characteristic equation considering the flux paths of Figure 8.

$$\begin{cases} H_m t_m + 2H_g g_0 = 0 \\ B_m S_m = B_g \frac{S_g}{2} \\ B_g = \mu_0 H_g \\ B_m = B_{rem} + \mu_0 \mu_{rec} H_m \end{cases} \quad (18)$$

where H_m and B_m are the values of the magnetic field and flux density within the magnet, H_g and B_g are the values of the magnetic field and flux density in the air gap, B_{rem} is the residual flux density, μ_0 and μ_{rec} are the magnetic permeability of the void and the relative

permeability of the air, t_m is the thickness of the magnet, S_g is the pole air-gap surface and S_m is the magnet surface orthogonal to the flux lines:

$$S_g = \frac{\pi D_{is}}{12} L_{stk} \quad S_m = h_m L_{stk} \quad (19)$$

By solving the system of equations, the magnet flux density is obtained as

$$B_m = \frac{B_{rem}}{1 + \frac{4g_0}{t_m} \frac{S_m}{S_g} \mu_{rec}} \quad (20)$$

Then, the air-gap flux density is

$$B_g = B_m \frac{2S_m}{S_g} \quad (21)$$

Finally, the maximum value of the magnets MMF is obtained as

$$MMF_{m,max} = \frac{B_g}{\mu_0} g_0 \quad (22)$$




Figure 8. Flux paths associated with the permanent magnets.

By assuming a linear variation in the MMF within the PMs, the MMF function is represented by a trapezoidal wave with a period of τ_s , as shown in Figure 9, which can be expressed by means of the Fourier series expansion as

$$MMF_m(\vartheta_m^s) = MMF_{m,max} \left[\sum_{k=1}^{+\infty} a_k \cos(z \vartheta_m^s) \right] \quad (23)$$

where

$$z = k \frac{\pi}{\tau_s} \quad (24)$$

$$a_k^{(m)} = \frac{8}{\Delta \vartheta_{tm} \tau_s z^2} \left[\cos\left(z \frac{\tau_s - \Delta \vartheta_{tm}}{2}\right) - \cos\left(z \frac{\tau_s}{2}\right) \right] \quad (25)$$




Figure 9. Equivalent magnetomotive force distribution produced by the permanent magnets.

3.3.4. Windings MMF

Similarly to what has just been described, the magnetomotive force produced by each phase winding is represented by trapezoidal half-waves, see Figure 10, where the amplitude at a given time is

$$MMF_{x,max}(t) = \frac{n_{cs} i_x(t)}{2} \quad x = a, b, c \quad (26)$$

where $i_x(t)$ is the amplitude of the phase x current at a given time, while n_{cs} is the number of series-connected conductors per slot. The MMF distribution of each phase is described by the following Fourier series expansion:

$$MMF_x(\theta_m^s, t) = MMF_{x,max}(t) \left[a_0^{(x)} + \sum_{k=1}^{+\infty} a_k^{(x)} \cos(z\theta_m^s + \phi_q) \right] \quad (27)$$




Figure 10. Equivalent magnetomotive force distribution produced by the stator windings.

$$z = k \frac{2\pi}{3\tau_s} \quad (28)$$

In which ϕ_q is the phase shift of the j -th phase:

$$\phi_q = q \frac{2\pi}{3} \quad q = 0, 1, 2 \quad (29)$$

And the series coefficients are

$$a_0^{(x)} = \frac{\gamma_s + \tau_s}{6\tau_s} \quad (30)$$

$$a_k^{(x)} = \frac{4}{2\tau_s} \left[\frac{\sin\left(z\frac{\gamma_s}{2}\right)}{2} \left(1 + \frac{\gamma_s - \tau_s}{\theta_{s0}}\right) - \frac{2\cos\left(z\frac{\tau_s}{2}\right)}{\theta_{s0}z^2} + \frac{2\cos\left(z\frac{\gamma_s}{2}\right)}{\theta_{s0}z^2} \right] \quad (31)$$

A graphical representation of the magnetomotive force distribution produced by the stator windings is given by Figure 11.




Figure 11. MMF distributions of the three phases: phase a ($i_a = \sqrt{2}I$) in red, phase b ($i_b = -\frac{\sqrt{2}I}{2}$) in green, phase c ($i_c = -\frac{\sqrt{2}I}{2}$) in black.

4. Analytical and FEM Results Comparison

4.1. Geometry

The stator and rotor geometries obtained iteratively through the analytical model are reported in Table 2. These parameters were subsequently implemented in a FEM aimed at validating the results of the analytical simulations. To ensure a fair comparison with the analytical model, the nonlinear magnetic behavior of the iron is not included in this validation stage. Therefore, the comparison presented in Section 4 is performed between the analytical model and a linear FEM.

Table 2. FSPM motor geometry.

Parameter	Symbol	Value	Unit
Outer diameter	$D_{e,s}$	123	mm
Inner diameter	$D_{i,s}$	70	mm
Stack length	L_{stk}	48	mm
Stator teeth width	w_t	5.6	mm
Magnet grade	-	N45UH	-
Magnet thickness	t_m	2.5	mm
Magnet height	h_m	26.4	mm
Air-gap length	g	0.5	mm
Rotor teeth angle	$\Delta\theta_{rt}$	15	deg
Rotor inner diameter	$D_{i,r}$	53	mm

4.2. Comparison

No-load flux and torque are evaluated over a rotor rotation corresponding to 360 electrical degrees (θ_e^s). As shown in Figure 12, the agreement between analytical and FEM

results is generally good, although some differences are observed. These deviations are primarily due to the analytical model computing the air-gap flux density at the sampled points, based on the local values of permeance and MMF functions. As a result, only the radial component of the flux density is captured. However, as shown in the following figures, the analytical approach still allows the prediction of the machine performance with acceptable precision.




Figure 12. Air-gap flux density: analytically computed (blue) and FEM-computed (dashed red).

The no-load flux linkage obtained with the two models is compared in Figure 13. The maximum value analytically computed is $\lambda_m = 20.21$ mWb, while the value obtained through the FEM is $\lambda_m = 21.41$ mWb, resulting in a relative error of 5.88%.




Figure 13. No-load flux linkage: analytically computed (blue) and FEM-computed (dashed red).

The torque at nominal operating conditions has also been compared, considering a stator current of $I_s = 41$ A_{rms}. In order to have a fair comparison between the analytical model and FEM, the torque computed using the d - q flux components is considered:

$$T_{dq} = \frac{3}{2}p(\lambda_d I_q - \lambda_q I_d) \quad (32)$$

This formulation represents the average electromagnetic torque of the machine. It does not account for the oscillating (torque ripple) component because the analytical model computes the flux linkage as the sum of the flux contributions in each segment into which the air gap has been discretized. With this approach, only the average component of the air-gap flux density contributes to the torque calculation, while the spatial harmonics responsible for torque ripple are inherently neglected. Moreover, although in FSPM machines a reluctance torque component can appear under high magnetic saturation, the comparison shown in Figure 14 is carried out at the nominal current condition. Under these operating conditions, the d -axis and q -axis inductances are expected to assume similar values, resulting in a negligible reluctance torque contribution compared to the permanent magnet contribution. This assumption is further confirmed by the nonlinear FEM analysis reported in Appendix A.2. Therefore, the torque expression is simplified as

$$T_{dq} = \frac{3}{2} p \lambda_m I_q \quad (33)$$

The average torque computed through the FEM is 18.58 N m, while the analytical model predicts an average value of 17.39 N m. This leads to a relative error of 5.61%, which confirms the good agreement between the two approaches. To further assess the predictive capability of the analytical model over the whole operating range, an evaluation of the torque error over the dq-plane has been carried out. The resulting error distribution is reported in Appendix A.1.




Figure 14. Torque under constant load operation: analytically computed (blue); FEM-computed (dashed red).

5. Design Improvement and Validation

5.1. Magnet Slot Design

In the design refinement and performance evaluation stage, a nonlinear FEM is adopted in order to account for magnetic saturation effects. This nonlinear model is used in all the analyses presented in this section. A proper slot–magnet design is essential to ensure the manufacturability of the prototype. During the curing process, the laminations must remain precisely aligned so that the magnets can be correctly inserted into the slots afterward. In addition, the laminations must provide sufficient structural integrity to allow safe handling during assembly, without the risk of deformation.

To guarantee adequate mechanical robustness, a 1 mm magnetic bridge is introduced on top of each magnet slot. Although this solution leads to a limited amount of flux leakage, the resulting increase in back-iron thickness reduces magnetic saturation in the machine. As a consequence, the diameter is increased to 125 mm and the overall torque output remains essentially unchanged, see Figure 15.




Figure 15. Magnet slot before (right) and after (left) manufacturing-oriented design improvement.

Furthermore, to ensure proper alignment of the laminations during assembly, four notches with a radius of 1 mm are introduced along the outer diameter of the machine.

5.2. Magnet Design: Eddy-Current Mitigation and Demagnetization Analysis

5.2.1. Eddy-Current Mitigation

Unlike conventional machine configurations, in which the permanent magnets are located on the rotor and experience flux variations mainly due to the slotting effect, stator-mounted magnet topologies exhibit significantly larger flux fluctuations. In Flux-Switching Permanent Magnet (FSPM) motors, the magnets are fully subjected to the time-varying magnetic field produced by the stator currents flowing in the concentrated windings around the teeth. These pronounced flux variations induce substantial eddy currents within the magnets, see Figure 16, leading to increased magnet losses, excessive heating, and a consequent degradation of machine performance. For this reason, the use of specially laminated or segmented permanent magnets is of considerable importance in FSPM machines, as it effectively limits eddy-current circulation and significantly reduces associated losses.




Figure 16. Schematic comparison between non-laminated and laminated magnets: eddy-current mitigation.

Consequently, a 10-layer laminated magnet is adopted for this machine. As shown in Figures 17 and 18, this solution results in a dramatic reduction in magnet losses.




Figure 17. Direct comparison of permanent magnet losses as a function of rotational speed between non-laminated magnets (blue) and laminated magnets (red).




Figure 18. Focus on laminated permanent magnet losses as a function of rotational speed.

5.2.2. Demagnetization Analysis

The resistance of the magnets to demagnetization was evaluated by applying increasing direct-axis (d-axis) demagnetizing currents. The applied current ranged from 0 up to twice the peak current delivered by the inverter.

As shown in Figure 19, the average flux density in the magnets remains positive throughout the entire current range. Considering that the B–H characteristic of N45UH magnets exhibits a knee point—beyond which irreversible demagnetization occurs at 120 °C for negative flux densities—the magnets are effectively protected under the tested conditions.




Figure 19. Average flux density in the magnets as a function of applied direct-axis current.

5.3. Rotor Geometry Improvement

The rotor geometry was optimized by selectively removing iron from magnetically inactive regions, as shown in Figure 20. This design refinement allowed a reduction in the rotor mass by approximately 270 g.




Figure 20. Rotor geometry before iron mass reduction (left); rotor geometry after iron mass reduction (right).

However, the resulting lightweight structure introduces areas where the cross-sectional thickness (in the plane) is significantly reduced. Considering the high rotational speed of the machine, a mechanical stress analysis becomes necessary to verify the structural integrity of the rotor and to ensure safe operation under centrifugal loading. The results shown in Figure 21 indicate that the maximum stress at 20 krpm remains well below the material yield strength reported in the datasheet (260 MPa).

An additional performance improvement is achieved through the adoption of a six-step skewing of the rotor laminations.




Figure 21. Stress distribution within the rotor.

Once the electromagnetic and mechanical design of the machine is finalized, the total mass of the active parts is evaluated. Table 3 reports the mass breakdown of the stator iron, rotor iron, copper windings, and permanent magnets, together with the resulting total active mass of the machine. Although the stator diameter has been increased to enhance manufacturability, the overall weight of the machine's active components remains well within the typical range for machines employed in this type of application.

Table 3. FSPM motor active mass contributions.

Parameter	Value	Unit
Stator Iron Mass	1.79	kg
Rotor Iron Mass	0.722	kg
Copper Mass	0.521	kg
Magnet Mass	0.289	kg
Total Mass	3.32	kg

5.4. Performance Evaluation

Figure 22 reports the FEM-derived torque and power versus speed characteristics under both overload and nominal operating conditions (dashed).

For the overload condition,

- Continuous torque: $T = 26.4 \text{ N m}$.
- Base speed: $n_{\text{base}} = 10.4 \text{ krpm}$.
- Maximum power: $P_{\text{max}} = 31.2 \text{ kW}$.

For nominal load operation,

- Continuous torque: $T = 13.6 \text{ N m}$.
- Base speed: $n_{\text{base}} = 16 \text{ krpm}$.
- Maximum power: $P_{\text{max}} = 26.2 \text{ kW}$.

The machine demonstrates strong torque and power performance, with overload operation enabling higher torque at reduced speed and peak power suitable for the dynamic demands of Formula Student racing.




Figure 22. Torque (black) and power (red) versus speed: continuous lines indicate overload; dashed lines indicate nominal operation.

6. Manufacturing Aspects

6.1. Iron Laminations

The laminations used in the prototype are coated with a backlack insulation layer, which enables the stack to be bonded through thermal curing without the need for welding

or soldering. This process ensures that the laminations remain fully electrically insulated from each other. However, the adoption of this solution requires special care during manufacturing, particularly to ensure proper alignment of the iron laminations throughout the thermal curing process. For the rotor, this issue is solved by curing the laminations while they are already stacked onto the shaft. In the stator, lamination alignment is finally achieved by inserting 12 mm pins into the slots instead of using notches on the outer diameter, as this provides greater structural solidity. High-temperature-resistant insulating spacers, similar to those commonly used in transformer winding manufacturing, are also inserted to preserve the magnet slot geometry during the curing process.

The stator and rotor stacks were subsequently compressed, resulting in the core packs shown in Figure 23. A torque wrench was used to ensure the uniform application of the pressure required for the proper activation of the backlack coating. This step was essential to guarantee intimate contact between adjacent laminations during the bonding process and to promote consistent adhesion across the entire stack height.




Figure 23. Lamination- pressed packs before curing process.

The curing process was performed in a controlled oven environment, starting from an ambient temperature of 30 °C. The thermal cycle was defined to ensure complete polymerization of the coating while avoiding thermal stresses that could compromise stack alignment. The procedure was carried out according to the following sequence:

- Temperature ramp: 19 min at a controlled heating rate of $10 \frac{^{\circ}\text{C}}{\text{min}}$.
- Dwelling phase: 30 min for the rotor stack and 40 min for the stator stack, in order to ensure full curing of the adhesive layer.
- In-oven cooling: 30 min for the rotor and 50 min for the stator stack, allowing gradual temperature reduction under controlled conditions.
- Open-air cooling until ambient temperature was reached.

The results of the curing process is shown in Figure 24

After the curing process, the permanent magnets are inserted into the stator slots and fixed using structural adhesive, see Figure 25.



Figure 24. Curing process results.




Figure 25. Magnet insertion into the stator core.

6.2. Stator Windings

Since the manufactured prototype is intended for laboratory investigations and experimental analysis, not all tooth coils belonging to each phase have been permanently connected by soldering. Some coil connections have intentionally been left open, as shown in Figure 26, in order to allow the implementation of a six-phase asymmetrical winding configuration, enabling further studies on multiphase applications.

During the winding manufacturing process, the stator stack was fixed on a dedicated support. Each coil was then inserted individually; to simplify the procedure, the coils were wound in an alternating sequence. Once the insertion process was completed, see Figure 27, the fixed connections between coils were soldered and insulated. Meanwhile, the phase terminals were connected to external cables with an equivalent cross-sectional area.




Figure 26. Stator winding connections.



Figure 27. Stator winding process results.

6.3. Design Improvement Summary

Table 4 summarizes the design changes proposed based on the considerations in Section 5, along with their corresponding effects on the machine performance and manufacturability.

Table 4. Design features and corresponding effects.

Feature	Effect
Magnet slot design	Optimized magnet slots and reinforced stator core, facilitating magnet insertion and winding manufacturing.
Laminated permanent magnets	Magnet losses heavily reduced (e.g., from 544 W to 10.7 W at 10 krpm).
Rotor weight reduction	≈27% rotor mass reduction (270 g) while ensuring mechanical robustness at 20 krpm

7. Conclusions and Future Developments

This paper has presented the design, analysis, and prototyping of a Flux-Switching Permanent Magnet (FSPM) motor specifically developed for Formula SAE electric vehicle applications. An analytical modeling approach based on air-gap permeance and magnetomotive force distributions was employed as a preliminary design tool, enabling a rapid definition of the main geometrical parameters while significantly reducing computational effort compared to full finite element simulations.

The analytical results were validated through finite element analysis, showing good agreement in terms of air-gap flux density, no-load flux linkage, and electromagnetic torque, with relative errors consistently below 6%. This confirms the suitability of the proposed analytical model for early-stage FSPM motor design. Based on the validated model, several design refinements were introduced to improve manufacturability, mechanical robustness, and electromagnetic performance, including optimized magnet slot geometry, laminated permanent magnets for eddy-current loss mitigation, and a lightweight rotor design verified through mechanical stress analysis.

The machine achieves torque and power levels that are compatible with the electromagnetic performance requirements typically encountered in Formula SAE, while maintaining a compact structure and a low active mass. However, complete validation would require detailed loss and thermal analyses.

Future work will focus on extensive experimental testing of the manufactured prototype to validate the predicted performance under real operating conditions. In addition, the intentionally flexible stator winding connections will enable further investigations into six-phase asymmetrical winding configurations, fault-tolerant control strategies, and advanced multiphase drive solutions, with the aim of further enhancing reliability and performance in high-demand motorsport and automotive applications.

Author Contributions: Conceptualization: F.C.; investigation: F.C.; methodology: F.C.; writing—original draft: F.C.; writing—review and editing: N.B., F.C.; supervision: N.B. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A

Appendix A.1

To evaluate the analytical model accuracy over the operating region, the electromagnetic torque error has been computed across the dq current plane. The comparison is performed between the analytical model and the linear FEM.

The resulting error distribution is shown in Figure A1. The Maximum Torque per Ampere (MTPA) and Maximum Torque per Voltage (MTPV) trajectories are also indicated to highlight the machine operating region.

The analytical model maintains good agreement with the FEM results over the operating region of interest, with the largest discrepancies occurring in the deep flux-weakening region ($\approx 11.3\%$).

Finally, it should be noted that a small difference with respect to the error specified in Section 4.2 is present, since in the mapping procedure for this diagram, the rotor is kept still while the stator currents are changed; meanwhile, the error reported in Figure 14 refers to the average torque values from a full rotation.




Figure A1. Analytical model error compared to linear FEM.

Appendix A.2

To clarify the contribution of reluctance torque, the direct-axis and quadrature-axis inductances have been evaluated over the dq current plane using the nonlinear FEM.

Figures A2 and A3 show the distributions of L_d and L_q , respectively, while Figure A4 reports the percentage difference between the two inductances.




Figure A2. Nonlinear FEM L_d map over the dq current plane.




Figure A3. Nonlinear FEM L_q map over the dq current plane.




Figure A4. Percentage difference between L_d and L_q over the dq current plane obtained from nonlinear FEM simulations.

The inductance behavior over the dq domain results in non-negligible reluctance torque values for high-load working points; see Figure A5. However under nominal current condition ($|I| = \sqrt{2} \cdot 41$ A), the reluctance torque component is inherently very small, and the electromagnetic torque is therefore mainly determined by the permanent magnet contribution.




Figure A5. Reluctance torque mapping over the dq -plane.

References

1. Wang, P.; Hua, W.; Zhang, G.; Wang, B.; Cheng, M. Principle of Flux-Switching Permanent Magnet Machine by Magnetic Field Modulation Theory—Part I: Back-Electromotive-Force Generation. *IEEE Trans. Ind. Electron.* **2022**, *69*, 2370–2379. [CrossRef]
2. Liu, C.; Xu, D.; Wu, W.; Yang, B. The Study of the Salient Pole Geometry Optimization of the Flux Switching Permanent Magnet Machine. *Actuators* **2024**, *13*, 398. [CrossRef]
3. Sarshar, M.R.; Jalali Kondelaji, M.A.; Asef, P.; Mirsalim, M. Electromagnetic Investigation of Innovative Stator–Permanent Magnet Motors. *Energies* **2025**, *18*, 2400. [CrossRef]
4. Chen, H.; El-Refaie, A.M.; Demerdash, N.A.O. Flux-Switching Permanent Magnet Machines: A Review of Opportunities and Challenges—Part I: Fundamentals and Topologies. *IEEE Trans. Energy Convers.* **2020**, *35*, 684–698. [CrossRef]

5. Wang, Q.; Zhao, X.; Niu, S. Flux-Modulated Permanent Magnet Machines: Challenges and Opportunities. *World Electr. Veh. J.* **2021**, *12*, 13. [CrossRef]
6. McFarland, J.D.; Jahns, T.M.; El-Refaie, A.M. Analysis of the Torque Production Mechanism for Flux-Switching Permanent-Magnet Machines. *IEEE Trans. Ind. Appl.* **2015**, *51*, 3041–3049. [CrossRef]
7. Zhu, Z.Q.; Liu, Y. Analysis of Air-Gap Field Modulation and Magnetic Gearing Effect in Fractional-Slot Concentrated-Winding Permanent-Magnet Synchronous Machines. *IEEE Trans. Ind. Electron.* **2018**, *65*, 3688–3698. [CrossRef]
8. Wang, P.; Hua, W.; Zhang, G.; Wang, B.; Cheng, M. Principle of Flux-Switching PM Machine by Magnetic Field Modulation Theory—Part II: Electromagnetic Torque Generation. *IEEE Trans. Ind. Electron.* **2022**, *69*, 2437–2446. [CrossRef]
9. Li, Y.; Li, S.; Yang, Y.; Sarlioglu, B. Analysis of Flux-Switching Permanent Magnet Machine Design for High-Speed Applications. In Proceedings of the IEEE Energy Conversion Conference and Expo (ECCE), Pittsburgh, PA, USA, 14–18 September 2014; pp. 302–309. [CrossRef]
10. Chen, J.T.; Zhu, Z.Q.; Howe, D. Stator and Rotor Pole Combinations for Multi-Tooth Flux-Switching Permanent-Magnet Brushless AC Machines. *IEEE Trans. Magn.* **2008**, *44*, 4659–4667. [CrossRef]
11. Yang, H.; Li, Y.; Lin, H.; Zhu, Z.Q.; Lyu, S. Principle Investigation and Performance Comparison of Consequent-Pole Switched Flux PM Machines. *IEEE Trans. Transp. Electr.* **2021**, *7*, 766–778. [CrossRef]
12. Bi, Y.; Fu, W.; Niu, S.; Huang, J. Comparative Analysis of Consequent-Pole Flux-Switching Machines with Different Permanent Magnet Arrangements for Outer-Rotor In-Wheel Direct-Drive Applications. *Energies* **2023**, *16*, 6650. [CrossRef]
13. Zhu, Z.Q.; Pang, Y.; Howe, D.; Iwasaki, S.; Deodhar, R.; Pride, A. Analysis of Electromagnetic Performance of Flux-Switching Permanent-Magnet Machines by Nonlinear Adaptive Lumped Parameter Magnetic Circuit Model. *IEEE Trans. Magn.* **2005**, *41*, 4277–4287. [CrossRef]
14. Gysen, B.L.J.; Ilhan, E.; Meessen, K.J.; Paulides, J.J.H.; Lomonova, E.A. Modeling of Flux Switching Permanent Magnet Machines With Fourier Analysis. *IEEE Trans. Magn.* **2010**, *46*, 1499–1502. [CrossRef]
15. Gaussens, B.; Hoang, E.; de la Barriere, O.; Saint-Michel, J.; Lecrivain, M.; Gabsi, M. Analytical Approach for Air-Gap Modeling of Field-Excited Flux-Switching Machine: No-Load Operation. *IEEE Trans. Magn.* **2012**, *48*, 2505–2517. [CrossRef]
16. Bianchi, N. Analysis and Design Remarks of Variable Flux Reluctance Motors. In Proceedings of the IEEE Workshop Electrical Machines Design, Control and Diagnosis (WEMDCD), Valletta, Malta, 9–10 April 2025; pp. 1–6. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Article

A Novel Dynamic Surge Modeling Framework for Gas Turbines: Integration of Compressor Variable Geometry

Jinshi Du, Yu Zhang *, Miguel Martínez García and Adrian Spencer

Department of Aeronautical and Automotive Engineering, Loughborough University, Loughborough LE11 3TU, UK; j.du2@lboro.ac.uk (J.D.); m.martinez-garcia@lboro.ac.uk (M.M.G.); a.spencer@lboro.ac.uk (A.S.)

* Correspondence: y.zhang@lboro.ac.uk

Abstract

Gas turbines are complex mechatronic systems that require reliable dynamic models to support automated operation under varying aerodynamic conditions. This study presents a novel dynamic surge modeling framework that integrates compressor variable geometry into a gas turbine component-level model. A physics-based formulation is developed in which the influence of inlet guide vane (IGV) deflection is incorporated through sensitivity-based parameterization and a physics-informed extension of compressor performance characteristics. The proposed framework captures the nonlinear interaction between compressor surge dynamics and component-level system behavior, enabling consistent prediction of instability onset and dynamic stability margins over a wide range of operating conditions. Model verification through stability analysis, phase-space characterization, and time-domain simulations demonstrates that the framework reproduces key features of classical compressor surge and quantifies the impact of variable geometry on system stability. The results show that the proposed model provides a practical and computationally efficient basis for control-oriented surge analysis, including stability monitoring and surge delay assessment. By coupling the IGV-aware surge dynamics with a gas turbine component-level model, the proposed method enables control-oriented, automation-ready simulation for gas turbine design and control.

Keywords: gas turbine; compressor surge modeling; variable geometry; component-level coupling; energy conversion system

1. Introduction

Gas turbines are widely utilized in modern energy conversion systems owing to their high power-to-weight ratio and operational flexibility. As the core component of a gas turbine, the compressor increases the air pressure through successive stages of rotors and stators, thereby ensuring efficient energy transfer within the gas turbine [1]. Under unstable operating conditions, the compressor would experience a surge phenomenon, in which the outlet pressure and mass flow undergo significant low-frequency oscillations. These oscillations result in a sharp reduction in overall compression capability and rotational speed, marking a loss of aerodynamic stability within the system [2–4]. Compressor surge represents a typical nonlinear dynamic instability characterized by large-scale flow reversal, intense backflow through the compressor stages, and strong periodic pressure and velocity fluctuations. Such behaviors disrupt the stable operation of the gas turbine, leading to cyclic mechanical and thermal loads on the compressor components [5]. In addition, compressor surge could accelerate material fatigue, induce mechanical failures, and degrade the overall

performance and reliability of the gas turbine [6]. Hence, a fundamental understanding of surge mechanisms and accurate dynamic modeling are indispensable for bridging the gap between instability prediction and control, which forms the basis for the efficient and autonomous operation of modern gas turbine systems.

To develop effective enough methods for suppressing or mitigating compressor surge, it is essential to first investigate its underlying physical mechanisms and dynamic characteristics, which are currently studied mainly through experimental testing and digital-twin model simulation [7,8]. Although extensive experimental studies have been conducted on surge phenomena, several limitations remain. Each experimental test involves high cost and considerable risk, as surge testing requires operating the compressor near or beyond its stability limit, where mechanical or thermal failures may occur [9]. From the perspective of digital twin modeling, the establishment of a reliable model can effectively reduce experimental expenses and shorten the validation cycle. Meanwhile, recent data-driven and machine learning approaches have also been explored for compressor health monitoring and instability prediction, aiming to detect incipient surge or stall precursors from measurement signals. For instance, Carrattieri et al. developed machine learning models for radial compressor monitoring with instability detection, demonstrating the feasibility of data-driven surveillance under realistic operating variations [10]. Thapa et al. further proposed a data-driven dynamic modeling approach for spike stall precursor detection in an axial compressor, highlighting the potential of learning-based methods for early-warning applications [11]. Despite their promise, such approaches typically rely on the representativeness of training data and may face challenges in extrapolation across unseen configurations or variable geometry schedules, which motivates physics-based models that retain interpretability and are readily integrated with system-level simulation and control. Therefore, developing and establishing a model that can accurately capture and reproduce compressor surge behavior is of great importance. With the increasing adoption of hybrid energy architectures, improvements in battery thermal management facilitate more dynamic power sharing and load-following operations [12,13]. These trends motivate robust gas turbine stability monitoring and control, since frequent transients may move compressor operation closer to surge boundaries. In parallel, Internet-distributed energy management and multi-objective design optimization for fuel-cell systems further highlight the trend toward coordinated multi-source control under time-varying conditions [14]. These developments motivate robust gas turbine stability monitoring and control, since frequent power reallocation and transients in hybrid operation can drive the compressor closer to its stability limit, making accurate surge-onset modeling practically valuable.

Research on compressor surge modeling has mainly focused on using reduced-order models to simulate the dynamic behavior of a surge. Greitzer [15,16] developed a one-dimensional nonlinear model to describe the rotating stall behavior in axial compressors and introduced the well-known parameter B to characterize the degree and type of instability. Then, Moore [17] proposed a two-dimensional model of a rotating stall to investigate the conditions that allow flow disturbances to propagate steadily through multistage compressors, even when the upstream and downstream pressures remain constant. Building on these individual contributions, Moore and Greitzer [18,19] collaboratively developed an unsteady nonlinear Moore–Greitzer (MG) model that could capture the essential dynamic characteristics of compression systems. This work represented a significant advancement in the dynamic modeling of compressor instability. Based on their work, many researchers have developed modified or extended models to address specific aspects of compressor surge and rotating stall. Haynes et al. further refined the model by incorporating blade row time-lag effects, improving its ability to capture the dynamic response [20]. Li and Zheng took humidity into account in the model [21]. Zhang et al. established separate

characteristic models for inlet distortion and blade-tip loss [22], further improving the comprehensiveness of the model. Shahriyari et al. designed an LQR controller for the MG model [23] and compared the closed-loop performance of the system. However, these studies are confined to the compressor characteristics or its external structure, without considering the strong aerodynamic coupling between the variable geometry compressor and the overall gas turbine performance. In addition, the issues of model coordination and compatibility, as well as the corresponding mapping relationship between variable geometry and surge margin, have not been further derived. In variable geometry compressors, adjustments of inlet guide vanes directly modify the flow and pressure characteristics, leading to strong nonlinear coupling between compressor aerodynamics and the overall gas turbine system. To position the proposed modeling framework within mainstream surge modeling approaches, a concise comparison is provided in terms of computational efficiency, surge reproduction capability, and engineering applicability, as summarized in Table 1.

Table 1. Comparison between mainstream surge modeling approaches and the proposed framework from computational efficiency, surge reproduction capability, and engineering applicability.

Model Family	Computational Efficiency	Surge Reproduction Capability	Engineering Applicability (Pros/Cons)	Variable Geometry Integration
MG models	High	Yes	Pros: physics-based, interpretable. Cons: not directly deployable for practical control; typically focuses on compressor dynamics.	Limited
Component-level models	Moderate	Limited (needs additional surge submodels)	Pros: high system-level fidelity. Cons: heavy calibration; surge not natural without extra modeling.	Possible but intensive
Data-driven models	High	Data-dependent (within training domain)	Pros: effective for surge monitoring/early warning. Cons: requires extensive training data; insufficient physical interpretability.	Possible (needs IGV data)
Proposed modeling framework	High	Yes	Pros: mechanism-driven and interpretable; variable geometry integrated; control-oriented. Cons: depends on characteristic parameterization.	Explicit

According to existing research, the component-level model has been widely applied to simulate the dynamic processes of gas turbines at design points and during acceleration or deceleration due to its high accuracy performance [24–26]. However, when a gas turbine enters the surge condition, a standalone component-level model cannot accurately capture the parameter variations occurring during the compressor surge process. Although relatively few studies have focused on describing the aerodynamic effects during compressor surge, the continuous evolution of modern gas turbine technology and the introduction of variable geometry structures inherently enhance overall aerodynamic performance. Therefore, developing a real-time dynamic model capable of simulating compressor surge under various gas turbine operating conditions is of great significance for understanding surge mechanisms, resolving surge faults, and advancing active stability control of such electromechanical systems. In this context, this paper proposes a real-time surge modeling framework for gas turbines that explicitly incorporates the influence of variable geometry

into the compressor–turbine coupling framework, enabling multi-configuration mapping between the surge margin and IGV setting angle. First, a component-level model of the gas turbine and a classical MG model are established, with the variable geometry mechanism introduced into the system. Then, through sensitivity analysis and physics-based modeling, a surge dynamic model under variable geometry conditions is developed. The compressor’s full-speed characteristics are extended to multiple geometric configurations, forming a mapping relationship between surge margin and geometric angle. Based on the coupling between the variable geometry surge model and the CLM, the proposed approach integrates the surge dynamics model with the gas turbine component-level model capable of steady-state and transient simulations. This integration enables the computation of key characteristic parameters throughout the entire surge process of the gas turbine. Finally, a widely used surge detection method is employed to validate the credibility of the proposed model. The developed surge model provides a solid foundation for model-based surge monitoring and control, avoiding the high risks and costs associated with surge experiments, and offers essential support for the technological advancement and stability enhancement of modern advanced gas turbines.

This paper is organized as follows. Section 2 outlines the gas turbine component-level model. Section 3 describes the proposed real-time surge modeling framework incorporating compressor variable geometry. Section 4 presents the validation and simulation results to validate the coupled CLM–surge modeling framework under various operating conditions. Finally, Section 5 concludes the study.

2. Gas Turbine Component-Level Model

The structure and sectional definitions of the gas turbine developed in this study are shown in Figure 1, where the numbers 1, 2, 31, 41, etc., represent the inlet or outlet sections of different components. The main components include the inlet duct, compressor, combustor, gas generator turbine, power turbine, and exhaust system. Meanwhile, all components satisfy the mass flow balance equations. In particular, under steady-state conditions, the components connected by the same rotor also satisfy the corresponding power-balance relationships.




Figure 1. Structural schematic of the gas turbine.

The gas turbine system is formed by linking each component model according to its physical relationships. The coupled governing equations, covering both steady-state and dynamic conditions, describe the interactions among components. Once the component-level models are established, the aerodynamic and thermodynamic parameters at each section are computed and combined into a nonlinear system of equations to obtain the system solution.

Under steady-state conditions, the rotating components are required to satisfy the power balance, while the mass flow at each section fulfills the continuity constraint. Based

on these relationships, a residual vector F is formulated, and the steady-state operating parameters are determined by solving $F = 0$.

$$F = \begin{bmatrix} F_1 \\ F_2 \end{bmatrix}, \quad F \in \mathbb{R}^5, \quad F_1 \in \mathbb{R}^3, \quad F_2 \in \mathbb{R}^2. \quad (1)$$

where the residual vector F comprises the flow-continuity residuals F_1 and the power-balance residuals F_2 .

$$F_1 = \begin{bmatrix} \frac{W_{41} - W_{41c}}{W_{41c}} \\ \frac{W_{44} - W_{44c}}{W_{44c}} \\ \frac{W_6 - W_{6c}}{W_{6c}} \end{bmatrix}, \quad F_2 = \begin{bmatrix} \frac{P_{gt}}{P_c / \eta_{gspool}} - 1 \\ \frac{P_{pt}}{P_{load} / \eta_{pspool}} - 1 \end{bmatrix}. \quad (2)$$

where the numeric subscripts denote the corresponding stations along the gas turbine flow path. Variables with the subscript “c” (e.g., W_{41c} , W_{44c} , W_{6c}) denote the calculated mass flow rates, whereas those without “c” (e.g., W_{41} , W_{44} , W_6) represent the actual flow rates at the respective sections. P denotes output power, with the subscripts gt , pt , and $load$ referring to the gas turbine, the power turbine, and the external load, respectively. η represents the shaft mechanical efficiency, where $gspool$ and $pspool$ correspond to the high-speed and low-speed spools. During the dynamic process, the flow continuity across each section of the gas turbine remains valid, whereas the power-balance condition of the rotor is no longer satisfied. The corresponding dynamic equation is expressed as follows:

$$\begin{cases} \frac{dn_g}{dt} = \left(\frac{30}{\pi}\right)^2 \frac{(P_{gt}\eta_{gspool} - P_c)}{J_{gt}n_g} \\ \frac{dn_p}{dt} = \left(\frac{30}{\pi}\right)^2 \frac{(P_{pt}\eta_{pspool} - P_{load})}{J_{pt}n_p} \end{cases} \quad (3)$$

where n_g and n_p denote the rotational speeds of the gas generator and power turbine, respectively; η_{gspool} and η_{pspool} are the mechanical efficiencies of the high-speed and low-speed shafts; and J_g and J_p are the rotational inertias of the corresponding shafts.

In steady-state conditions, the residual equations $F = 0$ described above are expected to be equal to zero. The nonlinear system is solved using the Newton–Raphson (N–R) iterative method. The iteration process can be expressed as:

$$Y^{(i+1)} = Y^{(i)} - \lambda \left[\frac{\partial F}{\partial Y} \right]_{Y^{(i)}}^{-1} F(Y^{(i)}) \quad (4)$$

where

$$Y = \begin{bmatrix} n_g & n_p & \pi_c & \pi_{gt} & \pi_{pt} \end{bmatrix}^T$$

Here, Y represents the iterative variable vector, which consists of the normalized gas generator speed n_g , normalized power turbine speed n_p , compressor pressure ratio π_c , gas generator pressure ratio π_{gt} , and power turbine pressure ratio π_{pt} . $F(Y)$ denotes the residual vector composed of the governing equations, and $\frac{\partial F}{\partial Y}$ is the Jacobian matrix of F with respect to Y . The parameter λ is the iteration step size, which controls the convergence rate of the Newton–Raphson process. The iteration continues until all elements of the residual vector converge below a predefined threshold (e.g., 10^{-5}), ensuring that both the flow-continuity and power-balance conditions are satisfied under steady-state operation.

3. Real-Time Variable Geometry Surge Modeling Framework

3.1. Variable Geometry Extended Surge Model

Conventional compressor surge models aim to describe the fundamental dynamics of compressor surge using relatively simplified physical assumptions. These models are often based on lumped-parameter formulations and are able to capture key system behaviors, such as pressure-flow oscillations. Among them, the MG model stands out as a fundamental model, providing a system-level explanation of surge dynamics. Due to its clarity and mathematical simplicity, it has been widely used in surge analysis and control design.

3.1.1. Classical MG Model

The classical MG model could capture the dynamic characteristics of surge and rotating stall in a compression system, as illustrated in Figure 2.




Figure 2. Scheme diagram of the compressor system.

The main contribution of this study is the incorporation of the static pressure-rise characteristic, the throttle valve characteristic and the B parameter to characterize the dimensionless compressor performance. The definitions of these three quantities are given as follows:

$$\Psi_c(\Phi) = \Psi_{c0} + H \left[1 + \frac{3}{2} \left(\frac{\Phi}{W_c} - 1 \right) - \frac{1}{2} \left(\frac{\Phi}{W_c} - 1 \right)^3 \right], \quad (5)$$

$$\Phi_T(\Psi) = \gamma_T \sqrt{\Psi}, \quad (6)$$

$$B = \frac{U}{2a_s} \sqrt{\frac{V_p}{A_c L_c}}, \quad (7)$$

where Ψ_c is the steady pressure-rise coefficient, Φ is the dimensionless mass flow coefficient, H and W_c denote the height and width of the static pressure-rise curve, respectively, and B is the non-dimensional Greitzer parameter that characterizes the system stability. U is the rotor tip speed, a_s is the speed of sound at the compressor inlet, and V_p , A_c , and L_c represent the plenum volume, duct cross-sectional area, and duct length, respectively.

However, the MG model examines compressor instability under the assumption of a constant rotational speed, whereby the compressor characteristics are only represented by a single axisymmetric curve. To overcome this limitation and enhance model generality, the original model was later extended to account for multiple operating speeds by [22]. The resulting multi-speed compressor surge model can be formulated in a state-space form as

$$\dot{z} = f(z), \quad z = (\Phi, \Psi, J_n, B)^T \in \mathbb{R}^4. \quad (8)$$

The governing equations are given by:

$$\frac{d\Phi}{d\xi} = \frac{H}{l_c(B)} \left(-\frac{\Psi - \Psi_{c0}}{H} + 1 + \frac{3}{2} \left(\frac{\Phi}{W} - 1 \right) \left(1 - \frac{J}{2} \right) - \frac{1}{2} \left(\frac{\Phi}{W} - 1 \right)^3 - A \right), \quad (9)$$

where

$$A = \frac{l_E U_d \Gamma \Lambda_1}{bH} \Phi. \quad (10)$$

In Equations (9) and (10), H is the height of the static pressure-rise curve, W is the width of the characteristic curve, and J denotes the amplitude. $l_c(B)$ is a length-scale function of the non-dimensional flow coefficient B , while A is an empirical parameter. Other parameters include l_E , U_d , Γ , Λ_1 , b , and H , which are characteristic geometric and aerodynamic coefficients associated with the compressor configuration.

$$\frac{d\Psi}{d\xi} = \frac{bR}{U_d B l_c} (\Phi - \Phi_T) - \frac{2\rho R^3 A_c b \Gamma}{U_d} \Psi \quad (11)$$

Equation (11) governs the transient behavior of the compressor pressure-rise coefficient Ψ . The first term reflects the influence of the mass flow deviation $(\Phi - \Phi_T)$ on the pressure response, while the second term represents the damping contribution arising from the plenum and duct dynamics. Together, these terms characterize the pressure evolution within the compression system during unsteady operation.

In this study, the speed-dependent characteristic parameters are identified based on compressor performance data, allowing the static pressure-rise and throttle characteristics to be explicitly expressed as functions of rotational speed n_g . This formulation enables the multi-speed surge model to capture variations in compressor behavior across a range of operating conditions.

Based on Equations (5) and (6), the static pressure-rise characteristics of the compressor and the throttle valve characteristics can be extended to multiple gas turbine rotational speeds n_g as follows:

$$\Psi_c(\Phi, n_g) = \Psi_{c0}(n_g) + H(n_g) \left[1 + \frac{3}{2} \left(\frac{\Phi}{W_c(n_g)} - 1 \right) - \frac{1}{2} \left(\frac{\Phi}{W_c(n_g)} - 1 \right)^3 \right], \quad (12)$$

$$\Phi_T(\Psi, n_g) = \gamma_T(n_g) \sqrt{\Psi}, \quad (13)$$

Then, the static pressure-rise curves under various rotational speeds are shown in Figure 3.




Figure 3. Static pressure-rise curves cluster at multiple rotational speeds.

3.1.2. Quantitative Sensitivity Analysis

The inlet guide vane (IGV) serves as a key variable geometry component that directly alters the incidence angle and aerodynamic loading of the compressor blades, thereby influencing both the steady-state performance and dynamic stability characteristics of the compression system. Small variations in the IGV angle can significantly affect the compres-

sor characteristics, surge margin, and overall system stability. Therefore, a quantitative sensitivity analysis is conducted to evaluate the degree of influence of the IGV angle on the compressor's aerodynamic and dynamic behaviors. This analysis provides a quantitative explanation for incorporating variable geometry effects into the surge model. The nonlinear dynamic system of the compressor can be generally expressed as:

$$\dot{z} = f(z, \theta), \quad \theta = \begin{bmatrix} \theta_p \\ \theta_n \end{bmatrix}. \quad (14)$$

$\theta = [\theta_p, \theta_n]^T$ represents the IGV deflection vector, where θ_p and θ_n denote the positive (opening) and negative (closing) deviations from the engineering design setting, respectively.

Then, considering the coupling between the surge model and the CLM, the sensitivity coefficient matrix can be further defined as:

$$S_\theta = \begin{bmatrix} \frac{\partial \Phi}{\partial \theta_p} & \frac{\partial \Psi}{\partial \theta_p} & \frac{\partial W_{31}}{\partial \theta_p} & \frac{\partial P_{31}}{\partial \theta_p} & \frac{\partial T_{31}}{\partial \theta_p} \\ \frac{\partial \Phi}{\partial \theta_n} & \frac{\partial \Psi}{\partial \theta_n} & \frac{\partial W_{31}}{\partial \theta_n} & \frac{\partial P_{31}}{\partial \theta_n} & \frac{\partial T_{31}}{\partial \theta_n} \end{bmatrix}^T. \quad (15)$$

where θ_p and θ_n denote the positive and negative IGV deflection angles, respectively. The bidirectional sensitivity matrix S_θ quantifies the asymmetric influence of IGV opening and closing on the system states, reflecting the nonlinear geometric coupling effects within the compressor flow path. The temporal evolution of the sensitivity matrix is obtained by differentiating the governing equations with respect to θ :

$$\frac{dS_\theta}{dt} = \frac{\partial f}{\partial z} S_\theta + \frac{\partial f}{\partial \theta}. \quad (16)$$

where $\frac{\partial f}{\partial z}$ represents the system Jacobian matrix describing the coupling among state variables, and $\frac{\partial f}{\partial \theta}$ characterizes the direct impact of IGV angle variations on the system dynamics.

To allow a consistent comparison between aerodynamic parameters of different dimensions, a normalized sensitivity index (NSI) is introduced to describe the relative effect of inlet guide vane (IGV) angle variations on compressor performance. This dimensionless index removes the scale differences among parameters such as flow, pressure. The NSI is defined as:

$$SI_{x_i, \theta} = \frac{\theta}{x_i} \frac{\partial x_i}{\partial \theta}, \quad x_i \in \{\Phi, \Psi, W_{31}, P_{31}, T_{31}\}. \quad (17)$$

Here, $SI_{x_i, \theta}$ denotes the normalized sensitivity of the variable x_i with respect to the IGV angle θ . A positive value of $SI_{x_i, \theta}$ indicates that the parameter increases with IGV opening, while a negative value signifies a decrease. The magnitude of $SI_{x_i, \theta}$ represents the relative strength of the variable's response, allowing direct comparison across different quantities. This formulation provides a unified basis for assessing how IGV adjustment affects the overall aerodynamic and thermodynamic behavior of the coupled compressor–turbine system.

To visualize the quantitative influence of the IGV deflection on each variable, the sensitivity coefficient matrix is illustrated in the form of a heatmap, as shown in Figure 4. The analysis is performed under standard atmospheric conditions to ensure consistency and comparability of results, with two representative IGV angles ($\theta = -10^\circ$ and $+10^\circ$) selected to introduce controlled, symmetric perturbations around the nominal setting. These values correspond to a typical moderate IGV adjustment range in practical operation and remain

within the validity envelope of the employed compressor characteristics, thereby enabling a clear and fair comparison between negative and positive deflections.

These deflection values correspond to typical vane opening and closing conditions observed in variable geometry compressors, allowing the bidirectional effects of IGV adjustment to be clearly examined. For each deflection case, the resulting changes in the aerodynamic and thermodynamic parameters (Φ , Ψ , W_{31} , P_{31} , and T_{31}) are computed, and the corresponding sensitivity coefficients are normalized and visualized.




Figure 4. Normalized sensitivity coefficient matrices of key aerodynamic and thermodynamic parameters with respect to inlet guide vane (IGV) deflection at different corrected speeds. (a) IGV closing (-10°); (b) IGV opening ($+10^\circ$). Each cell represents the normalized sensitivity of a parameter (Φ , Ψ , W_{31} , P_{31} , T_{31}) at a given corrected speed ($n_{cor} = 0.6$ – 1.0), illustrating the bidirectional influence of IGV adjustment on compressor flow, pressure, and temperature characteristics.

Figure 4 reveals distinct sensitivity distributions associated with opposite IGV deflection directions. Under IGV closing (-10°), parameters related to pressure rise (Ψ , P_{31}) exhibit strong positive sensitivities, while flow-related variables (Φ , W_{31}) respond negatively. Physically, IGV closing reduces the inlet flow incidence and effective flow area, resulting in a lower mass flow rate and a higher static pressure rise across the compressor. As the incoming air is turned more toward the axial direction, the compressor operates at a lower flow coefficient and closer to the high-pressure, low-flow region of the performance map. In this regime, the flow becomes more throttled and less responsive to further variations in the IGV angle, leading to smaller overall sensitivity magnitudes.

In contrast, IGV opening ($+10^\circ$) increases the inlet incidence and mass flow rate, shifting the compressor operating point toward the high-flow region. The aerodynamic loading on the blades is reduced, and the pressure ratio decreases accordingly. This change reverses the sensitivity signs compared with the IGV closing case. Moreover, the system becomes more responsive to geometric perturbations, but the overall sensitivity magnitude decreases because the flow field is more uniform and less prone to separation.

3.1.3. Physical Mechanism of Variable Geometry Effects

In order to reveal the impact mechanism of inlet guide vane (IGV) adjustment on the overall performance of the compressor, a physical deduction model of the influence of guide vane angle change on flow rate, pressure ratio and efficiency was established based on the velocity triangle relationship and energy equation. The variable geometry of the compressor, mainly achieved through the adjustment of the inlet guide vanes, alters the

inlet flow direction and thus changes the flow incidence on the rotor blades. This section derives the physical mechanism by which the IGV angle influences the overall compressor performance in terms of flow rate, pressure ratio, and efficiency.

At a constant rotational speed, the velocity triangles at the inlet and outlet can be expressed as:

$$\tan \beta_1 = \frac{V_x}{U - V_{\theta 1}}, \quad \tan \alpha_1 = \frac{V_{\theta 1}}{V_x}, \quad (18)$$

$$\tan \beta_2 = \frac{V_x}{U - V_{\theta 2}}, \quad \tan \alpha_2 = \frac{V_{\theta 2}}{V_x}, \quad (19)$$

where U is the blade peripheral velocity, V_x the axial velocity component, V_θ the tangential velocity component, and α , β the absolute and relative flow angles, respectively. A change in IGV setting angle θ_{IGV} modifies the outlet flow direction α_1 , thereby changing the tangential velocity $V_{\theta 1}$ and the rotor inlet incidence angle β_1 .

The geometric relationships of the velocity components and flow angles in Equations (18) and (19) are illustrated in Figure 5. The figure shows the velocity triangles at the rotor inlet and outlet, where C denotes the absolute velocity, W the relative velocity, and U the blade peripheral velocity. The angles α and β represent the absolute and relative flow angles, respectively. A variation in the inlet guide vane (IGV) angle changes the outlet flow direction α_1 , thereby altering the tangential velocity component $V_{\theta 1}$ and the rotor inlet incidence angle β_1 .




Figure 5. Velocity triangles at the rotor inlet and outlet of the compressor.

According to the Euler turbomachinery equation, the specific work input is given by:

$$\Delta h = U(V_{\theta 2} - V_{\theta 1}), \quad (20)$$

and the corresponding stage pressure ratio can be expressed as:

$$\pi_c = \left(1 + \frac{\eta_c U (V_{\theta 2} - V_{\theta 1})}{c_p T_1} \right)^{\frac{\gamma}{\gamma-1}}, \quad (21)$$

where c_p is the specific heat at constant pressure, T_1 the inlet total temperature, η_c the efficiency, and γ the specific heat ratio.

The variation of the IGV angle also affects the rotor incidence, defined as the deviation between the actual and design relative inlet angles:

$$i(\theta_{IGV}) = \beta_1(\theta_{IGV}) - \beta_{1,d}, \quad \beta_1(\theta_{IGV}) = \arctan\left(\frac{V_x}{U - V_{\theta 1}(\theta_{IGV})}\right), \quad (22)$$

Following the sign convention in Figure 5, β_1 is measured from the axial direction and $V_{\theta 1}$ is defined as positive in the rotor rotational direction; hence the rotor-relative tangential

component satisfies $W_{\theta 1} = U - V_{\theta 1}$. Here, $\beta_{1,d}$ denotes the design incidence angle, and a positive incidence ($i > 0$) implies increased deviation from the design inlet angle and is typically associated with higher aerodynamic losses due to boundary-layer separation and wake thickening.

Combining incidence and diffusion effects, the compressor efficiency can be approximated by an empirical loss model:

$$\eta_c(\theta) = \eta_0 - k_i[i(\theta)]^2 - k_D[D(\theta) - D_0]^2, \quad (23)$$

where η_0 is the design-point efficiency, and k_i, k_D are the empirical coefficients associated with incidence and diffusion losses, respectively. The diffusion factor $D(\theta)$ follows the Lieblein approximation:

$$D(\theta) = 1 - \frac{V_2}{V_1} + \frac{V_{\theta 2} - V_{\theta 1}(\theta)}{2V_1}, \quad V_1 \simeq \sqrt{V_x^2 + V_{\theta 1}^2}. \quad (24)$$

From Equations (20)–(24), it can be inferred that the IGV deflection influences the compressor performance through coupled aerodynamic mechanisms:

- Changing the inlet flow incidence and thereby the mass flow rate;
- Modifying the tangential momentum exchange and therefore the stage pressure ratio;
- Affecting diffusion and loss characteristics, which determine the overall efficiency.

In a simplified normalized form, the tendencies can be summarized as:

$$\frac{\Delta \dot{m}}{\dot{m}} \propto -\Delta \theta_{IGV}, \quad \frac{\Delta \pi_c}{\pi_c} \propto \Delta \theta_{IGV}, \quad \frac{\Delta \eta_c}{\eta_c} \propto -(\Delta \theta_{IGV})^2 \quad (25)$$

These relations reveal that IGV closure increases pressure ratio and efficiency but reduces mass flow rate, while excessive deflection (either opening or closing) degrades efficiency due to strong incidence and diffusion losses.

3.2. Coupled Framework of the Surge Model and the CLM

The variable geometry surge model is coupled with the component-level model to represent the interaction between compressor surge dynamics and the overall system response. The surge model describes the compression system in terms of dimensionless aerodynamic variables, while the component-level model is formulated using measurable thermodynamic quantities. To enable a consistent exchange of information between the two models, transformation relations are introduced to link the corresponding flow and pressure representations. The dimensionless flow and pressure-rise coefficients are

$$\Phi = \frac{\dot{m}_c}{\rho_a A_c U}, \quad \Psi = \frac{\Delta p_c}{\frac{1}{2} \rho_a U^2}, \quad (26)$$

where $\dot{m}_c$ is the compressor mass flow rate, Δp_c the total pressure rise, ρ_a the inlet-air density, A_c the flow area, and U the blade-tip speed.

The inverse relations that convert the surge model variables back to the CLM form are

$$\dot{m}_c = \Phi \rho_a A_c U, \quad \pi_c = \left(1 + \frac{\Psi U^2}{2c_p T_1}\right)^{\frac{\gamma}{\gamma-1}}, \quad (27)$$

where T_1 is the compressor inlet total temperature, c_p the specific heat, and γ the specific heat ratio. These relations allow real-time data exchange between the aerodynamic and thermodynamic subsystems.

The coupled dynamic equations of the integrated model can be expressed as

$$\begin{cases} \dot{z}_S = f_S(z_S, \theta_{IGV}, n_g), \\ \dot{z}_{CLM} = f_{CLM}(z_{CLM}, \dot{m}_c, \pi_c), \end{cases} \quad (28)$$

where $z_{MG} = (\Phi, \Psi)^T$ and z_{CLM} contains the gas turbine thermodynamic states. The variables $\dot{m}_c$ and π_c serve as the coupling interface between the two models, forming a closed-loop compressor–gas turbine system capable of reproducing surge-onset and recovery behavior.

As shown in Figure 6, the coupling is achieved through bidirectional data exchange: the surge model supplies the transient aerodynamic response, while the CLM updates the thermodynamic boundary conditions in real time. This interaction forms a self-consistent closed-loop representation of the entire gas turbine–compressor system. This framework ensures consistency between the dimensionless surge dynamics and the physical model, providing a unified basis for simulation, fault diagnosis, and control of compressor instabilities.




Figure 6. Schematic of the coupled framework between the surge model and CLM.

4. Simulation and Verification

Simulation results are presented in this section to validate the proposed variable geometry surge modeling framework.

4.1. Verification of Surge Reproduction Capability

The classical surge behavior reproduced by the proposed model is examined in this section through phase-space and time-domain analyses under different operating conditions.

4.1.1. Stability Characteristics of Surge Onset

To verify whether the proposed surge model reproduces the correct instability mechanism associated with compressor surge, a stability-based verification is carried out using a Hopf bifurcation criterion. In compressor surge modeling, Hopf bifurcation analysis is commonly used to identify the transition from steady operation to self-sustained oscillations and to define the onset of surge instability [21,27]. Within this framework, surge inception corresponds to the loss of local stability of the equilibrium point, indicated by a pair of complex-conjugate eigenvalues crossing the imaginary axis.

For verification purposes, the nonlinear surge dynamics of the proposed model are expressed in a state-space form,

$$\dot{z} = f(z, B, n_g, \theta_{IGV}), \quad (29)$$

where $z = (\Phi, \Psi)^T$ denotes the flow and pressure state vector. For a given operating condition, the equilibrium state $z_e = (\Phi_e, \Psi_e)$ is determined by the steady coupling between the compressor characteristic and the throttle relation, satisfying

$$f(z_e, B, n_g, \theta_{IGV}) = 0. \quad (30)$$

Local stability of the equilibrium point is evaluated by linearizing the nonlinear system around z_e , yielding

$$\dot{\xi} = J(z_e, B, n_g, \theta_{IGV}) \xi, \quad \xi = z - z_e, \quad (31)$$

where J denotes the Jacobian matrix of the system. Under variable geometry conditions, the Jacobian matrix explicitly depends on both the rotational speed and the IGV angle and can be expressed as

$$J(\Phi_0, \Psi_0, n_g, \theta_{IGV}) = \begin{bmatrix} \frac{1}{l_c} \frac{\partial \Psi_c(\Phi, n_g, \theta_{IGV})}{\partial \Phi} \Big|_{\Phi_0} & -\frac{1}{l_c} \\ \frac{1}{4B^2 l_c} & -\frac{1}{4B^2 l_c} \frac{\partial \Psi_T(\Psi, n_g)}{\partial \Psi} \Big|_{\Psi_0} \end{bmatrix}, \quad (32)$$

By varying the parameter B , the eigenvalues of the Jacobian matrix are tracked to evaluate the stability evolution of the equilibrium point, as shown in Figure 7. A Hopf bifurcation is identified when the real part of a complex-conjugate eigenvalue pair crosses zero while the imaginary part remains nonzero.




Figure 7. Eigenvalue evolution from bifurcation analysis: (a) three-dimensional eigenvalue trajectories illustrating the migration of the dominant eigenvalue pair with varying rotational speeds; (b) two-dimensional eigenvalue trajectories showing the real-part crossing associated with Hopf bifurcation.

The three-dimensional trajectories show that the dominant eigenvalue pair gradually migrates toward the imaginary axis as B increases. The corresponding two-dimensional projection further highlights the crossing of the real part through zero, which marks the onset of surge instability. This behavior is consistently observed across all relative rotational speeds and follows the same qualitative trend reported in [22]. These results provide further confirmation that the proposed surge model reproduces the classical surge instability mechanism via a Hopf bifurcation.

To quantify the stability boundary implied by the eigenvalue evolution, the critical Hopf bifurcation parameters are extracted for each normalized rotational speed and IGV setting. Specifically, the critical value B_{cr} is identified at the point where the dominant complex-conjugate eigenvalue pair crosses the imaginary axis, and the associated equilibrium operating point (Φ_{cr}, Ψ_{cr}) and oscillation frequency ω_{cr} are reported simultaneously. For clarity and engineering relevance, the bifurcation points are tabulated at representative discrete IGV settings $\theta_{IGV} = \{-12^\circ, -6^\circ, 0^\circ, +6^\circ, +12^\circ\}$ commonly used in practical IGV scheduling and control. The resulting bifurcation point dataset is summarized in Table 2, enabling a direct comparison of the surge stability margin and the corresponding dynamic characteristics across different operating conditions.

Table 2. Critical Hopf bifurcation points under different normalized rotational speeds and IGV settings.

N/N_{ref}	IGV Angle	B_{cr}	Φ_{cr}	Ψ_{cr}	ω_{cr}
0.6	-12°	0.4651	0.9385	2.4464	0.8415
0.6	-6°	0.4171	0.9130	2.3158	0.9184
0.6	0°	0.4191	0.8878	2.1894	0.9099
0.6	$+6^\circ$	0.4404	0.8655	2.0806	0.8701
0.6	$+12^\circ$	0.4717	0.8465	1.9906	0.8187
0.7	-12°	0.3855	0.9957	2.7540	0.9955
0.7	-6°	0.3791	0.9660	2.5917	1.0032
0.7	0°	0.3955	0.9392	2.4500	0.9637
0.7	$+6^\circ$	0.4241	0.9167	2.3342	0.9067
0.7	$+12^\circ$	0.4599	0.8983	2.2414	0.8446
0.8	-12°	0.3484	1.1210	3.4906	1.1044
0.8	-6°	0.3384	1.0878	3.2875	1.1252
0.8	0°	0.3511	1.0577	3.1072	1.0861
0.8	$+6^\circ$	0.3749	1.0320	2.9583	1.0250
0.8	$+12^\circ$	0.4058	1.0108	2.8386	0.9563
0.9	-12°	0.2952	1.2568	4.3881	1.2912
0.9	-6°	0.2982	1.2185	4.1247	1.2728
0.9	0°	0.3159	1.1853	3.9022	1.2092
0.9	$+6^\circ$	0.3407	1.1578	3.7239	1.1317
0.9	$+12^\circ$	0.3704	1.1357	3.5825	1.0510
1.0	-12°	0.2637	1.3697	5.2114	1.4390
1.0	-6°	0.2738	1.3277	4.8969	1.3865
1.0	0°	0.2941	1.2928	4.6425	1.3023
1.0	$+6^\circ$	0.3198	1.2647	4.4428	1.2104
1.0	$+12^\circ$	0.3498	1.2422	4.2861	1.1184

B_{cr} denotes the critical stability parameter at the Hopf bifurcation, and ω_{cr} is the corresponding critical oscillation frequency.

As summarized in Table 2, B_{cr} exhibits a clear dependence on the IGV setting at each normalized speed. In particular, B_{cr} reaches a minimum around $\theta_{IGV} \approx -6^\circ \sim 0^\circ$ and increases as $|\theta_{IGV}|$ becomes larger, indicating an enlarged stability margin against surge. With increasing rotational speed, the overall B_{cr} level decreases, whereas the critical frequency ω_{cr} increases, implying a faster oscillatory mode at the onset of instability. Meanwhile, the equilibrium coordinates (Φ_{cr}, Ψ_{cr}) shift systematically with IGV deflection, reflecting the IGV-induced redistribution of the steady operating point on the compressor characteristic.

4.1.2. Phase-Space Characteristics and Limit-Cycle Formation

The nonlinear surge behavior beyond the linear stability boundary is characterized through phase-space analysis in the Φ – Ψ plane. The resulting phase portraits could directly reveal the evolution of system trajectories and indicate whether the post-instability dynamics remain bounded or develop into sustained oscillations.

In this subsection, a baseline configuration without geometric variation is considered. The inlet guide vane angle is fixed at its nominal setting (0°), and the parameter B is fixed at $B = 0.7$, which is above the critical stability threshold. Under this condition, the rotational speed is varied to assess the nonlinear surge behavior across different operating points. The corresponding phase portraits for different rotational speeds are shown in Figure 8.




Figure 8. Phase portraits in the Φ – Ψ plane illustrating the formation of stable limit cycles under different rotational speeds.

As shown in Figure 8, closed trajectories are observed for all rotational speeds, indicating the formation of stable limit cycles after the loss of equilibrium stability. The size of the limit cycle increases with rotational speed, reflecting enhanced pressure–flow interaction at higher operating levels. This trend is physically reasonable, as higher rotational speeds correspond to increased compressor pressure rise and stronger coupling between flow inertia and system compressibility, leading to larger oscillation amplitudes in the surge regime. The bounded limit cycles observed across all cases are characteristic of classical compressor surge dynamics and are consistent with the expected post-bifurcation behavior of a Hopf instability, confirming that the proposed model captures the essential nonlinear mechanisms governing the transition to sustained surge oscillations.

4.1.3. Time-Domain Evolution of Surge Dynamics

While the phase-space analysis confirms the existence of stable limit-cycle oscillations after the loss of equilibrium stability, the time-domain response provides a direct view of how surge oscillations develop and persist in time. Time-domain simulations are used to examine the coupled evolution of flow and pressure during surge.

Figure 9 shows the evolution of the flow and pressure coefficients under surge conditions. Following the loss of equilibrium stability, both variables exhibit sustained oscillations with finite amplitude, indicating the establishment of a stable nonlinear oscillatory state.




Figure 9. Time-domain surge responses under different rotational speeds.

The oscillations are characterized by low frequency and large amplitude, consistent with the classical definition of compressor surge. The flow and pressure coefficients oscillate in a strongly coupled manner, reflecting the interaction between flow inertia and system compressibility that governs surge dynamics.

As the rotational speed increases, the amplitudes of both flow and pressure oscillations increase, while the overall waveform and oscillation pattern remain similar. This indicates that the same instability mechanism governs the surge dynamics across the investigated operating conditions, with higher rotational speeds intensifying the pressure–flow interaction. The presence of bounded and repeatable oscillations in the time domain confirms that the instability saturates through nonlinear effects rather than leading to divergence. Together with the linear stability and phase-space analyses presented earlier, these results verify that the proposed model reproduces the essential dynamic features of classical compressor surge.

4.1.4. Comparison with Experimental Surge Characteristics

To further corroborate the model, the simulated surge response is compared with experimental surge characteristics reported in the open literature. In practice, publicly available experimental records that provide complete transient waveforms and fully specified test conditions are limited. Therefore, the comparison focuses on commonly accepted *surge signatures* that are consistently observable across platforms, including: (i) a rapid transient trigger, (ii) finite-amplitude low-frequency oscillations in pressure and mass flow after onset, and (iii) a bounded closed-loop operating-point trajectory on the compressor map.

Figure 10 provides two typical experimental surge signatures from the literature. The time-domain pulsations of compressor outlet pressure and mass flow reported by [9] are shown in Figure 10a, and a characteristic closed-loop trajectory in the pressure ratio–mass flow plane discussed in [28] is shown in Figure 10b. These signatures are widely used as qualitative identifiers of surge in experiments and transient analyses.




Figure 10. Typical surge signatures reported in the literature: (a) typical surge pulsations reported in experiment [9] and (b) typical surge operating-point trajectory reported in experiment [28].

Figure 11 shows the simulated surge transient induced by a rapid fuel step. After the trigger, the compressor outlet pressure and mass flow exhibit sustained oscillations with finite amplitude (Figure 11b,c), which is consistent with the experimentally reported time-domain surge behavior in Figure 10a. In addition, the simulated operating-point trajectory forms a bounded closed loop on the compressor map (Figure 11d), consistent with the classical surge loop illustrated in Figure 10b.




Figure 11. Simulated surge transient induced by a rapid fuel step.

Then, a frequency-domain check is performed for the post-step outlet mass flow signal (Figure 12) to provide an additional characterization of the post-onset oscillations. The spectrum shows a distinct dominant peak at $f_{\text{surge}} = 1.072$ Hz, together with comparatively weaker broadband components, indicating a predominantly periodic low-frequency response in the post-surge regime. The identified dominant frequency lies within the surge frequency band commonly reported in the literature (0–15 Hz) [22], providing a consistent

time-scale reference for the limit-cycle oscillations observed in the time domain. Together with the bounded closed-loop trajectory of the operating point on the compressor map, these results suggest that the model captures the principal dynamic signatures of surge reported in experiments.




Figure 12. Frequency spectrum of the post-step compressor response.

4.2. Effect of Variable Geometry on Surge Stability

4.2.1. IGV-Dependent Modification of Compressor Characteristics

The inlet guide vane (IGV) angle is introduced as a geometric control parameter to parametrically modify the steady compressor characteristic, thereby affecting the equilibrium operating condition and the associated stability boundary. The analysis here is restricted to the onset-related stability behavior implied by the steady characteristic changes; detailed in-surge flow structures are not pursued. For engineering relevance, the representative IGV deflections are selected as $\theta_{IGV} = \{-12^\circ, -6^\circ, 0^\circ, +6^\circ, +12^\circ\}$, which correspond to typical discrete settings used in practical IGV scheduling and control implementations. Figure 13 shows representative static pressure-rise characteristics at three relative rotational speeds (0.6, 0.8, and 1.0 n_{des}) under different IGV deflection angles.




Figure 13. Static pressure-rise characteristics under different IGV angles and rotational speeds.

As shown in Figure 13, changing the IGV angle produces systematic and repeatable shifts of the static characteristic family across the investigated speed lines. In particular,

the characteristic curves exhibit a noticeable displacement in the low-flow region, which is the regime most relevant to surge inception. From a stability perspective, such changes influence not only the attainable operating range but also the local slope of the pressure-rise characteristic near the low-flow end, which directly enters the linearized stability properties in the present framework. These observations motivate treating the IGV setting as a control input at the level of the compressor characteristic.

To connect the modeled IGV effects with published experimental observations, we use the IGV-dependent compressor characteristics reported by Li et al. [29] as a literature reference, as shown in Figure 14. Since the original results are available only in figure form and the corresponding test boundary conditions are not fully documented, the assessment is conducted at a trend level after nondimensionalization. Specifically, characteristic curves at representative IGV angles are digitized, and the extracted experimental data points are listed in Table 3. These points are then used to reconstruct the corresponding static pressure-rise curves. The reconstructed curves are subsequently nondimensionalized and normalized with respect to the nominal-IGV case at the same corrected speed, enabling a consistent visualization of the IGV-induced characteristic shifts.




Figure 14. Experimental IGV-dependent compressor characteristics [29].

Table 3. Compressor characteristic points at three IGV settings.

$\theta_{IGV} = -10^\circ$		$\theta_{IGV} = 0^\circ$		$\theta_{IGV} = +10^\circ$	
W_c	π_c	W_c	π_c	W_c	π_c
0.986	3.341	0.954	3.495	0.934	3.547
0.988	3.318	0.956	3.477	0.940	3.507
0.991	3.294	0.959	3.455	0.943	3.481
0.993	3.269	0.961	3.433	0.949	3.429
0.996	3.242	0.964	3.410	0.954	3.376
0.998	3.215	0.967	3.386	0.960	3.320
1.000	3.188	0.969	3.362	0.965	3.262
1.002	3.160	0.971	3.337	0.969	3.201
1.003	3.131	0.973	3.312	0.973	3.137
1.005	3.102	0.976	3.286	0.976	3.072
1.006	3.072	0.977	3.260	0.979	3.005

Table 3. Cont.

$\theta_{IGV} = -10^\circ$		$\theta_{IGV} = 0^\circ$		$\theta_{IGV} = +10^\circ$	
W_c	π_c	W_c	π_c	W_c	π_c
1.007	3.042	0.979	3.233	0.981	2.903
1.008	3.013	0.981	3.206	0.988	2.591
1.009	2.862	0.986	3.063	0.994	2.404
1.007	2.747	0.988	2.856	0.980	2.937

W_c denotes corrected mass flow and π_c denotes compressor pressure ratio.

Figure 15 compares the normalized IGV-dependent static characteristics obtained from the published reference data (red curves) with those predicted by the proposed model (blue curves). The IGV characteristic modification in the proposed framework is derived from physically interpretable mechanisms, where IGV deflection changes the effective flow capacity and the pressure-rise distribution, leading to a systematic deformation of the compressor characteristic. Aerodynamically, IGV closing reduces flow capacity and shifts the characteristic toward lower corrected flow and higher loading, whereas IGV opening increases flow capacity and moves the characteristic toward higher flow with reduced loading. As shown in Figure 15, the model predictions reproduce the same IGV regulation signatures observed in the reference trends: the characteristic family shifts monotonically with θ_{IGV} , and the relative separation among the $\{-10^\circ, 0^\circ, +10^\circ\}$ cases is preserved over the main operating range. This agreement provides literature-based validation that the proposed physics-based formulation captures the correct direction and magnitude order of IGV-induced steady characteristic variations, thereby supporting the subsequent stability-boundary evaluation under variable geometry conditions.




Figure 15. Comparison with published IGV-adjustment characteristics.

4.2.2. Quantitative Evaluation of IGV Influence on the Stability Boundary

To quantify how IGV deflection affects the onset threshold over the investigated operating envelope, the critical Greitzer parameter B_{cr} is evaluated as a function of IGV angle and relative rotational speed. Within the present model, B_{cr} denotes the dynamic stability threshold associated with surge onset under the adopted operating-point definition.

Figure 16 summarizes the variation of B_{cr} with IGV deflection along different speed lines. The results indicate that IGV regulation leads to a consistent shift of the stability boundary, while the degree of sensitivity depends on the rotational speed. In addition, the influence of IGV deflection becomes less pronounced at higher relative speeds, suggesting

that the coupled compressor–system dynamics increasingly dominate the onset threshold in this regime.




Figure 16. Critical B_{cr} versus IGV angle for different relative rotational speeds.

Since the Greitzer parameter is related to the characteristic flow velocity and therefore varies with operating speed under fixed system geometry, changes in B_{cr} can be interpreted as an equivalent shift in the admissible operating range before instability. To facilitate comparison across speed lines, an equivalent stability margin is defined as the relative change of the critical parameter with respect to the nominal IGV setting:

$$M_{eq}(n_g) = \frac{B_{cr}(n_g, \theta_{IGV}) - B_{cr}(n_g, 0)}{B_{cr}(n_g, 0)} \times 100\%, \quad (33)$$

where $B_{cr}(n_g, 0)$ denotes the critical value at the nominal IGV setting for the same rotational speed.

The mapping in Figure 17 provides a compact description of how IGV deflection shifts the modeled stability boundary over the investigated speed range. In this sense, M_{eq} serves as a normalized indicator of the IGV-induced boundary movement within the present framework.




Figure 17. Equivalent stability margin M_{eq} versus IGV angle for different rotational speeds.

5. Conclusions

This study proposes a variable geometry surge modeling framework integrated within a gas turbine component-level model, with a specific focus on the dynamic role of inlet guide vanes (IGVs). By embedding IGV effects into a physics-based surge formulation, the proposed approach establishes a direct link between compressor variable geometry and system-level surge dynamics. A physically motivated parameterization is introduced to describe how IGV deflection modifies flow capacity, pressure rise, and the local slope of the static characteristic near surge inception. This extension enables the model to capture geometry-induced changes in dynamic stability, rather than treating variable geometry as a purely steady-state adjustment. The resulting framework reproduces key features of classical compressor surge, including Hopf bifurcation at instability onset, stable post-bifurcation limit cycles, and low-frequency large-amplitude oscillations in the time domain. In addition, the surge transient signatures and dominant oscillation time scale predicted by the model are shown to be consistent with representative experimental surge characteristics reported in the open literature.

A central contribution of this work is the formulation of a stability-based surge boundary that explicitly depends on both rotational speed and IGV setting. Using the critical B parameter as a unified stability indicator, the influence of IGV deflection on surge onset is quantified across the operating envelope. Within the present modeling framework, IGV opening leads to a systematic increase in the critical threshold, indicating delayed surge inception. To facilitate engineering interpretation, an equivalent stability margin is introduced and expressed in terms of an effective increase in allowable rotational speed prior to surge. This formulation provides a clear and quantitative mapping between IGV angle and surge delay capability. The predicted IGV-induced characteristic shifts are further validated at a trend level against published IGV-regulated compressor measurements, supporting the physical soundness of the proposed IGV characteristic modification. The proposed framework is primarily intended for surge-onset prediction and control-oriented stability analysis under variable geometry, and it captures the post-onset limit-cycle dynamics in the vicinity of the stability boundary. Other instability modes, such as deep surge and rotating stall, are not explicitly represented in the present lumped-parameter formulation; covering these phenomena would require additional nonlinear and spatial dynamics, which will be investigated in future work. In addition, future work will include validation against independent compressor rig test measurements under varied IGV schedules, focusing on surge-onset thresholds, oscillation characteristics, and transient signatures. It is worth noting that, although this study focuses on a gas turbine, the proposed modeling framework can be extended to other rotating turbomachinery given appropriate characteristic-map and parameter calibration.

Overall, this work provides a unified and computationally efficient tool for analyzing the dynamic impact of compressor variable geometry on surge behavior. By shifting the role of variable geometry from static characteristic modification to dynamic stability regulation, the proposed framework offers new insight into how geometry scheduling can extend the surge-free operating envelope during transient operation. The results establish a foundation for model-based surge monitoring and active stability control design in gas turbine systems.

Author Contributions: Conceptualization, J.D., Y.Z.; validation, J.D., Y.Z.; writing—original draft, J.D.; writing—review and editing, Y.Z., M.M.G. and A.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

CLM	Component Level Model
IGV	Inlet Guide Vane
NSI	Normalized Sensitivity Index

References

1. Boyce, M.P. *Gas Turbine Engineering Handbook*; Elsevier: Oxford, UK, 2011; pp. 51–56, ISBN 978-0-12-383842-1.
2. Park, Y.; Choi, M.; Kim, D.; Lee, J.; Choi, G. Performance analysis of large-scale industrial gas turbine considering stable combustor operation using novel blended fuel. *Energy* **2021**, *236*, 121408. [CrossRef]
3. Yang, Y.; Nikolaidis, T.; Jafari, S.; Pilidis, P. Gas turbine engine transient performance and heat transfer effect modelling: A comprehensive review, research challenges, and exploring the future. *Appl. Therm. Eng.* **2024**, *236*, 121523. [CrossRef]
4. Tan, C.S.; Day, I.; Morris, S.; Wadia, A. Spike-type compressor stall inception, detection, and control. *Annu. Rev. Fluid Mech.* **2010**, *42*, 275–300. [CrossRef]
5. Wang, W.; Lu, J.; Luo, X.; Huang, R.; Chu, W. Failure mechanism of casing treatment in improving stability of a highly loaded two-stage axial compressor. *Aerosp. Sci. Technol.* **2020**, *105*, 105979. [CrossRef]
6. Aprea, C.; Mastrullo, R.; Renno, C. Determination of the compressor optimal working conditions. *Appl. Therm. Eng.* **2009**, *29*, 1991–1997. [CrossRef]
7. Liu, A.X.; Zheng, X.Q. Methods of surge point judgment for compressor experiments. *Exp. Therm. Fluid Sci.* **2013**, *51*, 204–213. [CrossRef]
8. Galindo, J.; Serrano, J.R.; Climent, H.; Tiseira, A. Experiments and modelling of surge in small centrifugal compressor for automotive engines. *Exp. Therm. Fluid Sci.* **2008**, *32*, 818–826. [CrossRef]
9. Munari, E.; Morini, M.; Pinelli, M.; Spina, P.R.; Suman, A. Experimental investigation of stall and surge in a multistage compressor. *J. Eng. Gas Turbines Power* **2017**, *139*, 022605. [CrossRef]
10. Carrattieri, L.; Cravero, C.; Marsano, D.; Valenti, E.; Sishtla, V.; Halbe, C. The Development of Machine Learning Models for Radial Compressor Monitoring with Instability Detection. *J. Turbomach.* **2025**, *147*, 051004. [CrossRef]
11. Thapa, A.; Li, J.; Schoen, M.P. Spike Stall Precursor Detection in a Single-Stage Axial Compressor: A Data-Driven Dynamic Modeling Approach. *Machines* **2025**, *13*, 338. [CrossRef]
12. Qi, W.; Yu, J.; Wang, X.; Zhang, R.; Wu, X.; Xue, S.; Yang, J.; Ge, S.; Zhang, Y.; He, Z. Thermal Management Performance Enhancement of Cylindrical Power Battery Modules via Cooling Strips with Reflux Microchannel Structure. *Energy* **2026**, *346*, 140312. [CrossRef]
13. Sheng, L.; Wang, H.; Zhang, C.; Zhang, X. Experimental study on immersion cooling performance of a lithium-ion battery module for commercial/industrial energy storage at different ambient temperatures. *Int. J. Therm. Sci.* **2026**, *221*, 110473. [CrossRef]
14. Qi, W.; Yu, L.; Tang, X.; Wu, J.; Zhang, Y.; He, Z. Multi-objective optimization of a hydrogen-fueled PEMFC with multi wavy channels via machine learning and CFD simulation. *Int. J. Hydrogen Energy* **2026**, *199*, 152748. [CrossRef]
15. Greitzer, E.M. Surge and rotating stall in axial flow compressors—Part I: Theoretical compression system model. *J. Eng. Power* **1976**, *98*, 190–198. [CrossRef]
16. Greitzer, E.M. Surge and rotating stall in axial flow compressors—Part II: Experimental results and comparison with theory. *J. Eng. Power* **1976**, *98*, 199–217. [CrossRef]
17. Moore, F.K. A theory of rotating stall of multistage axial compressors: Part I—Small disturbances. *J. Eng. Gas Turbines Power* **1984**, *106*, 313–321. [CrossRef]
18. Moore, F.K.; Greitzer, E.M. A theory of post-stall transients in axial compression systems: Part I—Development of equations. *J. Eng. Gas Turbines Power* **1986**, *108*, 68–76. [CrossRef]
19. Greitzer, E.M.; Moore, F.K. A theory of post-stall transients in axial compression systems: Part II—Application. *J. Eng. Gas Turbines Power* **1986**, *108*, 231–239. [CrossRef]
20. Haynes, J.M.; Hendricks, G.J.; Epstein, A.H. Active stabilization of rotating stall in a three-stage axial compressor. *J. Turbomach.* **1994**, *116*, 226–239. [CrossRef]

21. Li, M.; Zheng, Q. Wet compression system stability analysis: Part I—Wet compression Moore–Greitzer transient model. In *Proceedings of the ASME Turbo Expo: Power for Land, Sea, and Air*; ASME: Vienna, Austria, 2004; Volume 41693, pp. 703–712. [CrossRef]
22. Zhang, X.; Zhang, T.; Sheng, H. A novel aeroengine real-time model for active stability control: Compressor instabilities integration. *Aerosp. Sci. Technol.* **2024**, *145*, 108844. [CrossRef]
23. Shahriyari, M.J.; Firouzabadi, A.; Khaleghi, H.; Esmailifar, S.M. A new model for compressor surge and stall control. *Sci. Rep.* **2024**, *14*, 5347. [CrossRef]
24. Al-Hamdan, Q.Z.; Ebaid, M.S.Y. Modeling and simulation of a gas turbine engine for power generation. *J. Eng. Gas Turbines Power* **2006**, *128*, 302–311. [CrossRef]
25. Larsson, E.; Åslund, J.; Frisk, E.; Eriksson, L. Gas turbine modeling for diagnosis and control. *J. Eng. Gas Turbines Power* **2014**, *136*, 071601. [CrossRef]
26. Xu, M.; Liu, J.; Li, M.; Geng, J.; Wu, Y.; Song, Z. Improved hybrid modeling method with input and output self-tuning for gas turbine engine. *Energy* **2022**, *238*, 121672. [CrossRef]
27. Zhang, X.; Li, L.; Zhang, T. Research on real-time model of turboshaft engine with surge process. *Appl. Sci.* **2022**, *12*, 744. [CrossRef]
28. Chung, K.; Leamy, K.; Collins, T. A turbine engine aerodynamic model for in-stall transient simulation. In *Proceedings of the 21st Joint Propulsion Conference*, Monterey, CA, USA, 8–11 July 1985; p. 1429. [CrossRef]
29. Li, P.; Zuo, Z.; Li, J.; Tao, H.; Guo, W.; Chen, H. Characteristics of inlet guide vane adjustment of multi-stage axial compressor in compressed air energy storage system. *J. Energy Storage* **2023**, *72*, 108342. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Article

Mechanical Design and Experimental Study of a Small-Scale Wind Turbine Model

Eduardo Muñoz-Palomeque ^{1,*}, Segundo Esteban ¹ and Matilde Santos ²

¹ Department of Computer Architecture and Automation, Complutense University of Madrid, 28040 Madrid, Spain; sesteban@ucm.es

² Institute of Knowledge Technology, Complutense University of Madrid, 28040 Madrid, Spain; msantos@ucm.es

* Correspondence: edumun04@ucm.es

Abstract: The advancement of onshore and offshore wind turbines depends on the experimental validation of new technologies, novel component designs, and innovative concepts. However, full-scale models are typically very expensive, have limited functionality, and are difficult to adapt to diverse research needs. To address this shortcoming, this article presents the design of a low-cost, modular 3D-printed small prototype of a wind turbine. It includes a multi-hollow platform for marine environments configuration and stabilization, the turbine tower, and three blades with active pitch control, not always included in wind turbine prototypes. The modular tower design allows for easy height extensions, while the rotor incorporates custom blades optimized for the prototype geometry and experimental setup. Tests were conducted to evaluate the system's operational response and verify the proper functioning of the assembled components at various wind speeds and blade pitch angles. The results confirm that the rotor speed with the prototype's onshore configuration is highly pitch-dependent, reaching a maximum efficiency of approximately 5°. The tower displacement, measured with an IMU, remained within a narrow range, oscillating around 2° and reaching up to 4° at higher wind speeds due to elastic deflections of the PLA structure. These results, consistent with the prototype scale, validate its usefulness in capturing essential aerodynamic and structural behaviors of the wind turbine. They also demonstrate its relevance as a new tool for experimental studies of wind turbines and open up new research, validation, and control possibilities not considered in previous developments by incorporating blade pitch control.

Keywords: mechanical design; physical prototype; experimental tests; wind energy; wind turbine

1. Introduction

The planet is immersed in a continuous process of evolution, and global energy generation is one of them. Energy demand is constantly increasing due to rising human needs, population expansion, industrialization, and technological progress. Meeting these needs has become a priority and a challenge, placing great importance on sustainable solutions as viable alternatives. Fossil fuels, primarily coal and natural gas, have been the main sources of energy until recently, but they produce adverse effects as they contribute to greenhouse gas emissions and climate change, causing serious environmental impacts such as air and water pollution and biodiversity loss [1,2].

To address these challenges, renewable energy sources such as wind, solar, hydro-electric, and geothermal energy are used, and these solutions have experienced significant

expansion in recent decades. Using these natural resources presents notable advantages, as they are abundant on the planet and produce few or no emissions during operation [3]. Due to these characteristics, they can help mitigate the negative effects that polluting procedures have on life and the planet. Furthermore, they are less expensive and competitive compared to fossil fuels. Recent data shows that most new renewable energy generation projects are now cheaper than traditional alternatives [4,5].

The growing need for renewable energy sources has driven the development of more efficient and adaptable technologies. Among them, wind energy has proven to be a suitable option due to advances in turbine design, improved control systems [6] and thanks to the reduction in manufacturing, installation and maintenance costs [7]. All this has made wind energy today one of the largest sources of renewable energy [8]. In search of the best locations for wind turbines, production has shifted to marine areas, both on the coast (bottom offshore wind turbines) and offshore (floating offshore wind turbines, FOWT). The latter find better conditions in deep waters, with more consistent and stronger wind flows, more available space, and less visual and noise impact [9,10].

These systems have experienced massive expansion in installed capacity over the past few decades [11]. By the end of 2024, the cumulative global installed wind power capacity reached approximately 1136 GW, of which 83 GW corresponded to offshore wind. Figure 1 shows the percentage distribution of this cumulative offshore capacity by country, with China, the UK, and Germany together accounting for 80.4%. In terms of annual additions during 2024, approximately 109 GW of new onshore wind and 8 GW of new offshore wind were installed [12]. Offshore deployments are expected to increase, with predictions exceeding 500 GW by 2050, reflecting the maturity of floating offshore wind turbines (FOWTs) [13].




Figure 1. Percentage distribution of worldwide cumulative offshore wind energy capacity by country (Data source: [12]).

Marine wind turbines are classified into two groups according to their foundation type: anchored to the seabed and floating ones. This categorization has implications for the water depth and distance from the coast where they can be installed [14]. The first group is located in shallow waters, with fixed-bottom foundations up to approximately 60 m deep. These structures commonly include monopile foundations, which consist of large steel tubes inserted into the seabed, and jacket structures, which consist of grid structures supported by several piles.

On the other hand, floating offshore wind turbines (FOWTs) are installed in deeper waters (more than 60 m), where fixed foundations cannot be installed due to the coastal

slope. According to the structure of these turbines, which are anchored to the seabed by mooring lines to prevent drift, three types can be identified. The spar buoy platform is characterized by a long, cylindrical tower, which provides stability by ballast and is connected to the seabed [15]. The second group corresponds to Tension Leg Platforms (TLPs), which are designed as structures with reduced vertical movement achieved with tensioned tendons anchored to the seabed [16]. In the third category are semi-submersible platforms as well as barge-type platforms, popular designs that present greater stability in different marine conditions and are suitable for large wind turbines (LOWT) with high production capacity [17].

The large size of FOWTs and the harsh environment in which they operate have prompted research into electrical stability and grid connection, energy tracking and pitch control strategies, platform stabilization, wind prediction, and vibration reduction [18–20]. Another important aspect to consider is the design of the components of these large turbines [21]. The structure provides the support and functionality necessary for the device's normal operation.

For the process of perfecting the various parts of a wind energy device, simulations, isolated tests, and more complete representations of this device in software are usually run. While numerical simulations for refining the design remain indispensable, it is necessary to advance in the analysis and study of the wind turbine in physical models that integrate elements of reality. The use of pilot wind turbines for testing, demonstrations, or research is a necessary step to bring solutions closer to commercialization [22]. The use of small-scale prototypes allows not only the structural representation of the wind turbine, but also the incorporation of the corresponding electronics and control systems, and the testing of the operation of these devices.

Recent works have addressed these gaps by using experimental setups with wave-tank [23], additive manufacturing [24], and modular platforms [25]. Although these approaches are highly relevant, some challenges are associated with scaling and cost in wave-tank testing, reduced mechanical strength of 3D-printed parts, and structural integrity challenges in modular platforms. For example, Bianchini et al. [26] reviewed the current status of small wind turbine prototypes, discussing some of the limitations. One of the most crucial ones is the lack of active pitch implementation due to its complexity, something that we address in this study. In [27], the research highlights the viability of low-scale wind turbines to be integrated into buildings for energy production, taking advantage of its modularity, reduced cost and simple structure.

Similarly, several studies [28–30] have shown that computer-aided design (CAD) software and 3D printing can be successfully used for prototyping wind turbine blades and other small-scale components of wind turbines. The use of these tools also facilitates aerodynamic model analysis, design refinement, and experimental testing. In this regard, most of the work is conducted in the laboratory or on a small scale, but it offers promising results that can be extrapolated to other levels to improve the mechanical performance and structural design of commercial wind devices.

Some small-scale wind turbine prototypes are reported in the literature, generally focusing on stall or yaw control. In addition to the limitations of all pilot wind turbines, such as high manufacturing costs, lack of modularity, and insufficient functionalities, they are rarely built to allow active pitch control due to the complexity of the elements required for their implementation [26,31,32]. This limitation restricts their ability to reproduce the aerodynamic and operational characteristics of large-scale wind turbines.

To address the need to optimize the design and configuration and improve the performance and functionality of the wind system, this work aims to design a small-scale wind

turbine that is compact and easily modifiable so that the analysis of its operation allows the transfer of study conclusions to real turbines.

To this end, the mechanical design of a prototype onshore/offshore wind turbine was performed, focusing on its structural topology and geometry, prioritizing the versatility of the system for modular modifications. Furthermore, the prototype addresses a common limitation in small-scale wind turbine models by incorporating blade pitch control as an additional degree of freedom. CAD software (FreeCAD 1.0) was used for geometric modeling of the device, while manufacturing was achieved using 3D printing.

The main contributions of this work can be summarized as follows:

- Design of a small-scale land-based/floating wind turbine prototype, composed of mechanically coupled parts, suitable for laboratory experimentation.
- Development of a modular system that facilitates reconfiguration and resizing. This modularity allows for rapid assembly, disassembly, and replacement of parts, without the need to redesign the entire structure.
- Integration of blade pitch motion into the prototype, something not very common, allowing for the implementation of pitch control strategies in addition to rotor speed regulation.
- The design allows for the adaptation of two motor/generators of different sizes and capacities, allowing the assembly to be used as systems with different power ratings (capacities).
- Experimental study of the modular prototype by acquiring key operational variables, including voltage, rotor speed, and tower pitch angle, the latter serving as an indicator of structural vibrations.

The manuscript is organized into different sections that describe the mechanical design proposed in this work. Section 2 presents the methodological approach used in the development of the wind turbine prototype. Section 3 details the design of each component of the wind device, considering both the geometry of the parts and their assembly feasibility. Section 4 presents the model validation based on data acquired from the prototype. Finally, Section 5 summarizes the main conclusions and potential directions for future research.

2. Methodology

The wind turbine prototype design was developed using a component-based, modular approach, aiming to ensure scalability, reconfigurability, and flexibility. Each component was designed as an independent module to facilitate rapid replacement and allow modification and experimentation with different configurations and geometries, without requiring a complete turbine redesign.

The methodology used in this study aligns with a systematic process often applied in physical system design. Initially, background research was performed and essential operational requirements for the system were defined, followed by the conceptual design stage and CAD modeling of each component. Then, the models were then fabricated using 3D printing technology and assembled to evaluate modularity, compatibility, and functionality, checking how well the parts fit together and how easily the system could be reconfigured. Finally, experimental validation was performed to assess the operational behavior of the prototype under controlled wind conditions. This approach enabled iterative improvements and is consistent with conceptual design methodologies for mechatronic systems, particularly those applied to wind turbine prototypes [33].

The modules designed for the small-scale wind turbine model are described below:

- **Blades:** Three blades spaced 120° apart are integrated into the horizontal axis of the wind turbine. The three-blade geometry features a 15° curvature along its length to allow for different angles of attack as the blade changes its pitch motion. The

choice was based on the aerodynamic behavior of small-scale turbines at low Reynolds numbers, where blade curvature improves the lift-to-drag ratio and compensates for the lower aerodynamic efficiency [34].

- Tower: A segmented tower was implemented to allow height adjustments and facilitate transport and reconfiguration, consistent with recent reviews highlighting modular architectures as effective for exploring structural adaptability in offshore wind turbine prototypes [35].
- Platform: The barge platform is a component where the tower is attached and allows the wind turbine to be adapted to an offshore configuration. Additionally, hollows have been implemented in the platform to simulate aspects of flotation and conduct stability experiments when necessary.
- Connection system: A bayonet-type connector was selected to join the tower segments, as it provides torsional rigidity and ease of assembly/disassembly.

The prototype was developed with a focus on modularity in key elements such as the blades, rotor, and tower. Three-dimensional representations of these geometric components were created using CAD software (FreeCAD 1.0) to allow for the evaluation of different configurations.

From the digital model, two physical prototypes were manufactured using fused deposition 3D printing (FDM). Printing was carried out on a Bambu Lab X1C printer (Bambu Lab, Shenzhen, China) using generic 1.75 mm PLA material. This filament was chosen due to its availability, low cost and environmental friendliness, as well as its good dimensional accuracy under optimized 3D printing parameters, making it suitable for iterative prototyping [36].

Indeed, PLA has proven to be an ideal material for the initial phase of prototyping, thanks to its easy and stable printing, as it does not warp upon cooling, allowing for highly detailed parts with good surface finishes from the first attempt. Its low cost and biodegradability make it suitable for rapid iteration in research environments. However, it has some mechanical limitations, such as fragility and low impact resistance, which can lead to cracking under sudden stress or continuous tension. It also has low thermal resistance (softening around 60 °C). Nevertheless, under the operating conditions of this study and the proposed future research, these conditions, which might deform the model, are not expected to occur.

Some printing parameters used for this work were:

- Nozzle temperature: 220°
- Bed temperature: 55°
- Layer height: 0.2 mm
- Infill density: 15%
- Print speed: 100%
- Assembled parts were designed with a dimensional tolerance of ± 0.25 mm, enough to ensure proper fitting of bayonet connectors, blades, tower and motor housing.

The modular components are assembled manually to verify ease of construction and compatibility with the electrical components integrated into the prototype.

The two small-scale models assembled are one with a 502 mm diameter rotor and the other with a 438 mm diameter, with the tower segments being changed for each device.

The smaller model was used for initial experimentation. Variables such as voltage, rotor speed, and tower angular displacement were measured in a fixed-base configuration.

The prototype is equipped with sensors to capture data and a microcontroller, an Arduino NANO board (Arduino, Somerville, MA, USA), to compile it. The Arduino device is located on the tower, ensuring that the integrated optical sensor points to the plane of the blades swept area to measure their intersection and calculate the rotational speed.

In addition, an Inertial Measurement Unit (IMU) mounted on top of the tower is used to detect tower movement while the wind turbine is operating. During the experiments, ambient conditions were also monitored, including airflow (measured with an anemometer) and temperature.

The electromechanical system includes a small DC motor connected directly to the rotor. No external loads were applied during this testing phase, as the objective was to observe the system's natural response and verify component assembly, rather than optimize power output. Testing focused solely on structural response and functionality, not power optimization or full aerodynamic characterization.

The experiments were conducted in a laboratory environment at 26 °C, with steady wind speeds of 3.7 to 5.5 m/s.

The design and testing strategy adopted in this work follows the rationale of recent experimental methodologies for floating offshore wind turbines, which emphasize the importance of scaled physical models to validate structural responses before advancing to more complex environments [37].

3. Component Topology and Geometry

3.1. Rotor Blades

A customized NACA 4412 airfoil was used as a basis for designing the blade geometry. This choice was motivated by the widespread use of airfoils from this series in small-scale wind turbine research, as they provide reliable aerodynamic performance at low wind speeds with acceptable lift-to-drag characteristics [38]. The blade design incorporates a geometric twist that progressively decreases from 15° to 0° from root to tip, improving aerodynamic response across the span by locally adjusting the angle of attack.

The blades were designed to attach to the hub with adequate clearance from the tower, facilitating their replacement without the need for specialized tools.

Two sets of blades were manufactured. The first extends 231 mm from the blade root connection hole to the tip (Figure 2a). For the second, smaller configuration, the blades have a length of 199 mm (Figure 2b). The second configuration, assembled with the remaining wind turbine components described below, was designed considering the spatial constraints of the experimental facilities. In particular, this configuration was sized to fit into a wind tunnel and wave tank (0.65 m high and 0.55 m wide) [21], while maintaining geometric consistency with the rotor and the overall system. This ensures compatibility with future experimental work. It should be noted that the blades were not aerodynamically optimized; the goal was to reproduce a representative rotor geometry on a small scale suitable for modular prototyping.

3.2. Nacelle and Supports

The nacelle incorporates a prefabricated 90° bevel gear system, which acts as the main mechanical transmission of motion between the rotor and the generator. The horizontal shaft, aligned with the rotor axis, transfers rotational motion through the bevel gears to a perpendicular vertical shaft connected to the generator. The gear system maintains a 1:1 transmission ratio, directly connecting both gearbox shafts without altering the angular velocity. This system was chosen to preserve the kinematic behavior of the rotor, maintaining the mechanical simplicity and robustness necessary for small-scale operation in a laboratory. Bevel gears are commonly used in wind turbines and industrial systems for their ability to transfer motion between perpendicular shafts while maintaining efficient compactness [39].




Figure 2. CAD model of the blade with front, side, and section views: (a) First configuration; (b) Second configuration.

This mechanism connects to the blade connecting piece through the blade-hub connector, which also allows blade pitch adjustment. This arrangement allows the angular position of the blades to be modified with a simple actuating lever, facilitating preliminary studies of pitch variation. Figure 3 illustrates the transmission mechanism, showing both the front (a) and side (b) views, which constitute the central component of the nacelle.




Figure 3. View of the motion transmission mechanism: (a) Frontal view; (b) Lateral view.

To visualize the integration of the blades, Figure 4 presents a CAD model of the three-blade rotor mounted on a simplified representation of the hub. This highlights the modular

design strategy, which facilitates assembly, disassembly, and possible reconfiguration of the rotor system without the need to redesign the hub.



Figure 4. CAD model of the three-blade rotor.

3.3. Modular Tower

The tower was designed as a segmented modular structure, composed of tubular sections connected by a quarter-turn bayonet mechanism, reinforced with a 3D-printed pressure coupling ring. This configuration allows the height to be modified by replacing the top segment, adapting the structure to different experimental scenarios, or using a different wind turbine model. Modular tower concepts have also been addressed in the literature; for example, in [9], the authors review advances in tower design, including segmented architectures, emphasizing structural integrity, scalability, and adaptability in offshore and onshore towers.

The use of bayonet connectors was due to their ability to provide torsional rigidity while ensuring rapid assembly and disassembly. The lower tower segment is designed to fit onto the platform and accommodate a commercial 36 mm diameter direct current (DC) motor. Its geometry is shown in Figure 5. The height is set at 117.5 mm with an outer diameter of 13 mm. The tower segment mounting mechanism is located at the top, where the internal path of the quarter-turn hole and the external rail for the connecting ring can be identified. The tower base therefore features an adjustable system that allows it to be attached to different test stands, such as floating platforms or onshore wind energy benches. The top section features a connection system to other segments to modify the overall tower height.

To complete the tower assembly, the upper section is connected. This segment extends the tower's height and can be designed with different geometries to modify the wind turbine's topology and allow for the coupling of components using other arrangements, thus creating a unique model.

In this paper, two wind turbine model configurations of different tower heights and blade sets were designed to demonstrate the flexibility of this approach.

In the first wind turbine configuration, the prototype was built with a taller tower and 231 mm blades. In this case, the top section (Figure 6) directly connects the vertical drive shaft to the motor using an aluminum bar inside the tower to transfer angular rotation from the gear shaft.




Figure 5. Bottom section of the tower: (a) Drawing; (b) 3D isometric view.




Figure 6. Upper section of the tower for the first configuration: (a) Drawing and section view (A-A); (b) 3D isometric view.

In the second WT configuration, the prototype employed a shorter tower and 199 mm blades. Here, the upper segment (Figure 7) was designed as an extension that incorporates a motor housing on top of the tower, allowing direct coupling to the gearbox. This new design reduces shaft alignment limitations and provides greater flexibility for testing a smaller motor-generator in the nacelle area.




Figure 7. Upper section of the tower for the second configuration: (a) Drawing and section view (A-A); (b) 3D isometric view.

The structure that houses the motor and connects to the gearbox mechanism is shown in Figure 8a,b. This module was designed to attach a generic 5 V DC motor to the top of the tower and connect to the mechanical transmission module that supports the rotating parts of the wind turbine. Additionally, a top bracket was included to locate and attach a servomotor used to control the blade pitch angle. The gearbox and generator are coupled by a 3D-printed shaft (Figure 8c,d), which transmits rotational motion. This component was printed in PLA with a fine mesh structure and 80% infill, providing adequate strength to withstand the stresses that may arise during normal wind turbine operation.




Figure 8. Motor-gearbox housing structure for the second configuration: (a) Drawing; (b) 3D isometric view; (c) Drawing and section view (A-A) of the motor-gearbox connection axis; (d) 3D isometric view of the axis.

The two tower segments (base and top) are joined by a quarter-turn bayonet joint (Figure 9a), where the bottom of the upper tower segment is attached to the top of the lower segment by a connecting ring (Figure 9b). This system was selected to strengthen the connection between the two segments and provide rigidity to the tower structure.





Figure 9. Junction system: (a) Quarter-turn bayonet system; (b) Connection ring.

3.4. Base Platform


The platform is designed as a robust barge that provides stable support for the tower. The WT platform design is intended to be suitable for future studies of both onshore and floating turbines, so it has been equipped with features that can facilitate the analysis of both its stability and hydrodynamics.

Two key factors were taken into account when defining the platform size. First, the limitations of the facilities where the prototype will be experimentally tested, and second, the height of the platform to ensure its buoyancy. Based on these criteria, the platform geometry shown in Figure 10 was constructed, in which some of its main dimensions are specified. It is a rectangular barge-platform measuring $240 \times 310 \times 60$ mm. The vertical length was set at 60 mm, based on the total mass of the wind turbine, to ensure that submersion does not exceed 30 mm. It also includes a space to securely house the electronic components inside.

It has eight 30 mm long square moonpools or hollows, evenly distributed at the front and rear. These hollows are designed to regulate wave-induced forces, thereby reducing wave thrust on the structure and contributing to the platform's overall stability. While the hydrodynamics of a platform with these hollow sections is a complex problem requiring specific analysis, previous studies have shown that they can effectively mitigate wave-induced motions on floating structures [40].



(a)



(b)

Figure 10. Wind turbine platform design: (a) Drawing and section views (A-A, B-B, C-C); (b) 3D isometric view.

After all the components of the wind turbine prototypes were designed, they were assembled. The fully assembled wind system, including all its subsystems, is shown in

Figure 11a. For a better understanding of the overall structure, an exploded view of the elements is presented in Figure 11b. This illustration helps visualize how the subsystems and individual parts are organized and integrated into the final assembly. Using the different component configurations designed, such as smaller or larger blades, and the upper tower section to house the generator above or below the structure, two different wind turbine models can be built.




Figure 11. Design of the wind device: (a) Fully assembled floating offshore wind turbine prototype; (b) Exploded view showing the structural components.

The dimensions of the assembled physical models are summarized in Table 1 for each wind turbine version designed. It can be seen that the size and configuration can be varied to easily create different models with the modular structure.

Table 1. Geometric dimensions of the wind turbine models.

Characteristic	Value	
	Configuration 1	Configuration 2
Bounding x dimension (width)	502 mm	438 mm
Bounding y dimension (height)	610 mm	546 mm
Bounding z dimension (depth)	310 mm	310 mm
Platform to rotor center	300 mm	267 mm
Radius of the turbine	251 mm	219 mm
Platform width	240 mm	240 mm
Platform height	60 mm	60 mm
Platform depth	310 mm	310 mm
Pitch motor position	Platform level	Tower-top level
Power generator position	Under the tower	At the tower-top

4. Results from Experimental Tests

The experimental evaluation focused on verifying the operational capability of the prototype wind turbine (using the second tower configuration) under controlled wind conditions. Available sensors were used for data acquisition to assess its performance, focusing on voltage/power generation. The experiments were conducted using a land-based wind turbine with a fixed platform. Data were captured in real time for visualization and analysis, with wind speeds as the external input to the system. The purpose of this initial study is to validate the design and integration of the subsystems (rotor, pitch mechanism, sensors, and modular tower) and assess whether the prototype reproduces the expected aerodynamic trends.

The physical specifications of the system are shown in Table 2. In addition, Figure 12 provides a real view of the assembled prototype with a length scale, illustrating its main dimensions and proportions.

Table 2. Technical specifications of the wind turbine model.

Label	Value
Number of blades	3
Maximum Voltage	60 mV
Maximum Power	0.036 mW
Rotor speed	335 RPM
Rotor diameter	0.43800 m
Sweep area	0.15067 m ²
Center of mass	(0, 0.02434, 0.23232) m from tower base
Gearbox ratio	1:1




Figure 12. Photograph of the assembled prototype with reference scale.

In these experiments, the wind speed and blade pitch angle are varied over a range of values. Specifically, three wind speeds and five angular positions are combined:

Wind speeds = [3.70, 4.50, 5.50] m/s

Pitch angles: = [0, 5, 10, 15, 20] degrees

The variables measured during the experimental tests include rotor speed and generated voltage, both related to power generation, as well as the tower's pitch motion, which can serve as an indicator of structural deflection and vibrations.

The generated, measured and filtered voltage from the pilot wind turbine is shown in Figure 13. This variable is considered equivalent to the power produced by the WT prototype for this initial analysis.




Figure 13. Generated voltage signals at different wind speeds and blade pitch angles.

The color code in Figure 13 and following corresponds to wind speeds of 3.70 m/s (blue line), 4.50 m/s (red line), and 5.50 m/s (orange line) for pitch angles of 0, 5, 10, 15, and 20 degrees (from top to bottom). The average value of each signal is presented in Table 3, where the maximum energy obtained with the prototype is 51.06 mV, with a wind speed of 5.50 m/s and a pitch angle of 5 degrees, indicating the approximate position of the blade for maximum wind energy capture.

Table 3. Voltage (mV) and speed (rpm) measured under different wind speeds and pitch angles.

		Pitch Angle (deg)				
		0°	5°	10°	15°	20°
Wind speed (m/s)	3.70 m/s	21.39 mV	27.6 mV	25.75 mV	25.99 mV	19.97 mV
		173.39 rpm	229.28 rpm	216.64 rpm	172.32 rpm	136.05 rpm
	4.50 m/s	28.76 mV	40.8 mV	29.97 mV	29.8 mV	21.85 mV
		222.75 rpm	278.13 rpm	226.07 rpm	193.09 rpm	146.58 rpm
	5.50 m/s	34.25 mV	51.06 mV	34.69 mV	33.6 mV	25.27 mV
		249.17 rpm	290.99 rpm	245.22 rpm	202.44 rpm	161.91 rpm

The results in Figure 13 confirm how power generation varies according to wind speed and blade pitch angle, as expected, obtaining reasonable and relatively stable values. As is well known, adjusting the pitch angle is crucial as it modifies the angle of attack of the blades with respect to the incoming wind. This variation directly influences the aerodynamic lift and drag forces acting on the blades and affects the torque produced on the rotor and, thus, the amount of wind energy extracted by the turbine [41].

Similarly, Figure 14 shows the rotor speed for the same wind speed and pitch angle scenarios (the average speeds are presented in Table 3). When pitch angles are set too low, the effective angle of attack is reduced, limiting the lift effect and, therefore, the torque on the rotor. On the other hand, at higher pitch angles, the angle of attack is excessive, increasing the drag force and reducing aerodynamic efficiency.




Figure 14. Rotor speed signals at different wind speeds and blade pitch angles.

While this small-scale prototype is not designed to replicate the full aerodynamic complexity of a commercial wind turbine, the results confirm the expected trend: higher wind speeds, combined with an appropriate pitch angle, generally produce higher rotational speeds and greater energy extraction.

It is important to note that, in a full-scale turbine, the actual power output depends not only on the rotor speed, but also on the torque generated by the blades and the characteristics of the generator and control system. Under the test conditions, the maximum rotor speed for the designed prototype occurred with a blade pitch angle of approximately 5° ,

which was the best configuration within the limits of the experimental setup, reproducing the general aerodynamic trends reported in the literature for small-scale wind turbines.

On the other hand, it was possible to measure a variable that provides information about the structure's deflection, specifically the tower's angular displacement. In each test, data were collected on the tilt angle at the top of the tower. Figure 15 represents this angle ($^{\circ}$) under different wind conditions. They are quite similar for all experiments with different blade pitches. The value of this metric has an average amplitude of approximately 0.5° , reaching a maximum of 4° when the rotor operates at maximum speed, a condition in which the system experiences significant dynamic behavior.




Figure 15. Tower pitch deflection at different wind speeds and blade pitch angles.

These measurements correspond to the elastic deflection of the PLA tower during tests with the platform fixed and are therefore not related to the platform's motion. The slight inclination it experiences is due to the prototype's structural flexibility under the nacelle and rotor loads and shows that the structure is sufficiently stable for aerodynamic testing.

It is important to emphasize that, while the PLA prototype proved to be structurally suitable for testing at this scale, the structural deflection results are not directly scalable to full-size turbines or even larger wind tunnel models because stiffness does not follow linear scaling law. Future work should explore alternative composite materials or manufacturing techniques for studies requiring dynamic structural similarity.

Beyond these structural considerations, evaluating the floating platform is essential for offshore adaptation. As a first approximation of the floating wind turbine configuration, a simple buoyancy test was performed with the platform. For this purpose, the prototype was placed in a water tank.

The mass of the designed platform is 1.250 kg, and the mass of the wind turbine for the tested configuration is 0.260 kg. The density of water, ρ_{water} , is 1000 kg/m³. Therefore, according to the buoyancy principle, the weight (W_{WT}) of the system and the buoyancy force (F_b) must be equivalent to ensure buoyancy:

$$F_b = W_{WT} \quad (1)$$

$$\rho_{water} \cdot V \cdot g = m_{wt} \cdot g \quad (2)$$

$$V = \frac{m_{wt}}{\rho_{water}} \quad (3)$$

where m_{wt} (Kg) is the mass of the whole device and g (m/s²) is the gravity. Using this expression, the volume of water displaced by the system, V (m³), can be obtained as follows:

$$V = \frac{m_{wt}}{\rho_{water}} = \frac{1.510 \text{ Kg}}{1000 \text{ Kg/m}^3} = 0.00151 \text{ m}^3 \quad (4)$$

The buoyancy force, F_b (Nw), acting on the wind turbine is:

$$F_b = \rho_{water} \cdot V \cdot g = 1000 \text{ Kg/m}^3 \cdot 0.00151 \text{ m}^3 \cdot 9.81 \text{ m/s}^2 = 14.81 \text{ Nw} \quad (5)$$

Based on this volume and the geometry of the platform containing the voids, the draft h (m) of the submerged part of the platform is calculated as follows:

$$V = V_{solid} - V_{moonpool} \quad (6)$$

$$0.00151 \text{ m}^3 = (h \cdot 0.31 \text{ m} \cdot 0.24 \text{ m}) - 8 \cdot (h \cdot 0.03 \text{ m} \cdot 0.03 \text{ m}) \quad (7)$$

$$h = 0.02247 \text{ m} \quad (8)$$

where V_{solid} is the volume of the platform submerged in the water and $V_{moonpool}$ is the volume of the eight open rectangular cavities.

The experimental results showed that the FOWT was submerged in water approximately 22.47 mm from the surface. This represents less than half the height of the platform, which is sufficient to demonstrate the buoyancy of the prototype design in a wave-free marine environment and provides sufficient headroom to add additional mass to the prototype, such as sensors or complementary components that can be used in more advanced tests.

It can be summarized that the experimentation has demonstrated the ability of the wind turbine scale model to adequately reflect the trends of some variables involved in the device's operation. These data validate the possible integration of the platform for a FOWT, confirming the expected aerodynamic trends. They also demonstrate that the prototype is functional and meets the initial working requirements for future, more advanced tests, such as full characterization of the model or its implementation as a floating offshore wind turbine in more realistic conditions.

Versatility and Adaptation

Beyond the electrical and structural responses, the experiments confirmed the robustness of the modular design. The wind system allowed for safe assembly and disassembly of its components. Element replacement and reconfiguration were performed quickly, with

assembly tests taking less than 5 min for tower and blade swaps. This demonstrates the feasibility of iterative prototyping.

This adaptability is crucial for scaling the prototype to different test facilities and wind turbine configurations. Therefore, this modularity facilitates device customization without the need to redesign the entire system to build a new model.

Furthermore, the printed structure can be easily scaled and modified using CAD software, allowing for deeper analysis of the parts and their behavior, as well as functional adaptation to specific geometric and mechanical requirements and experimental environments. In this way, the prototype provides a flexible research tool that can be adapted to future research on floating offshore wind turbines.

5. Conclusions

This work presents the design and initial validation of a small-scale modular offshore wind turbine prototype, developed using CAD tools and 3D printing. The modular approach proved effective for rapid assembly and reconfiguration, with the ability to replace components in less than five minutes. This highlights its suitability for laboratory testing and differentiates it from more rigid prototyping approaches. Furthermore, the elements for implementing pitch control are integrated, a lack in many wind turbine prototypes of the size proposed here.

The prototype was developed based on the geometric characteristics of a wind system, the coupling of components, and considering the physical characteristics associated with the standard layout of a general-purpose horizontal-axis wind turbine, without applying external loads or advanced control strategies.

Experimental results confirmed that the prototype reproduces the expected aerodynamic tendencies and operational behaviors of a wind device. The WT responded consistently to changes in wind speed and blade pitch angle, with maximum performance observed at a pitch of 5° and a wind speed of 5.5 m/s, reaching 335 RPM and an output of 51.06 mV, the configuration identified as having the maximum capacity of the model. Structural stability was also verified, as tower oscillations remained within a 4° range, consistent with the expected elastic behavior of PLA and validating the robustness of the modular joints. These results demonstrate that, although simplified, the prototype successfully captures the fundamental interactions between aerodynamics and structure in scaled-down wind turbines.

Therefore, the prototype is presented as a flexible experimental tool for analyzing wind turbine energy production and structural responses under controlled wind conditions. However, this study represents an initial stage of development, with successful tests focused on system integration and functional verification.

After confirming the prototype functionality, it is necessary to advance this line of development. Future work will focus on improving the model through detailed mechanical analysis of each component and the assembly to optimize the design. It is also necessary to characterize the prototype and design active control systems for optimal power generation. In this line, the electrical characterization of the wind turbine generator, encompassing voltage, current, and connection parameters, along with the power curve associated with the power coefficient (C_p), wind speed vs. power, and tip speed ratio (TSR) curves, will be determined in the next step. We plan to implement a measurement circuit with a known load to characterize the generator and this way, derive some characteristics such as the power coefficient. Similarly, when applying it to marine turbines, it is necessary to delve deeper into the hydrostatic stability and hydrodynamic responses of the floating platform, adapting the system to marine environments in wave tank experiments.

Author Contributions: Conceptualization, S.E.; Methodology, M.S.; Software, E.M.-P.; Validation, E.M.-P., S.E. and M.S.; Investigation, E.M.-P.; Writing—original draft, E.M.-P.; Writing—review & editing, S.E. and M.S.; Supervision, S.E. and M.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research was partially funded by the Spanish Ministry of Science and Innovation under the project MCI/AEI/FEDER number PID2021-123543OB-C21 and PID2024-155653OB-C21.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. International Energy Agency—IEA. *World Energy Outlook 2023*; International Energy Agency: Paris, France, 2023; Available online: <https://www.iea.org/reports/world-energy-outlook-2023> (accessed on 5 June 2025).
2. Intergovernmental Panel on Climate Change (IPCC). *Climate Change 2022: Mitigation of Climate Change. Contribution of Working Group III to the Sixth Assessment Report of the IPCC*; Shukla, P.R., Skea, J., Slade, R., Al Khourdajie, A., van Diemen, R., McCollum, D., Pathak, M., Some, S., Vyas, P., Fradera, R., et al., Eds.; Cambridge University Press: Cambridge, UK; New York, NY, USA, 2022. [CrossRef]
3. Farghali, M.; Osman, A.I.; Chen, Z.; Abdelhaleem, A.; Ihara, I.; Mohamed, I.M.A.; Yap, P.-S.; Rooney, D.W. Social, environmental, and economic consequences of integrating renewable energies in the Electricity Sector: A Review. *Environ. Chem. Lett.* **2023**, *21*, 1381–1418. [CrossRef]
4. REN21. *Renewables 2024 Global Status Report*. Paris: REN21 Secretariat. 2024. Available online: <https://www.ren21.net/gsr-2024/> (accessed on 5 June 2025).
5. IRENA. *Renewable Power Generation Costs in 2024*; International Renewable Energy Agency: Abu Dhabi, United Arab Emirates, 2025.
6. Muñoz-Palomeque, E.; Sierra-García, J.E.; Santos, M. Enhancing offshore wind turbines performance with hybrid control strategies using neural networks and conventional controllers. *J. Comput. Des. Eng.* **2025**, *12*, 80–97. [CrossRef]
7. Shafiee, M. Maintenance logistics organization for offshore wind energy: Current progress and future perspectives. *Renew. Energy* **2015**, *77*, 182–193. [CrossRef]
8. Desalegn, B.; Gebeyehu, D.; Tamrat, B.; Tadiwose, T.; Lata, A. Onshore versus offshore wind power trends and recent study practices in modeling of wind turbines' life-cycle impact assessments. *Clean. Eng. Technol.* **2023**, *17*, 100691. [CrossRef]
9. Ribeiro, J.A.; Ribeiro, B.A.; Pimenta, F.; Tavares, S.M.; Zhang, J.; Ahmed, F. Offshore wind turbine tower design and optimization: A review and AI-Driven Future Directions. *Appl. Energy* **2025**, *397*, 126294. [CrossRef]
10. Bayat, S.; Lee, Y.H.; Allison, J.T. Nested control co-design of a spar buoy horizontal-axis floating offshore wind turbine. *Ocean Eng.* **2025**, *328*, 121037. [CrossRef]
11. Jung, C.; Schindler, D. The properties of the global offshore wind turbine fleet. *Renew. Sustain. Energy Rev.* **2023**, *186*, 113667. [CrossRef]
12. Global Wind Energy Council (GWEC). *Global Wind Report 2025*; Global Wind Energy Council: Brussels, Belgium, 2025.
13. Musial, W.; Beiter, P.; Spitsen, P.; Duffy, P.; Shields, M.; Hernando, D.M.; Hammond, R.; Marquis, M.; King, J.; Sathish, S. *Offshore Wind Market Report: 2023 Edition*; U.S. Department of Energy, National Renewable Energy Laboratory (NREL): Golden, CO, USA, 2023.
14. Ma, X.; Li, M.; Li, W.; Liu, Y. Overview of offshore wind power technologies. *Sustainability* **2025**, *17*, 596. [CrossRef]
15. Barooni, M.; Ashuri, T.; Sogut, D.V.; Wood, S.; Taleghani, S.G. Floating offshore wind turbines: Current status and future prospects. *Energies* **2022**, *16*, 2. [CrossRef]
16. Boo, S.Y.; Ha, Y.-J.; Shelley, S.A.; Park, J.-Y.; Lim, C.-H.; Kim, K.-H. Concept design of a 15 MW TLP-type floating wind platform for Korean offshore installation. *J. Mar. Sci. Eng.* **2024**, *12*, 796. [CrossRef]
17. Su, X.; Wang, X.; Xu, W.; Yuan, L.; Xiong, C.; Chen, J. Offshore wind power: Progress of the edge tool, which can promote sustainable energy development. *Sustainability* **2024**, *16*, 7810. [CrossRef]
18. Muñoz-Palomeque, E.; Sierra-García, J.E.; Santos, M. Wind turbine maximum power point tracking control based on unsupervised neural networks. *J. Comput. Des. Eng.* **2023**, *10*, 108–121. [CrossRef]
19. Sierra-García, J.E.; Santos, M. Deep learning and fuzzy logic to implement a hybrid wind turbine pitch control. *Neural Comput. Appl.* **2022**, *34*, 10503–10517. [CrossRef]
20. Galán-Lavado, A.; Santos, M. Analysis of the effects of the location of passive control devices on the platform of a floating wind turbine. *Energies* **2021**, *14*, 2850. [CrossRef]

21. Aboutalebi, P.; Garrido, A.J.; Schallenberg-Rodriguez, J.; Garrido, I. Validation of vibration reduction in barge-type floating offshore wind turbines with oscillating water columns through experimental and numerical analyses. *Front. Built Environ.* **2024**, *10*, 1497123. [CrossRef]
22. Pegalajar-Jurado, A.; Bredmose, H.; Borg, M. Multi-level hydrodynamic modelling of a scaled 10MW TLP wind turbine. *Energy Procedia* **2016**, *94*, 124–132. [CrossRef]
23. Du, S.; Zhou, J.; Li, F. Active motion control of platform and rotor coupling system for floating offshore wind turbines. *Mech. Syst. Signal Process.* **2025**, *228*, 112484. [CrossRef]
24. Wang, Y.; Ahmed, A.; Azam, A.; Bing, D.; Shan, Z.; Zhang, Z.; Tariq, M.K.; Sultana, J.; Mushtaq, R.T.; Mehboob, A.; et al. Applications of Additive Manufacturing (AM) in sustainable energy generation and battle against COVID-19 pandemic: The knowledge evolution of 3d printing. *J. Manuf. Syst.* **2021**, *60*, 709–733. [CrossRef]
25. Lerche, J.; Enevoldsen, P.; Seppänen, O. Application of takt and Kanban to modular wind turbine construction. *J. Constr. Eng. Manag.* **2022**, *148*, 05021015. [CrossRef]
26. Bianchini, A.; Bangga, G.; Baring-Gould, I.; Croce, A.; Cruz, J.I.; Damiani, R.; Erfort, G.; Ferreira, C.S.; Infield, D.; Nayeri, C.N.; et al. Current status and grand challenges for small wind turbine technology. *Wind. Energy Sci.* **2022**, *7*, 2003–2037. [CrossRef]
27. Calautit, K.; Johnstone, C. State-of-the-art review of micro to small-scale wind energy harvesting technologies for building integration. *Energy Convers. Manag. X* **2023**, *20*, 100457. [CrossRef]
28. O'Neill, F.; Mehmanparast, A. A review of additive manufacturing capabilities for potential application in offshore renewable energy structures. *Forces Mech.* **2024**, *14*, 100255. [CrossRef]
29. Wulandana, R.; Werner, G.; Gardner, B. Integrated Wind Turbine Blade Design Education: Combining theory, simulation, CAD, and experimental testing. In Proceedings of the ASEE Annual Conference & Exposition, Montreal, QC, Canada, 22–25 June 2025. [CrossRef]
30. Leong, W.Y. 3D printing for Renewable Energy Technologies. *ASM Sci. J.* **2025**, *20*, 1–10. [CrossRef]
31. Nagai, B.M.; Ameku, K.; Roy, J.N. Performance of a 3 kW wind turbine generator with Variable Pitch Control System. *Appl. Energy* **2009**, *86*, 1774–1782. [CrossRef]
32. Mohammadi, E.; Fadaeinedjad, R.; Naji, H.R. Platform for design, simulation, and experimental evaluation of small wind turbines. *IET Renew. Power Gener.* **2019**, *13*, 1576–1586. [CrossRef]
33. Hamza, G.; Hammadi, M.; Barkallah, M.; Choley, J.-Y.; Riviere, A.; Louati, J.; Haddar, M. Conceptual design methodology for the preliminary study of a mechatronic system: Application to wind turbine system. *Mech. Ind.* **2017**, *18*, 413. [CrossRef]
34. Verma, N.; Baloni, B.D. Influence of Reynolds number consideration for aerodynamic characteristics of airfoil on the blade design of small horizontal axis wind turbine. *Int. J. Green Energy* **2021**, *19*, 733–746. [CrossRef]
35. Chen, J.; Kim, M.-H. Review of recent offshore wind turbine research and optimization methodologies in their design. *J. Mar. Sci. Eng.* **2021**, *10*, 28. [CrossRef]
36. Li, Y.; Molazem, A.; Kuo, H.-I.; Ahmadi, V.; Shastri, V.P. Comparative analysis of Dimensional Accuracy in PLA-based 3D printing: Effects of key printing parameters and related variables. *Polymers* **2025**, *17*, 1698. [CrossRef]
37. Mahmoud, M.; Semeraro, C.; Abdelkareem, M.A.; Olabi, A.G. Designing and prototyping the architecture of a digital twin for wind turbine. *Int. J. Thermofluids* **2024**, *22*, 100622. [CrossRef]
38. Kumar, V.; Pavan, M. Modal Analysis of NACA 4412 Airfoil Based Air-Wing Using Different Composite Materials. *Int. J. Eng. Sci. Adv. Technol.* **2018**, *18*, 26–35; ISSN 2250-3676.
39. Sakthi, S.; Muthukumaran, G.; Balasubramanian, S. Windmill Power Generation Using Mult-Generator and Single Rotor (Horizontal and Vertical Blade). *J. Energy Technol. Policy* **2012**, *2*, 11–20.
40. Xu, H.; Jia, K.; Wang, E.; Xiao, Q.; Incecik, A. Effects of moonpools and heave plates on hydrodynamic performance of a barge-type floating offshore wind turbine platform in regular waves. *Phys. Fluids* **2025**, *37*, 087166. [CrossRef]
41. Firoozi, A.A.; Hejazi, F.; Firoozi, A.A. Advancing Wind Energy Efficiency: A systematic review of aerodynamic optimization in wind turbine blade design. *Energies* **2024**, *17*, 2919. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Article

Graded Evaluation and Optimal Scheme Selection of Mine Rock Diggability Based on the Multidimensional Cloud Model

Shibin Yao ¹, Xiaoyuan Li ^{1,*}, Jian Zhou ¹ and Manoj Khandelwal ^{2,3,4,*}

¹ School of Resources and Safety Engineering, Central South University, Changsha 410083, China; shibin.yao@csu.edu.cn (S.Y.); j.zhou@csu.edu.cn (J.Z.)

² Institute of Innovation, Science and Sustainability, Federation University Australia, Ballarat 3350, Australia

³ Laboratory of Sustainable Development in Natural Resources and Environment, Institute for Advanced Study in Technology, Ton Duc Thang University, Ho Chi Minh City 700000, Vietnam

⁴ Faculty of Environment and Labour Safety, Ton Duc Thang University, Ho Chi Minh City 700000, Vietnam

* Correspondence: 225502046@csu.edu.cn (X.L.); m.khandelwal@federation.edu.au or khandelwamanoj@tdtu.edu.vn (M.K.)

Abstract: With the advancement of mining technologies, the evaluation of rock diggability has become a critical research topic for ensuring both safety and efficiency in mining operations. This study establishes a comprehensive evaluation system for mine rock diggability and proposes corresponding grading criteria. For the determination of indicator weights, a combination of subjective and objective methods is employed, integrating expert knowledge and data characteristics to identify optimal weights, thereby providing a reliable basis for comprehensive evaluation. The single-indicator cloud model effectively mitigates the difficulties associated with defining transitional values between adjacent intervals. The multidimensional cloud model, by considering the interactions among indicators, enables the optimization of indicator interactions and enhances the interpretability of diggability grades. Comparison with the Diggability Index (DI) method shows a high consistency between the two approaches ($R^2 = 0.991$). The absolute accuracy of diggability levels reaches 74%, while the accuracy based on cloud model fuzzy evaluation reaches 100%, demonstrating the effectiveness of the cloud model in handling transitional intervals and capturing uncertainty. This study provides a novel methodology and theoretical foundation for the scientific evaluation of mine rock diggability, offering practical guidance for reasonable grading, optimization of mining parameters, and interpretation of diggability levels in engineering practice.

Keywords: multidimensional cloud model; rock diggability; comprehensive evaluation; fuzzy evaluation

1. Introduction

In mining operations and geotechnical engineering, the selection of appropriate excavation methods is of critical importance; therefore, research on rock diggability holds significant theoretical and engineering value. From a construction design perspective, diggability grading provides guidance for the selection of excavation methods and the optimization of process workflows, helping to avoid schedule delays and cost overruns caused by unreasonable designs. From an equipment allocation standpoint, different diggability grades correspond to varying requirements for crushing, loading, and transportation equipment. A scientifically sound diggability assessment facilitates rational equipment selection and efficient utilization. In terms of construction safety, diggability evaluation enables the identification of potential adverse geological conditions prior to excavation,

thereby reducing operational risks. Accordingly, the study of rock diggability is a crucial measure to ensure the overall economic efficiency, safety, and design rationality of excavation projects [1].

Depending on the application context, terms such as “excavatability,” “drillability,” and “rippability” have been proposed to evaluate excavation methods and optimize equipment selection [2–6]. From the perspective of rock characteristics, diggability is influenced by multiple factors, including lithology, weathering degree, joint development, bedding structure, porosity, and water content [7,8]. The heterogeneity and discontinuity of rock masses lead to significant variations in rock behavior under different mining conditions, thereby increasing the complexity of diggability prediction [9–11].

In the study of diggability grading, various traditional empirical methods based on rock mechanical parameters and rock mass structural characteristics have been employed to evaluate diggability and provide intuitive engineering guidance. Scoble and Muftuoglu [6] proposed a diggability index based on weathering degree, uniaxial compressive strength, joint spacing, and bedding spacing, classifying excavation difficulty into seven categories and recommending corresponding equipment types. The effectiveness of this approach was validated using ten sandstone bench cases from open-pit coal mines. Rotimi et al. and Ismail et al. suggested that seismic wave measurements could also be used to characterize rock tearability (Rippability) [12,13]. Given the numerous factors influencing rock diggability, which cannot be comprehensively covered, and the incomplete understanding of their mechanisms, incorporating the fuzziness of unknown factors has emerged as a promising approach for diggability grading. Iphar and Goktan [14] addressed the issues of sharp boundaries between adjacent categories and subjective uncertainty of boundary data in traditional diggability rating methods by integrating fuzzy set theory with the Mamdani algorithm, thereby overcoming the limitations of conventional grading methods. Zhou et al. [15] proposed a hybrid model based on uncertainty measures and information entropy, employing multiple membership functions to compute single-indicator values, determining weights through information entropy, and introducing the reliability coefficient (Rec) to adjust the evaluation results. This approach provides a practical method for assessing rock diggability, suitable for guiding equipment selection and cost optimization in excavation projects.

With the increase in computational power, accurate simulation and in-depth analysis of geotechnical conditions have become feasible. By reflecting actual geological conditions based on geological surveys and laboratory tests, combined with numerical simulations of blasting and mechanical rock-breaking processes, researchers can investigate geological diggability and optimize excavation procedures accordingly [16–19]. Meanwhile, artificial intelligence methods have been widely applied in the study of rock diggability. Artificial neural networks (ANNs) have been employed to predict the cuttability of rocks [20]. Huang and Zhou [21] utilized an expanded database to predict rock diggability by optimizing multilayer perceptrons (GA-MLP) with genetic algorithms and standards based on gene expression programming (GEP).

Despite the progress achieved using traditional empirical methods, mechanical parameter-based evaluations, and artificial intelligence in the study of rock diggability, several limitations remain [22–24]. First, empirical and mechanical parameter-based approaches largely rely on manually defined evaluation rules and indicator weights, which cannot fully capture the heterogeneity and complex structure of rock masses, resulting in evaluations that are highly subjective. Second, although numerical simulations and AI methods can model or predict rock failure processes, they typically require large amounts of high-quality data, and their outputs are often singular, providing insufficient handling

of data uncertainty and fuzziness. Consequently, their reliability may be compromised under unclear boundary conditions or ambiguous indicators [25–27].

The cloud model, as a mathematical tool capable of characterizing both randomness and fuzziness, organically integrates qualitative language with quantitative data. It has been widely applied in geotechnical and related engineering evaluations to address issues such as ambiguous indicators, unclear boundaries, and complex data distributions, thereby enhancing the scientific rigor and reliability of evaluation results [28–30].

In summary, this study establishes a systematic evaluation framework for rock diggability. Following the principle of “comprehensive accessibility,” weathering degree (WD), uniaxial compressive strength (UCS), joint spacing (JS), and bedding spacing (BS) were selected as core evaluation indicators. A five-level grading standard was constructed and mapped to established diggability evaluation results to verify its rationality. To determine indicator weights, a combination of subjective and objective approaches was employed. Subjective weights were obtained by consulting five experts who constructed judgment matrices, while objective weights were calculated and optimized using the Entropy Weight Method (EW), Correlation-based Criteria Importance (CRITIC), Coefficient of Variation (CV), Principal Component Analysis (PCA), and Factor Analysis (FA). The final comprehensive indicator weights were subsequently derived.

2. Data Sources and Evaluation Framework

2.1. Data Description

This study collected and compiled rock diggability-related data from the literature in the United Kingdom and Türkiye (Figure 1), comprising a total of 56 open-pit coal mine records—46 from Türkiye and 10 from the United Kingdom [6,14]. Figure 1 illustrates the geographical distribution of these sites, highlighting the spatial coverage and representativeness of the dataset. The dataset consists of four core indicators (WD, UCS, JS, and BS) sourced from field geological surveys and laboratory rock mechanics tests. These data provide a comprehensive representation of the physical and mechanical characteristics of rock masses under different geological conditions and are highly accessible for engineering applications. The qualitative indicator WD was defined according to the International Society for Rock Mechanics and Rock Engineering (ISRM, 1978) standard [31]. In addition, diggability indices (DI) calculated based on existing evaluation methods were collected as an important reference for validating the reliability of subsequent assessment results. The calculation of DI used in this study follows the method described in [6].

2.2. Data Analysis

DI was treated as the output indicator, while the other four indicators (WD, UCS, JS, and BS) were used as input indicators for correlation analysis (Figure 2). As illustrated in Figure 2, since DI serves as the output, its correlation with each input indicator reflects the explanatory power of the indicator system on the evaluation results. The analysis shows that JS and BS exhibit the highest correlation with DI ($r > 0.90$), indicating that geometric parameters have a significant impact on the output. UCS also demonstrates a strong correlation with DI ($r = 0.77$), suggesting that rock strength is closely related to overall performance. In contrast, WD shows a lower correlation with DI ($r = 0.59$), but still provides complementary information for the assessment.




Figure 1. Geographic distribution of data sources.




Figure 2. Correlation analysis of indicators.

On the other hand, strong correlations exist among the input indicators themselves, such as between JS and BS ($r = 0.87$), which may indicate potential redundancy and the need to avoid duplicate measurements. However, existing studies have shown that joints

influence rock mass stability by reducing rock strength and altering permeability, whereas bedding affects rock anisotropy and weathering patterns. Therefore, both JS and BS should be considered simultaneously to comprehensively account for the influence of joint and bedding distributions.

2.3. Rock Diggability Evaluation Framework

Building upon the original evaluation system, this study developed an adjusted framework for rock diggability assessment, which mainly comprises two components: indicator fuzzification and diggability grading. First, using interval scoring, a fuzzy conversion table for the indicators was established (Table 1) based on the four key diggability control factors (WD, UCS, JS, and BS). This table maps qualitative descriptions of each indicator to corresponding quantitative ranges, enabling the quantitative expression of qualitative features under different conditions as well as the standardized representation of all rock mass characteristics. Subsequently, by relating the fuzzy scoring results to the comprehensive indicators, rock diggability levels were further classified. Typical excavation methods corresponding to each level were incorporated to form a complete evaluation framework. This system not only reflects the relative ease or difficulty of rock excavation but also provides a scientific basis for the rational selection of excavation methods.

Table 1. Fuzzy conversion of diggability indicators.

Indicators	Transformed Score				
	0–20	20–40	40–60	60–80	80–100
Weathering degree	Completely 10	Highly 30	Moderately 50	Slightly 70	Unweathered 90
Uniaxial compressive strength (MPa)	0–20	20–40	40–60	60–100	>100
Joint spacing (m)	0–0.3	0.3–0.6	0.6–1.5	1.5–2	>2
Bedding spacing (m)	0–0.1	0.1–0.3	0.3–0.6	0.6–1.5	>1.5

Building upon the existing evaluation system [6], this study further optimized and adjusted the framework to develop a rock diggability assessment system, which mainly comprises two components: indicator fuzzification and diggability grading. In the indicator fuzzification stage, interval scoring was applied to the four key diggability control factors (WD, UCS, JS, and BS) to construct the fuzzy conversion table for the indicators (Table 1). This process maps qualitative descriptions of each indicator to corresponding quantitative ranges, enabling the quantitative expression of qualitative features under different conditions and ensuring standardized representation of rock mass characteristics on a unified scale.

In the diggability grading stage, a grading standard was established based on the comprehensive scores obtained from the fuzzy conversion and their correspondence with the diggability index (Table 2). On this basis, the original seven-level grading system from previous studies was merged and optimized to form a five-level system, ranging from Level I to Level V, corresponding to excavation difficulty from “Very Easy” to “Marginal Without Blasting.” This grading system not only ensures effective linkage between the fuzzy conversion results and diggability classification but also clearly reflects the diggability characteristics of rock masses under engineering conditions. Furthermore, each diggability level is associated with typical excavation methods, providing a practical reference for method selection and construction planning.


Table 2. Diggability levels and corresponding excavation methods.

Diggability Level	I	II	III	IV	V
Fuzzy score	0–30	30–40	40–55	55–70	70–100
Index	0–40	40–50	50–70	70–100	>100
Ease of digging	Very Easy	Easy	Difficult	Very Difficult	Marginal Without Blasting
Excavation method	1. Ripping 2. Dragline 3. Shovel digging	1. Ripping 2. Dragline 3. Shovel digging	1. Ripping 2. Shovel digging	Shovel digging	Shovel digging

3. Methodology

3.1. Multidimensional Cloud Model

The cloud model is a cognitive tool that characterizes the uncertain mapping between qualitative concepts and quantitative data. It consists of expectation (Ex), entropy (En), and hyper-entropy (He), where Ex denotes the central tendency of the cloud drops, En characterizes uncertainty and fuzziness, and He describes the volatility, reflecting the stability of uncertainty, as illustrated in Figure 3 [32–35]. When multiple indicators are considered simultaneously, the cloud model can be extended to a d-dimensional form (multidimensional cloud model) [36–38]. As shown in Figure 3, given the expectation, entropy, and hyper-entropy as Ex_j , En_j , He_j , respectively, the generation process of cloud drops in a d-dimensional cloud model is as follows:

**Figure 3.** Introduction to the principles of the multidimensional cloud model.

(1) Generate perturbed entropy

$$En'_j \sim N(En_j, He_j^2), j = 1, 2, \dots, d \quad (1)$$

where En'_j denotes the perturbed entropy.

- (2) Determine the positions of cloud drops

$$x_j \sim N(Ex_j, En'_j), j = 1, 2, \dots, d \quad (2)$$

where Ex_j denotes the expectation of the j -th indicator.

- (3) Membership calculation

$$\mu(x_1, x_2, \dots, x_d) = e^{-\sum_{j=1}^d \frac{(x_j - Ex_j)^2}{2(En'_j)^2}} \quad (3)$$

where $\mu \in (0, 1]$ is the membership degree. When $d = 1$, it corresponds to the cloud drop generation process of a one-dimensional (single-indicator) cloud model; when $d \geq 1$, it corresponds to the cloud drop generation process of a multidimensional (multi-indicator) cloud model (Figure 3).

3.2. Subjective Weights (Expert Evaluation)

The subjective weights of the indicators are determined by experts through pairwise comparisons of the relative importance between different indicators, as outlined below:

- (1) Expert evaluation

As shown in Figure 4, invited experts evaluate the relative importance of indicators A and B to be weighted, using a 1–9 scale, where 1 denotes equal importance and 2–9 represent increasing levels of importance. By conducting pairwise comparisons of all indicators, the evaluation of the importance of all indicators is completed [39,40].




Figure 4. Evaluation of the interaction matrix and its importance relationships.

- (2) Construction of the judgment matrix

Referring to the principle shown in Figure 4, the experts' evaluation results are transformed into a normalized judgment matrix: $M = [m_{pq}]_{n \times n}$. In this matrix, m_{pq} denotes the relative importance of indicator p compared to indicator q as judged by the experts.

(3) Weight calculation

$$w_p = \frac{(\prod_{q=1}^n m_{pq})^{1/n}}{\sum_{r=1}^n (\prod_{q=1}^n m_{rq})^{1/n}} \quad (4)$$

where $p = 1, 2, \dots, n$; w_p denotes the weight of the p -th indicator.

(4) Consistency test

To ensure the consistency of expert evaluations, the consistency ratio (CR) can be used as an indicator [41].

$$CI = \frac{\lambda_{\max} - n}{n - 1} \quad (5)$$

$$CR = \frac{CI}{RI} \quad (6)$$

where n is the number of indicators, and RI denotes the random consistency index. When $CR < 0.1$, the judgment matrix is considered consistent.

3.3. Objective Weight

Compared with subjective weighting methods that rely on expert judgment, objective weighting methods determine weights based on the intrinsic data characteristics of the indicators, reducing the influence of subjective factors and offering strong objectivity and reproducibility. The objective weighting methods applied in the study area include the entropy weight method, coefficient of variation method, CRITIC method, principal component analysis, and factor analysis, with the specific procedure as follows.

3.3.1. Entropy Weight Method (EW)

Based on the concept of information entropy, the greater the variability of an indicator, the more information it provides, and accordingly, the higher its assigned weight. Let there be m sets of data and n indicators; the steps for calculating weights using the entropy weight method are as follows [42–44]:

(1) Standardization

$$a = x_{ij} = \frac{a_{ij} - \min_i a_{ij}}{\max_i a_{ij} - \min_i a_{ij}}, j = 1, \dots, n \quad (7)$$

where a_{ij} denotes the original value of the j -th indicator in the i -th dataset, and x_{ij} denotes the corresponding value after standardization.

(2) Calculation of proportions

$$a = p_{ij} = \frac{x_{ij}}{\sum_{i=1}^m x_{ij}}, \sum_{i=1}^m p_{ij} = 1 \quad (8)$$

where p_{ij} denotes the proportion of the j -th indicator in the i -th dataset.

(3) Calculation of entropy

$$a = e_j = -k \sum_{i=1}^m p_{ij} \ln(p_{ij}) \quad (9)$$

where $k = \frac{1}{\ln(m)}$; e_j denotes the information entropy of the j -th indicator.

(4) Calculation of the difference coefficient

$$d_j = 1 - e_j \quad (10)$$

where d_j denotes the information utility value of the j -th indicator.

(5) Weight calculation

$$w_j = \frac{d_j}{\sum_{j=1}^n d_j} \quad (11)$$

where $\sum_{j=1}^n w_j = 1$; w_j denotes the final weight of the j -th indicator.

3.3.2. Coefficient of Variation Method (CV)

The coefficient of variation method allocates weights based on the degree of variability of the indicators [45].

(1) Calculation of standard deviation

$$\sigma_j = \sqrt{\frac{1}{m} \sum_{i=1}^m (x_{ij} - \bar{x}_j)^2} \quad (12)$$

where σ_j denotes the standard deviation (dispersion) of the j -th indicator, and $\bar{x}_j$ denotes the mean value of the j -th indicator.

(2) Calculation of the coefficient of variation

$$V_j = \frac{\sigma_j}{\bar{x}_j} \quad (13)$$

where V_j denotes the coefficient of variation of the j -th indicator.

(3) Weight calculation

$$w_j = \frac{V_j}{\sum_{j=1}^n V_j} \quad (14)$$

3.3.3. CRITIC Method

The CRITIC method determines weights by comprehensively considering both the variability of indicators and their conflicts with other indicators [46].

(1) Calculation of correlation

$$\rho_{jk} = \frac{\sum_{i=1}^m (x_{ij} - \bar{x}_j)(x_{ik} - \bar{x}_k)}{\sqrt{\sum_{i=1}^m (x_{ij} - \bar{x}_j)^2 \cdot \sum_{i=1}^m (x_{ik} - \bar{x}_k)^2}} \quad (15)$$

where ρ_{jk} denotes the correlation coefficient between indicator j and indicator k (degree of conflict).

(2) Calculation of information quantity

$$C_j = \sigma_j \cdot \sum_{k=1}^n (1 - \rho_{jk}) \quad (16)$$

where C_j denotes the comprehensive information quantity of indicator j .

(3) Weight calculation

$$w_j = \frac{C_j}{\sum_{j=1}^n C_j} \quad (17)$$

3.3.4. Principal Component Analysis (PCA)

By extracting principal components through dimensionality reduction, indicator weights are transformed into principal component contribution functions [47].

- (1) Construction of the covariance matrix

$$S = \frac{1}{m} X^T X \quad (18)$$

where X denotes the standardized indicator matrix, and S denotes the covariance matrix.

- (2) Solve for eigenvalues and eigenvectors

$$S a_k = \lambda_k a_k \quad (19)$$

- (3) Calculation of contribution rate and cumulative contribution rate

$$\eta_k = \frac{\lambda_k}{\sum_{k=1}^n \lambda_k} \quad (20)$$

$$\eta = \sum_{k=1}^p \eta_k \quad (21)$$

where η_k denotes the variance contribution rate of the k -th principal component, and p is the number of selected principal components.

- (4) Calculation of indicator weights

$$w_j = \frac{\sum_{k=1}^p \eta_k |a_{jk}|}{\sum_{j=1}^n C_j \sum_{k=1}^p \eta_k |a_{jk}|} \quad (22)$$

where a_{jk} denotes the coefficient of indicator j in the k -th principal component vector.

3.3.5. Factor Analysis Method (FA)

Factor analysis determines indicator weights by extracting common factors to reveal shared information among indicators [48].

- (1) Establishment of the factor model

$$x_j = \mu_j + \sum_{k=1}^p l_{jk} f_k + \varepsilon_j \quad (23)$$

where x_j denotes the j -th indicator; μ_j denotes the constant term, which is 0 since the data are standardized; l_{jk} denotes the factor loading coefficient; f_k denotes the common factor; ε_j denotes the unique factor.

- (2) Extraction of factors and loading matrix

$$L = (l_{jk})_{n \times p} \quad (24)$$

- (3) Calculation of weights

$$w_j = \frac{\sum_{k=1}^p \eta_k |l_{jk}|}{\sum_{j=1}^n \sum_{k=1}^p \eta_k |l_{jk}|} \quad (25)$$

3.4. Comprehensive Weights (Combination of Subjective and Objective Weights)

To comprehensively consider both expert judgment and data characteristics, the final indicator weights are calculated using a linear combination of subjective weights and the optimized objective weights. Let W_j^S and W_j^O denote the subjective and optimized objective weights of the j -th indicator, respectively. The comprehensive weight W_j^C is calculated as:

$$W_j^C = \alpha \cdot W_j^S + (1 - \alpha) \cdot W_j^O, j = 1, 2, \dots, n \quad (26)$$

where $\alpha \in [0, 1]$ is the coefficient reflecting the relative importance of subjective versus objective weighting. In this study, $\alpha = 0.5$ is used, indicating equal consideration of expert judgment and data characteristics.

4. Rock Excavatability Modeling and Evaluation

In rock diggability modeling, determining the weights of indicators is a crucial step in constructing a comprehensive evaluation model. The weights identified here are divided into objective and subjective weights to ensure that the model both accurately reflects the characteristics of the existing data and incorporates expert macro-level judgment.

4.1. Comprehensive Weights

4.1.1. Subjective Wights (Expert Evaluation)

First, the four identified indicators (WD, UCS, JS, BS) and their relevant research backgrounds were provided to five invited experts from different professional fields with extensive mining experience, who independently evaluated the importance of each indicator. Specifically, Experts 1 and 2 specialize in mining equipment development, Experts 3 and 4 specialize in mining and extraction research, and Expert 5 specializes in mining systems engineering research, ensuring that different professional perspectives on the diggability indicators are adequately represented. The experts' pairwise comparison matrices are constructed using Saaty's 1–9 scale method, resulting in an independent judgment matrix for each expert:

$$\text{Expert 1 : } \begin{bmatrix} 1 & 1 & 1/3 & 1/5 \\ 1 & 1 & 1/3 & 1/3 \\ 3 & 3 & 1 & 2 \\ 5 & 3 & 1/2 & 1 \end{bmatrix} \quad (27)$$

$$\text{Expert 2 : } \begin{bmatrix} 1 & 1/2 & 1/3 & 1/5 \\ 2 & 1 & 1/3 & 1/3 \\ 3 & 3 & 1 & 1 \\ 5 & 3 & 1 & 1 \end{bmatrix} \quad (28)$$

$$\text{Expert 3 : } \begin{bmatrix} 1 & 2 & 1 & 1/4 \\ 1/2 & 1 & 1/2 & 1/2 \\ 1 & 2 & 1 & 1/2 \\ 4 & 2 & 2 & 1 \end{bmatrix} \quad (29)$$

$$\text{Expert 4 : } \begin{bmatrix} 1 & 1 & 2 & 1 \\ 1 & 1 & 1 & 1/2 \\ 1/2 & 1 & 1 & 1/2 \\ 1 & 2 & 2 & 1 \end{bmatrix} \quad (30)$$

$$\text{Expert 5 } \begin{bmatrix} 1 & 2 & 1 & 1/7 \\ 1/2 & 1 & 1/2 & 1/7 \\ 1 & 2 & 1 & 1/4 \\ 7 & 7 & 4 & 1 \end{bmatrix} \quad (31)$$

As shown in Table 3, based on the independent judgment matrices of the five experts (26–30), the indicator weights for each expert are calculated using Equation (4) (Figure 5). Subsequently, Equations (5) and (6) are used to perform a consistency check on each expert's judgment results to ensure the rationality of the evaluation matrices and weight assignments (Table 3).

Table 3. Summary of subjective weights.

Indicators		Expert 1	Expert 2	Expert 3	Expert 4	Expert 5
Weathering degree		0.135	0.120	0.202	0.279	0.158
Uniaxial compressive strength		0.151	0.166	0.163	0.216	0.111
Joint spacing		0.398	0.344	0.237	0.188	0.180
Bedding spacing		0.316	0.370	0.398	0.317	0.552
Consistency test	CI	0.046	0.016	0.061	0.02	0.017
	RI	0.900	0.900	0.900	0.900	0.900
	CR	0.051	0.018	0.068	0.022	0.019




Figure 5. Summary of subjective weights.

As shown in Figure 5, the five experts exhibit certain tendencies in their allocation of indicator weights: overall, JS and BS are generally regarded as the most important indicators with relatively higher weights, reflecting that the experts focus more on the influence of rock structure features on diggability. Conversely, WD and UCS have relatively lower weights, indicating that experts consider their influence comparatively minor. Specifically, Expert 1 and Expert 2 emphasize JS more, whereas Expert 4 gives greater weight to WD; Expert 5 assigns significantly higher weight to BS, indicating a stronger focus on bedding

structures. Overall, experts tend to concentrate the core weights of diggability evaluation on the geometric structural features of the rock (BS, JS), while the rock mechanical properties (UCS) and weathering degree (WD) are of secondary importance. It is considered that the invited experts exhibit no obvious professional bias, and each contributes reasonably to the allocation of weights.

The final subjective weights were obtained by averaging the independent indicator weights provided by the five experts (Table 3), resulting in WD (0.18), UCS (0.16), JS (0.27), and BS (0.39). Although the expert evaluation method is straightforward, it is combined with objective weighting and the multi-dimensional cloud model to ensure a balanced, reliable, and scientifically grounded assessment of rock diggability.

4.1.2. Objective Weights

Based on objective data characteristics, this study calculates indicator weights using Equations (7)–(11) for EWM, Equations (12)–(14) for CV, Equations (15)–(17) for CRITIC, Equations (18)–(22) for PCA, and Equations (23)–(25) for FA. The indicator weights obtained by each method and their corresponding classification accuracies are presented in Table 4 and visualized in Figure 6. The indicator weights corresponding to the method with the highest classification accuracy are selected as the optimal weights under the objective weighting approach.

Table 4. Summary of objective weights.

Indicators	EWM	CRITIC	CV	PCA	FA
Weathering degree	0.14	0.36	0.13	0.25	0.29
Uniaxial compressive strength	0.20	0.24	0.19	0.25	0.30
Joint spacing	0.41	0.24	0.39	0.23	0.15
Bedding spacing	0.25	0.16	0.29	0.27	0.26
Classification accuracy (%)	67.90	60.70	64.30	60.70	60.70




Figure 6. Summary of objective weights.

Accordingly, the optimal weights under the objective methods are: WD (0.14); UCS (0.20); JS (0.41); BS (0.25). It can be seen that, based on data characteristics, the indicators' importance ranking from high to low is JS > BS > UCS > WD. Compared with subjective

weights, there is some consistency in the ranking of core indicators, but differences exist in the emphasis on secondary indicators (UCS and WD), providing a multi-perspective reference for comprehensive evaluation.

4.1.3. Optimal Weights

Based on the integration of subjective and objective weights, this study determined the optimal weight of each indicator using the linear combination method (Equation (26)). The resulting optimal weights are WD (0.16), UCS (0.18), JS (0.34), and BS (0.32), ensuring that the data characteristics reflect objective patterns while also incorporating expert judgment, thus providing a reliable basis for subsequent diggability assessment.

4.2. Cloud Model

4.2.1. Single-Indicator Benchmark Cloud Model

Based on the fuzzy scores corresponding to different diggability levels in Table 2, the cloud model parameters for the five levels are obtained, as shown in Table 5. Further, using these cloud model parameters and Equations (1)–(3) under the condition $d = 1$, cloud droplets are generated iteratively to ultimately construct the single-indicator benchmark cloud model (Figure 7). This cloud model serves as a reference for comparison in single-indicator evaluation of specific cases, and is used for assessing the diggability of actual cases.

Table 5. Single-indicator benchmark cloud model parameters.

Level	I	II	III	IV	V
<i>Ex</i>	15	35	47.5	62.5	85
<i>En</i>	5	1.67	2.5	2.5	5
<i>He</i>	0.5	0.167	0.25	0.25	0.5




Figure 7. Single-indicator benchmark cloud model.

4.2.2. Multi-Dimensional Benchmark Cloud Model

In the multi-indicator case, the parameters of the single-indicator benchmark cloud model are still used as the basis. Cloud droplets are iteratively generated using Equations (1)–(3) under the condition $d = 2$, following the generation principles of the multi-dimensional cloud model (Figure 3), ultimately producing the multi-dimensional

benchmark cloud model (Figure 8). This cloud model serves as a reference for multi-indicator interaction evaluation.




Figure 8. Multi-dimensional benchmark cloud model.

4.3. Diggability Evaluation Based on Multi-Dimensional Cloud Model

4.3.1. Diggability Evaluation Results Based on Multi-Dimensional Cloud Model

The 56 collected sets of qualitative and quantitative data were first defuzzified and normalized according to Table 1 to obtain diggability scores. Using the optimally determined weights combining subjective and objective methods (WD: 0.16, UCS: 0.18, JS: 0.34, BS: 0.32), weighted calculations were applied to derive the multi-dimensional cloud model evaluation score for each case. Furthermore, the diggability classification for each case was obtained by dividing levels according to the interval of the symmetry axis of the corresponding generated cloud model (Table 6).

Table 6. Diggability classification evaluation results.

Site	Diggability Index Rating Method		Multidimensional Cloud Model	
	Rating	Ease of Digging Classification	Rating	Ease of Digging Classification
1	90	Very Difficult	62.49	Very Difficult
2	85	Very Difficult	63.48	Very Difficult
3	95	Very Difficult	64.57	Very Difficult
4	90	Very Difficult	63.63	Very Difficult
5	90	Very Difficult	71.34	Marginal Without Blasting
6	70	Very Difficult	50.01	Difficult
7	85	Very Difficult	59.98	Very Difficult
8	60	Difficult	43.33	Difficult
9	90	Very Difficult	68.63	Very Difficult
10	50	Difficult	37.59	Easy
11	110	Marginal Without Blasting	65.67	Very Difficult
12	80	Very Difficult	60.21	Very Difficult
13	70	Very Difficult	48.87	Difficult
14	85	Very Difficult	55.24	Very Difficult
15	105	Marginal Without Blasting	70.08	Marginal Without Blasting
16	30	Very Easy	34.57	Easy
17	80	Very Difficult	54.82	Difficult
18	85	Very Difficult	59.42	Very Difficult
19	80	Very Difficult	60.54	Very Difficult
20	125	Marginal Without Blasting	87.07	Marginal Without Blasting
21	45	Easy	45.04	Difficult

Table 6. Cont.

Site	Diggability Index Rating Method		Multidimensional Cloud Model	
	Rating	Ease of Digging Classification	Rating	Ease of Digging Classification
22	45	Easy	39.48	Easy
23	45	Easy	40.73	Difficult
24	70	Very Difficult	50.66	Difficult
25	70	Very Difficult	52.32	Difficult
26	60	Difficult	47.62	Difficult
27	80	Very Difficult	52.01	Difficult
28	100	Marginal Without Blasting	70.27	Marginal Without Blasting
29	110	Marginal Without Blasting	72.31	Marginal Without Blasting
30	75	Very Difficult	58.93	Very Difficult
31	60	Difficult	46.94	Difficult
32	90	Very Difficult	63.87	Very Difficult
33	75	Very Difficult	51.31	Difficult
34	80	Very Difficult	54.23	Difficult
35	120	Marginal Without Blasting	80.79	Marginal Without Blasting
36	60	Difficult	48.36	Difficult
37	40	Easy	38.25	Easy
38	35	Very Easy	29.18	Very Easy
39	45	Easy	35.49	Easy
40	65	Difficult	41.71	Difficult
41	35	Very Easy	31.65	Easy
42	40	Easy	32.65	Easy
43	55	Difficult	46.62	Difficult
44	35	Very Easy	29.11	Very Easy
45	30	Very Easy	28.35	Very Easy
46	25	Very Easy	27.87	Very Easy
47	110	Marginal Without Blasting	77.74	Marginal Without Blasting
48	105	Marginal Without Blasting	64.29	Very Difficult
49	115	Marginal Without Blasting	78.16	Marginal Without Blasting
50	120	Marginal Without Blasting	77.53	Marginal Without Blasting
51	55	Difficult	40.21	Difficult
52	110	Marginal Without Blasting	81.70	Marginal Without Blasting
53	105	Marginal Without Blasting	70.69	Marginal Without Blasting
54	85	Very Difficult	62.63	Very Difficult
55	95	Very Difficult	67.97	Very Difficult
56	125	Marginal Without Blasting	86.43	Marginal Without Blasting

To verify the rationality and effectiveness of the proposed method, the evaluation results of the multi-dimensional cloud model were compared with the diggability grades and ratings obtained from the validated and reliable Diggability Index Rating Method. As shown in Table 6, the evaluation results of the two methods are consistent in most cases, indicating that the multi-dimensional cloud model can effectively reflect the diggability characteristics of rock masses. Furthermore, the diggability scores obtained by both methods were fitted against the Diggability Index, as illustrated in Figure 9, where Figure 9a presents the unrestricted linear fitting and Figure 9b shows the fitting constrained to pass through the origin.

The coefficient of determination R^2 is calculated as:

$$R^2 = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y}_i)^2} \quad (32)$$

where y_i is the observed value, $\hat{y}_i$ is the corresponding fitted value, $\bar{y}_i$ is the mean of the observed values, and n is the number of data points. Here, the numerator represents the residual sum of squares (RSS) and the denominator represents the total sum of squares (TSS). An R^2 value closer to 1 indicates a better fit.




Figure 9. Validation of the reliability of Multi-dimensional cloud model evaluation: (a) Unconstrained fitting, (b) Zero-intercept Fitting.

The zero-intercept fitting exhibited higher quantitative performance ($R^2 = 0.991$), demonstrating that this origin-constrained approach provides a more reasonable and accurate evaluation relationship. Accordingly, a refined conversion formula between the diggability scores obtained by this method and the Diggability Index was established. To clearly present this correlation, the zero-intercept fitting is also expressed as a separate equation:

$$DI = 0.709 \times S \quad (33)$$

where DI is the Diggability Index, S is the diggability score obtained by the proposed method, and 0.709 is the fitting coefficient.

4.3.2. Diggability Evaluation Analysis Based on Single-Indicator Cloud Model

For cases where the multi-dimensional cloud model grading results differ from the diggability index grading (site5, site17, site23, and site41), single-indicator cloud models were plotted for interpretive analysis (Figure 10).

According to Table 6, Sites 5, 17, 23, and 41 were classified as misclassified under the absolute Diggability Index (DI) grading. However, the cloud model analysis in Figure 10 indicates that fuzzy evaluation of diggability is not a simple “correct/incorrect” issue. Taking Site 5 as an example, it is classified as “Marginal Without Blasting” in Table 6, yet its probability distribution in Figure 10 clearly also includes the “Very Difficult” grade, which corresponds to the actual DI classification. Therefore, while the traditional absolute classification appears incorrect, the cloud model-based fuzzy evaluation accurately reflects the site’s diggability and can be considered basically correct. The same applies to Sites 17, 23, and 41.

A confusion matrix heatmap of diggability levels was plotted for all evaluation cases in Table 6 (Figure 11). For the absolute accuracy of individual levels, level III (Difficult) achieved the highest accuracy at 0.88, while level IV (Very Difficult) had the lowest accuracy at 0.61, with errors primarily distributed in adjacent intervals: 0.35 in level III (Difficult) and 0.04 in level III (Marginal Without Blasting). The accuracy of other levels falls between these two values.

It was observed that all classification errors occurred only in level intervals adjacent to the actual levels, with no severe deviations spanning multiple levels. Accordingly, the cloud model method can reasonably reflect the uncertainty and fuzziness of samples during evaluation, offering stronger rationality and interpretability compared with traditional interval division methods.




Figure 10. Diggability grading cloud models for site 5, 17, 23, and 41.




Figure 11. Confusion matrix heatmap of diggability grades.

Regarding overall evaluation accuracy, when measured by absolute level accuracy (Table 6), the overall accuracy was 71.4%. However, when assessed using the cloud model-based fuzzy evaluation (distribution shown in Figure 10), the comprehensive accuracy reached 100%, meaning that all cases classified as failures in absolute evaluation still had cloud model distributions that included the correct diggability levels (Table 7). This further highlights the potential value of adjacent-level classification errors observed in Figure 11 and demonstrates the superiority of the cloud model in handling transitional issues between adjacent levels and providing richer evaluation information.

Table 7. Grading error analysis based on the cloud model.

Site	Diggability Index Rating Method		Multidimensional Cloud Model Distribution	
	Rating	Ease of Digging Classification	Rating	Ease of Digging Classification
5	90	Very Difficult	71.34	Marginal Without Blasting
6	70	Very Difficult	50.01	*—Marginal Without Blasting
10	50	Difficult	37.59	Difficult *—Very Difficult
11	110	Marginal Without Blasting	65.67	Easy *—Difficult
13	70	Very Difficult	48.87	Very Difficult *—Marginal Without Blasting
16	30	Very Easy	34.57	Difficult *—Very Difficult
17	80	Very Difficult	54.82	Very Easy—Easy *
21	45	Easy	45.04	Difficult *—Very Difficult
23	45	Easy	40.73	Easy—Difficult *
24	70	Very Difficult	50.66	Easy—Difficult *
25	70	Very Difficult	52.32	Difficult *—Very Difficult
27	80	Very Difficult	52.01	Difficult *—Very Difficult
33	75	Very Difficult	51.31	Difficult *—Very Difficult
34	80	Very Difficult	54.23	Difficult *—Very Difficult
41	35	Very Easy	31.65	Very Easy —Easy *
48	105	Marginal Without Blasting	64.29	Very Difficult *—Marginal Without Blasting

* In the Multidimensional Cloud Model Distribution column, expressions such as Difficult—Very Difficult indicate the range of the diggability level covered by the cloud model distribution. The value marked with an asterisk (e.g., Difficult *) represents the grade corresponding to the absolute classification of the sample.

4.3.3. Analysis of Diggability Evaluation Based on Multi-Dimensional Cloud Model

In the further analysis stage, the evaluation results of the corresponding case (site 5) were introduced into the multi-dimensional cloud model to construct pairwise interactions among the four main indicators (WD, UCS, JS, and BS), forming the diggability multi-dimensional cloud model for the case (Figure 12).

The black point cloud corresponds to the multi-dimensional cloud model reference cloud (Figure 8), while the red point cloud represents the two-indicator interaction cloud used to reveal specific characteristics of the evaluation results.

Taking the WD-UCS interaction cloud model as an example, it can be observed that the cloud droplets are distributed noticeably higher along the UCS dimension, indicating that the case has relatively higher diggability scores in this indicator dimension. Under the combined effect in this region, although both WD and UCS affect the diggability level, the WD indicator demonstrates greater sensitivity in this interaction, thus requiring particular attention in practical evaluation and decision-making. Similarly, in this case, the BS-JS interaction indicates that the BS indicator plays a more critical role than JS in increasing diggability difficulty. Other indicator combinations can similarly be interpreted through such interaction analyses, thereby providing strong support for comprehensively understanding the synergistic mechanisms among different indicators and optimizing evaluation indicators.




Figure 12. Multi-dimensional cloud model evaluation of site 5.

5. Conclusions

This study established a scientific evaluation system for the diggability of mining rock masses and provided corresponding quantitative parameters. A new method for comprehensive diggability evaluation and grading of mining rock masses was proposed. For indicator weighting, a combination of subjective and objective methods was used, taking into account both expert judgment and data characteristics, providing a reliable basis for comprehensive evaluation. Single-indicator cloud models were used to alleviate difficulties in grading diggability at transitional values between adjacent intervals, while multi-dimensional cloud models considered interactions among indicators, achieving optimized indicator interactions and enhancing the interpretability of diggability grades. Comparison with the Diggability Index method indicates that the proposed approach has high reliability and validity. The main conclusions are as follows:

- (1) A diggability evaluation system for rock masses was established, and diggability grade criteria under different conditions were defined.
- (2) Indicator weights were determined using a combination of subjective and objective methods, with the final optimal weights being WD (0.16), UCS (0.18), JS (0.34), and BS (0.32), balancing expert experience and data characteristics.
- (3) Visualization using single-indicator cloud models effectively mitigated the issue of transitional intervals, addressing the difficulty of grading transitional values.
- (4) Multi-dimensional cloud models considered interactions among indicators, optimizing indicator interactions and enhancing the interpretability of diggability grades.

- (5) Evaluation results from multi-dimensional cloud models were highly consistent with the Diggability Index method ($R^2 = 0.991$), with an absolute level accuracy of 74% and a cloud-model-based accuracy of 100%, validating the reliability and applicability of the proposed method.

Although the proposed method demonstrates good reliability and applicability in evaluating the diggability of rock masses, certain limitations should be acknowledged. The dataset used in this study is primarily derived from 56 open-pit mines in Türkiye and the United Kingdom, which may constrain the model's applicability under other geological conditions, mineral types, and underground mining scenarios. Moreover, the combined subjective-objective weighting and cloud model approaches, while effective in addressing indicator uncertainty, may still be affected by expert subjectivity and parameter sensitivity. Future research will focus on expanding data sources and optimizing the model to further enhance its robustness and generalization.

Author Contributions: Conceptualization, S.Y. and X.L.; methodology, S.Y. and X.L.; software, S.Y.; validation, X.L. and M.K.; formal analysis, X.L. and M.K.; data curation, X.L.; writing—original draft preparation, S.Y.; writing—review and editing, J.Z. and M.K.; visualization, S.Y. and X.L.; supervision, J.Z. and M.K.; funding acquisition, J.Z. All authors have read and agreed to the published version of the manuscript.

Funding: This research was partially supported by the Deep Earth Probe and Mineral Resources Exploration—National Science and Technology Major Project (2025ZD1010703) and the National Natural Science Foundation Project of China (52474121 and 42177164).

Data Availability Statement: The data used in this study were obtained from previously published sources [6,14], which are openly available and properly cited within the manuscript.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Stille, H.; Palmström, A. Practical Use of the Concept of Geotechnical Categories in Rock Engineering. *Tunn. Undergr. Space Technol.* **2018**, *79*, 1–11. [CrossRef]
2. Goktan, R.M.; Eskikaya, S. Prediction of Ripping Machine Performance in Terms of Rockmass Properties: Technical Paper. *Civ. Eng. S. Afr.* **1991**, *1991*, 13–14.
3. Karpuz, C. A Classification System for Excavation of Surface Coal Measures. *Min. Sci. Technol.* **1990**, *11*, 157–163. [CrossRef]
4. Howarth, D.F. Mechanical Rock Excavation—Assessment of Cuttability and Borability. *Int. J. Rock Mech. Min. Sci. Geomech. Abstr.* **1989**, *26*, A197. [CrossRef]
5. Bilim, N. Determination of Drillability of Some Natural Stones and Their Association with Rock Properties. *Sci. Res. Essays* **2011**, *6*, 382–387.
6. Scoble, M.J.; Muftuoglu, Y.V. Derivation of a Diggability Index for Surface Mine Equipment Selection. *Min. Sci. Technol.* **1984**, *1*, 305–322. [CrossRef]
7. Santos, V.; Da Silva, P.F.; Brito, M.G. Estimating RMR Values for Underground Excavations in a Rock Mass. *Minerals* **2018**, *8*, 78. [CrossRef]
8. Rusydy, I.; Al-Huda, N. New Rock Mass Classifications for Limestone of the Woyla Group and Its Empirical Relationship in Aceh Province, Indonesia. *Carbonates Evaporites* **2021**, *36*, 15. [CrossRef]
9. Xu, J.; Luo, S.; Xiao, X. Review of the Experimental Studies of the Cracking Behaviors of Fractured Rocks under Compression. *Geohazard Mech.* **2024**, *2*, 59–82. [CrossRef]
10. Hekmatnejad, A.; Rojas, E.; Saavedra, C.; Crespin, B. Presentation of the Universal Discontinuity Index (UDi) System and Its Application to Predict the Geometry of over-Excavation along a Tunnel at New El Teniente Mine. *Eng. Geol.* **2022**, *311*, 106901. [CrossRef]
11. Meng, Y.; Jing, H.; Liu, X.; Yin, Q.; Zhang, L.; Liu, H. Experimental and Numerical Investigation on the Effect of Bedding Plane Properties on Fracture Behaviour of Sandy Mudstone. *Theor. Appl. Fract. Mech.* **2021**, *114*, 102989. [CrossRef]
12. Rotimi, O.J.; Atunbi, J.F.; Enaworu, E.; Ameliko, A.A. Rippability Potential of the Near Surface Deposits of Jubilee Homes Parkland, Southwest, Nigeria. *J. Environ. Earth Sci.* **2015**, *5*, 85.

13. Ismail, M.A.M.; Kumar, N.S.; Abidin, M.H.Z.; Madun, A. Rippability Assessment of Weathered Sedimentary Rock Mass Using Seismic Refraction Methods. *J. Phys. Conf. Ser.* **2018**, *995*, 012105. [CrossRef]
14. Iphar, M.; Goktan, R.M. An Application of Fuzzy Sets to the Diggability Index Rating Method for Surface Mine Equipment Selection. *Int. J. Rock Mech. Min. Sci.* **2006**, *43*, 253–266. [CrossRef]
15. Zhou, J.; Chen, C.; Armaghani, D.J.; Ma, S. Developing a Hybrid Model of Information Entropy and Unascertained Measurement Theory for Evaluation of the Excavatability in Rock Mass. *Eng. Comput.* **2022**, *38*, 247–270. [CrossRef]
16. Bendezu, M.; Romanel, C.; Roehl, D. Finite Element Analysis of Blast-Induced Fracture Propagation in Hard Rocks. *Comput. Struct.* **2017**, *182*, 1–13. [CrossRef]
17. Pinheiro, M.; Emery, X.; Miranda, T.; Lamas, L.; Espada, M. Modelling Geotechnical Heterogeneities Using Geostatistical Simulation and Finite Differences Analysis. *Minerals* **2018**, *8*, 52. [CrossRef]
18. Zhang, X.; Xia, X.; Li, H.; Fu, S. Damage Mechanism and Parameter Optimization of Rock Mass Blasting Under High In-Situ Stress. In Proceedings of the 58th U.S. Rock Mechanics/Geomechanics Symposium, Golden, CO, USA, 23–26 June 2024; OnePetro: Richardson, TX, USA, 2024.
19. Yaghoubi, E.; Yaghoubi, E.; Khamees, A.; Vakili, A.H. A Systematic Review and Meta-Analysis of Artificial Neural Network, Machine Learning, Deep Learning, and Ensemble Learning Approaches in Field of Geotechnical Engineering. *Neural Comput. Appl.* **2024**, *36*, 12655–12699. [CrossRef]
20. Tiriyaki, B. Application of Artificial Neural Networks for Predicting the Cuttability of Rocks by Drag Tools. *Tunn. Undergr. Space Technol.* **2008**, *23*, 273–280. [CrossRef]
21. Huang, S.; Zhou, J. Cutting—Edge Soft Computing Technologies for Rock Mass Excavatability: Transforming Prediction with Hybrid GA—MLP and GEP—Based Criteria. *Rock Mech. Rock Eng.* **2025**. [CrossRef]
22. Kuznetsova, N.; Zhgun, D.; Golovanevskiy, V. Plasma Blasting of Rocks and Rocks-like Materials: An Analytical Model. *Int. J. Rock Mech. Min. Sci.* **2022**, *150*, 104986. [CrossRef]
23. Niu, G.; He, X.; Xu, H.; Dai, S. Development of Rock Classification Systems: A Comprehensive Review with Emphasis on Artificial Intelligence Techniques. *Eng* **2024**, *5*, 217–245. [CrossRef]
24. Hartlieb, P.; Bock, S. Theoretical Investigations on the Influence of Artificially Altered Rock Mass Properties on Mechanical Excavation. *Rock Mech. Rock Eng.* **2018**, *51*, 801–809. [CrossRef]
25. Zhao, H.; Chen, B.; Li, S.; Li, Z.; Zhu, C. Updating the Models and Uncertainty of Mechanical Parameters for Rock Tunnels Using Bayesian Inference. *Geosci. Front.* **2021**, *12*, 101198. [CrossRef]
26. Liu, G.; Zhou, C.; Feng, K.; Jiang, Q.; Li, S.; Bao, H.; Li, A. Probabilistic Evaluation Method for the Stability of Large Underground Cavern Considering the Uncertainty of Rock Mass Mechanical Parameters: A Case Study of Baihetan Underground Powerhouse Project. *Eng. Geol.* **2024**, *340*, 107660. [CrossRef]
27. Alipenhani, B.; Bakhshandeh Amnieh, H.; Majidi, A. Evaluation of Effective Geomechanical Parameters in Rock Mass Cavability Using Different Intelligent Techniques. *Int. J. Min. Geo-Eng.* **2024**, *58*, 307–313. [CrossRef]
28. Yao, S.; Zhou, J.; Khandelwal, M.; Lawal, A.I.; Li, C.; Onifade, M.; Kwon, S. Intelligent Decision Framework for Booster Fan Optimization in Underground Coal Mines: Hybrid Spherical Fuzzy-Cloud Model Approach Enhancing Ventilation Safety and Operational Efficiency. *Machines* **2025**, *13*, 367. [CrossRef]
29. Zhou, J.; Chen, C.; Wei, C.; Du, K. An Improved Connection Cloud Model of an Updated Database: A Multicriteria Uncertainty Model for Coal Burst Liability Evaluation. *Nat. Resour. Res.* **2022**, *31*, 1687–1704. [CrossRef]
30. Guo, D.-S.; Meng, F.-Y.; Wu, H.-N.; Yang, X.-X.; Liu, Z. Risk Assessment of Shield Tunneling Crossing Building Based on Variable Weight Theory and Cloud Model. *Tunn. Undergr. Space Technol.* **2024**, *145*, 105593. [CrossRef]
31. Barton, N. Suggested Methods for the Quantitative Description of Discontinuities in Rock Masses: International Society for Rock Mechanics. *Int. J. Rock Mech. Min. Sci. Geomech. Abstr.* **1978**, *16*, 22.
32. Li, D.; Liu, C.; Gan, W. A New Cognitive Model: Cloud Model. *Int. J. Intell. Syst.* **2009**, *24*, 357–375. [CrossRef]
33. Li, S.; Wang, G.; Yang, J. Survey on Cloud Model Based Similarity Measure of Uncertain Concepts. *CAAI Trans. Intell. Technol.* **2019**, *4*, 223–230. [CrossRef]
34. Gong, S.; Yao, S.; Xi, F.; Liu, J.; Ren, X.; Niu, D.; Ding, J.; Zhang, H.; Wang, J.; Wang, S. A Comparison of Research Methods to Determine the Sustainability of Mineral Resources in Henan Province Based on Cloud Analysis. *Sustainability* **2022**, *14*, 15834. [CrossRef]
35. Yang, J.; Han, J.; Wan, Q.; Xing, S.; Chen, F. A Novel Similarity Measurement for Triangular Cloud Models Based on Dual Consideration of Shape and Distance. *PeerJ Comput. Sci.* **2023**, *9*, e1506. [CrossRef]
36. Deng, Y.; Liu, S.; Zhang, W.; Wang, L.; Wang, J. General Multidimensional Cloud Model and Its Application on Spatial Clustering in Zhanjiang, Guangdong. *J. Geogr. Sci.* **2010**, *20*, 787–798. [CrossRef]
37. Wang, M.; Wang, X.; Liu, Q.; Shen, F.; Jin, J. A Novel Multi-Dimensional Cloud Model Coupled with Connection Numbers Theory for Evaluation of Slope Stability. *Appl. Math. Model.* **2020**, *77*, 426–438. [CrossRef]

38. Zheng, Y.; Shi, Z.; Peng, X.; Li, J.; Feng, X. Multidimensional Generalized Normal Cloud Model for Uncertainty Propagation. *Comp. Appl. Math.* **2025**, *44*, 272. [CrossRef]
39. Vidal, L.-A.; Marle, F.; Bocquet, J.-C. Using a Delphi Process and the Analytic Hierarchy Process (AHP) to Evaluate the Complexity of Projects. *Expert. Syst. Appl.* **2011**, *38*, 5388–5405. [CrossRef]
40. Saaty, T.L. *The Analytic Hierarchy Process: Planning, Priority Setting, Resource Allocation*; McGraw-Hill: New York, NY, USA, 1980; ISBN 978-0-07-054371-3.
41. Pant, S.; Kumar, A.; Ram, M.; Klochov, Y.; Sharma, H.K. Consistency Indices in Analytic Hierarchy Process: A Review. *Mathematics* **2022**, *10*, 1206. [CrossRef]
42. Zhou, K.; Lin, Y.; Deng, H.; Li, J.; Liu, C. Prediction of Rock Burst Classification Using Cloud Model with Entropy Weight. *Trans. Nonferrous Met. Soc. China* **2016**, *26*, 1995–2002. [CrossRef]
43. Qiao, C.; Wang, Y.; Li, C.; Yan, B. Application of Extension Theory Based on Improved Entropy Weight Method to Rock Slope Analysis in Cold Regions. *Geotech. Geol. Eng.* **2021**, *39*, 4315–4327. [CrossRef]
44. Wang, N.; Zhang, Y. Research on Evaluation of Wuhan Air Pollution Emission Level Based on Entropy Weight Method. *Sci. Rep.* **2024**, *14*, 5012. [CrossRef]
45. Ghorbani, S. A Hybrid Quantitative Approach for Assessment of Geotechnical Hazards in Rock Tunnels Using Finite Element and Variation Coefficient Methods. *Sci. Rep.* **2025**, *15*, 28802. [CrossRef]
46. Ecer, F.; Haseli, G.; Krishankumar, R.; Hajiaghahi-Keshteli, M. Evaluation of Sustainable Cold Chain Suppliers Using a Combined Multi-Criteria Group Decision-Making Framework under Fuzzy ZE-Numbers. *Expert Syst. Appl.* **2024**, *245*, 123063. [CrossRef]
47. Jolliffe, I.T.; Cadima, J. Principal Component Analysis: A Review and Recent Developments. *Philos. Trans. R. Soc. A Math. Phys. Eng. Sci.* **2016**, *374*, 2065. [CrossRef]
48. Tripathi, M.; Singal, S.K. Allocation of Weights Using Factor Analysis for Development of a Novel Water Quality Index. *Ecotoxicol. Environ. Saf.* **2019**, *183*, 109510. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Article

A Study of Tau-Robot Configuration for Friction Stir Welding

Despoina Almpani and George-Christopher Vosniakos *

School of Mechanical Engineering, National Technical University of Athens, 157 72 Athens, Greece;
albanidespina@yahoo.com

* Correspondence: vosniak@central.ntua.gr

Abstract

This paper examines the use of high-rigidity Tau-robots in friction stir welding, where process loads are very high. The rigidity of Tau-robots increases at the expense of the workspace. Therefore, the right configuration of the Tau-robot is sought to reconcile rigidity and workspace requirements. This is studied by use of kinematics, followed by static and modal analysis. In particular, by extending an existing kinematic model employing free vectors, the robot workspace was derived in non-dimensional parametric form and was then maximized through evolutionary optimization. However, finite element static and modal analysis that were carried out subsequently may prove, as in a case demonstrated here, that the optimized configuration may not withstand high loads, typically axial forces of 15 kN and torques of 80 Nm, and it may also be susceptible to forced vibrations in the typical spindle rotation range up to 3000 rpm. As a rectification measure, it was shown how a modified configuration by placing robot kinematic chain bases further apart and shortening robot links achieves higher rigidity, axial displacement being reduced by one or two orders of magnitude to below 1 mm and increases critical modal frequency 3 to 5 times depending on the workspace position, of course sacrificing part of the workspace, i.e., reducing it 3-fold to enclose welding lines in a rectangle of dimensions 700×800 mm. In the quest for the appropriate robot configuration desired dimensions of parts to be welded and available standard components are briefly considered, too.

Keywords: Tau-robot; friction stir welding; closed kinematic chain; workspace; modal analysis; optimization

1. Introduction

Friction Stir Welding (FSW), since its invention by The Welding Institute (TWI) in 1991, has become a key technology for joining lightweight materials, particularly aluminum alloys, due to its ability to produce high-quality welds with minimal distortion and excellent mechanical properties since the material remains largely in a solid state throughout the welding process [1]. Traditionally, FSW has been performed using rigid machine tools such as dedicated FSW machines, modified milling machines, or CNC machining centres. These platforms provide the high stiffness and precision needed to sustain the large axial forces and precise tool positioning required during welding [2]. However, machine-based FSW systems are typically expensive, inflexible, and require significant floor space, which limits their appeal for low-volume or large, complex parts [3]. To overcome these limitations, recent research has explored the use of industrial robots for FSW. Robots offer significant advantages: high flexibility, the ability to access complex geometries, reconfigurability for different tasks, and the potential for cost-effective automation in low-to-medium production volumes [4].

At a first approximation, the following requirements for the robot can be defined:

- Strength and power: The robot should be able to compensate for the developed loads.
- Robustness and precision: The compliance of the robot should be minimal so that the tool does not deviate from the desired path and is not affected by vibrations.
- Dexterity: At least 5 degrees of freedom are required to allow adjustment of both the position of the tool and the welding angles.
- Measurement and control: The robot should have position, force and temperature measurement sensors to adjust welding parameters through a control loop.

Despite its potential, robotic FSW has to overcome several technical issues. The primary limitation is the relatively low stiffness of standard industrial robots compared to machine tools [5], leading to path deviations under high process loads, especially during plunge and initial welding phases [6]. This compliance can result in defects such as voids, tunnel defects, or inconsistent weld properties [7]. In addition, thermal effects during welding can exacerbate robot positioning errors due to thermal expansion and changes in the mechanical properties of the links [8].

Parallel kinematic robots, such as the Tau type, offer a promising alternative. Their closed kinematic chains provide higher dynamic stiffness and better load distribution, potentially overcoming the compliance issues faced by serial robots at the expense of comparatively severely limited workspace [9]. In addition, their design, while stiff, introduces complex singularities and workspace boundaries that complicate trajectory planning and real-time control to manage both position and force simultaneously under dynamic loading, especially under high process loads [10]. Addressing these issues is critical to making Tau-robots a viable alternative to traditional FSW machine tools.

While parallel robot architectures have been proposed for FSW and isolated studies involving Tau or Gantry-Tau mechanisms exist, the gap identified concerns the investigation of an appropriate Tau-robot configuration that fulfils FSW requirements, balancing structural performance and workspace size in connection to intended welding seam shape and size. In this direction, the following three interrelated aspects are investigated in this work: (i) a parametric (dimensionless) geometric configuration of the Tau-robot open to optimization rather than predefined Tau-robot configurations, (ii) explicit stiffness–workspace trade-off for FSW conditions showing how geometry modification can restore stiffness at the expense of workspace rather than seeking workspace maximization or stiffness improvement separately, (iii) static and dynamic behaviour under representative FSW loads based on finite element static and modal analysis.

Following this introduction, an in-depth literature review is provided in Section 2. A succinct analysis of FSW is presented in Section 3. In Section 4, kinematic analysis and workspace determination of a parametrically defined Tau-robot are provided, leading to an optimal configuration. For this configuration, static and dynamic analysis under loading is presented in Section 5, resulting in a modified configuration. Its suitability for FSW is discussed in Section 6, including some implementation considerations. Conclusions and outlook are given in Section 7.

2. Literature Review

FSW process windows, tool design, and metallurgy are intertwined with robotics issues such as positioning accuracy, structural stiffness, and real-time control [4,8].

Early robot design for FSW focused on serial links; more recent work aims at optimization of structural members [11] and explores parallel robot topologies to push up stiffness without sacrificing workspace [12]. Redundant actuation and careful dimensional synthesis can make process parameter errors less important and redistribute internal forces [13].

Model-based stiffness prediction enables dimensioning of links and joints such that the end-effector deflection stays within tolerance across the envelope [14].

A comprehensive stiffness modelling framework for a 5-DoF hybrid manipulator designed for friction stir welding, incorporating both serial and parallel subsystem compliances, is presented in [15]. The model is analytically and experimentally evaluated to assess stiffness distribution and its impact on machining accuracy. A model-based approach for predicting the stability of robotic machining processes by coupling robot structural dynamics with manufacturing process force models enables identification of stable operating regions [16]. The performance of a Gantry-Tau parallel mechanism is evaluated in [17] through kinematic and experimental analysis.

Robotic FSW is dominated by load-induced deflection and thermal drift over trajectories [18]. Parameter studies reveal how feed rate, rotational speed, and tool geometry map to axial force transients that controllers must gracefully absorb to avoid chatter or wormhole defects [19]. Considering measured force in position loops enhances seam tracking and plunge robustness in the presence of fixture or geometry uncertainty [20,21]. By anticipating force-dependent deviations and pre-deforming the path, weld seam geometry and plunge consistency improve, especially on long curvilinear seams [22].

Digital-twins now embed finite-element thermal-mechanical models with robot dynamics, allowing verification of contact conditions, plunge depth, and force margins before hardware trials [23]. For complex space-curve welds (e.g., on monocoque structures), Tau-robot installation optimization using dynamic dual particle swarm methods determines the workpiece placement that maximizes stiffness margins and avoids kinematic singularities along the seam [24]. Kinematic sensitivity analysis and dimensional synthesis further guide the early-stage design of parallel robots with redundant degrees of freedom so that isotropy, transmission factors, and singularity loci are acceptable within the anticipated weld envelope [12].

In summary, robotic FSW has been improving through: (a) increased structural stiffness and improved load distribution achieved by parallel or hybrid kinematic architectures, (b) reduction in end-effector deflection through model-based stiffness optimization, (c) improved trajectory accuracy and seam quality through force-based control and feed-forward compensation methods, (d) enhanced process stability through dynamic modelling and digital-twin approaches. However, a practical trade-off is that workspace enlargement often leads to reduced structural stiffness and lower natural frequencies, which may compromise suitability for FSW. In addition, control-based compensation approaches may be effective in improving path accuracy, but do not eliminate the underlying mechanical compliance and therefore cannot fully replace structural optimization.

Robotic FSW applications continue to expand. Battery pack assembly in lightweight electric vehicles increasingly uses robotic FSW for extrusion-cast aluminum, where narrow flanges, high thermal mass, and dissimilar interfaces demand both stiffness and dexterous reach [25]. FSW using Temperature-Controlled Robotics and dissimilar materials joining has been reported for demanding applications, such as for high-pressure die-cast cylinder blocks [26]. Repair and retrofit use cases, e.g., of train sets, large aerostructures, etc., benefit from mobile or gantry-mounted Tau-robots with compliant fixtures maintaining weld quality even on thin skins and curved geometries [13,27]. Case studies confirm the existing practicality gap, namely, maintaining force and pose fidelity over large workspaces, on complex seams, under temperature-dependent material properties.

3. Load and Accuracy Requirements in Robotic FSW

The FSW tool consists of a specially shaped pin, usually with some type of helix on its perimeter and a shoulder which is of larger diameter and particularly robust. As the tool

rotates, the pin of the tool penetrates the interface of the two materials to be welded. Due to the frictional forces, the temperature locally increases, and the materials enter the plastic domain, yet without melting. Due to the special geometry of the pin, the material is stirred and is transferred from the front (advancing side) to the back (retreating side) of the tool, where it cools due to the significant temperature difference in these two regions. There, in the presence of the pressure exerted by the shoulder of the tool, it solidifies completely [3]. The process is carried out by simultaneously advancing the tool along the interface. With additional material removal and/or grinding, the welding trace can literally disappear.

The so-called ‘rotating plug’ model is most commonly used for estimating the welding torque and power in FSW: $M = 2\pi R^3 (1 + 3 r^2 t/R^3) Y/3$ and $P = M\omega$, where M/P : the torque/power, Y : the yield strength of the material, r : the radius of the tool pin, R : the radius of the tool shoulder, t : the depth of the pin, and ω : the rotational speed of the tool [28]. Two additional parameters are the tilt angle (φ_t) and yaw angle (φ_y), see Figure 1.




Figure 1. FSW process setup (φ_t : tilt angle, φ_y : yaw angle, AS: advancing side, RS: retreating side).

Figure 1 also shows the developed forces on the workpiece in the steady state. The axial component F_z is responsible for the developed friction between the tool and the workpiece, contributing to the heat generation. In addition, it exerts pressure on the newly welded area, thereby improving weld quality. The frontal component F_x results from the resistance of the cold material along the seam path. The lateral component F_y is due to the asymmetry of the process, as the material on the advancing side is hotter and therefore softer compared to that on the retreating side [18].

The temperature in the yield zone varies little over a wide range of operating conditions. Therefore, the spindle power remains approximately constant over a wide range of speeds. Hence, the torque may be reduced if the process is carried out at high speeds [29], provided that the increase in frictional shear rate does not take over the thermal softening. Note that the axial force during the initial penetration of the tool into the workpiece also decreases with speed, but it can be up to three times higher than that observed during the welding seam generation by the tool following the welding path [6].

Typically, to cover a wide range of materials and configurations to be welded in practice, the spindle power should exceed 4 kW, and the working angle should be possible to vary between 0 and 5°. The developed loads reach up to 30 kN and 20 kN for the axial and lateral force, respectively [4,8], which are extremely high. However, for the case of aluminum alloys and polymers, load requirements are significantly lower, typically by a factor of 5–10 [30]. In [3], for instance, an extensive table is provided presenting welding parameters and tool dimensions for various cases of materials. As far as accuracy is

concerned, the maximum deviations from the nominal size for weldments to be acceptable are position: ± 1 mm, tilt angle: $\pm 1^\circ$, yaw angle: $\pm 2^\circ$ [31,32].

4. Tau-Robot Parametric Kinematic and Workspace Analysis

The Tau-robot configuration can be implemented using either linear or rotary actuators. In the second case, the tool centre point (TCP) has only 3 degrees of freedom, i.e., its orientation is predetermined, and therefore, the type of welds that can be achieved is limited. In addition, the linear action is easier to construct; it corresponds to distributed loads on the actuator bases, enhancing robustness, and the corresponding kinematics solutions are known [33].

4.1. Forward Kinematics

In Figure 2, the configuration of a linearly actuated Tau-robot is presented. Its kinematics analysis is summarized next, following the detailed treatise presented in [33]. The robot consists of three individual kinematic chains: the first having one arm (rod), the second having two arms parallel to each other, and the third having three arms, also parallel to each other. Each arm is connected via a 2 DoF Universal Joint (UJ) to the base, and via a 3 DoF Spherical Joint (SJ) to the final frame accommodating the TCP. The base of each chain is connected via a Prismatic Joint (PJ) to a linear guide. The actuation is performed on each of the three PJs along the linear guides.




Figure 2. Configuration of linearly actuated Tau-robot and its vectorial analysis (UJ: universal joint, PJ: prismatic joint).

The three DoFs, which constitute the control variables of the system, are denoted as $Q_i, i = 1, 2, 3$. To simplify the kinematic equations, the global coordinate system is placed so that the global X-axis is parallel to the three linear guides, i.e., the three control variables take only positive values. The TCP is placed on the final frame, where all three individual kinematic chains end up and constitutes the Reference Point. The latter is defined in the global coordinate system by the position vector $\vec{R}$:

$$\vec{R} = \vec{a}_i + \vec{Q}_i + \vec{p}_{ij} + \vec{L}_i + \vec{b}_{ij}, i = 1, 2, 3 \quad j = 1, \dots, i \quad (1a)$$

where (see Figure 2)

$\vec{a}_i$: The position vector of the linear guide i ;

$\vec{Q}_i$: The linear DoF of guide i ;

$\vec{p}_{ij}$: The vector starting at the linear joint i and ending at spherical joint j of link i ;

$\vec{L}_i$: The length vector of link i ;

$\vec{b}_{ij}$: The vector from the spherical joint j , where link i ends, to the TCP.

Due to the Tau-configuration for each link i , the sums $\vec{p}_{ij} + \vec{b}_{ij}, \forall j$ are equivalent to each other. Thus, all (known) constant terms add up to the vector $\vec{G}_i = \vec{a}_i + \vec{p}_{ij} + \vec{b}_{ij}$ hence Equation (1a) becomes:

$$\vec{R} = \vec{G}_i + \begin{Bmatrix} L_{ix} & L_{iy} & L_{iz} \end{Bmatrix} + \begin{Bmatrix} Q_i & 0 & 0 \end{Bmatrix} \quad (1b)$$

4.2. Dimensionless Equations

The robot workspace is mapped by assigning numerical values to the (constant) geometric entities involved. This is also necessary because it is impossible to express analytically the geometric constraints that ensure the closed nature of kinematic chains. Given the lack of experience in designing and using Tau-robots, it was decided to optimize a dimensionless configuration of the robot. The optimization criterion is the increase in the robot workspace.

The maximum travel of any DoF is defined as Q , thus $Q_i \in [0, Q]$. Then, new dimensionless variables are defined as follows:

$$q_i = \frac{Q_i}{Q}, q_i \in [0, 1], r_j = \frac{R_j}{Q}, g_{ij} = \frac{G_{ij}}{Q}, l_{ij} = \frac{L_{ij}}{Q} \quad j = x, y, z \quad i = 1, 2, 3 \quad (2)$$

Thus, the new system of equations defining the robot's kinematics is expressed as follows (compare with Equations (1a) and (1b)):

$$\begin{cases} r_x = g_{ix} + l_{ix} + q_i \\ r_y = g_{iy} + l_{iy} \\ r_z = g_{iz} + l_{iz} \end{cases}, i = 1, 2, 3 \quad (3)$$

and

$$l_{ix}^2 + l_{iy}^2 + l_{iz}^2 = c_i, i = 1, 2, 3 \quad (4)$$

where

$$c_{is} = \sqrt{c_i} = \frac{|L_i|}{Q} \quad (5)$$

The unknowns are $r_x, r_y, r_z, l_{ix}, l_{iy}, l_{iz}, i = 1, 2, 3$.

Similarly, the corresponding dimensionless geometric parameters are:

$$g_{ix}^2 + g_{iy}^2 + g_{iz}^2 = k_i, i = 1, 2, 3 \quad (6)$$

where

$$k_{is} = \sqrt{k_i} = \frac{|G_i|}{Q} \quad (7)$$

An additional dimensionless factor that quantifies the proportionality of link size with respect to other structural quantities is defined as:

$$m_i = \frac{c_{is}}{k_{is}} \quad (8)$$

Note that (a) each vector $\vec{L}_i$ has a constant magnitude but its orientation is variable and generally unknown, and (b) each vector $\vec{G}_i$ has a constant magnitude and con-

stant and known orientation. Therefore, by introducing the unit vector $\hat{G}_i$, the following expression holds:

$$\vec{G}_i = \hat{G}_i \|\vec{G}_i\| \Rightarrow g_{ix} = \frac{\hat{G}_{ix} \|\vec{G}_i\|}{Q} = \hat{G}_{ix} k_{is}, \quad i = 1, 2, 3 \quad (9)$$

In Figure 3a, the vectors involved in defining the constant geometry of the configuration are depicted. As a default, $\vec{G}_1$ is positioned between $\vec{G}_2, \vec{G}_3$. Furthermore, using a suitable global coordinate system definition, $\vec{G}_{ix}$ can be neglected, yielding $\hat{G}_i = \{0, \cos(\theta_i), \cos(\theta_i - 90^\circ)\}$; see Figure 3b.




Figure 3. (a) Constant vectors of the robot's geometry (view from behind). (b) Vector $\vec{G}_i$.

4.3. Optimization of Dimensionless Variables

The vector to be optimized is:

$$\vec{x} = \{\theta_1, \theta_2, \theta_3, m_1, m_2, m_3, c_{1s}, c_{2s}, c_{3s}\} \quad (10)$$

The applicable constraints are: an upper bound: $\{85, 85, 85, 100, 100, 100, 2, 2, 2\}$, a lower bound: $\{-85, -85, -85, 0.1, 0.1, 0.1, 0.1, 0.1, 0.1\}$ and $\theta_3 \leq \theta_1 \leq \theta_2$.

The upper and lower bounds were set as follows. Preliminary exploratory runs and CAD-based geometric trials showed that very small parameter values lead to fragmented or extremely limited workspaces, whereas excessively large values generate elongated configurations with poor geometric conditioning and reduced structural realism. In addition, the adopted intervals were chosen so that the nonlinear kinematic system remains numerically stable and yields a sufficient percentage of real (physically meaningful) solutions during optimization. Although the model is dimensionless, the selected bounds also reflect typical proportional ranges encountered in Tau-type mechanisms, avoiding configurations that would be impractical. Therefore, the optimization bounds represent a compromise between workspace enlargement, computational efficiency and realistic geometric proportions.

Trying out a sufficiently large number of vector values in a 'brute force' mode would be computationally very intensive since at each and every point, a 12×12 system of equations needs to be solved, corresponding to a total computation time of the order of 24 h for a dense mesh of points. Therefore, a relatively sparse mesh of points is adopted, and solutions for Q_i are sought not in the interval $[0, 1]$ but in its subset $[0.3, 0.7]$. This

selection resulted from repeated trials using a selection of different robot dimensions to define robot models in a CAD environment. The choice was validated by considering the shape of the workspace at the end of this section.

Furthermore, instead of maximizing the workspace volume, it was chosen to minimize the percentage of unacceptable solutions in a reference volume discretized as a reasonably sparse mesh.

Optimization was performed through a classic genetic algorithm that was set up in MatlabTM R2020b optimization toolbox, using mostly the system default parameter values except for the following:

- (a) Population set to 50 instead of the default 200 chromosomes to reduce execution time;
- (b) Elitism set to 1 chromosome;
- (c) Termination condition: Either 30 generations or minimization of the objective function within 0.001, whichever occurs first.

Upper and lower bounds for variables were defined with an accuracy of 0.01.

The objective function simply counts the number of inaccessible points. To do so, the end-effector is positioned at each and every point of the mesh. Then, the inverse kinematics solutions are derived, and their acceptability is evaluated with the aim of labelling them accessible or inaccessible. Note that the 12×12 system of equations that is solved at each point always yields two mathematical solutions for each end-effector position coinciding with the point under consideration. This occurs because the closure equations of the Tau-robot allow two geometrically feasible configurations (poses) that satisfy the same end-effector coordinates. However, these are just mathematically admissible solutions and not necessarily physically meaningful. The counter of accessible points is therefore increased by one if both conditions hold. A complex number solution is rejected by definition. Furthermore, the workspace should not contain points that correspond to any collision between links, but this is expensive and difficult to express. Thus, duplicate solutions are avoided by just setting $L_{ix} > 0$, i.e., links are turned towards the positive global $\hat{x}$ and not the other way around.

The convergence history of the genetic algorithm is shown in Figure 4. The fitness of the average chromosome does not fall so fast, but the best chromosome does indeed. The difference between the mean (77.9) and best (0) fitness values is large enough to prove the existence of diversity in the population. Since the absolutely best solution was found, i.e., the one corresponding to all points examined being accessible with that particular robot configuration, it was thought to be superfluous to allow the genetic algorithm to run for more generations. This is due to the computational cost of the inverse kinematics solution, being applied exhaustively for all the points of the potential workspace.




Figure 4. Evolution history of the genetic algorithm solution.

The results are shown in histogram form in Figure 5, for each parameter. Overall, a low non-accessibility percentage is achieved ($P < 3\%$). The best solution turns out to be $\{70^0, 75^0, -83^0, 60, 8, 3, 2, 2, 2\}$.




Figure 5. Optimization of angles $\vec{\theta}_i$, coefficients m_i and coefficients c_{is} .

Furthermore, $m_1 = 60 \gg m_2 = 3 > m_3 = 2$, i.e., link 1 is much longer than links 2 and 3, which is in agreement with the position of the coordinate system origin on the guide of link 1, i.e., with a low contribution of $\vec{a}_1$ and $\vec{p}_{11}$.

As far as coefficients c_{is} are concerned, they tend to their maximum allowable values, thereby pointing to the need to use as long as possible links.

4.4. Workspace Computation

Based on the optimized Tau-robot configuration, as determined in the previous section, it is now possible to compute the coordinates (r_x, r_y, r_z) of the robot end-effector by using an appropriately small step equal to 0.1 for each DoF $Q_i \in [0, 1]$, thereby yielding the dimensionless workspace; see Figure 6a.




Figure 6. Dimensionless workspace (a) $Q_i \in [0, 1]$ and (b) $Q_i \in [0, 3]$.

The domain is arc-shaped, and its volume is $v_{wspc} = 5.0625$. If the length of the links is increased by using $c_{is} = 5$ or even $c_{is} = 10, i = 1, 2, 3$, keeping all the rest parameters invariable, the workspace volume is computed as $v_{wspc} = 14.5970$ or 26.3137 correspondingly, i.e., significantly higher. In fact, it is not necessary for Q to correspond to the maximum length of the links. For example, in Figure 6b, the dimensionless workspace corresponding to $Q_i \in [0, 3]$ is $v_{wspc} = 22.3$, i.e., 5.4 times larger than the initial one.

The shape of the robot workspace seems suitable for FSW of planar or near-planar surfaces extending along Z or along X axes. However, the suitability of the Tau-robot is still to be confirmed from both strength and dynamics points of view.

5. Tau-Robot Static and Modal Analysis

An important feature of the robot is its robustness, which for complex models can only be determined by numerical/finite element analysis (FEA) methods [23,33]. For this reason, it is chosen to perform static analysis and modal analysis of a dimensioned instance of the robot in some indicative poses while maintaining the optimized ratio coefficients of the dimensionless case.

The robot instance's loads correspond to welding of aluminum alloys, i.e., $F_z = 15$ kN, $F_x = 2$ kN, as well as $M_z = 80$ Nm for the case of maximum power of 4 kW, and maximum speed of 3000 rpm. The first eigenfrequency must be greater than the critical operating frequency, i.e., $f_c = 50$ Hz (3000 rpm). It is noted that the modal analysis is simplified as internal damping of the components is not considered.

Thus, applying $Q = 1$ m to Equations 1–9, the following values are obtained: $c_{is} = 2$, $L_1 = L_2 = L_3 = 2$ m, $k_{is} = \{0.03 \ 0.25 \ 0.67\}$, $\{G_1 \ G_2 \ G_3\} = \{0.03 \ 0.25 \ 0.67\}$.

From Equations (6) and (9), the geometric constants $\vec{G}_i$ are obtained, and the individual components that make it up are free to choose. Thus, seeking the greatest possible separation of linear vectors, the vectors of Figure 7 are selected. The End-Effector (EE) is the point of reference and the intersections $\vec{b}_i$ ending on it are transferred to the end of the corresponding $\vec{a}_i$ to compute the respective $\vec{G}_i$.




Figure 7. Selection of geometric fixed vectors (EE: end-effector).

A space of 400×400 mm is provided at the EE to accommodate a 4 kW/3000 rpm motor, and a typical FSW tool of diameter of $\varnothing 8$ mm is placed at its end. The links and bases are simplified for kinematic modelling; see Figure 2.

The 3D models of the robot were created in the Siemens NXTM 3D CAD environment, whilst FEA was conducted in the ANSYSTM R19 environment. Note that, due to the nature of the welding, there are no large inertial loads, and the main stress is due to the large imposed loads. Therefore, the static study in critical poses of the device offers essential information on the suitability of the robot.

The following quantities are determined from the study:

- Displacement of the EE under the applied loads. This is the open-loop behaviour, which should be in the order of a few mm before applying any control.
- Maximum loads on the bases aiming to determine the requirements on the supports and actuators.
- Maximum stresses on the links aim to check the adequacy of the cross-section.

5.1. Static Analysis

The link rods are assigned a diameter of 75 mm and are made of tool steel DIN 1.2738 ($\sigma_Y = 980$ MPa). They are modelled as deformable bodies, whereas the bases, U- and S-Joints, end-plate and workpiece are modelled as rigid bodies. Rigid joints correspond to the respective DoFs, whereas the bases and the workpiece are fixed. The loads on the tool are modelled as joint loads. The loads are F_z (local axial) = 15 kN, F_x (local side) = 2 kN,

M_z (local axial) = 80 Nm. Two alternatives for the welding line direction are considered, namely: along the global Z-axis (WLZ) and along the global X-axis (WLX).

In case WLZ, the main load in the axial direction is in the global X direction; therefore, the stiffness is higher than in case WLX. However, in case WLX, the welding path can be of any desired length and, in addition, no significant changes in the robot pose are expected (hence in system dynamics, too).

Adaptive Meshing of tetrahedra and trihedra of general size ‘fine’ is adopted in AnsysTM. Two element sizes were tried, namely 40 mm and 15 mm, resulting in an average number of elements per link equal to 10,253 and 30,047, respectively. The results corresponding to the same loading scenario (see Figure 8) are shown in Table 1. The largest difference between large and small element sizes emerged in the max von Mises stress cases. This reflects the need to use finer elements when geometry becomes more intricate at the end of the links, interfacing with connectors; the adoption of large elements there is likely to lead to stress singularities in view of idealized joint or boundary condition modelling. However, note that the already very low level of stress becomes even lower when smaller elements are employed, which is expected to approach reality more closely.




Figure 8. Model case WLZ.

Table 1. Mesh size influence (legend: Max_Stress_N-M: Max von Mises stress at link N, rod M).

Property	Element Size		Property	Element Size	
	40 mm	15 mm		40 mm	15 mm
Max_Stress_1-1 (MPa)	8.8297	4.6525	F_{total} Base 1 (kN)	9.668	9.675
Max_Stress_2-1 (MPa)	10.537	5.7171	F_{total} Base 2 (kN)	25.827	25.831
Max_Stress_2-2 (MPa)	13.411	7.2219	F_{total} Base 3 (kN)	1.820	1.818
Max_Stress_3-1 (MPa)	0.6619	0.8245	M_{total} Base 1 (kNm)	6.4553	6.4601
Max_Stress_3-2 (MPa)	6.5340	3.5953	M_{total} Base 2 (kNm)	2.2888	2.3292
Max_Stress_3-3 (MPa)	5.0508	2.8647	M_{total} Base 3 (kNm)	1.7672	1.7638
Side Displacement (mm)	1.4024	1.4169	F_{total} Clamp (kN)	17.361	17.369
Axial Displacement (mm)	−0.4191	−0.4238	M_{total} Clamp (kNm)	11.455	11.462

Figure 9a depicts the workspace points studied in both WLZ and WLX cases of the original configuration. Note that of the six characteristics points studied, only points A, B, and C are visible since the rest are occluded. Table 2 presents the coordinates of the six characteristic points studied.




Figure 9. Robot end-effector points studied: (a) original configuration, (b) modified configuration.

Table 2. Workspace points studied in the original configuration.

End-Effector Vector	Points					
	A	B	C	D	E	F
R_x (mm)	1746.0	1674.0	2876.0	1978.0	2300.0	2000.0
R_y (mm)	346.8	−594.6	−434.0	−434.0	−434.0	0.0
R_z (mm)	1175	−1522	−198.6	−198.6	−198.6	500.0

5.2. FEA Results

Figure 10 presents FEA results, as an example, for pose A in case WLX.




Figure 10. Pose A—case WLX. (a) von Mises stress of rod 2 of link 2, (b) modal analysis.

Figure 10a depicts the von Mises stress in one of the links of the assembly. It is observed that a very low average value is maintained throughout, while the maximum of about 130 MPa is located at the connection points with the respective joints. These local peaks are mainly associated with boundary-condition effects and modelling simplifications, since the joints are represented only by kinematic degrees of freedom and are not evaluated from a strength perspective. Therefore, stress values near the joint interfaces are interpreted qualitatively, whereas the stress distribution along the main body of the rod is considered representative for structural assessment. Even if maximum stresses observed are unrealistically high due to the simplifications adopted in modelling, they still remain well below the material yield limit; the main focus of the static analysis is on the calculated displacements.

Figure 10b shows the results of the natural frequency analysis.

Analogous results appear in all poses and cases, and Table 3 sums them up.

Table 3. Static analysis results for the original configuration for 6 end-effector coordinates.

Property	Case											
	WLZ						WLX					
	Pose											
	A	B	C	D	E	F	A	B	C	D	E	F
Max von Mises Stress—Rods (MPa)	10	40	30	30	30	20	130	120	60	50	50	30
Max Reaction Force—Bases (kN)	25.0	43.9	62.9	61.5	61.5	23.7	111.4	129.0	78.5	65.0	65.0	46.0
Max Reaction Moment—Bases (kNm)	6.5	28.8	69.5	51.0	51.0	15.0	102.0	91.0	43.0	33.8	33.7	18.3
Max Reaction Force—Clamping System (kN)	17.0	17.0	26.0	27.8	27.8	18.8	52.9	39.0	29.4	21.2	21.2	75.0
Max Reaction Moment—Clamping System (kNm)	11.5	2.58	26.0	14.9	14.9	2.0	54.9	2.1	15.3	14.1	14.1	39.4
Axial Relative Displacement—Tool (mm)	−0.4	−0.2	−0.1	−0.2	−0.2	0.0	−5.6	−10.4	1.8	−0.4	−0.4	−0.2
Side Relative Displacement—Tool (mm)	1.4	−0.5	−0.9	−0.9	−0.9	−0.5	45.4	89.3	−4.0	−0.3	−0.3	3.5
First Mode (Hz)	11.6	17.1	9.7	10.7	10.7	8.0	2.1	1.8	18.0	24.2	24.2	5.0

Regarding maximum von Mises stress on links, since the yield stress of 1.2738 steel (DIN) is 900 MPa, there are no concerns associated with any of the poses. Regarding the reactions at the bases, both forces and moments are considered excessive (case WLZ) or even formidable (case WLX). As far as reactions in the clamping system are concerned, both forces and moments seem manageable, except for case WLX, where a very high force value of 75 kN is shown. Note that the bases and the clamping system can be made of many individual components that, if appropriately dimensioned, can collectively manage the required loads studied in separate FEA, which is outside the purpose of this study.

Regarding the relative displacement of the end-effector, results are satisfactory for case WLZ, as the largest displacement is barely 1.4 mm. However, in case WLX, the largest deviations exceed 10 mm.

As far as the first eigenfrequency is concerned, this is lower than 50 Hz in all poses, which hints at possible resonance in case of high rotary speeds of the spindle. Better performance is observed at the inner points of the workspace. In general, case WLX leads to much worse performance, mainly due to low stiffness.

Note that the increase in the workspace has led to small distances among the bases and, at the same time, to proportionally very long links. This configuration makes the distribution and possible mutual partial cancelling out of the developing forces and moments problematic. Thus, it is worth conducting a corresponding investigation into a modified arrangement to increase robustness at the expense of the workspace.

6. Tau-Robot Modified Configuration

The new modified configuration has larger distances between the bases and decreased length of the links compared to the original, i.e., from 2 down to 1.5 m; see Figure 11.




Figure 11. Original configuration (a) versus modified (b) configuration.

A comparison of the dimensionless parameters with those of the original configuration is presented in Table 4. As for the angles of the geometric vectors $(\vec{G}_i, \hat{y})$, the comparison is presented only for completeness, as the new global coordinate system has been placed at a different, more favourable point, with orientation still being maintained.

Table 4. Comparison of dimensionless parameters between initial and modified configurations.

Quantity	Initial	Modified
angles $(\vec{G}_i, \hat{y})$	{70,75,−83}	{29,96,6.5}
m_i	{60,8,3}	{7.1,1.2,1.4}
c_{is}	{2,2,2}	{1.5,1.5,1.5}

Referring to Table 4, a very significant reduction is observed in the size factors (m_i) compared to the reduction in the length factors (c_{is}). Consequently, the size factors do not owe their large change to the new length of the links, but mainly to the enlargement of the layout, which has led to a greater contribution of the $\vec{G}_i$ vectors.

Figure 9b presents the new (dimensionless) workspace and the new points studied, their coordinate being given in Table 5. The new dimensions and geometric proportions of the layout lead to a very significant reduction in the workspace by 92% (from $v_0 \simeq 5$ to $v_{\text{wide}} \simeq 0.8$). Therefore, case WLX becomes necessary to exploit the $\hat{x}$ -axis.

Table 5. Workspace points studied in the modified configuration.

End-Effector Vector	Points					
	A	B	C	D	E	F
R_x (mm)	1150.0	1320.0	2004.0	1104.0	1650.0	1450.0
R_y (mm)	500.0	280.0	346.7	357.8	300.0	500.0
R_z (mm)	1150.0	50.0	512.3	683.4	500.0	700.0

Static Analysis Results for the Modified Configuration

Table 6 presents the results for the two welding cases of the modified configuration. At first glance, a reduction in the loads is indeed observed for both cases in the modified configuration compared to the initial one. Again, case WLX is more demanding than case WLZ; however, the loads are now much more manageable than in the initial configuration. An additional improvement is also shown in the first natural frequency for both cases, which is due to the reduction in the length of the links, yet it is again marginally unsuitable regarding the 50 Hz limit.

Table 6. Static analysis results of the modified robot configuration for 6 end-effector coordinates and associated robot poses.

Property	Workspace Points											
	Case WLZ						Case WLX					
	A	B	C	D	E	F	A	B	C	D	E	F
Max von Mises—Rods (MPa)	10	10	10	10	10	10	10	20	40	30	30	30
Max Reaction Force—Bases (kN)	9.2	10.2	10.5	10.9	11.1	8.7	15.3	23.6	45.5	34.9	36.7	27.2
Max Reaction Moment—Bases (kNm)	13.1	11.1	6.8	7.4	5.4	11.7	21.9	35.4	65.2	49.6	52.7	38.5
Max Reaction Force—Clamping System (kN)	18.4	15.9	15.1	15.2	15.1	15.5	22.9	47.0	50.9	53.7	21.5	36.9
Max Reaction Moment—Clamping System (kNm)	8.1	1.1	4.1	5.0	4.0	5.2	9.8	13.9	21.2	23.6	8.1	14.4
Axial Relative Displacement—Tool (mm)	0.0	0.0	0.0	0.0	0.0	0.0	−0.1	−0.1	−0.2	−0.1	−0.2	−0.1
Side Relative Displacement—Tool (mm)	0.0	−0.0	−0.0	−0.0	−0.0	−0.0	−0.1	−0.1	−0.2	−0.5	−0.3	−0.0
First Mode (Hz)	33.9	53.0	44.5	42.0	46.0	37.5	54.5	35.7	27.8	38.9	23.7	53.0

Table 7 presents the comparison of the best, worst and mean values between the initial and the modified configuration for both cases. There is indeed an improvement in the modified configuration, as all load values show a decrease on average. The relative displacement is within the desired accuracy, even in case WLX, which confirms the increase in stiffness. The main frequency is significantly higher compared to the initial one, yet exceeding only in some cases the limit of 50 Hz. Thus, further reduction in the length of the links and introduction of damping elements are necessary.

Table 7. Configuration comparison (Org: original optimized, Mod: modified).

Properties	Case WLZ							Case WLX						
	Best		Worst		Mean			Best		Worst		Mean		
	Org	Mod	Org	Mod	Org	Mod	%	Org	Mod	Org	Mod	Org	Mod	%
Max von Mises (MPa)—Rods	10 (A)	10 (A)	40 (B)	10 (C)	30	10	76	30 (F)	10 (A)	130 (A)	40 (C)	70	30	60
Max Reaction Force (kN)—Bases	23.7 (F)	8.7 (F)	62.9 (C)	11.1 (E)	46.4	10.1	78	46.0 (F)	15.3 (A)	129.0 (B)	45.5 (C)	82.5	30.5	63
Max Reaction Moment (kNm)—Bases	6.5 (A)	7.4 (D)	69.5 (C)	13.1 (A)	36.9	9.3	75	18.3 (F)	21.9 (A)	102.0 (A)	65.2 (C)	53.6	43.9	18
Max Reaction Force (kN)—Clamping System	17.0 (A/B)	15.1 (C/E)	27.8 (D/E)	18.4 (A)	22.4	15.9	29	21.1 (D/E)	21.5 (E)	75.0 (F)	53.7 (D)	39.8	38.8	2
Max Reaction Moment (kNm)—Clamping System	2.0 (F)	1.1 (B)	26.0 (C)	8.1 (A)	11.9	4.6	62	2.1 (B)	8.1 (E)	54.9 (A)	23.6 (D)	23.3	15.2	35
Axial Relative Displacement (mm)—Tool	0.0 (F)	0.0	−0.4 (A)	0.0	−0.2	0.0	100	−0.2 (F)	−0.1 (F)	−10.4 (B)	−0.2 (E)	−2.5	−0.1	96

Table 7. Cont.

Properties	Case WLZ							Case WLX						
	Best		Worst		Mean			Best		Worst		Mean		
	Org	Mod	Org	Mod	Org	Mod	%	Org	Mod	Org	Mod	Org	Mod	%
Side Relative Displacement (mm)—Tool	−0.5 (F)	−0.0 (B)	1.4 (A)	0.0 (A)	−0.4	−0.0	100	−0.3 (D/E)	−0.0 (F)	89.3 (B)	−0.5 (D)	22.3	−0.2	99
First Mode (Hz)	17.1 (B)	53.0 (B)	8.0 (F)	33.9 (A)	11.3	42.8	-	24.2 (D/E)	54.5 (A)	1.8 (B)	23.7 (E)	12.6	38.9	-

7. Discussion

The dimensionless model essentially serves to capture and increase the workspace. For the study of rigidity, it is obvious that different scales, i.e., values of Q , affect mass distribution, hence rigidity and natural frequencies. This is the reason why the relevant study was done for $Q = 1$ m, as a practical engineering choice to generate a realistic robot size compatible with commercially available components and typical FSW workspaces. For this reason, too, a modified configuration is subsequently investigated with shorter robot links and larger base spacing to restore stiffness and improve modal performance while sacrificing workspace, demonstrating that structural performance is governed primarily by dimensional scaling rather than the dimensionless optimization itself.

Although the modified configuration increases the first eigenfrequency, some poses still yield values close to the nominal spindle excitation frequency (~50 Hz at 3000 rpm). Therefore, resonance cannot be excluded solely by comparing eigenfrequencies, although damping at joints has not been modelled. In practice, vibration risk can be mitigated by (i) further stiffness-oriented tuning (shorter links, higher bending stiffness cross-sections, and increased joint stiffness), (ii) introducing damping at joints and/or at the spindle/end-effector interfaces to reduce resonance amplification, and (iii) planning the welding operations appropriately, i.e., selecting spindle speed ranges that avoid modal frequencies for the specific pose as well as trajectory/pose planning.

The robot configuration studied can indeed meet the requirements of FSW, especially by introducing a closed control loop, since the open loop performance is satisfactory. Note that the study was conducted for maximum loads corresponding to aluminum alloys. Consequently, the robot is expected to be capable of welding aluminum and other ductile materials such as brass alloys, mild steels and polymers.

Furthermore, by introducing linear guides in two independent links, the extension to 5 DoFs can be achieved, providing control for both tool and welding angles. Moreover, since this is a small range of angle change, short-stroke linear actuators are sufficient. This observation, in combination with the mean stress in the rods in the modified model, points to feasible solutions based on commercially available components.

Regarding the base supports, these can be improved by increasing the distances between the linear carriers or by reducing the length of the links, at the expense of the workspace. Linear carriers are widely used in practice involving high loads, as in this case. Furthermore, solutions can be designed that combine multiple carriages capable of compensating for lateral loads and overturning moments.

Modifications are made interactively to increase rigidity, starting from the kinematically optimum configuration that is universally applicable since it is dimensionless. Thus, the designer can introduce geometry modifications at will, which will reduce, of course, the workspace. Such a reduction may impact the size of the sheets, notably when the welding line is along the Z direction (case WLZ). Consequently, welding in the X direction seems indispensable.

Figure 12 shows an FSW example of two sheets with a non-flat welding line in the X direction. Thus, the robot workspace must suffice to cover the displacement range in the two other directions, i.e., the side and axial offsets. In the modified configuration, the maximum side offset is about 720 mm, whereas the maximum axial offset is about 810 mm. By comparison, in the original optimized configuration, the offsets would be 2800 mm and 2100 mm, respectively. These dimensions are extracted from the workspace as the largest straight line segments in Y and Z that can be drawn continuously between two points within the workspace. As for the X-axis, it can potentially cover a distance of typically up to 2m, depending on the availability of reliable robot components.




Figure 12. Sheet metal parts to be welded, showing side and axial offsets.

Regarding the implementation potential for such a robot, first, linear slides with appropriate carriages are a standard, commercially available solution. For the extension to 5 DoFs, short-stroke hydraulic actuators with high resistance to lateral loads can be used, their use probably requiring additional bases for proper support. A significant challenge is the implementation of the joints. U-Joints that are available on the market are mainly aimed at power transmission, and they cannot handle lateral bending loads. Ideally, a cross-shaped U-Joint should be used.

8. Conclusions

In this work, significant steps were taken to investigate the kinematics, the workspace and the robustness of the Tau-robot, thereby clarifying its suitability for FSW, which is a very demanding manufacturing application in terms of loads on the robot.

The robustness can be practically increased interactively by constraining the distances of the bases as well as reducing the length of the links, i.e., practically by reducing the workspace, starting from the kinematically ideal situation determined first and serving as the baseline.

To exploit the configuration as much as possible, the welding line should be parallel to the linear guides of the device base.

In cases of relatively small welding lines, the Tau-robot is capable of providing a robust and reliable solution, given that the accuracy already achieved from the open loop is sufficient. The limitation of the dimensions of workpieces to be welded refers to the so-called offsets of the welding line and not to the workpiece dimensions, as the robot

moves with respect to the welding line and not to the rest of the part's surface. The welding line is approximately aligned with the X-axis of the global coordinate system.

In the future, we aim to investigate closed-loop control for the Tau-robot and test it on a suitable prototype, initially at a smaller scale in order to prove the concept. In addition, a new optimization exercise might be introduced to balance mechanical strength, vibration immunity and workspace. Since a number of computational tools will be needed in such an exercise, one of the implementation challenges would be to integrate them on the same optimization platform.

Author Contributions: Conceptualization, G.-C.V. and D.A.; methodology, D.A.; software, D.A.; formal analysis, D.A.; resources, G.-C.V.; writing—original draft preparation, D.A.; writing—review and editing, G.-C.V.; supervision, G.-C.V.; project administration, G.-C.V. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author(s).

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Kaygusuz, E.; Karaomerlioglu, F.; Akinci, S. A review of friction stir welding parameters, process and application fields. *Turk. J. Eng.* **2023**, *7*, 286–295. [CrossRef]
2. Kilic, S.; Ozturk, F.; Demirdogen, M.F. A comprehensive literature review on friction stir welding: Process parameters, joint integrity, and mechanical properties. *J. Eng. Res.* **2025**, *13*, 122–130. [CrossRef]
3. Rai, R.; De, A.; Bhadeshia, H.K.D.H.; DebRoy, T. Review: Friction stir welding tools. *Sci. Technol. Weld. Join.* **2011**, *16*, 325–342. [CrossRef]
4. Cook, G.E.; Crawford, R.; Clark, D.E.; Strauss, A.M. Robotic friction stir welding. *Ind. Robot* **2004**, *31*, 55–63. [CrossRef]
5. Yue, W.; Liu, H.; Huang, T. An approach for predicting stiffness of a 5-DoF hybrid robot for friction stir welding. *Mech. Mach. Theory* **2022**, *175*, 104941. [CrossRef]
6. Cheng, B.; Dong, B.; Chen, S.; Xiao, J. Effect of process parameters during the insertion phase of robotic friction stir welding on axial force. *Int. J. Adv. Manuf. Technol.* **2025**, *137*, 951–963. [CrossRef]
7. Ibannai, A.I. Review of common defects in friction stir welding. *Int. J. Sci. Technol. Res.* **2020**, *9*, 318–329.
8. Guillo, M.; Dubourg, L. Impact and improvement of tool deviation in friction stir welding: Weld quality and real-time compensation on an industrial robot. *Robot. Comput.-Integr. Manuf.* **2016**, *39*, 22–31. [CrossRef]
9. Taghirad, H.D. *Parallel Robots: Mechanics and Control*; CRC Press: Boca Raton, FL, USA, 2025.
10. Mo, S.; Hu, Y.; Huang, T.; Yao, B.; Chen, S.; Huang, Y.; Zhang, W. Dynamic modeling and performance analysis of a new friction stir welding parallel robot. *Mech. Mach. Theory* **2025**, *210*, 106030. [CrossRef]
11. Zhao, J.; Duan, Y.; Xie, B.; Zhang, Z. FSW robot system dimensional optimization and trajectory planning based on soft stiffness indices. *J. Manuf. Process.* **2021**, *63*, 88–97. [CrossRef]
12. Chai, X.; Zhang, N.; He, L.; Li, Q.; Ye, W. Kinematic sensitivity analysis and dimensional synthesis of a redundantly actuated parallel robot for friction stir welding. *Chin. J. Mech. Eng.* **2020**, *33*, 1. [CrossRef]
13. Luo, H.; Zhao, F.; Guo, S.; Yu, C.; Liu, G.; Wu, T. Mechanical performance of friction stir welding robots for aerospace applications. *Int. J. Adv. Robot. Syst.* **2021**, *18*, 1729881421996543. [CrossRef]
14. Luo, H.; Wu, T.; Fu, J.; Zhao, F. Analysis of typical working conditions and experimental research of friction stir welding robot for aerospace applications. *Proc. Inst. Mech. Eng. Part C J. Mech. Eng. Sci.* **2021**, *235*, 1045–1056.
15. Sun, T.; Wu, H.; Lian, B.; Qi, Y.; Wang, P.; Song, Y. Stiffness modeling, analysis and evaluation of a 5-DoF hybrid manipulator for friction stir welding. *Proc. Inst. Mech. Eng. Part C J. Mech. Eng. Sci.* **2017**, *231*, 4441–4456. [CrossRef]
16. Mousavi, S.; Gagnol, V.; Bouzgarrou, B.C.; Ray, P. Model-based stability prediction of a machining robot. *Mech. Mach. Sci.* **2017**, *46*, 379–387.
17. Chalak Qazani, M.R.; Pedrammehr, S.; Abdi, H.; Nahavandi, S. Performance evaluation and calibration of Gantry-Tau parallel mechanism. *Iran. J. Sci. Technol. Trans. Mech. Eng.* **2020**, *44*, 1013–1027. [CrossRef]
18. Akbari, M.; Esfandiari, M.; Abdollahzadeh, A. Role of force and torque in friction stir welding: A detailed review. *J. Adv. Join. Process.* **2025**, *11*, 100289. [CrossRef]

19. Prabhakar, D.A.P.; Korgal, A.; Shettigar, A.K.; Herbert, M.A.; Chandrashekhara, M.P.G.; Pimenov, D.Y.; Giasin, K. Review of optimization and measurement techniques of the friction stir welding process. *J. Manuf. Mater. Process.* **2023**, *7*, 181. [CrossRef]
20. Yan, W.; Yu, Y. Study on force/position hybrid control strategy for eight-axis robotic friction stir welding. *Metals* **2025**, *15*, 442.
21. Walz, D.; Sommer, K.; Werz, M.; Weihe, S. Development and evaluation of a robot-guided friction stir welding gun. *Int. J. Precis. Eng. Manuf.-Green Technol.* **2025**, *12*, 1795–1808. [CrossRef]
22. Kolegain, K.; Leonard, F.; Zimmer-Chevret, S.; Attar, A.B.; Abba, G. Feedforward deflection compensation coupled with offline path planning for robotic friction stir welding. *IFAC-PapersOnLine* **2018**, *51*, 728–733. [CrossRef]
23. Chen, S.; Zong, G.; Kang, C.; Jiang, X. Digital twin virtual welding approach for robotic friction stir welding based on FEA–robot co-simulation. *Sensors* **2024**, *24*, 1001.
24. Zong, G.; Kang, C.; Chen, S.; Jiang, X. Optimization of installation position for complex space-curve weldments in robotic friction stir welding using dynamic dual particle swarm optimization. *Processes* **2024**, *12*, 536. [CrossRef]
25. Patel, V.; Wouters, H.; Baghdadchi, A.; De Backer, J.; Igestrand, M.; Azimi, S.; Andersson, J. Robotic friction stir welding in lightweight battery assembly of extrusion-cast aluminum alloys. *J. Adv. Join. Process.* **2023**, *8*, 100156. [CrossRef]
26. Silva-Magalhaes, A.; Cederqvist, L.; De Backer, J.; Håkansson, E.; Ossiansson, B.; Bolmsjö, G. Friction stir welding case study using temperature-controlled robotics with HPDC cylinder blocks and dissimilar materials. *J. Manuf. Process.* **2019**, *46*, 177–184. [CrossRef]
27. Jin, T.; Cui, X.; Qi, C.; Yang, X. Development of a nonrigid-support friction stir welding repair robot for aluminum alloy trains. *Ind. Robot* **2024**, *51*, 490–500. [CrossRef]
28. Raghulapadu, J.K.; Peddieson, J.; Buchanan, G.R.; Nunes, A.C. Rotating plug model of friction stir welding heat transfer. *Heat Transf. Eng.* **2008**, *29*, 321–327. [CrossRef]
29. Ikumapayi, O.M.; Akinlabi, E.T.; Majumdar, J.D. Review on thermal, thermo-mechanical, and thermal stress distribution during friction stir welding. *Int. J. Mech. Eng. Technol.* **2018**, *9*, 534–548.
30. Christy, J.V.; Mourad, A.H.I.; Sherif, M.M.; Shivamurthy, B. Recent trends in friction stir welding of aluminum alloys and aluminum matrix composites. *Trans. Nonferrous Met. Soc. China* **2021**, *31*, 3281–3309. [CrossRef]
31. Takahara, H.; Motoyama, Y.; Tsujikawa, M.; Oki, S.; Chung, S.W.; Higashi, K. Allowance of deviation and gap in butt joints during friction stir welding. *Adv. Mater. Res.* **2007**, *15–17*, 375–380.
32. Zimmer-Chevret, S.; Jemal, N.; Langlois, L.; Ben Attar, A.; Hatsch, J.; Abba, G.; Bigot, R. FSW process tolerance according to tool position and orientation: Requirements for production system design. *Mater. Sci. Forum* **2014**, *783–786*, 1820–1825. [CrossRef]
33. Arian, A.; Danaei, B.; Abdi, H.; Nahavandi, S. Kinematic and dynamic analysis of the Gantry-Tau 3-DoF translational parallel manipulator. *Appl. Math. Model.* **2017**, *51*, 217–231. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Article

Generalized Design Methodology for Dual-Arm Robotic Platforms: From Conceptualization to Experimental Validation Within the MANiBOT Framework

Mario Peñacoba-Yagüe * and Jesús Enrique Sierra-García *

Department of Digitalization, University of Burgos, 09001 Burgos, Spain

* Correspondence: mpenacoba@ubu.es (M.P.-Y.); jesierra@ubu.es (J.E.S.-G.)

Abstract

This work proposes a general methodology for the design and experimental validation of dual-arm robotic platforms intended for intelligent manipulation tasks in real-world environments. The proposed framework formalizes the complete engineering process, from the definition of functional requirements to the structural validation of the final prototype, ensuring reproducibility and adaptability across different applications. The methodology is organized into five main stages: (i) requirement analysis and context characterization; (ii) conceptual architecture definition; (iii) detailed mechanical design and structural analysis; (iv) prototype construction and integration; and (v) experimental validation and iterative refinement. Each stage defines its expected deliverables, evaluation metrics, and decision criteria to support systematic design progression. The approach is demonstrated through its implementation within the European project MANiBOT, where the framework guided the development of a modular bimanual robotic platform capable of integrating collaborative manipulators and conveyor subsystems for dual-arm manipulation. Structural testing, deflection measurements, and stability analyses confirmed the robustness and safety of the resulting design. Beyond this specific case, the proposed methodology provides a replicable and extensible design reference for research and industrial teams developing modular robotic structures, supporting the standardization of engineering practices in bimanual mobile robotics.

Keywords: methodology; mechanical design; modular robotic platform; bimanual manipulation; collaborative robotics; mobile manipulators

1. Introduction

Although robotic perception, cognition, and control have progressed considerably in recent years, current collaborative robots remain constrained in terms of physical capabilities, especially when benchmarked against human manipulation skills. These limitations are especially evident in tasks that demand safe, efficient, and robust manipulation of objects in real-world, human-shared environments. Industrial robots, on the other hand, deliver speed and precision primarily in controlled, predefined settings where object properties and task conditions are fully modeled, and human presence is minimized. Such constraints hinder the deployment of collaborative robotic systems in sectors such as retail logistics, warehousing, and airport operations, where environments are dynamic, objects are heterogeneous, and task requirements are not always predefined in advance [1].

Humans, by contrast, manage complexity through an exceptional combination of perceptual awareness, adaptive reasoning, and coordinated motor actions, enabling fluid

responses to uncertain or changing environments [2]. Everyday manipulation tasks rely not only on dexterous bimanual coordination, but also on the use of non-prehensile strategies—such as pushing, sliding, or rotating—and the seamless integration of multimodal sensory feedback [3,4]. These capabilities allow humans to adapt in real time to unexpected changes, deformable or occluded objects, and spatial constraints. Replicating such levels of adaptability and robustness in service robotics remains one of the most pressing challenges in the field [5,6].

In this context, research has increasingly explored the integration of dual-arm configurations on mobile robotic bases to approximate human-like versatility. However, the mechanical embodiment of such systems remains underexplored [7]. Most studies emphasize perception, control, or learning aspects, while the design methodology of the physical platform itself—which governs stiffness, reach, balance, and modularity—has rarely been formalized. As a result, many prototypes are developed in isolation, lacking common reference metrics, reproducibility, or design traceability [8]. Establishing standardized design workflows would allow future researchers to evaluate, replicate, and benchmark new robotic architectures with greater scientific consistency.

To address this gap, there is a growing need for methodologies that standardize the design of robotic platforms capable of bimanual manipulation. While numerous research efforts have focused on improving perception and control, comparatively less attention has been devoted to the systematic mechanical design process that underpins the stability, reachability, and safety of dual-arm robotic structures [9]. The absence of shared guidelines often leads to ad hoc designs that are difficult to reproduce or adapt across different domains.

In response to safety and reliability concerns, the design of collaborative robotic systems is governed by international standards such as [10–12], which define essential safety requirements, interaction modes, and force or pressure limits for human–robot collaboration. These standards play a fundamental role in ensuring safe operation and regulatory compliance. However, they are intentionally technology-agnostic and prescriptive in nature: they specify what must be fulfilled, but do not address how complex robotic platforms—such as mobile dual-arm systems—should be systematically designed, dimensioned, and integrated from a mechanical and system-level perspective.

In particular, existing standards do not provide guidance on (i) structured design workflows for mobile bimanual platforms, (ii) the integrated mechanical embodiment of dual manipulators mounted on autonomous mobile robots, (iii) the adaptation of the platform architecture to specific operational use cases, (iv) modular and detachable designs enabling the reuse of heterogeneous robotic arms and mobile bases, or (v) practical validation strategies illustrated through complete, real-world systems.

This work contributes to filling that gap by proposing a structured, reproducible, and domain-independent methodology for the design and validation of dual-arm robotic platforms. The methodology formalizes each step of the mechanical development process—from requirement identification to structural validation—specifying the deliverables, evaluation metrics, and decision criteria that ensure compliance with functional and safety constraints. This provides a scientific foundation for engineering future dual-arm systems with predictable performance and documented design rationale.

The approach is demonstrated through its application to the design of the MANiBOT platform [13], a dual-arm robotic structure tailored for real-world retail and airport environments.

The structure of the paper is as follows: Section 2 reviews related work and the technological background. Section 3 defines the robotic platform design methodology. Section 4 describes the platform operational scenarios. Section 5 details how the design methodology is applied to the mechanical requirements derived from these scenarios. Section 6 describes

the tests performed on the robotic platform. Section 7 presents the experimental results, while Section 8 concludes with reflections and future research directions.

2. State-of-the-Art

The development of mobile dual-arm robotic platforms represents a rapidly evolving frontier in robotics engineering, combining advances in mechanical design, system integration, and validation methodologies. Multi-arm systems extend manipulation capabilities beyond single-arm baselines by enabling cooperative lifting, coordinated handovers, stabilization of large or deformable objects, and non-prehensile strategies [14]. Recent studies have reported significant gains in efficiency and adaptability when coordinated planning is combined with dual-arm execution [15,16]. However, these advantages introduce new challenges related to kinematic redundancy, inter-arm synchronization, collision avoidance, and structural rigidity—all of which are strongly influenced by the mechanical architecture that supports the manipulators.

Early developments focused on fixed dual-arm cells designed for assembly or laboratory automation, demonstrating the feasibility of high-precision bimanual manipulation under controlled conditions [17,18]. Subsequent research extended these systems toward mobile or semi-mobile embodiments capable of performing navigation, inspection, and telemanipulation tasks [19,20]. Collectively, these works underline a persistent design trade-off: the supporting frame must be compact and AMR-compatible, yet sufficiently stiff and well-balanced to withstand the forces and dynamic interactions generated by dual manipulators operating in confined or human-shared environments [21,22].

To explore this balance between stability and adaptability, many research groups have developed prototype platforms based on lightweight aluminum profiles, which accelerate early iterations and reduce fabrication costs [23]. While suitable for validating motion coordination and control concepts, such prototypes often lack the rigidity and endurance required for real industrial deployment. Consequently, more advanced designs have incorporated steel-reinforced or hybrid frames, improving stiffness, vibration resistance, and fatigue performance [24]. In parallel, several studies have integrated conveyors, linear axes, and height-adjustable mechanisms to extend workspace coverage and enable multi-height object transfer. These additions increase task versatility but also complicate mass distribution, modularity, and mechanical integration, leading most reported systems to remain task-specific prototypes rather than reusable, documented design frameworks.

In parallel to research-driven developments, the design of industrial and collaborative robotic systems has been strongly influenced by international safety standards such as [10–12], which define mandatory requirements for safe operation, human–robot interaction, and system integration. These standards constitute an essential reference framework and are widely adopted in both industrial practice and applied research. However, by design, they focus on compliance and risk mitigation rather than on mechanical engineering methodology. As such, they do not provide guidance on the systematic mechanical design of mobile dual-arm platforms, nor on architectural decisions related to arm placement, structural dimensioning, modularity, or the integration of task-specific subsystems such as conveyors or lifting mechanisms. Consequently, while standards ensure safety boundaries, they leave open the question of how complex dual-arm robotic platforms should be designed, documented, and validated as reproducible engineering artifacts.

Across the literature, existing works consistently show that the success of dual-arm systems depends as much on methodical mechanical design as on control and sensing [25]. Yet despite numerous case studies, there is still no unified, reproducible methodology that guides the design, construction, and validation of dual-arm robotic platforms as engineering artifacts. In particular, the community lacks a framework capable of: (i) structuring requirements into

functional and operational categories; (ii) translating them into architectural decisions such as arm placement, conveyor geometry, and maintainability features; (iii) defining standardized analysis and test artifacts (e.g., CAD/FEA models, stability and reachability envelopes, vibration and multi-height tests); and (iv) establishing quantitative acceptance criteria and traceability from requirement to experimental evidence.

To make this gap explicit, Table 1 provides a concise version of the comparative analysis of representative dual-arm robotic platforms reported in the literature, focusing on whether and how design methodologies are formalized, their degree of generality, and their treatment of mobility, adaptability, and validation. This comparison highlights that most existing works prioritize task- or application-specific solutions, while lacking a unified and reproducible mechanical design framework. The complete and detailed comparison is provided in Appendix A.1 (Table A1).

Table 1. Comparative analysis of dual-arm robotic platforms with respect to design methodology, mobility, and adaptability.

Reference	Platform Design Methodology	Arm Configuration	Mobility (AMR/Mobile Base)	Robot-Agnostic (Adaptability)	Main Focus
Peñacoba and Sierra-García (This work)	Yes	Dual-arm	Yes	Yes	Standardization of mechanical engineering
Li et al. (2023) [21]	No	Multi-arm	Yes	No	Task planning and perception
Park et al. (2014) [17]	No	Dual-arm	Yes	No	Electronic device assembly
Di Castro et al. (2017) [19]	Partial	Dual-arm	Yes	Yes	Control architecture
Chen and Hsiao (2023) [22]	No	Dual-arm	No	No	Joint modules
Duan et al. (2023) [25]	No	Dual-arm	No	Partial	Visual monitoring and digital twin
Moore et al. (2007) [18]	No	Dual-arm	No	No	Flexible programming

In response to this gap, the present work proposes a structured and reproducible design methodology for mobile dual-arm robotic platforms—spanning from requirement identification through conceptual and detailed mechanical design, prototype construction, and experimental validation.

3. Generalized Design Framework for Dual-Arm Robotic Platforms

This section introduces a generalized, requirement-driven design framework for the development of dual-arm robotic platforms, explicitly structured around an iterative reduction in design abstraction. The proposed methodology is applicable to bimanual robotic systems intended for industrial, service, or research environments, and aims to ensure structural robustness, functional compliance, safety, and adaptability across diverse operational scenarios.

Unlike purely sequential design approaches, the proposed framework formalizes the design process as an iterative loop governed by abstraction levels, where successive design proposals progressively evolve from high-level concepts to fully validated physical prototypes. The overall workflow is illustrated in Figure 1.

The methodology starts from an explicit Requirement List (RL) derived from use cases and contextual constraints, and iteratively generates and validates design proposals (DP_i), mechanical designs (MD_i), and physical prototypes (P_i), until the abstraction level is fully resolved and the final prototype (P_0) is obtained.




Figure 1. Design methodology workflow.

Each stage of the framework is defined in terms of inputs, outputs, and decision criteria, ensuring traceability and reproducibility throughout the design lifecycle.

Each stage is defined in terms of inputs, outputs, and design objectives, in order to facilitate its application to different robotic platforms.

Stage 1. Requirement Identification and Context Analysis

The framework starts with the system definition stage, where both functional and operational requirements are collected and structured.

Inputs: target use cases, environmental constraints, expected payloads, safety standards, and integration requirements.

Outputs: a comprehensive requirement specification document (RL) that defines design limits and performance indicators.

Key aspects include:

- I. Defining mission objectives and target tasks.
- II. Quantifying load, stiffness, workspace, and reach requirements.
- III. Establishing safety margins following ISO 10218 and ISO/TS 15066.
- IV. Identifying interoperability constraints with AMRs, conveyors, or perception systems.

Stage 2. Conceptual Architecture Definition.

This stage transforms abstract requirements into a conceptual architecture of the robotic platform.

Inputs: requirement list (RL) and environmental constraints.

Outputs: Conceptual architectural definition and an initial set of design proposals $DP_{AL_{max}}$ corresponding to the highest abstraction level, AL_{max} .

At this stage, the system's global architecture is outlined. The main objectives are to:

- I. Select candidate robotic manipulators and determine their relative positioning, ensuring workspace overlap and minimizing interference.
- II. Define the support structure geometry, including height, footprint, and accessibility.
- III. Select preliminary materials and joining strategies (e.g., extruded aluminum profiles for modularity, welded steel for strength).
- IV. Explore modular configurations for future adaptability.

Conceptual evaluation tools, digital mock-ups, reachability maps, or ergonomic analyses are recommended to validate feasibility before moving to detailed design.

Stage 3. Detailed Mechanical Design and Structural Analysis.

The core of the methodology lies in an iterative mechanical design loop, explicitly governed by the abstraction level AL_i .

For a given iteration i :

Inputs: design proposal (DP_i) for the current abstraction level, AL_i

Outputs: mechanical design instance (MD_i).

Using CAD and simulation tools, the design proposal is translated into a consistent mechanical model by:

- I. Optimize mass distribution and center of gravity for stability.
- II. Conduct finite element analyses (FEA) to assess stiffness and stress concentration, and real testing when feasible.
- III. Design of subsystems such as conveyors, lifting mechanisms, or cable routing paths.
- IV. Definition of a fixed space for sensors, actuators, and computing units.

If the resulting mechanical design does not satisfy design constraints or validation criteria, design changes are introduced, generating an updated design proposal (DP_i) and restarting the iteration at the same abstraction level.

Stage 4. Prototype Construction and Assembly.

The mechanical design (MD_i) at the current abstraction level is used to construct the corresponding physical prototype, P_i .

Inputs: mechanical design (MD_i).

Outputs: physical prototype (P_i).

Prototyping activities include:

- I. Manufacturing and assembly of full or partial prototypes, depending on the abstraction level.
- II. Select materials based on trade-offs between stiffness, weight, and cost.
- III. Verification of manufacturability and assembly tolerances.
- IV. Verify the integration of electrical, pneumatic, and communication systems within the mechanical frame.

Stage 5. Validation and Iterative Refinement.

Each prototype (P_i) undergoes a structured validation process.

Inputs: physical prototype (P_i).

Outputs: Validation outcome (Yes/No) and new design proposal for design improvement or abstraction reduction, DP_i .

Typical evaluations include:

- I. Static and dynamic load testing to ensure the frame remains within elastic limits [26,27].

- II. Vibration and deflection analysis under representative manipulation motions [28,29].
- III. Stability and rollover assessments, verifying that the center of gravity remains within the support polygon [30,31].
- IV. Functional testing with manipulators and payloads to confirm accessibility and safety [32].

If validation is successful, a decision is made regarding the abstraction level:

- If $AL_i > 0$, the abstraction level is reduced ($i = i - 1$), and a new design proposal (DP_{i-1}) is generated with increased design fidelity.
- If $AL_i = 0$, the design process converges to the final validated prototype (P_0).

If validation fails, the process loops back to design proposal modification without changing the abstraction level.

3.1. Quick Reference Guide

To facilitate reproducibility and promote methodological consistency, the following structured checklist summarizes the key checkpoints for each design phase. Each design stage is accompanied by its corresponding evaluation goals and a set of guiding questions. This list helps designers to trace requirements, assess design maturity, and ensure completeness across all development stages of the robotic platform.

1. Requirements

Goals:

Establish measurable functional and safety boundaries that guide all subsequent design decisions; formalize operational and environmental assumptions for reproducibility.

Guiding Questions:

- Have the payload, reach, and stiffness targets been quantitatively specified?
- Are cycle time, positioning accuracy, and repeatability requirements defined?
- What environmental conditions (temperature, vibration, humidity) must the structure withstand?
- Which international standards and safety directives (ISO 10218-2, ISO/TS 15066, CE marking) apply?
- Are human–robot interaction constraints, such as contact forces and protective distances, clearly established?
- Has interoperability with AMRs, conveyors, or external controllers been considered?
- Are measurable Key Performance Indicators (KPIs) defined for later validation?

2. Conceptual Design

Goals:

Define a coherent mechanical architecture that balances reachability, compactness, and modularity while ensuring ergonomic compatibility and task feasibility.

Guiding Questions:

- Are the manipulators adequately selected for the payload, reach, and dexterity required by the target tasks?
- Do both arms achieve sufficient workspace overlap for cooperative manipulation?
- Is the structure compact, accessible, and modular for the defined tasks?
- Is the geometry compact enough for navigation and accessibility?
- Are accessibility and maintenance zones clearly defined?
- Are anthropomorphic considerations (height, reachability for human co-workers) integrated?
- Does the proposed concept allow modular extensions or reconfiguration for future use cases?

3. Material and Structural Design

Goals:

Validate that the mechanical structure satisfies stiffness, strength, stability, and longevity requirements; confirm that material selection supports durability, modularity, and compliance with safety factors.

Guiding Questions:

- What material combinations best balance stiffness, mass, and manufacturing cost?
- How will loads and vibrations be absorbed and distributed?
- Is the center of gravity always inside the support polygon?
- Are welds, bolted joints, or modular extrusions appropriately dimensioned?
- Has finite element analysis (FEA) confirmed that stresses remain below allowable limits under static and dynamic loads?
- Have fatigue, wear, and corrosion effects been considered for long-term use?
- Is the design compliant with industrial safety factors (>1.5 under maximum load)?

4. Integration

Goals:

Demonstrate manufacturability, maintainability, and interoperability; guarantee that the assembled system can be safely operated, serviced, and upgraded with minimal redesign.

Guiding Questions:

- How are cables, power lines, and controllers routed and protected?
- Are the power, data, and pneumatic/hydraulic lines safely routed and accessible for maintenance?
- Are electromagnetic compatibility (EMC) and grounding strategies incorporated?
- Have standardized electrical and mechanical interfaces been defined for sensors, actuators, and controllers?
- Are accessibility and safety preserved during assembly and disassembly?
- Have ergonomic factors for operators (height, reach, tool access) been reviewed?

5. Validation

Goals:

Provide experimental evidence of compliance with design specifications; close the design loop through quantitative correlation between analytical models and empirical data, ensuring readiness for real-world deployment.

Guiding Questions:

- Does the structure remain within the elastic range under maximum static and dynamic loading?
- Has rollover stability been tested for the most critical arm configurations?
- Can the platform be safely lifted, transported, and mounted on different AMR bases if needed?
- Do experimental deflection measurements remain within the specified limits?
- Are test logs, torque monitoring, and deformation data recorded for traceability?

3.2. Recommended Tests

To ensure that robotic platforms meet the required standards of mechanical robustness, stability, and operational safety, a set of standardized tests is recommended prior to prototype deployment. These tests, summarized below, provide a methodological framework applicable to any modular robotic structure intended for mobile manipulation.

1. Static Load Tests

Objective: Evaluate the capacity of the structural frame to support its nominal payload without exceeding elastic deformation limits.

Methodology: Apply uniformly distributed and concentrated loads corresponding to 1.0–1.5 times the rated payload for periods of 24 h (short-term) and 7–20 days (long-term).

Acceptance criteria: The maximum allowable deflection is defined by the classical beam design criterion (Equation (1)), widely adopted in mechanical and structural engineering handbooks such as Shigley's Mechanical Engineering Design [27] and Mechanics of Materials [33]:

$$\delta_{max} = \frac{L}{250} \quad (1)$$

Here, δ_{max} is the maximum deflection tested, and L is the length of the platform. No yielding or permanent deformation shall occur.

Typical instrumentation: calibrated torque wrench (25–50 Nm range) and dial gauges (0.01 mm precision).

2. Dynamic Load Tests

Objective: Assess the structural integrity of joints and fasteners under cyclic or vibration-induced stresses that simulate real operating conditions.

Methodology: Subject the loaded platform to repeated acceleration-deceleration cycles (0.2–0.5 g) on ground-supported and AMR-mounted configurations, monitoring torque variation before and after testing.

Acceptance criteria: Torque variation is evaluated using the criterion shown in Equation (2) [34].

$$\frac{\tau_{after} - \tau_{before}}{\tau_{before}} \times 100 \leq 5\% \quad (2)$$

And no crack initiation or residual displacement was observed.

It should be noted that the relative importance of static and dynamic analyses strongly depends on the intended operational profile of the robotic platform. For applications involving slow, infrequent, and quasi-static manipulation tasks, static load, deflection, and material yielding analyses typically represent the dominant design constraints, as gravitational and payload-induced loads prevail over inertial effects.

Conversely, for platforms intended to perform fast, repetitive, or highly dynamic motions, inertial forces associated with rapid accelerations and decelerations of the manipulators and payload may significantly increase the stresses transmitted to the supporting structure. In such cases, dynamic effects should be explicitly considered as part of the recommended test methodology, either through time-dependent dynamic simulations or by conservatively incorporating inertial loads derived from worst-case trajectories into the structural analysis, as discussed in Section 6.

3. Deflection Tests

Objective: Quantify structural stiffness by measuring displacement under static loading.

Methodology: Measure vertical deflection at critical points using dial gauges while applying distributed and concentrated loads equivalent to the nominal payload.

Acceptance criteria: It must be the same as in the recommended test 1 (Equation (1)).

4. Material Yielding Analysis

Objective: Verify analytically that induced stresses remain below the elastic limit of the material under the most demanding loading conditions.

Methodology: Compute bending moments and section modulus of the main beams; compare induced stress with the yield strength of the material, applying a recommended safety factor $SF \geq 1.5$.

Acceptance criteria: According to Navier's law for simple flexion and applying a safety factor of 1.5, the acceptance criteria are shown in Equations (3) and (4).

$$\sigma = \frac{My_{max}}{I} \quad (3)$$

$$\sigma < \frac{\sigma_y}{SF} \quad (4)$$

5. Lifting System Verification

Objective: Validate the mechanical safety of the lifting or docking systems used to transfer the platform between AMRs or maintenance stations.

Methodology: Apply analytical calculations for maximum bending moment on lifting bars and verify experimentally by gradually applying static loads equivalent to $SF \times SW$. Being SW the self-weight of the structure.

Acceptance criteria: The maximum bending stress on the lifting bars is calculated in Equations (5) and (6) following standard strength-of-materials relations [27].

$$\sigma_y = \frac{32M_{max}}{\pi d^3} \quad (5)$$

$$\sigma_{max} = \frac{\sigma_y}{SF} \quad (6)$$

No permanent deformation shall occur.

6. Rollover and Stability Tests

Objective: Confirm that the combined center of gravity (CoG) of the platform and manipulators remains within the support polygon under all operational configurations.

Methodology: Determine CoG position via CAD and physical measurements; perform tilt tests by incrementally inclining the base until the onset of tipping; record the critical angle.

Acceptance criteria: The center of gravity projection must remain inside the polygon formed by the supporting legs or AMR footprint [32].

7. Fastening and Torque Retention Tests

Objective: Verify the reliability of bolted or modular joints under repeated load cycles and vibration.

Methodology: Record torque before and after 100 load cycles using calibrated torque sensors; evaluate any loss exceeding 5% as a sign of insufficient preload or thread relaxation.

Acceptance criteria: Equation (2) must be verified, as well as no visible loosening, slippage, or cracking at joint interfaces.

The above test set provides a comprehensive methodological framework covering the most critical aspects of mechanical validation: static and dynamic loading, deflection, yielding, lifting, stability, fastening reliability, and environmental robustness. These procedures collectively ensure that the robotic platform satisfies the structural and safety criteria required for real-world operation, in alignment with recognized mechanical engineering standards and current literature.

4. Use Case: Design of Bimanual Robotic Platform Within the MANiBOT Project

4.1. MANiBOT European Project

The proposed framework is validated through its implementation within the European project MANiBOT [13], which aims to advance the physical and cognitive capabilities of collaborative mobile manipulators for real-world applications. MANiBOT serves as an

exemplary context to demonstrate the methodology in practice, encompassing demanding scenarios such as:

- Autonomous shelf restocking in supermarkets, involving bimanual manipulation in human-populated and confined spaces.
- Airport baggage-handling operations, requiring robust, safe, and efficient loading and unloading between conveyor belts and carts.

To this aim, it is critical to design a flexible robotic platform able to transport collaborative robots for dexterous bimanipulation tasks [9].

4.2. Functional Requirements of the Robotic System

The design of the MANiBOT platform was driven by a set of functional requirements derived directly from the operational scenarios defined for supermarket restocking and airport baggage handling. Although these domains differ in layout and workflow, both impose demanding mechanical constraints on the robotic system: manipulation of objects with varied mass and geometry, operation in confined and cluttered spaces, and physical interaction with shelves, carts, and conveyors.

From the analysis of these environments, the requirements essential for guiding the mechanical architecture of the platform were identified. These requirements focus exclusively on the physical capabilities, structural performance, and layout characteristics needed to ensure safe, stable, and reliable manipulation in real conditions. The critical functional requirements are the following:

1. **Dual-arm integration and coordinated operation:** the platform must support the installation of two collaborative manipulators capable of performing coordinated bi-manual tasks. The mechanical frame must provide rigid and precise mounting points that guarantee repeatable positioning and alignment.
2. **Adequate workspace and arm overlap:** the structure shall enable both manipulators to reach shared and individual workspaces, ensuring sufficient overlap for hand-over actions and cooperative manipulation. This includes maintaining clear, unobstructed areas around shelves, carts, and conveyor interfaces.
3. **Conveyor system accommodation:** the platform must incorporate a conveyor subsystem, mechanically integrated into the frame, enabling object transfer between the environment and both arms. The mechanical design must ensure stability during loading/unloading and allow configurations such as tilting or height variation.
4. **Structural rigidity and payload endurance:** the frame must maintain elastic, deformation-free behavior when subjected to payloads up to 11 kg per arm at maximum reach, including the additional forces arising from dynamic movements or bi-manual lifting.
5. **Stability when mounted on mobile bases:** the combined structure—manipulators, conveyor, and payload—must preserve static and dynamic stability when installed on different AMR platforms. This includes resisting tipping when the conveyor is inclined or when arms operate asymmetrically.
6. **Modularity, accessibility, and maintainability:** the mechanical design shall be modular, allowing easy assembly, reconfiguration, and servicing. The structure must include accessible panels, ventilation paths, and protected internal routing for cables and auxiliary systems.
7. **Dedicated internal space for equipment:** the platform must reserve adequate, mechanically protected internal volume for computing units, controllers, power electronics, and communication hardware. This includes ensuring organized cable routing and avoiding interference with arm workspaces.

8. **Geometric symmetry and balanced mass distribution:** the upper structure shall follow a symmetrical layout to ensure balanced load distribution, reduce torsional stresses on the frame, and simplify dynamic behavior when mounted on an AMR.

Collectively, these mechanical requirements establish the foundation for the platform's structural design and drive the engineering decisions taken throughout the conceptual and detailed design stages. They ensure that the MANiBOT platform is physically capable of performing the intended tasks in realistic retail and airport environments, supporting reliable bi-manual manipulation and stable interaction with shelves, carts, and conveyors.

4.3. Scenarios

To ground the functional requirements in real practice, the MANiBOT consortium identified four industrially significant and operationally demanding scenarios. These cover collaborative manipulation tasks in unstructured retail and airport domains, characterized by diverse object types, limited workspaces, and constant human presence. Each SC is broken down into detailed sub-scenarios (Sub-SCs), mapping the sequence of robotic actions required to accomplish the task.

4.3.1. Scenarios 1 and 2: Product Restocking in Retail Environments

The first two scenarios address automated product restocking in supermarket and pilot-shop settings. In both situations, the robot must operate in confined spaces, interact physically with a wide variety of items, and execute manipulation sequences that involve picking, orienting, and placing products onto shelves.

In Scenario 1, products are supplied on a conventional replenishment cart, while Scenario 2 extends this context to mixed pallets, where goods may arrive in heterogeneous packaging such as shrink-wrapped bundles or plastic crates. In both cases, the robot must correctly extract each item, manage its placement according to store-specific criteria (including FIFO policies when required), and perform the shelf arrangement in a stable and repeatable manner.

These scenarios emphasize the mechanical and operational demands placed on the platform: coordinated bi-manual actions, safe manipulation close to shelves and customer areas, and the ability to handle items of different sizes, orientations, and rigidity without requiring a redesign of the robotic structure. They serve as a benchmark to validate the platform's reach, stability, task execution workflow, and integration with real retail layouts (Figure 2).



Figure 2. Combined representation of Scenarios 1 (a,b) and 2 (c,d).

4.3.2. Scenarios 3 and 4: Airport Baggage Loading and Unloading

Scenarios 3 and 4 address complementary luggage-handling operations in an airport environment. Scenario 3 focuses on loading suitcases from a conveyor belt onto a transport cart, requiring the robot to interface physically with both fixed and mobile airport infrastructure while manipulating luggage of varying mass, geometry, and deformability. The robot must grasp each item securely, reorient it when required, and place it in a stable and ordered arrangement inside the cart, ensuring efficient use of space and proper load distribution. This operation serves to validate the structural rigidity of the platform, the effective workspace of both manipulators, and the coordination needed for bi-manual handling, particularly when managing large or heavy items.

Scenario 4 extends this evaluation by addressing the reverse operation: unloading luggage from the cart and depositing it onto a conveyor belt. In this phase, the robot must extract items from partially filled carts, avoid contact with surrounding structural elements, and ensure that each suitcase is positioned correctly before release. Additionally, the scenario includes moving between different carts and returning to the initial location, which helps assess the platform's compatibility with typical airport layouts and its ability to support uninterrupted task sequences.

Together, these two scenarios provide a comprehensive assessment of the platform's capacity to execute repetitive, physically demanding manipulation tasks in realistic airport workflows (Figure 3).




Figure 3. Representation of scenario 3 (a) and scenario 4 (b).

5. Application of the Design Methodology to the Use Case

This section demonstrates the practical implementation of the methodology presented in Section 3 through its application to the MANiBOT robotic platform. The process followed the cyclic workflow illustrated in Figure 1, evolving from requirement identification to the validation of a fully functional prototype. The test protocols and quantitative results of the validation campaign are presented in Section 7.

5.1. Stage I—Requirement Identification and Context Analysis

The MANiBOT platform was conceived to operate in two demanding, real-world domains: supermarket shelf restocking and airport baggage handling. From these scenarios, a set of functional and operational requirements was derived, forming the foundation of the design process.

In accordance with systems-engineering principles, the requirements were classified into Functional Requirements (FR) and Operational Requirements (OR).

Functional requirements describe what the system must do—that is, the intrinsic capabilities, behaviors, and design features needed to fulfill its intended tasks.

Operational requirements, in contrast, define under what conditions and constraints the system must operate, including environmental, regulatory, and maintainability aspects.

This distinction provides a clear structure for translating high-level mission objectives into verifiable engineering specifications, ensuring that both the internal functionality of the platform and its interaction with real-world environments are consistently addressed.

The MANiBOT robotic platform's Functional Requirements are as follows:

- I. **FR1—Dual-arm capability:** The platform shall host two collaborative manipulators enabling coordinated bimanual operation.
- II. **FR2—Workspace overlap:** The manipulators shall provide sufficient workspace overlap to enable cooperative tasks and safe handover operations.
- III. **FR3—Conveyor subsystem integration:** The structure shall include a conveyor system allowing coordinated object transfer between manipulators or external stations.
- IV. **FR4—Structural rigidity under payload:** The mechanical frame shall maintain elastic behavior under payloads up to **11 kg** per arm at maximum reach.
- V. **FR5—Stability on AMR base:** The system shall preserve static and dynamic stability when mounted on autonomous mobile robots (AMRs), including tilting conveyor configurations.
- VI. **FR6—Modularity and maintainability:** The platform shall support modular assembly and easy access for servicing, ventilation, and subsystem replacement.
- VII. **FR7—Allocation of internal components:** Adequate internal space shall be reserved for computing units, controllers, and sensors, ensuring organized cable routing and protection.
- VIII. **FR8—Symmetry and balance:** The overall design shall maintain a geometrically symmetrical configuration to simplify control, equalize loads, and minimize torsional stresses.

On the other hand, the following Operational Requirements for the MANiBOT robotic platform have been identified:

- I. **OR1—Spatial validation and AMR adaptability:** The footprint and mounting interface shall be compatible with multiple AMR bases (e.g., P204, P604) while preserving geometric stability.
- II. **OR2—Working height and task range:** The manipulators and conveyor shall jointly enable operation between **0.5 m and 2.5 m**, covering all expected manipulation heights.
- III. **OR3—Safety and regulatory compliance:** The design shall conform to collaborative-robot standards, including [11,12], ensuring safe human–robot coexistence.
- IV. **OR4—Maintainability and service access:** The platform shall feature accessible panels, guided cable routing, and proper ventilation to facilitate maintenance and minimize MTTR.
- V. **OR5—Integration of power and computing systems:** The platform shall provide sufficient energy capacity and internal integration for computing resources (NUCs, Jetsons, industrial PC), ensuring electromagnetic compatibility (EMC) and operational autonomy.

Together, these requirements define the baseline constraints and performance objectives guiding the conceptual, detailed, and validation stages of the MANiBOT platform, as summarized in Table 1 and later traced across successive design iterations.

5.2. Stage II—Conceptual Design of the Platform: Space Exploration and Initial Configuration

The mechanical development of the MANiBOT platform followed a stepwise, iterative refinement cycle explicitly aimed at aligning the structure with the demanding requirements of retail and airport scenarios. Over twelve main iterations, the system gradually evolved from a purely volumetric concept into a mature, modular, and robust solution ready for experimental validation.

In accordance with the conceptual architecture definition stage, this phase focused on transforming the initial system requirements and environmental constraints into a tangible structural concept. The process explicitly addressed the selection and positioning of manipulators, the geometric definition of the support frame, the preliminary choice of materials and joining strategies, and the exploration of modular configurations to ensure long-term adaptability. Conceptual tools such as digital mock-ups, reachability analyses, and ergonomic evaluations were applied to validate feasibility prior to detailed design.

Version 1—Conceptual volumetric mock-up.

The very first iteration (Figure 4) was conceived as a volumetric envelope corresponding to the P204 AMR footprint. At this stage, the focus was on defining the global geometry of the platform, evaluating height, footprint, and workspace accessibility (Objective II), taking into account that the manipulators will be on top (Objective I), or at least a common solution for the possible ones. Its purpose was not yet to host detailed components, but rather to provide an initial spatial framework for discussing the integration of manipulators, conveyor belts, and auxiliary modules. Reachability maps were generated to verify the potential workspace overlap between future manipulators. At this stage, the search for robotic manipulators was narrowed down to three possible candidates: CRB15000/95, CRB15000/127, and CRB15000/152.




Figure 4. Robotic Platform version 1.

This conceptual model enabled an early evaluation of the physical footprint, highlighted potential constraints for actuator placement, and served as a collaborative reference for aligning design perspectives across project partners.

Version 2—Adaptation to multiple AMR bases.

Building upon the volumetric concept, the second design iteration (Figure 5) addressed the need to ensure compatibility with different AMR geometries. The frame geometry was redefined to balance accessibility and stiffness, while an intermediate structure was introduced to guarantee ergonomic working heights (Objective II). A new intermediate frame was introduced as a structural interface, effectively elevating the upper platform while creating an intermediate compartment for cabling and electronics. This modular connection allowed flexible coupling to several mobile bases without compromising rigidity. At this stage, preliminary considerations for ventilation, power interfaces, and service panels were introduced, marking a transition from purely conceptual modeling to functional design. This version directly addressed, for the first time, the requirements of selecting preliminary materials and joining strategies (Objective III) and exploring adaptable configurations for future evolution (Objective IV). Future iterations will be the result of improving these objectives.

Version 3—Extended structure and internal reorganization.

The third iteration (Figure 6) incorporated additional components identified as essential for operability, which required a redefinition of the external dimensions and vertical organization of the platform. This design ensured that the relative positioning of the two

temporarily chosen manipulators (CRB15000/95) would provide sufficient workspace overlap while minimizing interference, in direct fulfillment of Objective I. The structure was extended in height, introducing a new upper level envisioned as a potential working surface or mounting point for manipulators and sensors. Internally, compartments were reorganized to improve accessibility to control and communication modules, simplifying future maintenance, while maintaining the modularity (Objective IV).




Figure 5. Robotic Platform version 2: with AMR P204 (a); with AMR P604 (b).




Figure 6. Robotic Platform version 3: with AMR P204 (a); with AMR P604 (b).

Version 4—Integration of manipulators and central conveyor.

The fourth design iteration (Figure 7) marked the first major step from conceptual modeling to functional prototyping. For the first time, active robotic components were introduced, with two ABB CRB15000/95 manipulators selected as a baseline following the earlier comparative analysis of potential arms (Objective I). Their installation allowed an initial evaluation of dual-arm feasibility, particularly regarding workspace overlap, potential interferences, and collaborative reach. In parallel, a central conveyor belt was incorporated as a placeholder subsystem, enabling early assessment of transfer operations between the manipulators or toward external stations. Although its dimensions were not yet finalized, the conveyor offered valuable insights into accessibility, interaction with moving objects, and workflow organization.

Structurally, this version consolidated anchoring reinforcements (Objective III) and introduced a detachable second level (Objective IV), partially elevating the upper platform to facilitate coupling with multiple AMRs. This modular capability transformed a key principle from earlier iterations into a tangible design feature, while also improving accessibility to control panels and ventilation systems.

Version 5—Virtual integration, structural reinforcement, and lifting system.

The fifth iteration (Figure 8) represented a substantial leap in maturity, as the MANI-BOT platform was integrated for the first time into a simulated airport baggage-handling environment. Within this virtual setup, manipulator reach, accessibility, and coordination

were validated against a conveyor belt and representative test objects such as suitcases. This virtual integration validated both the global geometry of the support structure and the manipulators' spatial coordination under realistic constraints (Objectives I and II).




Figure 7. Robotic Platform version 4: with AMR P204 (a); with AMR P604 (b).

From a mechanical perspective, the structure was redesigned using standardized extruded aluminum profiles, enhancing rigidity while maintaining modularity and compatibility with industrial fastening systems (Objective III). A critical innovation of this version was the introduction of the first lifting system, located at the lower part of the frame, which allowed safe detachment and transfer of the platform between different AMR bases. Together, these advances consolidated the platform as a modular, structurally robust prototype (Objective IV), bridging the gap between conceptual development and operational validation.




Figure 8. Robotic Platform version 5: operating mode in scenarios 3 and 4 (a); lifting system (b).

Version 6—Sheet-metal feasibility study.

The sixth design iteration (Figure 9) explored an alternative manufacturing pathway based on folded sheet-metal panels. This approach addressed Objective III by testing a different material, and therefore a new joining strategy, that minimized welding while embedding reinforcement features directly into the bending process. Unlike previous profile-based designs, this approach aimed to simplify production, reduce costs, and embed reinforcement features directly into the bending process, thereby minimizing welding and mechanical joints.

Functionally, the layout preserved the core architecture defined in earlier versions: elevated manipulators, a centrally positioned conveyor, front access to electronics, and compatibility with the lifting system. Simulation in the airport scenarios was maintained, ensuring continuity in evaluating interaction with objects and carts. Although conceived as an exploratory branch rather than the definitive design, this version demonstrated the feasibility of transitioning toward pre-industrial manufacturing while preserving all

essential functionalities. With this approach, the design lost some modularity (Objective IV was not as fully satisfied as in previous iterations), but it served to verify scalability and manufacturability.




Figure 9. Robotic Platform version 6.

Version 7—Linear system integration and transition to GoFa 12 manipulators.

The seventh design iteration (Figure 10) introduced two decisive advancements: the implementation of linear motion capabilities and the adoption of higher-performance manipulators. Manipulator positioning was revisited to confirm workspace overlap across extended trajectories (Objective I). Mechanically, this version refined the previously explored sheet-metal architecture while equipping one robotic arm with a linear track, substantially extending its operational envelope. This addition enabled the evaluation of tasks requiring longitudinal reach—such as handling distributed objects along a conveyor belt—and provided valuable insights into trajectory planning and kinematic constraints in extended workspaces. Material selection and stiffness calculations were reassessed, confirming the structural feasibility of supporting higher payloads (Objective III).




Figure 10. Robotic Platform version 7.

Simultaneously, the original GoFa 5 manipulator, which was not mounted on the linear track, was replaced with a GoFa 12. This upgrade, motivated by the need to address more demanding payload and reach requirements, allowed the team to perform more realistic simulations involving heavier objects and greater distances from the central platform. As a

result, this iteration effectively converged structural robustness, functional flexibility, and workspace analysis, marking a turning point toward readiness for physical prototyping and validation.

Version 8—Consolidated design with folded sheet-metal structure, linear tracks, and GoFa 12 arms.

The eighth iteration (Figure 11) represents the most advanced stage of the sheet-metal-based design pathway and served as the foundation for critical final decisions regarding both structure and manipulation hardware. At this stage, all four conceptual objectives were demonstrably satisfied: (I) optimal dual-arm positioning with minimal interference, (II) a well-defined support geometry balancing footprint, height, and accessibility, (III) material and joining strategies combining folded sheet-metal and bolted reinforcements, and (IV) a modular configuration adaptable to future AMRs and use cases.




Figure 11. Robotic Platform version 8.

Structurally, the geometry was carefully refined for efficient fabrication through bending processes, ensuring rigidity while maintaining lightweight properties. Mounting points, ventilation paths, and access to electronic compartments were systematically addressed, consolidating the design's maintainability and manufacturability. Functionally, the introduction of two linear tracks allowed each manipulator to extend its workspace significantly, supporting complex collaborative tasks and enabling simultaneous access to multiple zones around the central conveyor. Moreover, the linear motion system was re-engineered into a set of three coordinated elements—one actuator flanked by two supporting linear guides—that together absorb the moments generated at the base of the manipulators. Compared to a single heavy-duty linear track, this configuration provides a much more reliable and cost-effective solution, while still offering the stiffness required to withstand the significant loads induced during extended-reach operations.

The decision to adopt two GoFa 12 arms as the definitive manipulators confirmed the trajectory initiated in the previous iteration, providing the payload, precision, and reach required for the project's most demanding scenarios. Their integration with the linear tracks produced a highly versatile configuration, capable of addressing both retail and airport scenarios with increased adaptability and operational efficiency.

5.3. Stage III—Detailed Mechanical Design and Structural Analysis

In this stage, the previously defined conceptual architecture was refined into a fully parameterized and manufacturable mechanical system. Using CAD and simulation tools, the design team optimized mass distribution and center of gravity (Objective I), conducted finite element analyses to evaluate stiffness and stress (Objective II), performed dynamic and kinematic simulations to ensure collision-free operation (Objective III), detailed the

conveyor and lifting subsystems (Objective IV), and defined interfaces for sensors, actuators, and computing units (Objective V). The outcome of this stage was an optimized CAD model, a complete structural analysis, and an integration plan ready for prototyping.

Version 9—Modular architecture with aluminum profiles and optimized linear motion system.

The ninth design iteration (Figure 12) marked a strategic shift toward a fully modular configuration. Departing from the folded sheet-metal concept, this version adopted an extruded aluminum profile structure, enabling a platform that could be easily reconfigured to accommodate different use cases, environments, or future modifications without requiring extensive redesign. The use of standardized profiles provided multiple advantages: reduced weight, high mechanical strength, simplified assembly, and compatibility with widely available industrial fastening systems.




Figure 12. Robotic Platform version 9.

Within this framework, all critical subsystems were integrated for the first time in a coherent assembly: the two GoFa 12 manipulators, the central conveyor belt, and an improved linear track system. The latter was reinforced with dual linear guides, significantly enhancing stability under extended-reach maneuvers and dynamic loads. This redesign ensured accurate positioning of devices, maintaining the center of gravity at the midpoint of the platform by preserving symmetry (Objective I).

Equally important, this version introduced the first functional prototype of the conveyor subsystem, fully embedded into the aluminum frame (Objective III). While still subject to refinements in speed and geometry, it enabled preliminary experiments on coordinated manipulator–conveyor interaction, including luggage-handling scenarios representative of airport operations. In addition, the first reserved spaces for devices and interfaces were incorporated. These elements can be observed in Figure 12, located below the conveyor belt and between the two manipulators (Objective IV). Collectively, these improvements positioned version 9 as a robust, versatile, and highly modular prototype, ideally suited for integration testing and advanced control development.

Version 10—Airport trolley adaptation and conveyor tilting mechanism.

The tenth iteration (Figure 13a–d) addressed one of the most stringent requirements emerging from the airport scenarios: the need to interact with Fraport’s luggage transport carts. The incorporation of precise cart dimensions revealed that the conveyor system had to deliver objects at multiple heights, prompting a significant mechanical redesign.




Figure 13. Robotic Platform version 10: without cart (a); with cart (b); with conveyor belt in the highest height (c); with conveyor belt in the lowest height (d).

To meet this requirement, one longitudinal structural bar was removed to accommodate a lateral tilting mechanism for the conveyor (Objective III). This feature enabled controlled unloading at three distinct vertical levels, replicating real baggage-loading conditions inside the airport cart. The conveyor was designed to operate in both horizontal and inclined configurations, ensuring robust alignment with the manipulators' linear tracks and workspace envelopes.

Structural reinforcements were implemented to preserve stability despite the asymmetrical loading introduced by the inclined conveyor. These reinforcements were calculated through additional FEA simulations (Objective II), verifying that the asymmetric loading introduced by the tilted conveyor did not compromise global stiffness or safety margins. The coordination of the robotic arms with the newly angled surface was also validated through simulations, ensuring safe and precise handling at all levels.

All these changes made it necessary to redistribute the devices (Objective IV), and center-of-gravity adjustments were implemented (Objective I) to maintain stability despite the lateral tilt. With this functionality, version 10 bridged the gap between laboratory prototypes and the operational constraints of the airport environment.

Version 11—Physical validation and transition to steel-reinforced frame.

The eleventh iteration (Figure 14) represented a decisive milestone: the transition from purely digital validation to the first physical prototype. A partial assembly, built from aluminum profiles, was manufactured and tested to evaluate rigidity, stability, and dynamic response under simulated loads (Objective II). These experiments revealed critical limitations in structural stiffness and stability, particularly given the projected increase in total mass once both the manipulators and conveyors were fully integrated. Furthermore, range tests showed that the linear motion of the arms provided only marginal functional benefits relative to their high cost and limited applicability, leading to the strategic decision to discontinue their integration in subsequent designs.




Figure 14. Robotic Platform version 11.

Based on these findings, the design strategy shifted to a reinforced steel-tube frame, providing the robustness and safety margins necessary for real-world deployment. The new structure offered superior load-bearing capacity, improved vibration damping, and a significantly reduced risk of tipping during high-payload operations as the COG is lowered (Objective I).

In parallel, a custom conveyor system was developed and integrated, replacing the generic placeholder models of earlier iterations (Objective III). With an extended width of 1800 mm (Objective IV), this conveyor accommodated larger suitcases and enabled smoother transfer to external carts or conveyors. Together, these modifications brought the platform to a new level of maturity, combining physical test results, structural optimization, and operational alignment with the MANiBOT airport scenario.

5.4. Prototype Construction and Assembly

The twelfth and final design iteration (Figure 15) consolidated the MANiBOT platform into a structurally optimized, modular system ready for experimental deployment. Building directly on the limitations observed during physical testing, this version introduced reinforced 80 mm square steel profiles for the primary load-bearing elements. The enlarged cross-section significantly improved stiffness and mechanical robustness, ensuring reliable operation under dynamic conditions such as rapid manipulator movements or payload handling at extended reaches.

A second major innovation of this version was the incorporation of a dedicated lower technical compartment. This newly designed volume provided safe and accessible housing for essential electronics, including power supplies, communication units, controllers, and sensors. By internalizing these subsystems, the platform achieved cleaner cable routing, enhanced protection, and simplified maintenance. In addition, the redistribution of mass toward the lower frame contributed to a lower center of gravity, directly improving stability during mobile operation and interaction with external objects such as luggage carts.

Weighing 437.3 kg, the system guarantees stability against accelerations and dynamic stresses. The center of gravity was deliberately adjusted to reduce tipping risks and improve maneuverability.

From a workspace perspective, the platform's geometry was optimized to maximize the overlap of the manipulators' operational ranges. Both simulations and physical analyses confirmed the absence of collisions during collaborative operations, even when the arms work simultaneously on the same object. Each manipulator can reach up to 1 m beyond the

platform boundaries, covering vertical ranges from 0.5 to 2.5 m. This enables the system to interact effectively with carts, conveyors, shelves, and counters at multiple levels.

Furthermore, the design provides dedicated mounting space for all essential computational resources, including three NUCs, two Jetson devices, and one industrial-grade PC, which collectively support the perception, planning, and control layers of the system.

Functionally, the twelfth iteration retained all major advancements from previous versions: the extended-width tilting conveyor, side compartments, and universal AMR mounting interfaces. At the same time, it created new opportunities for scalability, as the technical compartment allows for the integration of additional devices without compromising accessibility or structural integrity.




Figure 15. Final version (v12): without AMR (a); with AMR (b).

To complement the presentation of the final platform design, additional diagrams have been included to illustrate both the overall reach of the manipulators and the interaction of the platform with the airport trolley (Figure 16a,b).




Figure 16. Final Robotic Platform reachability: upper view (a); isometric view (b).

Figure 17 shows the fully assembled MANiBOT prototype. This laboratory-ready platform integrates all structural and functional components—dual GoFa 12 manipulators, tilting conveyor, reinforced steel frame, and technical compartment—mounted on the AMR base. The image provides a direct view of the final system, bridging the iterative design process with the physical prototype now prepared for experimental validation.




Figure 17. Final Robotic Platform: front view (a); side view (b).

In parallel, a summary table has been added to explicitly relate the functional requirements defined in Section 3 with the design features implemented throughout the iterative development process. For each requirement, the table indicates whether and how it is fulfilled in the final configuration of the MANiBOT platform. This structured analysis provides a direct link between the original project objectives and the mechanical decisions adopted, demonstrating that the resulting system not only meets but, in several cases, exceeds the operational demands identified in real-world scenarios. For clarity and readability, Table 2 presents two representative requirements as illustrative examples of the proposed design methodology. The complete and detailed version of this table, including all functional and design requirements, is provided in Appendix A.2 (Table A2).

Table 2. Summarized design requirements across the iterative development of the MANiBOT platform.

Requirement ID	v1–3 (Conceptual)	v4–6 (Functional Prototyping)	v7–8 (Advanced Integration)	v9–10 (Detailed Design)	v11–12 (Final Structure)
FR1—Dual-arm capability	Conceptual positioning of manipulators is explored in a volumetric model.	First integration of dual GoFa 5 arms with a placeholder conveyor.	Upgrade to dual GoFa 12 arms with linear tracks to extend reach.	Verified manipulator–conveyor coordination and motion overlap.	Final dual GoFa 12 integration consolidated in reinforced frame.
FR2—Workspace overlap	Early assessment of arm reach envelopes.	Preliminary evaluation of cooperative workspace in simulation.	Linear-track architecture increased shared workspace coverage.	Detailed CAD and kinematic verification of collision-free overlap.	Experimental validation confirmed full overlap without interference.
OR1—Spatial validation and AMR adaptability	Volumetric envelope defined on P204 AMR base.	A detachable upper module was introduced for modular mounting.	Linear-track layout is considered to expand the workspace.	Structural redesign to accommodate airport trolley; conveyor tilt added.	Standardized AMR interface confirmed for P204/P604 bases.

Table 2. Cont.

Requirement ID	v1–3 (Conceptual)	v4–6 (Functional Prototyping)	v7–8 (Advanced Integration)	v9–10 (Detailed Design)	v11–12 (Final Structure)
OR2—Working height and task range	Not yet defined.	Conveyor height is conceptually positioned for accessibility.	Linear motion extended manipulation range.	Tilting conveyor enabled operation between 0.5 and 2.5 m.	Final design achieved full vertical reach, validated experimentally.
DR1—Spatial validation and AMR adaptability	Yes —volumetric envelope defined on P204 AMR; generic geometry adaptable.	Yes —detachable upper module introduced for modular mounting on multiple AMRs.	Yes —linear tracks considered with GoFa 12, ensuring coverage of extended workspaces.	Yes —structural redesign to adapt to Fraport luggage cart; conveyor tilting implemented.	Yes —standardized 80 mm steel profiles and universal AMR mounting interfaces.
DR2—Manipulator integration	No —only conceptual positioning studied.	Partial —first integration of dual GoFa 5 arms and placeholder conveyor.	Yes —upgrade to GoFa 12 arms with linear axes, validated for higher payload and reach.	Yes —manipulator–conveyor coordination verified under multi-level loading scenarios.	Yes —final dual GoFa 12 integration consolidated in reinforced frame.

6. Validation and Iterative Refinement: Description of the Evaluation Tests

To validate the structural robustness, stability, and safety of the MANiBOT platform, a comprehensive evaluation campaign was carried out. The tests were designed to replicate representative operational conditions while focusing exclusively on structural behavior, fastening reliability, and overall integrity of the system. Table 3 summarizes how the main design requirements were addressed through these tests, providing the rationale behind their selection. For clarity and readability, Table 3 presents a reduced set of validated design requirements. The complete validation matrix is included in Appendix A.3 (Table A3). The following subsections then describe in detail the procedures and methodologies employed.

Table 3. Design requirements validated through structural tests.

Design Requirement (DR)	Related Tests	Purpose of the Tests
DR4	Static load tests (1–4); Deflection tests (8–9); Finite Element Analysis tests (10); Material yielding analysis (11).	Confirm the platform’s ability to sustain the maximum operational load (250 kg) over short and long durations, ensuring that deflection remains within elastic limits and that no yielding occurs.
DR1	Static load on AMR-supported structure (3–4); Dynamic AMR-supported test (7)	Validate that the structure remains stable and reliable when mounted on different AMRs, replicating real deployment conditions.
DR5	Lifting system analysis and validation (11–12)	Ensure the feasibility of safely lifting and transferring the platform between AMRs, verifying both analytical sizing and experimental performance of the lifting system.
DR6	Rollover analysis (13); Anti-rollover solutions (14–18)	Assess tipping risk with the platform alone or mounted on the AMR, and evaluate corrective measures such as widened legs or counterweights to guarantee operational safety.

6.1. Static Load Tests

A series of static tests was performed to assess the platform's ability to sustain the maximum expected operational load of 250 kg under both short- and long-term conditions. The evaluation was carried out with the platform resting on the ground and mounted on the AMR. In all cases, bolt torque was initially adjusted to 25 Nm and monitored using a calibrated torque wrench.

Static load on ground-supported structure (24 h and 20 days):

The objective was to confirm that the frame could withstand the full load without exceeding an allowable deformation of 0.5 mm or a torque variation greater than 1% (24 h) or 5% (20 days). Weights were uniformly distributed across the structure (Figure 18a,b).



Figure 18. Static load tests on structure: Static Maximum Load Test on Ground-Supported Structure over a short period (24 h) (a); Static Maximum Load Test on Ground-Supported Structure over a large period (20 Days) (b).

Static load on AMR-supported structure (24 h and 20 days):

The same procedure was repeated with the platform mounted directly on the AMR, removing ground supports. Acceptance criteria were identical: ≤ 0.5 mm deformation and torque variations below 1% (24 h) and 5% (20 days) (Figure 19a,b).



Figure 19. Static load tests on AMR: Static Maximum Load Test on AMR-Supported Structure over Short Term (24 h) (a); Static Maximum Load Test on AMR-Supported Structure over Extended Period (20 Days) (b).

6.2. Dynamic Load Tests

Dynamic evaluations reproduced representative operating conditions in which the platform is subjected to cyclic accelerations while supporting the maximum operational load of 250 kg. In all cases, bolts were preloaded to 25 Nm and monitored with a calibrated torque wrench. The acceptance criterion was defined as a maximum torque variation $\leq 5\%$, together with the absence of permanent structural deformations.

Ground-supported dynamic test (V1):

The baseline configuration of the platform was placed on its support legs and subjected to 10 forward and backward cycles over a 10 m track. This first test aimed to detect potential weaknesses in the structural joints (Figure 20a).

Ground-supported dynamic test (V2, reinforced):

Following observations of deformation in V1, reinforcement brackets were added at the junctions between longitudinal beams and support legs. The test was then repeated under the same loading and motion conditions to validate the effectiveness of the reinforcement (Figure 20b).




Figure 20. Dynamic load tests: Visual deformation after dynamic test (a); Reinforced structure (b).

AMR-supported dynamic test:

The platform was mounted directly on the AMR, with ground supports removed, and subjected to 10 forward and backward cycles over a 5 m track. Bolt torque was measured before and after to assess fastener integrity under mobile dynamic stresses (Figure 21).



Figure 21. Dynamic Maximum Load Test on AMR-Supported Structure.

6.3. Deflection Tests

The deformation of the structure under maximum static load was quantified using dial gauges. The acceptance criterion was defined as a maximum allowable deflection of $L/250$.

Distributed load deflection: With a total load of 250 kg uniformly distributed across the platform, deflection was measured between the central point and the lateral support points using a dial gauge (Figure 22a,b).




Figure 22. Static Maximum Load Deflection Test with distributed force: side measure (a); Center measure (b).

Central point load deflection: A second configuration added an additional 90 kg concentrated load at the midpoint of the frame, while maintaining the distributed 250 kg load. Deflection was again measured between central and lateral points to assess the system under more critical loading (Figure 23a,b).




Figure 23. Static Maximum Load Deflection Test with principal force at the center of the structure: side measure (a); Center measure (b).

Although these deflection tests were carried out on versions 10 and 11 of the platforms, which were built with aluminum profiles, the results are considered directly transferable to the final design. The definitive steel structure employs square hollow sections of greater thickness and a material with a significantly higher yield strength, providing more than four times the bending stiffness of the earlier aluminum profiles. This combined increase in sectional inertia and material strength ensures that the final platform exhibits superior resistance to deformation under load, thereby reinforcing the validity of the conclusions drawn from the aluminum prototypes.

6.4. Finite Element Analysis of the Final Chassis

In addition to the experimental tests performed on the aluminum-based prototypes, a finite element analysis (FEA) was conducted to evaluate the structural behavior of the final steel chassis under representative loading conditions. The objective of this numerical study was not to replace experimental validation, but to complement it by providing a systematic and repeatable assessment of the platform response under different load distributions.

6.4.1. Load Application and Modeling Assumptions

The numerical model was derived directly from the final CAD geometry of the chassis, while assuming idealized connections, i.e., all joints were modeled as perfectly bonded. To ensure numerical robustness and to avoid non-physical local stress concentrations, the load application was intentionally simplified by consolidating it into a reduced set of representative regions associated with the dominant mass contributions of the system. These regions were conservatively sized, slightly overestimating the effective load transfer areas and thereby providing a robust and safe-side representation of the structural loading.

As illustrated in Figure 24, three main load application zones were considered:

- **Robot mounting areas (left and right):** representing the weight transmitted by each robotic arm to the platform.
- **Central load area:** representing the conveyor belt assembly and its payload.
- **Controller compartments:** included in the load distribution to account for auxiliary equipment mass.




Figure 24. Static maximum-load deflection test with the principal force applied at the center of the structure.

All loads were applied as vertical forces normal to the corresponding surfaces, ensuring a stable and physically consistent load introduction.

6.4.2. Parametric Load Definition

A systematic parametric study was carried out by varying both the magnitude and distribution of the applied loads:

- **Robot loads:** each robot mounting point was subjected to a vertical load ranging from 1000 N to 2000 N, applied independently at both the left and right robot locations.
- **Central load:** the central region was loaded from 500 N, corresponding to the nominal weight of the conveyor belt, up to 720 N, representing the combined weight of the conveyor belt and a maximum payload (luggage) placed on top.

This load definition enabled the evaluation of multiple representative operating scenarios, covering both nominal and worst-case conditions, while maintaining a consistent modeling framework.

6.4.3. Boundary Conditions

The chassis was constrained at the lower support regions to replicate the contact conditions with the AMR. Vertical displacements were restricted at these locations, while allowing in-plane degrees of freedom to avoid artificial over-constraining of the structure.

6.4.4. Mesh Characteristics

A solid finite element mesh was employed for the analysis. Mesh generation was based on curvature refinement to adequately capture geometric details while preserving computational efficiency. Quadratic high-order solid elements were used throughout the model.

The main mesh characteristics are summarized below:

- **Element type:** Quadratic solid elements (high order)
- **Meshing strategy:** Curvature-based combined mesh
- **Maximum element size:** approximately 52 mm
- **Minimum element size:** approximately 52 mm
- **Total number of nodes:** approximately 18,700
- **Total number of elements:** approximately 9500

This mesh configuration was selected as a compromise between accuracy and computational cost, and it was found to be adequate for capturing the global deformation and stress patterns of the structure.

6.4.5. Scope of the Numerical Analysis

The FEA focused on extracting global deflection and maximum von Mises stress values for each load configuration. These outputs were later used to construct load–deflection and stress–load curves, which are presented and discussed in Section 7. The analysis was restricted to linear elastic behavior, as all evaluated scenarios remained well below the material yield limit.

6.5. Material Yielding Analysis

In addition to the experimental deflection measurements, complementary analytical calculations were performed to assess whether the structure could experience plastic deformation under maximum operational loading. The objective was to compare the stresses induced by representative load cases with the elastic limit of the structural steel, applying a safety factor of 1.5.

For the specific use cases addressed in this work, the robotic manipulators operate at low velocities and perform manipulation actions that are relatively infrequent and quasi-static in nature. Consequently, inertial forces associated with rapid accelerations or decelerations of the robot arms were considered negligible with respect to the dominant gravitational and payload-induced loads, and the material yielding analysis was therefore conducted under conservative static loading assumptions.

It should be noted, however, that for applications involving faster motions, high cycle frequencies, or highly dynamic manipulation, inertial effects may significantly increase the stresses transmitted to the supporting structure. In such scenarios, a more detailed dynamic-aware design methodology should be adopted.

To this aim, the following procedure could be used. First, the worst-case manipulation trajectory is identified, in terms of joint motion and payload handling, using a robot

simulation environment. From this trajectory, the maximum joint accelerations and tool-center-point (TCP) accelerations are extracted. If the simulation tool provides interaction forces or reaction loads at the robot base, these can be directly used. Otherwise, inertial forces can be conservatively estimated by modeling the payload and relevant link masses as concentrated masses subjected to the computed accelerations. The resulting equivalent inertial forces and moments can then be applied to the supporting structure as external loads, either within a finite element model or through analytical stress evaluation. This approach allows dynamic effects to be incorporated in a conservative manner, ensuring that the supporting structure remains within the elastic regime under combined gravitational, payload, and inertial loading conditions.

The maximum induced stresses were calculated for the two representative cases in Equations (7)–(9).

$$I = \frac{(bh^3 - (b - 2t)(h - 2t)^3)}{12} \quad (7)$$

With b , h , t are width, height, and thickness, respectively.

$$y_{max} = \frac{h}{2} \quad (8)$$

Finally, induced stress is calculated:

$$\sigma = \frac{My_{max}}{I} \quad (9)$$

where M corresponds to the bending moment, y_{max} to the maximum distance from the neutral axis and I to the moment of inertia.

To test yielding, the calculated stresses were compared with the material's elastic limit ($\sigma_{elastic}$) of 275 MPa. The structure remains in the elastic range if $\sigma_{elastic} < \sigma_{plastic}$.

If $\sigma_{elastic} > \sigma_{plastic}$, the structure may enter plastic deformation, leading to increased deflection over time.

6.6. Lifting System Evaluation

The lifting mechanism, designed to enable safe detachment and transfer of the platform between different AMR bases, was evaluated analytically and experimentally.

First of all, a conservative lifted mass (m_{lift}) is considered, including a safety factor $SF = 1.5$ (Equation (10)), and lifted force (F_{lift}) is calculated (Equation (11)).

$$m_{lift} = m_{structure} \cdot SF \quad (10)$$

$$F_{lift} = m_{lift} \cdot g \quad (11)$$

The load is applied over a span L (lifting bar clear distance). Treating the load as uniformly distributed as shown in Equation (4)

$$\omega = \frac{F_{lift}}{L} \quad (12)$$

Because the platform is lifted by four supports, the load per side is halved:

$$\omega_{bar} = \frac{\omega}{2} \quad (13)$$

Now, using the maximum bending moment, M_{max} , the minimum solid circular bar is calculated (Equations (9)–(11)).

$$I = \frac{\pi d^4}{64} \quad (14)$$

$$c = \frac{d}{2} \quad (15)$$

$$\sigma_{max} = \frac{M_{max}c}{I} = \frac{32M_{max}}{\pi d^3} \quad (16)$$

To test the acceptance criteria, the calculated stresses were compared with the material's elastic limit ($\sigma_{elastic}$) of 316 MPa. The structure may remain in the elastic range if $\sigma_{max} < \sigma_{plastic}$ to accept the proposed bar diameter.

6.7. Rollover Stability Tests

To ensure safe operation, rollover stability tests were conducted to evaluate whether the platform's center of gravity (CoG) remains within the support polygon under the most unfavorable manipulator configurations. This analysis was performed on version 10 of the platform, both with and without the AMR attached.

Although dynamic effects were not explicitly modeled through time-dependent equations of motion, inertial effects associated with manipulator operation were conservatively taken into account by considering the maximum torques that the robotic arms can exert in their most unfavorable configurations.

Specifically, these torques were modeled in the CAD environment by applying equivalent concentrated masses at representative distances from the platform reference frame, reproducing the worst-case bending moments transmitted by the manipulators to the structure. This conservative static equivalence allows the center of gravity (CoG) to be evaluated under the most unfavorable manipulator configurations.

The acceptance criterion required the projection of the CoG to remain within the polygon defined by the platform's supporting legs. These rollover stability tests were conducted under two different conditions: with and without the AMR attached to the platform.

If stability could not be guaranteed, anti-rollover measures would be explored and tested.

6.8. Summary

To consolidate the description of the evaluation campaign, Table 4 provides a structured overview of all tests performed on the MANiBOT platform. A detailed description of all tests and conditions can be found in Appendix A.4 (Table A4). For each test, the table specifies its main objective, the methodology and instruments applied, and the corresponding figure where the setup is documented. This summary highlights the systematic nature of the validation process, covering static and dynamic load assessments, deflection measurements, material yielding analysis, lifting system verification, and rollover evaluations. Beyond serving as a compact reference, the table establishes the methodological link between the detailed procedures described above and the presentation of results in Section 7.

Table 4. Summary of evaluation tests performed on the MANiBOT platform.

Test Category	Test Name	Objective	Validation Approach	Figure(s)
Static loading	Static load tests (ground/AMR)	Validate structural capacity and long-term integrity under nominal payload	Distributed static loading, short- and long-duration monitoring	Figures 18 and 19
Dynamic loading	Dynamic load tests (ground/AMR)	Assess structural integrity under cyclic motion and operational loads	Cyclic displacement tests under maximum payload	Figures 20 and 21

Table 4. Cont.

Test Category	Test Name	Objective	Validation Approach	Figure(s)
Deflection	Static deflection tests	Quantify vertical deflection under distributed and concentrated loads	Experimental deflection measurements at representative points	Figures 22 and 23
Numerical validation	Finite Element Analysis (FEA)	Verify stress levels and stiffness under worst-case loading scenarios	Linear elastic FEA with parametric load cases	Figure 24
Material strength	Material yielding analysis	Assess the risk of plastic deformation under bending moments	Analytical stress evaluation with safety factor	Figure 25
Lifting system	Lifting system evaluation	Validate lifting bar sizing and operational feasibility	Analytical sizing and experimental lifting tests	Figures 31 and 32
Stability	Rollover stability analysis	Evaluate tipping risk under worst-case manipulator configurations	CAD-based center of gravity analysis	–

7. Results

The testing campaign provided a comprehensive evaluation of the MANiBOT platform under static, dynamic, and stability conditions. The following subsections summarize the outcomes of each test, highlighting the extent to which the structural design met the acceptance criteria established in Section 6.

7.1. Static Load Tests

The platform demonstrated robust performance under all static load conditions. No permanent deformations were detected, and torque variations remained well below the acceptance thresholds (Table 5).

Table 5. Results of static load tests on ground-supported structure (Tests 1–2) and on AMR-supported structure (Tests 3–4).

Test	Condition	Acceptance Criteria	Measured Result	Compliance
Test 1	24 h, ground-supported	Deformation ≤ 0.5 mm; Torque variation $\leq 1\%$	No deformation; Torque variation = 0%	Passed
Test 2	20 days, ground-supported	Deformation ≤ 0.5 mm; Torque variation $\leq 5\%$	No deformation; Torque variation < 1%	Passed
Test 3	24 h, AMR-supported	Deformation ≤ 0.5 mm; Torque variation $\leq 1\%$	No deformation; Torque variation < 1%	Passed
Test 4	20 days, AMR-supported	Deformation ≤ 0.5 mm; Torque variation $\leq 5\%$	No deformation; Torque variation < 1%	Passed

These results confirm that the platform maintains its structural integrity and fastening reliability both on the ground and on the AMR, for both short- and long-term operation. As a preventive measure, periodic bolt retightening every three months is recommended to guarantee long-term reliability.

7.2. Dynamic Load Tests

Dynamic evaluations revealed important insights into the platform's performance under cyclic stresses. The results obtained from these tests are presented in Table 6.

Table 6. Results of dynamic load tests on ground-supported structure (V1 and V2) and on AMR-supported structure.

Test	Condition	Acceptance Criteria	Observed Results	Compliance
Ground-supported V1	10 cycles, 10 m, baseline frame	No deformation; Torque variation $\leq 5\%$	Deformation at beam–leg junctions; Max torque variation = 4%	Failed (structural weakness)
Ground-supported V2	10 cycles, 10 m, reinforced with brackets	No deformation; Torque variation $\leq 5\%$	No deformation; 4 bolts loosened to 24 Nm (3% variation)	Passed (with maintenance note)
AMR-supported	10 cycles, 5 m	No deformation; Torque variation $\leq 5\%$	No deformation; No significant torque loss	Passed

These results show that the initial structure (V1) exhibited critical weaknesses under cyclic loads, with visible deformation at the beam–leg junctions. The reinforced version (V2) effectively solved this issue, eliminating structural deformation. However, slight loosening of four bolts suggests the need for preventive retightening after transport. When mounted on the AMR, the platform demonstrated robust stability and fastening reliability, validating its suitability for mobile dynamic operations.

7.3. Deflection Tests


The results confirm that structural deflections remained below the allowable limit of 7.2 mm in both loading configurations, demonstrating that the platform operates within the elastic range and maintains its structural integrity under maximum static loading (Table 7).

Table 7. Results of the distributed and central point load deflection test (Test 8).

Measurement	Values (mm)	Mean Deflection (mm)	Acceptance Criterion	Compliance
Lateral point	6.73–6.76	6.74	≤ 7.2 mm	Passed
Lateral point (with 90 kg at center)	6.99–7.01	7.01	≤ 7.2 mm	Passed

7.4. Material Yielding Analysis

The analytical verification confirmed that the steel structure remains well within the elastic range under the most demanding loading scenarios. Using the properties of the $80 \times 80 \times 3$ mm (b, h, and t, respectively), square hollow section ($I = 4.62 \times 10^6 \text{ mm}^4$, $y_{max} = 40$ mm) and considering a safety factor of 1.5, two possible loading configurations were considered (Figure 25a,b).

**Figure 25.** Static maximum load deflection test: uniform distributed force (a); distributed force with higher appliance in the central part of the structure (b).

- I. **Case 1—Distributed structural mass:** The total lifted mass of 437.5 kg was increased by a safety factor to 656.25 kg. When distributed over the full 1.8 m span,

this corresponded to a line load of 364.6 kg/m (≈ 3.58 kN/m). The resulting maximum bending moment was calculated as $M_{max,1} = 1450$ Nm.

- II. **Case 2—Combined lateral and central loading:** Two lateral loads of 200 kg were considered, increased to 300 kg after applying the safety factor, corresponding to a distributed load of 600 kg/m (≈ 5.89 kN/m). In addition, a central mass of 137.5 kg was increased to 206.25 kg. This combined configuration produced a maximum bending moment of $M_{max,1} = 1440$ Nm.

Since Case 1 generates the highest bending moment, it was adopted as the design reference for structural verification. Results are presented in Table 8.

Table 8. Analytical calculation of stresses under bending moments.

Case	Applied Bending Moment (Nm)	Safety Factor	Induced Stress (MPa)	Elastic Limit (MPa)	Compliance
1—Distributed load	970	1.5	12.5	275	Elastic
2—Combined lateral + central load	1400	1.5	13.0	275	Elastic

Both cases demonstrate a large safety margin, with induced stresses more than one order of magnitude below the elastic limit of S275 steel (275 MPa). Case 2, representing the most unfavorable load configuration, reached only 12.1 MPa, confirming that no yielding or progressive plastic deformation is expected. These results validate the robustness of the selected structural section under real operational conditions.

7.5. Finite Element Analysis of the Final Chassis

The finite element analysis provided a complementary numerical assessment of the structural behavior of the final steel chassis under a wide range of representative loading conditions.

Figure 26 shows the von Mises stress distribution obtained for a representative worst-case loading configuration. As expected, the highest stress concentrations appear at the robot mounting regions, where the applied loads are directly transferred to the structural frame. Lower stress levels are observed along the longitudinal beams and at the central section of the chassis. In all evaluated scenarios, the maximum von Mises stress remains significantly below the elastic limit of the structural steel, confirming that the structure operates entirely within the linear elastic regime.




Figure 26. Von Misses tension calculated with Finite Element Analysis.

Figure 27 presents the corresponding resultant displacement (URES) field for the same loading condition. The maximum deflection occurs at the robot mounting platforms, while

the central region exhibits a smoother deformation profile. This deformation pattern is consistent with the experimental deflection tests, where the largest displacements were measured near the most heavily loaded areas. The absolute displacement levels remain low and compatible with the acceptance criteria defined in Section 6.3 (≤ 7.2 mm).




Figure 27. Deflection calculated with Finite Element Analysis.

To provide a clearer and more quantitative interpretation, the numerical results were post-processed to generate load–response curves. Figure 28 illustrates the evolution of maximum von Mises stress as a function of the total applied load. A nearly linear relationship is observed, which is consistent with the assumption of linear elastic behavior and confirms the absence of non-linear effects or material yielding within the investigated load range.




Figure 28. Maximum stress vs. total applied load.

Figure 29 presents the maximum deflection as a function of the robot load for different central load levels. The curves illustrate the influence of the suitcase weight applied at the central section of the platform, while the load applied on each side of the platform is varied from 1000 N to 2000 N. The results highlight how increasing the central load leads to higher global deflection for the same robot load, emphasizing the combined effect of robot weight and payload distribution on the structural response.

These results highlight that not only the magnitude but also the spatial distribution of the loads plays a significant role in the global structural response of the platform.

Finally, Figure 30 summarizes the global deformation behavior by representing the maximum deflection as a function of the total applied load for all evaluated scenarios.




Figure 29. Load-deflection curves.




Figure 30. Deflection vs. total applied load.

Despite the different load distributions, the overall trend remains consistent, showing a stable and predictable increase in deflection with increasing load. This global view reinforces the robustness of the chassis design and confirms that the platform maintains adequate stiffness under all considered operating conditions.

Overall, the FEA results are in good agreement with the experimental findings and analytical calculations. The numerical analysis confirms that the final steel chassis exhibits sufficient stiffness and structural safety margins under representative operational loads, thereby supporting the validity of the adopted design.

7.6. Lifting System Evaluation

The lifting system was validated analytically and experimentally to ensure its capacity to safely elevate the platform during AMR exchange.

With the Analytical calculations, the aim was to determine the minimum diameter of the transverse steel bars capable of supporting the weight of the final structure ($437.5 \text{ kg} \approx 4291.88 \text{ N}$). A safety factor of 1.5 was applied, and the material's elastic limit was defined as 316 MPa. (Figure 31). The acceptance criterion required the induced stress to remain below the elastic limit for all scenarios.




Figure 31. Lifting system evaluation.

Analytical calculations determined the minimum required diameter of the transverse lifting bar under the worst-case bending moment of $M_{max} = 720 \text{ N}$. Applying the stress criterion with the elastic limit of $\sigma_{elastic} = 316 \text{ MPa}$, the minimum diameter was obtained as $d_{min} = 28.5 \text{ mm}$.

In conclusion, the 30 mm transverse bar was selected as the standard, exceeding the minimum required values and ensuring operational safety. Among the tested configurations, the one-point support design provided the most favorable results, requiring the lowest stress capacity and thus offering the safest operational margin.

To confirm the analytical findings, a practical lifting test was conducted in which four lifting legs supported a total load of 437.5 kg (Figure 32). This experiment empirically demonstrated that the calculated bar neither deflects excessively nor fails under load, and that the lifting mechanism is both capable and sufficient to safely perform the platform elevation required for AMR exchange.



Figure 32. Experimental Lifting Test.

7.7. Rollover Stability Tests

The rollover analyses revealed critical stability limitations in platform version 10 when evaluated without the AMR:

Standalone configuration: The CoG (shown as the blue–white marker) was found to lie dangerously close to the support polygon, posing a tipping risk under unfavorable manipulator postures (Figure 33a,b).




Figure 33. Rollover analysis of version 10: Evaluation of the rollover in platform version 10 (isometric view) (a). Evaluation of the rollover in platform version 10 (side view) (b).

With AMR: The additional weight shifted the CoG downward, fully eliminating the rollover risk (Figure 34).




Figure 34. Rollover analysis with AMR attached.


To mitigate the instability observed in standalone mode, several corrective strategies were evaluated, including false AMR anchoring (Figure 35a), widened support legs (Figure 35b), and counterweight systems with steel or concrete masses (Figure 35c,d). Each alternative was experimentally validated in simulation to examine feasibility and operational impact (Table 9).

Table 9. Evaluation of anti-rollover solutions.


Solution	Description	Effectiveness	Practical Limitations	Outcome
False AMR anchoring	Tie the platform to a heavy block simulating AMR	Eliminates rollover	Not mobile; arms must be folded for movement	Rejected
Widened support legs	Extend the support polygon beyond the current footprint	Works in the lab	Not feasible in real airport/supermarket layouts	Rejected (valid only in the lab)

Table 9. Cont.


Solution	Description	Effectiveness	Practical Limitations	Outcome
Hybrid stabilizer system	Combination of anchoring + leg extension	Structurally safest option	Poor maneuverability, impractical in real use	Rejected
Counterweights—steel blocks	Adds mass low on the structure to lower CoG	Effective, removable for AMR lifting	Material is costly, but feasible	Accepted as the best compromise
Counterweights—concrete cylinders	Cheaper than steel, lower density	Partially effective	Bulkier, less efficient	Discarded




(a)



(b)



(c)



(d)

Figure 35. Antirollover solutions: false AMR anchoring (a); widened support legs (b); counterweight systems with steel masses (c); counterweight systems with concrete masses (d).

8. Conclusions

This paper has presented the methodological framework and its practical application to the design and realization of the MANiBOT dual-arm robotic platform. Rather than describing a single prototype in isolation, the contribution provides a generalizable design methodology that guides the development of bimanual robotic structures from requirement definition to experimental validation.

The process followed a requirement-driven and iterative design cycle, in which each stage—requirement identification, conceptual architecture, detailed mechanical design, prototype construction, and validation—generated traceable deliverables and decision criteria. The application of this workflow to the MANiBOT platform demonstrated its capacity to ensure structural robustness, modularity, and adaptability while maintaining compliance with safety and interoperability standards.

The methodological approach enabled the systematic transformation of functional and operational requirements—derived from supermarket and airport scenarios—into verified engineering solutions. Critical outcomes of the process include:

- A validated requirement specification framework, translating operational tasks into measurable design constraints and performance indicators.
- A conceptual architecture supporting dual-arm cooperation, workspace overlap, and modular reconfiguration for different environments.
- A structurally optimized detailed design, using reinforced steel and standardized interfaces to balance rigidity, manufacturability, and maintainability.
- A prototyping and validation stage confirming that analytical predictions and physical tests are consistent within the defined acceptance criteria.

Beyond the specific case of MANiBOT, the results confirm that the proposed methodology can serve as a replicable guideline for the development of future dual-arm robotic platforms. The integration of requirement traceability, simulation-based verification, and iterative refinement ensures reproducibility and design transparency—critical aspects for technology transfer and industrial adoption.

Future work will extend this methodology toward the electromechanical and cognitive integration of perception modules, tactile and proximity sensors, and high-level control frameworks. The upcoming trials in real supermarket and airport environments will further validate not only the mechanical design but also the platform's capacity for perception-driven and human-aware manipulation.

In conclusion, the MANiBOT platform embodies both a proof of concept and a proof of methodology. By establishing a structured and verifiable workflow for designing dual-arm robotic systems, this work contributes a practical blueprint for developing next-generation collaborative platforms that combine mechanical robustness, safety compliance, and adaptive intelligence in dynamic, human-shared environments.

Author Contributions: M.P.-Y.: Conceptualization, Methodology, Software, Validation, Writing—Original Draft Preparation, Review and Editing. J.E.S.-G.: Conceptualization, Methodology, Formal Analysis, Supervision, Writing—Original Draft Preparation, Review and Editing. All authors have read and agreed to the published version of the manuscript.

Funding: This work was partially supported by the European Commission, under the European Project MANiBOT, grant number 101120823.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding authors.

Acknowledgments: The authors acknowledge the use of AI-based language tools to assist with English language polishing and stylistic improvements. The authors take full responsibility for the content of the manuscript.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A. Complete Tables

This appendix includes the complete versions of the tables summarized in the main body of the paper. It provides the full comparative analysis of dual-arm robotic platforms, the detailed design requirements considered during the iterative development of the MANiBOT platform, the requirements validated through structural testing, and the complete set of evaluation tests performed.

Appendix A.1.

Table A1. Comparative analysis of dual-arm robotic platforms with respect to design methodology, mobility, and adaptability.

Reference	Platform Design Methodology	Arm Configuration	Mobility (AMR/Mobile Base)	Robot-Agnostic (Adaptability)	Main Focus
Peñacoba and Sierra-García (This work)	Yes (generalized and structured). Proposes a five-stage framework spanning from requirements definition to experimental validation, applicable to different platforms.	Dual-arm (bimanual)	Yes. Mobility is considered from the early design stages.	Yes. Framework applicable to different manipulators and AMRs.	Standardization of mechanical engineering and structural validation of mobile dual-arm platforms.
Li et al. (2023) [21]	No. Describes a hardware/software integration tailored to a specific application.	Multi-arm (four Cartesian arms)	Yes. Wheeled mobile platform (Ackermann).	No. Application-specific design for agricultural harvesting.	Task planning, perception, and efficient apple harvesting.
Park et al. (2014) [17]	No. Relies on kinematic and dynamic analysis for an ad hoc design.	Dual-arm (custom 7-DoF arms)	Yes. Omnidirectional mobile platform.	No. Custom-designed arms integrated into the torso.	Electronic device assembly and workspace analysis.
Di Castro et al. (2017) [19]	Partial. Presents a modular control architecture, but not a step-by-step mechanical design methodology.	Dual-arm (configurable)	Yes. CERNbot tracked/wheeled platform.	Yes (modular hardware configuration).	Control architecture for teleoperation and navigation in hazardous environments.
Chen and Hsiao (2023) [22]	No. Focuses on the development and integration of joint modules.	Dual-arm (14 DoF, proprietary modules)	No. Stationary system cooperating with an external AMR.	No. Based on proprietary all-in-one joint modules.	Development of joint modules and control for collaborative industrial assembly.
Duan et al. (2023) [25]	No. Proposes a six-dimensional Digital Twin architecture rather than a physical construction methodology.	Dual-arm (UR5 + UR10)	No. Fixed experimental setup.	Partial (software-level framework).	Visual monitoring and digital twin for assembly processes.
Moore et al. (2007) [18]	No. Describes the implementation of a fixed workcell in an existing facility.	Dual-arm (Stäubli arms)	No. Large-scale fixed installation.	No. Designed for a specific pharmaceutical screening application.	Flexible programming and workflow management for high-throughput screening.

Appendix A.2.

Table A2. Design requirements across the iterative development of the MANiBOT platform.

Requirement ID	v1–3 (Conceptual)	v4–6 (Functional Prototyping)	v7–8 (Advanced Integration)	v9–10 (Detailed Design)	v11–12 (Final Structure)
FR1—Dual-arm capability	Conceptual positioning of manipulators is explored in a volumetric model.	First integration of dual GoFa 5 arms with a placeholder conveyor.	Upgrade to dual GoFa 12 arms with linear tracks to extend reach.	Verified manipulator–conveyor coordination and motion overlap.	Final dual GoFa 12 integration consolidated in reinforced frame.
FR2—Workspace overlap	Early assessment of arm reach envelopes.	Preliminary evaluation of cooperative workspace in simulation.	Linear-track architecture increased shared workspace coverage.	Detailed CAD and kinematic verification of collision-free overlap.	Experimental validation confirmed full overlap without interference.
FR3—Conveyor subsystem integration	Conveyor not yet included.	Placeholder conveyor introduced for feasibility checks.	The linear conveyor prototype enabled coordinated manipulation.	Inclinable conveyor implemented for multi-height unloading.	Custom 1.8 m tilting conveyor integrated; repeatability validated.
FR4—Structural rigidity under payload	Conceptual frame only, no load validation.	Aluminum prototype built; reinforcements added after deflection tests.	Sheet-metal alternative studied to improve stiffness.	Transition to extruded-aluminum profiles; FEA confirmed deflection within limits.	Reinforced 80 mm steel profiles validated under static/dynamic loading; CoG lowered.
FR5—Stability on AMR base	Basic stability analysis conceptually reviewed.	Anchor reinforcements improved static equilibrium.	Linear-track kinematics assessed for dynamic stability.	Rollover analyses identified risks; counterweights tested.	Steel reinforcement and low-mounted mass ensured stable behaviour on AMRs.
FR6—Modularity and maintainability	Internal distribution not defined.	Preliminary ventilation and access panels designed.	Compartments reorganized to improve accessibility.	Modular aluminum profiles are adopted for easy subsystem replacement.	Final steel frame with removable panels and standardized cable routing.
FR7—Allocation of internal components	Only external volume estimated.	First internal layout trials for electronics and power supply.	Compartments reorganized; clearer electronics allocation.	Conveyor and electronics integration refined; sensors included.	Lower technical compartment defined for NUCs, Jetsons, PSUs, and controllers.
FR8—Symmetry and balance	A symmetrical layout is proposed for control simplicity.	Partial geometric symmetry was achieved in prototypes.	Bilateral configuration completed with mirrored supports.	Structural reinforcements preserved geometric symmetry.	Final mirrored layout validated; balanced torque distribution confirmed.
OR1—Spatial validation and AMR adaptability	Volumetric envelope defined on P204 AMR base.	A detachable upper module was introduced for modular mounting.	Linear-track layout is considered to expand the workspace.	Structural redesign to accommodate airport trolley; conveyor tilt added.	Standardized AMR interface confirmed for P204/P604 bases.
OR2—Working height and task range	Not yet defined.	Conveyor height is conceptually positioned for accessibility.	Linear motion extended manipulation range.	Tilting conveyor enabled operation between 0.5 and 2.5 m.	Final design achieved full vertical reach, validated experimentally.

Table A2. Cont.

Requirement ID	v1–3 (Conceptual)	v4–6 (Functional Prototyping)	v7–8 (Advanced Integration)	v9–10 (Detailed Design)	v11–12 (Final Structure)
OR3—Safety and regulatory compliance	Safety distances are qualitatively considered.	Preliminary HRI clearances applied.	Layout aligned with ISO 10218-2 workspace rules.	Safety components (E-Stop, force limiters) integrated.	Compliance with ISO/TS 15066 confirmed through tests.
OR4—Maintainability and service access	Not addressed.	Early consideration of ventilation and serviceability.	Panels reorganized for accessibility.	Cable routing and ventilation paths are defined.	Service access optimized; MTTR reduced below specification.
OR5—Power and computing integration	Internal electronics not yet planned.	Space reserved for power distribution.	First integration of computing devices in a mock-up.	Dedicated areas for controllers and communication units.	The lower compartment hosts power and computing hardware with EMC compliance.
DR1—Spatial validation and AMR adaptability	Yes —volumetric envelope defined on P204 AMR; generic geometry adaptable.	Yes —detachable upper module introduced for modular mounting on multiple AMRs.	Yes —linear tracks considered with GoFa 12, ensuring coverage of extended workspaces.	Yes —structural redesign to adapt to Fraport luggage cart; conveyor tilting implemented.	Yes —standardized 80 mm steel profiles and universal AMR mounting interfaces.
DR2—Manipulator integration	No —only conceptual positioning studied.	Partial —first integration of dual GoFa 5 arms and placeholder conveyor.	Yes —upgrade to GoFa 12 arms with linear axes, validated for higher payload and reach.	Yes —manipulator–conveyor coordination verified under multi-level loading scenarios.	Yes —final dual GoFa 12 integration consolidated in reinforced frame.
DR3—Conveyor integration and evolution	No —conveyor not yet included.	Partial —placeholder conveyor introduced to simulate transfer.	Yes —linear motion enabled cooperative use of the conveyor.	Yes —inclined conveyor implemented for airport cart loading.	Yes —custom conveyor (1.8 m width) integrated for large luggage handling.
DR4—Structural robustness and load capacity	No —conceptual frame only, no load validation.	Partial —aluminum prototypes tested; reinforcements added after dynamic deformation.	Partial —sheet-metal exploration; improved rigidity, but not a final solution.	Yes —extruded aluminum profiles adopted; deflection tests confirmed acceptable limits.	Yes —reinforced steel frame validated under static/dynamic tests; center of gravity lowered.
DR5—Maintainability and modularity	No —internal distribution not defined.	Partial —initial ventilation and panel layouts explored.	Partial —compartments reorganized for improved access.	Yes —modular aluminum profiles allowed flexible integration of subsystems.	Yes —a lower technical compartment introduced for safe, accessible housing of electronics.
DR6—Stability and rollover prevention	No —only conceptual considerations of stability.	Partial —anchor reinforcements improved static stability.	Partial —linear track analysis ensured a workspace without collisions.	Partial —rollover analyses identified risk; countermeasure strategies tested.	Yes —steel reinforcement, counterweight systems, and lower center of gravity ensured robust stability.

Table A2. Cont.

Requirement ID	v1–3 (Conceptual)	v4–6 (Functional Prototyping)	v7–8 (Advanced Integration)	v9–10 (Detailed Design)	v11–12 (Final Structure)
DR7—Allocation and definition of internal components	No —only external volume estimated, internal layout undefined.	Partial —first distribution trials: preliminary space for panels, power supply, and ventilation.	Partial —compartments reorganized, internal electronics more clearly arranged.	Yes —conveyor and electronics integration refined; space for controllers and sensors accounted for.	Yes —dedicated lower compartment defined for NUCs, Jetsons, PSU, and communication units, with protected access and organized cabling.

Appendix A.3.

Table A3. Design requirements validated through structural tests.

Design Requirement (DR)	Related Tests	Purpose of the Tests
DR4—Structural robustness and load capacity	Static load tests (1–4); Deflection tests (8–9); Finite Element Analysis tests (10); Material yielding analysis (11).	Confirm the platform’s ability to sustain the maximum operational load (250 kg) over short and long durations, ensuring that deflection remains within elastic limits and that no yielding occurs.
DR1—Spatial validation and AMR adaptability	Static load on AMR-supported structure (3–4); Dynamic AMR-supported test (7)	Validate that the structure remains stable and reliable when mounted on different AMRs, replicating real deployment conditions.
DR5—Maintainability and modularity	Lifting system analysis and validation (11–12)	Ensure the feasibility of safely lifting and transferring the platform between AMRs, verifying both analytical sizing and experimental performance of the lifting system.
DR6—Stability and rollover prevention	Rollover analysis (13); Anti-rollover solutions (14–18)	Assess tipping risk with the platform alone or mounted on the AMR, and evaluate corrective measures such as widened legs or counterweights to guarantee operational safety.

Appendix A.4.

Table A4. Summary of evaluation tests performed on the MANiBOT platform.

Test No.	Test Name	Objective	Methodology/Tools	Figure(s)
1	Static load (ground-supported, 24 h)	Validate structural capacity under nominal load (250 kg) over a short duration	Distributed load; torque wrench monitoring	Figure 18a
2	Static load (ground-supported, 20 days)	Assess long-term integrity of structure and fasteners	Same as Test 1, monitored over 20 days	Figure 18b
3	Static load (AMR-supported, 24 h)	Validate load capacity when mounted on AMR, short duration	Platform on AMR; torque wrench monitoring	Figure 19a
4	Static load (AMR-supported, 20 days)	Assess long-term integrity of AMR	Same as Test 3, monitored over 20 days	Figure 19b

Table A4. Cont.

Test No.	Test Name	Objective	Methodology/Tools	Figure(s)
5	Dynamic load (ground-supported, V1)	Evaluate integrity under cyclic motion with maximum load	10 cycles (10 m forward/backward); torque wrench readings	Figure 20a
6	Dynamic load (ground-supported, V2, reinforced)	Validate reinforced frame with corner brackets	Same as Test 5, after reinforcement	Figure 20b
7	Dynamic load (AMR-supported)	Assess performance under cyclic motion when mounted on AMR	10 cycles (5 m forward/backward); torque wrench readings	Figure 21
8	Static deflection (distributed load)	Quantify deflection under 250 kg uniformly distributed	Dial gauge at the central and lateral points	Figure 22
9	Static deflection (central load)	Quantify deflection with 250 kg + 90 kg central point	Same as Test 8 with additional central load	Figure 23
10	Finite Element Analysis tests of the final chassis	Verify that stress levels remain below the elastic limit under worst-case loading. Evaluate global stiffness and deformation trends under varying load distributions	CAD-based finite element model; simplified load regions; parametric sweep: robot loads (1000–2000 N) and central load (500–720 N); Von Mises stress evaluation using linear elastic FEA.	Figure 24
10	Material yielding analysis	Assess the risk of plastic deformation under bending moments	Analytical stress calculation with safety factor 1.5	Figure 25a,b
11	Lifting system—bar sizing	Determine the minimum diameter of the transverse bar for lifting	Analytical bending stress calculation	Figure 31
12	Lifting system—experimental validation	Validate the lifting concept under a 200 kg partial load	Physical lifting with 2 legs; visual inspection	Figure 32
13	Rollover analysis (platform v10)	Evaluate tipping risk with arms extended, with/without AMR	CAD-based CoG analysis	-

References

1. Callari, T.C.; Segate, R.V.; Hubbard, E.M.; Daly, A.; Lohse, N. An ethical framework for human-robot collaboration for the future people-centric manufacturing: A collaborative endeavour with European subject-matter experts in ethics. *Technol. Soc.* **2024**, *78*, 102680. [CrossRef]
2. Xu, J.; Mawase, F.; Schieber, M.H. Evolution, biomechanics, and neurobiology converge to explain selective finger motor control. *Physiol. Rev.* **2024**, *104*, 983–1020. [CrossRef] [PubMed]
3. Baker, J.; Barton, R.A.; Venditti, C. Human dexterity and brains evolved hand in hand. *Commun. Biol.* **2025**, *8*, 1257. [CrossRef] [PubMed]
4. Ford, C.J.; Li, H.; Catalano, M.G.; Bianchi, M.; Psomopoulou, E.; Lepora, N.F. Shear-Based Grasp Control for Multi-Fingered Underactuated Tactile Robotic Hands. *IEEE Trans. Robot.* **2025**, *41*, 3113–3128. [CrossRef]

5. Luo, J.; Xu, C.; Wu, J.; Levine, S. Precise and dexterous robotic manipulation via human-in-the-loop reinforcement learning. *Sci. Robot.* **2025**, *10*, eads5033. [CrossRef] [PubMed]
6. Gao, B.; Fan, J.; Zheng, P. Empower dexterous robotic hand for human-centric smart manufacturing: A perception and skill learning perspective. *Robot. Comput. Integr. Manuf.* **2025**, *93*, 102909. [CrossRef]
7. Mu, Y.; Chen, T.; Peng, S.; Chen, Z.; Gao, Z.; Zou, Y.; Luo, P. RoboTwin: Dual-arm robot benchmark with generative digital twins (early version). In *Computer Vision—ECCV 2024 Workshops; Lecture Notes in Computer Science*; Del Bue, A., Canton, C., Pont-Tuset, J., Tommasi, T., Eds.; Springer: Berlin/Heidelberg, Germany, 2024; Volume 15630. [CrossRef]
8. Jiang, Y.; Fakhruddin, H.; Pizzuto, G.; Longley, L.; He, A.; Dai, T.; Clowes, R.; Rankin, N.; Cooper, A.I. Autonomous biomimetic solid dispensing using a dual-arm robotic manipulator. *Digit. Discov.* **2023**, *2*, 1733–1744. [CrossRef]
9. Peñacoba, M.; Sierra-García, J.E.; Santos, M. Human-intelligent trajectory optimization for robotic manipulators with hybrid PSO-PS algorithm. *Adv. Eng. Inform.* **2026**, *69*, 103941. [CrossRef]
10. ISO 10218-1:2011; Robots and Robotic Devices—Safety Requirements for Industrial Robots—Part 1: Robots. International Organization for Standardization (ISO): Geneva, Switzerland, 2011.
11. ISO 10218-2:2011; Robots and Robotic Devices—Safety Requirements for Industrial Robots—Part 2: Robot Systems and Integration. International Organization for Standardization (ISO): Geneva, Switzerland, 2011.
12. ISO/TS 15066:2016; Robots and Robotic Devices—Collaborative Robots. International Organization for Standardization (ISO): Geneva, Switzerland, 2016.
13. MANiBOT. MANiBOT Project. 2023. Available online: <https://manibot-project.eu/> (accessed on 1 January 2026).
14. Lv, N.; Liu, J.; Jia, Y. Dynamic modeling and control of deformable linear objects for single-arm and dual-arm robot manipulations. *IEEE Trans. Robot.* **2022**, *38*, 2341–2353. [CrossRef]
15. Xie, F.; Chowdhury, A.; De Paolis Kaluza, M.; Zhao, L.; Wong, L.; Yu, R. Deep imitation learning for bimanual robotic manipulation. *Adv. Neural Inf. Process. Syst.* **2020**, *33*, 2327–2337.
16. Xie, F.; Guo, Z.; Li, T.; Feng, Q.; Zhao, C. Dynamic Task Planning for Multi-Arm Harvesting Robots Under Multiple Constraints Using Deep Reinforcement Learning. *Horticulturae* **2025**, *11*, 88. [CrossRef]
17. Park, D.I.; Park, C.; Do, H.; Choi, T.; Kyung, J. Development of dual arm robot platform for automatic assembly. In Proceedings of the 2014 14th International Conference on Control, Automation and Systems (ICCAS 2014), Gyeonggi, Republic of Korea, 22–25 October 2014; pp. 319–321.
18. Moore, K.W.; Newman, R.; Chan, G.K.; Leech, C.; Allison, K.; Coulson, J.; Simpson, P.B. Implementation of a high specification dual-arm robotic platform to meet flexible screening needs. *JALA J. Assoc. Lab. Autom.* **2007**, *12*, 115–123. [CrossRef]
19. Di Castro, M.; Buonocore, L.R.; Ferre, M.; Gilardoni, S.; Losito, R.; Lunghi, G.; Masi, A. A dual arms robotic platform control for navigation, inspection and telemanipulation. In Proceedings of the 16th International Conference on Accelerator and Large Experimental Control Systems (ICALEPCS'17), Barcelona, Spain, 8–13 October 2017; pp. 8–13.
20. Xu, R.; Qian, L.; Zhao, X. Development of dual-arm mobile robot platform based on ROS: [version 1; peer review: 2 approved]. *Cobot* **2022**, *1*, 4. [CrossRef]
21. Li, T.; Xie, F.; Zhao, Z.; Zhao, H.; Guo, X.; Feng, Q. A multi-arm robot system for efficient apple harvesting: Perception, task plan and control. *Comput. Electron. Agric.* **2023**, *211*, 107979. [CrossRef]
22. Chen, C.C.; Hsiao, J.C. The Development and Application of an Industrial Collaborative Dual-Arm Robot for Tool Assembly Cell. *Int. J. Robotics* **2023**, *6*, 31–37.
23. Naveed, A.; Bukhari, S.M.A.; Husnain, N.; Balouch, N.; Noor, S.; Ahmad, W.; Sarfraz, S.; Wajahat, N.; Siddiqui, F.A. Structural Integrity and Cost Analysis of a Connecting Rod for Articulated Robotic Arm using various Stainless Steel and Aluminium Alloys. *Tech. J.* **2024**, *29*, 24–34.
24. Dörfler, K.; Dielemans, G.; Leutenegger, S.; Jenny, S.E.; Pankert, J.; Sustarevas, J.; Lachmayer, L.; Raatz, A.; Lowke, D. Advancing construction in existing contexts: Prospects and barriers of 3D printing with mobile robots for building maintenance and repair. *Cem. Concr. Res.* **2024**, *186*, 107656. [CrossRef]
25. Duan, J.; Gong, X.; Zhang, Q.; Qin, J. A digital twin-driven monitoring framework for dual-robot collaborative manipulation. *Int. J. Adv. Manuf. Technol.* **2023**, *125*, 4579–4599. [CrossRef]
26. Young, W.C.; Budynas, R.G.; Sadegh, A.M. *Roark's Formulas for Stress and Strain*, 9th ed.; McGraw-Hill: Columbus, OH, USA, 2020.
27. Budynas, R.G.; Nisbett, J.K. *Shigley's Mechanical Engineering Design*, 11th ed.; McGraw-Hill: Columbus, OH, USA, 2020.
28. Galloway, A.L. *Mechanical Vibrations: Types, Testing and Analysis*; Nova Science Publishers: Hauppauge, NY, USA, 2011.
29. Inman, D.J.; Singh, R.C. *Engineering Vibration*, 5th ed.; Pearson: London, UK, 2021.
30. Arteaga, M.A.; Gutiérrez-Giles, A.; Pliego-Jiménez, J. *Local Stability and Ultimate Boundedness in the Control of Robot Manipulators*; Springer: Cham, Switzerland, 2022.
31. Mustary, S.; Kashem, M.A.; Sony, J.M.; Hossain, N.; Chowdhury, M.A. Advancing Stability in Robot Manipulators: A Review of Recent Progress and Parameters. *Eng. Rep.* **2025**, *7*, e70207. [CrossRef]

32. Meriam, J.L.; Kraige, L.G.; Bolton, J.N. *Engineering mechanics: Dynamics*; John Wiley & Sons: Hoboken, NJ, USA, 2020.
33. Craig, R.R., Jr.; Taleff, E.M. *Mechanics of Materials*; John Wiley & Sons: Hoboken, NJ, USA, 2020.
34. Link, R.E. An introduction to the design and behavior of bolted joints. *J. Test. Eval.* **1991**, *19*, 417–418. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Article

Design and Development of an Autonomous Mobile Robot for Unstructured Indoor Environments

Ameur Gargouri ^{1,2}, Mohamed Karray ^{1,*}, Bechir Zalila ³ and Mohamed Ksantini ²

¹ ESME Research Lab, 94200 Ivry Sur Seine, France; gargouri.ameur@enis.tn

² ATISP Laboratory, ENETCom, University of Sfax, Sfax 3029, Tunisia; mohamed.ksantini@ipeis.usf.tn

³ ReDCAD Laboratory, ENIS, University of Sfax, Sfax 3029, Tunisia; bechir.zalila@enis.tn

* Correspondence: mohamed.karray@esme.fr

Abstract

This research work presents the design and the development of a cost-effective autonomous mobile robot for locating misplaced objects within unstructured indoor environments. The tools integrated into the proposed system for perception and localization are a hardware architecture equipped with LiDAR, an inertial measurement unit (IMU), and wheel encoders. The system also includes an ROS2-based software stack enabling autonomous navigation via the NAV2 framework and Adaptive Monte Carlo Localization (AMCL). For real-time object detection, a lightweight YOLO11n model is developed and implemented on a Raspberry Pi 4 to enable the robot to identify common household items. The robot's motion control is achieved by a fuzzy logic-enhanced PID controller that dynamically modifies gain values based on navigation conditions. Remote supervision, task management, and real-time status monitoring are provided by a user-friendly Flutter-based mobile application. Simulations and real-world experiments demonstrate the robustness, modularity, and responsiveness of the robot in dynamic environments. This robot achieves a 3 cm localization error and a 95% task execution success rate.

Keywords: autonomous mobile robot; ROS2; YOLO; SLAM; fuzzy logic PID; indoor navigation

1. Introduction

The last 10 years have witnessed the emergence of autonomous mobile robots (AMRs) as a pivotal technology across domestic, logistic, and industrial fields because of their capabilities of operating without dedicated infrastructure and performing tasks such as navigation, detection, and manipulation autonomously [1]. The rising demand for robotic service applications in hospitals, warehouses, smart manufacturing, and smart homes, where they enhance efficiency and reduce human intervention, has increased their adoption substantially [2,3]. Current autonomous mobile robots (AMRs) typically adopt a modular architecture that possesses three basic components: perception, localization, and navigation. Environmental perception is achieved by using sensors such as light detection and ranging (LiDAR), inertial measurement units (IMUs) and cameras. Simultaneous Localization and Mapping (SLAM) algorithms are also applied in order to construct environmental maps and localize the robot at specific locations within them. In navigation, global and local path planning algorithms such as A*, employed in global path planning, and the dynamic window approach (DWA), applied in local obstacle avoidance, guarantee smooth and adaptive mobility. The latest improvements in the field, such as the optimized integration of A* with DWA in Sánchez-Ibáñez et al. [4], Ishigami et al. [5] works, have aided in enhancing

path efficiency and reducing the computational overhead on embedded platforms, enabling real-time performance in complex environments.

In spite of these advancements, AMRs still encounter some flaws that restrict their widespread adoption. Because the computational performance of real-time SLAM and perception algorithms usually cannot compete with that of resource-constrained embedded systems, it is necessary to balance execution speed and precision. At a more basic level, it is also imperative to consider robustness in dynamic, unstructured environments, because it is typical for SLAM and navigating platforms to be initially deployed under static or semi-static conditions, with proportionally dropping performance in reaction to dynamic obstacles or varying lighting conditions. When deep learning models such as YOLO models are included, the computational demands are further increased, particularly in low-power hardware such as the Raspberry Pi. Moreover, achieving seamless human–robot interaction in household settings requires the consideration of user needs, as well as providing intuitive interfaces and ensuring that task execution capabilities are reliable, where current systems often fall short due to limited semantic understanding and adaptability. Finally, most AMRs are far from being cost-effective, which further limits the number of users who can afford them or use them in everyday environments.

This paper addresses these difficulties as it presents a modular AMR that is cost-effective and designed for unstructured indoor environments and is characterized by high-precision navigation and robust object localization. The system integrates a lightweight YOLO11n model for real-time object detection with an ROS2-based navigation stack using Cartographer and AMCL for SLAM and localization, complemented by a custom grid-based search algorithm. Both simulated and actual testing confirms the effectiveness and precision of the AMR. User interaction, including real-time monitoring and task management, is enhanced through the integration of a Flutter-based mobile application. To offer insights into the optimization of embedded systems for domestic robotics applications, this work critically assesses the trade-offs between computational efficiency and performance.

The remainder of this paper is organized as follows. Section 2 reviews related works on autonomous mobile robots, briefly summarizing the current designs, functions, and limitations. Section 3 gives an outline of the designed robot architecture, the hardware components, and the software platform, with an emphasis placed upon normal ROS2-based components as well as hand-designed parts. Section 4 presents the methodology, covering the modeling process, localization and path planning strategies, object detection module, custom searching strategy, actuator control by both classic PID and fuzzy logic-improved mechanisms and the development of a Flutter-based application that allows users to interact with the system and complete tasks. Section 5 presents the results and discussion, including simulation experiments, real-world experiments, and comparisons with other robotic platforms. Finally, Section 6 concludes this paper by summarizing key contributions as well as future research avenues.

2. Related Works

This work analyzes the performance and limitations of autonomous mobile robots (AMRs) available on the market that are designed for indoor navigation and object localization, characterized by their system architectures.

Singh et al. [6] tested the TurtleBot3, developed by Robotis [7], is an open-source AMR widely used in academic and research settings for indoor navigation tasks, such as basic object detection, path planning, and mapping. Equipped with a 360-degree LiDAR, a Raspberry Pi 4, and optional cameras, it leverages ROS2 for sensor integration and navigation, making it suitable for prototyping and education, as it is a relatively low-cost platform (approximately USD 500–700). Although this low-cost hardware has a modular

design and can achieve tasks like SLAM and obstacle avoidance in small-scale indoor environments, it suffers from substantial limitations. For example, the Raspberry Pi creates navigation delays of up to 1 s, resulting from its limited processing power when it comes to complex tasks like real-time SLAM or visual detection in dynamic settings. Furthermore, its standard sensor set suffers in low-feature scenes (e.g., feature-poor corridors), leading to localization inaccuracies of up to 0.5 m. Another significant weakness lies in the short time of operation based on its low-rated battery life, which does not surpass 1.5 h.

The Clearpath Jackal, developed by Clearpath Robotics [8], is a high-power AMR for research as well as industry applications, has the capabilities of both indoor and outdoor navigation through high-end onboard inertial-aided LiDAR, as well as an onboard Intel NUC microcomputer. Its localization errors are less than 0.1 m in structured environments due to the utilization of ROS2-based high-resolution 2D mapping SLAM. It offers reliable navigation under controlled indoor conditions, e.g., labs or warehouses, thanks to the device's robust construction as well as powerful hardware. However, its accessibility for domestic or small-scale applications is limited because of its high cost (approximately USD 10,000). Drift errors of 0.4 m or more occur when using the LiDAR-based SLAM in feature-scarce environments like long highways. Moreover, the Jackal consumes excessive power, which reduces its battery life to 2 h under heavy loads, and its bulky design makes it less suitable for navigation in tight indoor spaces [9].

The NVIDIA JetBot, developed by NVIDIA Corporation [10] is a low-cost, school- and hobbyist-oriented AMR that utilizes a Jetson Nano to enable GPU-assisted object detection and YOLO algorithms to execute functions such as tracking objects and navigating. The AMR is cost-effective because of its perception-based cameras and its low price (around USD 250). The JetBot achieves object detection accuracies of 80–85% for household items, with latencies of around 1 s in simple environments. However, when it operates in cluttered settings, its inference rates are limited to 0.5 Hz, rendering operation unstable at higher speeds because its processing power is limited. Because the JetBot relies on vision-based detection, its accuracy is reduced to (50–70%), according to Ramesh et al. [11], Kawakura and Shibasaki [12] works, in unstructured indoor spaces with varying lighting or occlusion. Moreover, its practical deployment is constrained due to its small battery capacity, which limits operation to under 1 h.

The Pozyx UWB-based AMR developed by Pozyx [13] uses ultra-wideband (UWB) positioning technology to precisely locate objects indoors, with its location accuracy ranging from 10 to 30 cm according to Crețu-Sîrcu et al. [14] work. This makes it useful for industrial activities like tracking inventories. Implemented as an ROS-integrated unit, it also accommodates sensor fusion, including with sensors such as LiDAR and IMUs, used for guidance under structured conditions. However, as Pozyx's AMR relies on pre-installed UWB beacons, it is costly and complex to set up and therefore unsuitable for domestic environments that do not have specialized infrastructure. Its localization performance becomes low in environments with metallic obstructions or multipath interference, where the errors rise to 0.3 m. The system has limited scalability for cost-sensitive applications like smart homes due to its high cost (nearly USD 5000 for a complete setup) and dependency on external hardware.

The MIR100, developed by Mobile Industrial Robots, tested by Gülденring et al. [15] and Soheilifar [16], is a commercial AMR that is designed for warehouse and industrial navigation and uses ROS-based navigation with A* for global planning and DWA for local obstacle avoidance. Equipped with LiDAR and cameras, it is able to perform tasks like material transport. However, because of the MIR100's high cost (approximately USD 20,000), its use is limited to industrial settings, making it inaccessible for domestic applications. The robot's utilization drops to 58.56% when it navigates at a 1.5 m/s speed,

which leads to excessive idle time, often waiting for the next item to be available. Moreover, its high power consumption, caused by its heavy reliance on powerful onboard computers, reduces the operational time to 8 h under heavy loads, and it navigates with difficulty in tight spaces because of its large footprint.

3. Proposed Robot Architecture

The autonomous mobile robot is designed to locate lost objects in cluttered indoor environments through the use of a tightly coupled hardware–software system. It integrates a robust hardware foundation with a Raspberry Pi 4 (Sony UK Technology Centre in Pen-coed, Wales), ESP32 microcontroller (Espressif Systems in Shanghai, China, at ZhangJiang Hi-Tech Park), and sensor suite running the ROS2 software (ROS 2 Humble Hawksbill available for Ubuntu Jammy (22.04)) to enable perception, movement, and human interaction. Figure 1 outlines the robot architecture as a whole and explicitly illustrates the interconnecting software/hardware components.




Figure 1. Robot architecture of the autonomous mobile robot.

3.1. Hardware Architecture

The system of an autonomous robot consists of a powerful collection of microcontrollers, sensors, actuators, and powering systems that allow it to achieve object localization and navigation effectively in unstructured indoor spaces. This work presents a detailed examination of all parts of the design, including their nature, interlinkages, and functions in the complete system.

Raspberry Pi 4 Model B Board: This is a quad-core Cortex-A72 (1.5 GHz) master microcontroller combined with a Broadcom VideoCore VI GPU. It runs with Ubuntu 22.04 with ROS2 Humble and is responsible for high-level processing tasks including sensor data fusion, SLAM, NAV2 navigation, and YOLO11n object detection. It communicates through the universal serial bus (USB) and the universal asynchronous receiver/transmitter (UART), processes data in real time, and runs ROS2 middleware to execute communication among nodes based on multicore processing to complete tasks in parallel.

ESP32 Microcontroller: The auxiliary microcontroller is an ESP32-WROOM-32, which has a dual-core Tensilica LX6 CPU (240 MHz) as well as 520 KB SRAM. It schedules real-time tasks through the Free Real-Time Operating System (FreeRTOS), as well as managing low-

level interfacing such as motor interfacing and sensor interfacing. The ESP32 receives data from the wheel encoders, as well as from the IMU sensor; produces pulse width modulation (PWM) signals to facilitate motor interfacing, and interfaces with the Raspberry Pi via UART.

LiDAR (YDLIDAR TSA, Shenzhen EAI Technology Co., Ltd., Shenzhen, China): Developed by YDLIDAR [17] and used by Hakkim et al. [18], consists of an 8-m rangefinder and a 360-degree time-of-flight (ToF) 2D LiDAR. Its ranging distance ranges from 0.12 to 8 m (80% reflectivity), its scanning frequency ranges from 5 to 8 Hz, and its ranging frequency is 3000 Hz. Cartographer-based SLAM and collision avoidance depend on the high-accuracy environmental perception provided by the LiDAR's high angular resolution of up to 0.72 degrees.

Stream Cam: YOLO11n-based object classification and detection uses images captured by a Logitech StreamCam (Logitech, Model VU0054, made in China) with a 78° FOV, 1080 p, 60 fps, and auto-focus. It is mounted 12 cm above the chassis and is connected to the Raspberry Pi via USB. At 5 Hz, an ROS2 camera node publishes raw image messages. Despite changing indoor light intensities, the high frame rate and high-resolution camera enable strong visual perception.

IMU (MPU6050, InvenSense Inc., San Jose, CA, USA): A 6-DOF inertial measurement unit (IMU), used by Sultan et al. [19], consists of a 3-axis accelerometer and a 3-axis gyroscope, which record orientation and angular velocity data at 100 Hz. It is connected to the ESP32 via the I2C protocol and transmits data at 50 Hz to the Raspberry Pi, used for localization and attitude estimation. The precision of pose estimation in dynamic settings is improved thanks to the incorporation of data from the IMU with wheel encoders and LiDAR data via AMCL.

Wheel Encoders: These incremental magnetic encoders are placed on the right and left drive wheels to obtain odometry information with a resolution of 204 pulses per rotation. They are connected to the ESP32 through digital GPIO pins and output the wheel velocity and position, which are sent to the Raspberry Pi at a rate of 50 Hz through UART and fed into an ROS2 topic. This high-accuracy odometry greatly helps in the precise localization of differential drive vehicles.

DC Motors (JGA25-370 DC Gear Motor with Encoder 600RPM) and Motor Driver: When employing a differential drive system composed of two 12V DC, 600 RPM motors, the top speed achieved will be 0.5 m s^{-1} as long as the wheel diameter is 65 mm. An L298N dual H-bridge DC motor driver (STMicroelectronics, Plan-les-Ouates, Switzerland) is employed to drive the motor and is connected to the ESP32 through PWM signals that operate at a frequency of 5 kHz such that smooth control of the speed is facilitated. Control of the ESP32 through the driver allows the vehicle to operate smoothly and safely.

Power Supply System: The system is powered by a 12 V, 1300 mAh LiPo battery, which allows about 2 h of operation under nominal load conditions. We have employed a buck converter to give a regulated 5V supply to the Raspberry Pi and ESP32 such that an uninterrupted supply is ensured. A 1-8S LiPo battery voltage detector with a buzzer is also placed to monitor the power system and prevent possible damage of the battery due to situations of overdischarge or overcharge.

Chassis: The robot chassis is a 30-cm-diameter circular Plexiglass structure with a dual-layer configuration. The top level is covered by the top layer, placed at a 10 cm height to keep the camera and LiDAR above ground-level obstacles, with the highest coverage for the sensors. The bottom layer is covered with the motors, cell battery, and microcontrollers. A low center of gravity is maintained as a condition for stability during navigation. A low weight is also maintained throughout the design (about 1.4 kg).

3.2. Software Architecture

The hierarchical mobile robot software architecture is composed of the ROS2 Humble middleware running Ubuntu 22.04 and combines standard and homegrown frameworks to preserve this hierarchical organization. The standard frameworks are the ROS2 middleware, the navigation stack, SLAM and localization modules, and object detection and classification using YOLO11n. The homegrown modules include the fuzzy logic-enhanced PID motor control and search algorithm with a user-definable grid system. The Flutter-based user interface constitutes an additional custom-developed module. Figure 2 illustrates how these frameworks combine to maintain the robot's sensor perception, planning, and control functionalities with fast and real-time operation across the Raspberry Pi 4 and ESP32 platforms.




Figure 2. Hierarchy of the software architecture for AMRs.

3.2.1. Standard Frameworks

ROS2 Middleware: The core of this system is the ROS2 Humble framework, which enables communication among nodes across the Data Distribution Service (DDS) to facilitate publish–subscribe interactions among distributed nodes with minimal latency. Central nodes are navigation, sensor drivers, perception, and control, which are executed on the Raspberry Pi 4 within a real-time executor to allow deterministic running. The configuration of the ROS2 environment uses Quality of Service (QoS) policies that ensure the reliability of important information, such as sensor information, and the best-effort transmission of unimportant information, such as visualization, while ensuring fault-tolerant communications [20].

Navigation Stack: The NAV2 system, developed by Open Navigation LLC [21], is an integrated module within the ROS2 package that offers a global and a local planner to achieve safe navigation. While the global planner achieves optimal paths in a two-dimensional occupancy grid costmap of resolution 0.05 m, which is published at 5 Hz, through the use of the A* algorithm, the local planner utilizes the dynamic window approach (DWA) to evade obstacles dynamically while following kinematically the robot's differential drive parameters of a maximum linear speed of 0.5 m s^{-1} and an angular velocity limit of 1.0 rad s^{-1} . Furthermore, NAV2 integrates LiDAR data to detect and avoid moving obstacles dynamically and publishes navigation commands through the `/cmd_vel` topic [22].

SLAM and Localization: Cartographer, developed by Google [23], an ROS2-compatible SLAM program, generates 2D occupancy grid maps of resolution 0.05 m from LiDAR scans. To predict the pose in an accurate manner, another ROS2 package named Adaptive Monte Carlo Localization (AMCL) combines information from the LiDAR and wheel encoders along with the IMU [24].

Object Detection and Classification: YOLO11n, developed by Ultralytics [25] and tested by Cherubin Szymon and Michał [26], is a lightweight deep neural network that is trained with the COCO dataset and runs from TensorFlow Lite on the Raspberry Pi and accepts RGB images of size 128×128 from the camera to detect and label common objects in daily life. The ROS2 perception node (yolo_detector) subscribes to the /camera/image_raw topic and publishes detection results (the /yolo/detections message containing bounding box information along with class labels and confidence scores) to engage with the search algorithm.

3.2.2. Custom-Developed Components

Search Algorithm: A costmap-optimized grid-based search algorithm is programmed as an ROS2 node (search_node) and divides the environment into $0.5 \text{ m} \times 0.5 \text{ m}$ cells to discretize and structure exploration and create waypoints. It is paired with NAV2 to provide navigation and YOLO11n to detect obstacles with time complexity of $O(n)$, where n is the number of waypoints. Failed points are re-explored three times through the action server of ROS2 to enhance navigation as it navigates through dynamically changing environments.

Motor Control: The ESP32 runs a fuzzy logic-enhanced PID controller with FreeRTOS tasks using C++. The controller is run at 50 Hz with wheel encoder feedback and provides PWM speed commands to the motors. Fuzzy logic rules with seven linguistic variables (e.g., Negative Big, Zero, Positive Small) for the error and derivative of error cause the PID gains to change dynamically to allow adaptable operation in changing circumstances.

User Interface: A Flutter-based mobile application, developed for the Android and iOS platforms, provides a user-friendly interface to control and monitor tasks. The app allows the visualization of maps in real time, the specification of navigation targets through touch feedback, and the initiation of searches for objects. The application interacts with ROS2 using the rosbridge_websocket node at 10 Hz with a low-latency update rate and with a JSON protocol to transfer commands and statuses.

The module structure was simulated and tested through Gazebo 11 and RViz 2 to ensure that all modules could run smoothly prior to field deployment.

4. Proposed Methodology

Robot development was carried out using the Agile Scrum methodology with the application of iteration cycles of design, development, and testing with short development sprints. Modeling was performed using URDF and Xacro in ROS2 with the realistic simulation of the chassis, sensors, and kinematics. The LiDAR-based SLAM and AMCL and the NAV2 stack were used for localization and navigation with dynamic path planning in real time. Object detection was performed with YOLO11n on a Raspberry Pi, performing object recognition of household items, while executing a grid-based search algorithm with systematic search. Motor control on the ESP32 used a fuzzy PID controller for precise speed regulation. A Flutter application was used for the real-time display of maps, setting goals, and object search control. Modular development was used so as to enable the iteration and integration of all subsystems with ease. Figure 3 illustrates the stepwise methodological workflow for the design, development, and realization of the proposed system.




Figure 3. Robot design, development, and realization flowchart.

4.1. Robot Modeling

The definition of the robot's physical structure, as illustrated in Figure 4, employed the Unified Robot Description Format (URDF) and Xacro within ROS2 Humble, allowing a modular and parametric representation of the robot's components.

- **Links:** The two-layer, circular chassis, composed of lightweight Plexiglass, is 30 cm in diameter. Maneuverability and stability are provided by a spherical link of two drive wheels (6.2 cm in diameter and 2 cm in width) plus a caster wheel. Sensor connections have YDLiDAR mounts. Components include the TSA LiDAR, Logitech StreamCam, MPU6050 IMU sensor, and wheels.
- **Joints:** The components are mounted on the chassis through constant joints to allow rotation, powered by 12V DC motors of 600 RPM. We use fixed joints, securing the LiDAR system (after elevation at 10 cm above the vehicle body) and camera (at 12 cm above the chassis, frontward-facing), and the IMU is placed close to the center of the chassis for high-accuracy orientation measurement.
- **Sensors:** The YDLiDAR TSA LiDAR is configured with a 360° field of view and a 10 m range, outputting point cloud data through an ROS2 sensor plugin. The Logitech StreamCam, with a field of view of 78° and a resolution of 1080 p, is parameterized with a plugin for a camera to take pictures. The MPU6050 IMU produces output of 6 DOF (accelerometer and gyroscope) and is interfaced as an ROS2 plugin for the IMU. Wheel encoders, at 204 pulses/rev, are defined to produce odometry data via a differential drive plugin.
- **Coordinate Frames:** The transform tree has base_link as the origin of the vehicle, odom as the odometry tracker, and map as the global frame. Sensor-specific frames (lidar_link, camera_link, imu_link) are defined with offset with respect to base_link so that spatial transformation is correct, and the transformation is published at 50 Hz through tf2.

The differential drive system is modeled kinematically. The wheelbase is 25 cm and the maximum linear velocity is 0.5 m s^{-1} and is computed based on the 600 RPM motors and wheel diameter. Parameterization is achieved via Xacro macros, facilitating the easy modification of the dimensions and sensor locations. The model is also simulated dynam-

ically for correct collision and physics (e.g., friction, inertia) using Gazebo. RViz is also used for the visualization of the robot state, checking for correct transform propagation and alignment of the sensor state. URDF/Xacro files are also streamlined for interaction with the navigation and SLAM stacks of ROS2, ensuring seamless integration with Cartographer and NAV2.




Figure 4. URDF model of the robot.

4.2. Localization and Navigation

The robot operates autonomously in unstructured indoor environments thanks to the Simultaneous Localization and Mapping (SLAM) and navigation system, while robot perception, localization, and path planning are achieved by leveraging ROS2 Humble's navigation stack. The ROS2 localization and navigation flowchart, shown in Figure 5, describes the sequential pipeline of data processing and decision-making in the robot's navigation system.




Figure 5. ROS2 localization and navigation flowchart.

- **SLAM:** A LiDAR-based 2D SLAM algorithm generates high-resolution occupancy grid maps from streamed data from the YDLiDAR TSA LiDAR, which provides the point cloud data (blue color variation in Figure 6) at 7 Hz with a 360° view angle and 10 m range. Cartographer employs scan matching and loop closure to construct maps at a 0.05 m resolution with a local SLAM window of 5 s and a submap resolution of 0.05 m, being optimized for indoor environments with dynamic obstacles. Generated maps are published onto the ROS2 topic /map for use by the navigation stack.
- **Localization:** Adaptive Monte Carlo Localization (AMCL) provides robust pose estimation by combining data from the YDLiDAR TSA LiDAR, MPU6050 IMU (6-DOF orientation at 100 Hz), and wheel encoders (odometry at 200 Hz, sent at 50 Hz over the ESP32). AMCL runs a particle filter of 2000 particles updated at 10 Hz with a probabilistic motion model motivated by the differential drive kinematics (25 cm wheelbase, with a maximum speed of 0.5 m s^{-1}), yielding a localization error of 0.03 m in real-world tests. The tf2 library also manages the transform tree with coordinate frames (e.g., map, odom, base_link, lidar_link), with transformations published at 50 Hz.
- **Global Planning:** The NAV2 framework's global planner computes optimal, collision-free trajectories from the current pose of the robot to meet user-specified goals by employing the A* algorithm. A* acts on a 0.05 m resolution 2D occupancy grid costmap from the Cartographer map (black outline in Figure 6), which the costmap_server updates at 5 Hz. The planner uses a Euclidean distance heuristic, generating paths at a 0.1 m resolution, accounting for the robot's footprint by scaling up obstacles by 0.2 m for a safety margin. The global path is published onto the ROS2 topic /plan for use as input for local planning.
- **Local Planning:** The dynamic window approach (DWA) is the sole local planner used by NAV2, evaluating feasible velocity commands within a dynamic window constrained by the robot's kinematic limits (maximum linear velocity 0.5 m s^{-1} , maximum rotational velocity 1.0 rad s^{-1}), as well as by how closely the robot is positioned to the obstacles. DWA evaluates a sampling of potential velocity commands (linear and rotational) within a dynamic window based on criteria such as how closely the robot approaches obstacles, convergence with the global path, and how much progress the command makes toward the goal. The command that makes the maximum progress while ensuring safety within the environment is chosen. Real-time LiDAR updates the local costmap used as the movement planner around dynamic obstacles.
- **Feedback and Monitoring:** NAV2 provides mechanisms for providing feedback for the monitoring and debugging of the navigation process.
 - **Path Visualization:** The local and global plans are published in RViz2 on the /plan topic, as shown in Figure 6, which allows developers to see the planned path (green line) and identify issues.
 - **Costmap Visualization:** Both global and local costmaps are published and portrayed in RViz, including the costs for obstacles, inflation fields, and the robot's footprint, thereby facilitating the diagnosis of navigation failures in simulations.
 - **Navigation Status:** The bt_navigator provides status reporting by means of the /navigate_to_pose action feedback, progress, and current state, as well as several error conditions (e.g., "goal unreachable" errors)




Figure 6. Visualization in Rviz2.

4.3. Object Detection

The object detection system specifically classifies and detects common objects in real time and allows the robot to locate lost items in disordered interior settings. This work addresses the specific problem of restricted mobility by enabling the recovery of goods that are beyond reach. YOLO11 is the latest variant of the You Only Look Once (YOLO) computer vision models. The model is optimized for real-time processing and has an unparalleled trade-off in terms of processing speed, efficiency, and accuracy that cannot be matched by any other algorithm. The YOLO11n model has the lightest and fastest architecture. YOLO11n can successfully identify common objects like bottles, books, remote controllers, and small tools, which can help assistive robots to retrieve objects that are out of reach for a person with mobility limitations. Object recognition by YOLO11n enables the search algorithm to focus on target spots. YOLO11n runs on the Raspberry 4 Model B (4 GB RAM, quad-core cortex-A72 processor running at 1.5 GHz) as a compromise considering computational power and resource constraints.

To assess the computational efficiency of the current state-of-the-art object detection algorithms to be used in embedded platforms, we analyzed and contrasted the performance results obtained in various studies and community challenges on the Raspberry Pi 4 Model B with 4 GB of memory. The performance of several algorithms, namely YOLO11, YOLOv8, EfficientDet-Lite, Faster R-CNN, and SSD-MobileNet, in terms of their average latency and FPS throughput is summarized in Table 1. The performance data clearly indicate that there is a trade-off between the accuracy of object detection and real-time performance while implementing AI algorithms on low-power hardware platforms. Out of the evaluated models, YOLO11-Nano offers the best compromise between mean latency

and throughput in object detection tasks and can easily fulfil the performance needs of our proposed misplaced object search algorithm on the Raspberry Pi 4.

The proposed architecture introduces a distributed processing system through cooperation between the Raspberry Pi and ESP32 microcontroller, as shown in Figures 2 and 3, whereby ESP32 runs FreeRTOS and performs three tasks at once, including managing the motors of the robot, receiving motor control signals from the Raspberry Pi, and sending sensor information such as IMU and encoder readings back to the Raspberry Pi.

The delegation of low-level control and sensor data acquisition to the ESP32 optimizes the processing time on the Raspberry Pi, thereby making it more efficient than in the present scenario, as it is based on a monolithic architecture and has high system latency due to processing on the Raspberry Pi itself. Additionally, the system achieves flawless, dynamic, and continuous data transfer through the combination of ROS2 on the Raspberry Pi and FreeRTOS on the ESP32, making the overall robotic system free from latency and more efficient due to its ability to run low-level actuation tasks effectively at low latency.

Table 1. Comparison of object detection model performance on Raspberry Pi 4 Model B (4 GB).

Model	Mean Latency (ms)	Throughput (FPS)
YOLO11-Small (YOLO11s) [27]	300–1200	0.8–3.3
YOLOv8-Nano (YOLOv8n) [28]	400–700	1.4–2.5
EfficientDet-Lite0 (TFLite, INT8) [29,30]	50–150	6–20
Faster R-CNN (ResNet-50 + FPN) [30]	2000–6000	<0.5
SSD MobileNet v1 (TFLite, INT8) [28]	≈209	≈4.8
The proposed model	100–300	3.5–10

Model Selection and Implementation: YOLO11n, as a result of its optimality for devices with limited resources, is used without modification through TensorFlow Lite for streamlined Raspberry Pi inference. The pretrained YOLO11n model, which can be downloaded from Ultralytics and is trained on the COCO dataset comprising 80 object classes, provides the capabilities for detecting and recognizing home objects such as telephones, remote controls, keys, and cups. The model accepts input images of 128×128 pixels for a smaller computational overhead while providing sufficient detail for detection. The model provides class-labeled bounding boxes and confidence tensors, reaching a mean precision level of up to 95% for target objects when deployed in indoor settings, as verified against a custom test collection of domestic objects.

Data Acquisition and Preprocessing: The Logitech StreamCam takes images, providing a 78° wide angle of view at a 1080 p resolution at 60 fps. The camera is located 12 cm above the chassis and tilted upward to capture the forward scene. Images from it are scaled down by bilinear interpolation to the 128×128 pixel input required by YOLO11n, with colour normalization used for further detection robustness when the indoor lighting varies. The system publishes visual information through the /camera/image_raw topic at a rate of 5 Hz.

Inference Pipeline: The object detection code from the YOLO11n model serves as an ROS2 node (yolo_detector) optimized to run on the Raspberry Pi. This node subscribes to the /camera/image_raw topic, processes images into the necessary input format of 128×128 pixels, and conducts inference directly from the unmodified YOLO11n model. The inference runs at 1.25 Hz, with average latency of 800 ms, once every frame from the beginning of preprocessing until the termination of postprocessing. The model produces bounding boxes, class labels, and confidence scores filtered by a confidence filter of 0.5 and a non-maximum suppression (NMS) filter using an IoU score of 0.4, removing duplicate detections. Results are published onto the /yolo/detections topic as a custom ROS2 mes-

sage bearing bounding box coordinates, class IDs, and confidence scores, complemented by the grid-based search algorithm for target localization.

4.4. Custom Search Algorithm

Current methods, like the ones used by Cui [31], typically utilize hybrid methods that blend SLAM-oriented global planning and local avoidance schemes like the dynamic window approach (DWA). While efficient for goal-oriented movement, these tend to focus on path optimization and not thorough environment coverage. Hence, they can neglect certain areas and lower the search accuracy in systematic search scenarios. For this reason, a custom grid-based search algorithm is introduced, guaranteeing equal exploration, ease of waypoint handling, and sound integration with perception and navigation modules. This custom search algorithm permits the comprehensive exploration of unstructured indoor scenes for the purpose of detecting lost objects, while being suitable for use within the robot's navigation and perception components. It applies a grid-based approach to generating waypoints. A grid-based approach in search algorithms typically entails a grid or matrix representation of the problem's search space. Each state or location in the problem's environment is represented by a corresponding node or cell in the grid. The strategy for the proposed solution relies on a search via a grid pattern that divides the environment into traversable waypoints. Unlike similar search algorithms that explore the environment by only utilizing geometry patterns to search for waypoints or targets of interest, the strategy presented integrates perception and search by utilizing the YOLO11n detector to analyze images while planning with the NAV2 technique for motion to explore each waypoint. This offers comprehensive environment coverage, employing YOLO11n for object detection purposes and NAV2 for navigation.

Waypoint Generation: The environment is discretized into a $0.5\text{ m} \times 0.5\text{ m}$ grid of cells according to the 2D occupancy grid map from Cartographer (resolution of 0.05 m). Waypoints are provided to the center of each of the free cells (e.g., cells that are not marked as obstacles in the map), allowing for navigable goals. The grid is created on the fly from the `/map` topic, per waypoint, represented as a list of (x, y) coordinates in the `map` frame. For an environment of 5 m by 3 m , approximately 11 waypoints are generated (yellow points in Figure 7), but this is reduced when obstacles such as furniture or walls are present. The waypoints are shown in Figure 7, with their spread over an example indoor map with obstacles.

Algorithm Implementation: The search algorithm, which is executed as an ROS2 node (`search_node`), is given by Algorithm 1. It initializes two lists of visited and failed waypoints. It iterates over unvisited waypoints, choosing the closest waypoint using the Euclidean distance from the robot's current position (parsed from AMCL by the `tf2` library). The NAV2 stack with A* for the global planner and DWA for the local planner controls the robot towards the selected waypoint. At each waypoint, the `yolo_detector` node processes images from the Logitech StreamCam, publishing the detection results on the `/yolo/detections` topic. If there is a target object detection (the confidence score of the class of an object is greater than 0.5), the algorithm publishes its position in the `map` frame and terminates. If the robot cannot reach the waypoint due to dynamic obstacles or kinematic constraints, the waypoint is recorded as failed, and up to three attempts are made before skipping. The time complexity of the algorithm is $O(n)$, where n is the number of waypoints, with constant-time detection by YOLO11n (800 ms per waypoint).

Integration and Operation: The `search_node` lists itself as a subscriber of the `/map` topic of the occupancy grid, `/yolo/detections` of object detection outputs, and `/tf` of the robot's pose. The `search_node` publishes navigation goals on the `/nav2/goal` topic by communicating with NAV2.

Algorithm 1 Grid-Based Search Algorithm

```

1: procedure SEARCHLOOP
2:   Generate waypoints for grid cells
3:   Initialize visited and failed waypoint lists
4:   while target not found do
5:     Select unvisited waypoint
6:     Navigate to waypoint using NAV2
7:     Process YOLO11n detections
8:     if target detected then
9:       Publish target location
10:      break
11:    end if
12:    if navigation fails then
13:      Mark waypoint as failed
14:      Retry up to 3 times
15:    end if
16:  end while
17: end procedure

```

Validation and Performance: The algorithm was also validated in actual test experiments on an indoor setup. The grids enclosed the volume within less than 3 min at a 100% success rate for target object detection.

The coverage map, illustrated in Figure 7, shows the overlapping circular fields of view (FOVs) (blue circles) that are utilized to guarantee full room coverage.




Figure 7. Coverage map: circular FOVs overlapping for room coverage.

4.5. Motor Control

The smooth movement of the indoor autonomous robot requires motor control to ensure effective and precise motion. Smooth following of the path, collision prevention, stability despite changing scenarios, and the prevention of wild riding or high power usage are ensured by effective speed control. The speed of the DC motor is calculated by DC motor encoders that are connected to the ESP32 by two pins (encA and encB) for the purpose of detecting rotational movement. Accurate speed data are needed for PID control, which enables real-time motor adjustment for smooth robot motion.

- **Quadrature Encoding and Position Tracking:** The quadrature encoder offers two signals, A (encA) and B (encB), to monitor the motor's rotation. They are two waves phase-shifted by each other. Imagine two friends clapping their hands, but one a

fraction of a second after the other. This phase shift enables the ESP32 (a very small microcontroller, or a minicomputer) to determine both the speed and direction of turning, as illustrated by Figure 8. The Init routine configures encA and encB as input pins and attaches an interrupt (encoderISR) to encA. An interrupt occurs when there is a state change in encA, so that the system will notice at once rather than continuously checking. This interrupt occurs for any change in encA's state, so that encoder ticks (small rotational steps) are quickly recognized. The Update routine, prompted by the interrupt, reads encA's and encB's states to determine the direction of turning. The direction detection logic is given below.

- If encB is high (e.g., “on” or 1) and encA is high, then the wheel rotates clockwise (the direction when turning a clock's hands to the right).
- If encB is high and encA is low (e.g., “off” or 0), the wheel rotates counterclockwise (opposite direction), so the position count decreases by -1 .
- If encB is low and encA is high, the wheel turns counterclockwise; hence, the position decreases by -1 .
- If encB is low and encA is low, the wheel turns clockwise; hence, the position increases by $+1$.

The position variable is the encoder's cumulative count of ticks, increasing as we turn clockwise and decreasing as we turn counterclockwise, giving a real measurement of the motor's position. It is analogous to maintaining a bidirectional step count for wheel rotation.




Figure 8. Quadrature encoder signals for detection of rotational direction.

- **Speed Calculation:** The GetfilteredVelocity procedure calculates the motor speed in revolutions per minute (RPM). RPM is the number of complete revolutions per minute that the motor undertakes, similar to the number of turns per minute in which a wheel performs a full revolution. It obtains the current position (GetPosition) by means of a critical region with a mutex (muxMotor) to ensure thread safety. A mutex is similar to a lock on a door—only one process can enter at a time to avoid conflicts when multiple tasks are running simultaneously. The position change (Δpos) is then calculated by subtracting the previous one from the current one. The time elapsed (Δt) in seconds is obtained by the use of micros() (a procedure yielding the time in microsec, being very precise). The initial speed in counts per second (counts/s) is determined by Formula (1):

$$v = \frac{\Delta \text{pos}}{\Delta t} \quad (1)$$

This is converted to RPM using Formula (2):

$$v = \frac{\text{counts/s}}{204} \times 60 \quad (2)$$

where 204 represents the encoder's counts per revolution (a parameter specific to this model), and 60 is used to convert s^{-1} to min^{-1} .

- **Low-Pass Filtering for Speed Smoothing:** The measured speed is then filtered by a low-pass filter with a cut-off frequency of 25 Hz, thereby reducing measurement noise, and smooth control of the motor speed will be ensured. Here, “noise” denotes undesired signal fluctuations like radio static that could occur due to electrical noise or mechanical vibrations. A low-pass filter is similar to a sieve that accepts slow, gradual variations but rejects rapid, jerky variations, providing stable input to the PID controller. The filtering process uses the recursive Equation (3):

$$v_{\text{filt}} = 0.854 \times v_{\text{filt}} + 0.0728 \times (v + v_{\text{prev}}) \quad (3)$$

where v_{filt} is the filtered speed (the smoothed version), v is the current speed (raw measurement), and v_{prev} is the previous speed reading. In this formula, the old filtered value is combined with a fragment of the new data, creating a gradual smoothing effect.

- **State Update:** The algorithm continually reverts back to the previously noted position, velocity, and time (previousPosition, previousVelocity, previousTime), ensuring continuous speed computation. This process is analogous to a reset operation on a stopwatch, performed after the completion of every lap in preparation for the next one.

This procedure provides a constant speed measure, required for regulation by the PID controller.

4.6. Motor Speed Regulation Process

The procedure of motor speed regulation involves maintaining the DC motors at the desired rotational speed by utilizing speed information from the measurement of the encoders. In short, automatically changing the motor's power maintains the motor at the desired speed, similarly to a car's cruise control that maintains the speed despite the fact that the driver does not hold the pedal continuously. A manual PID tuning procedure was initially applied, followed by an adaptive fuzzy logic PID controller for higher efficiency when the situation changes. The following sections describe the procedure of regulation by comparing the two while describing the fuzzy logic implementation.

4.6.1. Manual PID Tuning

The initial solution was manual PID tuning, with the PID gains, K_p (proportional gain), K_i (integral gain), and K_d (derivative gain), tuned by trial and error towards a stable response. PID is a classic control system resembling a thermostat that will activate the heating depending upon how far from the set point the temperature is. The tuning consisted of the following:

- Initializing gains (e.g., $K_p = 1.0$, $K_i = 0.0$, $K_d = 0.0$) and observing the motor's response while being driven by a setpoint speed (the desired speed).
- Setting K_p until the motor speed was jittery around the setpoint, after which it was reduced by a small amount until a stable steady-state response was achieved. K_p will react to instantaneous errors, e.g., pushing harder if farther away.

- Tuning K_i to decrease the steady-state error (ongoing deviation from the aim) as the motor approached the setpoint over time. K_i is responsible for accumulated errors, such as providing an additional thrust if slow for a while.
- Fine-tuning K_d to dampen oscillations but remaining low enough so as not to enhance noise. K_d takes the rate of change into consideration, such as braking if entering too quickly.

The result of this strategy is shown in Figure 9, where we observe the motor speed often moving beyond the goal setpoint, causing overshoot (up to 220 RPM while the goal is 200 RPM). This produced rather rough speed transitions, notably at varying loads (such as when we ascend), with a settling time of approximately 0.6 s and noticeable oscillations.

There are shortcomings of manual PID tuning. Gains that are constant do not adjust rapidly when there are quick load transfers or changing terrains, causing the following:

- Overshoot and oscillations, like those observed in the above chart, which will destabilize the robot (rendering the robot jerky or unstable);
- Inconsistencies when in the presence of non-linear motor responses with unpredictable responses or with shifting environmental parameters;
- A labor-intensive tuning process that requires continual fine-tuning for different situations.




Figure 9. Manual PID tuning results.

4.6.2. Adaptive Fuzzy Logic PID Controller

In response to the drawbacks of manually tuning the PID, an adaptive fuzzy logic-based PID controller was devised. This technique, inspired from Nassim and Abdelkader [32] work, changes the PID gains automatically by varying the error (e) and its rate of change (dedt), providing a more robust solution. Fuzzy logic is akin to embedding human-like reasoning into programs, operating with “shades of gray” instead of binary yes/no values, by virtue of rules such as “if it’s warm, cool a little”. This enables the system to cope with robot load variations or surface variations.

Selection of Membership Functions: Trapezoidal membership functions were chosen for the system. Their flat top ensures a stable and firm response for large errors, necessary for quick corrective action in motor speed control. Smooth variation in the control output is ensured by their overlapping parts to prevent jerky motion during load transients. Their low complexity qualifies them for the resource-constrained ESP32-based platform, yet offers stability, responsiveness, and efficiency.

4.6.3. Selection of Fuzzy Rules

The fuzzy logic system begins by selecting the fuzzy rules, designed to make the response aggressive for quick error correction and a stable motor speed. They follow if-then structures, similar to natural language reasoning, for example, “if the error is large and changing rapidly, then adjust strongly.” The error (e) and its derivative (dedt), scaled to the range $[-1.0, 1.0]$ (scaled to lie within -1 and 1 for consistency), are defined via

seven linguistic terms: Negative Big (NB), Negative Medium (NM), Negative Small (NS), Zero (Z), Positive Small (PS), Positive Medium (PM), and Positive Big (PB). They are of the like-label type: NB means a large negative error (far too slow), and PB means far too fast. They are defined by the trapezoidal membership function (*trapezoidalMF*), with parameters selected such that there is a small amount of overlap for smooth operation (no jumpiness). An example is NB from -1.0 through -0.6 and NM from -0.8 through -0.2 , defined by the *calculateMemberships* routine. The membership functions are depicted in Figure 10.




Figure 10. Membership functions.

4.6.4. Fuzzy Rule Table

The fuzzy rules are implemented as 7×7 matrices (rule_Kp and rule_Ki) in the fuzzyInference procedure to adjustably set the PID controller gains for motor speed control such that membership pairs for e and $dedt$ are converted into proportional (Kp) and integral (Ki) gains with $Kd = 0.0$ (no derivative term here). They are selected such that they increase extremely aggressively for large errors (e.g., 6.0 for NB-NB, i.e., harsh correction for large deviations) but decrease rapidly for small errors (e.g., 0.6 for PB-PB, with subtle corrections for close values), scaled by certain factors (1.5 for Kp, 7.0 for Ki) for increased responsiveness. The tables are as follows.

- **Rule Table for Kp ($\times 1.5$):**

dedt e	NB	NM	NS	Z	PS	PM	PB
NB	6.0	5.8	5.6	5.4	5.2	5.0	4.8
NM	4.8	4.6	4.4	4.2	4.0	3.8	3.6
NS	3.6	3.4	3.2	3.0	2.8	2.6	2.4
Z	2.4	2.2	2.0	1.8	1.6	1.4	1.2
PS	2.2	2.0	1.8	1.6	1.4	1.2	1.0
PM	2.0	1.8	1.6	1.4	1.2	1.0	0.8
PB	1.8	1.6	1.4	1.2	1.0	0.8	0.6

- **Rule Table for Ki ($\times 7.0$):**

dedt e	NB	NM	NS	Z	PS	PM	PB
NB	6.0	5.8	5.6	5.4	5.2	5.0	4.8
NM	4.8	4.6	4.4	4.2	4.0	3.8	3.6
NS	3.6	3.4	3.2	3.0	2.8	2.6	2.4
Z	2.4	2.2	2.0	1.8	1.6	1.4	1.2
PS	2.2	2.0	1.8	1.6	1.4	1.2	1.0
PM	2.0	1.8	1.6	1.4	1.2	1.0	0.8
PB	1.8	1.6	1.4	1.2	1.0	0.8	0.6

The rules favor fast error correction for large errors (e.g., high Kp and Ki for NB-NB, such as quickly using the brakes if out of control) and weaker gain strength as the error

becomes close to zero (e.g., low Kp and Ki for PB-PB), offering stability around the setpoint (target speed). The error, error derivative, and Kp gain interaction are graphically depicted in Figure 11.




Figure 11. Fuzzy rule surface for Kp.

4.6.5. Fuzzy Inference System (FIS) Rules

The FIS rules operate as follows. Membership degrees for e and $dedt$ are computed by the “fuzzyInference” function using “calculateMemberships”, and the minimum operator is applied as shown in Formula (4)

$$\text{weight} = \min(\text{eMembership}[i], \text{dedtMembership}[j]) \quad (4)$$

to determine the firing strength of each rule (the extent to which the rule applies, selecting the smaller of the two memberships). The weighted sums are calculated as shown in Formulas (5)–(7):

$$\text{sum_kp} = \sum_{i=0}^6 \sum_{j=0}^6 (\text{weight} \times \text{rule_kp}[i][j] \times 1.5) \quad (5)$$

$$\text{sum_ki} = \sum_{i=0}^6 \sum_{j=0}^6 (\text{weight} \times \text{rule_ki}[i][j] \times 7.0) \quad (6)$$

$$\text{sum_weights} = \sum_{i=0}^6 \sum_{j=0}^6 \text{weight} \quad (7)$$

The final gains are calculated using Formula (8):

$$\text{kp} = \frac{\text{sum_kp}}{\text{sum_weights}}, \quad \text{ki} = \frac{\text{sum_ki}}{\text{sum_weights}}; \quad \text{sum_weights} > 0 \quad (8)$$

If $\text{sum_weights} = 0$, the previous gains are retained. This center-of-gravity method ensures a smooth gain transition.

The fuzzy logic PID controller controls the motor speed to an even greater extent through the calculation of the Kp and Ki gains from the normalized change in error with the aid of trapezoidal membership functions for input clustering into seven linguistic sets (NB through PB). It also includes gain determination rule tables with the aggressive setting of gains and a weighted average defuzzification technique (converting fuzzy outputs into

crisp numerical values), which is superior to interactive PID tuning under fluctuating loads.

4.6.6. Implementation Process

The procedure then continues as follows.

- **Error Calculation:** The ComputeFuzzy method uses error (e) calculation by subtracting the desired setpoint (from ROS 2/cmd_vel, robot's navigation system command) from the filtered speed (from GetfilteredVelocity). The time difference from the last call is obtained with micros() for appropriate derivative calculation accuracies; there is also a safeguard against division by zero (to prevent mathematical errors).
- **Error Derivative Calculation:** The derivative of the error (dedt) is computed using Formula (9):

$$\text{dedt} = \frac{\text{currentError} - \text{previousError}}{\text{deltaTime}} \quad (9)$$

This makes the controller capable of responding to rapid speed changes.

- **Normalization:** To fit the error (e) and error derivative (dedt) to the fuzzy logic system and to allow corner cases like zero or negative setpoints (where the scaling values will be inside the fuzzy "box"), the Map function normalizes them to range from -1.0 to 1.0 .
- **Membership Calculation:** Function *calculateMemberships* utilizes trapezoidal membership functions (trapezoidalMF) to project the normalized error (e) and error derivative (dedt) into seven fuzzy sets as a representation of each one's degree of membership of NB.
- **Fuzzy Inference:** The fuzzy inference mechanism utilizes already defined rules to compute adjusted gains (Kp, Ki, Kd), in which Kd is 0.0.
- **PID Computation:** The ComputeFuzzy method calculates the control signal using the adjusted gains according to Formula (10):

$$\text{controlSignal} = Kp \times e + Kd \times \frac{de}{dt} + Ki \times \int e dt \quad (10)$$

There is anti-windup protection on the integral term and a limit on the output so that it does not saturate (it holds signals within safe ranges).

- **Output of Motor Control:** The control signal from the L298N motor driver (a chip controlling the direction and speed of the motors) is transmitted to the DC motors, with the PWM signal being changed. The ESP32 adjusts the last error (e) and time before the next iteration.

4.6.7. Results and Performance

The fuzzy PID outputs according to Formula (10) are graphed in Figure 12 and compared to those of manual PID. There is no overshoot in the plot, with the speed closely tracking the setpoint (e.g., stable at 200 RPM without exceeding it).




Figure 12. Fuzzy logic PID results.

The fuzzy PID allows smoother speed transitions, especially with loads, with the settling time now being less than 0.3 s and almost 0 oscillations, which enhances the navigation performance (the robot becomes more reliable and efficient in moving).

This adaptive mechanism is appropriate in terms of motor speed control, efficiently dealing with real-time changes as well as continuous communication with the ROS 2 navigation stack.

4.7. Flutter Application Development

Developing a Flutter app is essential for good user interaction when dealing with an autonomous robot, offering a user-friendly interface to monitor tasks in real time. This provides instructions, robot statuses, and movement checks, being easy to integrate into houses. It becomes challenging when using an inefficient app to perform tasks or inspect responses. The proposed mobile application offers a complete control and monitoring interface for the autonomous robot. Users can remotely control the robot, track its motion on the real-time map, and follow sensor information updates from the robot. Using the application, users can start the navigation task or execute object searches. The application also enables users to receive real-time updates on the environment of the robot and the progress of its tasks. Flutter, developed by Google, is an open-source user interface development kit used to create natively compiled mobile, web, and desktop applications from a single codebase. It comes packaged with the Dart programming language to handle fast apps with richly styled, smooth motion interfaces [33]. Flutter's widget-based architecture makes it possible to easily create novel, responsive interfaces, which is critical in developing an intuitive robot control interface.

Flutter Application Development for Robot Control:

The Flutter app was designed to present an easy-to-use interface from which to instruct and follow the self-driving robot. It communicates with the robot through a WebSocket server on a laptop, which converts commands to the ROS2 stack running on the robot. The app's development included the following notable aspects.

- **User Interface Design:** The easy-to-understand and user-friendly interface of the app allows users to start tasks (e.g., begin navigating, search objects), see the online status (e.g., current real position), and see the robot's map. The interface, shown in Figure 13, provides buttons to move the robot in manual mode, provides a view of the map to navigate, and has displays of the status, so users without technical expertise can find it easy to communicate with the robot.
- **Task Management:** There is a type of task management offered by the software that enables users to assign tasks to the robot, such as initializing the robot position (green

- arrow), moving to a specific location (red arrow) or beginning an object search. These tasks are then encapsulated in the form of ROS2 commands (e.g., `/cmd_vel` when dealing with moving) and are transmitted through the WebSocket server. This is similar to providing a robot with a list of tasks that it executes on its own.
- **Real-Time Monitoring:** The program retrieves the status of the robot and its location information (blue triangle), sensor information, and task accomplishment information like remained waypoints to finish the search task (yellow dots) in real time through ROS2 topics that are subscribed to through WebSocket subscriptions, so that users can view the movement of the robot, such as monitoring a tracked vehicle across a surveyed area.
 - **ROS2 Integration:** The application utilizes an ROS2-WebSocket bridge to enable connectivity with the controller program that is integrated into the robot system. To ensure easy and smooth interactions, this bridge translates application commands into ROS2 syntactic form and then forwards them to the robot.




Figure 13. Main interface of the Flutter application for robot control.

The app has been tested and proven in simulations and in practical scenarios with adequate user communication and interaction. It demonstrated the low-latency issuance of commands and stable reporting of feedback, so that the robot could be issued commands by non-technical users.

5. Results and Discussion

5.1. Simulation Tests

The simulation tests, with Gazebo and RViz2, confirmed the efficacy of the ROS2-based autonomous mobile robot's control. Gazebo is a 3D physics simulation software program with realistic environment simulations and robotic sensor modeling for system testing, while RViz2 is an ROS2-specific visualization software program with ROS2 real-time displays of sensor information, maps, and robot poses. Gazebo was chosen given its strong physics engine and realistic robotic modeling capabilities, with the accurate emulation of LiDAR, IMU, and wheel encoder data, while RViz2 was chosen given its real-time displays, enabling efficient debugging and the monitoring of performance, a combination that is routinely applied in robotics research. Takaya et al. [34] also used Gazebo and RViz as ROS simulation softwares. As shown in Figure 14, the integration of these sensors ensured

correct localization via Adaptive Monte Carlo Localization (AMCL), with a 95% success rate during mapping and navigation through dynamic indoor simulation worlds with moving objects. The dynamic window-based (DWB) controller ensured effective path planning, with the robot succeeding in target waypoints, with an average positional error of 3 cm. The YOLO-based object detection platform, with testing through simulations, properly identified domestic items with a 95% success rate under varying lighting levels. The Flutter app with the WebSocket bridge interface ensured low-latency command responses (<100 ms) and a full real-time view of the robot's map and status with the smooth handling of tasks. However, the simulation sluggishness under very complex worlds suggests the need to use optimally equipped computing hardware.




Figure 14. Initial simulation test in Gazebo.

5.2. Real-World Tests

As illustrated in Figure 15, real-world tests validated the robot's performance in a realistic indoor environment, with the results broadly replicated in simulations but also with understandable practical restrictions. The test environment was set up following an inverted L-shape arrangement. Room dimensions were 4 m by 2.5 m, with a recess of 1.25 m by 1.25 m at one corner, making its geometry an inverted L. Environment illumination was achieved by employing two white LED lamps of 3000 lumens. Real-sized objects were used to simulate realistic scenarios (doors, bottle, walls, etc.). The ROS2 platform using LiDAR and ultrasonic sensing achieved a 95% success rate in navigating uneven surface floors and cluttered rooms, with the adaptive fuzzy logic PID controller implemented on the ESP32 reducing the motor settle time to less than 0.3 s and the overshoot to less than 5 RPM from a setpoint of 200 RPM. The YOLO model with the Raspberry Pi achieved a 95% detection rate for domestic items but fell slightly short in poor-light scenarios due to noise on the sensing peripherals. Successful remote monitoring and operation were ensured with commands running within an average time of 120 ms via the Flutter app. Problems were found with occasional WebSocket latency during high network utilization and negligible localization drift with highly reflective surface locations, with potential for improved integration with sensors and noise reduction.




Figure 15. Real-world testing.

5.3. Comparison

Compared with commercially available autonomous mobile robots such as the Clearpath Jackal UGV (Clearpath Robotics, Kitchener, ON, Canada), MiR100 (Mobile Industrial Robots, Odense, Denmark), NVIDIA JetBot (NVIDIA, Santa Clara, CA, USA), and TurtleBot3 (ROBOTIS, Seoul, Korea), Table 2 shows that the constructed robot achieves similar performance in indoor domestic applications while having a considerably cheaper and more lightweight structure.

According to Robotis (2023), the TurtleBot3 achieves a 94% indoor navigation success rate with ROS2 and 360° LiDAR, but it costs USD 1800 and weighs 1 kg, limiting its applicability to longer domestic missions [7]. Alternatively, the suggested robot with a YDLiDAR TSA provides slightly better performance, with a 95% success rate along the course and a 3 cm position error, while having the same 1.4 kg weight but with a fraction of the cost (USD 369).

Clearpath Robotics (2023) reports that the Jackal UGV possesses 96% reliability in lab operation given the high-precision LiDAR and IMU implementation, but, given its high cost of USD 15,000+ and heavy 17 kg weight, it is impractical for home applications [8]. Likewise, the MiR100, according to Mobile Industrial Robots (2023), performs best in industry, with a 97% success rate in navigation, but with a 70 kg weight and USD 20,000+ cost; thus, domestic application is highly limited [15].

For perception, YOLO11n, with 90% accuracy, is also adopted in the developed robot, which is an improvement compared to the NVIDIA JetBot, which employs SSD MobileNetV2, with 85% accuracy in cluttered indoor settings [11].

However, the MiR100 still possesses the greatest communication latency (50 ms) due to the proprietary interface, while the constructed system implements a Flutter application using a WebSocket interface with 120 ms latency, which reflects optimization potential within the communication layer.

Overall, the designed robot is impressive, with a low price (USD 369), light weight (1.4 kg), and equal balance of navigation and perception capabilities, making it well suited for unstructured indoor environments as compared with heavier, more expensive commercial alternatives.

Table 2. Comparison of autonomous mobile robots.

Robot	Navigation Success Rate (%)	Positional Error (cm)	Weight (kg)	Cost (USD)	Object Detection Accuracy (%)	Latency (ms)
TurtleBot3 [7]	94	3 *	1	1800	N/A	120 *
Clearpath Jackal [8]	96	5 *	17	15,000+	N/A	80 *
MiR100 [15]	97	3 *	70	20,000+	92 *	50
NVIDIA JetBot [11]	69.0	8 *	0.8 *	250	85	150 *
Proposed robot	95	3	1.4	369	95	120

(*) Estimated values derived from similar systems or general robotics data owing to missing reference information.

6. Conclusions

Applying autonomous robots to human assistance is challenging and requires user-friendly interfaces to help users to perform everyday tasks such as picking and navigation, as well as affordability to ensure large-scale adoption and adaptation to unstructured domestic environments. This work resolves these challenges by designing a new autonomous domestic robot tailored specifically to assistance indoors. Cost-competitive hardware consisting of a YDLiDAR TSA scanner, an ESP32 microcontroller, an ultrasonic sensor unit, and a Raspberry Pi-based sensor-actuator unit is combined with a high-level sophisticated software framework on ROS2. The solution results in **95% navigation success**, a **3 cm positional error**, and **95% YOLO-based object detection accuracy**, all while having a cost of just **USD 369**, which is much lower than those of comparable commercially available robots such as the TurtleBot3 or MiR100.

Moreover, the system integrates a customized grid-based search method, an adaptive fuzzy logic PID motor controller to control the momentum, and a Flutter-based mobile app that allows easy interaction even when the communication latency is **120 ms**. The experiments demonstrate that the proposed robot is competitive in terms of its cost-performance-extendability ratio, which qualifies it to be deployed in living rooms.

In the future, improving sensor fusion methods, network performance refinement, and the utilization of machine learning to enable real-time task prioritization will further improve the robot's performance, opening the door to broader applications in home environments.

Author Contributions: Conceptualization, A.G.; Software, A.G.; Validation, M.K. (Mohamed Karray); Writing—original draft, A.G.; Writing—review & editing, M.K. (Mohamed Karray), B.Z. and M.K. (Mohamed Ksantini); Supervision, M.K. (Mohamed Karray), B.Z. and M.K. (Mohamed Ksantini). All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Keith, R.; La, H.M. Review of Autonomous Mobile Robots for the Warehouse Environment. *arXiv* **2024**, arXiv:2406.08333. [CrossRef]

2. Gonzalez-Aguirre, J.A.; Osorio-Oliveros, R.; Rodríguez-Hernández, K.L.; Lizárraga-Iturralde, J.; Morales Menendez, R.; Ramírez-Mendoza, R.A.; Ramírez-Moreno, M.A.; Lozoya-Santos, J.d.J. Service Robots: Trends and Technology. *Appl. Sci.* **2021**, *11*, 10702. [CrossRef]
3. Siciliano, B.; Sciavicco, L.; Villani, L.; Oriolo, G. *Robotics: Modelling, Planning and Control*, 1st ed.; Advanced Textbooks in Control and Signal Processing, Springer: London, UK, 2009; pp. XXIV, 632. eBook ISBN 978-1-84628-642-1/Hardcover ISBN 978-1-84628-641-4/Softcover ISBN: 978-1-84996-634-4. [CrossRef]
4. Sánchez-Ibáñez, J.R.; Pérez-del Pulgar, C.J.; García-Cerezo, A. Path Planning for Autonomous Mobile Robots: A Review. *Sensors* **2021**, *21*, 7898. [CrossRef] [PubMed]
5. Ishigami, G.; Nagatani, K.; Yoshida, K. Path Planning and Evaluation for Planetary Rovers Based on Dynamic Mobility Index. In Proceedings of the 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems, San Francisco, CA, USA, 25–30 September 2011; pp. 601–606. [CrossRef]
6. Singh, D.; Trivedi, E.; Sharma, Y.; Niranjan, V. TurtleBot: Design and Hardware Component Selection. In Proceedings of the 2018 International Conference on Computing, Power and Communication Technologies (GUCon), New Delhi, India, 28–29 September 2018; pp. 805–809. [CrossRef]
7. Robotis. TurtleBot3: Specifications and Performance. Available online: <https://www.robotis.us/turtlebot-3/> (accessed on 6 November 2025).
8. Clearpath Robotics. Jackal UGV: Technical Overview. 2025. Available online: <https://clearpathrobotics.com/jackal-small-unmanned-ground-vehicle/> (accessed on 7 October 2025).
9. Mobile Manipulation Platform for Autonomous Indoor Inspections in Low-Clearance Areas. In Proceedings of the 25th International Conference on Advanced Vehicle Technologies (AVT), Boston, MA, USA, 20–23 August 2023; Volume 1. [CrossRef]
10. NVIDIA Corporation. JetBot: An Open Source AI Robot Platform. 2024. Available online: <https://jetbot.org> (accessed on 9 October 2025).
11. Ramesh, G.; Jeswin, Y.; Rao, D.R.; Suhaag, B.; Uppoor, D.; Kiran Raj, K.M. Real Time Object Detection and Tracking Using SSD Mobilenetv2 on Jetbot GPU. In Proceedings of the 2024 IEEE International Conference on Distributed Computing, VLSI, Electrical Circuits and Robotics (DISCOVER), Mangalore, India, 18–19 October 2024. [CrossRef]
12. Kawakura, S.; Shibasaki, R. Deep Learning-Based Self-Driving Car: JetBot with NVIDIA AI Board to Deliver Items at Agricultural Workplace with Object-Finding and Avoidance Functions. *Eur. J. Agric. Food Sci.* **2020**, *2*, 1–9. [CrossRef]
13. Pozyx. Industrial Automation: UWB RTLS for Robotics. 2025. Available online: <https://www.pozyx.io/solutions/industrial-automation> (accessed on 7 October 2025).
14. Crețu-Sircu, A.L.; Schiøler, H.; Cederholm, J.P.; Sircu, I.; Schjørring, A.; Larrad, I.R.; Berardinelli, G.; Madsen, O. Evaluation and Comparison of Ultrasonic and UWB Technology for Indoor Localization in an Industrial Environment. *Sensors* **2022**, *22*, 2927. [CrossRef] [PubMed]
15. Gündenring, R.; Görner, M.; Hendrich, N.; Jacobsen, N.J. Learning Local Planners for Human-aware Navigation in Indoor Environments. In Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Las Vegas, NV, USA, 25–29 October 2020. [CrossRef]
16. Soheilifar, S. Optimizing MiR100 Performance in Internal Logistics Through Digital Twin-Based Dynamic Speed Adjustment. 2025. Available online: <https://webthesis.biblio.polito.it/secure/36032/1/tesi.pdf> (accessed on 6 November 2025).
17. YDLIDAR. YDLIDAR TSA — 2D LiDAR Product Specifications. Product Datasheet (PDF); Detailed Specs: 0.12–8 m Range, 6 Hz Scan, 0.72° Resolution. Available online: <https://www.worldrobotconference.com/uploads/exfile/video/s4tbfh.pdf> (accessed on 7 October 2025).
18. Hakkim, F.; Rasheed, H.; Varsha, M.C.; Amal, K.A.; Minimol, B. Lidar Based Autonomously Navigating Robot. In Proceedings of the 2024 1st International Conference on Trends in Engineering Systems and Technologies (ICTEST), Kochi, India, 11–13 April 2024; pp. 1–6. [CrossRef]
19. Sultan, J.M.; Zani, N.H.; Azuani, M.; Yusop, A.M. Analysis of Inertial Measurement Accuracy using Complementary Filter for MPU6050 Sensor. *J. Kejuruter.* **2022**, *34*, 959–964. [CrossRef]
20. Dust, L.J.; Persson, E.; Ekström, M.; Mubeen, S.; Secleanu, C.; Gu, R. Experimental Evaluation of Callback Behavior in ROS 2 Executors. In Proceedings of the 2023 IEEE 28th International Conference on Emerging Technologies and Factory Automation (ETFA), Sinaia, Romania, 12–15 September 2023; pp. 1–8. [CrossRef]
21. Open Navigation LLC. Navigation2 (Nav2)—Tutorials and Documentation. 2025. Available online: <https://docs.nav2.org/> (accessed on 7 October 2025).
22. M J, A.K.; Babu, A.V.; Damodaran, S.; James, R.K.; Murshid, M.; Warriar, T.S. ROS2-Powered Autonomous Navigation for TurtleBot3: Integrating Nav2 Stack in Gazebo, RViz and Real-World Environments. In Proceedings of the 2024 IEEE International Conference on Signal Processing, Informatics, Communication and Energy Systems (SPICES), Kottayam, India, 20–22 September 2024; pp. 1–6. [CrossRef]

23. Google Cartographer ROS Integration—Documentation. 2022. Available online: <https://google-cartographer-ros.readthedocs.io> (accessed on 7 October 2025).
24. Dwijotomo, A.; Abdul Rahman, M.A.; Mohammed Ariff, M.H.; Zamzuri, H.; Wan Azree, W.M.H. Cartographer SLAM Method for Optimization with an Adaptive Multi-Distance Scan Scheduler. *Appl. Sci.* **2020**, *10*, 10347. [CrossRef]
25. Ultralytics. Model Comparisons (YOLO Family). 2025. Available online: <https://www.ultralytics.com> (accessed on 7 October 2025).
26. Cherubin Szymon, K.W.; Michał, S. YOLO object detection and classification using low-cost mobile robot. *Przegląd Elektrotechniczny* **2024**, *100*, 29–33. [CrossRef]
27. Ultralytics. Quick Start Guide: Raspberry Pi with Ultralytics YOLO11. 2025. Available online: <https://docs.ultralytics.com/guides/raspberry-pi/> (accessed on 30 October 2025).
28. Alqahtani, D.K.; Cheema, A.; Toosi, A.N. Benchmarking Deep Learning Models for Object Detection on Edge Computing Devices. *arXiv* **2024**, arXiv:2409.16808. [CrossRef]
29. Kamath, V.; Renuka, A. Performance Analysis of the Pretrained EfficientDet for Real-time Object Detection on Raspberry Pi. In Proceedings of the 2021 International Conference on Circuits, Controls and Communications (CCUBE), Bangalore, India, 23–24 December 2021; pp. 1–6. [CrossRef]
30. Zagitov, A.; Chebotareva, E.; Toshev, A.; Magid, E. Comparative analysis of neural network models performance on low-power devices for a real-time object detection task. *Comput. Opt.* **2024**, *48*, 242–252. [CrossRef]
31. Cui, J. A search and rescue robot device realization based on ROS operating system combined with visual algorithm. *J. Phys. Conf. Ser.* **2023**, *2492*, 012020. [CrossRef]
32. Nassim, M.; Abdelkader, A. Speed Control of DC Motor Using Fuzzy PID Controller. *arXiv* **2021** arXiv:2108.05450. [CrossRef]
33. Boukhary, S.; Colmenares, E. A Clean Approach to Flutter Development through the Flutter Clean Architecture Package. In Proceedings of the 2019 International Conference on Computational Science and Computational Intelligence (CSCI), Las Vegas, NV, USA, 5–7 December 2019; pp. 1115–1120. [CrossRef]
34. Takaya, K.; Asai, T.; Kroumov, V.; Smarandache, F. Simulation environment for mobile robots testing using ROS and Gazebo. In Proceedings of the 2016 20th International Conference on System Theory, Control and Computing (ICSTCC), Sinaia, Romania, 13–15 October 2016; pp. 96–101. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Article

Belt Sanding Robot for Large Convex Surfaces Featuring SEA Arms and an Active Re-Tensioner with PI Force Control

Hongjoo Jin ¹, Chanhyuk Moon ², Taegyun Kim ^{3,*} and TaeWon Seo ^{2,*}

¹ Department of Convergence Mechanical Engineering, Hanyang University, Seoul 04763, Republic of Korea; icacorp471@gmail.com

² School of Mechanical Engineering, Hanyang University, Seoul 04763, Republic of Korea; andrew04004@hanyang.ac.kr

³ Department of Mechanical Engineering, BK21 FOUR ERICA-ACE Center, Hanyang University, Ansan 15588, Republic of Korea

* Correspondence: taegyunkim@hanyang.ac.kr (T.K.); taewonseo@hanyang.ac.kr (T.S.)

Abstract

This study presents a belt sanding robot for large convex surfaces together with a proportional–integral force control method. Sanding belt tension strongly affects area coverage and spatial normal-force uniformity on large curved surfaces; existing approaches typically use fixed tool positions or lack active tension regulation, which limits coverage and makes force distribution difficult to control. The mechanism consists of two series elastic actuator arms and an active re-tensioner that adjusts belt tension during contact. In contrast to a conventional belt sander, the series elastic configuration enables indirect estimation of the reaction force without load cells and provides compliant interaction with contact transients. The system is evaluated on curved steel plates using vertical scans with a belt width of 50 mm and a drive wheel speed of 300 rpm. Performance is reported for two target curvature values, namely 0.47 and 1.37, with five trials for each condition. The control objective is a constant normal force along the contact, achieved through proportional–integral control of the arms for normal-force tracking and the re-tensioner for belt tension regulation. To quantify spatial force uniformity, the distribution rate is defined as the ratio of the difference between the maximum and minimum normal forces to the maximum normal force measured across the belt–workpiece contact region. Compared with a simple belt sander baseline, the proposed system increased the sanded area coverage by 31.85%, from 62.20% to 94.05%, at the curvature value of 0.47, and by 8.49%, from 81.21% to 89.70%, at the curvature value of 1.37. The distribution rate improved by 113% at the curvature value of 0.47 and by 16.7% at the curvature value of 1.37. Under identical operating conditions of 50 mm belt width, 300 rpm, and five repeated trials, these results indicate higher area coverage and more uniform force distribution relative to the baseline.

Keywords: sanding robot; series elastic actuator; force regulation; proportional–integral control

1. Introduction

Sanding, which is the process of making something smooth by rubbing it with something rough, is a crucial process in various industries, ranging from manufacturing to maintenance and repair. It serves multiple purposes, such as improving the aesthetic appeal of surfaces, preparing for subsequent painting or coating, and removing unwanted materials such as rust or other contaminants for safety. However, traditional sanding processes rely predominantly on manual labor, which presents several challenges. First, manual sanding

exposes workers to continuous vibrations that can have detrimental musculoskeletal health effects over time. Second, maintaining consistency while manually sanding large areas is difficult because of the inherent variability in human performance, which often results in uneven finishes and makes it almost impossible to restore over-sanded surfaces. Finally, manual sanding generates dust that can harm the respiratory health of workers if not properly managed with additional facilities or protective gear. Owing to the challenges associated with manual sanding, there is an increasing need for automation. Automated solutions can not only address the aforementioned issues but also enhance productivity and reduce operational costs. To meet such needs, this study proposes a mechanism adapting a series elastic actuator (SEA) and a re-tensioner as a new way to control the tension and normal force of a sanding robot for large-area convex surfaces, as shown in Figure 1.




Figure 1. Proposed large-area convex belt sanding robot.

Abrasion robots were grouped by mechanism into belt, orbit, and wheel modules. Prior work spans belt grinding and polishing, orbit polishing heads, and wheel grinding for seams and edges [1–3]. Most reported systems target small parts or locally planar and strip-like regions such as turbine blades, weld seams, and furniture panels [4,5]. By contrast, large, doubly curved surfaces demand both high structural capacity and precise coordination of position and contact force over a wide interface, which remains difficult in practice [6–8]. Recent work on compliant/SEA-equipped grinding end-effectors and force-planning on complex curved parts further underscores the need for active force/tension regulation in large-area finishing [9–11]. Performance is often summarized by coverage, roughness, or mark suppression; only a few studies explicitly quantify how normal force is distributed spatially across large contact areas [12]. This gap matters because non-uniform force across the belt–workpiece interface degrades coverage and produces local over-removal on convex geometries.

Against this backdrop, our study frames a concrete control problem for large-area convex sanding: regulating belt tension and the spatial distribution of normal force across the full belt width. We adopt series elastic actuators on the sanding arms to estimate reaction force indirectly at the contact without load cells and to provide compliance to contact transients under vibration and dust typical of belt abrasion. In our previous passive-arm ex-

periments, distinct groove marks indicated sudden jumps in normal force at the contact [13]. Integrating SEAs reduces the mechanical output impedance of the arms and improves tolerance to impact loads, mitigating those jumps. In addition, an active re-tensioner is introduced to maintain appropriate belt tension along the contacting span, addressing belt tracking and tension variability that conventional arrangements do not actively control. These design choices directly target the two limitations above—limited instantaneous coverage and non-uniform spatial force—and set up the measurable objectives evaluated in this paper.

The re-tensioner mitigates the challenges of belt detachment and escapes from its designated track when transitioning to a multi-wheel belt system. Owing to the limitations stemming from their primary purpose, only two wheels that did not contact the target surface were left capable of guiding the belt track. Along with the attached idlers, this configuration allowed the maintenance of optimal belt tension, ensuring that the belt remained securely on track throughout the sanding process. In doing so, we not only mitigated the risks associated with belt detachment but also optimized the overall efficiency of our multi-wheel belt system for improved performance.

Large-area finishing on doubly curved workpieces is limited by two concrete factors. Prior systems typically present a narrow instantaneous contact patch and operate from fixed tool positions, which constrains the achievable coverage across the belt width during a single pass. In addition, the spatial distribution of normal force across the belt–workpiece interface is often non-uniform, leading to locally excessive removal and visible banding rather than consistent surface quality.

This study addresses these gaps with measurable objectives evaluated under controlled operating conditions. The first objective is to achieve full-width contact across the nominal belt span during each vertical pass so that the belt engages the entire width rather than a narrow strip. The second objective is to promote uniform normal-force distribution across the contact region, reducing spatial variation relative to a conventional belt sander.

The evaluation focuses on steel plates with convex curvature representative of the target application. Operating conditions reported in the following sections include a belt width of 50 mm, vertical feed of 50 mm/min, and drive wheel speed of 300 rpm. Stating the objectives in terms of belt width coverage and force uniformity links the mechanism and control choices to practical finishing outcomes on large convex surfaces.

The sanding belt tension is controlled using the proportional–integral (PI) control method to ensure optimal belt tightness throughout the sanding process. This allows consistent and uniform sanding and avoids surface warping. PI control is also used to regulate the reaction force applied to the target surface. This enables adaptive and compliant interactions between the sanding belt and convex target surfaces, accommodating variations in curvature, and preventing uneven surfaces from excessive forces. By adjusting the applied force based on real-time feedback, damage to delicate surfaces can be prevented while maintaining effective and well-distributed material removal.

The remainder of this paper is organized as follows. Section 2 describes the robot mechanism, including the kinematics and dynamic modeling. Section 3 describes the robot control. Section 4 presents the experimental setup and corresponding results. Finally, Section 5 concludes the paper.

2. Robot Design

The proposed sanding robot is described in this section. After explaining the overall structure, the sanding module, re-tensioner, and the three-degrees-of-freedom (3-DOF) moving stage are described.

2.1. Overall Structure

Figure 2 illustrates the overall robot design suggested for large-area convex sanding. The robot is largely composed of the following three components: a belt sanding module, a re-tensioner, and a moving stage. Sanding begins when the target object approaches the moving stage. Once contact is made, the SEA arms of the sanding module are actuated to follow the target surface curvature while tracking the reference force. During this process, the tension throughout the contacting side of the belt is adjusted to obtain the proper tension for sanding by moving the re-tensioner idlers with a ball screw.




Figure 2. Overall design of the proposed sanding robot.

2.2. Sanding Module

Figure 3 shows the sanding module, which is capable of sanding the target with a sanding belt. The sanding belt is rotated by a driving BLDC motor (TM13-D2021-S-4P-K-SG DC24V 2KW, motorbank, Seoul, Republic of Korea), controlled by an MD2K motor driver, and supported by the two contact wheels, two idlers on the sanding arms, and two additional idlers on the re-tensioner. A DC power supply (TS6015A 60 V/15 A 900 W, Toyotech, Seoul, Republic of Korea) is used to ensure stable power delivery to the driving wheel. For the sanding arm, a reaction force-sensing SEA mechanism (RFSEA), in which the spring is located before transmission, is applied to measure and control the torque applied to the arms [14]. The sanding arms measure the torque exerted on both arms, calculate the reaction force from the target surface using linear springs and rotary encoders, and track the desired force. The upper parts, indicated in blue in Figure 3, are responsible for the angle input, while the lower parts (indicated in red) measure the actual rotated angle. The servo motors on each arm (B250 Brushless Servo, Highest, Seoul, Republic of Korea) were chosen such that they could generate more than 2 Nm at a voltage of 6 V for an estimated reaction force of 10 N. Four linear springs with lengths of 40 mm and spring constant k of 3.34 N/mm with an initial tension F_0 of 12.75 N were used and aligned in the rotational direction. A rotary encoder (E50S8-8000-3-V-24, Autonics, Busan, Republic of Korea) with a resolution of 8000 was used to measure the arm rotation angle and calculate the reaction force at each arm.




Figure 3. Overall sanding module design.

2.3. Re-Tensioner

As shown in Figure 4, the re-tensioner pulls the belt from the target object or vice versa to maintain the proper tension. An adjuster moving back and forth with a ball screw with a pitch lead of 4 mm, a diameter of 12 mm, and a DC motor (PGM42-4566, 1/4, 12 V, motorbank, Seoul, Republic of Korea) was used. Two idlers are attached to the actuator to guide the belt and prevent it from detaching.




Figure 4. Overall re-tensioner module design.

2.4. Moving Stage

In addition to the horizontal motion that is shown approaching the target surface, vertical motion was also considered because it would be unreasonable to neglect it when the operating time is an important factor in sanding large-area convex surfaces. A BLDC motor (PG42BL4261E1, 1/49, 24 V, motorbank, Republic of Korea) with a deceleration rate of 1/49 was used to achieve a high torque of 5 Nm, given it should move the entire robot by itself.

2.5. Sanding Robot Kinematics

Figure 5 shows how the angles in each SEA arm are structured. θ_1 is the total angle rotated by the SEA arm on wheel D while θ_2 indicates the same on wheel E. θ_{SV1} is the angle rotated by the servo motor and θ_{SP1} is the angle generated by the lengthened spring. θ_1 and θ_2 can be separately and actively adjusted for a more evenly distributed surface

normal force, leading to uniform sanding on the contact area between the sanding belt and convex target surface.




Figure 5. Angles on the SEA arm.

Figure 6 shows how the torque from the servo motor is transferred to the contact point on point D. The torque applied by a servo motor, M_{sv} , is transferred through the planetary gear set. In this gear set, a ring gear with 90 teeth is stationary and fixed to the frame. The sun gear has 60 teeth, and two planet gears with 15 teeth each form a 1:2.5 rotational angular speed ratio for the carrier, which is a line connecting the centers of the planet gears and sun gear. This enables the carrier angle to range from 0° to 64° and remains sufficiently strong to maintain θ_{sv} against the moment generated from F_T .




Figure 6. SEA arm force balance with two-point support, with annotations marking the applied, belt, and reaction forces: (a) side view and (b) top view

The angle rotated by the servo motor (θ_{sv}) and the angle measured by the rotary encoder at the base (θ_1) differ because of the presence of springs, which contribute to an additional angle (θ_{sp}). The notation and units used in this section are summarized in Table 1. These springs connect the lower and upper parts of the SEA arm. The springs also help the tip of the arm to not press the target surface too tightly, which can lead to a jump in the normal force applied, resulting in uneven sanding quality.

Table 1. Symbols and units used in the sanding robot kinematics.

Symbol	Description	Units
R_A	Measured force on drive wheel A	N
R_D, R_E, R_M	Measured normal forces at D, E, and midpoint M on the target surface	N
F_D, F_E	Calculated forces at D and E on the SEA arm	N
F_T	Resultant belt force on wheel D or E due to tension	N
T	Belt tension	N
k	Spring constant	N/mm
X_{sp}	Spring length	mm
ΔX_{sp}	Spring elongation	mm
L	Distance between re-tensioner and drive wheel center	mm
r_t	Radius of the target surface	mm
r_D, r_E	Radii of wheels D and E	mm
r_{gear}	Distance between sun and planet gear centers	mm
θ_1	SEA arm rotation at wheel D (base/encoder angle)	rad
Ψ_1	Contact/wrap geometry angle	rad
α	Spring-linkage angle defined	rad
M_{sv}^*	Gear-amplified servo torque	N·mm
M_{sp}	Torque due to spring force	N·mm
M^*	Net internal torque at the SEA arm joint	N·mm
M_T	Torque induced by F_T	N·mm

The relationship between the torques asserted on one SEA arm is expressed as follows:

$$M_{sv}^* + M_{sp} = M^* \quad (1)$$

$$M^* + M_D + M_T = 0 \quad (2)$$

where M_{sv}^* is the torque amplified by the gear, M_{sp} is the torque caused by spring force, M_T is the torque caused by F_T , and M_D is torque originated from force F_D .

Figure 7 shows the overall belt configuration during the sanding process, where point O is the center of the target surface. F_D and F_E are the forces exerted on D and E, respectively, on the contact points with the target surface; R_A is the force exerted on the drive wheel from the belt tension; and R_D , R_E , and R_M are the reaction forces on the target surface, where M is the midpoint between D and E on the target surface. For the analysis associated with Equations (3)–(7) and the free-body diagrams, we adopt three working assumptions that align with the reported operating regime. The belt is treated as inextensible on each straight or wrapped segment during a pass, so tension within a given segment is uniform and variations arise from geometry rather than stretch. The support wheels at C and D are bearing-mounted and do not transmit appreciable resisting moments; therefore, support moments are neglected in the two-point contact case. The contact geometry is evaluated within the convex curvature range represented by the tested plates in this study; behavior outside this range may require additional compliance or contact modeling. These assumptions justify the quasi-static simplifications used in the force balance and are referenced alongside Figure 7.

$$\Psi_1 = \sin^{-1}\left(\frac{\overline{OD}}{\overline{DD_Y}}\right) = \sin^{-1}\left(\frac{378 - 90 \sin \theta_1}{r_t + r_E}\right) \quad (3)$$

$$\alpha = \frac{\pi}{3} + \Delta\theta_{sp} \quad (4)$$

$$\Delta X_{sp} = (40^2 - 15^2)\left(1 - \sqrt{2 - 2 \cos \alpha}\right) \quad (5)$$

$$F_T = 2T \sin \Psi_1 \left(\cos^{-1} \left(\frac{R+r}{90} \right) - \theta_1 - \Psi_1 \right) \quad (6)$$

$$F_D = \frac{M^* + M_T}{90 \sin(\theta_1 - \psi_1)} \\ = \frac{2(k\Delta X_{sp} + F_0)r_{gear} \cos\left(\frac{1}{2}\alpha\right) + F_T \left(\cos\left(\frac{1}{2}\alpha\right) - \Psi_1 \right)}{90 \sin(\theta_1 - \Psi_1)} \quad (7)$$

Wang et al. [15] introduced a method to calculate the force F_T exerted on an abrasion belt when a circular target object is pressed against it. The bearings on wheels C and D allow the moment caused by belt tension to be negligible in the two-point contact situation. Using this method, M_D can be calculated, leading to the calculation of F_D and F_E , which are with respect to the servo angles on each arm, θ_1 and θ_2 , and the radius of the target surface, r_t , in (7), where 90 mm is the distance between the centers of wheels C, D, E, and F, and 37.5 mm is the distance between the center of the sun gear and that of one planet gear. Because linear coil springs were used, the equation for estimating F_D is nonlinear.




Figure 7. Overall belt configuration of the two-point support during sanding.

Figure 8 shows details of the reaction force with respect to the sanding arm rotated angle for two springs on a planar surface. The SEA arms affect the force exerted on the plane target surface at different θ_1 and k . The SEA arm moved toward a 3 mm thick steel plate for 9 s and then moved in the opposite direction at the same constant speed. The reaction force was measured using four load cells (SBA-100L, CAS) at each corner of the rectangular target surface at 70 signals per second. The appropriate sanding operation time and force were determined in the pre-conducted experiments. There was a force oscillation in the 0.6 N gap when the spring constant was low, resulting in the use of at least two springs, as described above. It also shows that the measured force value is within an error range of 5% of the predicted value, proving its suitability for sensing.

Figure 9 shows a comparison between the mechanism's calculated force F_E in the y -direction and the measured R_E value with respect to time, which represents the SEA arms' sensing ability on convex target surfaces without any control method applied. With a maximum gap of 2% in the maximum value area where actual sanding occurs, these results prove the effective sensing ability of the SEA sanding arms.




Figure 8. Reaction force with respect to the rotated angle for the two-spring cases.




Figure 9. Comparison of calculated and measured R_E over time in the y -direction for two convex target surfaces. Dotted lines indicate calculated values and solid lines indicate experimental measurements.

3. Control System

Simple PI control associated with an active SEA system is more desirable than impedance control when the goal is to maintain a constant force over time with limited dynamic interaction [16]. In classical interaction control, impedance control shapes a desired dynamic relation between motion and force at the contact, which is advantageous for tasks requiring active rendering of inertia, damping, or stiffness and for broadband interaction [17]. Our application targets constant-force tracking under limited interaction dynamics: the belt–workpiece contact is operated in a quasi-static regime, the series elastic actuators provide intrinsic compliance, and the control bandwidth is restricted by belt dynamics and structural vibration. Under these assumptions, a force-regulating PI outer loop on the SEA arms, combined with PI tension regulation for the re-tensioner, achieves steady-state accuracy without the parameter identification and gain scheduling typically required to render a stable target impedance over a wide frequency range. The chosen architecture therefore prioritizes low-frequency force accuracy and repeatability rather than interaction rendering, which aligns with the finishing objective of uniform material removal over large convex areas. Furthermore, without the need to change the input variables, such

as the material stiffness or damping of the target surface, PI control is more suitable for robust control.

PI controllers are mainly used for processes where precise and steady-state accuracy is required, whereas rapid change is insignificant [18,19]. Adding an integral component to the proportional error terms, as in P control, eliminates the steady-state error and provides balance between the response speed and error correction. Table 2 summarizes the selected PI gains for the SEA arm and the re-tensioner together with the unified control targets.

Table 2. PI gains and control targets.

Loop	K_P	K_I	Control Targets
SEA arm	0.3	0.1	Steady-state error $\leq 10\%$, Settling time ≤ 0.5 s
Re-tensioner	0.3	0.05	

The proportional and integral gains were determined by experiment-driven step-response tuning on the instrumented setup. We first applied P control and increased K_P until fast tracking was achieved without sustained oscillation, then introduced a small integral term to eliminate steady-state error under the reported operating conditions. The gains used in all experiments are $K_P = 0.3$, $K_I = 0.1$ for the SEA arm force loop and $K_P = 0.3$, $K_I = 0.05$ for the belt tension loop. During belt operation, evaluation used a simple moving-average filter as described to address periodic vibration from the belt joint.

The relationship between R_A , which is the force exerted on the drive wheel by the belt tension, and R_M , which is the force at the midpoint of the target surface from the belt tension, was empirically determined, as shown in Figure 10. There was insufficient tension on the belt to rotate when R_A was below 15 N, and the belt rotation was stopped because the tension was too high when R_A exceeded 130 N.

For each curvature, the mid-contact force R_M was fitted as a quadratic function of the drive wheel proxy R_A (both in newtons). The fitted relations are

$$R_M = -0.002 R_A^2 + 0.591 R_A + 7.262 \quad (\rho = 1.37, R^2 = 0.988), \quad (8)$$

$$R_M = -0.00002 R_A^2 + 0.016 R_A + 10.933 \quad (\rho = 0.47, R^2 = 0.992). \quad (9)$$

where ρ denotes the curvature value used in the tests.

These second-order least-square fits are used within the observed operating range ($R_A = 15\text{--}130$ N) reported in Figure 10. These relationship functions were chosen and used as an empirical block on the block diagram shown in Figure 11.

The drive wheel proxy R_A and mid-contact force R_M were related by a quadratic fit, explicitly limited to the observed operating range ($R_A = 15\text{--}130$ N). Although the number of sweeps per curvature was limited, the fits exhibited high quality and the residuals were approximately zero-mean without systematic trends. The mapping was then used as the reference-generation block in the control architecture, and its consistency was verified in subsequent closed-loop experiments using independent load cell measurements at points D, M, and E. Within the stated range, the mapping provides practically sufficient accuracy for reference generation.

The control of the force acting on the contact point on the SEA arm with the target surface and the PI control applied to the BLDC motor that controls L (the distance between the re-tensioner adjuster and the center of the drive motor) follows the form shown in Figure 11. The belt tension reference $R_{A,\text{ref}}$ was obtained from the empirical $R_A - R_M$ map in Figure 10 and applied within its observed range. Using this reference, PI control was

executed while the achieved forces were assessed by the load cells at D, M, and E; the measurements track the targets under the stated operating conditions.




Figure 10. Empirical relationship between R_A and R_M in cases of two target surfaces.




Figure 11. Control of the SEA arms and L .

The force-control loops are designed for quasi-static regulation rather than broadband interaction. In the proposed mechanism, the effective contact stiffness is high but the arm deflections are small, which therefore generates large reaction forces. To avoid exciting structural and belt dynamics, proportional and integral gains were intentionally kept small, yielding a low closed-loop bandwidth relative to the dominant mechanical time constants and to the digital sampling rate. In addition, the series elastic elements reduce the reflected output impedance of the arms and attenuate high-frequency dynamics at the belt-workpiece interface. Under these conditions, the controller operates in a conservative regime. The commanded corrections evolve slowly compared with the fastest plant modes, and the SEA compliance filters contact transients. Empirical step-response tuning produced fast but non-oscillatory tracking, and no limit cycles or chatter were observed during the sanding trials reported in Section 4. The combination of high plant stiffness, small PI gains, SEA-induced compliance, and a control bandwidth chosen well below structural and belt resonances provides practical robustness for the operating conditions evaluated in this study.

4. Experiments

This study prioritizes mechanism characterization and feasibility of large-area convex sanding rather than full process optimization. All trials used roll paper of nominal width

50 mm with abrasive grit #100, and the vertical feed was held at 50 mm/min. The belt was driven by a DC motor at a commanded drive wheel speed of 300 rpm. Force regulation was validated with independent load cells at points D, M, and E under the reported conditions.

4.1. Lab Test-Bench

During the abrasion process, including sanding, the surface roughness of the target surface converges to a certain value under a controlled feed rate and force [20]. This, along with the evenly distributed normal force, allowed the vertical movement of the robot to evenly cover the entire convex surface. Figure 12 shows the target surface dimensions, which were selected to meet the standard outlined in the European Union's ADR Agreement for the International Carriage of Dangerous Goods by Road, with a long axis diameter of at least 1.5 times the short axis diameter [21].




Figure 12. Steel convex target surfaces.

Vertical motion was also incorporated into the large-area sanding process. The target height was set as twice the belt width to achieve uniform sanding across the entire target surface. A vertical speed of 50 mm/min was maintained constant throughout the sanding process with a sanding belt width of 50 mm, sanding the entire area for 1 min, mirroring the conditions of the vertically stationary sanding method. Uniform sanding is achieved when the target surface is evenly sanded vertically.

4.2. Results

Unlike a typical contact situation, where the two forces exerted on the target surface at the contact points with the SEA arm wheels D and E are the same ($R_D = R_E$) and the target surface is symmetric, R_D becomes smaller than R_E once the drive motor is initiated. This is because of the rotating belt and its direction, which makes R_D larger than R_E if the belt direction is opposite. Therefore, it is important to control the servo angles to maintain both R_D and R_E separately and actively. PI control was used to maintain the reaction forces measured on R_D , R_E , and R_M at a reference force of $R_{ref} = 15$ N, while $R_{A,ref}$ was selected from the empirical graph such that $R_M = R_{D,ref}$.

In the first case, without belt rotation, the force tracking ability between the SEA arm with and without PI control was measured. K_P was set to 0.3 for both the SEA arm and the re-tensioner, while K_I was set to 0.1 and 0.05 for the SEA arm and re-tensioner, respectively. Figure 13 shows the forces measured at contact point E with and without PI control. As shown in the figure, a settling time of 0.096 s is observed for the PI-applied SEA arm, following a 4% error range around the reference force of 15 N. However, owing to the overlapped belt joint, vigorous vibration was inevitable even with a drive wheel of 300 rpm. Accordingly, a simple moving average (SMA) filter was used for evaluation, as in Duong et al.'s work [22] for a vibration sensor application using (10).

$$\text{SMA} = \frac{A_1 + A_2 + \dots + A_n}{n} \quad (10)$$

where n is the window length and A_i is the i -th sample of the measured force signal. Lotysh et al. [23] and Zakamulin [24] also expounded on the SMA and window or unit size effect.




Figure 13. Forces under no-belt-rotation with and without PI control. Red area marks the target control interval.

Figure 14 shows the force measured on contact point D with PI control and the SMA filter with 4% of the total number of samples—20 in this case—during belt operation. The belt joint generates equally spaced spikes in the force trace. To suppress this periodic component without distorting the low-frequency envelope, we use the simple moving average defined in Equation (10) with a window of $n = 20$ (4% of samples). All quantitative metrics (peak force and settling time) are computed from the unfiltered signals, the SMA curve is for visualization only and therefore does not bias those estimates. The robot approached the target surface with the belt rotating and then departed from it after 4%. The time gaps between the sudden jumps in the graph are constant, indicating that this is owing to the overlapping part of the belt. The gray part represents the data measured on the target surface with the load cells, while the black line represents the same data filtered via SMA. The blue area represents the force exerted by the SEA arm. In this study, PI control is used to validate the SEA arm and re-tensioner architecture under the reported operating conditions; detailed controller optimization and stability-margin analysis are left for future work.




Figure 14. Belt-operation force with SMA. Red area marks the target control interval.

Figure 15 shows the reaction force measured on points D, M, and E on the target surface with a curvature of 0.47. Figure 16 shows the same but for a target surface with a curvature of 1.37 during the sanding operation, with the same SMA window size percentage as before. These figures show that the force-tracking ability of R_D is superior to that of R_E . This is because of the belt rotation direction. For the SEA arm at point D, the belt moves in the direction that extends the spring while compressing the spring on the SEA arm at point E. However, because extension springs are used in this mechanism, the belt cannot be further compressed, making the force-tracking ability of the SEA arm at point E slightly inferior to that at point D.

Table 3 consolidates the operating and evaluation conditions used for the force-tracking results reported in Figures 15 and 16.

Table 3. Parameters used for force-tracking experiments.

Parameter	Value / Units
Abrasive	Grit #100 (roll paper)
Belt width	50 mm
Vertical feed	50 mm/min
Drive wheel speed (command)	300 rpm
Force reference	15 N
Curvature values	0.47; 1.37
Repetitions	$n = 5$




Figure 15. Forces at D, M, and E on the target with curvature 0.47 during the belt sanding operation. Thin lines indicate raw data and thick lines indicate SMA-filtered data.




Figure 16. Forces at D, M, and E on the target with curvature 1.37 during the belt sanding operation. Thin lines indicate raw data and thick lines indicate SMA-filtered data.

4.3. Discussion

To verify the effectiveness of the proposed mechanism, the gaps in the maximum and minimum normal forces exerted on the plate during operation, as well as the finished quality evaluated with the vision system, were compared with those of a passive arm belt sanding robot in Figure 17, which was used in a previous study. That robot has fixed idler and passive arm positions and cannot calculate any specific force on the target surface or itself. This imaging pipeline, including its latency considerations, is consistent with real-time inspection modules [25].




Figure 17. Previous passive robotic arm without a re-tensioner and SEA system.

Figure 18 illustrates the vision inspection concept and process. The plate was first colored with a white primer and then black to show the sanded area edge, as the sanded surface reflects light, making it difficult to distinguish between the untouched and sanded

surfaces. The finished surface was moved into a black jig to take the pictures, which were processed as described in a previous study, but in opposite colors. White represents the sanded area and black represents the opposite.

Figure 19 shows the vertical sanding process results with SEA and passive arms on each target surface. For the target surface with a curvature of 0.47, a 31.85% improvement in the sanded surface area was achieved, from 62.2% to 94.05%. For a target surface with a curvature of 1.37, an 8.49% improvement was achieved, from 81.21% to 89.7%.




Figure 18. Vision inspection concept: (a) test-bench for vision inspection; (b) vertical sanding process result; (c) image resized and processed into 320×640 pixels.


	ρ [m ⁻¹]	Arm	Result	Sanded area [%]	Distribution rate [%]
(a)	0.47	Passive		62.15	133
(b)		SEA		94.05	20
(c)	1.37	Passive		81.21	30
(d)		SEA		89.70	13.3

Figure 19. Vertical sanding process results: (a) Curvature of 0.47 with passive arms; (b) curvature of 0.47 with SEA arms; (c) curvature of 1.37 with passive arms; and (d) curvature of 1.37 with SEA arms.

The distribution rate, defined as the gap between the maximum and minimum normal forces exerted on the target plate throughout the contact area divided by the maximum force, was no less than 20% when the SEA arms were used. However, without a re-tensioner on the passive arm robot, even though having the same normal force on points D and E as 15 N was possible, the force at point M was unpredictable, ranging from 10 to -5 N, because of the surface deformation caused by the high normal forces on D and E. This resulted in a distribution rate improvement of 113% and 16.7% for the 0.47- and 1.37-curvature target surfaces, respectively.

Under the operating conditions (belt width 50 mm and vertical feed 50 mm/min), the area rate for a single vertical pass is $0.05 \text{ m} \times 0.05 \text{ m/min} = 0.0025 \text{ m}^2/\text{min}$, i.e., $0.15 \text{ m}^2/\text{h}$. Force uniformity across the contact follows the distribution rate results already reported for the two convex curvatures (0.47 and 1.37) under identical settings. A principal limitation observed in the experiments is the compression asymmetry of the extension spring at wheel E, which slightly degrades tracking when the belt drives the spring into compression. A practical mitigation within the present mechanism is to select an appropriate spring preload and to implement asymmetric force loop gains between D and E. Although the R_A – R_M fit is validated by closed-loop performance within $R_A = 15$ – 130 N,

statistical uncertainty will be quantified in future work after increasing the number of repeated sweeps.

5. Conclusions

In this study, a novel robotic hardware capable of sanding large convex surfaces while in motion was proposed. To achieve this, SEA sanding arm and re-tensioner were designed. A specially designed rotary SEA was introduced to enhance robustness against noise and effectively absorb shocks during the sanding process. Additionally, using PI control, the robot simultaneously sanded a width of 1 m while in motion. To prove its performance, the estimated normal force value was compared with the measured value at the two main contact points on the target surface as well as at the midpoint of the contact area. Furthermore, a visual method was used to evaluate the evenness of the normal force applied to the target surface. It was experimentally confirmed that the proposed mechanism is effective in tracking forces and maintaining a well-dispersed force such that the sanded area is evenly distributed throughout the contact area compared to the previous passive-arm sanding robot, which was optimized using the Taguchi method.

Author Contributions: Conceptualization, T.K. and T.S.; methodology, H.J. and C.M.; validation, H.J., and C.M.; formal analysis, H.J. and T.K.; investigation, H.J.; writing—original draft preparation, H.J.; writing—review and editing, T.K. and T.S.; supervision, T.S.; funding acquisition, T.K. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (RS-2025-00521072).

Data Availability Statement: Dataset available on request from the authors.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Zhang, X.; Cabaravdic, M.; Kneupner, K.; Kuhlenkoetter, B. Real-Time Simulation of Robot Controlled Belt Grinding Processes of Sculptured Surfaces. *Int. J. Adv. Robot. Syst.* **2004**, *1*, 12. [CrossRef]
2. Song, Y.; Yang, H.; Lv, H. Intelligent Control for a Robot Belt Grinding System. *IEEE Trans. Control Syst. Technol.* **2013**, *21*, 716–724.
3. Brecher, C.; Tuecks, R.; Zunke, R.; Wenzel, C. Development of a Force Controlled Orbital Polishing Head for Free Form Surface Finishing. *Prod. Eng.* **2010**, *4*, 269–277. [CrossRef]
4. Ge, J.; Deng, Z.; Li, Z.; Li, W.; Lv, L.; Liu, T. Robot Welding Seam Online Grinding System Based on Laser Vision Guidance. *Int. J. Adv. Manuf. Technol.* **2021**, *116*, 1737–1749. [CrossRef]
5. Nagata, F.; Kusumoto, Y.; Fujimoto, Y.; Watanabe, K. Robotic Sanding System for New Designed Furniture with Free-Formed Surface. *Robot. Comput.-Integr. Manuf.* **2007**, *23*, 371–379. [CrossRef]
6. Iglesias, I.; Sebastián, M.A.; Ares, J.E. Overview of the State of Robotic Machining: Current Situation and Future Potential. *Procedia Eng.* **2015**, *132*, 911–917. [CrossRef]
7. Jinno, M.; Ozaki, F.; Yoshimi, T.; Tatsuno, K.; Takahashi, M.; Kanda, M.; Tamada, Y.; Nagataki, S. Development of a Force Controlled Robot for Grinding, Chamfering and Polishing. In Proceedings of 1995 IEEE International Conference on Robotics and Automation, Nagoya, Japan, 21–27 May 1995 Volume 2, pp. 1455–1460.
8. Xu, X.; Chen, W.; Zhu, D.; Yan, S.; Ding, H. Hybrid Active/Passive Force Control Strategy for Grinding Marks Suppression and Profile Accuracy Enhancement in Robotic Belt Grinding of Turbine Blade. *Robot. Comput.-Integr. Manuf.* **2021**, *67*, 102047. [CrossRef]
9. Chen, H.; Yang, Z.; Ding, H. Robotic Compliant Grinding of Curved Parts Based on a Designed Active Force-Controlled End-Effector with an Optimized Series Elastic Component. *Robot. Comput.-Integr. Manuf.* **2024**, *86*, 102646. [CrossRef]
10. Li, D.; Yang, J.; Zhao, H.; Ding, H. Contact Force Plan and Control of Robotic Grinding Towards Accurate Material Removal on Complex Surfaces. *Int. J. Mech. Sci.* **2022**, *227*, 107449. [CrossRef]
11. Min, K.; Ni, F.; Liu, H. Robotic Abrasive Belt Grinding of Complex Curved Blades Based on a Novel Force-Control Architecture Integrating Smooth Trajectories. *J. Manuf. Process.* **2023**, *107*, 447–458. [CrossRef]
12. Wen, Y.; Jaeger, D.J.; Pagilla, P.R. Uniform Coverage Tool Path Generation for Robotic Surface Finishing of Curved Surfaces. *IEEE Robot. Autom. Lett.* **2022**, *7*, 4931–4938. [CrossRef]

13. Kim, H.; Jin, H.; Lee, W.; Chae, S.; Kim, T.; Seo, T. Experimental Design of 4-Point Supported Belt Robot for Sanding Large Convex Surfaces. *Sci. Rep.* **2024**, *14*, 6861. [CrossRef]
14. Lee, C.; Kwak, S.; Kwak, J.; Oh, S. Generalization of Series Elastic Actuator Configurations and Dynamic Behavior Comparison. *Actuators* **2017**, *6*, 26. [CrossRef]
15. Wang, Y.-H.; Lo, Y.-C.; Lin, P.-C. A Normal Force Estimation Model for a Robotic Belt-Grinding System. In Proceedings of the 2020 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM), Boston, MA, USA, 6–9 July 2020; pp. 1922–1928.
16. Siciliano, B.; Sciavicco, L.; Villani, L.; Oriolo, G. *Robotics: Modelling, Planning and Control*; Springer: New York, NY, USA, 2009.
17. Hogan, N. Impedance Control: An Approach to Manipulation. *ASME J. Dyn. Syst. Meas. Control* **1985**, *107*, 1–7. [CrossRef]
18. Åström, K.J.; Hägglund, T. *PID Controllers: Theory, Design, and Tuning*; ISA: Research Triangle Park, NC, USA, 1995.
19. Seborg, D.E.; Edgar, T.F.; Mellichamp, D.A. *Process Dynamics and Control*; John Wiley & Sons: Hoboken, NJ, USA, 2003.
20. Horvath, M.; Kundrak, J.; Mamalis, A.G.; Gyani, K. On the Precision Grinding of Advanced Ceramics. *Int. J. Adv. Manuf. Technol.* **2002**, *20*, 255–258. [CrossRef]
21. Bergesen, H.O.; Parmann, G.; Thommessen, Ø. B. European Agreement Concerning the International Carriage of Dangerous Goods by Road (ADR). In *Yearbook of International Cooperation on Environment and Development 1998–99*; Routledge: London, UK, 2018; pp. 94–95.
22. Duong, T.T.; Nguyen, B.H.Q.; Maneetien, N. Implementation of Moving Average Filter on STM32F4 for Vibration Sensor Application. In Proceedings of the 2018 4th International Conference on Green Technology and Sustainable Development (GTSD), Ho Chi Minh City, Vietnam, 23–24 November 2018; pp. 627–631.
23. Lotysh, V.; Gumeniuk, L.; Humeniuk, P. Comparison of the Effectiveness of Time Series Analysis Methods: SMA, WMA, EMA, EWMA, and Kalman Filter for Data Analysis. *Inform. Autom. Pomiar. Gospod. Ochr. Środowiska* **2023**, *13*, 71–74. [CrossRef]
24. Zakamulin, V. Moving Averages for Market Timing. *SSRN* **2016**, 1–19. [CrossRef]
25. Hoshino, Y.; Shimasaki, M.; Rathnayake, N.; Dang, T.L. Performance verification and latency time evaluation of hardware image processing module for appearance inspection systems using FPGA. *J. Real-Time Image Process.* **2024**, *21*, 20. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Article

Design and Experimental Validation of an Adaptive Robust Control Algorithm for a PAM-Driven Biomimetic Leg Joint System

Feifei Qin ¹, Zexuan Liu ¹, Yuanjie Xian ², Binrui Wang ¹, Qiaoye Zhang ³ and Ye-Hwa Chen ^{4,*}

¹ College of Mechanical and Electrical Engineering, China Jiliang University, Hangzhou 310018, China; feifei.qin@cjlj.edu.cn (F.Q.); p23010854073@cjlj.edu.cn (Z.L.); wangbinrui@cjlj.edu.cn (B.W.)

² School of Mechanical Engineering, Hefei University of Technology, Hefei 230009, China; xianyuanjie@sztu.edu.cn

³ Department of Clinical Medical Engineering, Zhejiang Provincial People's Hospital (Affiliated People's Hospital), Hangzhou Medical College, Hangzhou 310058, China; zqy1983@zju.edu.cn

⁴ The George W. Woodruff School of Mechanical Engineering, Georgia Institute of Technology, Atlanta, GA 30318, USA

* Correspondence: yehwa.chen@me.gatech.edu

Abstract

Biomimetic quadruped robots, inspired by the musculoskeletal systems of animals, employ pneumatic artificial muscles (PAMs) as compliant actuators to achieve flexible, efficient, and adaptive locomotion. This study focuses on a pneumatic artificial muscle (PAM)-driven biomimetic leg joints system. First, its kinematic and dynamic models are established. Next, to address the challenges posed by the strong nonlinearities and complex time-varying uncertainties inherent in PAMs, an adaptive robust control algorithm is proposed by employing the Udwadia controller. Rigorous theoretical analysis of the adaptive robust control algorithm is verified via the Lyapunov stability method. Finally, numerical simulations and hardware experiments are conducted on the PAM-driven biomimetic leg joints system under desired trajectories, where the adaptive robust control algorithm is systematically compared with three conventional control algorithm to evaluate its control performance. The experimental results show that the proposed controller achieves a maximum tracking error of within 0.05 rad for the hip joint and within 0.1 rad, highlighting its strong potential for practical deployment in real-world environments.

Keywords: biomimetic robots; adaptive control algorithm; uncertainties; Udwadia controller

1. Introduction

In recent years, biomimetic robots have emerged as a promising frontier in legged locomotion, offering new possibilities for achieving agile, robust, and energy-efficient movement in complex environments [1]. Unlike conventional rigid actuators, pneumatic artificial muscles (PAMs) offer a high power-to-weight ratio, inherent compliance, and are safe and lightweight, which is particularly important for energy-efficient movement. As illustrated in Figure 1, the canine-inspired knee joint does not follow a fixed-axis rotation but exhibits a variable instantaneous center of rotation (ICR) [2], which is realized via a four-bar linkage mechanism to achieve variable-center kinematics and enhanced mechanical flexibility. To further mimic the softness and adaptability of biological muscles, PAMs are employed to actuate both the hip and knee joints [3], enabling compliant and natural movement patterns. However, PAMs exhibit pronounced nonlinear and hysteretic behaviors [4,5]. For pneumatic artificial muscles, the hysteresis is mainly manifested as lag in the contraction force length

relationship and lag in the pressure length relationship. In addition, the variable ICR structure introduces strong kinematic nonlinearities, which are specifically reflected in the nonlinear mapping between the joint rotation angle and the PAM end displacement. These nonlinear effects, together with modeling uncertainties, external disturbances, and payload variations, render the system dynamics strongly nonlinear and time-varying. However, the Udwadia controller computes only the nominal joint torques of the model and does not account for unknown model uncertainties, thereby posing substantial challenges for accurate modeling and high-precision motion control in PAM-driven biomimetic leg-joint systems.




Figure 1. The key control challenges of the PAM-driven biomimetic leg joints system.

To overcome these challenges, advanced control systems such as sliding mode control (SMC) [6], adaptive control [7,8], and robust control [9] can be employed. These methods have been extensively developed to address the nonlinearities, uncertainties, and time-varying disturbances in mechanical systems, effectively suppressing their adverse effects while improving tracking accuracy and system robustness. Among the various promising techniques explored in related studies, it is noteworthy that Udwadia adopted a new perspective on desired trajectories [10] and proposed a novel Udwadia controller [11]. The controller inputs required to enforce the desired trajectories are obtained in a fully explicit closed form, and the associated computations reduce to simple matrix multiplications and additions, thereby avoiding heavy analytical derivations and numerical optimization. However, the Udwadia controller only computes the nominal joint torques of the model and does not account for unknown model uncertainties. In practical applications, such uncertainties are inevitable and can significantly degrade system performance if not properly handled. To address this issue, Chen et al. employed the Udwadia controller and further developed an adaptive robust control scheme in [12], contributing to the theoretical advancement of this methodology. Since then, numerous studies have extended the Udwadia controller to uncertain nonlinear systems. For instance, Chen et al. incorporated the Udwadia controller into a model-based robust controller to address modular joint dynamics in dual-arm robot [13], while Huang et al. proposed an adaptive robust strategy for uncertain inverted pendulum systems within the same framework [14].

In addition, robust Udwadia controller have been reported for systems operating under uncertainties [15,16]. Beyond electromechanical systems, Chen et al also proposed a three-level interaction force control architecture for a hydraulic underactuated stance-leg exoskeleton, employing a MIMO adaptive robust controller to cope with high-order nonlinearities and uncertainties [17]. Despite these advances, the practical control limitations of existing Udwadia controllers remain insufficiently resolved for PAM-driven systems. Most prior studies primarily focus on rigid actuated platforms and often rely on assumptions such as reasonably accurate nominal rigid body dynamics, bounded varying uncertainties, or the availability of reliable state derivatives. However, PAM actuation introduces distinctive challenges strong rate-dependent hysteresis, pressure dynamics and delay which can lead to significant model plant mismatch and degrade tracking performance in real hardware. Therefore, to address these unresolved practical issues, this paper develops a controller that preserves the structural advantages of the Udwadia controller while explicitly accommodating PAM-specific nonlinearities and implementation constraints.

Several studies have also focused on the Udwadia controller. For instance, Ma et al. proposes a flexible adaptive nonsingular recursive terminal sliding mode control that combines uncertainty analysis and nonlinear robust control design [18], and Yu et al. A novel Udwadia approach for redundant parallel manipulator dynamics [19]. In parallel, Zheng et al. developed a barrier function-based adaptive SMC method to provide robustness for nonholonomic wheeled mobile robots [20]. On the theoretical side, Mylapilli and Udwadia examined hyperelastic beam control [21], while Udwadia continued to refine the mathematical underpinnings of constrained Hamiltonian dynamics [22], experimental investigations have also been conducted. For instance, Dong et al proposed an analytical modeling method in harmonic drive with tunable parameters that can be used to compensate for the uncertainty [23]. Zhao et al. proposed a cooperative control for the unmanned ground vehicle swarm system [24], and Yin et al. studied possibility-based robust controllers for structural-acoustic systems [25]. Moreover, the application scope has expanded to the sky-hook model [26], hydraulic exoskeleton under dynamic model uncertainty [27], and even vehicle active anti-roll system [28]. Collectively, these studies highlight the versatility of the Udwadia controller and its strong potential for adaptive, robust, and analytically tractable control of complex nonlinear systems.

Building upon recent advances in Udwadia controller and its successful applications in various nonlinear systems [29,30], this paper proposes an adaptive robust control algorithm under a PAM-driven biomimetic leg joints system. The control algorithm is built upon the Udwadia controller, where an adaptive term and a robust term is introduced. Specifically, the nominal part of the system dynamics is handled by the Udwadia controller, which generates the joint torques, a robust term is introduced to suppress deviations caused by initial errors and external disturbances, and an adaptive term with online parameter adjustment is employed to compensate for dynamic modeling errors and time-varying external disturbances. Moreover, the adaptive law is derived in a fully explicit closed form without requiring a priori upper bounds of the uncertainties, thereby significantly enhancing the engineering applicability of the adaptive robust control algorithm under complex operating conditions.

The main contributions of this paper are as follows. Firstly, a kinematic and dynamic model of the PAM-driven biomimetic leg joints system is established, which provides a theoretical foundation for high-accuracy joint motion control. Secondly, in view of the inherent modeling uncertainties and unmeasurable external disturbances in the PAM-driven biomimetic leg joints system, an adaptive robust control algorithm is proposed. The control algorithm consists of Udwadia controller, a robust term for initial deviation suppression, and an adaptive term designed to handle unknown time-varying uncertainties.

An adaptive law is constructed to ensure online adjustment of unknown parameters without requiring prior knowledge of their bounds. Through the stability analysis, the adaptive robust control algorithm is uniform boundedness (UB) and uniform ultimate boundedness (UUB). Finally, to validate the feasibility and effectiveness of the adaptive robust control algorithm, numerical simulations and hardware experiments are conducted based on a typical sinusoidal desired trajectory. The adaptive robust control algorithm is comprehensively compared with three mainstream control algorithms, including SMC, genetic algorithm-tuned proportional-integral-derivative control (GA-PID), and the torque compensation control. In the simulation data analysis, we used seven metrics to evaluate the adaptive robust control algorithm. The results demonstrate that the adaptive robust control algorithm achieves accuracy, robustness and practical applicability.

2. Kinematics and Dynamics Analysis

2.1. The Structure of PAM-Driven Biomimetic Leg Joints System

The PAM-driven biomimetic leg joints system design of the robotic joints considers both actuation and structural requirements. Four proportional valves are installed to control the airflow supplied to the muscles. Two angle encoders are mounted to measure the joint motion. Proportional valves are connected to regulate the supply pressure, while a switching power supply is located to provide electrical power for the system. The PAM serves as the system's primary actuator. Two tension sensors are mounted to measure the muscle force output, and the Beckhoff controller is used as the central control unit to manage system operation as shown in Figure 2.

Structurally, both the hip and knee joints are configured with a single rotational degree of freedom to satisfy the kinematic demands of bipedal locomotion. The hip joint features two symmetrically arranged PAMs of equal initial length with a fixed moment arm of $r_1 = 0.05$ m. In contrast, the knee joint adopts a crossed four-bar linkage mechanism, which provides a variable instantaneous center of rotation and thereby enhances its kinematic adaptability during gait.




Figure 2. PAM-driven biomimetic leg joints system.

2.2. Kinematic Model

The PAM-driven biomimetic leg joints system consists of a mechanical torso, thigh, and calf. The kinematic model of the PAM-driven leg joints system is illustrated in Figure 3a, where label 1 denotes the hip joint and label 2 denotes the knee joint. The hip joint rotation

center is defined as the origin O of a cartesian coordinate system, and point O_2 approximates the ICR of the knee joint.

The joint angles θ_1 and θ_2 represent the hip and knee joint angles, respectively, with the clockwise direction defined as positive. The corresponding joint torques are denoted as τ_1 for the hip and τ_2 for the knee. Each actuator M_i denotes a PAM. The subscript $i = a, b, c, d$ indexes the four PAMs used in the antagonistic configuration. x_i and F_i represent the muscle's end displacement and tensile force, respectively.




Figure 3. (a) Kinematic model of PAM-driven biomimetic leg joints system. (b) Dynamic model of PAM-driven biomimetic leg joints system.

The position $\theta_{h1} = 1.22$ rad corresponds to the state where the PAM M_b is fully contracted and the PAM M_a is fully relaxed. Conversely, the position $\theta_{h2} = -0.35$ rad corresponds to the state where the PAM M_b is fully relaxed and the PAM M_a is fully contracted. Therefore, the end displacements x_a and x_b of the PAMs M_a and M_b can be expressed in terms of the hip joint angle.

$$x_a = r_1(\theta_{h1} + \theta_1), \quad (1)$$

$$x_b = r_1(\theta_{h2} - \theta_1). \quad (2)$$

Assumption 1. The four-bar knee mechanism is used to derive the kinematic mapping from θ_2 to the PAM end displacements, and the resulting mapping is assumed to be valid within the designed knee motion range.

The four-bar linkage of the knee joint enables a rotational range of up to 1.75 rad. Based on the geometric parameters of the knee-joint linkage and the predefined attachment locations, a kinematic mapping is established to relate the joint rotation to the end displacements of the two PAM actuation paths. Specifically, x_c and x_d can be expressed as functions of θ_2 as follows:

$$x_c = a_2\theta_2^2 + a_1\theta_2 + a_0, \quad (3)$$

$$x_d = b_2\theta_2^2 + b_1\theta_2 + b_0. \quad (4)$$

By substituting the knee joint mechanism parameters into the system kinematics model, the results can be obtained separately as $a_2 = 0.0006466$, $a_1 = 0.1821$, $a_0 = -0.309$, $b_2 = 0.001009$, $b_1 = -0.4291$, $b_0 = 43.15$. It is worth noting that this polynomial mapping is derived from an analytical geometry-based solution of the four-bar linkage kinematics, and is introduced to provide an explicit closed-form relationship between the joint rotation and the PAM end displacements. According to Equations (3) and (4), the force arm of PAMs tension on the ICR is obtained as follows:

$$l_c = c_3\theta_2^3 + c_2\theta_2^2 + c_1\theta_2 + c_0, \quad (5)$$

$$l_d = d_3\theta_2^3 + d_2\theta_2^2 + d_1\theta_2 + d_0, \quad (6)$$

where $c_3 = -1.347$, $c_2 = 1.806$, $c_1 = 2.509$, $c_0 = 35.34$, $d_3 = 0.9856$, $d_2 = -6.325$, $d_1 = 3.844$, $d_0 = 39.41$, l_c and l_d are the knee joint lever arms by the tensile force F_c and F_d .

2.3. Dynamic Model

The PAM-driven biomimetic leg joints system is modeled as a simplified two-degree-of-freedom planar dual-link system, as illustrated in Figure 3b. The thigh and calf segments are represented as rigid links with concentrated masses.

Let m_1 and m_2 denote the masses of the thigh and calf, respectively. The lengths of the corresponding fixed segments are defined as L_1 and L_2 . The distances from the center of mass of each link to its proximal joint are given by C_1 and C_2 , respectively. The moments of inertia of the thigh and calf about their respective centers of mass are denoted by I_1 and I_2 .

Assumption 2. *The system is modeled as a planar two-DOF rigid-link mechanism, while out-of-plane motions, structural compliance, and joint backlash are neglected.*

The system's equations of motion can be obtained by using the Lagrangian dynamic modeling approach.

$$\tau_j = \frac{d}{dt} \left(\frac{\partial L}{\partial \dot{\theta}_j} \right) - \frac{\partial L}{\partial \theta_j}, \quad (7)$$

where $j = 1, 2$ represents the corresponding joint. L is the Lagrangian function defined as

$$L = T - U, \quad (8)$$

in which T and U represent the system's kinetic and potential energy, respectively, as follows:

$$T = \frac{1}{2}(\dot{\theta}_1 + \dot{\theta}_2)^2[m_2C_2^2 + I_2] + \frac{1}{2}\dot{\theta}_1[m_1C_1^2 + m_2L_2^2 + I_1] + \dot{\theta}_1(\dot{\theta}_1 + \dot{\theta}_2)m_2L_2I_2\cos\theta_2, \quad (9)$$

$$U = -gm_1C_1\cos\theta_1 - gm_2[L_2\cos\theta_1 + C_2\cos(\theta_1 + \theta_2)]. \quad (10)$$

By substituting (8)–(10) into (7), the dynamic equation of the system can be obtained:

$$\begin{bmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{bmatrix} \begin{bmatrix} \ddot{\theta}_1 \\ \ddot{\theta}_2 \end{bmatrix} + \begin{bmatrix} H_{11} \\ H_{21} \end{bmatrix} + \begin{bmatrix} G_{11} \\ G_{21} \end{bmatrix} = \begin{bmatrix} \tau_1 \\ \tau_2 \end{bmatrix}, \quad (11)$$

where the submatrices are shown in Table 1, the specific physical meanings of the submatrix symbols are shown in Table 2.

Table 1. Model matrix elements.

Submatrix	Component
M_{11}	$m_1 C_1^2 + m_2 (L_1^2 + C_2^2 + 2L_1 C_2 \cos \theta_2) + I_1 + I_2$
M_{12}	$m_2 (C_1^2 + L_1 C_2 \cos \theta_2) + I_2$
M_{21}	$m_2 (C_1^2 + L_1 C_2 \cos \theta_2) + I_2$
M_{22}	$m_2 C_2^2 + I_2$
H_{11}	$m_2 L_1 C_2 \sin \theta_2 \dot{\theta}_1 \dot{\theta}_2 - (\dot{\theta}_1 + \dot{\theta}_2) m_2 L_1 C_2 \sin \theta_2 \dot{\theta}_2$
H_{21}	$m_2 L_1 C_2 \sin \theta_2 \dot{\theta}_1^2$
G_{11}	$m_1 g C_1 \sin \theta_1 + m_2 g L_1 \sin \theta_1 + m_2 g C_2 \cos(\theta_1 + \theta_2)$
G_{12}	$m_2 g C_2 \sin(\theta_1 + \theta_2)$

Table 2. Table of variables.

Submatrix	Component
θ_1	The hip joint rotation angle
θ_2	The knee joint rotation angle
m_1	The weight of the hip joint
m_2	The weight of the knee joint
L_1	The length of the thigh
L_2	The length of the shank
C_1	Distance from thigh centroid to hip joint
C_2	Distance from calf centroid to knee joint
I_1	The moment of inertia of the thigh
I_2	The moment of inertia of the calf

The driving torque of the hip joint τ_1 is related to the antagonistic PAMs tensile force F_a and F_b . The PAMs tensile force F_c and F_d drive the knee joint to rotate through a four-bar linkage structure. τ_2 is the driving torque of the knee joint, l_c and l_d represent the approximate arm of tensile force. The driving torques of hip and knee joint are as follows:

$$\tau_1 = (F_a - F_b)r_1, \quad (12)$$

$$\tau_2 = F_c l_c - F_d l_d. \quad (13)$$

PAMs can only generate unidirectional tensile forces. The mapping relationship between the hip joint torque τ_1 and the PAM tensile forces F_a and F_b is designed as follows:

$$F_a = \begin{cases} \frac{\tau_0 + \tau_1}{r_1} (\tau_1 \geq 0), \\ \frac{\tau_0}{r_1} (\tau_1 < 0), \end{cases} \quad (14)$$

$$F_b = \begin{cases} \frac{\tau_0}{r_1} (\tau_1 \geq 0), \\ \frac{\tau_0 - \tau_1}{r_1} (\tau_1 < 0), \end{cases} \quad (15)$$

where τ_0 is the F_a , F_b , F_c and F_d applied initial torque. Enforcing $\tau_0 \geq 0$ ensures that the antagonistic PAM pair remains in tension at all times; otherwise, one muscle may slacken and buckle, which would compromise the accuracy of the mapping from the joint angle to the PAM end displacement. Similarly, the mapping relationship between the hip joint torque τ_2 and the PAM tensile forces F_c and F_d is designed as follows:

$$F_c = \begin{cases} \frac{\tau_0 + \tau_2}{l_c} (\tau_2 \geq 0), \\ \frac{\tau_0}{l_c} (\tau_2 < 0), \end{cases} \quad (16)$$

$$F_d = \begin{cases} \frac{\tau_0}{l_d} (\tau_2 \geq 0), \\ \frac{\tau_0 - \tau_2}{l_d} (\tau_2 < 0). \end{cases} \quad (17)$$

3. Controller Design

3.1. Constraint Equation

From a universal perspective, the control algorithm design in this section is developed based on a common dynamic equation of a PAM-driven biomimetic leg joints system. Consider a PAM-driven biomimetic leg joints system as follows:

$$\mathbf{M}(\mathbf{q}(t), t) \ddot{\mathbf{q}}(t) + \mathbf{H}(\dot{\mathbf{q}}(t), \mathbf{q}(t), t) + \mathbf{G}(\mathbf{q}(t), t) = \boldsymbol{\tau}(t), \quad (18)$$

where the generalized coordinate $\mathbf{q}$ is an $n \times 1$ vector and n represents the number of degrees of freedom. Equation (11) is the same case of (18) when $n = 2$, and $\mathbf{q} = [\theta_1, \theta_2]^T$, $\dot{\mathbf{q}} = [\dot{\theta}_1, \dot{\theta}_2]^T$, $\ddot{\mathbf{q}} = [\ddot{\theta}_1, \ddot{\theta}_2]^T$.

Let the system (18) follow a set of desired trajectories in a trajectory tracking problem. The tracking trajectory θ_i^d of the i -th joint can be formulated as servo constraints:

$$\boldsymbol{\Lambda}(t) \mathbf{q}(t) = \mathbf{c}_0(t), \quad (19)$$

where $\boldsymbol{\Lambda}$ is an $n \times n$ identity matrix, $\mathbf{c}_0(t) = [\theta_1^d(t), \theta_2^d(t), \dots, \theta_n^d(t)]^T$. By differentiating Equation (19) twice with respect to time, the first and second-order matrix forms of the constraint equation can be obtained as:

$$\boldsymbol{\Lambda}(t) \dot{\mathbf{q}}(t) = \mathbf{c}_1(t), \quad (20)$$

$$\boldsymbol{\Lambda}(t) \ddot{\mathbf{q}}(t) = \mathbf{c}_2(t), \quad (21)$$

where c_i ($i = 0, 1, 2$) denotes the i -th order constraint, $\mathbf{c}_1(t) = [\dot{\theta}_1^d(t), \dot{\theta}_2^d(t), \dots, \dot{\theta}_n^d(t)]^T$, $\mathbf{c}_2(t) = [\ddot{\theta}_1^d(t), \ddot{\theta}_2^d(t), \dots, \ddot{\theta}_n^d(t)]^T$.

The constraint-following in this paper adheres to second-order servo constraints. The second-order constraint (21) can fully express the constraint information with a given initial condition. Practically all the performance requirements of systems can be formulated as servo constraints. The explicit closed-form constraint can be obtained, which enables control algorithm design. Second-order constraints can help to account for uncertainties and disturbances in the system dynamics. The proposed control algorithm is integrated with the Udwadia controller for constraint-following control, leading to an advanced control structure capable of dealing with time-varying uncertainties simultaneously.

3.2. Adaptive Robust Controller with Uncertainties

In practical situations, uncertainty and interference are inevitable. To account for the uncertainty, the motion Equation (18) is re-described as follows:

$$\mathbf{M}(\mathbf{q}(t), \sigma(t), t) \ddot{\mathbf{q}}(t) + \mathbf{H}(\mathbf{q}(t), \dot{\mathbf{q}}(t), \sigma(t), t) + \mathbf{G}(\mathbf{q}(t), \sigma(t), t) = \boldsymbol{\tau}(t), \quad (22)$$

where $\sigma \in R^p$ is time-varying uncertain parameters. Based on whether they contain uncertainty or not, the model parameters can be resolved into two parts:

Assumption 3. PAM non-idealities are not explicitly modeled, but lumped into the uncertainty term, which is compensated online by the proposed adaptive robust controller.

$$\mathbf{M}(\mathbf{q}(t), \sigma(t), t) = \bar{\mathbf{M}}(\mathbf{q}(t), t) + \Delta \mathbf{M}(\mathbf{q}(t), \sigma(t), t), \quad (23)$$

$$\mathbf{H}(\mathbf{q}(t), \dot{\mathbf{q}}(t), \sigma(t), t) = \bar{\mathbf{H}}(\mathbf{q}, \dot{\mathbf{q}}(t), t) + \Delta \mathbf{H}(\mathbf{q}(t), \dot{\mathbf{q}}(t), \sigma(t), t), \quad (24)$$

$$\mathbf{G}(\mathbf{q}(t), \sigma(t), t) = \bar{\mathbf{G}}(\mathbf{q}(t), t) + \Delta \mathbf{G}(\mathbf{q}(t), \sigma(t), t), \quad (25)$$

where $\bar{\mathbf{M}}(\cdot)$, $\bar{\mathbf{H}}(\cdot)$, and $\bar{\mathbf{G}}(\cdot)$ are the nominal matrices, $\Delta\mathbf{M}(\cdot)$, $\Delta\mathbf{H}(\cdot)$, $\Delta\mathbf{G}(\cdot)$ are the related ambiguity matrices which represents the uncertainty. In an ideal and uncertainty-free system, the uncertainty matrix term is 0.

Design explicit constraint-following term and robust terms based on the Udwadia controller as below:

$$\mathbf{p}_1 = \bar{\mathbf{M}}^{\frac{1}{2}}(\Lambda\bar{\mathbf{M}}^{-\frac{1}{2}})^+(c_2 - \Lambda\bar{\mathbf{M}}^{-1}\bar{\mathbf{Q}} - k\Lambda\dot{\mathbf{q}}), \quad (26)$$

$$\mathbf{p}_2 = -k\bar{\mathbf{M}}\Lambda(\Lambda\Lambda^T)^{-1}\mathbf{P}^{-1}\boldsymbol{\beta}, \quad (27)$$

where $k \in \mathbb{R}^+$. An increase in the value of k results in a higher degree of robustness.

However, given the presence of uncertainty and disturbances, it is highly unlikely that all constraint conditions can be fulfilled. Thus, we propose a vector $\boldsymbol{\beta}$ to quantify the level of dissatisfaction with regard to ideal constraints.

$$\begin{aligned} \boldsymbol{\beta}(\mathbf{q}(t), \dot{\mathbf{q}}(t), t) = & (\Lambda(\mathbf{q}(t), \dot{\mathbf{q}}(t), t)\dot{\mathbf{q}}(t) - \mathbf{c}_1(\mathbf{q}(t), \dot{\mathbf{q}}(t), t)) \\ & + \lambda(\Lambda(\mathbf{q}(t), \dot{\mathbf{q}}(t), t)\dot{\mathbf{q}}(t) - \mathbf{c}_0(\mathbf{q}(t), \dot{\mathbf{q}}(t), t)), \end{aligned} \quad (28)$$

where $\boldsymbol{\beta}$ takes into account both the first-order and second-order constraints, which can adequately capture the influence of uncertainties and disturbances.

Assumption 4. Given $\mathbf{P} \in \mathbb{R}^{m \times n}$, $\mathbf{P} > 0$, let

$$\begin{aligned} \mathbf{W}(\mathbf{q}(t), \sigma(t), t) = & \mathbf{P}\Lambda(\mathbf{q}(t), t)\bar{\mathbf{M}}(\mathbf{q}(t), t)\mathbf{M}(\mathbf{q}(t), \sigma(t), t) \\ & \times \bar{\mathbf{M}}(\mathbf{q}(t), t)\Lambda^T(\mathbf{q}(t), t)(\Lambda(\mathbf{q}(t), t)\Lambda^T(\mathbf{q}(t), t))^{-1}\mathbf{P}^{-1}. \end{aligned} \quad (29)$$

A constant $\rho_E > -1$ can be found, for all $(\mathbf{q}, t) \in \mathbb{R}^n \times \mathbb{R}$,

$$\frac{1}{2} \min \lambda_m(\mathbf{W}(\mathbf{q}(t), \sigma(t), t)) \geq \rho_E. \quad (30)$$

Assumption 5. (1) There exists a positive constant vector $\boldsymbol{\alpha}$, and a function $\Pi(\cdot) \in \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}$, such that

$$(1 + \rho_E)^{-1} \max[\|\mathbf{P}\Lambda(\mathbf{M}^{-1} - \bar{\mathbf{M}}^1)(-\mathbf{H} - \mathbf{G} + \mathbf{p}_1 + \mathbf{p}_2)\|] \leq \Pi. \quad (31)$$

(2) Every $(\boldsymbol{\alpha}, \mathbf{q}(t), \dot{\mathbf{q}}(t), t)$, the function $\Pi(\cdot)$ is possible to perform a linear factorization of $\boldsymbol{\alpha}$: there is a function $\Pi(\cdot): \mathbb{R}^2 \times \mathbb{R}^2 \times \mathbb{R} \rightarrow \mathbb{R}_+^k$ such that:

$$\Pi(\boldsymbol{\alpha}, \mathbf{q}(t), \dot{\mathbf{q}}(t), t) = \boldsymbol{\alpha}^T \tilde{\Pi}(\mathbf{q}(t), \dot{\mathbf{q}}(t), t), \quad (32)$$

where $\boldsymbol{\alpha}$ is unknown. The selection of the known function Π is not unique, as long as Assumption 4

By incorporating the above assumption and system dynamics, the following control law is proposed to ensure trajectory tracking under uncertainties

$$\boldsymbol{\tau}(t) = \mathbf{p}_1(\mathbf{q}(t), \dot{\mathbf{q}}(t), t) + \mathbf{p}_2(\mathbf{q}(t), \dot{\mathbf{q}}(t), t) + \mathbf{p}_3(\tilde{\boldsymbol{\alpha}}, \mathbf{q}(t), \dot{\mathbf{q}}(t), t), \quad (33)$$

where

$$\begin{aligned} \mathbf{p}_3 = & -[\bar{\mathbf{M}}(\mathbf{q}(t), t)\Lambda^T(\mathbf{q}(t), t)(\Lambda(\mathbf{q}(t), t)\Lambda^T(\mathbf{q}(t), t))^{-1}\mathbf{P}^{-1}]\gamma(\tilde{\boldsymbol{\alpha}}, \mathbf{q}(t), \dot{\mathbf{q}}(t), t) \\ & + \mu(\tilde{\boldsymbol{\alpha}}, \mathbf{q}(t), \dot{\mathbf{q}}(t), t)\Pi(\tilde{\boldsymbol{\alpha}}, \mathbf{q}(t), \dot{\mathbf{q}}(t), t), \end{aligned} \quad (34)$$

$$\gamma(\tilde{\alpha}, q(t), \dot{q}(t), t) = \begin{cases} \frac{1}{\varepsilon}, & \text{if } \|\mu(\tilde{\alpha}, q(t), \dot{q}(t), t)\| \leq \varepsilon, \\ \frac{1}{\|\mu(\tilde{\alpha}, q(t), \dot{q}(t), t)\|}, & \text{if else} \end{cases} \quad (35)$$

$$\mu(\tilde{\alpha}, q(t), \dot{q}(t), t) = \beta \Pi(\alpha, q(t), \dot{q}(t), t). \quad (36)$$

The parameter $\tilde{\alpha}$ is governed by the following adaptive law

$$\dot{\tilde{\alpha}} = \frac{\Gamma}{1 + \|\tilde{\Pi}(q(t), \dot{q}(t), t)\|^2} \tilde{\Pi}(q(t), \dot{q}(t), t) \|\beta\| - k_2 \tilde{\alpha}. \quad (37)$$

The initial conditions are selected to be positive, which ensures that $\tilde{\alpha}$ are positive for all t . By introducing a normalization factor $\frac{\Gamma}{1 + \|\tilde{\Pi}(q(t), \dot{q}(t), t)\|^2}$, the adaptation step is automatically limited, preventing parameter blow-up and unfavorable transients. As a result, the control algorithm avoids aggressive torque commands near saturation, mitigating chatter and prolonging hardware lifetime. In addition, reduced high-frequency compensation lowers compressed-air usage and peak power, yielding smoother valve actuation and decreased pneumatic consumption.

3.3. Stability Analysis

Theorem 1. Let $\tilde{\delta} := [\beta^T(\tilde{\alpha} - \alpha)^T]^T \in \mathbb{R}^{m+k}$. Subject to Assumptions 4 and 5, consider the system (22). The control (33) renders the following performance:

(i) UB: For any $r > 0$, there is a $d(r) < \infty$ such that if $\|\tilde{\delta}(t_0)\| \leq r$, then $\|\tilde{\delta}(t)\| \leq d(r)$ for all $t \geq t_0$.

(ii) UUB: For any $r > 0$ with $\|\tilde{\delta}(t_0)\| \leq r$, there exists a $\bar{d} > 0$ such that $\|\tilde{\delta}\| \leq \bar{d}$ for any $\bar{d} \geq d$ as $t \geq t_0 + T(\bar{d}, r)$, where $T(\bar{d}, r) < \infty$.

Let $\Delta \mathbf{D}(\mathbf{q}, \sigma, t) := \mathbf{M}^{-1}(\mathbf{q}, \sigma, t) - \bar{\mathbf{M}}^{-1}(\mathbf{q}, t)$. Then,

$$\mathbf{E}(\mathbf{q}, \sigma, t) = \bar{\mathbf{M}}(\mathbf{q}, t)(\mathbf{M}^{-1}(\mathbf{q}, \sigma, t) - \bar{\mathbf{M}}^{-1}(\mathbf{q}, t)), \quad (38)$$

$$\Delta \mathbf{D}(\mathbf{q}, \sigma, t) = \mathbf{D}(\mathbf{q}, t) \mathbf{E}(\mathbf{q}, \sigma, t), \quad (39)$$

$$\mathbf{M}^{-1}(\mathbf{q}, \sigma, t) = \Delta \mathbf{D}(\mathbf{q}, \sigma, t) + \mathbf{D}(\mathbf{q}, t). \quad (40)$$

Proof. A legitimate Lyapunov function candidate is given:

$$V(\beta, \tilde{\alpha}) = \beta^T \mathbf{P} \beta + \Gamma^{-1}(1 + \rho_E)(\tilde{\alpha} - \alpha)^T(\tilde{\alpha} - \alpha). \quad (41)$$

For given controlled system with uncertainty and desired trajectory, the derivative of V is evaluated as follows:

$$\dot{V} = 2\beta^T \mathbf{P} \dot{\beta} + 2\Gamma^{-1}(1 + \rho_E)(\tilde{\alpha} - \alpha)^T \dot{\tilde{\alpha}}. \quad (42)$$

The first term of (41) can be evaluated as follows:

$$\begin{aligned} 2\beta^T \mathbf{P} \dot{\beta} &= 2\beta^T \mathbf{P} [\Lambda \ddot{q} - c_2 + \lambda(\Lambda \dot{q} - c_1)] \\ &= [\Lambda \mathbf{D}(-\bar{H} - \bar{G}) + \Lambda \mathbf{D} p_1 - c_2 + \lambda(\Lambda \dot{q} - c_1)] + \Lambda \mathbf{D} p_2 \\ &\quad + \Lambda(\Delta \mathbf{D} + \mathbf{D}) p_3 + [\Lambda \mathbf{D}(-\Delta H - \Delta G) + \Lambda \Delta \mathbf{D}(-H - G) \\ &\quad + \Lambda \Delta \mathbf{D}(p_1 + p_2)]. \end{aligned} \quad (43)$$

With the special case $\sigma = 0$, (26) can be written as

$$\Lambda \mathbf{D}(-\bar{H} - \bar{G}) + \Lambda \mathbf{D} p_1 - c_2 + \lambda(\Lambda \dot{q} - c_1) = 0, \quad (44)$$

From the above formula that (43) can be evaluated as

$$\begin{aligned} & 2\beta^T \mathbf{P}[\Lambda \mathbf{D}(-\Delta \mathbf{H} - \Delta \mathbf{G}) + \Lambda \Delta \mathbf{D}(-\mathbf{H} - \mathbf{G}) + \Lambda \Delta \mathbf{D}(\mathbf{p}_1 + \mathbf{p}_2)] \\ & \leq 2\|\beta\| \|\mathbf{P}[\Lambda \mathbf{D}(-\Delta \mathbf{H} - \Delta \mathbf{G}) + \Lambda \Delta \mathbf{D}(-\mathbf{H} - \mathbf{G}) + \Lambda \Delta \mathbf{D}(\mathbf{p}_1 + \mathbf{p}_2)]\| \\ & \leq 2(1 + \rho_E) \|\beta\| \Pi(\alpha, q, \dot{q}, t). \end{aligned} \quad (45)$$

By (27), we have

$$2\beta^T \mathbf{P} \Lambda \mathbf{D} \mathbf{p}_2 = -2k\beta^T \beta = -2k\|\beta\|^2. \quad (46)$$

By (33), (46), (41), we have

$$2\beta^T \mathbf{P} \Lambda (\mathbf{D} + \Delta \mathbf{D}) \mathbf{p}_3 = 2\beta^T \mathbf{P} \Lambda \mathbf{D} \mathbf{E} [-\tilde{\mathbf{M}} \Lambda^T (\Lambda \Lambda^T)^{-1} \mathbf{P}^{-1} \gamma \mu \Pi] - 2\gamma \|\beta\|^2 \|\mu\|. \quad (47)$$

Based on Assumption 4 and Rayleigh's principle, we can show that

$$\begin{aligned} & 2\beta^T \mathbf{P} \Lambda \mathbf{D} \mathbf{E} [-\tilde{\mathbf{M}} \Lambda^T (\Lambda \Lambda^T)^{-1} \mathbf{P}^{-1} \gamma \mu \Pi] \\ & = -2\gamma \mu^T [\tilde{\mathbf{M}} \Lambda^T (\Lambda \Lambda^T)^{-1} \mathbf{P}^{-1}] \mu \\ & = -2\gamma \mu^T \frac{1}{2} [\tilde{\mathbf{M}} \Lambda^T (\Lambda \Lambda^T)^{-1} \mathbf{P}^{-1} + \tilde{\mathbf{M}} \Lambda^T (\Lambda \Lambda^T)^{-1} \mathbf{P}^{-1}] \\ & = -2\gamma \mu^T \frac{1}{2} (\mathbf{W} + \mathbf{W}^T) \mu \\ & \leq -2\gamma \rho_E \|\mu\|. \end{aligned} \quad (48)$$

Using (47) and (48), we obtain

$$2\beta^T \mathbf{P} \Lambda (\mathbf{D} + \Delta \mathbf{D}) \mathbf{p}_3 \leq -2\gamma(1 + \rho_E) \|\mu\|. \quad (49)$$

(i) If $\|\mu\| > \hat{\varepsilon}$

$$-2\gamma(1 + \rho_E) \|\mu\| = -2\gamma(1 + \rho_E) \frac{1}{\mu} \|\mu\|^2 = -2\gamma(1 + \rho_E) \|\mu\|. \quad (50)$$

(ii) If $\|\mu\| \leq \hat{\varepsilon}$

$$-2\gamma(1 + \rho_E) \|\mu\| = -2\gamma(1 + \rho_E) \frac{1}{\hat{\varepsilon}} \|\mu\|^2. \quad (51)$$

These lead to the following: for $\|\mu\| > \hat{\varepsilon}$, we have

$$2\beta^T \mathbf{P} \dot{\beta} \leq -2k\|\beta\|^2 + (1 + \rho_E) \{-2\|\beta\| \Pi(\tilde{\alpha}, q, \dot{q}, t) + 2\|\beta\| \Pi(\alpha, q, \dot{q}, t)\}. \quad (52)$$

For $\|\mu\| \leq \hat{\varepsilon}$

$$\begin{aligned} 2\beta^T \mathbf{P} \dot{\beta} & \leq -2k\|\beta\|^2 + (1 + \rho_E) \frac{\hat{\varepsilon}}{2} \\ & + (1 + \rho_E) \{-2\|\beta\| \Pi(\tilde{\alpha}, q, \dot{q}, t) + 2\|\beta\| \Pi(\alpha, q, \dot{q}, t)\}. \end{aligned} \quad (53)$$

Based on Assumption 5, we obtain

$$2\|\beta\| \Pi(\alpha, q, \dot{q}, t) - 2\|\beta\| \Pi(\hat{\alpha}, q, \dot{q}, t) = 2\|\beta\| (a - \hat{a})^T \tilde{\Pi}(q, \dot{q}, t). \quad (54)$$

For all $\|\mu\|$, the following inequality holds:

$$2\beta^T \mathbf{P} \dot{\beta} \leq -2k\|\beta\|^2 + (1 + \rho_E) \frac{\hat{\varepsilon}}{2} + 2(1 + \rho_E) \|\beta\| (a - \hat{a})^T \tilde{\Pi}(q, \dot{q}, t). \quad (55)$$

By using adaptive law (37), we have

$$\begin{aligned} & 2\Gamma^{-1}(1 + \rho_E)(\tilde{\alpha} - \alpha)^T \dot{\tilde{\alpha}} \\ & \leq 2(\tilde{\alpha} - \alpha)^T \tilde{\Pi}(q, \dot{q}, t) \|\beta\| - 2\Gamma^{-1}k_2 \|\tilde{\alpha} - \alpha\|^2 + 2\Gamma^{-1}k_2 \|\tilde{\alpha} - \alpha\| \|\alpha\|. \end{aligned} \quad (56)$$

Combining (19), (20), and (42), we derive

$$\begin{aligned}\dot{V} &\leq -2\kappa\|\beta\|^2 + (1 + \rho_E)\frac{\hat{\varepsilon}}{2} - 2\Gamma^{-1}k_2(1 + \rho_E)\|\tilde{\alpha} - \alpha\|^2 \\ &\quad + 2\Gamma^{-1}k_2(1 + \rho_E)\|\tilde{\alpha} - \alpha\|\|\alpha\| \\ &\leq -k_1\|\tilde{\delta}\|^2 + k_2\|\tilde{\delta}\| + k_3,\end{aligned}\quad (57)$$

where

$$\Gamma^{-1} = \min\{2\kappa, 2\Gamma^{-1}k_2(1 + \rho_E)\}, \quad k_2 = 2\Gamma^{-1}k_2(1 + \rho_E)\|\alpha\|, \quad k_3 = \frac{(1 + \rho_E)\hat{\varepsilon}}{2}.$$

Upon invoking the standard arguments as in Khalil [31] and Chen [32], we conclude the UB with

$$d(r) = \begin{cases} \sqrt{\frac{\gamma_2}{\gamma_1}} R, & r \leq R, \\ \sqrt{\frac{\gamma_2}{\gamma_1}} r, & r > R, \end{cases} \quad (58)$$

where

$$R = \frac{1}{2k_1} \left(k_2 + \sqrt{k_2^2 + 4k_1k_3} \right), \quad (59)$$

and $\gamma_1 = \min\{\lambda_{\min}(P), \Gamma^{-1}(1 + \rho_E)\}$, $\gamma_2 = \max\{\lambda_{\max}(P), \Gamma^{-1}(1 + \rho_E)\}$.

Furthermore, UUB also follows with

$$d = \sqrt{\frac{\gamma_2}{\gamma_1}} R, \quad (60)$$

$$T(\bar{d}, r) = \begin{cases} 0, & r \leq \bar{d}\sqrt{\frac{\gamma_1}{\gamma_2}}, \\ \frac{\gamma_2 r^2 - (\gamma_1/\gamma_2)\bar{d}^2}{k_1\bar{d}^2(\gamma_1/\gamma_2)^{1/2} - k_2\bar{d}(\gamma_1/\gamma_2)^{1/2} - k_3}, & \text{otherwise.} \end{cases} \quad (61)$$

Since the Lyapunov derivative is negative outside a region of $\tilde{\delta}$, one concludes UB and UUB. $\square$

4. Simulation

This section evaluates the performance of the proposed control algorithm described in Section 3 through numerical simulation. The joint control torques are computed based on the control law defined in (33). The parameters used for the model, including physical and control-related constants, are listed in Table 3.

Table 3. Table with parameters.

Parameters Name	Value
m_1	2.97 kg
m_2	0.54 kg
I_1	0.1175 kg·m ²
I_2	0.00708 kg·m ²
L_1	0.4 m
L_2	0.355 m
C_1	0.124 m
C_2	0.146 m
g	9.8 m/s ²

We introduced small sinusoidal perturbations to the mass parameters in simulation to emulate typical time-varying parametric uncertainties in a simple, controlled, and repeatable manner. This sinusoidal model is used only as a representative example for robustness evaluation. $m_1(t) = \bar{m}_1(t) + \Delta m_1(t)$, $m_2(t) = \bar{m}_2(t) + \Delta m_2(t)$, $\Delta m_1(t) = 0.1 \sin(2\pi t)$, $\Delta m_2(t) = 0.1 \sin(2\pi t)$. The initial conditions are set as: $\theta_1(0) = 0.0872$ rad, $\theta_2(0) = 0.4361$ rad, $\dot{\theta}_1(0) = 0.0017$ rad/s, $\dot{\theta}_2(0) = 0.0017$ rad/s. The desired trajectories of joints are set as below:

$$\theta_1^d(t) = (0.4361 \sin(4/3\pi * t) + 0.0017), \quad (62)$$

$$\theta_2^d(t) = (0.6106 \sin(4/3\pi * t) + 10 \sin(8/3\pi * t - \pi/4) + 0.4361). \quad (63)$$

The terms in $\mathbf{M}$ are either linear constants of position or linear constants of velocity. Consequently, Assumption 5 can be satisfied by

$$\begin{aligned} \Pi(\alpha, q, \dot{q}, t) &= \alpha_1 \|\dot{q}\| + \alpha_2 \|q\| + \alpha_3 \\ &\leq \alpha (\|\dot{q}\| + \|q\| + 1) \\ &=: \alpha \Pi(q, \dot{q}, t), \end{aligned} \quad (64)$$

where $\alpha_{1,2,3} > 0$, $\alpha_{1,2,3}$ are representing the first, second, and third components of vector α , respectively. $\alpha = \max [\alpha_1, \alpha_2, \alpha_3]$.

The control algorithm parameters for simulation are chosen as follows: $k = 10$, $\Gamma = 50$, $\lambda = 10$, $\varepsilon = 0.1$, $k_2 = 0.001$, $P = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$. Set the system simulation time $t = 20$ s. In practice, Γ is set to achieve a moderate adaptation rate without amplifying noise, k_2 is chosen as the smallest value that still guarantees robustness against the lumped uncertainty while avoiding excessive control effort/chattering, ε is selected according to sensor noise/resolution as a small positive regularization constant for numerical stability.

Under the proposed control algorithm, Figure 4 presents the desired and actual trajectories of the hip joint, while Figure 5 shows the corresponding results for the knee joint. In both cases, the actual trajectories closely follow the desired trajectories, indicating satisfactory tracking performance and effective control response.




Figure 4. Trajectory tracking of the hip joints.

Figure 6 comprises four subplots: the top two illustrate the tracking errors E_1 and E_2 under SMC, while the bottom two present the same error metrics under adaptive robust control algorithm. Under SMC, both error curves exhibit large initial deviations at $t = 0$ s, but rapidly converge to near zero within approximately one second and remain stable thereafter, demonstrating SMC's strong capability in rejecting parameter uncertainties and external disturbances.




Figure 5. The knee joints trajectories tracking.

In contrast, adaptive robust control algorithm achieves near-zero tracking errors at the initial moment, indicating superior transient response and ideal tracking performance. Quantitatively, the maximum errors under the adaptive robust control algorithm are only 0.036613 for E_1 and 0.1234 for E_2 , both of which are notably smaller than those observed under SMC. These results confirm that the adaptive robust control algorithm provides reduced maximum and steady-state errors compared to SMC, highlighting its improved tracking precision and robustness.




Figure 6. The hip and knee joint errors under the proposed control algorithm and SMC.

The time-domain responses of the joint torque signals τ_1 and τ_2 over a 20 s interval are depicted in Figure 7. The torque τ_1 exhibits a significantly larger amplitude, peaking at approximately $1.5 \text{ N} \cdot \text{m}$, while τ_2 reaches only around $0.4\text{--}0.5 \text{ N} \cdot \text{m}$. This disparity suggests that the hip joint experiences a higher inertial load or resistance, thereby requiring greater torque to track the same desired trajectory.

Both torque signals exhibit the same oscillation period, indicating that they are governed by a common desired frequency. A slight phase offset between τ_1 and τ_2 is evident at their peak values, likely due to inter-joint coupling effects and differing dynamic characteristics between the hip and knee joints. In summary, the larger magnitude of τ_1 reflects the increased actuation demand at the hip joint, the smaller amplitude of τ_2 corresponds to the relatively load or inertia at the knee joint.




Figure 7. Torques of the hip and knee joint.

Figure 8 illustrates the time evolution of the adaptive gain $\tilde{a}$ over the 0–20 s simulation interval. At the onset, $\tilde{a}$ rapidly increases from its initial value to approximately 18, followed by a slow and steady rise toward 19 as time progresses. After the initial transient phase, the gain trajectory flattens, indicating that the adaptive law has effectively converged and the control algorithm is making only minor adjustments thereafter. This behavior suggests that the online parameter estimation mechanism rapidly identifies a near-optimal gain and subsequently maintains it with minimal variation, ensuring stable and consistent control performance throughout the simulation.




Figure 8. The adaptive gain parameters.

Figure 9 presents the end-effector trajectory error in both the x - and y -directions over a 0–20 s interval, with the vertical axis indicating the position error in meters. At $t = 0$ s, both error curves start near -0.04 m and exhibit rapid convergence, reaching near-zero values within approximately 0.5 s. For $t > 1$ s, the errors remain confined within an extremely narrow oscillatory band around zero, indicating highly precise trajectory tracking.




Figure 9. End effector trajectory tracking errors plot.

As shown in the inset, the steady-state fluctuations are on the order of $\pm 10^{-5}$ m, demonstrating that the proposed control algorithm ensures ultra-high positioning accuracy

and robustness throughout prolonged operation. These results confirm the effectiveness of the control scheme in maintaining fine end-point regulation under dynamic conditions.

Figure 10 illustrates the air pressures of PAMs. Each joint is controlled by a set of antagonistic PAMs, with one side maintaining a relatively constant air pressure, while the other side undergoes periodic adjustments over time. As shown in Figure 10a is the hip joint, one muscle maintains an approximate constant pressure of about 0.19 MPa, while the opposing muscle exhibits significant periodic fluctuations between approximately 0.14 and 0.17 MPa, achieving a cyclical flexion-extension motion. Similarly, at the knee joint shown in Figure 10b, one muscle maintains a stable pressure of about 0.187 MPa, while the other oscillates periodically within the range of 0.181–0.188 MPa. This pattern reflects a biological antagonistic mechanism, where one muscle provides stable support and the other generates the dynamic changes needed for movement. This coordination ensures the stability of the joint while producing the rhythmic movements required for gait.




Figure 10. (a) PAMs air pressures values of hip joint. (b) PAMs air pressures values of knee joint.

Tables 4 and 5 respectively summarize the tracking errors of the first and second joint under four control algorithms. The errors include Maximum error (MAX), Mean Absolute Error (MAE), Root Mean Square Residual (RMSR), Mean Absolute Percentage Error (MAPE), Mean Relative Deviation (MRD), Normalized Root Mean Square Deviation (NRMSD). The proposed control algorithm, SMC, GA-PID, and the torque compensation control [33]. To facilitate a more intuitive comparison of each algorithm's identification and tracking capabilities, the computed performance indicators are visualized in radar chart form in Figures 11 and 12.

Table 4. Tracking error of the hip joint under four algorithms.

Metric	The Proposed Controller	SMC	GA-PID	Torque Compensation Control
Stabilization time (s)	0.0350	0.5700	0.5900	0.0400
MAX (rad)	0.0363	0.2553	0.0889	0.0590
MAE (rad)	0.0007	0.0062	0.0339	0.0104
RMSR (rad)	0.0045	0.0311	0.0377	0.0142
MAPE (%)	0.0661	0.6212	3.3937	1.0449
MRD (%)	0.0661	0.5140	3.3937	0.0344
NRMSD (%)	0.0045	0.0311	0.0377	0.0142

To ensure a fair comparison, The SMC parameters are selected as follows: the sliding-surface coefficients are set to $\tilde{\lambda} = [1, 1]$, and the reaching-law gains are chosen as $\mathbf{k} = [10, 10]$. The PID baseline is tuned using a genetic algorithm. The search ranges of the PID gains are set to $K_p \in [10^{-3}, 20]$, $K_i \in [10^{-3}, 1]$, and $K_d \in [10^{-3}, 1]$. The fitness function is constructed as a weighted combination of tracking error, overshoot, rise time, and settling time, with the corresponding weights selected as $w_e = 0.5$, $w_o = 0.4$, $w_r = 0.9$,

and $w_s = 0.4$, respectively. For fitness evaluation, the nonlinear plant is linearized at the equilibrium point, yielding $G(s) = \text{tf}(400, [1, 5, 0])$. The GA uses a population size of 50 and runs for 10 iterations, with a crossover probability of 0.8, a mutation probability of 0.1, and a replication probability of 0.2; the gene length is set to 10. The optimized (K_p, K_i, K_d) obtained by GA is then used for the GA-PID controller in all comparative tests.

Table 5. Tracking error of the knee joint under four algorithms.

Metric	The Proposed Controller	SMC	GA-PID	Torque Compensation Control
Stabilization time (s)	0.1940	0.2500	0.5970	0.2280
MAX (rad)	0.1234	0.3663	0.3129	0.2640
MAE (rad)	0.0015	0.0044	0.0238	0.0064
RMSR (rad)	0.0110	0.0260	0.0443	0.0187
MAPE (%)	0.1500	0.4400	0.2380	0.2200
MRD (%)	0.1500	0.2800	0.2380	0.2200
NRMSD (%)	1.1000	2.6000	4.4300	1.5600




Figure 11. Radar chart of four algorithms on error 1.




Figure 12. Radar chart of four algorithms on error 2.

In the Figures 11 and 12, the proposed control algorithm exhibits the lowest mean absolute error and root-mean-square error among all control algorithms. This indicates that adaptive robust control algorithm achieves the highest tracking accuracy for both joints, maintaining smooth convergence without overshoot. Compared with SMC, adaptive robust control algorithm reduces high-frequency chattering significantly while preserving robustness to dynamic uncertainties. Relative to GA-PID, the proposed control algorithm avoids sluggish responses and steady-state offsets. Although GA-PID achieves reasonable tracking under nominal conditions, its fixed gains cannot compensate for time-varying pneumatic nonlinearities; in contrast, adaptive robust control algorithm adaptive terms

through online parameter adjustment, yielding faster response and smaller steady-state error. Compared with the torque compensation control, adaptive robust control algorithm maintains stability and precision even under model mismatch and external disturbances, benefiting from its combined adaptive and robust structure. The torque compensation control relies heavily on accurate dynamic parameters, whereas adaptive robust control algorithm relaxes this requirement and still delivers superior accuracy and smoother control profiles.

5. Experiment

The structural components of the PAM-driven biomimetic leg joints system are primarily fabricated from 6-series aluminum alloy, while the leg shafts are constructed from steel. All mechanical parts were machined by a contracted Computer Numerical Control (CNC) manufacturing center to ensure precision and consistency. The torso frame of the leg is assembled using 20 mm square aluminum extrusions. The thigh segment is composed of square-section carbon fiber tubes, whereas the shank utilizes circular carbon fiber tubes. The combined use of aluminum alloy and carbon fiber provides a favorable balance between strength and weight, effectively reducing the overall mass of the system.

The foot is fabricated by cutting an elastic rubber sphere, which offers both high ground traction and impact absorption. This design not only enhances contact stability during motion but also attenuates high-frequency vibrations at the foot ground interface. As a result, the pressure sensor mounted between the foot and the shank experiences reduced signal noise, thereby improving measurement reliability.

Figure 13 illustrates the proposed experimental platform consists of four major components. First, MATLAB 2024a is employed for data analysis, while the dynamics and control algorithm are established and verified in Simulink 2024a. The validated control laws are subsequently executed by a Beckhoff controller, which communicates with the upper-level PC via EtherCAT, enabling real-time parameter monitoring and online tuning. The controller outputs voltage commands to the valve control system, thereby regulating the air pressure supplied by the compressor and generating the contraction force of the PAMs to drive the PAM-driven biomimetic leg joints system. The resulting joint motion and sensor feedback are transmitted back to the Beckhoff controller, forming a closed-loop control architecture. Overall, the system integrates modeling and simulation, real-time computation, pneumatic actuation, and feedback monitoring into a unified experimental platform, providing a controllable, adjustable, and verifiable platform for PAM-driven biomimetic leg joints system motion control research.

As the leg moves, sensors such as angle encoders measure joint angles. These signals are continuously fed back to the Beckhoff controller, which updates the control algorithm in real time. By adjusting the valve control voltage dynamically, the system maintains the desired trajectory and ensures stable and adaptive pneumatic actuation throughout operation. Note that the proposed controller does not rely on online optimization for system, and all required matrix operations are low-dimensional, making the computational burden compatible with real-time embedded implementation. Table 6 summarizes the technical specifications of the key equipment used in this study.

To ensure consistency with the simulation study, the uncertainty in the experimental setup is introduced into the mass terms, $m_1(t) = \bar{m}_1 + \Delta m_1(t)$ and $m_2(t) = \bar{m}_2 + \Delta m_2(t)$, where $\Delta m_1(t) = 0.1 \sin(2\pi t)$ and $\Delta m_2(t) = 0.1 \sin(2\pi t)$. The desired joint trajectories are defined as time-varying sinusoidal signals and are expressed as follows:

$$\theta_1 = (15 \sin(t) - 45) \frac{\pi}{180}, \theta_2 = (15 \sin(t) + 75) \frac{\pi}{180}. \quad (65)$$




Figure 13. Schematic design for the experiment platform.

By solving the joint-to-actuator mapping model in real time, the corresponding desired PAMs force F_i and muscle contraction displacement x_i can be obtained. The joint tracking performance of the PAM-driven biomimetic leg joints system is illustrated by the swing motion curves shown below.

An analog voltage signal was issued via the Beckhoff EL3068 to drive the proportional valve. By adjusting the valve opening, the internal pressure of the PAMs were regulated, enabling control of the muscle's generated force. Channel-wise voltage outputs are reported in Figure 14 and constitute the raw data for single PAM control. Figure 14a measured voltages of the two antagonistic PAMs at the hip, Figure 14b those at the knee. Apart from the voltage the measured PAM air pressures are reported during experiments. The hip antagonistic air pressures P_1 and P_2 vary periodically and approximately in anti-phase within $P_1 \in [2.23, 3.90]$ bar and $P_2 \in [1.80, 4.00]$ bar. For the knee joint, the air pressures vary within $P_3 \in [2.42, 3.22]$ bar and $P_4 \in [1.15, 2.98]$ bar. Overall, the air pressure trajectories remain smooth and repeatable over the 30 s tests, indicating stable pneumatic regulation during high-frequency motion.

Table 6. Technical specifications of key experimental equipment.

Device	Model	Technical Specifications
Hardware Controller	Beckhoff CX2040	Intel Core i7 quad-core @ 2.1 GHz; Microsoft WES7 OS; compact DIN-rail mount; supports real-time control and data processing.
Analog output terminal	Beckhoff EL4008	8-channel analog output; 0–10 V; used to generate proportional-valve command signals.
Analog input terminal	Beckhoff EL3068	8-channel analog input; 0–10 V; 12-bit resolution; used to acquire sensor signals.
Real-time communication	EtherCAT	High-speed real-time communication among the controller and I/O terminals to ensure synchronization.
Joint angle encoder	Single-turn absolute encoder (BRT)	Measurement range: 0–360°; 12-bit resolution; 24 V DC supply; power-off memory; IP54 protection.
Proportional pressure regulator	SMC ITV2050-212N	Electro-pneumatic regulator; command signal: 0–10 V; gain levels: 0–9.




Figure 14. Experimental results of joint output voltage, (a) Hip joint proportional valve input voltage value; (b) Knee joint proportional valve input voltage value.

Figures 15 and 16 present the experimental results of joint trajectory tracking under adaptive robust control algorithm and SMC for the hip and knee joints, respectively. In each Figure, the upper subplot displays the actual joint angle overlaid with the desired trajectory, while Figures 17 and 18 show the corresponding tracking errors.




Figure 15. Experimental results of the hip trajectories tracking.




Figure 16. Experimental results of the knee trajectories tracking.




Figure 17. Experimental results of the hip trajectories tracking errors.

For the hip joint, both controllers successfully follow the periodic desired trajectory. However, adaptive robust control algorithm demonstrates substantially improved tracking precision and smoothness. In particular, adaptive robust control algorithm achieves noticeably smaller error magnitudes near trajectory inflection points, as illustrated in the magnified insets. The tracking error under adaptive robust control algorithm consistently remains within ± 0.1 rad, whereas the SMC exhibits transient spikes exceeding 0.2 rad, highlighting its reduced robustness to dynamic disturbances.




Figure 18. Experimental results of the knee trajectories tracking errors.

A similar performance advantage is observed for the knee joint. The adaptive robust control algorithm achieves more accurate trajectory tracking with minimal overshoot and steady-state error compared to SMC. The SMC response shows clear deviations during rapid transitions, with errors reaching up to 0.25 rad, whereas the adaptive robust control algorithm maintains errors predominantly within 0.05 rad. These results confirm that the adaptive robust control algorithm provides superior tracking accuracy, improved robustness, and reduced sensitivity to model uncertainties and high-frequency disturbances, making it well suited for control algorithm of PAM-driven biomimetic leg joints system operating under dynamic motion conditions.

The proposed controller is inherently modular and can be extended from a single leg to multi-leg or full quadruped platforms by applying the same joint-level controller to each actuated DOF. For an n -DOF system, the main online computations include evaluating the nominal dynamics terms and updating the adaptive law. In practice, the required matrix operations remain of moderate size and can be implemented efficiently using precomputed symbolic expressions, recursive dynamics, and lightweight linear-algebra routines. Since the uncertainty compensation is computed locally from measured states and error signals, the algorithm scales approximately linearly with the number of joints when implemented in a per-joint fashion, while cross-leg coupling can be incorporated through the nominal model without changing the overall structure. Moreover, the controller does not require iterative optimization online; thus, its computational burden is suitable for real-time deployment on embedded hardware. With an appropriate sampling rate consistent with pneumatic bandwidth, the control law can be executed within each control cycle on common real-time processors, making the proposed approach a feasible candidate for quadruped locomotion control.

6. Conclusions

This paper proposes an adaptive robust control algorithm to address the intrinsic nonlinearities, time-varying uncertainties, and bandwidth limitations of PAM-driven biomimetic leg joints system. By explicitly embedding second-order trajectory constraints into the control law, the adaptive robust control algorithm enables precise end-effector trajectory tracking while inherently satisfying the system's dynamic constraints, eliminating the need for complex auxiliary control algorithm. Within this model, adaptive robust control algorithm decomposes the system dynamics into nominal and uncertain components. Unknown disturbances and model uncertainties are mitigated through online parameter estimation, while an adaptive law adjustment mechanism is employed to increase robustness when tracking deviations exceed predefined thresholds. This allows the control algorithm to achieve fast error convergence and maintain stable performance, even in the presence of significant modeling errors and external disturbances. Simulation studies are conducted to evaluate the adaptive robust control algorithm against SMC under trajectory tracking and disturbance rejection scenarios. The results demonstrate clear advantages in terms of steady-state accuracy, disturbance attenuation, and tracking precision. Furthermore,

experimental validation on PAM-driven biomimetic leg joints system shows that joint angle tracking errors remain on the order of 10^{-2} rad during prolonged high-frequency motion, with negligible chattering. The adaptive robust control algorithm exemplifies the integration of adaptiveness and robustness, thereby constituting an advanced control system that is applicable to a wide range of uncertain nonlinear mechanical systems. Its generality and flexibility make it a promising candidate for future intelligent robotic and cyber-physical systems.

Although favorable performance is observed in the tested high-frequency trajectories, the achievable tracking bandwidth is fundamentally constrained by pneumatic actuation, including valve flow limits, air compressibility, pressure dynamics, and transport delays. These factors may degrade performance when the desired motion contains higher-frequency components or rapid torque reversals. Second, the current experimental evaluation primarily focuses on joint-angle tracking; a more comprehensive assessment under broader operating conditions is still needed. Third, extending the proposed approach to human–robot interaction or load-bearing scenarios may introduce additional challenges, such as safety-critical compliance requirements, interaction-force regulation, actuator saturation management, and contact-induced uncertainties, which are not explicitly addressed in the present study. Future work will therefore focus on improving performance under pneumatic bandwidth constraints, reporting additional system-level metrics, and validating the approach in interaction-rich and load-bearing tasks, as well as investigating scalability to multi-leg or full quadruped systems with real-time embedded implementation.

Author Contributions: Conceptualization, F.Q. and Z.L.; methodology, Y.-H.C. and F.Q.; validation, B.W. and Z.L.; formal analysis, Y.-H.C. and Y.X.; investigation, Q.Z.; writing—original draft preparation, Z.L.; writing—review and editing, Y.-H.C. and Q.Z.; supervision, F.Q.; funding acquisition, Y.X. All authors have read and agreed to the published version of the manuscript.

Funding: This research was supported by the Natural Science Foundation of Zhejiang Province (QN25E050084) and the Zhejiang Province Medical and Health Science and Technology Program (2025KY605).

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Buchner, T.J.K.; Fukushima, T.; Kazemipour, A.; Gravert, S.-D.; Prairie, M.; Romanescu, P.; Arm, P.; Zhang, Y.; Wang, X.; Zhang, S.L.; et al. Electrohydraulic Musculoskeletal Robotic Leg for Agile, Adaptive, Yet Energy-Efficient Locomotion. *Nat. Commun.* **2024**, *15*, 7634. [CrossRef]
2. Zhu, G.; Guo, W.; Chu, S. Bionic Concept and Synthesis Methods of the Biomimetic Robot Joint Mechanism for Accurately Reproducing the Motion Pattern of the Human Knee Joint. *Mech. Mach. Theory* **2024**, *204*, 105832. [CrossRef]
3. Zielinski, A. *Pneumatic Actuating Systems for Automatic Equipment: Structure, Modeling and Control*; Springer: Berlin/Heidelberg, Germany, 2006.
4. Wang, Q.Q.; Wang, W.; Hao, D.X.; Yun, C. Hysteresis modeling and application of McKibben pneumatic artificial muscles. *Chin. J. Mech. Eng.* **2019**, *55*, 73–80. [CrossRef]
5. Xie, S.L.; Liu, H.T.; Mei, J.P. Achievements and developments of hysteresis and creep of pneumatic artificial muscles. *J. Syst. Simul.* **2018**, *30*, 809–823.
6. Yu, J.; Edwards, C. Sliding mode observer design for actuator fault detection. In *Advances and Applications in Sliding Mode Control Systems*; Edwards, C., Spurgeon, S., Eds.; Springer: Cham, Switzerland, 2017; pp. 123–140.
7. Zhou, J.; Krstic, M. Adaptive Predictor Control for Stabilizing Pressure in a Managed Pressure Drilling System Under Time-Delay. *J. Process Control* **2016**, *40*, 106–118. [CrossRef]

8. Lewis, F.L.; Vrabie, D.; Syrmos, V. *Optimal Adaptive Control and Differential Games by Reinforcement Learning Principles*; IET Press: London, UK, 2012.
9. Zhang, J.; Zhao, R.; Jiao, S. Optimization of robust formation tracking control for traffic cone robots with matching and mismatching uncertainties: A fuzzy-set theory-based approach. *Expert Syst. Appl.* **2023**, *227*, 120292. [CrossRef]
10. Udwadia, F.E. A new perspective on tracking control of nonlinear structural and mechanical systems. *Proc. R. Soc. A Math. Phys. Eng. Sci.* **2003**, *459*, 1783–1800. [CrossRef]
11. Udwadia, F.E. A new approach to stable optimal control of complex nonlinear dynamical systems. *J. Appl. Mech.* **2013**, *81*, 031001. [CrossRef]
12. Chen, Y.-H.; Zhang, X. Adaptive robust approximate constraint-following control for mechanical systems. *J. Frankl. Inst B* **2010**, *347*, 69–86. [CrossRef]
13. Chen, X.; Zhao, H.; Wang, F.; Zhen, S.; Fang, J. Cooperative Working Performance of a Dual-Arm Robot System Optimised by a Neural Network Adaptive Preset Control. *Int. J. Sen. Netw.* **2024**, *45*, 54–65. [CrossRef]
14. Huang, K.; Wang, M.; Zhen, S. Udwadia–Kalaba equation based adaptive robust control for two-wheeled inverted pendulum system: Underactuation and uncertainty. *Int. J. Control Autom. Syst.* **2021**, *19*, 2420–2430. [CrossRef]
15. Liu, X.; Wu, C.; Zhang, S.; Zhou, Y.; Zhen, S.; Sun, H. Nash Game Based Optimal Adaptive Robust Control for Dual-Arm Coordinated Robots. In Proceedings of the 2025 International Conference on Advanced Robotics and Mechatronics, Portsmouth, UK, 1–3 August 2025; pp. 924–929.
16. Liu, Z.; Dong, F.; Zhao, X.; Han, J.; Huang, X. Trajectory Generation Strategy for Robots with Redundant DOF Based on the Generalized Udwadia–Kalaba Method: A State-Constrained Natural Solution Approach. *J. Frankl. Inst.* **2026**, *363*, 108301. [CrossRef]
17. Chen, S.; Han, T.; Dong, F.; Lu, L.; Liu, H.; Tian, X.; Han, J. Precision interaction force control of an underactuated hydraulic stance leg exoskeleton considering the constraint from the wearer. *Machines* **2021**, *9*, 96. [CrossRef]
18. Ma, C.; Huang, K.; Zheng, J.; Shao, K.; Li, C.; Sun, H. Fuzzy Optimization of Adaptive Integral Terminal Sliding Mode Control for Mechanical Systems with Uncertainties and Disturbances. *IEEE Trans. Ind. Electron.* **2025**, *72*, 11822–11831. [CrossRef]
19. Yu, R.; Zhao, H.; Zhen, S.; Huang, K. A Novel Approach for 2-Degrees of Freedom Redundant Parallel Manipulator Dynamics. *Adv. Mech. Eng.* **2017**, *9*, 1–12. [CrossRef]
20. Zheng, Y.; Zheng, J.; Shao, K.; Zhao, H.; Xie, H.; Wang, H. Adaptive trajectory tracking control for nonholonomic wheeled mobile robots: A barrier function sliding mode approach. *IEEE/CAA J. Autom. Sin.* **2024**, *11*, 1007–1021. [CrossRef]
21. Mylapilli, H.; Udwadia, F.E. Control of three-dimensional incompressible hyperelastic beams. *Nonlinear Dyn.* **2017**, *90*, 115–135. [CrossRef]
22. Udwadia, F.E. Constrained motion of Hamiltonian systems. *Nonlinear Dyn.* **2016**, *84*, 1135–1145. [CrossRef]
23. Dong, F.; He, J.; Zhao, X.; Han, J.; Huang, X.; Zhang, X. Analytical Modeling Method for Joint Friction in Harmonic Drive Robot. *J. Mech. Sci. Technol.* **2024**, *38*, 3101–3111. [CrossRef]
24. Zhao, X.; Cui, Z.; Chew, C.M.; Dong, F. Rendering Compact Formation and Trajectory Tracking for Cyber Unmanned Ground Vehicles. *Fundam. Res.* **2025**, *in press*. [CrossRef]
25. Yin, H.; Yu, D.; Yin, S.; Xia, B. Possibility-Based Robust Design Optimization for the Structural-Acoustic System with Fuzzy Parameters. *Mech. Syst. Signal Process.* **2018**, *102*, 329–345. [CrossRef]
26. Wu, Q.; Liu, F.; Yin, H.; Huang, J. Constraint-based adaptive robust control for active suspension systems under sky-hook model. *IEEE Trans. Ind. Electron.* **2021**, *69*, 5152–5164.
27. Chen, S.; Lu, M.; Dong, F.; Liu, H.; Tian, X.; Han, J. Generalized Multiphase Dynamic Modeling and Precision Interaction Force Control of a Walking Lower Limb Hydraulic Exoskeleton. *Appl. Bionics Biomech.* **2022**, *2022*, 2801719. [CrossRef]
28. Huang, K.; Xia, W.; Shao, K.; Sun, H. Adaptive Sliding Mode Control for Vehicle Active Anti-Roll System Based on Property Adjustable Barrier Function. *IEEE Trans. Intell. Veh.* **2025**, *10*, 3124–3133. [CrossRef]
29. Liu, X.; Zhen, S.; Sun, H.; Zhao, H. A novel model-based robust control for position tracking of permanent magnet linear motor. *IEEE Trans. Ind. Electron.* **2020**, *67*, 7767–7777. [CrossRef]
30. Chen, X.; Zhao, H.; Sun, H.; Zhen, S.; Mamun, A.A. Optimal adaptive robust control based on cooperative game theory for a class of fuzzy underactuated mechanical systems. *IEEE Trans. Cybern.* **2022**, *52*, 3632–3644. [CrossRef] [PubMed]
31. Khalil, H.K. *Nonlinear Systems*, 3rd ed.; Prentice-Hall: Upper Saddle River, NJ, USA, 2002.
32. Chen, Y.H. On the deterministic performance of uncertain dynamical systems. *Int. J. Control* **1986**, *43*, 1557–1579. [CrossRef]
33. Wang, B.; Wang, Y.; Huang, J.; Zeng, Y.; Liu, X.; Zhou, K. Computed torque control and force analysis for mechanical leg with variable rotation axis powered by servo pneumatic muscle. *ISA Trans.* **2023**, *140*, 385–401. [CrossRef] [PubMed]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Article

Myoelectric Controlled Bionic Robotic Hand for Voluntary Finger Motion Driven by Neuromuscular Intent

André Moreira ¹, Marco Pinto ¹, Miguel Fernandes ^{1,2}, João Costa ^{1,2}, Jorge Fidalgo ¹
and Alessandro Fantoni ^{1,2,*}

- ¹ Departamento de Engenharia Electrónica e Telecomunicações e de Computadores, Instituto Superior de Engenharia de Lisboa, Rua Conselheiro Emídio Navarro 1, 1959-007 Lisboa, Portugal; a47125@alunos.isel.pt (A.M.); a47091@alunos.isel.pt (M.P.); mfernandes@deetc.isel.ipl.pt (M.F.); jcosta@deetc.isel.ipl.pt (J.C.); jfidalgo@deetc.isel.ipl.pt (J.F.)
- ² Center of Technology and Systems (CTS) and Associated Lab of Intelligent Systems (LASI), 2829-516 Caparica, Portugal
- * Correspondence: afantoni@deetc.isel.ipl.pt

Abstract

Reliable control of robotic hands using residual muscle activity is challenging due to low-amplitude myoelectric signals, susceptibility to noise, and the need for real-time actuation. This paper presents a myoelectric-controlled robotic hand capable of voluntary independent finger motion. Surface myoelectric signals from the forearm are processed via amplification, filtering, and digital analysis to enable accurate detection of muscle activity. The system achieves independent and simultaneous actuation of five fingers using a tendon-driven, servo-actuated mechanism in a lightweight ABS structure. Experimental evaluation demonstrates finger actuation delays ranging from 314 ms to 650 ms, maximum holding strengths between 1.75 N and 4.07 N, and minimum gripping distances between 22 mm and 49 mm across all five fingers, with peak motor currents remaining below 0.7 A. Results validate consistent muscle activity detection, successful execution of individual and combined finger movements, and the robustness of the proposed design.

Keywords: bionic hand; mechatronic design; digital signal processing; state-based control; tendon-driven actuation; finger coordination; myoelectric control; prosthetics

1. Introduction

Worldwide, the prevalence of limb amputation has increased steadily over recent decades, driven by trauma, vascular disease, cancer, and congenital conditions. Recent global estimates indicate that tens of millions of individuals are living with limb loss, with traumatic events remaining one of the leading causes worldwide, particularly in low- and middle-income regions [1–3]. These amputations are commonly classified into two principal categories, lower-limb (LL) and upper-limb (UL) amputations, each associated with distinct anatomical levels, functional demands, and rehabilitation challenges [4]. While lower-limb amputations primarily compromise mobility and gait, upper-limb amputation has a uniquely profound impact on an individual's ability to interact with the surrounding environment. The upper limb, and in particular the hand, plays a central role in object manipulation, self-care, communication, and social interaction [5,6].

This paper presents the design, implementation, and experimental validation of a complete EMG-driven robotic hand system capable of independent per-finger actuation in real time. The proposed system integrates a custom analog and digital signal conditioning

pipeline for surface electromyography (sEMG), a computationally efficient state-machine control framework, and a fully assembled tendon-driven mechanical hand platform, all deployed on resource-constrained embedded hardware. While the system is motivated by and evaluated in the context of prosthetic hand replacement, the underlying framework is broadly applicable to any domain requiring intuitive, myoelectric-driven multi-finger control. In teleoperation, replicating fine hand movements remotely is highly valuable for tasks in hazardous environments, remote surgery, and dexterous manipulation, where the availability of greater computational resources further relaxes the constraints inherent to embedded prosthetic platforms [7]. Rehabilitation robotics represents a further application domain, where EMG-driven devices can support motor relearning and assist patients with residual neuromuscular function [8,9].

Despite the clear functional relevance of prosthetic devices, prosthesis abandonment and non-use remain significant challenges within the upper-limb amputee population. Reviews of long-term acceptance indicate that approximately 20–50% of upper-limb amputees ultimately reject or discontinue prosthetic use, commonly citing limited functionality, lack of intuitive control, physical discomfort, and insufficient relevance to everyday tasks [10]. Cosmetic prostheses, which prioritize appearance over function, are particularly prone to rejection, whereas functional solutions, especially body-powered and myoelectric devices, tend to achieve higher acceptance rates when they provide meaningful improvements in task performance and usability [11–13]. The consequences of upper-limb loss are not limited to physical impairment alone. Amputation frequently triggers a cascade of psychological and social challenges, including disruption of body image, reduced self-esteem, and diminished perceived autonomy [14,15]. Elevated rates of anxiety, depression, and social withdrawal have been reported among upper-limb amputees, particularly in cases where prosthetic devices fail to meet functional expectations or are abandoned altogether [16–18]. The visibility of the upper limb during social interaction further amplifies these psychosocial effects, often leading to perceived stigma and reduced participation in social and occupational environments [19–22].

The primary contribution of this work is the design, implementation, and experimental validation of a complete, low-cost EMG-driven prosthetic hand system capable of independent per-finger actuation in real time. This is achieved through the integration of a custom analog and digital signal conditioning pipeline, a computationally efficient state-machine control framework, and a fully assembled tendon-driven mechanical platform, all deployed on resource-constrained embedded hardware. Unlike approaches relying on pattern recognition or deep learning, the proposed deterministic architecture enables predictable, low-latency multi-finger control without the computational overhead that typically limits deployment on embedded prosthetic devices. It is worth noting that the necessity of independent multi-finger control for the manipulation of everyday objects remains an open question in the prosthetics literature. While many common grasping tasks can be accomplished with simplified two or three degree-of-freedom systems, independent finger actuation becomes particularly relevant for precision tasks such as pinch grasps, keyboard interaction, and fine object manipulation, where dexterity and adaptability are paramount. Moreover, the psychosocial benefits of a more biomimetic hand, in terms of body image, perceived autonomy, and social acceptance, provide additional motivation beyond purely functional considerations [14,15].

The remainder of this paper is organized as follows. Section 2 reviews the relevant state of the art in myoelectric prosthetic control, mechanical design, and embedded implementation. Section 3 describes the methods adopted for signal acquisition, conditioning, and processing. Section 4 presents the system overview, detailing the control architecture and mechanical platform. Section 5 reports the experimental results and discusses the

performance of the proposed system. Finally, Section 6 draws conclusions and outlines directions for future work.

2. State of the Art

This section reviews the current state of research in EMG-driven prosthetic hand control, with a focus on three interconnected challenges: the acquisition and conditioning of surface myoelectric signals, the control architectures used to translate these signals into coordinated finger movements, and the mechanical design considerations that influence actuation performance and wearability. Together, these aspects define the technical landscape within which the present work is situated.

Among the various human–machine interfaces explored in the literature, surface electromyography (sEMG) remains one of the most widely adopted approaches due to its non-invasive nature and its capacity to capture voluntary neuromuscular activity [23,24]. By translating muscle activation into control commands, myoelectric control systems aim to provide intuitive interaction between the user and the prosthetic device. Nevertheless, achieving reliable and responsive control continues to present significant technical challenges, particularly in the presence of signal variability, noise, and real-time processing constraints [25–28]. In recognition of these limitations, recent research has explored the integration of onboard vision systems into prosthetic platforms as a complementary or alternative control modality [29–31]. By assessing environmental context, vision-based approaches can reduce the physical and cognitive demands placed on the user and mitigate performance degradation associated with prolonged sEMG use, such as muscle fatigue and perspiration-induced electrode impedance changes [32]. Nevertheless, sEMG-based control remains the most widely adopted interface in both research and clinical settings due to its non-invasive nature, direct correspondence with neuromuscular intent, and established integration with embedded prosthetic hardware [33].

Several control strategies have been proposed to expand the functional capabilities of myoelectric prostheses, each presenting distinct trade-offs between dexterity, robustness, and computational complexity [34,35]. Pattern-recognition approaches combine feature extraction with classification algorithms to identify intended gestures, enabling a broader repertoire of movements and more natural interaction [34,36,37]. However, their performance is often sensitive to electrode displacement, muscle fatigue, perspiration, and variations in skin impedance, frequently requiring recalibration and structured user training [38–41]. Deep learning models further enhance predictive capability by automatically capturing complex relationships within EMG signals and reducing reliance on handcrafted features, often achieving higher classification accuracy [42,43]. Despite these advantages, their deployment in embedded prosthetic platforms remains constrained by substantial computational demands, increased power consumption, and potential latency. Consequently, deterministic signal-processing pipelines remain attractive for resource-constrained systems due to their predictable behavior, reduced processing overhead, and implementation simplicity, even if they provide fewer controllable states [44]. Control architectures for EMG-driven robotic hands are commonly categorized as continuous or discrete. Continuous approaches map EMG features directly to joint trajectories or actuator forces, promoting smooth and proportional motion but remaining susceptible to instability caused by inter-muscle crosstalk and overlapping activation patterns, especially in systems with multiple degrees of freedom [45]. Discrete methods instead rely on the detection of predefined activation states associated with specific finger movements or grasp configurations, typically offering greater robustness at the expense of motion expressiveness [46]. State-based control models extend the discrete paradigm by structuring system behavior as a finite set of operational states with defined transitions triggered by neuromuscular events. This deterministic framework supports predictable timing, facilitates

verification, and enables coordinated multi-finger actuation while maintaining modest computational requirements, making it particularly suitable for embedded assistive devices [47]. Mechanical design also influences prosthetic performance, particularly in multi-finger hands where actuator placement, transmission mechanisms, and structural weight affect responsiveness and wearability. Multibody modeling approaches have proven valuable in this context, enabling simulation-based evaluation of mechanical behavior, contact dynamics, and control performance prior to physical prototyping, with model reliability further enhanced through experimental validation [48–50]. Tendon-driven mechanisms are widely adopted for their ability to emulate biological motion while supporting lightweight structures, and servo-based actuation continues to be favored in research prototypes due to its compactness and ease of control [51,52].

Despite significant progress across these areas, several gaps remain [26,53–56]. Pattern recognition and deep learning approaches achieve high gesture classification accuracy but introduce computational overhead and sensitivity to signal variability that limit their deployment on embedded hardware. Continuous control methods offer smooth actuation but struggle with crosstalk and instability in multi-finger configurations. Existing deterministic approaches, while computationally efficient, have rarely been demonstrated in fully integrated systems capable of independent per-finger actuation validated on a physical platform. Furthermore, the majority of embedded prosthetic systems reported in the literature either limit the number of independently controlled fingers or rely on simplified grasp modes rather than individual digit control, leaving a gap between the dexterity achievable in laboratory conditions and what is deployable in a low-cost embedded system.

To address these limitations, this work proposes a computationally efficient, state-machine-based control framework aimed at enabling precise, real-time, independent multi-finger actuation on resource-constrained embedded hardware. The framework is integrated with a custom sEMG signal conditioning pipeline and validated on a fully assembled tendon-driven robotic hand. Beyond prosthetics, the proposed approach is applicable to teleoperation and rehabilitation robotics, where replicating fine hand movements is valuable and the availability of greater computational resources further relaxes the constraints inherent to embedded prosthetic platforms.

3. Methods

Surface electromyography (sEMG) is a non-invasive technique used to measure the electrical potentials generated by muscle fibers during contraction. Unlike quasi-periodic biosignals such as electrocardiographic activity, sEMG exhibits a stochastic nature and continuously varies according to factors such as contraction intensity, muscle fatigue, electrode placement, and inter-subject variability.

These signals are primarily characterized by their amplitude and frequency content. The amplitude typically ranges between 0 and 10 mV peak-to-peak [57,58]. During contractions with lesser force, amplitudes may remain below 4 mV, which complicates their detection, requiring amplification circuitry. Due to their inherently low amplitude, sEMG signals are highly susceptible to environmental interference such as the 50 Hz power-line noise or motion artifacts caused by electrode-skin displacement [59,60], necessitating multiple filtering stages to attenuate unwanted frequency components.

The second key characteristic of sEMG signals is the frequency content. The signal commonly presents a spectrum ranging between 20 Hz and 500 Hz, with most power concentrated between 20 Hz and 150 Hz. However, depending on electrode configuration, filtering strategy, and amplification stages, usable frequency components may extend further [57,58,61].

In this work, independent control of the bionic hand requires reliable detection of intended finger movements. For individuals with upper-limb loss, control commands can be derived from sEMG signals generated by residual forearm musculature. Finger motion is primarily controlled by flexor and extensor muscle groups located within the anterior and posterior compartments of the forearm. The anterior compartment includes the flexor digitorum superficialis, flexor digitorum profundus, and flexor pollicis longus, while the posterior compartment contains the extensor digitorum and related extensor muscles. These muscles are anatomically dense, layered, and partially overlapping. Due to volume conduction through biological tissues, electrical activity from deeper or adjacent muscles propagates to the skin surface, limiting spatial selectivity. As a result, isolating the activation of a single finger using surface electrodes becomes inherently challenging.

This anatomical proximity introduces crosstalk. Electrodes positioned over a target muscle may detect activity originating from adjacent muscles, particularly during synergistic finger movements [62,63]. Crosstalk magnitude depends on electrode size, inter-electrode distance, muscle depth, and subcutaneous tissue thickness. Although reducing electrode spacing improves spatial selectivity, it also decreases signal amplitude, creating a trade-off between resolution and signal strength.

To mitigate these limitations, five surface electrodes were distributed over major forearm muscle regions associated with individual fingers, as shown in Figure 1. Electrode positions were determined experimentally to maximize signal amplitude while keeping crosstalk at a manageable level. Rather than attempting strict anatomical isolation, finger activation was inferred based on relative amplitude differences across channels. Although crosstalk remains present, the electrode exhibiting the highest amplitude during contraction provides a practical indicator of the predominantly activated finger. This amplitude-based selection strategy offers a robust and computationally simple approach for distinguishing finger activation despite overlapping muscle activity.




Figure 1. Electrode placement on the wrist and forearm.

4. System Overview

The following subsections detail the three main components of the bionic hand system: signal conditioning, power circuitry, and the microcontroller-driven servo interface.

Following electrode placement and signal acquisition considerations, the complete signal processing and actuation pipeline is illustrated in Figure 2. This high-level diagram summarizes the transformation of raw sEMG signals into controlled mechanical motion of the bionic hand.

The system design integrates carefully tuned analog and digital signal conditioning with finite-state machine-based control algorithms. This combination enables reliable capture and interpretation of sEMG signals on a compact microcontroller platform, allowing independent multi-finger actuation with near real-time response. By relying on determinis-

tic, low-complexity algorithms instead of computationally intensive approaches such as neural networks, the proposed design achieves precise finger movements while maintaining minimal hardware requirements. These design choices address common challenges in the literature, including signal noise, latency, and the difficulty of implementing robust control on small embedded platforms.




Figure 2. Overview of the EMG signal processing and actuation pipeline.

The acquired surface EMG signals first pass through the analog signal conditioning stage, which includes instrumentation amplification, band-pass filtering, notch suppression, and envelope detection. This stage increases signal amplitude and attenuates noise components prior to digitization. The conditioned analog signals are then sampled by the Arduino Nano’s ADC, where digital filtering and real-time processing are performed. The processed signals are evaluated using a state-based control algorithm that determines intended finger activation. Finally, control commands are transmitted via the I²C interface to a PCA9685 PWM driver, which generates precise hardware-based control signals for the servo motors, actuating the tendon-driven bionic hand.

4.1. EMG Signal Conditioning Circuitry

Surface electromyography (sEMG) signals acquired from forearm electrodes are inherently low in amplitude, typically ranging from 50 μ V, when at rest, to 5 mV, with higher-level contraction, and are highly susceptible to external interference, motion artifacts, and electrode-skin impedance variations. Consequently, an analog front-end with high noise immunity and adequate amplification is required prior to digital processing.

The first stage consists of an operational-amplifier instrumentation amplifier (INA125), selected for its high input impedance, excellent common-mode rejection ratio (CMRR), and precise gain control. High input impedance minimizes signal attenuation caused by the electrode-skin interface, while the elevated CMRR is essential for suppressing common-mode disturbances such as power-line coupling and electromagnetic interference. Additionally, the INA125 gain can be calculated with the following equation,

$$G_{INA} = 4 + \frac{60 \text{ k}\Omega}{R_G} \quad (1)$$

Since the value chosen for R_G is of 60 k Ω , represented in Figure 3, the gain of the INA125 is five, preventing saturation in the presence of noise and motion artifacts, with the remaining amplification applied downstream following signal conditioning.

Figure 3 illustrates the active filtering stages implemented after the instrumentation amplifier. A cascaded second-order Sallen-Key low-pass and high-pass topology was adopted to realize an active band-pass filter with a passband from 12 Hz to 300 Hz.

This frequency range captures the dominant spectral components of myoelectric signals while attenuating motion artifacts and high-frequency electronic noise.




Figure 3. Complete analog filtering chain implemented after the instrumentation amplifier, including cascaded second-order Sallen-Key low-pass and high-pass filters (12–300 Hz) and a second-order active notch filter for mains interference rejection.

The low-pass stage was designed using a Butterworth approximation due to its maximally flat magnitude response in the passband (Figure 3). This characteristic prevents artificial attenuation of relevant EMG components and preserves the overall signal amplitude, which is particularly important for reliable envelope extraction and muscle activation detection.

The transfer function of the adopted low-pass topology is expressed as:

$$A(s) = \frac{1}{1 + \omega_c C_1 (R_4 + R_5) s + \omega_c^2 R_4 R_5 C_1 C_2 s^2} \quad (2)$$

where $\omega_c = 2\pi f_c$ represents the angular cutoff frequency.

The resistor values for the second-order Butterworth low-pass filter were obtained from:

$$R_{4,5} = \frac{a_1 C_1 \pm \sqrt{a_1^2 C_4^2 - 4b_1 C_1 C_2}}{4\pi f_c C_1 C_2} \quad (3)$$

where $a_1 = 1.4142$ and $b_1 = 1$, corresponding to the normalized coefficients of a second-order Butterworth response, and f_c denotes the cutoff frequency.

To complement the amplitude-flat response of the low-pass stage, the high-pass filter was implemented using a Bessel approximation (Figure 3). The Bessel response was selected for its nearly linear phase characteristic and minimal group delay distortion, which helps preserve the temporal structure of the EMG waveform. Maintaining waveform morphology is advantageous for downstream processing, particularly when timing information is relevant for gesture onset detection.

The transfer function of the adopted high-pass topology is expressed as:

$$A(s) = \frac{1}{1 + \frac{2}{\omega_c R_7 C_3} \frac{1}{s} + \frac{1}{\omega_c^2 R_7 R_8 C_3 C_4} \frac{1}{s^2}} \quad (4)$$

The design parameters are defined as:

$$a = \frac{2}{\omega_c C_3 R_7}, \quad b = \frac{1}{\omega_c^2 R_7 R_8 C_3 C_4} \quad (5)$$

The cascade of both filters results in a fourth-order band-pass response, providing sufficient roll-off to suppress out-of-band interference while maintaining stability and low implementation complexity.

Following band-pass filtering, an active notch filter centered at the power-line frequency of 50 Hz was introduced to attenuate mains interference (Figure 3). Power-line contamination is one of the dominant noise sources in biopotential acquisition and can significantly degrade signal quality if not properly suppressed.

A second-order active notch topology was selected to provide narrowband attenuation while preserving adjacent spectral components of the EMG signal. Excessively wide rejection bands may remove physiologically relevant information, particularly because portions of the EMG spectrum overlap the power-line frequency.

The notch filter transfer function is expressed as:

$$H_{\text{notch}}(s) = \frac{s^2 + \omega_0^2}{s^2 + \frac{\omega_0}{Q}s + \omega_0^2} \quad (6)$$

where $\omega_0 = 2\pi f_0$ and Q is the quality factor controlling the selectivity of the filter. Higher values of Q produce a narrower rejection band, enabling targeted suppression of mains interference with minimal distortion to the remaining signal bandwidth.

The frequency response of the implemented filtering chain is shown in Figure 4. The cascaded second-order low-pass and high-pass stages produce a fourth-order band-pass characteristic spanning 12 Hz to 300 Hz, effectively constraining the bandwidth to the spectral region of myoelectric signals.




Figure 4. Bode magnitude response of the analog filtering chain.

After filtering, the signal is further amplified using two cascaded non-inverting operational amplifier stages with gains of 11 and a variable gain, since R_{17} is a potentiometer, meaning the gain of the second non-inverting amplifier can be up to 11, yielding a maximum gain of 121. Distributing the gain across multiple stages helps prevent saturation while preserving bandwidth.

To mitigate DC offsets originating from the instrumentation amplifier reference and component mismatches, each stage is preceded by an RC high-pass network that blocks the DC component prior to high amplification. This prevents offset propagation that could otherwise drive the amplifiers into saturation.

A final RC coupling stage is placed after the last amplifier to remove any residual offset introduced during amplification. The chain is terminated with a unity-gain buffer to isolate the amplification block from subsequent stages and maintain gain stability (Figure 5).

After amplification, the conditioned sEMG signal is full-wave rectified using a precision rectifier based on a dual-op-amp topology (Figure 6). This configuration ensures accurate rectification of the signal relative to the instrumentation amplifier reference, avoiding the voltage drops associated with diode-based rectifiers.




Figure 5. Cascaded non-inverting amplification with AC coupling and buffering.




Figure 6. Full-wave precision rectification and final gain stage for ADC compatibility.

Following rectification, the signal is scaled using a non-inverting amplifier with a gain of three to match the input range of the microcontroller ADC, resulting in a total circuit maximum gain of 1851, assuming that the R_{17} potentiometer is 10 k Ω . At this stage, the reference for the non-inverting amplifier is the supply ground rather than the instrumentation amplifier reference (Figure 6). This adjustment shifts the fully rectified signal into a range compatible with the ADC while maintaining signal integrity. The buffer provided by this stage also isolates the rectifier from the ADC input, minimizing loading effects.

Figure 7 presents a simulation of the sEMG signal as it passes through the complete analog front-end, from raw acquisition to the ADC-ready output. The figure highlights how the signal is conditioned, amplified, rectified, and scaled for digital conversion. The simulation of the full 30-s signal required approximately 1 min and 13 s of computation time.




Figure 7. Simulation of the sEMG signal through the complete processing chain, from raw acquisition to ADC-ready output.

4.2. Power Management and Regulation

A dedicated multi-stage power distribution system was implemented to provide stable and electrically isolated supply rails for actuation, digital control, and sensitive analog signal conditioning, presented in Figure 8. The system is powered through a transformer-based supply that provides a regulated 12 V rail, which is used primarily to drive the MG996R servo motors through the PCA9685 driver module. Supplying the servos from a dedicated higher-voltage rail ensures adequate torque availability while preventing excessive current draw from lower-voltage circuitry. Given the high transient currents associated with servo startup and load variations, the 12 V rail was dimensioned to tolerate peak demands without significant voltage sag, thereby reducing disturbances that could propagate into the signal acquisition stages.




Figure 8. Circuit power supply.

A 7805 linear regulator derived a stable 5 V supply from the 12 V rail for low-voltage subsystems, including the Arduino Nano and analog conditioning circuitry. To minimize noise coupling between digital components and millivolt-level sEMG signals, the 5 V output was split into separate branches for analog and digital loads. The analog branch includes an LC filter to attenuate high-frequency switching noise, and careful grounding practices further reduce common-impedance interference, ensuring stable operation and clean signals for the instrumentation amplifier and subsequent analog stages.

4.3. Digital Control and Actuation Architecture

The digital control subsystem converts processed sEMG signals into coordinated motion of the bionic hand. An Arduino Nano samples the conditioned signals through its ADC, evaluates activation thresholds, and executes a state-based control algorithm in real time.

Servo actuation is handled by a PCA9685 16-channel PWM driver (Figure 9), which offloads PWM generation from the Arduino, minimizing timing jitter and processor overhead. Five MG996R high-torque servo motors, with a torque of approximately 9–11 kg·cm at 6 V [64], actuate the tendon-driven fingers. Independent PWM channels allow precise angular positioning and synchronized multi-finger movements, ensuring adequate tendon tension while maintaining responsive and accurate motion under load.




Figure 9. Digital control and actuation architecture.

Communication between the Arduino Nano and PCA9685 occurs over the I²C bus, providing a scalable and modular architecture that decouples control logic from actuation timing, improves reliability, and facilitates future expansions of the system.

The complete electronic system, integrating power distribution, analog signal conditioning, and digital control, was implemented on a single custom PCB. Figure 10 shows the layout, illustrating the separation of analog and digital domains to minimize interference and ensure signal integrity.




Figure 10. Custom PCB integrating power management, analog signal conditioning, and digital control for the bionic hand.

4.4. Software Implementation

The software architecture responsible for processing EMG signals and translating them into motor commands is organized into three distinct layers, each with a well-defined role. This separation ensures that time-critical operations remain lean and predictable, while more complex processing is handled without strict timing constraints.

The first layer is the ISR routine, triggered by a hardware timer at a fixed frequency of 1 kHz. Every millisecond, the ISR reads the raw analog value from the ADC pin corresponding to a given finger and writes it into a circular ring buffer. It then sets a flag to notify the system that a new sample is available.

Once the ISR deposits a raw sample and raises the flag, the main loop handles the second layer. This layer retrieves the raw value from the buffer and passes it through a digital signal processing pipeline. Two filtering stages are applied sequentially. First, a sixth-order band-pass filter with the same cutoff frequencies as the analog filter, 10 Hz and 300 Hz, to isolate the EMG frequency range, and secondly, a notch filter to suppress power-line interference centered at 50 Hz. Figure 11 shows the frequency response of both filters, illustrating the frequencies effectively attenuated, while Figure 12 presents the frequency content of a real sEMG signal through FFT, and the corresponding frequency content after filtering, demonstrating that both digital filters will ensure minimal noise, readying the samples for usage.




Figure 11. Frequency response of the digital filters.




Figure 12. Frequency content of a real sEMG signal before, and after filtering.

The resulting signal is then processed over a sliding window to compute its root mean square value. RMS is preferred over a simple moving average for EMG signals because it reflects the signal's energy rather than its amplitude, making it more sensitive to genuine muscle contractions and more resilient to slow baseline drift. The computed RMS value, representing the current state of the muscle, is stored in a shared variable for the next layer to access.

The third layer consists of a non-blocking finite state machine, with one instance running independently for each finger, enabling parallelization and making sure all fingers can be activated simultaneously. Its operation relies on the clean, processed values from the pipeline and is entirely decoupled from raw ADC data, filter internals, or timing concerns. Its purpose is to interpret the signal as intentional user input and issue the corresponding motor command. Each FSM transitions through four states (IDLE, WAITING_FOR_SIGNAL, CONFIRMING_ACTIVATION, and CLASSIFYING_SIGNAL) based on whether the processed signal crosses a predefined threshold, the duration of that crossing, and the signal's trend. Figure 13 presents the corresponding state machine diagram.




Figure 13. Finite state machine for finger control.

5. Results and Discussion

This section presents the mechanical platform used to validate the electronic hardware and control strategies described in the previous sections. The results reported in this chapter therefore reflect the combined performance of the signal acquisition, processing, and actuation architectures when deployed on a physical robotic hand.

To enable experimental validation, a multi-finger robotic hand was manufactured and assembled as part of this work (Figure 14). The mechanical design is based on the open-source InMoov hand, originally developed by Gaël Langevin within the InMoov project [65]. The original design was adopted as a structural reference due to its well-documented geometry, tendon-driven architecture, and suitability for low-cost, reproducible fabrication.




Figure 14. Manufactured robotic hand used for system validation: (a) dorsal view of the robotic hand and (b) motor mounting and tendon routing.

The assembled hand comprises five independently actuated fingers, each driven by a single MG996R servo motor (TowerPro, Shenzhen, China) mounted in the hand base. Motion is transmitted from the servo horns to the finger joints via fishing line tendons, which were selected over textile cord to minimize friction and elastic compliance. Each finger is returned to its open position by a passive spring mechanism, resulting in a normally-open configuration that closes upon servo actuation. Power and control signals are routed from the embedded electronics to the servo motors via direct wiring, with the Arduino Nano generating PWM signals to command each servo independently.

All hand components were fully 3D printed and mechanically assembled by the authors, and the complete electronic system and embedded software were developed independently and integrated into the hand structure. Several practical adaptations were introduced relative to the original design to accommodate the developed actuation system and experimental constraints. These include the use of fishing line instead of textile tendon cord to reduce friction and elastic compliance, modifications to the hand base due to the absence of the forearm structure, and adjustments to the motor mounting and tendon routing to improve force transmission between the servo motors and the fingers.

In order to test the correct operation of the robotic hand, five pre-recorded EMG signals—one for each finger—were processed through the analog and digital filters (Figure 15b).

In the resulting traces, higher amplitudes indicate active finger closure, while lower amplitudes correspond to no effort, meaning the finger is open. Amplitude levels vary between fingers due to differences in EMG signal strength and electrode placement.

The red vertical line in the graph (Figure 15a) marks the instant when the processed signals were sent to the hand, and an accompanying image captures the hand at that precise moment. At this instant, the pinky remains open, consistent with its signal never closing; the index and middle fingers are already fully closed, reflecting signals that had been in the closed state for approximately 1.3 s. The thumb and ring fingers show shorter closure durations (0.3 s), so they appear partially closed, just beginning to move.




Figure 15. Test of the robotic hand using pre-recorded EMG signals: (a) processed EMG signals for each finger and (b) robotic hand at the instant indicated by the red line in the signals. Higher amplitudes indicate active finger closure, while lower amplitudes correspond to no effort. Amplitude levels vary between fingers due to differences in EMG signal strength and electrode placement.

In addition to using pre-recorded signals, the system is capable of acquiring and processing EMG signals in real time. As demonstrated in Video S1 (Supplementary Materials), the robotic hand responds near-instantaneously to the user's myoelectric activity, performing coordinated finger closures and openings in synchrony with the recorded signals.

Experimental evaluation was conducted on five able-bodied participants, namely the authors of this work. While this represents a limited sample, the primary objective of the evaluation was to validate the correct operation of the signal acquisition, processing, and actuation pipeline rather than to assess generalization across a broader population. Each participant was instructed to reproduce a standardized sequence of finger movements, including individual finger flexion, combined finger motions, and full-hand opening and closing.

Table 1 presents the results of several metrics measured for each finger. The results were obtained through a combination of mechanical, electrical, and software-based tests. Mechanical evaluations included measuring the minimum gripping distance and maximum holding strength using digital calipers with a ± 0.01 mm resolution (Digikey, Thief River Falls, MN, USA) and a handheld newton force gauge with a range of 0.1–5 N (ATO, Golden Springs, CA, USA).

Additionally, programming and electrical tests captured the delay between receiving the signal and actuating the fingers of the hand, as well as the peak motor currents each servo consumed.

Table 1. Finger characteristics for the EMG-controlled robotic hand, including example images.

	Thumb	Index	Middle Finger	Ring Finger	Pinky
Minimum Gripping Distance [mm]	22	35	37	31	49
Maximum Holding Strength [N]	1.75	4.07	3.45	3.79	3.88
Finger Actuation Time [ms]	314	650	516	588	529
Peak Motor Current [A]	0.43	0.31	0.35	0.37	0.7

Across the fingers, minimum gripping distances ranged from 22 mm for the thumb to 49 mm for the pinky, while maximum holding strength varied from 1.75 N to 4.07 N. The thumb and pinky represent slight outliers in these measurements, which can be attributed to differences in finger design, tendon routing, and spring tension. The peak motor current remained well below 1 A for all fingers, indicating that the selected MG996R servos provide sufficient torque without excessive electrical demand. Again, the thumb and pinky exhibited slightly higher currents, consistent with the additional torque required due to their positions on the extremities of the hand. These fingers experience more bending in the tendons as they pass through the wrist, whereas the middle three fingers sit more directly above the wrist, allowing more direct routing.

Actuation delay, measured from signal acquisition to finger movement, ranged from 314 ms to 650 ms. This delay includes approximately 20 ms attributable to the signal conditioning filters and the Arduino Nano’s internal ADC sampling. Variation between fingers is primarily due to differences in motor range settings. After testing, some motors would draw excessive current without improving gripping performance, so the range for each finger was adjusted to optimize efficiency.

It is of note that each test was repeated ten times to account for minor trial-to-trial variations, and the values reported in Table 1 represent the average across these tests.

Finally, several limitations of the proposed system should be acknowledged in the context of practical application. First, the inherent challenge of isolating individual finger activations from surface EMG electrodes introduces crosstalk, particularly during synergistic finger movements, which may reduce the reliability of independent finger control in unconstrained use. Second, the measured maximum holding strength, ranging from 1.75 to 4.07 N across fingers, and the minimum gripping distance, ranging from 22 to 49 mm, impose practical constraints on the range of objects and tasks the hand can accommodate. Third, the robustness of the system to real-world signal variability, such as electrode displacement, muscle fatigue, perspiration, and inter-subject differences, was not evaluated in this work, and performance degradation under these conditions remains an open question. Furthermore, the electrode placement strategy adopted in this work targets specific anatomical landmarks on the intact forearm, which may not be directly transferable to amputee users where residual muscle volume and anatomy vary considerably depending on amputation level and individual morphology.

6. Conclusions

This work presented the design and implementation of a modular EMG-controlled bionic hand, integrating analog signal conditioning, digital processing, and multi-channel actuation into a cohesive and scalable platform. The system translates neuromuscular intent into controlled finger motion using a cost-effective and reproducible architecture.

The analog front end reliably amplifies and filters low-amplitude sEMG signals through instrumentation amplification, band-pass filtering, notch suppression, and envelope detection. Combined with the layered digital processing architecture and interrupt-

driven sampling, this ensures stable gesture detection while maintaining timing consistency and modularity. Actuation is performed via a PCA9685 PWM driver and high-torque servo motors, allowing synchronized, independent control of five tendon-driven fingers while minimizing computational overhead.

Future improvements could include adaptive thresholding, advanced classification algorithms, optimized tendon mechanics, or higher-resolution ADC acquisition. Overall, this work demonstrates a functional EMG-controlled bionic hand and establishes a structured hardware–software framework for neuromuscular interfaces, providing a strong foundation for scalable prosthetic research and human–machine interaction applications.

Supplementary Materials: The following supporting information can be downloaded at: <https://www.mdpi.com/article/10.3390/machines14030355/s1>, Video S1: Experimental demonstration of real-time control of the robotic hand using myoelectric signals acquired from forearm electrodes, highlighting the system’s ability to translate live muscle activity into coordinated finger movements.

Author Contributions: Conceptualization, A.M., M.P., A.F., M.F., J.C. and J.F.; Methodology, A.M. and M.P.; Software, A.M. and M.P.; Validation, A.F., M.F., J.C. and J.F.; Formal analysis, A.F. and J.C.; Investigation, A.M. and M.P.; Resources, A.F., M.F. and J.C.; Data curation, A.M. and M.P.; Writing—original draft preparation, A.M. and M.P.; Writing—review and editing, A.F., M.F. and J.C.; Visualization, A.M. and M.P.; Supervision, A.F., M.F., J.C. and J.F.; Project administration, A.F.; Funding acquisition, A.F. and M.F. All authors have read and agreed to the published version of the manuscript.

Funding: This work was sponsored by FCT—Fundação para a Ciência e a Tecnologia, within the Research Unit CTS—Center of Technology and Systems, reference UIDB/00066/2025.

Data Availability Statement: The data supporting the findings of this study are openly available in the GitHub repository at <https://github.com/Marco-Andre-Labs/EMG-Controlled-Robotic-Hand> (accessed on 16 February 2026).

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Resnik, L.; Borgia, M.; Clark, M. Function and Quality of Life of Unilateral Major Upper Limb Amputees: Effect of Prosthesis Use and Type. *Arch. Phys. Med. Rehabil.* **2020**, *101*, 1396–1406. [CrossRef] [PubMed]
2. Yuan, B.; Hu, D.; Gu, S.; Xiao, S.; Song, F. The global burden of traumatic amputation in 204 countries and territories. *Front. Public Health* **2023**, *11*, 1258853. [CrossRef]
3. Choi, Y.; Kim, E.; Oh, B.M. Global Overview of Acquired Upper Limb Amputation Epidemiology: A Systematic Review of Prevalence, Incidence, Level, and Etiology. *Ann. Rehabil. Med.* **2025**, *49*, 335–359. [CrossRef]
4. McDonald, C.; Westcott-McCoy, S.; Weaver, M.; Haagsma, J.; Kartir, D. Global prevalence of traumatic non-fatal limb amputation. *Prosthet. Orthot. Int.* **2021**, *45*, 105–114. [CrossRef]
5. Shahsavari, H.; Matourypour, P.; Ghiyasvandian, S.; Ghorbani, A.; Bakhshi, F.; Mahmoudi, M.; Golestannejad, M. Upper limb amputation; Care needs for reintegration to life: An integrative review. *Int. J. Orthop. Trauma Nurs.* **2020**, *38*, 100773. [CrossRef]
6. Resnik, L.; Borgia, M.; Cancio, J.; Heckman, J.; Highsmith, M.J.; Levy, C.; Phillips, S.; Webster, J. Understanding Implications of Residual Limb Length, Strength, and Range-of-Motion Impairments of Veterans With Upper Limb Amputation. *Am. J. Phys. Med. Rehabil.* **2022**, *101*, 545–554. [CrossRef]
7. Toet, A.; Kuling, I.A.; Krom, B.N.; van Erp, J.B.F. Toward Enhanced Teleoperation Through Embodiment. *Front. Robot. AI* **2020**, *7*, 14. [CrossRef] [PubMed]
8. Huo, Y.; Wang, X.; Zhao, W.; Hu, H.; Li, L. Effects of EMG-based robot for upper extremity rehabilitation on post-stroke patients: A systematic review and meta-analysis. *Front. Physiol.* **2023**, *14*, 1172958. [CrossRef]
9. Kiper, P.; Kopp, C.; Nicolas, Z.; Taupin, S.; Meroni, R.; Calabrò, R.S.; Kiper, A.; Federico, S.; Cieřlik, B. EMG-Driven Robotic Therapy for Neurological Rehabilitation: A Systematic Review and Meta-Analysis. *Technologies* **2026**, *14*, 119. [CrossRef]
10. Brack, R.; Amalu, E.H. A review of technology, materials and R&D challenges of upper limb prosthesis for improved user suitability. *J. Orthop.* **2021**, *23*, 88–96. [CrossRef]
11. Carey, S.L.; Lura, D.J.; Highsmith, M.J. Differences in Myoelectric and Body-Powered Upper-Limb Prostheses: Systematic Literature Review. *J. Prosthet. Orthot.* **2017**, *29*, P4–P16. . [CrossRef]

12. Salminger, S.; Stino, H.; Pichler, L.H.; Gstoettner, C.; Sturma, A.; Mayer, J.A.; Szivak, M.; Aszmann, O.C. Current rates of prosthetic usage in upper-limb amputees – have innovations had an impact on device acceptance? *Disabil. Rehabil.* **2022**, *44*, 3708–3713. [CrossRef] [PubMed]
13. Resnik, L.; Borgia, M.; Acluche, F. Brief activity performance measure for upper limb amputees: BAM-ULA. *Prosthet. Orthot. Int.* **2018**, *42*, 75–83. [CrossRef]
14. de Putter, C.; Selles, R.; Haagsma, J.; Polinder, S.; Panneman, M.; Hovius, S.; Burdorf, A.; van Beeck, E. Health-related quality of life after upper extremity injuries and predictors for suboptimal outcome. *Injury* **2014**, *45*, 1752–1758. [CrossRef]
15. Trent, L.; Intintoli, M.; Prigge, P.; Bollinger, C.; Walters, L.S.; Conyers, D.; Miguelez, J.; Ryan, T. A narrative review: Current upper limb prosthetic options and design. *Disabil. Rehabil. Assist. Technol.* **2020**, *15*, 604–613. [CrossRef] [PubMed]
16. Singh, S.; Saini, R.; Mathur, R.; Sarkar, S.; Sagar, R. The prevalence of depression in people following limb amputation: A systematic review and meta-analysis. *J. Psychosom. Res.* **2024**, *181*, 111677. [CrossRef]
17. Mckechnie, P.S.; John, A. Anxiety and depression following traumatic limb amputation: A systematic review. *Injury* **2014**, *45*, 1859–1866. [CrossRef] [PubMed]
18. Collu, R.; Ferrazzano, E.; Murgia, V.; Salis, C.; Barbaro, M. The Technological and Psychological Aspects of Upper Limb Prostheses Abandonment: A Narrative Review. *Prosthesis* **2025**, *7*, 167. [CrossRef]
19. Cordella, F.; Ciancio, A.L.; Sacchetti, R.; Davalli, A.; Cutti, A.G.; Guglielmelli, E.; Zollo, L. Literature Review on Needs of Upper Limb Prosthesis Users. *Front. Neurosci.* **2016**, *10*, 209. [CrossRef] [PubMed]
20. Seren, R.; De Tilio, R. As vivências do luto e seus estágios em pessoas amputadas. *Rev. Spagesp* **2014**, *15*, 64–78.
21. Pomares, G.; Coudane, H.; Dap, F.; Dautel, G. Psychological effects of traumatic upper-limb amputations. *Orthop. Traumatol. Surg. Res.* **2020**, *106*, 297–300. [CrossRef]
22. Sahu, A.; Sagar, R.; Sarkar, S.; Sagar, S. Psychological effects of amputation: A review of studies from India. *Ind. Psychiatry J.* **2016**, *25*, 4–10. [CrossRef]
23. Avila, E.R.; Williams, S.E.; Disselhorst-Klug, C. Advances in EMG measurement techniques, analysis procedures, and the impact of muscle mechanics on future requirements for the methodology. *J. Biomech.* **2023**, *156*, 111687. [CrossRef]
24. Young, A.J.; Kuiken, T.A.; Hargrove, L.J. Analysis of using EMG and mechanical sensors to enhance intent recognition in powered lower limb prostheses. *J. Neural Eng.* **2014**, *11*, 056021. [CrossRef] [PubMed]
25. Chan, F.; Yang, Y.S.; Lam, F.; Zhang, Y.T.; Parker, P. Fuzzy EMG classification for prosthesis control. *IEEE Trans. Rehabil. Eng.* **2000**, *8*, 305–311. [CrossRef] [PubMed]
26. Yadav, D.; Veer, K. Recent trends and challenges of surface electromyography in prosthetic applications. *Biomed. Eng. Lett.* **2023**, *13*, 353–373. [CrossRef]
27. Benatti, S.; Donati, E.; Moin, A.; Zanghieri, M.; Orlandi, M.; Burrello, A.; Artoni, F.; Micera, S.; Benini, L.; Rabaey, J.M. EMG Acquisition and Processing for Hand Movement Decoding on Embedded Systems: State of the Art and Challenges. *Proc. IEEE* **2025**, *113*, 256–286. [CrossRef]
28. Taleshi, M.; Yeung, D.; Trong, M.D.; Negro, F.; Deny, S.; Vujaklija, I. Impact of Noise on Deep Learning-Based Pseudo-Online Gesture Recognition with High-Density EMG. In *Proceedings of the 47th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), Copenhagen, Denmark, 14–17 July 2025*; IEEE: Piscataway, NJ, USA, 2025; pp. 1–5. [CrossRef]
29. Cirelli, G.; Tamantini, C.; Cordella, L.P.; Cordella, F. A Semiautonomous Control Strategy Based on Computer Vision for a Hand–Wrist Prosthesis. *Robotics* **2023**, *12*, 152. [CrossRef]
30. Wang, S.; Zheng, J.; Huang, Z.; Zhang, X.; Prado da Fonseca, V.; Zheng, B.; Jiang, X. Integrating computer vision to prosthetic hand control with sEMG: Preliminary results in grasp classification. *Front. Robot. AI* **2022**, *9*, 948238. [CrossRef] [PubMed]
31. Zandigohar, M.; Han, M.; Sharif, M.; Günay, S.; Furmanek, M.; Yarossi, M.; Bonato, P.; Onal, C.; Padir, T.; Erdogan, D.; et al. Multimodal fusion of EMG and vision for human grasp intent inference in prosthetic hand control. *Front. Robot. AI* **2024**, *11*, 1312554. [CrossRef]
32. Cirelli, G.; Tamantini, C.; Zollo, L.; Cordella, F. Computer vision for high-level control of prosthetic limbs: A literature review. *Comput. Vis. Image Underst.* **2026**, *265*, 104669. [CrossRef]
33. Jiang, N.; Chen, C.; He, J.; Jianjun, M.; Pan, L.; Su, S.; Zhu, X. Bio-robotics research for non-invasive myoelectric neural-interfaces for prosthetic control—a ten-year perspective review. *Natl. Sci. Rev.* **2023**, *10*, nwad048. [CrossRef]
34. Li, W.; Shi, P.; Yu, H. Gesture Recognition Using Surface Electromyography and Deep Learning for Prostheses Hand: State-of-the-Art, Challenges, and Future. *Front. Neurosci.* **2021**, *15*, 621885. [CrossRef]
35. Ameri, A.; Ghorbani, M. Recent Advances in EMG Pattern Recognition for Prosthetic Control. *J. Biomed. Phys. Eng.* **2020**, *10*, 129. [CrossRef]
36. Scheme, E.; Englehart, K. Electromyogram pattern recognition for control of powered upper-limb prostheses: State of the art and challenges for clinical use. *J. Rehabil. Res. Dev.* **2011**, *48*, 643–659. [CrossRef] [PubMed]
37. Zhang, X.; Li, X.; Samuel, O.W.; Huang, Z.; Fang, P.; Li, G. Improving the Robustness of Electromyogram-Pattern Recognition for Prosthetic Control by a Postprocessing Strategy. *Front. Neurobot.* **2017**, *11*, 51. [CrossRef] [PubMed]

38. Raghu, S.T.P.; MacIsaac, D.T.; Scheme, E.J. Self-supervised learning via VICReg enables training of EMG pattern recognition using continuous data with unclear labels. *Comput. Biol. Med.* **2025**, *185*, 109479. [CrossRef]
39. Kulwa, F.; Tai, P.; Sarwatt, D.S.; Asogbon, M.G.; Khushaba, R.; Oyemakinde, T.T.; Aboyeji, S.T.; Li, G.; Samuel, O.W.; Li, Y. A NMF-Based Non-Euclidean Adaptive Feature Extraction Scheme for Limb Motion Pattern Decoding in Pattern Recognition System. *IEEE Trans. Biomed. Eng.* **2026**, *73*, 742–755. [CrossRef]
40. Polo, A.; Cárdenas-Bolaño, N.; Ripoll Solano, L.A.; Luengas-Contreras, L.A.; Robles-Algarín, C. TWISS: A Hybrid Multi-Criteria and Wrapper-Based Feature Selection Method for EMG Pattern Recognition in Prosthetic Applications. *Algorithms* **2025**, *18*, 633. [CrossRef]
41. Ferdousi, A.; Islam, M.J.; Ahmad, S.; Islam, M.R.; Chowdhury, M.N.H.; Haque, F.; Ali, S.H.M.; Reaz, M.B.I. Complicacy in Electrode Position Shift and Its Solution in sEMG Pattern Recognition: A Review. *IEEE Sens. J.* **2025**, *25*, 26269–26288. [CrossRef]
42. Yang, Y.; Tao, Q.; Li, S.; Fan, S. EMG-Based Dual-Branch Deep Learning Framework With Transfer Learning for Lower Limb Motion Classification and Joint Angle Estimation. *Concurr. Comput. Pract. Exp.* **2025**, *37*, e70263. [CrossRef]
43. Zaim, T.; Abdel-Hadi, S.; Mahmoud, R.; Khandakar, A.; Rakhtala, S.M.; Chowdhury, M.E.H. Machine Learning- and Deep Learning-Based Myoelectric Control System for Upper Limb Rehabilitation Utilizing EEG and EMG Signals: A Systematic Review. *Bioengineering* **2025**, *12*, 144. [CrossRef]
44. Karim, M.; Khalid, S.; Lee, S.; Almutairi, S.; Namoun, A.; Abohashrh, M. Next Generation Human Action Recognition: A Comprehensive Review of State-of-the-Art Signal Processing Techniques. *IEEE Access* **2025**, *13*, 135609–135633. [CrossRef]
45. Batista, P.G.; Vieira, A.C.; Gaspar, P.D. Robotic Prostheses and Neuromuscular Interfaces: A Review of Design and Technological Trends. *Machines* **2025**, *13*, 804. [CrossRef]
46. Liao, L.; Yan, D.; Zhang, G. An identification method of human joint interaction torque based on discrete EMG signals. *Front. Bioeng. Biotechnol.* **2025**, *13*, 1596180. [CrossRef] [PubMed]
47. Pérez-Reynoso, F.; Farrera-Vazquez, N.; Capetillo, C.; Méndez-Lozano, N.; González-Gutiérrez, C.; López-Neri, E. Pattern Recognition of EMG Signals by Machine Learning for the Control of a Manipulator Robot. *Sensors* **2022**, *22*, 3424. [CrossRef]
48. Cosenza, C.; Niola, V.; Savino, S. A mechanical hand for prosthetic applications: Multibody model and contact simulation. *Proc. Inst. Mech. Eng. Part H J. Eng. Med.* **2018**, *232*, 819–825. [CrossRef]
49. Cosenza, C.; Niola, V.; Savino, S. A simplified model of a multi-jointed mechanical finger calibrated with experimental data by vision system. *Proc. Inst. Mech. Eng. Part K J. Multi-Body Dyn.* **2021**, *235*, 146441932095511. [CrossRef]
50. Al-Shahrabi, A.; Javid, M.; Fahmy, A.; Griffiths, C.; Li, C. Modeling the DLR-HIT II Robotic Hand: A Dual-Platform Simulation Approach in MATLAB and CoppeliaSim. *IEEE Access* **2024**, *12*, 181332–181340. [CrossRef]
51. Rodrigues, R.; Miranda, D.; Carvalho, V.; Matos, D. Design and Development of an EMG Upper Limb Controlled Prosthesis: A Preliminary Approach. *Actuators* **2025**, *14*, 219. [CrossRef]
52. Castro, M.C.F.; Pinheiro, W.C.; Rigolin, G. A Hybrid 3D Printed Hand Prosthesis Prototype Based on sEMG and a Fully Embedded Computer Vision System. *Front. Neurobot.* **2022**, *15*, 751282. [CrossRef]
53. Fleming, A.; Stafford, N.; Huang, S.; Hu, X.; Ferris, D.P.; Huang, H.H. Myoelectric control of robotic lower limb prostheses: A review of electromyography interfaces, control paradigms, challenges and future directions. *J. Neural Eng.* **2021**, *18*, 041004. [CrossRef]
54. Lee, H.; Lee, Y.; Jiang, L.; Chen, J.; Lam, W.L.; Lo, Y.T. Advances in myoelectric neuroprosthetics: A narrative review of current trends, challenges, and future directions. *Adv. Technol. Neurosci.* **2025**, *2*, 204–215. [CrossRef]
55. Salazar, M.; Portero, P.; Zambrano, M.; Rosero, R. Review of Robotic Prostheses Manufactured with 3D Printing: Advances, Challenges, and Future Perspectives. *Appl. Sci.* **2025**, *15*, 1350. [CrossRef]
56. Piyathilaka, L.; Sul, J.H.; Dunu Arachchige, S.; Jayawardena, A.; Moratuwage, D. Advances in EMG Signal Processing and Pattern Recognition: Techniques, Challenges, and Emerging Applications. *Electronics* **2026**, *15*, 590. [CrossRef]
57. Lu, C.; Xu, X.; Liu, Y.; Li, D.; Wang, Y.; Xian, W.; Chen, C.; Wei, B.; Tian, J. An Embedded Electromyogram Signal Acquisition Device. *Sensors* **2024**, *24*, 4106. [CrossRef] [PubMed]
58. Intelligence, A.; R., M.; Yavari, F.; Schoen, M. Feature Extraction of Forearm EMG Signals for Prosthetics. *Expert Syst. Appl.* **2011**, *38*, 4058–4067. [CrossRef]
59. Phukpattaranont, P.; Thiamchoo, N.; Neranon, P. Real-time identification of noise type contaminated in surface electromyogram signals using efficient statistical features. *Med. Eng. Phys.* **2024**, *131*, 104232. [CrossRef]
60. Boyer, M.; Bouyer, L.; Roy, J.S.; Campeau-Lecours, A. Reducing Noise, Artifacts and Interference in Single-Channel EMG Signals: A Review. *Sensors* **2023**, *23*, 2927. [CrossRef] [PubMed]
61. De Luca, C.J.; Donald Gilmore, L.; Kuznetsov, M.; Roy, S.H. Filtering the surface EMG signal: Movement artifact and baseline noise contamination. *J. Biomech.* **2010**, *43*, 1573–1579. [CrossRef]
62. Islam, M.A.; Sundaraj, K.; Ahmad, R.B.; Sundaraj, S.; Ahamed, N.U.; Ali, M.A. Cross-talk in mechanomyographic signals from the forearm muscles during sub-maximal to maximal isometric grip force. *PLoS ONE* **2014**, *9*, e96628. [CrossRef] [PubMed]

63. Fortune, B.C.; McKenzie, L.R.; Chatfield, L.T.; Pretty, C.G. Crosstalk Removal in Stimulus-evoked Forearm Electromyography. *IFAC-PapersOnLine* **2018**, *51*, 315–320. [CrossRef]
64. Technology, H. MG996R Metal Gear Servo Motor. Available online: https://www.handsontec.com/dataspecs/motor_fan/MG996R.pdf (accessed on 16 February 2026).
65. Langevin, G. InMoov—Open Source 3D Printed Life-Size Robot. Available online: <http://www.inmoov.fr> (accessed on 15 February 2026).

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Article

Robotic Harvesting of Apples Using ROS2 [†]

Connor Ruybalid, Christian Salisbury and Duke M. Bulanon *

Department of Engineering and Physics, Northwest Nazarene University, Nampa, ID 83686, USA; cruybalid@my.nnu.edu (C.R.); christiansalisbury@my.nnu.edu (C.S.)

* Correspondence: dbulanon@nnu.edu

[†] This paper is an extended version of our published paper: Ruybalid, C.; Salisbury, C.; Bulanon, D.M. Robotic Fruit Harvesting with Dual Camera System. In Proceedings of the 2025 IEEE Opportunity Research Scholars Symposium (ORSS), Atlanta, GA, USA, 22–24 April 2025; pp. 1–4.

Abstract

Rising global food demand, increasing labor costs, and farm labor shortages have created significant challenges for specialty crop production, particularly in labor-intensive tasks such as fruit harvesting. Robotic harvesting offers a promising long-term solution, yet its adoption in orchard environments remains limited due to unstructured conditions, variable lighting, and difficulties in fruit recognition and manipulation. This study presents an improved robotic fruit harvesting system, Orchard roBot (OrBot), developed by the Robotics Vision Lab at Northwest Nazarene University, with the goal of advancing autonomous apple harvesting applications. The updated OrBot platform integrates a dual-camera vision system consisting of an eye-to-hand stereo camera with a wide field of view for fruit detection and an eye-in-hand RGB-D camera for precise manipulation. The control architecture was redesigned using Robot Operating System 2 (ROS2) and Python, enabling modular subsystem development and coordination. Fruit detection was performed using a YOLOv5 deep learning model, and visual servoing was employed to guide the robotic manipulator toward the target fruit. System performance was evaluated through laboratory experiments using artificial trees and field tests conducted in a commercial apple orchard in Idaho. OrBot achieved a 100% harvesting success rate in indoor tests and a 75–80% success rate in outdoor orchard conditions. Experimental results demonstrate that the dual-camera approach significantly enhances fruit search efficiency and harvesting efficiency. Identified limitations include sensitivity to lighting conditions, end effector performance with varying fruit sizes, and depth estimation errors. Overall, the results indicate a positive potential toward effective robotic fruit harvesting and highlight key areas for future improvement in vision, manipulation, and system robustness.

Keywords: apple; harvesting; machine vision; robotics

1. Introduction

Global food systems are facing increasing pressure from population growth, rising production costs, climate variability, and persistent labor shortages. It is projected that global food production will need to double by 2050 to meet the demands of a population expected to reach nearly 10 billion [1]. At the same time, agriculture—particularly specialty crop production—remains heavily dependent on manual labor for tasks such as harvesting. In the United States, and especially in Idaho, one of the nation's top five producers of specialty crops, growers of fruits and vegetables are increasingly challenged by a shrinking

labor force and rising wages. Although the H-2A Temporary Agricultural Worker Program [2] has helped mitigate labor shortages, it represents a short-term and increasingly costly solution, highlighting the need for sustainable, long-term alternatives.

Technological advancements in precision agriculture, automation, and robotics offer promising pathways to address these challenges [3,4]. While robotics has been successfully integrated into highly structured environments such as automotive manufacturing, its adoption in agriculture—particularly in orchards—remains complex. Orchards present unstructured and dynamic environments characterized by variable lighting, occlusion from foliage, and natural variability in fruit location and tree geometry. These factors make robotic harvesting significantly more challenging than industrial automation, where workpieces are consistently positioned, and lighting conditions are controlled [5].

Despite these challenges, robotic fruit harvesting has been an active area of research for more than three decades. Early systems demonstrated feasibility using black-and-white cameras and simple robotic manipulators, and subsequent studies have explored a wide range of sensing technologies, including color, stereo, thermal, multispectral, and depth cameras [6]. Among the key technical challenges are reliable fruit recognition, precise visual servoing of robotic manipulators, and the development of end effectors capable of removing fruit without damaging either the produce or the tree [7]. Fruit recognition, in particular, has received significant research attention, as accurate detection is a prerequisite for successful harvesting [8].

The Robotics Vision Lab of Northwest Nazarene University has been developing a robotic harvesting platform, Orchard roBot (OrBot), to investigate these challenges in real-world orchard environments [9]. OrBot integrates robotic manipulation with machine vision systems to locate and harvest apples and has demonstrated promising results in commercial orchards. Although its harvesting speed remains slower than that of human labor, robotic systems offer the potential to operate continuously and complement human workers, including during nighttime hours when labor is unavailable. This study builds on previous iterations of OrBot [10] by exploring improved vision systems, control systems, and harvesting strategies, with the goal of advancing robotic fruit harvesting toward a more practical and economically viable solution to the growing labor challenges faced by specialty crop producers. The first iteration of OrBot used MATLAB (R2022b) to develop its control system, and it used only one camera for finding and locating the fruit.

The objectives of this study are:

1. To improve the control system of OrBot using Robot Operating System 2 (ROS2);
2. To improve the fruit vision system by adding another camera with a wider field of view;
3. To evaluate the performance of OrBot in fruit harvesting.

2. Materials and Methods

2.1. Fruit Harvesting Robot: Orchard Robot (OrBot)

For the study of robotic fruit harvesting, the Robotics Vision Lab of Northwest Nazarene University developed a robotic platform called Orchard Robot, OrBot. OrBot is composed of the following: the manipulator, the eye-in-hand vision system, the eye-to-hand vision system, the end effector, a control system, a power supply, and a tank-tread-based navigation platform (Figure 1).




Figure 1. OrBot and its components.

2.1.1. Manipulator

The manipulator is a third-generation Kinova robotic arm (Kinova Robotics, Boisbriand, QC, Canada) [11]. This arm has six degrees of freedom and has a reach of 902 mm. The arm has a full-range continuous payload of 2 kg, which is more than the weight of an apple fruit, and it has a maximum Cartesian translation speed of 500 mm/s. The average power consumption of the arm is 36 W.

2.1.2. Eye-in-Hand Vision System

The first version of OrBot has only one vision system [10] and is referred to in this paper as the eye-in-hand vision system. This vision system is integrated with the end effector. The vision system, composed of the color sensor, has a resolution of 1920×1080 pixels with a field of view up to 65 degrees. The depth sensor, which is enabled by both stereo vision and active IR sensing, has a resolution of 480×270 with a field of view of 72 degrees. The color sensor is used to detect the fruits, and the depth sensor is used to estimate the distance of the fruit from the end effector.

2.1.3. Eye-to-Hand Vision System

In the current version of OrBot, another vision system was added [12]. This vision system is fixed and does not move with the end effector. In this paper, we refer to this vision system as the eye-to-hand vision system. The ZED 2 Stereo Camera (Stereolabs, San Francisco, CA, USA) is used as the eye-to-hand vision system [13]. Unlike the eye-in-hand

vision system, the ZED 2 camera has a wider field of view (FOV) of 110° (H) $\times$ 70° (V) and a resolution of $2 \times (1920 \times 1080)$. It has a depth range of 0.3 m to 20 m.

2.1.4. End Effector

The end effector is a standard two-finger gripper. The gripper has a stroke of 85 mm and a grip force of 20 to 235 N, which could be controlled depending on the target. Customized 3D-printed fingers of the gripper were designed and developed to allow for a larger stroke for larger-sized fruit and to conform to the spherical shape of the fruit.

2.1.5. Control System

The control system used ZED Box Orin NX (Stereolabs, San Francisco, CA, USA) is powered by an NVIDIA Jetson Orin NX with an Ampere GPU, Arm Cortex CPU, and LPDDR5 memory, as well as multiple USB ports, Ethernet, HDMI, and ports for cameras [14]. The control system used UBUNTU as the operating system and Python 3.10 as the main programming language. A foldable keyboard and monitor were mounted to the chassis of the robot for easy access in the field.

2.1.6. Tank-Tread Navigation Platform

The tank-tread navigation platform is an HD2 Treaded ATR Tank Robot Platform [15]. It is driven by 2 24 VDC Gear Motors with encoders and has a payload capacity of 45 kg. The platform has been configured to be able to mount OrBot, its components, the power supply, and other materials as needed.

2.2. Object Detection Algorithm

YOLO (You Only Look Once) is a series of real-time object detection models built on PyTorch 2.0.0 and developed by Ultralytics [16]. This project used YOLO version 5 (YOLOv5). As a brief overview, the architecture of YOLOv5 contains a backbone, neck, and head. The backbone is the main body of the model, designed using the CSPDarknet53 structure, which is a convolutional neural network. The head is in charge of generating a final output. The neck connects the backbone and head.

YOLOv5 provides tools and Python scripts for training the model on a custom dataset, changing the weights to serve custom purposes. These new weights can then be loaded into the YOLOv5 architecture and used to perform inference.

When evaluating an object detection model like YOLO, it is common to use the mean average precision. The average precision refers to the average precision level of the model over every recall level. Precision measures the accuracy of the positive predictions made by the model, calculated as the ratio of true positives to the sum of true positives and false positives. Recall measures the model's ability to identify all relevant instances, calculated as the ratio of true positives to the sum of true positives and false negatives. To determine if a predicted object is a true positive or a false positive, an intersection over union (IoU) threshold must be determined. The IoU measures the overlap between the prediction and the true label. If the overlap of an object prediction is above the IoU threshold, then it is counted as a true positive. Otherwise, it is counted as a false positive [17].

2.3. Control System for Fruit Harvesting

2.3.1. Robot Operating System

The Robot Operating System (ROS) is a set of software libraries and tools for building robot applications. From drivers and state-of-the-art algorithms to powerful developer tools, ROS has the open-source tools needed for any robotics development [18]. One of the core advantages of ROS2 is that it is centered around modularity. ROS2 allows the robot to be broken down into subsystems that can be developed and run asynchronously.

In a typical ROS2 environment, each subsystem is organized into a node, which can be tested and run asynchronously with the rest of the system [19]. Each of these subsystems, or nodes, can communicate through topics, which are unidirectional asynchronous communication channels. Nodes can either publish to a topic (send messages to the topic) or subscribe to a topic (listen for messages on the topic). This communication can be established through publishers or subscribers created within each node. Topics can be used in one-to-one, one-to-many, many-to-one, or many-to-many relationships.

For more complicated communication, a service or an action may be used. In these communications methods, one node must host a service or action server, and other nodes may become clients of that service or action. Services and actions may create one-to-one or one-to-many relationships, as illustrated in Figure 2.




Figure 2. Example of ROS2 nodes with associated communication.

2.3.2. Designing a ROS2 Application for OrBot

The first task of the project was to redesign OrBot's software to utilize ROS2. In doing so, the primary language of the software was switched from MATLAB to Python. The first version of OrBot's control system was developed using MATLAB. While MATLAB did have some ROS2 capabilities, Python allowed much more freedom with the tools.

To integrate ROS2 into the OrBot system, subsystems needed to be identified and programmed as nodes. Major subsystems that were identified were:

1. Cameras;
2. Image processing algorithms;
3. The robotic arm;
4. The tank treads;
5. Apple picking algorithm;
6. Display;
7. Gamepad connection.

Each node was written as a class that inherited from the node class from the rclpy library. Using methods inherited from the node class, publishers and subscribers, as well as service and action servers and clients, were added in the constructor. Additionally, timers were created in the constructor for continuous repeated actions. Class methods were then created that would be called when a message was received or when a timer finished. For example, a 0.1 s timer in the camera nodes would call a method to retrieve a frame and publish it. Another example is a subscriber in the tank treads node that would call a

method to change the speed of the motors whenever a message was received on the motor speed topic. To use this class, an instance of it was created, and the spin() function was called. A basic outline of this format can be seen in Figure 3.




Figure 3. ROS2 nodes for fruit harvesting using OrBot.

2.4. Two-Camera System

2.4.1. Eye-in-Hand and Eye-to-Hand Systems

The next objective of the project was to add a second camera to the system with a wider field of view that would remain stationary with respect to the navigation platform. This meant that the camera needed to have eye-to-hand coordination with the robotic arm, which was a departure from the eye-in-hand coordination used by the first camera located behind the gripper.

The camera attached to the gripper was an Intel RealSense camera (RealSense, Inc., Cupertino, CA, USA) [20]. This camera and the eye-in-hand coordination used with it were excellent for performing the precise motions needed to pick apples. However, the eye-in-hand camera had a very small field of vision of 72 degrees and was unlikely to see most of the apples on a tree. A search algorithm could be implemented to have the arm move and allow the camera to scan the entire tree, but this process would be tedious and inefficient. A new camera was needed with a field of view that would encompass the entire tree.

The ZED 2, a stereo vision camera developed by StereoLabs (San Francisco, CA, USA), was utilized for this project. This camera had a much wider field of view than the eye-in-hand camera, with 120 degrees. The ZED 2 required an Nvidia graphics card to be used.

The ZED SDK allowed for the calculation of a 3D point cloud, where each pixel would be mapped to a Cartesian coordinate relative to the camera. Using the point cloud, the location of an apple identified in the image could be found with reference to the camera. This facilitated the three-dimensional location of the camera using the eye-to-hand camera.

2.4.2. Eye-to-Hand Coordination

Creating eye-to-hand coordination meant translating what the camera saw to coordinates for the robotic arm to move to. The vector from the camera to the apple can be found using the 3D point cloud. However, to move the arm in front of the apple, the vector from the base of the robotic arm to the apple needed to be known. This can be calculated using the vector from the camera to the apple, $\vec{r}_{CA}$, and the vector from the base of the robot to the eye-to-hand camera, $\vec{r}_{RC}$, as shown in Figure 4.




Figure 4. Position vector relationship between the camera, fruit, and manipulator.

The vector from the robot to the camera was measured physically and subtracted from the vector obtained from the point cloud. By translating the coordinates, eye-to-hand coordination was established using the following equation,

$$\vec{r}_{RC} + \vec{r}_{CA} = \vec{r}_{RA}$$

Position vector $\vec{r}_{RC}$ defines the location of the eye-to-hand camera with respect to the origin of the robot, which is defined by the extrinsic camera parameter matrix,

$$\begin{bmatrix} 1 & 0 & 0 & -0.361 \text{ m} \\ 0 & 1 & 0 & +0.795 \text{ m} \\ 0 & 0 & 1 & -0.175 \text{ m} \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

Position vector $\vec{r}_{CA}$ defines the position of the fruit with respect to the eye-to-hand camera. When these two position vectors are added, we get the position vector $\vec{r}_{RA}$, which was used to direct the end effector towards the apple fruit, and the precise picking movement was controlled by the eye-in-hand camera, as shown by the yellow position vectors.

2.5. Evaluation of Fruit Harvesting

The performance of OrBot for fruit harvesting was evaluated, both in the laboratory and in a commercial field (Figure 5). In the laboratory, two artificial trees with five artificial fruits randomly positioned on the trees were used for the test. Lighting in the laboratory was provided by the room's fluorescent lights. During the test, the fruit picking success rate and the picking time were measured. The picking success rate was evaluated as the ratio of the total number of fruits harvested to the total number of fruits harvestable by the robot. When the eye-to-hand camera searches for fruit, it records the fruits that are reachable, and these reachable fruits are counted as harvestable fruits. OrBot was programmed to start from the first artificial tree, harvest the fruits, and move to the second artificial tree to

harvest the remaining fruits. The fruits were then re-attached to the trees, and the test was replicated five times.




Figure 5. Evaluation of fruit harvesting with (a) laboratory test and (b) commercial orchard test.

OrBot was also tested in a commercial apple orchard, Symms Fruit Ranch, which is located in Caldwell, Idaho, USA. The apple variety in the orchard is Pink Lady apples. The apple trees, which are semi-dwarf, are trained using a spindle-type structure and were spaced about two meters apart. Twenty trees on the same row were used for the harvesting test. Similar to the laboratory test, OrBot was programmed to start from the first tree in the row. The eye-to-hand camera searched the fruits in its field of view and identified harvestable fruits. When there were no more harvestable fruits in its field of view, OrBot moved forward until a harvestable fruit was identified, and the robot stopped. The harvesting test was conducted during midday and with fine weather conditions in November 2024 and November 2025.

3. Results

3.1. Indoor Test

In all the replications of the indoor test, OrBot was able to successfully recognize and pick all the fruits. Figure 6 shows the harvesting sequence of finding the fruits, centering on the fruit, and harvesting the fruit. It can be seen from the figure the difference between the field of view of the two cameras. The camera-to-hand has a wider field of view, while the camera-in-hand has a narrower field of view.

3.2. Outdoor Test

For the outdoor test, harvesting was considered a success if the target fruit was recognized and removed from the tree. Figure 7 shows (a) the eye-to-hand camera view when searching for a target fruit and (b) the eye-in-hand camera view for precision movement to pick fruit. Figure 7c,d shows the posture of the robot during a fruit search and fruit picking. OrBot moved autonomously along the row and stopped when a target fruit was

recognized. OrBot harvested 120 apple fruits from the row of twenty trees with a success rate of 75%. In the second harvesting season, OrBot harvested 150 apple fruits from the same row of trees with a success rate of 79%. The diameters of the fruits ranged from 6 cm to 8 cm, and the fruit weights ranged from 100 g to 200 g. There were several fruits that had indentations in them that were caused by the gripper, but these were still counted as a successful harvest.




Figure 6. Fruit harvesting sequence in the laboratory test.




Figure 7. Fruit harvesting test in a commercial orchard showing the (a) eye-to-hand view, (b) eye-in-hand view, (c) OrBot in search mode, and (d) OrBot picking fruit.

4. Discussion

The 100% harvesting success rate of OrBot in the indoor test proves the capability of OrBot to pick fruit autonomously. In the indoor test, OrBot moved from one tree to the other without any issues. The indoor test also showed that the two-camera system facilitated searching for the fruits. The eye-to-hand vision system searched for fruits while it was moving. Once a target fruit was identified, OrBot stopped and transferred control to the eye-in-hand and conducted a more precise motion by centering on the target fruit before it approached and picked the fruit. After the fruit was picked, control was transferred again to the camera-to-hand vision system.

When moving to the commercial orchard, the harvesting success rate of OrBot decreased to about 77%. The failed attempts can be attributed to hardware and software failures. The software failures are (a) failure to recognize fruit and (b) fruit was recognized, and OrBot tried to pick another fruit (tracking failure). The hardware failures are that (c) a fruit was recognized but not picked by the end effector (end effector failure) and (d) a fruit was recognized, but the end effector was not able to approach the fruit (manipulator failure). The distribution of these failed attempts due to these factors is shown in Table 1.

Table 1. Distribution of harvesting failures.

Factors for Harvesting Failures	Percentage (%)
(a) Fruit recognition failure	16
(b) Tracking failure	17
(c) End effector failure	50
(d) Manipulator failure	17

In the failure to recognize the fruit, there were a few instances where OrBot was in front of the tree with fruit, and there was no fruit that was recognized. The main reason for this was the lighting condition. A review on fruit robotic harvesting performance listed lighting variation as one of the exogenous disturbances that affected harvesting performance [21]. When the robot and the fruit position caused a backlighting condition, the fruit appeared dark in the image. This is unlike the indoor tests, where there were no issues with fruit recognition because the fruits were always front-lit. However, when the robot was repositioned, it was able to recognize the fruit. One of the solutions that will be tried in future tests will be to include all different types of lighting conditions, such as cloudy, sunny, front light, and back light, in the training of the fruit recognition algorithm.

The other failure due to the software was when the robot harvested a different fruit from the target fruit. This was considered a failed harvesting because the target fruit was not harvested. However, after the different fruit was harvested, the robot moved towards the target fruit and harvested it. There were also instances where the robot would center between the target fruit and a nearby fruit, and it would remain in this loop. Currently, the size of the fruit in the image is used for object-tracking purposes. The relative position of the fruit with respect to the robot origin should be added for object tracking.

In the instance where the fruit was recognized but not picked by the end effector, it was an issue of the end effector design [22]. The current end effector uses a standard two-finger gripper that has customized fingers designed to capture the curvature of the apple fruit. As mentioned in the results, the harvested apple fruits had a large variation in size. The robot had difficulty with harvesting smaller fruits as the fruits slipped out of its fingers during rotation. A study for an end effector for tomatoes noted that the fruit's major diameter affected fruit grasping [23]. The end effector needs to be modified to ensure that the fruit is picked. The use of vacuum pickers [24] was also tested during this harvesting season, as

shown in Figure 8. Initial tests showed that it improved the picking efficiency to more than 90%, but further tests and design modifications are needed to validate its performance.




Figure 8. OrBot harvesting apples using a vacuum-based end effector.

There were also a few times where the fruits were recognized, but the manipulator did not move. The robot estimated that the fruit was outside of its work envelope. The issue was with the eye-to-hand vision system. During the search section of the eye-to-hand vision system, the control system used the cloud point data combined with the color image to estimate the distance of the fruit. If one of the cloud points on the fruit is selected for distance estimation, and if this point falls outside the work envelope, then the robot will not move. The overall area of the fruit should be used and averaged to estimate the distance of the fruit. In addition, it will be useful to combine the infrared distance sensor of the camera-in-hand to double-check the target fruit's distance.

5. Conclusions

This study presented an improved robotic fruit harvesting system, Orchard roBot (OrBot), designed to address labor shortages in specialty crop production through automation. By integrating a dual-camera vision system, transitioning the control architecture to Robot Operating System 2 (ROS2), and implementing a modular, Python-based software framework, the updated OrBot platform demonstrated enhanced fruit detection, system coordination, and harvesting performance compared to previous iterations.

Experimental evaluation in both laboratory and commercial orchard environments confirmed the effectiveness of the proposed approach. OrBot achieved a 100% harvesting success rate in controlled indoor conditions and a 75–80% success rate in outdoor orchard tests, with improved performance observed after orchard pruning. The dual-camera configuration proved particularly effective, enabling an efficient wide-area fruit search using the eye-to-hand vision system and precise manipulation using the eye-in-hand system. These results validate the feasibility of combining complementary vision modalities to overcome the limited field of view and precision constraints inherent in single-camera harvesting systems.

Field testing also revealed key challenges that must be addressed to further improve reliability and commercial viability. Variable lighting conditions, particularly backlighting, reduced fruit recognition accuracy, highlighting the need for more robust training datasets and adaptive vision algorithms. Limitations in the current gripper design affected harvesting performance for smaller fruits, suggesting that alternative end effector designs, such as vacuum-based systems, may offer improved consistency. Additionally, depth estimation and object-tracking errors in cluttered environments emphasize the importance of integrating multi-sensor fusion and improved distance estimation strategies.

Overall, the results of this study demonstrate that robotic fruit harvesting is a viable and promising solution for augmenting human labor in orchard environments. While further refinement is required to increase speed, robustness, and economic feasibility, the advancements presented in this work represent a significant step toward practical deployment of autonomous harvesting systems capable of supporting sustainable food production in the face of growing labor constraints.

Author Contributions: Conceptualization, D.M.B.; methodology, D.M.B., C.R. and C.S.; software, C.R. and C.S.; validation, C.R., C.S. and D.M.B.; formal analysis, C.R., C.S. and D.M.B.; investigation, C.R., C.S. and D.M.B.; resources, D.M.B.; data curation, C.R., C.S. and D.M.B.; writing—original draft preparation, D.M.B., C.R. and C.S.; writing—review and editing, D.M.B., C.R. and C.S. supervision, D.M.B.; project administration, D.M.B.; funding acquisition, D.M.B. All authors have read and agreed to the published version of the manuscript.

Funding: This project was supported in part by the Idaho State Department of Agriculture through the Specialty Crop Block Grant Program.

Data Availability Statement: The data presented in this study are available on request from the corresponding author.

Acknowledgments: The authors would like to acknowledge the support of Symms Fruit Ranch for allowing the use of OrBot in their commercial orchards. GenAI (Copilot for Microsoft 365) was used to check on grammar and improve sentence structure. GenAI was not used to generate any data, figures, or results.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. van Dijk, M.; Morley, T.; Rau, M.L.; Saghai, Y. A meta-analysis of projected global food demand and population at risk of hunger for the period 2010–2050. *Nat. Food* **2021**, *2*, 494–501. [CrossRef] [PubMed]
2. Wei, X.; Campbell, B.L.; Khachatryan, H.; Brumfield, R.G. What Firms Hire H-2A Workers? Evidence from the US Ornamental Horticulture Industry. *HortScience* **2023**, *58*, 375–382. [CrossRef]
3. Karunathilake, E.M.B.M.; Le, A.T.; Heo, S.; Chung, Y.S.; Mansoor, S. The Path to Smart Farming: Innovations and Opportunities in Precision Agriculture. *Agriculture* **2023**, *13*, 1593. [CrossRef]
4. Botta, A.; Cavallone, P.; Baglieri, L.; Colucci, G.; Tagliavini, L.; Quaglia, G. A Review of Robots, Perception, and Tasks in Precision Agriculture. *Appl. Mech.* **2022**, *3*, 830–854. [CrossRef]
5. Kootstra, G.; Wang, X.; Blok, P.M.; Hemming, J.; van Henten, E. Selective Harvesting Robotics: Current Research, Trends, and Future Directions. *Curr. Robot. Rep.* **2021**, *2*, 95–104. [CrossRef]
6. Hou, G.; Chen, H.; Jiang, M.; Niu, R. An Overview of the Application of Machine Vision in Recognition and Localization of Fruit and Vegetable Harvesting Robots. *Agriculture* **2023**, *13*, 1814. [CrossRef]
7. Williams, H.A.M.; Jones, M.H.; Nejati, M.; Seabright, M.J.; Bell, J.; Penhall, N.D.; Barnett, J.J.; Duke, M.D.; Scarfe, A.J.; Ahn, H.S.; et al. Robotic kiwifruit harvesting using machine vision, convolutional neural networks, and robotic arms. *Biosyst. Eng.* **2019**, *181*, 140–156. [CrossRef]
8. Zhao, Y.; Gong, L.; Huang, Y.; Liu, C. A review of key techniques of vision-based control for harvesting robot. *Comput. Electron. Agric.* **2016**, *127*, 311–323. [CrossRef]
9. Waltman, J.; Buchanan, E.; Bulanon, D.M. Nighttime Harvesting of OrBot (Orchard RoBot). *AgriEngineering* **2024**, *6*, 1266–1276. [CrossRef]

10. Bulanon, D.M.; Burr, C.; DeVlieg, M.; Braddock, T.; Allen, B. Development of a Visual Servo System for Robotic Fruit Harvesting. *AgriEngineering* **2021**, *3*, 840–852. [CrossRef]
11. Discover Our Gen3 Robotic Arm. Available online: <https://www.kinovarobotics.com/product/gen3-robots> (accessed on 10 February 2026).
12. Ruybalid, C.; Salisbury, C.; Bulanon, D.M. Robotic Fruit Harvesting with Dual Camera System. In Proceedings of the 2025 IEEE Opportunity Research Scholars Symposium (ORSS), Atlanta, GA, USA, 1 April–1 July 2025; pp. 1–4.
13. ZED 2 Versatile Stereo Camera for Spatial Perception. Available online: <https://www.stereolabs.com/products/zed-2> (accessed on 10 February 2026).
14. ZED Box Orin. Available online: <https://www.stereolabs.com/docs/embedded/zed-box-orin> (accessed on 10 February 2026).
15. HD2 Treaded ATR Tank Robot Platform. Available online: <https://www.superdroidrobots.com/store/product=789> (accessed on 1 April 2026).
16. Ultralytics YOLOv5. Available online: <https://docs.ultralytics.com/models/yolov5/> (accessed on 10 February 2026).
17. Ruybalid, C. Improving OrBot. Bachelor’s Thesis, Northwest Nazarene University, Nampa, ID, USA, April 2025.
18. ROS—Robot Operating System. Available online: <https://www.ros.org/> (accessed on 10 February 2026).
19. ROS 2 Documentation. Available online: <https://docs.ros.org/en/foxy/index.html> (accessed on 10 February 2026).
20. Intel RealSense. Available online: <https://www.intel.com/content/www/us/en/products/sku/98319/intel-realsense-depth-module-d410/specifications.html> (accessed on 10 February 2026).
21. Zhou, H.; Wang, X.; Au, W.; Kang, H.; Chen, C. Intelligent robots for fruit harvesting: Recent developments and future challenges. *Precis. Agric.* **2022**, *23*, 1856–1907. [CrossRef]
22. Silwal, A.; Davidson, J.R.; Karkee, M.; Mo, C.; Zhang, Q.; Lewis, K. Design, integration, and field evaluation of a robotic apple harvester. *J. Field Robot.* **2017**, *34*, 1140–1159. [CrossRef]
23. Li, Z.; Miao, F.; Yang, Z.; Chai, P.; Yang, S. Factors affecting human hand grasp type in tomato fruit-picking: A statistical investigation for ergonomic development of harvesting robot. *Comput. Electron. Agric.* **2019**, *157*, 90–99. [CrossRef]
24. You, K. Development of an adaptable vacuum based orange picking end effector. *Agric. Eng. Int. CIGR J.* **2019**, *21*, 58–66.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Article

Design and Industrial Integration of Automated Coordinate Measuring Machines for Automotive Production

Eva M. Rubio ^{1,*}, Marian Sáenz-Nuño ², Marta M. Marín ¹ and David Gómez ¹

¹ Department of Manufacturing Engineering, Industrial Engineering School, Universidad Nacional de Educación a Distancia (UNED), St/Juan del Rosal 12, E-28040 Madrid, Spain; mmarin@ind.uned.es (M.M.M.); dgomez317@alumno.uned.es (D.G.)

² IIT Mechanical Engineering Department, Instituto Católico de Artes e Industrias (ICAI), Comillas Pontifical University, Alberto Aguilera 25, E-28015 Madrid, Spain; msaenz@iit.comillas.edu

* Correspondence: erubio@ind.uned.es

Abstract

Recent advances in machine design, automation, and industrial digitalization have transformed Coordinate Measuring Machines (CMMs) from standalone inspection devices into fully integrated elements of automated manufacturing systems. In the automotive sector, CMMs increasingly operate in workshop, near-line, and in-line environments, interacting with production equipment and contributing directly to process control and zero-defect manufacturing strategies. This paper presents a structured methodology for the industrial deployment of automated CMMs in automotive mechanical manufacturing. The proposed approach is illustrated through an industrial use case involving the dimensional inspection of mechanically machined components under real production conditions. The methodology addresses machine design selection, sensor configuration, environmental constraints, and multi-axis architectures, as well as validation and acceptance procedures based on the ISO 10360 series. Particular attention is given to the integration of CMMs within automated manufacturing systems, including robustness against thermal variations, vibrations, and contamination, and the use of metrological data for feedback to machining processes. Rather than introducing new metrological principles, the proposed approach focuses on the structured integration of established engineering practices into a coherent lifecycle-based deployment framework. Based on industrial experience, the proposed methodology is illustrated through an industrial case study to support the reliable of automated dimensional inspection, reduce measurement-related risks, and support the integration of CMMs as active components of modern automated manufacturing systems.

Keywords: Coordinate Measuring Machine (CMM); machine design; industrial automation; automotive manufacturing; dimensional inspection; metrology integration; smart manufacturing; zero defect manufacturing

1. Introduction

Coordinate Measuring Machines (CMMs) have traditionally been used for off-line dimensional verification in metrology laboratories [1]. In recent years, their use has expanded beyond laboratory environments [2–6], allowing CMMs to be applied closer to manufacturing processes and production equipment. As a consequence, dimensional inspection is increasingly positioned closer to the production process and contributes directly to manufacturing control activities [4,5].

This evolution is particularly relevant in the automotive sector, where high production volumes coexist with increasingly demanding dimensional and geometric tolerances. In such environments, CMMs are progressively deployed in workshops, near-line, and in-line locations, interacting with production equipment and supporting activities such as process monitoring and machine tool adjustment [3,4]. For this reason, the design and integration of CMMs can affect the robustness and performance of the production process.

Despite significant technological progress, the industrial deployment of CMMs is often addressed in a fragmented manner [3]. Key aspects such as machine selection, metrological validation, adaptation to industrial environments, and operation during series production are frequently addressed as independent stages. As a result, a unified lifecycle-based approach that considers the measurement system as a whole is often lacking [2]. This separation can result in inadequate design decisions, increased metrological risks, and limited automation capability when the CMM is deployed in a real production environment [4].

Accordingly, the aim of this work is to propose a structured methodology for the industrial deployment of automated CMMs in automotive manufacturing. The proposed approach integrates machine design decisions, validation procedures, and production operation requirements within a coherent lifecycle-based framework. The objective is to ensure reliable and robust dimensional inspection under real industrial operating conditions. This lifecycle-based approach has been represented in Figure 1.




Figure 1. Lifecycle-based methodology for the industrial deployment of automated Coordinate Measuring Machines (CMMs), integrating machine design decisions, metrological validation and acceptance, and series production operation within automated manufacturing systems.

This approach is based on international standards and industrial experience in the manufacturing of mechanical components for the automotive sector and is illustrated through a representative industrial case study. The objective is to enhance the robustness and consistency of dimensional metrology in alignment with the principles of industrial automation and smart manufacturing. Within this framework, CMMs are considered automated machines integrated into the production process rather than standalone inspection devices.

In line with this approach, it is necessary to analyze CMM from the perspective of their design as automated industrial systems. Aspects such as machine mechanical architecture, sensor configuration, the incorporation of additional axes, and adaptation to the working environment directly affect both metrological performance and the capability for integration into automated manufacturing systems. For this reason, the following section reviews the main CMM configurations and their evolution toward workshop and production-line environments, focusing on the design implications that enable reliable operation as active components of modern automated manufacturing systems.

Rather than addressing CMM selection solely as a specification-matching exercise, this work proposes a lifecycle-based engineering perspective that integrates dimensional requirements, environmental conditions, validation standards, and automation constraints into a unified decision framework [1,4]. By structuring the deployment of CMM as part of an industrial risk-management and automation strategy [3,5], the proposed approach bridges the gap between metrological validation and real production environments, ensuring that accuracy, robustness, and system availability are considered from the earliest project stages through series production [2].

It should be noted that the contribution of this work does not lie in the introduction of new metrological principles or fundamentally new selection criteria. Instead, the novelty resides in the structured integration of established engineering practices—such as tolerance-based selection, environmental assessment, and standardized validation—into a coherent lifecycle-based framework specifically adapted to the deployment of automated CMM systems in manufacturing environments.

Based on the identified gap, this work aims to structure the industrial deployment of automated CMM within a lifecycle-based engineering framework. The main contributions of this work can be summarized as follows:

- Integration of lifecycle phases into a unified framework for CMM deployment in automated manufacturing environments.
- Explicit linkage between tolerance– Maximum Permissible Error (MPE) criteria, environmental conditions, and automation requirements within the decision-making process.
- Application-oriented validation strategy that connects standardized metrological verification with industrial acceptance and production performance.

The remainder of this paper is organized as follows. Section 2 reviews the main CMM architectures and design considerations relevant for automated manufacturing environments. Section 3 presents the proposed lifecycle-based methodology for industrial deployment and validation of automated CMMs. Section 4 illustrates the methodology through an industrial case study and discusses the main results. Finally, Section 5 summarizes the conclusions and outlines future research directions.

2. Coordinate Measuring Machines as Industrial Automated Systems

In automated manufacturing environments, CMMs are not only used for metrological purposes. Their evolution has been driven not only by advances in measurement technology, but also by developments in machine design, control systems, and industrial production requirements [2,3,5]. Mechanical architecture, sensor configuration, and adaptation to the operating environment therefore influence metrological performance and the use of CMMs in production lines [1,4,6].

2.1. Machine Architectures

One of the primary design factors in assessing the suitability of a Coordinate Measuring Machine for integration into automated environments is its mechanical architecture [1]. Figure 2 illustrates the most common CMM architectures used in industrial contexts, highlighting their accessibility, metrological performance, and suitability for automated inspection [1,6].

In automotive manufacturing, bridge-type CMM configurations are the predominant solution, as they provide a well-balanced compromise between measurement volume, accessibility, and metrological performance [1]. Their structural rigidity and well-defined kinematic behavior enable the inspection of a wide range of machined components and support reliable automation under industrial operating conditions [1,6].




Figure 2. Typical Coordinate Measuring Machine (CMM) architectures used in industrial environments: (a) bridge-type CMM, (b) cantilever CMM, and (c) bridge CMM equipped with a rotary table, highlighting their implications for accessibility, metrological performance, and suitability for automated inspection.

Cantilever CMM configurations are increasingly used in workshop and near-line applications, especially for small- to medium-sized components [1]. While offering lower structural rigidity than bridge-type machines, they provide enhanced accessibility and support automated loading and unloading, making them suitable for applications where flexibility and production-line integration are key requirements [1,6].

The integration of additional axes, such as rotary tables, is a key design feature of CMMs intended for automated applications [1,4]. These axes enhance access to geometric features, limit the need for multiple probing configurations, and streamline measurement programs [4]. In automated inspection systems, the reduction in reconfigurations improves repeatability, shortens cycle times, and increases the overall robustness of the measurement process [4,6,7].

From a lifecycle perspective, CMM architecture selection directly influences accessibility, structural stability, and integration constraints, which are critical factors in defining selection criteria and validation requirements in automated production environments.

Table 1 summarizes the main characteristics of typical CMM architectures with respect to their suitability for automated industrial deployment [1].

Table 1. Comparison of typical Coordinate Measuring Machines (CMMs) architectures with respect to automation-oriented design criteria.

CMM Architecture	Metrological Accuracy	Structural Robustness	Accessibility to Measurement Volume	Suitability for Automation	Typical Industrial Application
Bridge-type CMM	High	High	Medium	High	Wide range of machined automotive components requiring tight tolerances
Cantilever CMM	Medium to High	Medium	High	Very High	Near-line and workshop inspection of small- to medium-sized parts

Table 1. *Cont.*

CMM Architecture	Metrological Accuracy	Structural Robustness	Accessibility to Measurement Volume	Suitability for Automation	Typical Industrial Application
Bridge CMM with Rotary Table	High	High	Very High	Very High	Automated inspection of complex geometries and multi-sided components
Gantry-type CMM	Very High	Very High	Medium	Medium	Large automotive components and structures

Note: The qualitative assessment presented in this table is based on industrial experience and literature review, and is intended as an indicative comparison rather than a standardized classification.

2.2. Sensors and Multisensor Configurations

The selection of sensors and their configuration play a decisive role in both the metrological performance and the automation capability of a Coordinate Measuring Machine [1,4]. In automotive mechanical manufacturing, contact-based probing systems remain the predominant solution due to their high repeatability and suitability for tight dimensional and geometric tolerances, as well as for features with limited accessibility, such as holes and internal surfaces [1,4,8].

Within contact probing systems, the choice between touch-trigger probes and continuous scanning probes involves a trade-off between robustness, flexibility, and measurement capability [4,6]. Touch-trigger probes are widely used in manufacturing environments because of their simplicity, reliability, and ease of implementation, whereas scanning probes enable a more comprehensive evaluation of form and profile characteristics [4]. However, scanning systems typically require more careful configuration and are more sensitive to environmental conditions and system dynamics [6].

The adoption of rotary probe heads and multisensor configurations has become a key enabler for automated dimensional inspection [1,4]. These systems reduce the need for manual probe changes and allow access to multiple measurement orientations without operator intervention [6]. In addition, the combination of contact probes with non-contact sensing technologies, such as optical or laser sensors, extends the measurement capabilities of CMMs and enables inspection strategies adapted to different stages of the manufacturing process [4,8].

From the perspective of automation, multisensor configurations contribute to reduced non-productive times, increased system flexibility, and improved integration within production lines characterized by high product variability [3,8]. As a result, sensor configuration must be considered a fundamental design decision when deploying CMMs as automated inspection systems in production environments [3,5].

From a lifecycle perspective, the selection of sensing and probing technologies affects measurement strategy definition, achievable accuracy, and inspection cycle time, thus playing a key role in both validation procedures and production integration.

2.3. Design Implications for Accuracy and Robustness

Design decisions in a Coordinate Measuring Machine (CMM) affect not only metrological accuracy but also the robustness of the system under real industrial operating conditions. Factors such as structural stiffness, material selection, axis kinematics, and compensation systems directly influence the ability of the machine to maintain stable performance in the presence of external disturbances [1,9].

In industrial production environments, thermal variations, vibrations, and contamination are major sources of measurement error. Modern CMMs address these effects through integrated temperature sensing, advanced compensation algorithms, and passive

and active vibration isolation, enabling reliable measurement performance under variable operating conditions [1,6].

From an automation perspective, robustness to environmental variations is a fundamental requirement. When a CMM is integrated into a production line, it must ensure reliable and predictable operation, as even minor deviations in measurement performance can directly affect manufacturing process stability [3,8].

From a lifecycle perspective, environmental conditions such as temperature variations, vibrations, and contamination directly impact measurement reliability and must be considered as key constraints in machine selection, validation, and long-term operational performance.

2.4. Evolution Toward Workshops and In-Line Environments

Dimensional metrology is no longer confined to laboratory environments, but is increasingly applied in workshops, near-line and in-line contexts [1,8]. This transition is driven by the need to shorten process feedback loops and to use metrological data as an active element of manufacturing control rather than solely for post-process verification [3].

CMMs designed for workshop, near-line, and in-line environments incorporate enhanced protection against contamination, advanced environmental compensation systems, and mechanical architectures that facilitate automated part loading and unloading [1]. In in-line measurement applications, CMMs can be integrated with robots, material handling systems, and machine tools, enabling them to operate as active components of automated manufacturing systems rather than as passive inspection devices [3,8]. This shift supports proactive process control strategies and contributes to the implementation of smart manufacturing and zero-defect production approaches [2,5].

This evolution highlights the need for structured engineering approaches that address the design, validation, and integration of CMMs as automated industrial systems throughout their entire lifecycle. In this context, the following section introduces a methodology for the industrial deployment of automated CMMs, with particular emphasis on design-oriented decision-making, metrological validation, and integration within automated manufacturing environments [3,5].

From a lifecycle perspective, automation requirements and production integration constraints influence system configuration, cycle time compatibility, and operational robustness, linking design decisions with validation and series production performance.

3. Methodology for Industrial Deployment of Automated CMMs

In this work, terms such as availability, robustness, and predictability are used from an industrial perspective to refer to the capability of the measurement system to maintain stable performance, consistent results, and operational continuity under real production conditions.

The deployment of an automated CMM should be addressed as a structured engineering process rather than as a simple equipment installation. In automated manufacturing environments, decisions taken during the early project stages have a significant impact on measurement reliability, system availability, and effective integration within the production process [3,6,8].

For this reason, machine design, metrological capability, environmental conditions, and automation requirements must be considered jointly from the initial project phase [6,8].

The methodology proposed in this work follows a lifecycle-based approach to the measurement system and is organized into three main phases [3,10]:

- the project phase, focused on the selection of the machine and the measurement system;

- the validation and acceptance phase, aimed at reducing industrial and metrological risk prior to production start-up; and
- the series production phase, focused on long-term robustness, maintenance, and integration with automated production systems.

3.1. Project Phase: Machine and Measurement System Selection

The project phase plays a key role in ensuring the reliable operation of the CMM as an automated industrial system. Design decisions are therefore guided by product requirements and process constraints, with particular emphasis placed on the criticality of the characteristics to be inspected [1,6]. To address this need, a structured decision framework is defined to ensure that the selected configuration supports automated dimensional inspection under real production conditions.

Based on these principles, the proposed framework supports machine and measurement system selection during the project phase by integrating tolerance requirements, MPE-based criteria, and industrial environment conditions, as illustrated in Figure 3 [6,7].




Figure 3. Decision framework for Coordinate Measuring Machines (CMMs) selection in the project phase, illustrating the progressive elimination of non-compliant configurations, the evaluation of trade-offs between competing criteria, and the identification of the final solution based on industrial decision drivers. The framework integrates product requirements, metrological validation constraints, and environmental conditions within a lifecycle-based decision process.

Figure 3 illustrates the proposed decision framework as a structured and iterative selection process. It incorporates successive filtering steps to eliminate non-compliant

configurations, followed by a trade-off analysis to balance conflicting criteria such as accuracy, productivity, and robustness. The final selection is driven by key industrial decision factors, ensuring compliance with tolerance requirements, MPE constraints, and environmental conditions.

3.1.1. Geometric and Metrological Criteria

The first step of the decision-making process consists of identifying the set of geometric characteristics that must be verified using the CMM. These include both final product characteristics and intermediate process characteristics that influence subsequent manufacturing operations. Each characteristic is associated with a specific geometric tolerance and a defined level of criticality, determined by its impact on product functionality and process stability [7]. For each characteristic to be measured, the Coordinate Measuring Machine (CMM) must provide sufficient resolution and discrimination capability to reliably detect critical features. This requirement extends beyond nominal accuracy and includes aspects such as repeatability, long-term stability, and sensitivity to environmental variations [1,4]. Consequently, geometric accessibility, probing strategies, and sensor configuration must be evaluated in conjunction with the mechanical architecture of the machine [1,4,6].

During the project phase, these product and process requirements are translated into quantifiable metrological constraints that support objective CMM selection. This approach ensures consistency between quality planning, measurement system capability, and the objectives of production automation [1,6].

3.1.2. Tolerance- and MPE-Based Selection

In industrial practice, when the maximum permissible error (MPE) of a Coordinate Measuring Machine is known, an additional capability criterion is frequently considered in which the task-specific measurement uncertainty (U 's) should remain below approximately 10% of the declared MPE. This condition is commonly used as a practical indicator of the adequacy of the measurement system for a given application [1,6]. However, for critical characteristics, a final verification of task-specific measurement uncertainty (e.g., according to ISO 15530-3 or simulation-based approaches such as ISO 15530-4) is recommended as part of the acceptance phase to ensure the adequacy of the measurement system for the intended application.

A practical and widely used approach for the preselection of CMM configurations consists of comparing the tolerance intervals (TI) of the characteristics to be inspected with the maximum permissible errors (MPE) declared in accordance with the ISO 10360 series [10–13]. This comparison allows an initial filtering of candidate machines during the project phase [1,6].

A commonly adopted guideline is the so-called “10% rule,” according to which the measurement uncertainty should not exceed 10% of the tolerance interval. More restrictive ratios, such as one-tenth of the tolerance interval, are typically applied to highly critical characteristics, whereas fractions such as one-eighth or one-quarter may be adopted when balancing metrological risk and industrial feasibility. These stricter margins aim to reduce the risk of incorrect conformity decisions [14].

In cases involving extremely tight tolerances, the strict application of a one-tenth ratio may require disproportionately expensive high-accuracy equipment. Therefore, industrial practice may intentionally adopt a one-eighth fraction as a risk-balanced compromise between measurement reliability and economic viability. In such situations, the associated risk must be explicitly assessed in accordance with the relevant decision rules defined for coordinate metrology [15].

Table 2 presents example tolerance-to-MPE decision criteria typically applied under controlled industrial risk assumptions. These criteria should be understood as structured decision guidelines rather than as absolute acceptance limits. The final selection must also consider characteristic criticality, production environment stability, and automation requirements.

Table 2. Example of tolerance- and MPE-based decision criteria derived from industrial practice and supported by standards (e.g., ISO 10360 series for Coordinate Measuring Machines (CMMs) performance verification, ISO 14253-1 for conformity decision rules, and ISO 15530-3/4 for task-specific uncertainty evaluation) and literature [1,6,15].

Geometric Tolerance Interval	Characteristic Criticality	Applied Rule *	Required MPE Range	Decision Implication
≥ 0.20 mm	Low	1/4 tolerance	$MPE \leq 25\%$ of tolerance	Broad range of CMM configurations acceptable
0.10–0.20 mm	Medium	1/8 tolerance	$MPE \leq 12.5\%$ of tolerance	Standard industrial CMMs suitable
0.05–0.10 mm	High	1/10 tolerance	$MPE \leq 10\%$ of tolerance	High-accuracy CMMs required
< 0.05 mm	Very High/Critical	1/10 (recommended) or risk-assessed alternative	$MPE \leq 5\text{--}10\%$ of tolerance	Only robust, high-performance configurations acceptable

* Note: The selection of the tolerance fraction is not determined solely by the width of the tolerance interval, but also by the criticality of the characteristic, the industrial environment, and the conformity decision rules adopted. In certain industrial contexts, less restrictive fractions may be applied provided that the associated measurement risk is formally assessed according to ISO 14253-1 [15].

It should be noted that this tolerance-based preselection does not replace a complete measurement uncertainty evaluation when extremely narrow tolerances or safety-critical features are involved. In such situations, experimental or model-based approaches for the identification of measurement uncertainty and error sources may be required [13,15].

This distinction highlights that the proposed approach is suitable for standard industrial applications, while critical or safety-related features may require additional uncertainty-based analysis.

The tolerance-to-MPE criteria presented in Table 2 are derived from established industrial practice, supported by relevant standards (e.g., ISO 10360, ISO 14253) and literature references. These values should be understood as practical engineering guidelines rather than strict normative limits.

3.1.3. Influence of the Industrial Environment

Working environmental conditions represent a determining factor in the deployment of automated CMMs. Temperature variations, vibrations, and contamination typical of manufacturing environments can significantly affect metrological performance and system availability if they are not adequately considered during the project phase [1,6].

In workshop and near-line installations, floor-transmitted vibrations from machining centers, material handling systems, or forklifts may significantly affect measurement repeatability [1]. In typical near-line automotive machining environments, vibration levels transmitted through the floor may range approximately from 0.5 to 2 m/s², depending on the proximity to machining processes and the dynamic behavior of the production system. Similarly, temperature variations in such environments are commonly within ± 2 °C to ± 5 °C over a production shift, depending on climate control conditions. These environmental variations can influence measurement repeatability and introduce systematic deviations if not properly compensated.

Passive isolation systems (e.g., pneumatic anti-vibration supports) or active damping solutions are frequently required to maintain performance within the MPE limits declared according to ISO 10360-2 [11]. The vibration spectrum and amplitude should therefore be evaluated during the project phase, particularly when the CMM is integrated close to high-dynamic machining processes [1,6].

The proposed methodology explicitly integrates the industrial environment as a decision criterion in machine selection. Regardless of the selected installation context, whether a metrology laboratory, a controlled room adjacent to the workshop, a near-line location, or an in-line production environment, each scenario imposes specific requirements. These requirements include the level of thermal compensation, vibration isolation, and protection against contamination that must be provided by the system [1,6].

Proper management of environmental conditions is essential for the design of machines with advanced compensation systems and robust structural configurations [1]. Although this may involve a higher initial investment, it typically results in improved system reliability over the long term. From an automation perspective, environmental robustness is directly related to system availability and operational continuity. A CMM integrated into an automated production system must exhibit stable and predictable behavior over time, minimizing unplanned downtime and corrective actions [3,8].

In this context, environmental characterization is not treated as a secondary constraint but as an integral element of the decision-making framework, directly influencing machine selection, validation strategy, and long-term production performance. Early consideration of these factors therefore constitutes a key element of the overall methodology.

3.2. Validation and Acceptance Framework

Once the machine and measurement system configuration has been defined, a structured validation and acceptance framework is required to verify that the CMM performs as expected under representative industrial conditions [1,6,8]. In automated environments, acceptance is therefore treated not merely as a formal compliance requirement, but as a critical step for reducing industrial and metrological risk prior to production start-up [1,6].

Figure 4 summarizes the proposed validation and acceptance framework, integrating ISO 10360 series [10–13] standardized verification procedures together with application-oriented decision criteria. When task-specific measurement uncertainty (U 's) is considered at the application level, approaches based on ISO 15530-3 [16] or ISO 15530-4 [17] may be applied depending on the availability of calibrated artifacts or simulation-based methods.

In addition to standardized geometric verification and uncertainty evaluation, application-level validation of measurement programs and probing strategies forms part of the industrial acceptance process [4,6,16].

3.2.1. Geometric Validation According to ISO 10360

Geometric validation of the CMM is performed in accordance with the ISO 10360 series [10–13], which defines standardized procedures for assessing the metrological performance of CMM. These procedures evaluate key parameters such as the maximum permissible error in length measurement, probing errors, and repeatability indicators [1,6].

Rather than considering these parameters in isolation, they are evaluated in relation to the specific industrial application. The validation results are analyzed to verify that the machine maintains sufficient margins with respect to the most demanding tolerances identified during the project phase, even under representative environmental operating conditions [8,18].




Figure 4. Validation and acceptance framework integrating ISO 10360 series [10–13] standardized tests with industrial decision rules to ensure the suitability of automated Coordinate Measuring Machines (CMMs) for operation in production environments.

3.2.2. Verification Strategies for Sensors and Additional Axes

Automated CMMs often incorporate complex sensor configurations, such as rotary probe heads, scanning systems, and additional axes, including rotary tables and other auxiliary axes [1,4]. The verification of these elements therefore constitutes an essential part of the acceptance framework [6].

The proposed methodology focuses on defining verification strategies that ensure consistency between the validated machine geometry and the behavior of sensors and additional axes. The influence of probe orientations, scanning modes, and auxiliary axis movements on system accuracy and repeatability is systematically evaluated to ensure that the complete configuration is compatible with the intended level of automation [6,12,14,19].

3.2.3. Industrial Acceptance Criteria and Decision Rules

Industrial acceptance criteria go beyond simple compliance with applicable standards. The proposed framework integrates validation results with decision rules derived from tolerance requirements, environmental conditions, and automation objectives [6,14]. Acceptance of the CMM is granted only when the verified performance enables reliable operation within the defined production context [3,5].

This approach ensures that the accepted CMM not only complies with relevant standards but is also demonstrably fit for purpose as an automated inspection system integrated into the production process [3,5,8].

3.3. Series Production Phase and Automation Integration

Following acceptance, the CMM enters the series production phase, in which long-term robustness and system availability become the predominant considerations [3,5].

The lifecycle perspective ensures that metrological validation is not limited to initial acceptance but is sustained through structured maintenance and performance monitoring strategies [2,6,8].

3.3.1. Metrological Maintenance

Metrological maintenance strategies are defined with the objective of preserving the validated performance of the measurement system over time [6,8]. These strategies include periodic verification, monitoring of key performance indicators (KPI), and the planning of controlled recalibration activities [6]. In automated environments, maintenance activities are closely coordinated with production planning to minimize their impact on system availability [3,5].

3.3.2. Robustness Under Production Conditions

During series production, the CMM must operate continuously under variable environmental conditions [1,8]. System robustness is achieved through a combination of appropriate mechanical design choices, compensation algorithms, and preventive maintenance measures [1,6]. Monitoring environmental parameters and machine behavior enables deviations to be detected and addressed proactively before measurement quality is affected [3,19].

3.3.3. Integration with Automated Production Lines and CNC Feedback

In advanced applications, the CMM is integrated into automated production lines and provides dimensional data for process feedback [3,5,8]. Figure 5 illustrates an example of the integration of an automated CMM within an automated production line, including the use of dimensional data for feedback to machining processes [2,18].




Figure 5. Integration of an automated Coordinate Measuring Machines (CMMs) within an automated production line, enabling dimensional feedback to machining processes and closed-loop control in zero-defect manufacturing strategies.

Measurement results can be used to adjust machining parameters, enabling closed-loop control strategies and supporting zero-defect manufacturing objectives [5,20].

In this context, the CMM acts as an active component of the automated manufacturing system rather than a passive inspection device, reinforcing the need for a deployment approach based on machine design criteria and structured management throughout the entire system lifecycle [2,3,8,20].

4. Industrial Case Study and Discussions

4.1. Industrial Case Study

To illustrate the practical application of the proposed methodology, a representative case study from the automotive sector is presented below. The case study is based on industrial experience in the manufacturing of machined mechanical components and has been formulated to exclude confidential information and product-specific data [3,8].

The considered component is a medium-sized machined part, manufactured through turning and milling operations and incorporating multiple functional geometric features, including holes, planar surfaces, and geometric references. From a functional perspective, the component can be classified as a typical powertrain-related mechanical part, characterized by the presence of precision bores, reference datum surfaces, and tight positional and dimensional relationships between critical features. These features are subject to tight dimensional and geometric tolerances, some of which are classified as critical due to their direct impact on assembly requirements and the functional performance of the final product [7]. Such components are representative of typical automotive machining applications where dimensional inspection is directly linked to assembly performance and process capability.

The approach adopted in this case study is based on the operational reality of a machining plant for powertrain components governed by a quality management system. The project team has the responsibility of choosing the inspection means to guarantee the product quality. The input for the selection of the CMM in a risk-based quality management logic is the Control Plan as the result of a product and process risk analysis, DFMEA and PFMEA [3,5].

As shown in Figure 6, this risk-management workflow establishes a structured linkage between Design Failure Mode and Effects Analysis (DFMEA) and Process Failure Mode and Effects Analysis (PFMEA) activities, the definition of dimensional requirements in the Control Plan, and the subsequent validation of the measuring process. This framework ensures that metrological decisions are directly aligned with product criticality and process risk considerations [3,5,6].




Figure 6. Example of quality management workflow linking Design Failure Mode and Effects Analysis (DFMEA), Process Failure Mode and Effects Analysis (PFMEA), and control plan definition to Coordinate Measuring Machines (CMMs) adequacy.

4.1.1. Dimensional Requirements and Industrial Context

Based on the dimensional requirements defined in the Control Plan, the first step in the CMM selection process is to ensure that the machine volume and architecture are compatible with the geometry of the component under study [1,6]. The spatial distribution, accessibility, and metrological characteristics of the critical features determine the required sensor types, probing strategies, and machine configuration [4,6].

One of the fundamental parameters of a CMM is its usable measurement volume. This must be sufficient to accommodate the part, the fixturing, the stylus configurations and the required head rotations (where applicable) [1].

Depending on the manufacturer, recommendations range from 20% to 100% additional travel beyond the maximum dimensions of the part in each axis. The most reliable approach is to use a virtual simulation (digital twin) of the machine and the part/fixture/probe system, a capability available in modern CMM software packages [21,22].

Regarding machine configuration, bridge-type CMMs are the most common in automotive mechanical environments due to their suitable size and strong metrological performance.

In near-line installations, floor-transmitted vibrations generated by machining centers, automated handling systems, or logistics equipment may significantly affect measurement repeatability and probing stability [1,8]. For this reason, vibration levels in the intended installation area were evaluated during the project phase. The selected CMM configuration incorporated passive anti-vibration supports and structural design features compatible with workshop environments, ensuring that dynamic disturbances remained within the limits compatible with the declared MPE for length measurement according to ISO 10360-2 [11].

The consideration of vibration behavior at the project stage contributed to ensuring stable measurement performance under real production conditions and avoided the need for corrective actions after installation [1,6,8].

4.1.2. Application of the Decision Framework in the Project Phase

The next step consists of selecting candidate machines from the approved supplier panel. In medium and large industrial organizations, this panel includes CMM configurations for which prior capability assessments and industrial performance data are available [3,8].

Applying the tolerance-to-MPE framework described in Section 3.1.2, the tightest tolerance identified in the control plan was 20 μm . Based on a one-eighth tolerance fraction, the target MPE for length measurement was therefore established below 2.5 μm [1,6].

The selection of a one-eighth tolerance fraction reflects a risk-balanced industrial decision. Although a one-tenth criterion would provide a larger metrological margin, the adoption of the 1/8 ratio was considered acceptable in this context due to the controlled near-line environment, the demonstrated stability of the manufacturing process, and the conformity decision rules applied in accordance with ISO 14253-1 [15]. This approach ensured adequate measurement reliability while maintaining economic and operational feasibility [1,6].

To illustrate the extraction of nominal values and tolerance intervals used for the tolerance-to-MPE comparison, Table 3 presents a simplified example of the control plan applied in the case study in Figure 7. The tightest tolerance interval identified in this document establishes the reference for defining the target metrological performance of the candidate CMM configurations [1,6].




Figure 7. Extraction of nominal values and tolerance intervals from the control plan, including tolerance interval calculation and decision-making based on a one-eighth tolerance fraction.

Table 3. Basic control plan example.

SIDE	Characteristic	Nominal	LTL	UTL	TI (mm)	1/8 TI (mm)
C100	Character 001	13.000	12.970	12.997	0.027	0.0034
	Character 002	0.000	0.000	0.200	0.200	0.0250
	Character 003	9.500	9.495	9.520	0.025	0.0031
	Character 004	0.000	0.000	0.200	0.200	0.0250
	Character 005	220.458	220.408	220.508	0.100	0.0125
	Character 006	9.500	9.495	9.520	0.025	0.0031
	Character 007	0.000	0.000	0.200	0.200	0.0250
	Character 008	164.940	164.897	164.997	0.100	0.0125
	Character 009	310.030	309.970	310.090	0.120	0.0150
	Character 010	12.500	12.400	12.600	0.200	0.0250
	Character 011	138.600	138.400	138.800	0.400	0.0500
	Character 012	8.700	8.500	8.900	0.400	0.0500
	Character 013	0.000	0.000	0.200	0.200	0.0250
	Character 014	0.000	0.000	0.140	0.140	0.0175
	Character 015	224.970	224.900	225.040	0.140	0.0175
	Character 016	8.000	7.994	8.019	0.025	0.0031
	Character 017	140.830	140.750	140.910	0.160	0.0200
	Character 018	176.100	176.000	176.200	0.200	0.0250
C500	Character 019	72.200	72.200	72.220	0.020	0.0025
	Character 020	0.000	0.000	0.200	0.200	0.0250
	Character 021	72.200	72.200	72.220	0.020	0.0025
	Character 022	0.000	0.000	0.200	0.200	0.0250
	Character 023	72.200	72.200	72.220	0.020	0.0025
	Character 024	0.000	0.000	0.200	0.200	0.0250
	Character 025	72.200	72.200	72.220	0.020	0.0025
	Character 026	0.000	0.000	0.200	0.200	0.0250

The main metrological characteristics of the candidate CMM configurations are summarized in Table 4. In addition to nominal MPE values, environmental robustness, thermal stability, and compatibility with the intended automation architecture were considered during the comparative evaluation [1,6,8].

Table 4. Main metrological features for three market Coordinate Measuring Machines (CMMs).

Supplier	Model	Measuring Volume (mm)	MPE E ₀ /20 °C (μm)	MPE E ₀ /40 °C (μm)	RONt (μm)
Supplier 01	Model 01	700 × 500 × 500	1.9 + L/300	2.9 + L/200	1.6
Supplier 02	Model 02	600 × 500 × 280	2.2 + 3.5L/1000	3.6 + 6.5L/100	NA
Supplier 03	Model 03	700 × 900 × 600	1.7 + L/280	NA	1.7

Note: NA = Not Available.

This structured filtering process reduced the risk of overspecification while ensuring that the selected configuration met both metrological and industrial integration requirements [1,3,6].

Based on the comparison presented in Table 4, only Model 01 and Model 03 are initially suitable for the proposed measuring tasks under standard environmental conditions (20 °C) [1,6].

However, the CMM is intended to operate in a near-line environment, which is less strictly controlled than a metrology laboratory while remaining compatible with integration into an automated production system. For this reason, MPE values under different expected working temperatures were compared with the target tolerance limits in order to verify whether the thermal stability characteristics of the candidate machines met the project requirements [1,6,8].

In this context, measurement system accessibility, environmental robustness, and thermal compensation capabilities become critical selection criteria. Machine architectures and sensor configurations incorporating enhanced structural rigidity and temperature compensation systems were therefore prioritized [1,6].

As a result of this extended evaluation, a CMM configuration compatible with automated inspection of critical characteristics and integration into the production system was selected [3,8]. In addition, execution time of measurement programs—compatible with machining cycle times—was considered a determining factor, making CMM travel speed another decisive parameter in the final decision [23].

Following the configuration selection phase, the validation and acceptance framework was implemented to confirm compliance with the predefined metrological targets [1,6].

4.1.3. Validation, Acceptance, and Start of Production

Once the system configuration was defined, the validation and acceptance framework described in Section 3.2 was applied. The validation was performed according to ISO 10360-2 [11] and ISO 10360-5 [13], confirming compliance with the predefined MPE targets established during the project phase [1,6].

The validation results confirmed that the CMM met the defined industrial acceptance criteria, maintaining adequate margins with respect to critical tolerances under representative working environment conditions [1,6]. Following acceptance, the system was integrated into the production flow, enabling automated and repeatable dimensional inspections [3,8].

During the initial production phase, additional operational indicators were monitored to support the assessment of system performance under real industrial conditions. Measurement repeatability was observed to remain within the variation range identified during the validation phase, ensuring consistency of results across repeated inspection cycles. Inspection cycle time was verified to be compatible with the machining cycle, avoiding bottlenecks in the production flow. In addition, no unplanned downtime related to metrological performance or environmental disturbances was reported during the initial implementation period, supporting the assessment of stable operation and high system availability [1,6].

A summary of the industrial case study and the main decisions and outcomes is provided in Table 5.

Table 5. Summary of the industrial case study illustrating the application of the proposed lifecycle-based methodology.

Aspect	Description
Industrial sector	Automotive manufacturing
Component type	Medium-sized machined mechanical component
Manufacturing processes	Turning and milling
Key geometric characteristics	Holes, planar surfaces, and functional geometric references
Tolerance level	Tight dimensional and geometric tolerances, including critical characteristics
Installation environment	Near-line industrial environment
Selected CMM configuration	Automation-oriented CMM with suitable architecture and sensor configuration
Validation approach	ISO 10360 series geometric validation and application-oriented acceptance criteria
Automation level	Automated loading/unloading and integration into production flow
Main outcomes	Robust operation supported by stable measurement results and high system availability

The validation process for the case study consists of several steps:

Geometry and Metrological Characteristics validation. The previous sections described how standardized criteria are applied to determine the maximum permissible errors (MPE) of a CMM and assess their adequacy with respect to the required geometric specifications and tolerance intervals [1,6]. For consistency, the same standards and criteria must be used during the acceptance phase to confirm that the selected equipment meets the project requirements [1]. Geometric verification was performed in accordance with the ISO 10360 series [10–13], including ISO 10360-2 [11] for length measurement performance, ISO 10360-3 [12] for CMMs with rotary tables, and ISO 10360-5 [13] for CMMs using probing systems, including scanning probes and multiple probe configurations. The geometric verification test was performed under controlled environmental and operational conditions in accordance with the ISO 10360 series [10–13]. The relevant test parameters, including material temperature, artifact positioning, and the applicable MPE model, are summarized in Table 6.

Table 6. Test conditions.

	Axis X	Axis Y	Axis Z	Test Artifact
Material temperature (°C)	20.00	20.00	20.00	24.07
Test artifact position (mm)	48	−1029	−155	

Note: Monitoring factor = 0.30; MPE model for length measurement defined as $E_0 = 1.7 + L/300$, according to ISO 10360-2 [11].

The measured nominal lengths, actual values, and resulting deviations are presented in Table 7. The mean deviation corresponds to the indication error E_0 , while the minimum and maximum values reflect the observed measurement dispersion at each test length.

Table 7. Measurement results.

Nominal Value (mm)	Actual Value (mm)	Mean Deviation (mm)	Minimum (mm)	Maximum (mm)
19.9583	19.9578	−0.0005	−0.0005	−0.0004
179.8421	179.8416	−0.0005	−0.0006	−0.0004
359.8433	359.8434	0.0001	−0.0001	0.0001
539.8572	539.8568	−0.0004	−0.0004	−0.0003
639.8172	639.8184	−0.0004	−0.0005	−0.0003

The graphical representation of the indication error as a function of the measured length, together with the corresponding MPE limits, is shown in Figure 8.

Measuring programs validation. Regardless of whether measurement routines are developed externally or in-house, they must undergo a structured validation [1,6]. The reliability of CMM results depends not only on the intrinsic metrological characteristics of the machine, but also on the robustness of the measurement strategy and program design [1,4,6]. Best-practice guidelines—both industrial and normative—emphasize: correct definition of datums and coordinate systems, appropriate part orientation and fixturing, probes qualification, adequate probing patterns and point distribution, and the correct selection of filtering, fitting algorithms, and evaluation methods (e.g., least squares, minimum circumscribed) [6,7,13].

Program validation was complemented by repeatability testing covering 100% of the controlled characteristics and by intercomparison of relevant characteristics of selected critical features with a reference measurement process, preferably an accredited laboratory or central metrology facility with significantly lower measurement uncertainty [14,16]. Par-

ticular attention was given to characteristics with tight tolerances, complex datum systems, safety-critical or regulatory relevance, low-capability processes, or features sensitive to tool wear or special manufacturing processes [7,14].




Figure 8. Graphical representation of the indication error as a function of the measured length. The solid line (green) represents the measured indication error (E_0), while the dashed lines indicate the upper (blue) and lower (orange) limits of the maximum permissible error (MPE) according to ISO 10360-2 [11]. The results confirm that the measured errors remain within the specified MPE limits.

Additional Validation Aspects. Additional validation aspects, such as safety systems, handling conditions, accessibility, and environmental requirements, were considered as part of the overall system assessment. The evaluation of these aspects was performed based on predefined industrial acceptance criteria, ensuring compliance with operational, safety, and integration requirements.

As an industrial asset, the CMM must also comply with regulations applicable to machinery, workplaces, and automated production systems. Additional verification included assessment of automation systems when integrated into production, acoustic and ergonomic considerations, safety systems (emergency stops and protective devices), environmental and utility requirements, compliance with specified cycle time, mechanical conformity of fixtures and handling systems [3,8,18,23]. These checks ensured full integration of the CMM into the industrial environment while maintaining operational safety and reliability [3,8].

The selection and validation of a CMM in industrial environments cannot be reduced to a purely metrological assessment. Instead, it requires the structured integration of product risk analysis, dimensional requirements, environmental constraints, and digital automation considerations [3,5,8,21,22]. Figure 9 summarizes this integrated decision framework, illustrating how technical specifications derived from DFMEA and PFMEA analyses are progressively translated into CMM requirements and industrial deployment criteria [3,5].




Figure 9. Integrated decision framework for Coordinate Measuring Machines (CMMs) deployment, connecting product risk analysis, technical requirements, environmental constraints, and digital integration in automated manufacturing environments.

To improve traceability between the decision framework and the industrial outcomes, Table 8 summarizes the key criteria, target values, validation results, and final acceptance decisions.

Table 8. Summary of decision criteria, validation results, and industrial acceptance.

Aspect	Criterion	Target Value/Requirement	Obtained Result	Decision/Outcome
Critical tolerance	Tightest tolerance (Control Plan)	20 μm	Identified in Table 3	Used as reference for CMM selection
Tolerance–MPE criterion	1/8 rule (industrial decision)	$\text{MPE} \leq 2.5 \mu\text{m}$	Model 01: $1.9 + L/300$; Model 03: $1.7 + L/280$	Both candidates initially acceptable
Environmental conditions	Near-line environment	Moderate thermal variation and vibration	Passive anti-vibration supports implemented	Suitable for industrial deployment
Thermal performance	MPE under non-ideal temperature	Maintain compliance with tolerance limits	Model 01 includes 40 °C MPE data; Model 03 limited	Model 01 preferred for robustness
Machine selection	Combined criteria (MPE + environment + accessibility)	Compatibility with automation and tolerances	Configuration selected based on robustness and integration capability	Final CMM configuration selected

Table 8. Cont.

Aspect	Criterion	Target Value/Requirement	Obtained Result	Decision/Outcome
Geometric validation	ISO 10360-2	Compliance with declared MPE	Max deviation within ± 0.0006 mm	Requirement fulfilled
Repeatability/dispersion	Measurement variation	Stable and low dispersion	Range approx. ± 0.0005 mm	Acceptable repeatability
Measurement program validation	Industrial best practices	Correct datum definition and probing strategy	Verified and validated	Approved for production
Cycle time	Compatibility with production cycle	Within machining cycle time	Confirmed during implementation	Accepted
Industrial integration	Automation, safety, and environment compliance	Full integration in production system	Achieved (loading, operation, safety checks)	System accepted
Production performance	Stability and availability	Stable operation under real conditions	No significant deviations observed	High system availability

4.2. Discussion

The industrial case study highlights the strong influence of early project decisions on both the feasibility of automation and the robustness of the measurement system. In particular, the results indicate that focusing exclusively on nominal metrological performance does not provide sufficient guarantees when CMMs are intended to operate under real production conditions [1,6,8]. From an industrial perspective, the joint consideration of machine design aspects and automation requirements supports stable dimensional inspection in production environments [3,5,8].

Compared to conventional approaches where selection, validation, and operation are treated independently, the proposed methodology introduces a structured integration of these phases within a unified lifecycle-based framework. In traditional industrial practice, CMM selection is often driven primarily by nominal metrological specifications, while validation and production performance are addressed in later stages. This fragmented approach may lead to inconsistencies between machine capability, environmental conditions, and automation requirements. In contrast, the proposed framework explicitly links tolerance-based selection criteria, environmental constraints, and validation procedures from the early project phase through to series production. This integrated perspective supports more robust decision-making and reduces the risk of performance deviations when the system is deployed under real industrial conditions.

The use of structured criteria based on tolerances and MPE values significantly reduces the risk associated with machine selection [1,6]. The comparison of different CMM configurations shows that machines with higher nominal performance are not necessarily the most suitable option for real production environments [1,6,8].

The case study illustrates that configurations offering slightly lower nominal accuracy but improved environmental robustness and automation compatibility may provide superior industrial performance when long-term system availability and operational continuity are considered [3,8]. In industrial production contexts, availability, predictability, and resistance to environmental disturbances become as critical as nominal accuracy values declared under laboratory conditions [1,6].

Furthermore, the validation phase proved to be more than a formal compliance exercise. The interpretation of results obtained according to ISO 10360 series [10–13] in

relation to the specific tolerance structure and environmental conditions was essential to ensure fitness for purpose. This confirms that standardized tests must be contextualized within the intended industrial application [1,8].

From a lifecycle perspective, the integration of environmental assessment, tolerance-based decision criteria, and structured validation procedures contributes to reducing industrial risk not only at start-up but throughout series production [3,5,8]. The methodology therefore supports the transformation of CMMs from standalone inspection devices into active and reliable components of automated manufacturing systems [3,8].

Nevertheless, certain limitations should be acknowledged. The rule-based linkage between tolerance intervals and MPE values provides a robust preliminary filtering tool but does not replace a full uncertainty analysis when extremely critical characteristics are involved [14,16]. Additionally, the methodology assumes a relatively stable industrial environment; highly dynamic or poorly controlled workshop conditions may require additional mitigation strategies [1,6].

Overall, the results indicate that a design-oriented and risk-informed deployment strategy enhances the reliability of dimensional inspection systems and supports the integration of metrology into automated production environments [3,8].

This lifecycle-based perspective aligns with current trends in digitalized manufacturing systems and digital twin-driven production strategies [21,22].

5. Conclusions and Future Developments

5.1. Main Contributions

The results obtained in this work underline the importance of addressing system robustness and availability from the early stages of the project. In this context, automated Coordinate Measuring Machines (CMMs) must be considered as integrated industrial systems rather than as standalone inspection devices.

In addition, the proposed approach reinforces the role of metrological validation as a risk-reduction tool rather than merely a formal compliance requirement, highlighting its contribution to achieving reliable industrial deployment of automated coordinate measuring systems.

A key contribution of this work lies in structuring and formalizing the integration of established engineering practices within a lifecycle-based framework for automated CMM deployment.

Rather than introducing new metrological concepts, the proposed approach focuses on the consistent integration of design, validation, and production considerations in real production environments.

5.2. Relevance for Automated Manufacturing Systems

The results of this work indicate that successful deployment of CMMs in automated manufacturing environments depends not only on nominal metrological accuracy but also on operational continuity, predictable system behavior, and effective integration with production equipment. In this context, CMMs operate as integral components of automated manufacturing systems rather than solely as inspection devices.

Machine design decisions, including architectural configuration, sensor selection, and consideration of environmental influences, are as critical as the metrological requirements themselves. The proposed methodology aligns these decisions with the requirements of industrial automation and structured risk management.

This perspective is particularly relevant in sectors such as the automotive industry, where high production rates and stringent manufacturing tolerances require measurement

systems capable of operating reliably in real industrial environments without compromising product quality or process efficiency.

Although the case study focuses on a near-line automotive application, the principles described may be extended to other manufacturing sectors. However, the proposed methodology is illustrated through a representative industrial case study, and further validation across additional scenarios would be beneficial to assess its broader applicability.

5.3. Future Directions: Artificial Intelligence, Digital Twins, and In-Line Metrology

Current trends in smart manufacturing point toward an even deeper integration of metrology within automated systems. In this context, the application of artificial intelligence (AI) techniques offers significant opportunities to optimize measurement strategies, detect deviations in system behavior, and anticipate maintenance needs through the analysis of metrological and environmental data. These techniques can be further applied during the series production phase to analyze metrological and environmental data, enabling predictive maintenance strategies and early detection of performance drift in measurement systems.

The development of digital twins [21,22] for CMMs and associated measurement systems enables the simulation of design alternatives, environmental variations, and production process changes, allowing their impact on system performance to be assessed in advance. From the perspective of the proposed lifecycle framework, digital twins can support the project phase by enabling simulation-based evaluation of alternative CMM configurations, environmental conditions, and measurement strategies prior to physical deployment.

This development reinforces the lifecycle-based design approach and supports the transition toward in-line metrology solutions that are fully integrated into the production flow. As a result, CMMs are expected to play an increasingly important role as decision-making nodes within automated manufacturing systems.

Taken together, these developments reinforce the lifecycle-based framework proposed in this work and highlight its potential extension toward data-driven and digitally integrated metrology systems.

Author Contributions: Conceptualization, E.M.R., M.S.-N. and D.G.; methodology, E.M.R., M.S.-N. and D.G.; software, E.M.R., M.S.-N. and D.G.; validation, E.M.R., M.S.-N., M.M.M. and D.G.; formal analysis, D.G.; investigation, D.G.; resources, E.M.R., M.S.-N., M.M.M. and D.G.; data curation, D.G.; writing—original draft preparation, E.M.R. and M.S.-N.; writing—review and editing, E.M.R., M.S.-N., M.M.M. and D.G.; visualization, E.M.R., M.S.-N., M.M.M. and D.G.; supervision, E.M.R. and M.S.-N.; project administration, E.M.R., M.S.-N. and M.M.M.; funding acquisition, E.M.R. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The raw data supporting the conclusions of this article will be made available by the authors on request.

Acknowledgments: The authors are grateful for the support of the Industrial Production and Manufacturing Engineering (IPME) Research Group, the Innovation and Teaching Group for Industrial Technologies in Productive Environments (TIA Plus UNED), the Master's Degree in Metrology from UNED–CEM, and HORSE Powertrain Spain (Process Engineering HTD).

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial Intelligence
CMM	Coordinate Measuring Machine
DFMEA	Design Failure Mode and Effects Analysis
ISO	International Organization for Standardization
KPI	Key Performance Indicator
MPE	Maximum Permissible Error
PFMEA	Process Failure Mode and Effects Analysis
ZDM	Zero-Defect Manufacturing

References

1. Hocken, R.J.; Pereira, P.H. *Coordinate Measuring Machines and Systems*; CRC Press: Boca Raton, FL, USA, 2012.
2. Weckenmann, A.; Estler, W.T.; Peggs, G.N.; McMurtry, D. Probing systems in dimensional metrology. *CIRP Ann.—Manuf. Technol.* **2004**, *53*, 657–684. [CrossRef]
3. Kusiak, A. Smart manufacturing must embrace big data. *Nature* **2017**, *544*, 23–25. [CrossRef] [PubMed]
4. Colledani, M.; Tolio, T.; Fischer, A.; Iung, B.; Lanza, G.; Schmitt, R.; Váncza, J. Design and management of manufacturing systems for production quality. *CIRP Ann.—Manuf. Technol.* **2014**, *63*, 773–796. [CrossRef]
5. Psarommatis, F.; Kiritsis, D. Identification of the inspection specifications for achieving zero defect manufacturing. In *IFIP Advances in Information and Communication Technology*; Springer New York LLC: New York, NY, USA, 2019; Volume 566, pp. 267–273. [CrossRef]
6. Flack, D. *CMM Measurement Strategies—Good Practice Guide No. 41*; National Physical Laboratory (NPL): Teddington, UK, 2011.
7. *ISO 1101:2017*; Geometrical Product Specifications (GPS)—Geometrical Tolerancing—Tolerances of Form, Orientation, Location and Run-Out. International Organization for Standardization: Geneva, Switzerland, 2017.
8. Majstorovic, V.D.; Durakbasa, N.; Takaya, Y.; Stojadinovic, S. Advanced Manufacturing Metrology in Context of Industry 4.0 Model. In *Proceedings of the 12th International Conference on Measurement and Quality Control—Cyber Physical Issue*; Majstorovic, V., Durakbasa, N., Eds.; IMEKOTC14 2019. Lecture Notes in Mechanical Engineering; Springer: Cham, Switzerland, 2019. [CrossRef]
9. Schwenke, H.; Schmitt, R.; Jatzkowski, P.; Warmann, C. On-the-fly calibration of linear and rotary axes of machine tools and CMMs using a tracking interferometer. *CIRP Ann.—Manuf. Technol.* **2009**, *58*, 477–480. [CrossRef]
10. *ISO 10360-1:2000*; Geometrical Product Specifications (GPS)—Acceptance and Reverification Tests for Coordinate Measuring Machines (CMM)—Part 1: Vocabulary. International Organization for Standardization: Geneva, Switzerland, 2000.
11. *ISO 10360-2:2010*; Geometrical Product Specifications (GPS)—Acceptance and Reverification Tests for Coordinate Measuring Machines (CMM)—Part 2: CMMs Used for Measuring Linear Dimensions. International Organization for Standardization: Geneva, Switzerland, 2010.
12. *ISO 10360-3:2000*; Acceptance and Reverification Tests for Coordinate Measuring Machines—Part 3: CMMs with Rotary Tables. International Organization for Standardization: Geneva, Switzerland, 2000.
13. *ISO 10360-5:2020*; Acceptance and Reverification Tests for Coordinate Measuring Machines—Part 5: CMMs Using Multiple Probing Systems. International Organization for Standardization: Geneva, Switzerland, 2020.
14. Nafi, A.; Mayer, J.R.R.; Woźniak, A. Novel CMM-based implementation of the multi-step method for the separation of machine and probe errors. *Precis. Eng.* **2011**, *35*, 318–328. [CrossRef]
15. *ISO 14253-1:2017*; Geometrical Product Specifications (GPS)—Inspection by Measurement of Workpieces and Measuring Equipment—Part 1: Decision Rules for Proving Conformance or Non-Conformance. International Organization for Standardization: Geneva, Switzerland, 2017.
16. *ISO 15530-3:2011*; Geometrical Product Specifications (GPS)—Coordinate Measuring Machines (CMM): Technique for Determining the Uncertainty of Measurement—Part 3: Use of Calibrated Workpieces or Standards. International Organization for Standardization: Geneva, Switzerland, 2011.
17. *ISO 15530-4:2008*; Geometrical Product Specifications (GPS)—Coordinate Measuring Machines (CMM): Technique for Determining the Uncertainty of Measurement—Part 4: Evaluating Task-Specific Measurement Uncertainty Using Simulation. International Organization for Standardization: Geneva, Switzerland, 2008.
18. Calvo, R.; D’Amato, R.; Gómez, E.; Domingo, R. Integration of Error Compensation of Coordinate Measuring Machine into Feature Measurement: Part I—Model Development. *Sensors* **2016**, *16*, 1610. [CrossRef] [PubMed]
19. Magdziak, M. Estimating Time of Coordinate Measurements Based on the Adopted Measurement Strategy. *Sensors* **2022**, *22*, 7310. [CrossRef] [PubMed]
20. Vacharanukul, K.; Mekid, S. In-process dimensional inspection sensors. *Measurement* **2005**, *38*, 204–218. [CrossRef]

21. Lindqvist, R.P.; Strand, D.; Nilsson, M.; Collins, V.; Torstensson, J.; Kressin, J.; Spensieri, D.; Archenti, A. 3D Model-Based Large-Volume Metrology Supporting Smart Manufacturing and Digital Twin Concepts. *Metrology* **2023**, *3*, 29–64. [CrossRef]
22. Tao, F.; Zhang, H.; Liu, A.; Nee, A.Y.C. Digital twin in industry: State-of-the-art. *IEEE Trans. Ind. Inform.* **2019**, *15*, 2405–2415. [CrossRef]
23. Lu, Y.; Liu, C.; Wang, K.; Huang, H.; Xu, X. Digital twin-driven smart manufacturing. *J. Manuf. Syst.* **2020**, *54*, 12–28. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Article

Design, Calibration, and Troubleshooting of a Modular Low-Cost 3D Printer Based on Open-Source Technologies

Mauricio Arturo Moreno-Gerena ¹, Luis Manuel Navas-Gracia ^{2,*} and Juan Gonzalo Ardila-Marín ^{2,3}

¹ SIMEX S.A.S, Carrera. 48 #48 Sur 181, Envigado 055422, Colombia; mauricio.moreno@simex.com.co

² TADRUS Research Group, Department of Agricultural and Forestry Engineering, Universidad de Valladolid, UVA Campus of Palencia, 34004 Palencia, Spain; juan.ardila@usco.edu.co

³ Program of Agricultural Engineering, Faculty of Engineering, Universidad Surcolombiana, USCO Sede Central, Pastrana Borrero Avenue, Carrera 1, Neiva 410001, Colombia

* Correspondence: luismanuel.navas@uva.es

Abstract

This paper presents the design, construction, and calibration of a modular low-cost 3D printer based on open-source technologies, developed as part of an academic research project. The printer utilises fused filament fabrication (FFF) and is built using locally available materials and components, including a T-slot aluminium frame, NEMA 23 stepper motors, and an Arduino Mega 2560 with RAMPS 1.4 control board. The system integrates Marlin firmware and CURA slicing software, enabling autonomous operation via an LCD panel and encoder interface. A detailed methodology is provided for mechanical assembly, electronic integration, firmware configuration, and calibration procedures. Special attention is given to the challenges encountered during the initial testing phase, including filament feeding issues, thermal inconsistencies, and mechanical misalignments. Solutions such as replacing inadequate components (e.g., fibreglass bushings with PTFE), adjusting spring tension, and refining firmware parameters are discussed. The results demonstrate successful printing of complex geometries after iterative calibration, validating the printer's performance and replicability. This work contributes to the democratisation of additive manufacturing by offering a replicable, open-source solution for educational and prototyping purposes. The findings are relevant to machine design, automation, and robotics communities seeking practical insights into low-cost fabrication systems.

Keywords: motion control; Human–Machine Interface (HMI); embedded firmware; parametric design; automated calibration; power electronics; modular architecture

1. Introduction

The potential of layer-based manufacturing to optimise design workflows, reduce development time and costs, and foster the emergence of enterprises based on automated technologies was already recognised by Ardila Marín in 2007 [1]. Since then, the proliferation of additive manufacturing through open-source platforms has significantly contributed to democratising access to digital fabrication and prototyping tools, particularly in low-resource environments. Nevertheless, technical and economic barriers continue to hinder its widespread adoption. Several studies have identified the high cost of commercial 3D printers as a central issue, with prices exceeding USD 500,000 for metal-based systems [2] and reaching up to USD 200,000 for soft-material printers [3], figures that effectively exclude laboratories, small enterprises, and rural communities. Although the RepRap movement has promoted low-cost solutions, critical challenges remain in the construction, calibration,

and operation of these systems. For instance, Laplume et al. (2016) report failure rates of up to 20% in open-source printers [4], while Khan et al. (2025) demonstrate that technical errors can compromise geometric accuracy by as much as 3.34% [5]. Furthermore, studies by Keller et al. (2025) [6] and Jadayel (2024) [7] show that insufficient calibration can lead to systematic errors of up to $\pm 150\ \mu\text{m}$, adversely affecting print quality in complex geometries. However, work such as that of Ahmad et al. (2020) [8] reinforced that additive manufacturing enables the accurate reproduction of complex forms, such as anatomical structures, while reducing fabrication time and eliminating intermediate processing steps, underscoring its growing relevance in technical and educational applications. In this context, there is a clear need to document replicable technical procedures that enable the construction and optimisation of low-cost modular 3D printers, enhancing their functionality and educational utility in non-specialised environments.

Since the earliest developments of open-source 3D printers, the RepRap movement has inspired numerous research initiatives aimed at democratising additive manufacturing. Anzalone et al. (2013) proposed a low-cost metal printer based on MIG welding, reducing the cost of commercial equipment to under USD 2000 [2]. Although their technical approach validated the feasibility of metal deposition, it did not address calibration procedures or educational replicability. In parallel, Wittbrodt et al. (2013) conducted a life-cycle economic analysis of a RepRap printer, demonstrating annual savings ranging from USD 300 to USD 2000 [9], yet without considering the technical challenges that may hinder the realisation of such benefits. Schuh (2024) documented the construction and calibration of an FFF printer using recycled materials, offering a pedagogical perspective that underscores the educational value of projects built around Marlin firmware and Arduino platforms [10].

In subsequent years, Gwamuri et al. (2016) emphasised the potential of open-source 3D printers in rural communities, proposing energy solutions such as solar power and plastic recycling [11]. While their focus was primarily socio-economic, they did not explore the technical design or operational calibration aspects. Laplume et al. (2016) and Petersen and Pearce (2017) delved into business models [4] and return-on-investment analyses [12], respectively, but left unexamined the technical factors that underpin printer functionality.

Geometric accuracy has been another key area of investigation. Minetola et al. (2018) assessed improvements in Prusa i3 printers using self-replicated components, achieving performance levels comparable to professional-grade equipment according to ISO IT classification [13]. Rendón-Medina et al. (2018) validated the clinical use of low-cost FDM printers, reporting a dimensional error of 1.96% [14]. However, both studies focused on geometric outcomes without addressing user experience or operational failures, such as filament jams or bed levelling issues. Rosli et al. (2018), Sood and Pradhan (2019), and Nath et al. (2020) explored metal-material printers [15], parameter optimisation for PLA [16], and FFF process optimisation under uncertainty using physics-informed models [17], respectively. These studies reinforced the need for precise calibration to enhance process reliability across conventional and non-conventional materials yet did not detail the practical challenges of assembly and configuration. Kechagias et al. (2022) provided a comprehensive review of the parameters influencing surface quality and dimensional accuracy in FFF, highlighting the complexity of their interactions and the necessity for systematic control methodologies [18]. Similarly, Erokhin et al. (2023) classified common defects in FFF printing and proposed mitigation strategies [19].

More recently, Weiss et al. (2025) developed a multimaterial printer for biofabrication, validating its precision ($\pm 5\ \mu\text{m}$) and cell viability ($>95\%$) [3], although they did not elaborate on the mechanical calibration processes. Jadayel (2024) and Keller et al. (2025) proposed computational and optical methods to improve geometric accuracy, achieving enhancements of up to 88.5% [7] and error reductions by a factor of 1.7 [6], respectively.

While their approaches are metrological and computational in nature, they do not address replicability in educational contexts or the use of accessible components.

Other studies have addressed dynamic control and automation. Weiss (2013) proposed a closed-loop control system to enhance trajectory precision [20]; Ctibor et al. (2018) improved stepper motor control in Arduino-based printers [21]; and Salazar-Serrano et al. (2018) automated optical mounts using printed components [22]. These works share a common interest in accessible and precise systems, yet do not apply their solutions directly to the geometric calibration of FFF printers. Praja et al. (2020) [23], Khona et al. (2022) [24], Schuh (2024) [10] and Khan et al. (2024) [5] documented construction and testing processes using Marlin firmware, but without in-depth technical analysis of thermal performance or evaluation of critical parameters.

Finally, studies such as those by Campana et al. (2021) and Dupont et al. (2021) have explored sustainability [25] and collaborative manufacturing [26], respectively, highlighting the environmental impact of self-replication and the potential of the “Do-It-Together” model. Although they share an interest in technological decentralisation, they do not incorporate energy systems or alternative materials such as agricultural residues. Tolentino et al. (2024) and Gibb and Pearce (2025) further reinforce the relevance of open-source hardware in education and research [27,28], yet fail to document replicable technical processes or applications in rural contexts.

In this context, this article documents the design, construction, calibration and troubleshooting of a modular low-cost 3D printer based on open-source technologies, developed within an academic project on technological development. The central hypothesis posits that it is possible to build a functional and replicable additive manufacturing system using exclusively accessible components and open-source platforms, without compromising print quality in complex geometries. To validate this hypothesis, the methodological procedures followed in the project are described, including the mechanical and electronic design of the system, firmware and control software integration, calibration of critical printing parameters, and the identification and resolution of recurrent faults during initial testing. The results obtained allow for an evaluation of the system’s performance and its potential as an educational and prototyping tool in resource-limited environments.

This work offers a comprehensive engineering development study. The primary contribution resides in the formalisation of practical engineering procedures and replicable workflows, aligning with fundamental benchmarks that prioritise technical documentation. Our approach is inspired by the seminal works of Anzalone et al. (2013) [2] and Hallgren and Wiberg (2015) [29], who validated that providing exhaustive documentation for “derivative” open-source designs is essential for technological democratisation and user-driven innovation. Similar to the case studies by Kayfi et al. (2015) [30] and Velu et al. (2024) [31], this paper serves as an educational reference and a practical implementation guide for integrating OSH components into functional R&D environments. We further integrate the perspective of Dunn et al. (2023) [32] regarding the viability of self-produced, complex hardware, demonstrating how community-driven modifications strengthen design stability. Furthermore, following the empirical findings of Omer et al. (2024) [33], we address the critical need for context-specific design modifications to ensure global replicability in resource-constrained settings, where “minor adjustments” are often the key to functional success. By aligning with these practical precedents, our research provides the necessary technical infrastructure to transform a modular design into a real-world tool.

2. Materials and Methods

The 3D printer developed in this project was conceived using a modular approach, which enabled its design to be divided into mechanical, electronic, and control subsystems.

This strategy not only facilitated the construction and assembly process, but also enhances its maintenance, scalability, and replicability in educational or low-budget contexts [34,35].

The mechanical design was developed using Autodesk Inventor 2014® software (Autodesk Inc., San Rafael, CA, USA), allowing for the 3D modelling of each component, the generation of technical drawings, and the anticipation of interferences or adjustments prior to fabrication. Figure 1 provides an annotated overview of the mechanical architecture of the system. Figure 1a highlights the structural layout of the T-slot aluminium frame, including the main junction points and alignment features that ensure geometric stability during assembly. Figure 1b shows the complete CAD assembly with labelled subsystems—X, Y, and Z carriages, extruder module, build platform, and electronics tray—to support clear identification of functional components. These annotations were added to enhance readability and to facilitate replication of the mechanical configuration by external users. This CAD-based methodology is widely employed in rapid prototyping and digital manufacturing projects [36–38].




Figure 1. Modular architecture of the 3D printer: (a) annotated T-slot aluminium frame highlighting structural joints and alignment elements; (b) complete CAD assembly with labelled subsystems, including axis carriages, extruder module, electronics tray, and build platform.

The printer features a usable print volume of $300 \times 300 \times 200$ mm, a dimension selected for its versatility in producing medium-sized functional parts. This working area falls within the typical range of desktop FFF printers aimed at educational and prototyping applications [5,39,40]. The assembled machine has overall external dimensions of approximately $502 \times 686 \times 571$ mm (width $\times$ depth $\times$ height).

From its initial conception, the system was designed to be low-cost, employing materials and components readily available in the regional market. The total material cost of the complete prototype was approximately USD 350–450, depending on the local availability of components such as T-slot aluminium profiles for the frame, NEMA 23 stepper motors (standard 23HS-series model supplied by Wantai Motor Co., Ltd., Beijing, China) and the ATX power supply (Shenzhen XinLong Electronics Co., Ltd., Shenzhen, China), and an electronics setup based on the Arduino Mega 2560 (Arduino S.r.l., Turin, Italy), paired with the RAMPS 1.4 board (Geeetech, Shenzhen, China) and A4988 drivers (Pololu Corporation, Las Vegas, NV, USA). This hardware combination is widely adopted within the RepRap community and has proven to be both reliable and economical [10,21,41,42].

Regarding software, open-source tools commonly used in the 3D printing community were selected. The slicing software, Ultimaker Cura 14.07 (Ultimaker B.V., Utrecht, The Netherlands) was chosen for its user-friendly interface and compatibility across multiple

platforms, while the Marlin firmware was configured to manage critical printing parameters such as steps per unit, travel limits, speeds, and sensor inputs [23,43–45].

For the fabrication of components, accessible processes such as laser cutting, conventional milling, and turning were employed. This choice was driven by the local availability of services and the need to keep costs low without compromising the system's functionality. The firmware folder, design files and supporting documentation used in this study are available in the public repository cited in the Data Availability Statement.

2.1. Mechanical Components

The mechanical system of the 3D printer was designed with a modular and functional approach, prioritising ease of assembly, maintenance, and replicability. Each subsystem was conceived to fulfil a specific role within the fused filament fabrication (FFF) process, ensuring precision, stability, and compatibility with low-cost components.

2.1.1. Frame

The main frame is constructed using T-slot aluminium profiles, which are widely used in automation and robotics applications due to their versatility and ease of assembly [41,46–48]. This type of profile allows for precise adjustments in component positioning and facilitates system expansion or modification.

The structure consists of a lower and upper frame, joined by perforated corner cubes and L-brackets, providing rigidity and stability to the assembly. Levelling feet are integrated into the base to ensure proper alignment of the print plane. This modular design enables the movement of the X, Y, and Z axis carriages with precision along linear guides. The designed structure is shown in Figures 1a and 2d.

2.1.2. X, Y, and Z Axis Carriages

The carriages are designed to move along hardened steel linear guides, driven by $M10 \times 1.5$ threaded rods and NEMA 23 stepper motors (Wantai Motor Co., Ltd., Beijing, China). This configuration enables precise positioning across the three Cartesian axes, a critical feature in FFF systems [23,40,47,49]. Each carriage comprises side mounting plates, guiding blocks, and nuts coupled to the threaded rods. The design includes flexible couplings between the motors and rods to compensate for potential misalignments, thereby reducing mechanical wear and enhancing system durability. The CAD models of the Y and Z carriages are shown in Figure 2a with labelled linear guides, carriage blocks, lead screws, flexible couplings, and motor mounts, allowing readers to clearly identify the elements responsible for multi-axis motion and alignment.




Figure 2. Cont.




Figure 2. Annotated mechanical subsystems of the 3D printer: (a) Y–Z axis carriages with labelled linear guides, lead screws, and couplings; (b) filament feeding mechanism showing drive gear, idler, and tension system; (c) extruder assembly with identified thermal components; (d) modular frame and control interface with marked modules.

2.1.3. Filament Feeding System

The filament feeding system consists of a stepper motor that drives the filament through a gear and pulley assembly with spring-loaded tension adjustment. This mechanism ensures consistent filament feeding to the extruder, minimising the risk of clogs or slippage. A CAD representation of the system is shown in Figure 2b, where the filament feeding mechanism includes annotations identifying the drive gear, idler pulley, tension adjustment system, filament guiding plate, support plate, feeder mounting plate, spacer washer. These labels clarify how filament is constrained, driven, and stabilised before entering the hotend.

2.1.4. Extruder

The extruder is based on the Budaschnozzle 1.3 model, with custom modifications to enhance its performance. It comprises a set of heat sinks, a heating block with a resistor and thermistor, and a 0.5 mm nozzle. The filament is guided through a PTFE tube, which acts as a thermal insulator and reduces internal friction. A schematic 3D model of the extruder is presented in Figure 2c with detailed labels indicating the heatsink front plate, cooling discs (heatsink fins), central mounting shaft, PTFE liner, thermal break tube, hotend mounting block, heater block, and nozzle. The design follows the open-source Budaschnozzle 1.3 specification, originally created by LulzBot/Aleph Objects, Inc. (Loveland, CO, USA), but the assembly used in this work was self-fabricated and modified.

2.1.5. Control Panel and Component Mounting

The control panel enables autonomous operation of the printer via an LCD screen, a rotary encoder, and a reset button. It is mounted on an adjustable stainless-steel plate. The electronic system (comprising the Arduino Mega 2560–Arduino S.r.l., Turin, Italy, RAMPS 1.4 board–Geetech, Shenzhen, China, and A4988 drivers–Pololu Corporation, Las Vegas, NV, USA) is housed in a ventilated metal tray, ensuring adequate heat dissipation. A front view of the structure, including the control panel layout, is shown in Figure 2d with labels highlighting the LCD interface, rotary encoder, reset button, and mounting plate. Structural frame elements already described in Figure 1 are omitted here to maintain focus on user-interaction and control features relevant for standalone operation.

2.2. Electronic Components

The electronic system of the 3D printer was designed based on open-source hardware platforms, which allowed for cost reduction, simplified programming, and ensured compatibility with a wide range of control software. The architecture comprises a controller board, an expansion shield, stepper motor drivers, position sensors, and a human-machine interface system.

2.2.1. Controller Board: Arduino Mega 2560

The main board used is an Arduino Mega 2560, based on the ATmega2560 microcontroller (Microchip Technology Inc., Chandler, AZ, USA). Figure 3a illustrates the board, highlighting its primary communication ports and pin groups. This view emphasises the interfaces most relevant to the printer's control architecture, including USB communication, power input, analogue inputs and PWM-enabled digital outputs, which serve as the foundation for signal routing within the RAMPS 1.4 shield. This platform is widely adopted in 3D printing projects due to its processing capacity, number of input/output pins, and compatibility with firmware such as Marlin [38,50–52].




Figure 3. Annotated electronic components of the control system: (a) Arduino Mega 2560 with labelled USB port, power jack, microcontroller, ICSP header, analogue inputs, and digital PWM pins; (b) RAMPS 1.4 shield showing power input terminals, stepper-driver sockets, heater outputs, fuse holders, and endstop connectors; (c) A4988 stepper driver with microstepping pins, current adjustment potentiometer, logic-signal pins, and heatsink; (d) NEMA 23 stepper motor with annotated shaft, coil wires, mounting holes, and specification label; (e) mechanical endstop switch with labelled signal pins, actuator lever, and mounting holes.

2.2.2. Expansion Shield: RAMPS 1.4

The Arduino board (Arduino S.r.l., Turin, Italy) is complemented by a RAMPS 1.4 shield (Geetech, Shenzhen, China), which follows the open-source RepRap Arduino Mega Pololu Shield specification, and acts as an interface between the microcontroller and the power components. Figure 3b annotated to identify its principal connection termi-

nals and driver sockets. This view clarifies the power distribution layout, heater outputs, fuse locations and endstop connectors, ensuring that readers can accurately replicate the wiring scheme and the integration of the motion-control subsystems. This shield enables the connection of motor drivers, temperature sensors, endstops, and other peripherals. Its modular design and widespread adoption within the RepRap community make it a standard choice for DIY (Do It Yourself) 3D printers [10,41,53–55].

2.2.3. Motor Drivers: A4988 Pololu

To control the stepper motors, A4988 drivers (Pololu Corporation, Las Vegas, NV, USA) were used. These drivers support microstepping up to 1/16 step, feature adjustable current control, and include overcurrent and overtemperature protection. As shown in Figure 3c with labelled functional elements, including microstepping configuration pins, the current-adjustment potentiometer and logic-signal terminals. This information is essential for correctly tuning torque, motion smoothness and thermal performance during axis and extruder operation. They are widely used in 3D printers due to their low cost and ease of integration [21,41,56,57].

2.2.4. Stepper Motors

To ensure precise and reliable movement of the printer's Cartesian axes, the required torque for the stepper motors was calculated, considering load conditions, friction, mechanical efficiency, and the characteristics of the leadscrew transmission system. The total required torque comprises two main components: load torque and acceleration torque. The expression used for the calculation is presented in Shigley et al. (2004) [58], is shown in Equation (1).

$$M = F \left(\frac{h}{2\pi\eta} + r_B \mu_B \right) \quad (1)$$

where M : Total required torque [N·cm]; F : Total axial force [N]; h : Leadscrew pitch [cm]; η : Transmission system efficiency; r_B : Mean bearing radius [cm]; and μ_B : Bearing friction coefficient. For the X-axis, with a load mass of approximately 1.5 kg, a leadscrew pitch of 1.5 mm, and typical friction coefficients for steel–bronze contact, the required torque was calculated to be approximately 2.70 N·cm. This value was obtained without considering gear reduction or preload, representing a conservative estimate. A NEMA 23 stepper motor was selected, with the following technical specifications: Step angle: 1.8°; Steps per revolution: 200; Rated current: 2.8 A; Holding torque: 0.9 N·m (90 N·cm); Inductance: 1.2 mH; Resistance: 0.8 Ω ; Shaft diameter: 6.35 mm; and Weight: 0.6 kg. This motor was chosen with a generous safety margin, enabling not only efficient axis movement but also the potential integration of additional tools such as cutting heads or dual extruders in future prototype versions.

Moreover, these bipolar NEMA 23 motors (Wantai Motor Co., Ltd., Beijing, China) were selected for their torque and compatibility with A4988 drivers (Pololu Corporation, Las Vegas, NV, USA). Figure 3d shows the external features of the stepper motor, indicating the shaft, mounting interface and coil-wire connections. These annotations support correct mechanical alignment and electrical phasing, which are critical for achieving stable and repeatable motion across the X, Y and Z axes. These motors allow for precise positioning control of the X, Y, and Z axes, as well as the filament feeding system [47,49,59,60]. Their operation is based on the switching of internal coils via control signals generated by the firmware.

2.2.5. Position Sensors: Endstops

Each axis is equipped with two endstop switches (generic KW11-type, sourced from Geeetech, Shenzhen, China) minimum and maximum, as shown in Figure 3e, which define

the travel limits and enable the homing operation. The Figure presents the mechanical endstop, with labels identifying the actuator lever and the signal–VCC–GND pin arrangement. This view assists users in correctly orienting and wiring the homing sensors, thereby ensuring reliable axis referencing during system calibration. These mechanical sensors are connected to the RAMPS 1.4 board (Geeetech, Shenzhen, China) and are managed by the firmware to prevent collisions and ensure system repeatability.

2.3. Firmware and Control Software

The printer’s control system is based on open-source tools widely adopted by the 3D printing community, ensuring flexibility, continuous updates, and compatibility with open hardware platforms.

2.3.1. Firmware: Marlin

The firmware selected for this project is Marlin (MarlinFirmware Project, GitHub Open-Source Community, USA), a robust and optimised solution for Arduino-based 3D printers [6,10,23,43]. Marlin enables the management of axis kinematics, extrusion, temperature control, and communication with slicing software. Its most notable features include: PID control for hotend and heated bed temperature regulation; Support for SD cards and LCD displays; Flexible configuration of steps per unit, travel limits, and speeds; and Algorithms for linear acceleration and curve compensation. Figure 4a illustrates the initial section of the Configuration.h file within the Arduino IDE, highlighting the firmware metadata, electronics board selection and baud-rate definition. These parameters form the foundation of the Marlin firmware configuration, and their correct setup is essential before adjusting motion parameters, endstop logic or travel limits in the subsequent sections of the file.

2.3.2. Slicing Software: CURA

In this study, Ultimaker Cura version 14.07 (released in 2014) was used to generate the G-code, converting 3D models into printer instructions. CURA is an open-source slicing software developed and maintained by Ultimaker B.V. (The Netherlands), widely adopted in the FFF 3D-printing community for its cross-platform compatibility and ease of configuration. CURA allows the configuration of critical parameters such as layer height, print speed, infill percentage and pattern, extruder and bed temperature, and support generation and auxiliary structures. Its intuitive interface and compatibility with multiple operating systems make it a standard tool for FFF printers [44,53]. Figure 4b shows the CURA Preferences window, highlighting the machine-dimension fields (build volume), extruder count, filament-density and cost parameters, and the serial-port and baud-rate communication settings. These configurations must remain consistent with the firmware settings to ensure correct printer communication and material handling. Figure 4c presents the JOG control window in CURA, showing the directional movement controls for the X, Y and Z axes, as well as the manual plastic extrusion and retraction buttons. The Connect, Print and Pause controls and the printing statistics panel are also visible. These elements allow users to verify axis response, movement direction and extrusion behaviour during initial setup, and homing.

2.3.3. Communication and Calibration

Communication between the firmware and software is established via USB connection using a serial protocol. The initial setup includes Serial port and baud rate configuration, Calibration of steps per unit for each axis, Adjustment of travel limits and maximum speeds, and Validation of endstop logic. These parameters are modified in the Configuration.h file of the Marlin firmware and verified through testing in the CURA environment and complementary tools such as Pronterface [61–63]. Figure 4d shows the Pronterface interface

used during calibration and diagnostic procedures. The annotated elements include the axis-movement controls, feed rate settings, temperature controls for the hotend and heated bed, manual extrusion buttons, printer-connection controls, and the G-code terminal for real-time communication with the firmware. These features enable precise verification of motion behaviour, temperature response and extrusion performance. The programmed extrusion length was then used to calibrate the feed rate and adjust the firmware steps-per-unit accordingly.




Figure 4. Annotated software interfaces supporting configuration and calibration: (a) Arduino IDE showing key firmware parameters (metadata, motherboard selection, baud rate); (b) Cura Preferences panel with build-volume fields, extruder count, filament settings, and communication parameters; (c) Cura JOG window with axis-movement controls and extrusion/retraction buttons; (d) Pronterface interface with jog controls, feed rate selection, temperature settings, manual extrusion panel, and G-code console.

Table 1 summarises the main technical specifications of the developed 3D printer, including mechanical components, electronics, firmware configuration, and operational parameters. This consolidated view aims to support reproducibility and complements the detailed subsystem descriptions provided in Sections 2.1–2.3.

Table 1. Technical specifications of the developed modular 3D printer.

Overall dimensions	~502 × 686 × 571 mm (W × D × H)
Usable printing volume	300 × 300 × 200 mm
Frame	T-slot aluminium profiles (modular assembly)
Axes	Cartesian architecture: X, Y, Z

Table 1. *Cont.*

Linear guides	Hardened steel rods with carriage blocks (X, Y, Z)
Transmission	M10 × 1.5 threaded rods + flexible couplings
Motors	NEMA 23 stepper motors (1.8°, 90 N·cm holding torque, 2.8 A)
Extruder model	Modified Budaschnozzle 1.3
Hotend components	Heatsink, thermal break, PTFE liner, heater block, 0.5 mm nozzle
Feeding system	Drive gear + idler + spring-loaded tensioner
Electronics	Arduino Mega 2560 + RAMPS 1.4 + A4988 drivers
Power supply	ATX 12 V, 750 W
Sensors	Mechanical endstops (NC), thermistor
Firmware	Marlin (steps/mm, travel limits, PID tuning configured)
Slicing software	CURA (layer height, speed, infill, temperature configured)
Communication	USB serial protocol @ 115200 baud
Bed surface	5 mm glass plate
Validated temperatures	70 °C warm-up/190 °C PLA printing
Measured motion accuracy	±0.2 mm (axes), ±0.1 mm (homing)
Motion parameters	50 mm/s print speed, 3000 mm/s ² acceleration, jerk setting of 20 mm/s
Post-calibration dimensional error	≤0.5%
Cost	USD 350–450

2.4. Assembly Procedure

The assembly of the 3D printer followed a logical and modular sequence, aimed at ensuring the correct integration of mechanical and electronic subsystems. The process began with the construction of the main frame using T-slot aluminium profiles, which facilitated the alignment and positioning of components. The Y and Z axis carriages were then installed, followed by the X-axis and the base plate. Once the mechanical structure was complete, the filament feeding system and extruder were mounted, along with the material spool holder. Finally, the electronic components (including the control board, drivers, and user interface) were integrated into the system.

The modular nature of the design significantly simplified the assembly process, allowing for independent installation and testing of each subsystem. This approach not only reduced assembly time but also enhanced accessibility for future maintenance and potential upgrades.

2.5. Adjustments and Calibration

Calibration is a critical stage to ensure dimensional accuracy and surface quality of the printed parts. This process involves mechanical, electrical, and software adjustments to synchronise axis movement, material extrusion, and thermal control.

2.5.1. Axis and Endstop Adjustment

The movement direction of the X, Y, and Z axes was verified using manual commands in the control software (see Figure 4c). In cases of reversed motion, the corresponding variables were modified in the Marlin firmware. Subsequently, the endstops were adjusted to ensure that each axis correctly identifies its home position.

2.5.2. Print Bed Levelling

The print bed was levelled using a 5 mm glass plate placed on the base platform. The extruder height was adjusted to maintain a uniform gap of 0.1 mm at all four corners, ensuring a consistent first layer and proper adhesion [18,64].

2.5.3. Extruder and Filament Feed Calibration

The Pronterface software (Printrun package, Kliment Yanev and contributors, open-source community, USA) was used to calibrate the amount of extruded filament (see Figure 4d) [65–67]. The procedure involved commanding the extrusion of a specific length (e.g., 100 mm) and measuring the actual filament displacement. In case of discrepancies, the steps per unit were adjusted in the firmware using a simple proportional calculation. Real measurements were taken using a calliper. The procedure consisted of sending a movement command from the interface software (CURA or Pronterface), measuring the actual physical displacement, and adjusting the steps per unit accordingly. The formula applied is shown in Equation (2).

$$\text{New value} = \frac{\text{Actual value} \times \text{Expected displacement}}{\text{Actual displacement}} \quad (2)$$

This calibration method follows the standard RepRap procedure, which relies on proportional correction of the extrusion steps-per-unit based on measured filament displacement. This approach is widely adopted in open-source FFF printers because the extrusion mechanism behaves approximately linearly over the range of operational loads and temperatures.

Additionally, mechanical adjustments were made to the feeding system to improve filament tension. Figure 5a illustrates the mechanical intervention performed on the filament-feeding support block, showing the installed tension springs, compression screws and lock nuts, along with the original spring that was replaced due to insufficient stiffness. This annotated view clarifies the specific components involved in the pressure-regulation mechanism, thereby improving the transparency and reproducibility of the corrective action taken to eliminate filament slippage. The calibration procedure is detailed in Section 3.4.3, where the results of the adjustments are presented.




Figure 5. Mechanical adjustment of the feeding system and dimensional-verification test piece: (a) annotated feeding-system support block showing the installed tension springs, compression screws, lock nuts and the original spring replaced due to insufficient stiffness; (b) calibration test piece with nominal dimensions (25 × 4 mm) and indicated measurement areas for dimensional and layer-uniformity verification.

2.5.4. Temperature and Print Parameter Adjustment

The thermal system was tested by setting the hotend temperature to 70 °C and then to 190 °C for PLA printing. Thermistor readings were compared with external measurements using a thermocouple to validate thermal stability. Subsequently, parameters such as layer

height, print speed, and infill percentage were configured in the CURA software (Ultimaker B.V., Utrecht, The Netherlands). To verify dimensional calibration, a test piece was printed. Figure 5b presents the calibration test piece used for dimensional verification following the mechanical and firmware adjustments. The annotated nominal dimensions (25×4 mm) and the highlighted measurement areas support the evaluation of dimensional accuracy and layer uniformity, ensuring that the validation procedure can be reliably reproduced by other users.

2.6. Validation Approach

The validation strategy adopted in this work focuses on internal functional verification of the developed prototype rather than comparative benchmarking against commercial or open-source printers. This decision aligns with the scope of engineering development studies, in which the objective is to assess the operability, calibration consistency, and troubleshooting processes of the newly developed system.

2.7. Use of Generative AI

During the preparation of this manuscript, the authors used Microsoft Copilot (GPT-4 architecture, Microsoft Corporation, Redmond, WA, USA) to assist in the drafting and refinement of specific sections. Additionally, some images were refined using AI-powered inpainting tools from OpenArt (OpenArt AI, San Francisco, CA, USA) to enhance resolution and legibility. The authors have reviewed and edited the output and take full responsibility for the content of this publication.

3. Results

3.1. Assembly and Connection of Electronic Components

The control system of the 3D printer is based on a modular open-hardware architecture, integrating an Arduino Mega 2560 board, a RAMPS 1.4 expansion shield, Pololu A4988 stepper motor drivers, endstop sensors, a thermistor, a heating resistor, an LCD control panel with rotary encoder, and a 12 V ATX power supply. The electrical design was structured to facilitate assembly, maintenance, and future upgrades. Figure 6 shows the complete electrical wiring diagram of the 3D printer, detailing the connections between the Printed Base Board, the stepper motors for the X, Y and Z axes, the extruder heater, the thermocouple input, the heated-bed circuit and the main 11A/5A power rails. This annotated schematic clarifies the organisation of the power and signal lines and provides a reproducible reference for the assembly and integration of the electronic subsystem.

Regarding the main connections, the stepper motors for the X, Y, Z axes and the extruder are connected to the dedicated outputs of the RAMPS 1.4 shield, which manage PWM signals. The endstop sensors are connected to the X-min, Y-min, and Z-min inputs, using a normally closed (NC) configuration for enhanced safety. The extruder thermistor is connected to the analogue input T0, while the heating resistor is connected to the digital output D10. The cooling fan is connected to D9, also configured as a PWM output. The LCD control panel and encoder are connected via the EXP1 and EXP2 ports, using SPI communication and digital signals. The SD card is accessed through the SDRAMPS module, and the ATX power supply is connected directly to the VMOT and GND terminals of the RAMPS, providing 12 V DC.

The electronic components were mounted on a 1.5 mm thick stainless-steel tray, specifically designed to house the power supply, Arduino Mega 2560 board, RAMPS 1.4 shield, Pololu A4988 drivers, cooling fan, and input/output connectors. This tray was fixed to the main structure using M6 screws and T-slot nuts, allowing for an organised

and accessible layout for maintenance. The power supply used was a 12 V, 750 W ATX unit, modified for standalone operation by bridging the PS-ON and GND pins on the Molex connector. Four output lines were selected (two 12 V and two GND) to power the RAMPS 1.4 shield directly, using tubular terminals to ensure a secure and stable connection. Figure 7a shows the ATX 12 V power supply mounted on the stainless-steel electronics tray, illustrating its physical orientation, ventilation outlet and AC-input interface. This view clarifies the mechanical integration of the power unit and its role as the primary source for the 11 A and 5 A power rails used in the RAMPS board.




Figure 6. Electrical wiring diagram of the control and power subsystems. Annotated schematic showing the Printed Base Board, stepper-motor connectors, extruder heater and thermocouple inputs, heated-bed circuit, auxiliary connector J1, and the main 11A/5A power-supply rails.

The Arduino Mega 2560 board was fixed to the tray using metal spacers and M2.5 screws, onto which the RAMPS 1.4 shield was mounted. The A4988 drivers were installed in their respective sockets, ensuring correct orientation of control and power pins. Figure 7b

shows the RAMPS 1.4 shield mounted on the Arduino Mega 2560, together with the installed A4988 driver modules and the power-input terminal block. This view clarifies the physical layout of the control electronics as assembled on the stainless-steel tray, supporting the reproducibility of the installation described in this section.

Once the boards were mounted, the NEMA 23 stepper motors were connected. The motor coils were identified using resistance measurements with a multimeter, and four-pin JUMPER terminals were soldered to facilitate connection to the RAMPS 1.4 shield. The cables were routed through cable glands and channels attached to the frame profiles, minimising interference and maintaining system order.




Figure 7. Final system assembly with highlighted components: (a) ATX 12 V power supply mounted on the stainless-steel electronics tray, showing the ventilation outlet, AC input and power switch; (b) RAMPS 1.4 shield mounted on the Arduino Mega 2560, showing the A4988 driver modules, power-input terminal block and signal headers; (c) mechanical endstop mounted on the carriage block, showing the switch, wiring, mounting screw, linear guides and lead screw; (d) overview of the fully assembled electronics and frame, showing the LCD control panel, ATX power supply, Z-axis motor, heated bed and modular T-slot structure; (e) fully assembled 3D printer in operational configuration, showing the filament spool holder.

Endstop sensors were installed on each axis (X, Y, Z) in a normally closed (NC) configuration, connected to the X-min, Y-min, and Z-min pins of the RAMPS. Shielded cables with FASTON terminals and two-pin female connectors were used. Continuity of each connection was verified with a multimeter before powering the system. Figure 7c shows the installation of a mechanical endstop on the carriage block, highlighting the switch, wiring, mounting screw and its alignment with the linear guides and lead screw.

This view supports the description of the endstop-installation procedure and clarifies the physical arrangement required for reliable axis homing.

The extruder was connected via a four-pin aerial terminal, allowing quick disassembly for maintenance. An extension was fabricated using JUMPER connectors for the thermistor and tubular terminals for the heating resistor, connected to the T0 and D10 pins, respectively.

Finally, the control panel, consisting of a 4×16 LCD screen, a rotary encoder, and a reset button, was connected. The screen was wired using jumpers for power and backlight pins, while the encoder was mounted on a universal board and connected to the RAMPS via a three-wire extension. The reset button was connected with a two-wire extension and fixed to the front panel. This set of connections enabled autonomous operation of the printer, as well as control via interface software such as CURA and Pronterface. The modularity of the design facilitates future expansions, such as the integration of multiple extruders or additional sensors. Figure 7d provides an overview of the fully assembled electronic subsystem integrated into the modular T-slot frame, showing the mounted LCD control panel, ATX power supply, Z-axis motors and the overall spatial arrangement of the wiring and mechanical interfaces. This complete view complements the detailed component-level images presented previously and supports the reproducibility of the assembly process described in this section.

3.2. Final System Assembly

The final assembly, illustrated in Figures 7 and 8, involved the physical integration of all components and the validation of practical aspects related to the modular design. The use of T-slot aluminium profiles proved advantageous for initial alignment and component positioning. However, it required careful planning to insert T-nuts into the slots beforehand, as retrofitting them later would necessitate partial disassembly, an important consideration for future iterations. The installation of the Y and Z axis carriages and the integration of the X-axis with the base plate are shown in Figure 8a,b. Figure 8a shows the assembled Y–Z carriage structure, illustrating the alignment of the linear guides and the integration of the lead-screw mechanism within the modular frame. Figure 8b shows the assembly of the X-axis subsystem, including the print-bed support plate, X-axis stepper motor, linear guides and carriage block. This view supports the description of the initial stages of the final assembly process. During this phase, it became evident that verifying the smooth movement of the linear guides before securing the side plates was essential to avoid friction and misalignment. Additionally, the coupling between leadscrews and stepper motors, while functional, required fine adjustments to compensate for minor misalignments, reinforcing the importance of using flexible couplings.

Several mechanical issues were observed during the assembly process:

- Filament feeding system slippage: The spring-loaded tension mechanism initially caused gear slippage, reducing effective extrusion. This issue was resolved by replacing the original springs ($\varnothing$ 6 mm, 18 mm length) with stiffer ones, increasing compression force by approximately 30%. The improvement is shown in Figure 5a.
- Inadequate bushing in the extruder: The original fibreglass bushing in the hotend generated excessive friction, leading to clogs and inconsistent filament flow. Figure 8c shows the original fibreglass bushing removed from the hotend, where the burnt end and surface degradation are clearly visible. This damage was responsible for excessive friction and clogging during filament feeding. It was replaced with a PTFE bushing (friction coefficient ≈ 0.04), which reduced resistance by 60% and stabilised extrusion and improved extrusion consistency.

- **Nozzle leakage:** Material leakage was observed at the junction between the nozzle and the threaded barrel, affecting surface quality. The issue, illustrated in Figure 8d, was resolved by applying high-temperature silicone sealant and tightening the joint to 2.5 N·m, eliminating the leak and improving surface finish.

The electronic components were mounted on a custom-designed 1.5 mm thick stainless-steel tray, which housed the ATX power supply (12 V, 750 W), Arduino Mega 2560 board, RAMPS 1.4 shield, A4988 drivers, cooling fan, and I/O connectors. The tray was secured to the frame using M6 screws and T-slot nuts, ensuring a stable and accessible layout. Cable routing was managed using adhesive cable ducts and cable glands fixed to the frame, which helped minimise electromagnetic interference and maintain system order.

The installation of the control panel, comprising a 4×16 LCD screen, rotary encoder, and reset button, also required careful cable management to prevent interference with axis movement. The modular layout facilitated these adjustments and is recommended for future designs aiming for ease of maintenance and expandability.




Figure 8. Assembly process and some identified issues: (a) assembled Y–Z carriage showing the carriage block, linear guides, lead-screw transmission and guide supports within the modular T-slot frame; (b) X-axis assembly showing the print-bed support plate, X-axis stepper motor, linear guides and T-slot frame integration; (c) burnt fibreglass bushing removed from the hotend, showing the heat-damage region that caused excessive friction and filament clogging; (d) nozzle leakage showing molten-plastic accumulation at the interface between the nozzle and the threaded barrel, indicating insufficient sealing at the heater-block junction.

Figure 7e shows the 3D printer fully assembled in its operational configuration, including the mounted filament spool, LCD interface, Z-axis motors, heated bed and final

cable-routing arrangement. This complete view illustrates how all mechanical, electronic and control subsystems integrate within the modular T-slot frame, complementing the assembly description provided in this section. These observations highlight the importance of anticipating assembly constraints in modular designs and provide practical insights for users aiming to replicate or adapt the system in educational or low-resource settings.

3.3. Functional Testing

Once the assembly was completed, initial tests were conducted to verify the operability of the system. These tests allowed for the evaluation of axis movement, sensor response, thermal stability of the extruder, and communication with the control software. They also enabled the adjustment of critical parameters and confirmed the overall stability of the system.

3.3.1. Movement of X, Y, Z Axes

Axis movement was tested using manual commands via the control interface (see Figure 4c). The direction of movement was validated and corrected through firmware parameters. Movements of 10 mm were commanded and measured using a digital calliper, yielding deviations of less than ± 0.2 mm from the nominal value. The homing speed was configured in the firmware as HOMING_FEEDRATE {2000, 2000, 160, 0} mm/min for the X, Y, and Z axes, respectively. This adjustment ensured fast positioning without compromising accuracy. Smooth motion was achieved after minor adjustments to guide tension and leadscrew alignment. This verification is essential to ensure precision in material deposition.

3.3.2. Endstop Activation

The correct operation of the endstops was validated during the homing process. Activation logic was adjusted in the firmware (e.g., X_ENDSTOPS_INVERTING = true) to prevent collisions and ensure repeatable positioning. The tests confirmed positioning repeatability with a tolerance of less than ± 0.1 mm.

3.3.3. Extruder Temperature Control

The thermal system was tested by increasing the hotend temperature from 70 °C to 190 °C for PLA printing. Thermistor readings were compared with external measurements using a thermocouple, showing deviations of less than ± 2 °C, thereby validating the stability of the PID control. The stabilisation time at 190 °C was approximately 3 min. This control is critical to prevent extrusion failures.

3.3.4. Communication with CURA and Pronterface Software

Communication between the Marlin firmware and CURA software was verified via USB connection at 115,200 baud using a serial protocol. Correct interpretation of G-code and synchronisation with the hardware were confirmed. Additionally, Pronterface was used for specific extrusion tests and parameter adjustments.

3.4. System Calibration

Calibration was carried out to ensure dimensional accuracy and stability throughout the printing process. The procedures described in the methodology were applied iteratively, and the results are summarised below.

3.4.1. Firmware Adjustment (Steps per Unit, Travel Limits, Axis Direction)

Critical parameters were configured in the Configuration.h file of the Marlin firmware. The initial values for steps per unit were DEFAULT_AXIS_STEPS_PER_UNIT {1958, 1958, 2148, 98} corresponding to the X, Y, Z axes and the extruder, respectively.

A command was issued to extrude 100 mm of filament, and 86 mm were measured, indicating a 14% error. This calibration of steps per unit [steps/mm] is a critical stage in the setup of a 3D printer, ensuring that movements commanded from the control software correspond to the actual physical displacements of the machine. The detected error was corrected using a proportional calculation, adjusting the extruder value from 98 to 714.28 steps/mm. This procedure was repeated for each axis, including the extruder, whose calibration is essential to prevent over-extrusion or under-extrusion.

Travel limits were set as follows: X_MAX_POS = 300 mm; Y_MAX_POS = 300 mm; Z_MAX_POS = 180 mm; to prevent collisions. Axis directions were verified and adjusted using the firmware variables INVERT_X_DIR, INVERT_Y_DIR, and INVERT_Z_DIR.

3.4.2. Bed Levelling

The print bed was levelled using a 5 mm glass plate placed on the base platform (see Figure 7). The extruder height was adjusted to maintain a uniform gap of 0.10 mm at all four corners, validated using a digital calliper. This adjustment is critical to ensure proper first-layer adhesion and to prevent warping.

3.4.3. Filament Feed System Calibration

Following the previously identified 14% error, the steps per unit parameter was adjusted, and the test was repeated, yielding 99.5 mm—within a tolerance of ± 0.5 mm. Additionally, the tension of the feeding system was reinforced by replacing the springs, reducing gear slippage (see Figure 5a).

3.4.4. Thermal Stability Validation

The hotend temperature was tested at 70 °C and 190 °C for PLA printing. Thermistor readings were compared with external thermocouple measurements, showing deviations of less than ± 2 °C. The system reached thermal stability at 190 °C within 3 min, confirming the effectiveness of the PID control.

3.5. Printing Results

Following the completion of calibration, printing tests were conducted to evaluate dimensional accuracy, layer adhesion, and system stability. Results were compared before and after calibration, using both simple and complex geometries. The dimensional values reported correspond to single-print measurements intended to verify functional calibration rather than statistical variability.

3.5.1. Initial Prints Without Calibration

The first printed parts were a square-based prism measuring $20 \times 20 \times 2$ mm and a cylinder of $\varnothing 25$ mm $\times$ 3 mm. Significant deviations were observed in the cube, with actual dimensions of 21.8 mm in X, 19.2 mm in Y, and 2.05 mm in Z, representing an average error of +6.5%. Figure 9a shows the initial test prints obtained before calibration, where several defects are evident, including under-extrusion, poor first-layer adhesion, blob formation and warping. These issues reflect the lack of synchronisation between extrusion parameters, bed levelling and axis motion, and justify the calibration procedures described in the following sections. The cylinder measured 26.4 mm in diameter, corresponding to an error of +5.6%, visible in Figure 9b that shows the cylindrical test piece printed before calibration, exhibiting severe extrusion inconsistencies, surface roughness, edge deformation, and first-

layer irregularities. These defects highlight the lack of synchronisation between extrusion flow, bed levelling, and axis motion, reinforcing the need for the calibration procedures described in the following section.




Figure 9. Printing results: (a) initial prints without calibration showing under-extrusion, first-layer detachment, blob formation and warping; (b) cylindrical test piece printed before calibration, showing over-extrusion, surface roughness, edge deformation and first-layer irregularities; (c) thin-wall test piece printed after calibration, showing improved extrusion continuity and bed adhesion, with minor surface irregularities still present; (d) post-calibration motion-validation print (“M” path), showing correct axis synchronisation and continuous toolpath execution, with minor non-critical stringing.

3.5.2. Prints After Calibration

After adjusting the steps per unit and levelling the bed, the same geometries were reprinted. The cube now measured 20.1 mm in X, 20.0 mm in Y, and 2.02 mm in Z, reducing

the average error to +0.5%. Figure 9c the thin-wall test performed after the mechanical and firmware calibration procedures. The print quality improved substantially compared to the initial results, exhibiting continuous extrusion, proper bed adhesion and the absence of severe defects such as warping or material blobs. Nevertheless, minor irregularities remain visible in the infill pattern and surface finish, reflecting the non-optimised nature of this prototype and the fact that the calibration process focused on achieving functional performance rather than fine-tuned quality optimisation. Compared to the prints shown in Figure 9a, where severe under-extrusion, detachment and deformation were present, Figure 9c demonstrates a clear improvement in print consistency following the calibration process. The cylinder measured 25.1 mm in diameter, with an error of +0.4% (see Figure 5b). Surface quality improved significantly, with uniform layers and complete infill. The corrected E-steps value remained stable throughout the subsequent calibration and test-print sessions, with no additional drift observed during operation.

3.5.3. Complex Geometry Prints

Figure 9d shows a motion-validation test carried out after calibration, in which the extruder follows a sharp directional pattern (“M” shape) without missed steps or positional drift. The continuous deposition along the angled segments confirms that the X- and Y-axis movements are properly synchronised and that the step-per-unit calibration ensures stable motion. Although minor artefacts such as slight stringing remain visible, these do not compromise the assessment of motion stability, which was the primary purpose of this test. The parts exhibited good edge definition and dimensional stability, confirming the effectiveness of the calibration. The average printing time for these parts was 45 min, using the following parameters: layer height: 0.25 mm; infill: 20%; the motion parameters used during validation were 50 mm/s print speed, 3000 mm/s² acceleration and a jerk setting of 20 mm/s, consistent with typical RepRap-class configurations.

4. Discussion

The results obtained in this study confirm the initial hypothesis: it is possible to build a functional, accurate, and replicable 3D printer using exclusively accessible components and open-source technologies. The reduction in average dimensional error from over 6% to less than 0.5% following calibration, along with improvements in surface quality and thermal stability, validates the technical feasibility of the proposed system. The dimensional accuracy obtained after calibration (average error below 0.5%) is consistent with the ranges reported for other low-cost or RepRap-based FDM systems. Minetola et al. (2018) observed improvements that brought Prusa i3 printers close to ISO IT7–IT8 accuracy classes [13], while Rendón-Medina et al. (2018) reported dimensional errors of approximately 1.96% in clinical-grade low-cost printers [14]. Similarly, Khan et al. (2025) reported geometric deviations up to 3.3% prior to parameter optimisation [5]. In this context, the post-calibration behaviour of our prototype aligns with the performance envelope of comparable open-source machines, despite relying exclusively on accessible components and simple calibration procedures. The motion-validation results shown in Figure 9d further indicate that axis synchronisation and extrusion stability reach levels suitable for educational and prototyping applications, even without the advanced optical or closed-loop methods described in recent metrological studies.

From the perspective of previous studies, these findings complement and extend the approaches documented by authors such as Schuh (2024) [10] and Khan et al. (2025) [5], who explored the construction of low-cost FFF printers but did not detail systematic calibration procedures or fault-resolution strategies. Likewise, the observations of Keller et al. (2025) [6] and Jadayel (2024) [7] regarding the importance of precise calibration to

reduce systematic errors are reinforced here. However, this study demonstrates that such improvements can be achieved without resorting to advanced optical or computational systems, but rather through mechanical adjustments and accessible firmware.

The modular design, based on T-slot profiles and an Arduino–RAMPS architecture, not only facilitated assembly and maintenance but also enabled an organised integration of mechanical, electronic, and control subsystems. This modular approach aligns with the recommendations of Kechagias et al. (2022) [18] concerning the need for systematic methodologies to control critical parameters in the FFF process.

The issues identified—such as gear slippage, inadequate bushings, and nozzle leakage—were addressed through practical solutions that can be replicated by non-specialist users. These interventions demonstrate that system reliability does not depend solely on hardware, but also on the ability to identify and resolve faults through technical observation and iterative adjustments.

In the broader context of additive manufacturing, this work provides empirical evidence supporting the democratisation of prototyping technologies through open hardware. As suggested by Gwamuri et al. (2016) [11] and Dupont et al. (2021) [26], technological appropriation in low-resource communities requires not only economic access but also clear and replicable technical documentation, as presented in this article.

Future research directions could include the incorporation of feedback sensors for closed-loop control, the development of interchangeable modules for multi-material printing, and the integration of alternative energy systems such as solar power. Additionally, it is recommended to explore the use of recycled or biocomposite materials as printing feedstock, in line with the sustainable approaches proposed by Campana et al. (2021) [25].

5. Limitations

The scope of this study is strictly centred on the engineering development and functional validation of a single prototype. It does not attempt to propose new design theories or optimisation mechanisms, nor does it pursue experimental generalisation beyond the documented implementation. Future studies could extend this work through comparative performance benchmarking or advanced modelling-based optimisation, which fall outside the development-focused remit of the present contribution. Additionally, no comparative evaluation against commercial or open-source printers was conducted, as the aim of the study was functional validation of the prototype rather than performance benchmarking.

A theoretical error model or nonlinear extrusion analysis was not included, as the aim of the study was functional calibration of a development-stage prototype. The proportional method provides sufficient accuracy for verifying operability without requiring full characterisation of thermal, rheological or mechanical parameters. No statistical analysis (such as standard deviations or confidence intervals) is reported because this study did not involve multiple replicates or an experimental design aimed at quantifying variability. The objective was to verify functional operation of a development-stage prototype rather than to conduct metrological or performance-variability assessment.

6. Conclusions

This work documents the design, construction, calibration, and fault resolution of a low-cost modular 3D printer based on open-source technologies, validating the hypothesis that it is possible to develop a functional and replicable system using exclusively accessible components. The results obtained, both in axis movement tests and print quality, demonstrate that the modular approach and the use of open-source hardware and software platforms enable levels of precision suitable for educational and prototyping applications.

The methodology that was employed, which integrates CAD modelling, electronic integration, firmware configuration and iterative calibration, provides a replicable guide for similar projects in resource-constrained contexts. Furthermore, the identification and resolution of mechanical and thermal faults offer practical recommendations that may be useful for non-specialist users.

Taken together, this study contributes to the democratisation of additive manufacturing, reinforcing the value of open hardware as a tool for education, innovation, and technological appropriation. It is recommended that future research should explore the incorporation of closed-loop control, multi-material printing, and the use of recycled or biocomposite materials, thereby expanding the technical, social, and environmental impact of such developments.

Beyond its technical validation, the developed machine has potential impact across several domains. In educational settings, its low cost, modularity and open-source configuration offer a practical platform for teaching engineering concepts related to automation, CAD, electronics and control. For researchers, the system provides a fully documented and reproducible tool that can be adapted for experimental prototyping, instrumentation or process-monitoring studies. In low-resource or rural environments, the accessibility of materials and the detailed assembly procedures presented here contribute to the democratisation of digital fabrication, enabling local technological appropriation. Finally, for small enterprises and industry, the printer represents a viable entry-level solution for rapid prototyping and customised manufacturing workflows, reinforcing the relevance of open hardware in emerging innovation ecosystems.

Author Contributions: Conceptualisation, J.G.A.-M. and L.M.N.-G.; methodology, M.A.M.-G.; software, M.A.M.-G.; validation, M.A.M.-G., J.G.A.-M. and L.M.N.-G.; formal analysis, M.A.M.-G.; investigation, M.A.M.-G. and J.G.A.-M.; resources, M.A.M.-G., J.G.A.-M. and L.M.N.-G.; data curation, M.A.M.-G.; writing—original draft preparation, M.A.M.-G. and J.G.A.-M.; writing—review and editing, J.G.A.-M. and L.M.N.-G.; visualisation, M.A.M.-G.; supervision, J.G.A.-M.; project administration, L.M.N.-G.; funding acquisition, M.A.M.-G., J.G.A.-M. and L.M.N.-G. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding. The Article Processing Charge (APC) was covered by the authors.

Data Availability Statement: A public GitHub repository has been created to provide open access to the core resources associated with the developed 3D printer. The repository includes the Marlin firmware (MarlinFirmware Project, GitHub Open-Source Community, USA) folder, the design files, and the AGPL-3.0 open-source licence, and is available at: <https://github.com/mauricioamg26/Impresora-3d> (accessed on 5 February 2026).

Acknowledgments: The authors would like to express their gratitude to the Universidad de Valladolid, Universidad Surcolombiana, Instituto Tecnológico Metropolitano (ITM), and Fundación Carolina for providing the academic, technical and institutional support that enabled the development of this work. Their collaboration was essential throughout the project. During the preparation of this manuscript, the authors used Microsoft Copilot (GPT-4 architecture, Microsoft Corporation, Redmond, WA, USA) for drafting and refining specific sections of the text and AI-powered inpainting tools from OpenArt (OpenArt AI, San Francisco, CA, USA) to enhance resolution and legibility of some images. The authors have reviewed and edited the output and take full responsibility for the content of this publication.

Conflicts of Interest: Authors Mauricio Arturo Moreno-Gerena was employed by the SIMEX S.A.S. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

3D	Three-Dimensional
ATX	Advanced Technology eXtended
CAD	Computer-Aided Design
DIY	Do It Yourself
FDM	Fused Deposition Modelling
FFF	Fused Filament Fabrication
HMI	Human–Machine Interface
IDE	Integrated Development Environment
LCD	Liquid Crystal Display
NC	Normally Closed
PID	Proportional–Integral–Derivative
PLA	Polylactic Acid
PTFE	Polytetrafluoroethylene
RAMPS	RepRap Arduino Mega Pololu Shield
SD	Secure Digital
SPI	Serial Peripheral Interface
USB	Universal Serial Bus

References

1. Marín, J.G.A. *Evaluación Preliminar del Potencial de la Manufactura por Capas para la Industria Colombiana*; Universidad Nacional de Colombia: Medellín, Colombia, 2007.
2. Anzalone, G.C.; Zhang, C.; Wijnen, B.; Sanders, P.G.; Pearce, J.M. A low-cost open-source metal 3-D printer. *IEEE Access* **2013**, *1*, 803–810. [CrossRef]
3. Weiss, J.D.; Mermin-Bunnell, A.; Solberg, F.S.; Tam, T.; Rosalia, L.; Sharir, A.; Rüttsche, D.; Sinha, S.; Choi, P.S.; Shibata, M.; et al. A Low-Cost, Open-Source 3D Printer for Multimaterial and High-Throughput Direct Ink Writing of Soft and Living Materials. *Adv. Mater.* **2025**, *37*, e2414971. [CrossRef] [PubMed]
4. Laplume, A.; Anzalone, G.C.; Pearce, J.M. Open-source, self-replicating 3-D printer factory for small-business manufacturing. *Int. J. Adv. Manuf. Technol.* **2016**, *85*, 633–642. [CrossRef]
5. Khan, I.; Abas, M.; Ahmad, S.; Al Rashid, A.; Koç, M. Fabrication of a low-cost fused filament fabrication-based 3D printer and investigation of the effects of process parameters on mechanical properties of 3D-printed samples. *J. Eng. Res.* **2025**, *13*, 1936–1946. [CrossRef]
6. Keller, W.; Girard, J.; Goldberg, M.W.; Zhang, S. Precise calibration for error detection and correction in material extrusion additive manufacturing using digital fringe projection. *Meas. Sci. Technol.* **2025**, *36*, 025203. [CrossRef]
7. Jadayel, M. *Improving Geometric Accuracy of Fused Filament Fabrication (FFF) 3D Printed Parts by Computational Analysis of Digitized Sacrificial Parts*; Polytechnique Montréal: Montreal, QC, Canada, 2024. Available online: <https://publications.polymtl.ca/59432/> (accessed on 25 August 2025).
8. Ahmad, M.N.; Tarmeze, A.A.; Rasib, A.A. Capability of 3D printing technology in producing molar teeth prototype. *Int. J. Eng. Appl.* **2020**, *8*, 64–70. [CrossRef]
9. Wittbrodt, B.T.; Glover, A.G.; Laureto, J.; Anzalone, G.C.; Oppliger, D.; Irwin, J.L.; Pearce, J.M. Life-cycle economic analysis of distributed manufacturing with open-source 3-D printers. *Mechatronics* **2013**, *23*, 713–726. [CrossRef]
10. Schuh, J.F. *Building a Low-Cost 3D Printer*; Regional University of Blumenau: Blumenau, Brazil, 2024.
11. Gwamuri, J.; Poliskey, J.E.; Pearce, J.M. Open-source 3D printers: An appropriate technology for building low-cost optics labs for the developing communities. In Proceedings of the 14th Conference on Education and Training in Optics and Photonics, ETOP 2017, Hangzhou, China, 29–31 May 2017; p. 187. [CrossRef]
12. Petersen, E.E.; Pearce, J. Emergence of Home Manufacturing in the Developed World: Return on Investment for Open-Source 3-D Printers. *Technologies* **2017**, *5*, 7. [CrossRef]
13. Minetola, P.; Galati, M.; Iuliano, L.; Atzeni, E.; Salmi, A. The Use of Self-replicated Parts for Improving the Design and the Accuracy of a Low-cost 3D Printer. *Procedia CIRP* **2018**, *67*, 203–208. [CrossRef]
14. Rendón-Medina, M.A.; Andrade-Delgado, L.; Telich-Tarriba, J.E.; Fuente-Del-Campo, A.; Altami-rano-Arcos, C.A. Dimensional error in rapid prototyping with open-source software and low-cost 3D-printer. *Plast. Reconstr. Surg. Glob. Open* **2018**, *6*, e1646. [CrossRef]

15. Rosli, N.A.; Alkahari, M.R.; Ramli, F.R.; Maidin, S.; Sudin, M.N.; Subramoniam, S.; Furumoto, T. Design and development of a low-cost 3d metal printer. *J. Mech. Eng. Res. Dev.* **2018**, *41*, 47–54. [CrossRef]
16. Sood, R.; Pradhan, S.K. Design and development of a low-cost open-source 3D printer and its single response optimization using polylactic acid (PLA) material. *Mater. Today Proc.* **2019**, *27*, 2981–2991. [CrossRef]
17. Nath, P.; Olson, J.D.; Mahadevan, S.; Lee, Y.T.T. Optimization of fused filament fabrication process parameters under uncertainty to maximize part geometry accuracy. *Addit. Manuf.* **2020**, *35*, 101331. [CrossRef] [PubMed]
18. Kechagias, J.; Chaldas, D.; Vidakis, N.; Salonitis, K.; Vaxevanidis, N.M. Key parameters controlling surface quality and dimensional accuracy: A critical review of FFF process. *Mater. Manuf. Process.* **2022**, *37*, 963–984. [CrossRef]
19. Erokhin, K.S.; Naumov, S.A.; Ananikov, V.P. Defects in 3D Printing and Strategies to Enhance Quality of FFF Additive Manufacturing. A Review. *ChemRxiv* **2023**, 1–29. [CrossRef]
20. Weiss, B.M. Closed-Loop Control of a 3D Printer Gantry. *J. Chem. Inf. Model.* **2013**, *53*, 1689–1699.
21. Ctibor, J.; Pazdera, I. The New Stepper Driver for Low-Cost Arduino Based 3D Printer with Dynamic Stepper Control. In *Advances in Intelligent Systems and Computing: Mechatronics 2017*; Březina, T., Jabłoński, R., Eds.; Springer International Publishing: Cham, Switzerland, 2018; pp. 458–466. [CrossRef]
22. Salazar-Serrano, L.J.; Jiménez, G.; Torres, J.P. How to automate a kinematic mount using a 3D printed arduino-based system. *Inventions* **2018**, *3*, 39. [CrossRef]
23. Praja, P.; Kanna, R.R.; Bazeer, E.; Arunpandiyan, S. *Design and Fabrication of FDM Based Portable 3D Printer*; Anna University: Chennai, India, 2020.
24. Khona, N.B. Design and Fabrication of Arduino Based Flexible Manufacturing Process on the Desk: 3D Printing. *Int. J. Res. Appl. Sci. Eng. Technol.* **2022**, *10*, 2603–2610. [CrossRef]
25. Campana, G.; Mele, M.; Ciotti, M.; Rocchi, A. Environmental impacts of self-replicating three-dimensional printers. *Sustain. Mater. Technol.* **2021**, *30*, e00335. [CrossRef]
26. Dupont, L.; Fedoua, K.; Pearce, J.M.; Ortt, R.J. ‘Do-It-Together’: Towards the Factories of the Future. In *Cosmolocal Reader*; Ramos, J.M., Bauwens, M., Ede, S., Wong, J.G., Eds.; Futures Lab: Atlanta, GA, USA, 2021; pp. 52–59.
27. Masa, A.T.; Ramos, R.A.R.; Recto, K.H. Impacts of Open-source Hardware Movements on Electronics Engineering Educational Outcomes. *Inst. Electron. Eng. Philipp. J.* **2024**, *6*, 1–7.
28. Seidle, A.G.; Pearce, J.M. Overcoming Limitations of Proprietary Scientific Hardware Funding. *J. Knowl. Econ.* **2025**, 1–18. [CrossRef]
29. Hallgren, J.; Wiberg, W. *Open Source Hardware: A Case Study of Userdeveloped Derivatives*; The Royal Institute of Technology: Stockholm, Sweden, 2015.
30. Ragab, D.; Tutunji, T.A. Mechatronic system design project: A 3D printer case study. In *Proceedings of the 2015 IEEE Jordan Conference on Applied Electrical Engineering and Computing Technologies (AEECT)*; IEEE: New York, NY, USA, 2015; pp. 1–6.
31. Velu, C.; Petzold, V.; Reichelt, D.; Stange, M.; Ihlenfeldt, S. A case study on the use of open-source hardware in mechanical engineering. *Procedia CIRP* **2024**, *128*, 680–685. [CrossRef]
32. Dunn, K.; Feng, C.; Peek, N. Jubilee: A case study of distributed manufacturing in an open-source hardware project. *J. Open Hardw.* **2023**, *7*, 1–15. [CrossRef]
33. Omer, M.; Kaiser, M.; Buxbaum-Conradi, S.; Moritz, M.; Redlich, T. Designing for replicability: A qualitative empirical study on the replication of open-source machine tools. *Des. Sci.* **2024**, *10*, e30. [CrossRef]
34. Kantaros, A.; Diegel, O.; Piromalis, D.; Tsaramirsis, G.; Khadidos, A.O.; Khan, F.Q.; Jan, S. 3D printing: Making an innovative technology widely accessible through makerspaces and outsourced services. In *Materials Today: Proceedings*; Elsevier Ltd.: Amsterdam, The Netherlands, 2021; pp. 2712–2723. [CrossRef]
35. Doloi, S.; Das, M.; Li, Y.; Cho, Z.H.; Xiao, X.; Hanna, J.V.; Osvaldo, M.; Tat, L.N.W. Democratizing self-driving labs: Advances in low-cost 3D printing for laboratory automation. *Digit. Discov.* **2025**, *4*, 1685–1721. [CrossRef]
36. Amri, A.A.N.; Sumbodo, W. Pengembangan 3D Printer tipe core XY berbasis fused deposition modeling (FDM) menggunakan software Autodesk Inventor 2015. *J. Din. Vokasional Tek. Mesin* **2018**, *3*, 110–115. [CrossRef]
37. Kun, K. Reconstruction and development of a 3D printer using FDM technology. In *Procedia Engineering*; Elsevier Ltd.: Amsterdam, The Netherlands, 2016; pp. 203–211. [CrossRef]
38. Darsin, M.; Jatisukamto, G.; Sabri-Hussin, M. Design of Portable Cartensia 3D Printer Portable Using Arduino Mega 2560. *J. 3D Print. Addit. Manuf.* **2021**, *1*, 1–7. Available online: <https://www.researchgate.net/publication/356171717> (accessed on 25 August 2025).
39. Solarte, J.; Hoyos, A.; Mercado, F.; Rojas, Á. Development of a 3D Printer with Dual Extrusion Head for Industrial Applications. *I+T+C-Res. Technol. Sci.* **2024**, *1*, 1–11. [CrossRef]
40. Oluwajobi, A.O.; Kolawole, F.O. Design of a Fused Filament Fabrication (FFF) 3D-printer. *Niger. J. Technol.* **2021**, *40*, 252–260. [CrossRef]
41. Harsha-Vardhan, U.; Bhavana, K.; Vamsi, G.; Ashish, D. *Fused Deposition Modelling 3D Printer Using ARDUINO*; Anil Neerukonda Institute of Technology and Sciences: Bheemunipatnam, India, 2018.

42. Yavuz, G.A. *The Production of Objects with Three Dimensional Printing Technique*; Dokuz Eylül University: İzmir, Türkiye, 2014.
43. Moore, S.; Armstrong, P.; McDonald, T.; Yampolskiy, M. Vulnerability Analysis of Desktop 3D Printer Software. In *Proceedings of the Resilience Week (RWS), Chicago, IL, USA, 2016*; IEEE: New York, NY, USA, 2016; pp. 46–51. [CrossRef]
44. Karakuş, S. Design and Manufacturing of a Two-Stage Reduction Gearbox with 3D Printers. *Int. J. 3D Print. Technol. Digit. Ind.* **2023**, *7*, 18–28. [CrossRef]
45. Altan, A.; Hacıoğlu, R. The Algorithm Development and Implementation for 3D Printers based on Adaptive PID Controller. *J. Polytech.* **2018**, *21*, 559–564. [CrossRef]
46. Nguyen, T. *Study, Design and Implementation of 3D Printer*; Häme University of Applied Sciences: Hämeenlinna, Finland, 2019.
47. Breault, S.; Godbole, V.; Kent, J.; Tarnowski, M. *Five Axis 3D Printer*; University of Rhode Island: South Kingstown, Rhode Island, 2021. Available online: <https://digitalcommons.uri.edu/mechanical-engineering-capstones> (accessed on 25 August 2025).
48. Noorani, R.; Lynch, M.; Hodgkiss, P.; Fumo, M.; Burke, C. Design and Manufacturing of a 3D Printer for Teaching and Research. In *IOP Conference Series: Materials Science and Engineering*; Institute of Physics Publishing: Bristol, England, 2019. [CrossRef]
49. Sawczuk, M.G. *Design and Control of a 3D Printed, 6DoF Robot Arm*; KTH Royal Institute of Technology: Stockholm, Sweden, 2021.
50. Putra, A.K.; Fidiyanto, F.; Prakoso, B.; Armantya, R.Z.; Sandi, M.K.; Al Haris, F.H.S. Perakitan 3D printer fused deposit modeling (FDM) berbasis Arduino Mega 2560. *Gaung Inform.* **2019**, *12*, 123–133.
51. Mardiyana, D.; Sulaiman, Z.; Ihsan, S.; Ridha, F.; Rahman, T. Rancang Bangun 3D Printer FDM Model Cartesian Berbasis Arduino. *J. Mater. Dan Proses Manufaktur* **2023**, *7*, 63–72. [CrossRef]
52. Promarin, K.; Somwang, P. Thermal behavior of FDM 3D printing by using Arduino Mega 2560 approach. In *Proceedings of the 16th International Conference on Electrical Engineering/Electronics, Computer, Telecommunications and Information Technology*; IEEE: New York, NY, USA, 2019; pp. 767–770.
53. Deshpande, S.P.; Kulkarni, S.; Shah, S.; Irwin, J. Developing an open source, inexpensive, large-scale polar configuration 3D printer. *Int. J. Eng. Res. Innov.* **2019**, *11*, 13–22.
54. Alam, M.; Kader, G.; Molla, H.-O.-R.; Rokonzaman, A.S.M. Design, Construction and Performance Test of a Fused Deposition Modeling FDM 3D Printe. In *Proceedings of the International Conference on Mechanical Engineering and Renewable Energy*, Yogyakarta, Indonesia, 28–29 August 2019. Available online: <https://www.researchgate.net/publication/349324030> (accessed on 25 August 2025).
55. Kamal, M.; Essam, K.; Gamal, M.; Ali, Z.; Samir, A.; Ahmed, E.; Mohamed, A.; Hesham, R.; Zayed, N.; Fouda, K. 3D Construction Printer. In *Proceedings of the 5th International Undergraduate Research Conference*, Cairo, Egypt, 9–12 August 2021; pp. 450–452.
56. Timperley, S.; Stec, G.; Miller, D. Timperley 1 Preliminary Design of a 3D Printing System for Creating Infinitely Large Structures in Space. In *Proceedings of the 38th Annual Small Satellite Conference*, Logan, UT, USA, 3–8 August 2024; pp. 1–9.
57. Sharma, S.; Harish; Nikam, S.B. XY-Drawing Robot using Arduino. *Int. J. Innov. Res. Technol.* **2019**, *6*, 170–173.
58. Shigley, J.E.; Mischke, C.R.; Budynas, R.G. *Mechanical Engineering Design*, 7th ed.; McGraw-Hill: New York, NY, USA, 2004.
59. Popescu, D.; Florea, A.; Popescu, R.; Roibu, H. 3D Printing for the Development of Robots by the Students. In *Proceedings of the 27th EAAEIE Annual Conference*; IEEE: New York, NY, USA, 2017; pp. 1–6.
60. Ismaranatasia, W.; Setiawan, B.; Subiyantoro, S. Kendali Motor Stepper untuk Pergerakan Sumbu X, Y, Z pada 3D Printer Simetris Bilateral. *J. Elkolind* **2021**, *8*, 66–75. Available online: https://www.researchgate.net/publication/308568286_Vulnerability_analysis_of_desktop_3D_printer_software(accessed on 25 August 2025). [CrossRef]
61. Hadeif, M.W. *Conception et Realisation d'une Imprimante 3D*; Université Mohamed Khider de Biskra: Biskra, Algeria, 2019.
62. Devijver, S. Building Your Own 3D Printer an Introduction. Available online: <https://www.scribd.com/document/534946917/Building-Your-Own-3D-Printer> (accessed on 25 August 2025).
63. Honiball, J.R.; Pepper, M.S.; Prinsloo, E. Step-by-step assembly and testing of a low-cost bioprinting solution for research and educational purposes. *MethodsX* **2021**, *8*, 101186. [CrossRef]
64. Butt, J.; Bhaskar, R.; Mohaghegh, V. Analysing the effects of layer heights and line widths on FFF-printed thermoplastics. *Int. J. Adv. Manuf. Technol.* **2022**, *121*, 7383–7411. [CrossRef]
65. Evans, B. Calibrating your printer. In *Practical 3D Printers: The Science and Art of 3D Printing*; Springer: Berlin/Heidelberg, Germany, 2012; pp. 49–74.
66. Li, L.; McGuan, R.; Isaac, R.; Kavehpour, P.; Candler, R. Improving precision of material extrusion 3D printing by in-situ monitoring & predicting 3D geometric deviation using conditional adversarial networks. *Addit. Manuf.* **2021**, *38*, 101695. [CrossRef]
67. Jaksa, L.; Pahr, D.; Kronreif, G.; Lorenz, A. Calibration Dependencies and Accuracy Assessment of a Silicone Rubber 3D Printer. *Inventions* **2022**, *7*, 35. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Article

Kinematic Synthesis of Planar Leg Mechanisms Through Large-Scale Dataset Generation, Geometric Filtering, and Optimization

Ray Tang ^{1,2}, Zhijie Lyu ¹ and Anurag Purwar ^{1,*}

¹ Computer-Aided Design and Innovation Lab, Stony Brook University, Stony Brook, NY 11794, USA; ray.tang@stonybrook.edu (R.T.)

² Shoreham-Wading River High School, Shoreham, NY 11786, USA

* Correspondence: anurag.purwar@stonybrook.edu

Abstract

Walking is one of many basic human motor functions, yet replicating it in robotic systems remains a complex problem. Historically, the design of walking mechanisms has relied on human intuition and iterative refining. Some well-known mechanisms, like Theo Jansen, have been invented by artists rather than engineers. In this paper, we present a novel, automated pipeline that includes dataset generation, filtering, and an optimization procedure for synthesizing 1-DOF geometrically feasible walking mechanisms. Four million mechanisms were simulated and evaluated for 25 mechanism types, for a total of 100 million mechanisms. Quantitative design criteria for walking motion were identified and applied to retain a total of 23,250 valid, stable walking mechanisms. We then apply a custom optimization process to adjust near-walking mechanisms whose joints run into the ground. A custom function is used to minimize the error related to ground intersection and step uniformity. The computational generation and optimization of walking linkages demonstrated in this work aims to systematically generate a large number of design concepts for walking mechanisms. While the focus of this work is on the synthesis of mechanisms for walking robots, the same methodology could be generalized to identify mechanisms for a wide range of applications beyond walking robots.

Keywords: kinematics; mechanisms; robotics; simulation; computational design; datasets; optimization; kinematic synthesis

1. Introduction

This paper presents a kinematic synthesis pipeline for planar single-degree-of-freedom (1-DOF) geometrically feasible walking mechanisms. Building on simulation-based mechanism synthesis and analysis frameworks [1,2], we introduce a design-criteria-driven workflow that (i) formalizes geometric constraints specific to walking, (ii) performs large-scale simulation and dataset filtering under these constraints, and (iii) refines near-feasible candidates via optimization using a custom design-driven objective. The pipeline produces both a reusable synthesis workflow (generator) and a curated dataset of feasible 1-DOF walking mechanisms. Throughout this paper, the term “walking mechanism” refers to linkages that satisfy geometric criteria for walking feasibility. Evaluation is conducted at the kinematic level, where walking potential is assessed from the resulting foot trajectory geometry; dynamic stability and contact-force effects are beyond the scope of the present study.

Planar 1-DOF walking linkages have a rich history and continue to attract interest in robotics and mechanism design. Classic examples include Chebyshev-type leg mechanisms, Klann six-bar linkages, and Theo Jansen's multi-link walkers [3–6]. Figure 1 shows these three existing walking linkages.




Figure 1. The figure shows three existing walking linkages drawn and simulated in MotionGen [1]: (a) Jansen Linkage, (b) Klann Linkage, (c) Trotbot Linkage.

Historically, many of these designs relied heavily on intuition, manual iteration, and physical prototyping. In contrast, modern robotics and mechanism design increasingly demand computational tools that can systematically explore large design spaces while enforcing task-specific constraints (e.g., ground contact behavior, clearance, and gait-like features). Walking mechanisms provide an especially useful testbed because successful designs must satisfy multiple interacting criteria beyond reproducing a single coupler curve.

Mechanism synthesis aims to determine mechanism topology and dimensional parameters to satisfy prescribed path, motion, or function requirements, as treated extensively in classical synthesis texts [7,8]. Analytic approaches such as kinematic mapping and algebraic fitting provide unified formulations for motion synthesis and have been extended toward simultaneous type-and-dimension synthesis [9]. To make these synthesis techniques practical, interactive tools have been developed, notably MotionGen for real-time four-bar motion generation with pose and geometric constraints [2].

More recently, there has been a growing interest in combining large-scale data generation with machine learning for mechanism design. Prior work proposed learning-based synthesis of defect-free mechanisms and methods to handle uncertain user inputs, including image-based representations of input curves and invariant/normalized representations of coupler curves [10–12]. Large datasets have accelerated data-driven design, e.g., the LINKS dataset of 100M planar linkage mechanisms [13] and a curated dataset of nearly 3M single-DOF planar 4-, 6-, and 8-bar mechanisms with open/closed coupler curves [14]. In rehabilitation-oriented motion design, representative pipelines combine data generation, learned representations, and retrieval/selection to produce multiple candidate mechanisms under practical constraints [15]. Beyond Variational Auto Encoder (VAE)-style embeddings and retrieval, other learning paradigms have also been explored, including conditional GANs under constraints [16], reinforcement-learning-based synthesis [17], graph-based generative approaches for linkage graphs [18], sequence-model-based mechanism design automation [19], and conditional generative modeling applied to six-bar synthesis [20].

Frequently, mechanism design problems are reduced to finding suitable combinations of links and joints that generate a given path, such as that of a foot point during average walking motion. However, path-matching alone is often insufficient for broader robotics

design tasks requiring satisfying multiple and often conflicting design specifications. First, many applications admit families of functional mechanisms rather than a single best-fitting curve; optimizing only for a target trajectory can overlook alternatives whose coupler curves differ slightly yet still satisfy the task. Second, existing pipelines do not provide a general method to filter or curate large datasets using application-specific geometric criteria, especially when those criteria involve more than curve similarity. Third, practical design problems impose hard feasibility constraints that must be satisfied in addition to producing a plausible path. For walking mechanisms, examples include requiring a closed foot trajectory, enforcing a sufficiently flat stance phase for ground contact and propulsion, ensuring adequate step height and step length, and preventing non-foot joints or links from intersecting the ground during motion. These requirements motivate a synthesis pipeline that treats “walking-ness” as a constrained objective rather than as a byproduct of path similarity.

To overcome these limitations while leveraging the scalability of data-driven approaches, we propose a design-criteria-driven synthesis pipeline that treats “walking-ness” as multi-criteria satisfaction rather than path similarity. Unlike pure learning methods that excel at efficient search but can struggle with hard constraints, our approach combines large-scale simulation with explicit geometric filters, offering complementary advantages: (1) **guaranteed geometric feasibility** through direct constraint checking, (2) **sufficient sampling** within the constrained region, (3) **explicit criteria formulation** allowing transparent tuning for specific applications, and (4) **interpretable failure modes** when mechanisms are rejected. This approach maintains scalability while introducing application-specific intelligence through targeted filtering criteria and soft optimization of near-feasible designs. As a result, our pipeline provides a verified training dataset with explicit feasibility boundaries that can bootstrap and validate learned models.

Implementation scope: We demonstrate this pipeline on four types of planar four-bar mechanisms with revolute (R) and prismatic (P) joints—RRRR, RRRP, RRPR, and PRPR—and five six-bar mechanisms, classified into two Watt types and three Stephenson types (Figure 2). Additionally, we include different variants of these five types where coupler points are placed on different links. For details on these kinematic structures, see [21]. The resulting dataset provides not just feasible mechanisms but also the geometric criteria that define their feasibility, enabling both direct application and further learning-based refinement.




Figure 2. Five types of six-bar mechanisms in our work.

Although developed and demonstrated for walking, the methodology is intended to generalize to other robotics mechanism design tasks where feasibility and performance are defined by geometric/kinematic criteria. The main contributions of this paper are:

1. **Geometric constraint formulation for walking mechanisms:** We propose a scalable computational pipeline to identify planar walking mechanisms from randomly generated linkages by enforcing five explicit geometric criteria.

2. **Large-scale simulation and dataset filtering:** We generate a large-scale dataset by sampling linkage joint coordinates from an explicit distribution (uniform within fixed bounds) and simulating each candidate across 25 structural subtypes. Then, we evaluate approximately 10^8 candidates in total (4 M per subtype) and retain 23,250 feasible walking mechanisms after filtering. Our implementation runs efficiently on High-Performing Computing resources (e.g., ~5–6 h per 4 million candidates for a representative subtype on the SeaWulf computing cluster from Stony Brook University (Website for SeaWulf: <https://rci.stonybrook.edu/HPC/understanding-SeaWulf> (accessed on 21 January 2026))).
3. **Optimization-based refinement of near-feasible designs:** We further improve walkability via a derivative-free optimization that adjusts only ground-violating joints. We use Differential Evolution [22] with penalty-based constraint handling to minimize a composite objective that discourages invalid stance detection, ground penetration, and poor trajectory closure, and we terminate after a fixed evaluation budget or when the best objective value stalls.

The remainder of this paper is arranged as follows. Section 2 describes the dataset generation, filtering procedure, and the design-driven objective used for fine-tuning. Section 3 reports filtered dataset statistics and representative walking results, and evaluates the effectiveness of the optimization step. Section 4 concludes with a discussion of limitations and implications for broader robotics mechanism design.

2. Methodology

The methodology developed in this work builds upon an existing mechanism simulation framework by Lyu et al. [2]. The overall workflow consists of three stages: simulation (Section 2.1), dataset filtering (Section 2.2), and fine-tuning (Section 2.3) (Figure 3).

Neither random sampling nor numerical optimization alone is sufficient. In a design space with at least ten dimensions (RRRR, for example), sampling only produces suboptimal solutions. Meanwhile, nonlinear optimization initialized from random parameters frequently fails to converge, incurring substantial computational cost. We therefore combine the two: sampling provides good initial guesses, and optimization refines them to best-fit mechanisms efficiently.

Therefore, in the simulation stage, we sample large numbers of randomized mechanism parameters and perform kinematic simulation to generate candidate motions for the selected mechanism families. In the filtering stage, we evaluate a set of walking-specific geometric criteria on the simulated motions and retain only candidates that satisfy the feasibility conditions, resulting in a curated dataset of potentially walkable mechanisms. Finally, the fine-tuning stage refines near-feasible candidates by adjusting design parameters to correct geometric violations that prevent walking (e.g., non-foot joints intersecting the ground). This refinement is formulated as an optimization problem with a custom, design-driven objective constructed from the same geometric criteria together with a measure of walking effectiveness. Figure 3 provides an overview of the complete pipeline.




Figure 3. Design pipeline for walking mechanism generation, filtering, and optimization. The dashed arrows represent an iterative refinement loop: near-feasible mechanisms from filtering are optimized, with successful candidates evaluated using the filters again.

2.1. Simulation

We use the simulator developed by Lyu et al. [2] to generate kinematic trajectories from randomized initial joint configurations. Given initial joint coordinates and a specification of the mechanism type, the simulator outputs the coupler curves (trajectories) of all joints over a full motion cycle.

To support large-scale sampling, we first generate random coordinate values and store them in a matrix $\mathbf{R} \in \mathbb{R}^{10^8 \times 16}$ prior to simulation. These values are between -10 to 10 , because we standardize the simulation result with respect to the coupler path using the method in [15] so that the curve is invariant to translation, rotation, scaling, and reflection. Additionally, unlike [23], which uses the Fourier Descriptor (FD), the method in this work does not require full rotation or a fixed number of points.

Each row of $\mathbf{R}$ contains 16 scalars, interpreted as 8 planar (x, y) pairs, which is sufficient to provide randomized coordinates for all joints required by the 4-bar and 6-bar families considered in this work (five joints for 4-bar mechanisms, and seven or eight joints for 6-bar mechanisms).

$$\mathbf{R} = \begin{bmatrix} x_1^{(1)} & y_1^{(1)} & x_2^{(1)} & y_2^{(1)} & \cdots & x_8^{(1)} & y_8^{(1)} \\ x_1^{(2)} & y_1^{(2)} & x_2^{(2)} & y_2^{(2)} & \cdots & x_8^{(2)} & y_8^{(2)} \\ x_1^{(3)} & y_1^{(3)} & x_2^{(3)} & y_2^{(3)} & \cdots & x_8^{(3)} & y_8^{(3)} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ x_1^{(10^8)} & y_1^{(10^8)} & x_2^{(10^8)} & y_2^{(10^8)} & \cdots & x_8^{(10^8)} & y_8^{(10^8)} \end{bmatrix}.$$

The kinematic structure of each mechanism type is encoded by three matrices, $\mathbf{B}$, $\mathbf{S}$, and $\mathbf{I}$, following the representation in [2]. These mechanisms are collectively saved with their matrices in a dictionary (See a definition of `BSIDict` in [21]).

The binary incidence matrix $\mathbf{B}$ specifies link–joint connectivity. Formally, the $\mathbf{B}$ matrix is defined as

$$\mathbf{B} \in \{0, 1\}^{\ell \times n},$$

where n is the number of joints, ℓ is the number of link-incidence rows (including the ground row), and b_{ij} denotes the element in row i and column j .

The first row in the matrix $\mathbf{B}$ is defined as

$$\mathbf{b}_1 = [b_{11}, b_{12}, \dots, b_{1n}], \quad b_{1j} = \begin{cases} 1 & \text{if joint } j \text{ is grounded,} \\ 0 & \text{otherwise.} \end{cases}$$

while the rows 2 through ℓ are defined as For each link row $i = 2, \dots, \ell$,

$$\mathbf{b}_i = [b_{i1}, b_{i2}, \dots, b_{in}], \quad b_{ij} = \begin{cases} 1 & \text{if joint } j \text{ lies on link } i, \\ 0 & \text{otherwise.} \end{cases}$$

For mechanism types that include prismatic joints (i.e., collinearity constraints), the corresponding $\mathbf{S}$ matrix is non-empty. Each row of $\mathbf{S}$ stores a triplet of joint indices that must remain collinear. For example, if joints p , q , and r are constrained to lie on the same line, then the row $[p, q, r]$ appears in $\mathbf{S}$.

The actuator specification is stored in the $\mathbf{I}$ matrix. Similar to $\mathbf{S}$, each row of $\mathbf{I}$ contains three joint indices defining the actuation constraint, together with a fourth entry indicating actuator type. For a rotary actuator that enforces the angle between vectors $\vec{v}_{qp}$ and $\vec{v}_{qr}$ (with joint indices p , q , and r), the corresponding row is $[p, q, r, 0]$. For a linear actuator, $\vec{v}_{qr}$ defines the reference direction and $\vec{v}_{qp}$ defines the actuated length, yielding the row $[p, q, r, 1]$.

Given a set of random joint coordinates together with the matrices $\mathbf{B}$, $\mathbf{S}$, and $\mathbf{I}$ for a selected mechanism type, the simulator outputs a tensor $\mathbf{P} \in \mathbb{R}^{2 \times m \times N}$, where m is the number of joints and N is the number of simulated time steps. The first dimension (of size 2) corresponds to (x, y) coordinates. We write

$$\mathbf{P} = [\mathbf{P}^{(0)}, \mathbf{P}^{(1)}, \dots, \mathbf{P}^{(N-1)}],$$

where each

$$\mathbf{P}_i = \begin{bmatrix} x_1^{(i)} & \cdots & x_m^{(i)} \\ y_1^{(i)} & \cdots & y_m^{(i)} \end{bmatrix} \in \mathbb{R}^{2 \times m}$$

contains all joint positions at simulation step i . From this tensor, we extract the trajectory of the coupler point and denote it by

$$P = \{P^{(0)}, P^{(1)}, \dots, P^{(n)}\}, \quad P_i \in \mathbb{R}^2,$$

where P_i is the planar position of the coupler at simulation step i , and $n = N - 1$.

2.2. Dataset Filtering

After simulating each mechanism from randomized parameters, we evaluate a set of geometric criteria to determine whether the candidate should be retained in the dataset. These criteria are derived from common design considerations for planar 1-DOF walking mechanisms reported in prior design-oriented studies and articles [24–26]. Only

mechanisms that satisfy all required feasibility conditions are saved to the database for downstream analysis and refinement.

The adopted criteria are:

1. **Closed foot trajectory:** the selected coupler curve (foot point) must form a closed, periodic curve over a full actuation cycle (Section 2.2.1).
2. **No self-intersection:** the coupler curve must be a simple closed curve (i.e., it must not intersect itself) (Section 2.2.1).
3. **Bounded curvature/smoothness:** the trajectory must not contain excessively sharp corners or abrupt direction reversals; equivalently, curvature and/or the turning angle between successive samples must remain below a specified threshold (Section 2.2.2).
4. **Existence of a stance segment:** the trajectory must contain a sufficiently long “flat” portion suitable for ground contact during stance (Section 2.2.3).
5. **Stance proportion constraint:** depending on the intended robot configuration (e.g., the number of legs operating in parallel), the stance segment must occupy at least a minimum fraction of the full cycle (duty factor) (Section 2.2.3).
6. **Ground clearance for non-foot joints:** when the mechanism is oriented such that the stance segment lies on the ground line, no other joint trajectory (and, if applicable, the mechanism geometry) may intersect the ground (Section 2.2.4).

Motivation for geometric criteria: Each filtering criterion corresponds to a physical requirement for walking. Closure ensures a periodic, continuous motion for repeated stepping. Self-intersection avoidance prevents ambiguous foot positions, which could interfere with the stance portion of a step. Smoothness corresponds to foot acceleration profiles; sharp, angled portions in the coupler curve would create uneven acceleration. Flat stance segments enable ground contact and propulsion. Ground clearance prevents mechanical interference. These geometric proxies allow efficient screening before expensive dynamic simulation.

Within these algorithms, many thresholds are used to calculate the different properties of the mechanisms. They are outlined in Appendix A.

2.2.1. Closed Coupler Curve and Self-Intersection Test

To ensure that the simulated coupler curve defines a valid closed foot trajectory for walking, we perform two preliminary checks: (i) closure and (ii) self-intersection (Algorithm 1). A walking trajectory must be periodic; otherwise, the mechanism cannot sustain continuous stepping over repeated cycles, even if other criteria are satisfied. In addition, we reject self-intersecting trajectories because a self-intersection implies that the curve is not a simple closed path; this can yield ambiguous contact/stance regions and non-unique horizontal positions for the same phase, which is undesirable for stable forward walking.

For the classic four-bar linkages, analytical criteria such as Grashof’s condition can be used to determine the existence of a continuously rotating crank. However, for six-bar topologies (e.g., Watt and Stephenson types), the Grashof condition is not applicable. Since the objective of this work is to identify mechanisms that generate closed and periodic walking trajectories, closure is verified directly at the trajectory level by simulating one full input cycle and ensuring periodic return to the initial configuration within numerical tolerance. This approach guarantees closed output behavior for all types of kinematic structures.

Closure test. Let $P = \{P^{(0)}, P^{(1)}, \dots, P^{(n)}\}$ denote the sampled points of the coupler trajectory returned by the simulator. Rather than comparing only the first and last samples, we search for a return-to-start event within the trajectory, since some simulations may contain multiple cycles or partial overlaps, in which case the final sample may not coincide with a true cycle boundary. Specifically, we compute the minimum distance from the

start point P_0 to later samples P_i while skipping an initial prefix to avoid trivial near-start matches:

$$Q_c = \min_{i \geq i_{\min}} \|P^{(i)} - P^{(0)}\|_2,$$

where $i_{\min} = 40$ in our implementation (Figure 4), and $n = 512$ is the total number of samples. We declare the trajectory to be closed if $Q_c < \tau$, where $\tau = 0.1$ is a tolerance that admits small numerical gaps while still rejecting clearly open curves. Although a self-intersecting trajectory may also yield a small value of Q_c , such cases are subsequently rejected by the self-intersection test.

Algorithm 1: Coupler Curve Closure and Self-Intersection Tests

Input: $P = \{P^{(0)}, P^{(1)}, \dots, P^{(n)}\}$, closure tolerance $\tau = 0.1$
Output: booleans *is_closed*, *is_self_intersecting*

```

1 Function IsSimple( $P$ ):
2    $S \leftarrow \{\overline{P^{(i)}P^{(i+1)}} \mid i \in [0, n-1]\}$ ;
3   foreach pair  $(s_a, s_b) \in S \times S$  with no shared vertices do
4     if  $s_a \cap s_b \neq \emptyset$  then
5       return False;
6   end
7 end
8 return True;
9 end
10 is_self_intersecting  $\leftarrow$  not IsSimple( $P$ );
11  $Q_c \leftarrow \infty$ ;
12 for  $i \leftarrow 40$  to  $n$  do
13    $d \leftarrow \|P^{(i)} - P^{(0)}\|$ ;
14    $Q_c \leftarrow \min(Q_c, d)$ ;
15 end
16 is_closed  $\leftarrow (Q_c < \tau)$ ;
17 return is_closed, is_self_intersecting;

```




Figure 4. The closure test performed on the coupler curve in Algorithm 1. The blue curve is the coupler curve that we are evaluating, the red line indicates the distance $P^{(i)} - P^{(0)}$, and the purple dashed line indicates the iteration through the curve.

Self-intersection test. Let $P = \{P^{(0)}, P^{(1)}, \dots, P^{(n)}\}$ denote the sampled planar trajectory, and define the associated polyline segments

$$S = \{ \overline{P^{(i)}P^{(i+1)}} \mid i = 0, \dots, n-1 \}.$$

The curve is declared self-intersecting if there exist two non-adjacent segments $\overline{P^{(i)}P^{(i+1)}}, \overline{P^{(j)}P^{(j+1)}} \in S$ with $|i - j| > 1$ such that

$$\overline{P^{(i)}P^{(i+1)}} \cap \overline{P^{(j)}P^{(j+1)}} \neq \emptyset.$$

Adjacent segments sharing a common endpoint are excluded from this test. In practice, this amounts to checking all pairs of non-adjacent segments for intersection using standard segment–segment intersection predicates. If any such intersection is found, the boolean flag `is_self_intersecting` is set to `True` and the curve is rejected.

2.2.2. Smoothness Test

We assess the smoothness of a coupler curve using its curvature. Large curvature values correspond to sharp turns or abrupt changes in direction, which are generally undesirable for walking trajectories.

To reduce sensitivity to sampling noise, we smooth the sampled (x, y) coordinates using a one-dimensional Gaussian filter. We then compute first- and second-order derivatives via numerical gradients and evaluate the curvature at each sample point as

$$\kappa = \frac{|x'y'' - y'x''|}{(x'^2 + y'^2)^{3/2}}.$$

If the maximum curvature along the curve exceeds a threshold α , the trajectory is rejected as non-smooth. In our experiments, $\alpha = 15$ provided an effective cutoff for eliminating trajectories with excessively sharp features. Algorithm 2 summarizes the procedure, and Figure 5 visualizes three example coupler curves with curvature values shown by color.

Algorithm 2: Smoothness Test via Curvature Thresholding.

Input: Coupler coordinates $P = \{(x_i, y_i)\}$, smoothness threshold α

Output: `has_sharp_corners`

- 1 $x, y \leftarrow \text{gaussian_filter_1D}(P, \text{smoothing } \sigma = 1)$
 - 2 Compute $x' = \nabla x, y' = \nabla y, x'' = \nabla x', y'' = \nabla y'$
 - 3 $\kappa_i = \frac{|x'_i y''_i - y'_i x''_i|}{(x'^2_i + y'^2_i)^{3/2}} \forall i$
 - 4 `has_sharp_corners` $\leftarrow$ `true_if_any`($\kappa_i > \alpha$)
 - 5 **return** `has_sharp_corners`
-

2.2.3. Flatness Detection: Stance, Contact, and Ground Model

In this work, we define stance and ground contact using a purely geometric approximation based on the shape of the foot trajectory. The selected coupler point is treated as the contact point, and the ground is modeled as a straight line inferred from the flat portion of the trajectory via least-squares fitting. To reason about flatness consistently, we analyze trajectories in a ground-aligned coordinate frame obtained by rotating the curve so that its dominant direction of motion is horizontal. The stance phase is then identified as the longest near-horizontal segment of the trajectory in this aligned frame. This definition provides a practical proxy for foot-ground contact and does not attempt to model contact forces, slip, or dynamics.




Figure 5. The figure shows three example coupler curves with curvature color-mapped along their lengths. A log scale for curvature is used for better visualization. (a) is the first example, (b) is the second example, and (c) is the third example.

The flat portion of the coupler curve corresponds to the stance phase of walking, during which the foot remains approximately in contact with the ground and provides support and propulsion.

We detect this stance segment by (i) rotating the trajectory so that the dominant direction of variation aligns with the x -axis, and (ii) identifying regions with sufficiently small slope in the rotated frame.

To align the curve, we apply Principal Component Analysis (PCA) to the sampled (x, y) points and rotate the trajectory such that the first principal component (the direction of greatest variance) is aligned with the x -axis. For walking-like trajectories, the stance segment is typically the longest near-linear portion and therefore dominates the variance, making this alignment a practical approximation. After alignment, we compute the slope, and the points with slope magnitude below a threshold are labeled as “flat”; in our implementation, we use $\epsilon = 0.0075$. See Figure 6 for illustration.

We then identify contiguous runs of flat points and define the flatness fraction as the length of the longest flat run normalized by the total curve length (Algorithm 3). The minimum required flatness fraction depends on the desired duty factor and the number of legs operating in parallel. As a simple guideline, if N identical legs are phase-shifted to provide continuous ground contact, each leg typically requires a stance fraction of at least $1/N$ of the cycle. In this study, we accept a trajectory if at least one quarter of the cycle is classified as flat (flatness fraction ≥ 0.25), corresponding to a nominal requirement of $N \leq 4$ legs for continuous contact under ideal phase scheduling. Additional legs may be needed in practice for stability and balance, which are hardware-level considerations beyond the scope of this work.

While vertical deviation constraints could limit absolute displacement along the y -axis, they do not directly suppress local oscillatory behavior within a tolerance band. By monitoring the derivative $\frac{dy}{dx}$ during the stance phase, the method enforces directional consistency and limits small oscillations that could otherwise occur while remaining within vertical bounds.

The derivative is computed using sufficiently fine discretization of the trajectory, and numerical stability was verified empirically during implementation.

Algorithm 3: Flat Region, indices, and flatness fraction detection via PCA Alignment

Input: Coupler coordinates $P = \{(x_i, y_i)\}$, slope threshold $\epsilon = 0.0075$, minimum flat fraction $f_{min} = 1/4$

Output: flat_indices, flat_fraction, is_flat

```

1  $P \leftarrow \text{resample\_uniform}(P, 512)$ ;
2 rotated_curve  $\leftarrow \text{pca\_align}(P)$ ;
3 flat_mask  $\leftarrow |\text{gradient}(\text{rotated\_curve}[:, 1])| < \epsilon$ ;
4 max_run  $\leftarrow 0$ ; run_indices  $\leftarrow []$ ;
5 current  $\leftarrow []$ ;
6 for  $i = 0$  to  $\text{length}(P) - 1$  do
7   if flat_mask[i] then
8     current.append(i);
9     if  $\text{len}(\text{current}) > \text{max\_run}$  then
10      max_run  $\leftarrow \text{len}(\text{current})$ ; run_indices  $\leftarrow \text{current.copy}()$ ;
11    end
12  else
13    current  $\leftarrow []$ ;
14  end
15 end
16 flat_fraction  $\leftarrow \text{max\_run} / \text{length}(P)$ ;
17 is_flat  $\leftarrow \text{flat\_fraction} \geq f_{min}$ ;
18 return run_indices, flat_fraction, is_flat;
```




Figure 6. Visualization of the flatness detection by aligning the principal axis and detecting low slope regions. The orange line indicates the principal axis, the dashed purple arrow shows the algorithm iterating over the coupler curve, and the red portion indicates the run of points with low slopes below the slope threshold ϵ .

2.2.4. Ground Intersection Test

The final test verifies ground clearance: when the mechanism is oriented such that the foot is supported by its flat stance segment, no other joint trajectory should intersect the ground. We approximate the ground by a line ℓ fitted in the least-squares sense to the detected stance-phase foot points. Specifically, ℓ is chosen to minimize the sum of squared orthogonal distances to these points. We then compute the signed distance of every joint position to ℓ , with the sign chosen so that the supporting foot points lie on the zero or negative side. The mechanism passes the test if, for all time steps, every non-foot joint remains strictly on the positive (clearance) side of ℓ . See Algorithm 4 and Figure 7. In this algorithm, m and b are the slope and y -intercept, respectively.




Figure 7. This figure shows an example mechanism that does not pass the ground intersection test (**left**), and one that passes the test (**right**).

Algorithm 4: Joint-Ground Intersection Test.

Input: Coupler path $P = \{(x_i, y_i)\}$, joint coordinates $J = \{(x_i, y_i)\}$, flat indices I_{flat}

Output: isGroundClear

- 1 $(m, b) \leftarrow \text{linear_regression}(P_{I_{\text{flat}}})$
- 2 $d_{\text{path}} \leftarrow P_y - (mP_x + b)$
- 3 $\text{majoritySign} \leftarrow \text{sign}(\sum d_{\text{path}})$
- 4 $d_{\text{joints}} \leftarrow J_y - (mJ_x + b)$
- 5 $\text{isGroundClear} \leftarrow \text{all}(\text{sign}(d_{\text{joints}}) \text{ is equal to } \text{majoritySign})$
- 6 **return** isGroundClear

2.3. Optimization Methodology

Even after passing the geometric filters, some candidates still fail functionally because one or more non-foot joints intersect the best-fit ground line estimated from the stance segment of the foot trajectory. Physically, this corresponds to linkages whose joints collide with the ground, preventing locomotion (Figure 8).




Figure 8. An example of a near-walking 6-bar mechanism. Initial joint positions are shown as green X's, and all joint trajectories are shown. In this specific mechanism, joints 0 and 2 are grounded. The coupler curve passes through all criteria; however, joints 3 and 4 pass through the ground line. The offending portion intersecting the ground line is marked in red.

In practice, many of these near-feasible mechanisms require only small geometric adjustments to resolve the collision. Modest perturbations of one or two joint locations can often restore ground clearance while preserving the overall walking-like trajectory. Discarding such candidates would therefore eliminate designs that are close to feasible and potentially valuable.

For this reason, we retain a subset of near-feasible mechanisms that do not strictly satisfy the ground-clearance condition and subsequently refine them through an optimization-based fine-tuning step (Algorithm 5).

Algorithm 5: Optimization of Mechanisms with Ground Intersections

Input: Initial parameters θ , mechanism representation matrices $\mathbf{B}$, $\mathbf{S}$ and $\mathbf{I}$, flatness threshold ϵ

Output: Optimized trajectories $\mathbf{P}^*$, and configuration θ^*

- 1 Simulate to obtain all joint trajectories $\mathbf{P} = \text{simulate}(\theta, \mathbf{B}, \mathbf{S}, \mathbf{I})$ (Section 2.1)
- 2 Resample $\mathbf{P}$ uniformly along its path.
- 3 Retrieve coupler curve CC from trajectory.
- 4 Compute ground line $(m, b, s) = \text{compute_ground_line_with_sign}(CC)$
- 5 Identify joints intersecting the ground:
 $J_{adj} = \text{find_joints_below_ground}(\mathbf{P}, m, b, s)$
- 6 Define optimization variables ϕ (offsets for joints in J_{adj}).
- 7 Set search bounds $\phi \in [-5, 5]^{2|J_{adj}|}$.
- 8 **Define error function** $E(\phi)$:
 - 9 For a given ϕ , obtain mechanism initial state $\theta' = \text{apply}(\theta, \phi, J_{adj})$.
 - 10 Obtain new trajectories $\mathbf{P}' = \text{simulate}(\theta', \mathbf{B}, \mathbf{S}, \mathbf{I})$.
 - 11 Compute geometric penalties (I, W, L, C) from $\mathbf{P}'$ and return total error
 $E(\phi) = I + W + L + C$.
- 12 Minimize $E(\phi)$ using Differential Evolution until convergence:
- 13 $\phi^* = \arg \min_{\phi} E(\phi)$
- 14 Apply optimal offsets ϕ^* to θ^* and regenerate optimized trajectories
 $\mathbf{P}^* = \text{simulate}(\theta^*, \mathbf{B}, \mathbf{S}, \mathbf{I})$.
- 15 **return** $\mathbf{P}^*, \theta^*$

Specifically, we identify the set J_{adj} of joints whose simulated trajectories violate the ground-clearance test. We optimize only these joints by applying planar offsets ϕ to their initial coordinates, keeping all other joints fixed. The search domain is $\phi \in [-5, 5]^{2|J_{adj}|}$. We use Differential Evolution (DE), a population-based, derivative-free optimizer suitable for non-smooth objectives [22], which is used in other linkage synthesis works [27]. Other optimization algorithms, like improved artificial bee colony (IABC) [28] and genetic algorithms [29], are also used in linkage mechanism design [30–32]. Constraints are handled implicitly by the kinematic simulator (which rejects invalid configurations) and explicitly via penalty terms added to the objective. Unless otherwise stated, we use standard DE hyperparameters (population size, mutation, and recombination) and stop after a fixed maximum number of objective evaluations or when the best objective does not improve for G generations.

To enable this refinement, the ground-intersection filter is implemented as a lightweight feasibility check rather than an exhaustive verification of all joint trajectories, thereby admitting imperfect designs for downstream correction.

We formulate fine-tuning as the minimization of a custom, design-driven objective that penalizes ground violations while encouraging preservation of walking-relevant trajectory

features. The objective is intentionally design-driven rather than physically based, reflecting the geometric nature of the filtering criteria used throughout the pipeline.

Specifically, the objective consists of four terms:

1. **Intersection penalty (I):** penalizes self-intersection of the foot trajectory.
2. **Walking penalty (W):** penalizes degradation of walking features, including insufficient stance flatness and overly sharp turns.
3. **Ground-violation penalty (L):** penalizes the extent of non-foot joint trajectories that lie below the estimated ground line.
4. **Closure penalty (C):** penalizes loss of periodicity (i.e., trajectories that become effectively open).

The overall objective is defined as

$$E = I + W + L + C. \quad (1)$$

The following subsections define each term.

Penalty weighting rationale: The penalty weights (2×10^7 , 1×10^4 , 1×10^2 , 4×10^8) reflect constraint hierarchy. Self-intersection (I) and closure (C) receive heaviest penalties as they represent fundamental infeasibility. Ground violation (L) uses quadratic scaling to strongly penalize persistent collisions while allowing small violations during optimization. Walking features (W) receive moderate penalties as they degrade performance but don't preclude functionality. These weights were tuned empirically to balance constraint satisfaction with optimizer convergence.

2.3.1. Intersection and Walking Penalties (I, W)

The intersection and walking penalties are auxiliary terms that discourage degeneracies which would invalidate the walking trajectory during optimization. Empirically, large penalty weights were required to reliably suppress self-intersection and to enforce the stance and smoothness constraints.

The intersection penalty is defined as

$$I = \begin{cases} 2 \times 10^7, & \text{if the foot trajectory is self-intersecting,} \\ 0, & \text{otherwise.} \end{cases}$$

The walking penalty W discourages (i) loss of stance flatness and (ii) loss of smoothness. We define

$$W = \begin{cases} 1 \times 10^4, & \text{if } \exists (x_i, y_i) \in P[\text{flat_indices}] \text{ such that } \left| \frac{dy}{dx}(x_i, y_i) \right| > \epsilon, \\ 1 \times 10^2, & \text{if } \exists i \text{ such that } \kappa_i > \alpha, \\ 0, & \text{otherwise,} \end{cases}$$

where ϵ is the flatness slope threshold, κ_i is the curvature at the i -th sampled point of the curve, and α is the curvature threshold.

2.3.2. Ground-Violation Penalty L

The ground-violation penalty encourages the optimizer to eliminate collisions between non-foot joints and the ground by reducing the portion of joint trajectories that lie below the estimated ground line.

Let the ground line be given by $y = mx + b$, as estimated from the stance segment of the foot trajectory. For each non-foot joint j , let $\{\mathbf{p}_{j,k}\}_{k=1}^T$ denote its sampled planar

trajectory. We define the ground-violation measure ξ as the total arc length of trajectory segments whose endpoints lie below the ground line:

$$\xi = \sum_{j \neq \text{foot}} \sum_{\substack{k=1 \\ \mathbf{p}_{j,k}, \mathbf{p}_{j,k+1} \in \mathcal{B}}}^{T-1} \|\mathbf{p}_{j,k+1} - \mathbf{p}_{j,k}\|,$$

where

$$\mathcal{B} = \{(x, y) \in \mathbb{R}^2 \mid y < mx + b\}$$

denotes the region below the ground line. The ground-violation penalty is then defined as

$$L = (10^8 \xi)^2.$$

This formulation penalizes both the extent and persistence of ground penetration while allowing small violations to be corrected through optimization.

2.3.3. Closure Penalty C

To encourage periodic walking motion, we penalize trajectories that fail to return sufficiently close to their starting point over one cycle. Let $\{\mathbf{p}_k\}_{k=1}^T$ denote the sampled foot (coupler) trajectory. We define the closure quality

$$Q_c = \min_{k \geq k_{\min}} \|\mathbf{p}_k - \mathbf{p}_1\|,$$

where $k_{\min}$ excludes the initial portion of the trajectory to avoid trivial near-matches (we use $k_{\min} = 40$). This is the same closure definition used earlier for hard filtering of open trajectories, here reused as a soft penalty during optimization.

The closure penalty is defined as

$$C = \begin{cases} 0, & Q_c \leq \tau, \\ 4 \times 10^8 (Q_c - \tau)^3, & Q_c > \tau, \end{cases}$$

where $\tau = 0.1$ is the closure tolerance. This formulation penalizes open trajectories while allowing small deviations from perfect closure.

3. Results

3.1. Dataset Generation

Dataset generation was performed on Stony Brook University's SeaWulf computing cluster. We simulated approximately 4×10^6 mechanisms for each of 25 mechanism subtypes (i.e., on the order of 10^8 total candidates). Among all simulated candidates, 23,250 mechanisms satisfied the proposed geometric criteria and were retained in the dataset. Figure 9 shows representative examples from the resulting dataset, and Table 1 shows the number of mechanisms generated per type, and aggregate statistics for each family. (For six-bar families, different variants define the coupler point on different links. The corresponding kinematic structures are publicly provided in the BSIdict at [21].) The filtered dataset of 23,250 mechanisms is publicly available [33] and contains data of mechanism type and initial joint coordinates.

3.2. Optimization Results

A subset of the retained mechanisms still exhibits geometric violations that prevent physical walking, most commonly due to non-foot joints intersecting the estimated ground line. We therefore evaluate the proposed fine-tuning procedure (Figure 3) for correcting

these near-feasible candidates using the custom objective defined in Section 2.3. Three representative near-walking mechanisms are selected, and ten independent optimization trials are performed for each. Table 2 reports the corresponding success rates, and Figure 10 compares the trajectories before and after optimization.

Across the evaluated cases, optimization success rates ranged from 30–90%. Failures were primarily associated with trajectories becoming effectively open (loss of closure) or with persistent ground intersections of non-foot joint trajectories. Consistent with intuition, candidates that were closer to feasibility at initialization typically required smaller adjustments and achieved higher success rates. Since differential evolution is stochastic, repeated runs can yield different outcomes; we therefore report results over multiple trials for each mechanism (Table 2).

Table 1. Number of Generated Mechanisms in Filtered Dataset and their Acceptance Rates.

Mechanism Type	Count	Mechanism Acceptance Rate (%)
Total (4-bar + 6-bar)	23,250	0.02325000
4-bar Mechanisms	12,783	0.07989375
RRRR	3926	0.09815000
RRRP	8219	0.20547500
RRPR	621	0.01552500
PRPR	17	0.00042500
6-bar Mechanisms	10,467	0.01246071
Stephenson (all variants)	3781	0.00859318
Watt (all variants)	6686	0.01671500

Table 2. Optimization Results for Selected Mechanisms.

Mechanism Type	Success	Fail
Example Mechanism 1	8	2
Example Mechanism 2	6	4
Example Mechanism 3	6	4

Table 3 shows the detailed error components for the three representative mechanisms for their original, successful and failed optimization runs. Successful optimization eliminates or significantly reduces error components: ground violation penalty L is reduced to zero in successful cases, while failed optimizations show L values one-to-two orders of magnitude higher, indicating persistent ground intersections and the optimizer getting stuck in local minima. Intersection penalty I and closure penalty C remain zero across all cases, confirming these hard constraints were maintained consistently. Walking penalty W is presented in some original mechanisms but is eliminated during successful optimization. The dramatic reduction in total error (up to 17 orders of magnitude for Mechanism 1) demonstrates the effectiveness of the penalty-based optimization approach.

Table 4 shows detailed improvement percentages for all error components from Table 3. Analysis of successful versus failed optimization runs reveals a clear trend: mechanisms with lower initial ground-violation penalties (L) exhibit higher convergence rates, whereas candidates with large initial L values are more likely to fail.

These observations suggest that the proposed differential evolution framework is effective as a refinement tool for mechanisms that are already close to geometric feasibility. Rather than transforming highly infeasible candidates, the optimizer reliably corrects minor violations while preserving closure and intersection constraints. This behavior aligns with the intended design of the filtering–optimization pipeline.




Figure 9. Selected walking mechanisms from the filtered dataset rendered using the MotionGen (2026) software. Subfigures (a–d) are 4-bar RRRR mechanisms, (e–h) are 4-bar RRRP mechanisms, (i–l) are 4-bar RRPR mechanisms, (m–p) are Stephenson 6-bars, and (q–t) are Watt 6-bars.




Figure 10. Optimization results for three example mechanisms. Each row corresponds to results for one example mechanism. Each column shows a stage of optimization: the original mechanism, a successful optimization, and a failed attempt. One representative success and failure are shown per mechanism. (a,d,g) are the original three mechanisms, (b,e,h) are the successfully optimized mechanism versions, and (c,f,i) are the failed optimization runs. Note the common features in the failed optimization trials shown—most failed runs were because the coupler curves remain under the ground line (c,f) or because the curve became open (c,f,i).

Table 3. Optimization Error Component Analysis for Representative Mechanisms.

Mechanism	State	I	W	L	C	Total Error
Mechanism 1	Original	2.00×10^8	1.00×10^4	5.49×10^{16}	0.0	5.49×10^{16}
	Success	0.0	0.0	0.0	0.0	0.0
	Fail	0.0	0.0	8.88×10^{18}	0.0	8.88×10^{18}
Mechanism 2	Original	0.0	1.00×10^4	3.59×10^{16}	0.0	3.59×10^{16}
	Success	0.0	1.00×10^4	0.0	0.0	1.00×10^4
	Fail	0.0	0.0	1.95×10^{18}	0.0	1.95×10^{18}
Mechanism 3	Original	0.0	0.0	6.05×10^{17}	0.0	6.05×10^{17}
	Success	0.0	0.0	0.0	0.0	0.0
	Fail	0.0	0.0	1.00×10^{18}	0.0	1.00×10^{18}

Table 4. Optimization Improvement Metrics for All Error Components.

Case	I (Intersection)		W (Walking)		L (Ground)		E (Total)	
	Initial	Final	Initial	Final	Initial	Final	Initial	Final
Mech 1 Success	2.00×10^8	0.0	1.00×10^4	0.0	5.49×10^{16}	0.0	5.49×10^{16}	0.0
Reduction	100%		100%		100%		100%	
Mech 1 Fail	2.00×10^8	0.0	1.00×10^4	0.0	5.49×10^{16}	8.88×10^{18}	5.49×10^{16}	8.88×10^{18}
Reduction	100%		100%		−16,172%		−16,172%	
Mech 2 Success	0.0	0.0	1.00×10^4	1.00×10^4	3.59×10^{16}	0.0	3.59×10^{16}	1.00×10^4
Reduction	—		0%		100%		>99.9999%	
Mech 2 Fail	0.0	0.0	1.00×10^4	0.0	3.59×10^{16}	1.95×10^{18}	3.59×10^{16}	1.95×10^{18}
Reduction	—		100%		−5331%		−5331%	
Mech 3 Success	0.0	0.0	0.0	0.0	6.05×10^{17}	0.0	6.05×10^{17}	0.0
Reduction	—		—		100%		100%	
Mech 3 Fail	0.0	0.0	0.0	0.0	6.05×10^{17}	1.00×10^{18}	6.05×10^{17}	1.00×10^{18}
Reduction	—		—		−65%		−65%	

3.3. Optimization Robustness Evaluation

To evaluate robustness beyond the three illustrative examples in Table 2, the optimization procedure was applied to a total of 10 near-walking mechanisms, including the three representative examples previously presented and seven additional cases randomly selected from the filtered dataset.

The mean final error across all 10 cases was $\mu = 8.85 \times 10^{17}$, with standard deviation $\sigma = 2.68 \times 10^{18}$. These statistics are summarized in Table 5. Considering only converged cases, the success rate was 50%, increasing to 70% when including partially improved solutions.

Among these 10 cases:

- 5 converged to zero total error,
- 2 achieved substantial reduction in geometric error but retained minor walk penalties,
- 3 did not converge to feasible closed configurations.

Across the ten evaluated mechanisms, five converged to zero total error, while two exhibited substantial error reduction but did not fully eliminate ground violations. Three cases diverged due to loss of trajectory closure. The mean final error therefore reflects a mixture of successful and failed trials, inflating the average despite complete convergence in half of the cases. When considering only successful runs, the final error is identically zero.

Table 5. Statistical summary of optimization performance over 10 additional near-walking mechanisms.

Metric	Value
Mean Initial Total Error	1.62×10^{18}
Mean Final Total Error	8.85×10^{17}
Standard Deviation (Final Error)	2.68×10^{18}
Full Convergence Rate	50%
Improvement Rate (Final < Initial)	70%

3.4. Comparative Analysis with Classic Walking Mechanisms

In this section, we compare a representative mechanism discovered through our geometric filtering pipeline with the classic Klann linkage (Stephenson III configuration). Figure 11 presents a side-by-side comparison of (left) the Klann linkage, (middle) a geometrically feasible Watt I mechanism from our dataset exhibiting similar walking characteristics, and (right) a modified version of the generated mechanism adjusted for manufacturability using straight links.




Figure 11. Comparison of classic and generated walking mechanisms. (a): Klann linkage (Stephenson III). (b): geometrically feasible Watt I mechanism discovered through the proposed filtering pipeline. (c): refined version of the generated mechanism with straight links for manufacturability.

All these three mechanisms produce feasible walking trajectories. However, mechanism in Figure 11b,c require only two ground joints, whereas the Klann linkage in Figure 11a needs three. This reduction may simplify structural support and anchoring requirements in practical implementations.

To further examine kinematic behavior, Figure 12 compares the foot trajectories and instantaneous velocity fields of the Klann linkage, the generated mechanism, and the refined generated mechanism. The resulting coupler curves exhibit comparable stride geometries and ground-contact behavior. Additionally, the velocity vectors of the generated mechanism indicate smoother kinematic transitions under the applied geometric constraints. This is especially seen throughout the stance phase, suggesting improved kinematic continuity under the geometric constraints enforced during filtering and optimization.

The rightmost configuration in Figure 11 illustrates a manually refined variant of the generated mechanism in which triangular links are replaced with straight members to reflect improved manufacturability. This modification was performed to demonstrate that geometrically feasible walking behavior can be preserved under structural simplification. While the present optimization framework does not explicitly enforce manufacturability constraints, this example suggests that additional objective terms or geometric penalties could be incorporated in future work to systematically account for fabrication-oriented design considerations.




Figure 12. Foot trajectory and velocity comparison between the three mechanisms. (a): The foot trajectory of the existing Klann Linkage. (b): Foot trajectory of the discovered Watt mechanism. (c): Foot trajectory of the modified Watt mechanism. Comparable stride envelopes are observed in (b,c). The generated mechanism exhibits smoother velocity direction transitions during stance.

While this comparison does not constitute a full dynamic or energetic evaluation, it demonstrates that the proposed geometric synthesis framework is capable of producing structurally distinct walking mechanisms with kinematic behavior comparable to established designs, and that these mechanisms can be further adapted to practical engineering considerations.

4. Conclusions and Future Outlook

This paper presented a design-criteria-driven kinematic synthesis pipeline for planar 1-DOF walking mechanisms within selected 4-bar and 6-bar mechanism families. The proposed workflow combines large-scale simulation, automated geometric filtering, and optimization-based refinement to construct a curated dataset of feasible walking candidates and to improve near-feasible designs.

Using the proposed filters, we evaluated approximately 100,000,000 mechanisms across 25 subtypes and obtained a dataset of 23,250 walking mechanisms. Only approximately 0.024% of the sampled designs satisfied the walking criteria, indicating that feasible walking behavior is rare under naive random sampling and that multi-criterion feasibility constraints significantly restrict the design space. To increase yield beyond filtering alone, we introduced a custom, design-driven objective for refinement of near-feasible candidates. Across the tested cases, the optimization success rate ranged from 30–90% (Table 2), demonstrating that targeted refinement can substantially improve feasibility relative to pure sampling and filtering.

Looking ahead, the central idea of treating application performance as computable geometric/kinematic criteria suggests a general path toward task-specific dataset construction and synthesis in mechanism design. Future work will focus on extending the approach to additional mechanism families and applications, improving sampling strategies to reduce computational cost, and integrating learning-based models to guide search toward high-feasibility regions of the design space while preserving strict feasibility checks via simulation.

Additionally, although the present work focuses on geometric feasibility and large-scale mechanism generation, the resulting database of over 23,000 walking-feasible designs presents an opportunity for future statistical analysis. Pattern extraction techniques could be applied to investigate correlations between link-length ratios, joint distributions, and qualitative walking characteristics. Such analysis may enable the derivation of higher-level structural design principles beyond the geometric criteria imposed in this study.

4.1. Limitations

A primary limitation is computational cost. In this study, only 0.024% of the 100 million evaluated mechanisms passed the filters, implying that brute-force dataset generation can require large-scale computation for narrow or strict criteria. On average, evaluating and filtering 4 million mechanisms for a single subtype required approximately 5–6 h on Stony Brook University’s SeaWulf computing resources. Running the same workload on typical personal hardware would be substantially slower, which may limit practicality for users without access to high-performance computing or for problems with similarly low feasibility rates.

A second limitation is that the filtering criteria are application-dependent. While the pipeline is general, each new robotics task requires the specification of appropriate feasibility and performance criteria (e.g., collision avoidance, workspace constraints, payload clearance, or timing requirements). Designing such criteria may require domain expertise and iterative validation to ensure that the filters reflect the true functional requirements rather than an overly restrictive proxy. Nevertheless, once formulated, these criteria can be evaluated systematically at scale within the same simulation–filter–refine framework.

4.2. Integration with Physical Design

While this work focuses on geometric feasibility, future work could potentially explore physical implementation. Generated mechanisms could be coupled with: (1) dynamic analysis via physics-based linkage simulation to evaluate stability, contact forces, and torque, (2) control synthesis for actuation timing, (3) morphological optimization for leg compliance and energy efficiency, and (4) manufacturing constraints (link thickness, joint clearances). Our filtered dataset provides a high-quality starting point for these downstream design stages. Additionally, transmission angles are not explicitly constrained in the present study. Incorporating transmission-angle bounds represents another important extension for future physically realizable designs.

Author Contributions: Conceptualization, R.T., Z.L. and A.P.; methodology, R.T. and Z.L.; software, Z.L.; validation, R.T.; formal analysis, R.T.; investigation, R.T.; resources, Z.L., A.P. and Stony Brook University; data curation, R.T.; writing—original draft preparation, R.T.; writing—review and editing, Z.L. and R.T.; visualization, R.T.; supervision, A.P.; project administration, A.P.; funding acquisition, A.P. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The dataset generated in this study is publicly available on Kaggle at: <https://www.kaggle.com/datasets/purwarlab/walking-mechanisms> (accessed on 21 January 2026).

Acknowledgments: The first author began this research through participation in the Simons Summer Research Program (Website: <https://www.stonybrook.edu/simons/> (accessed on 13 January 2026)) at Stony Brook University and continued this work in the Computer-Aided Design and Innovation Lab thereafter. The authors acknowledge the Simons Summer Research Program for providing the initial opportunity and mentorship. The authors also acknowledge the Institute for Advanced Computational Science (IACS) at Stony Brook University for providing access to the SeaWulf high-performance computing cluster, which enabled large-scale mechanism simulation and dataset generation. During the preparation of this manuscript, the author(s) used ChatGPT 5.2 for the purposes of text polishing. The authors have reviewed and edited the output and take full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. Variable Definitions and Thresholds

Table A1. Notation used in the paper.

Symbol	Meaning
τ	Closure tolerance (0.1)
α	Curvature threshold for smoothness (15)
ϵ	Slope threshold for flatness (0.0075)
Q_c	Closure quality (min distance from start)
κ_i	Curvature at point i
f_{min}	Minimum flat fraction (flat proportion of coupler curve)
P	Coupler coordinates $P = \{P^{(0)}, P^{(1)}, \dots, P^{(n)}\}$, $P^{(i)} = \{(x_i, y_i)\}$
J	Initial joint coordinates $J = \{(x_i, y_i)\}$

Threshold Justification: Threshold values were empirically determined through iterative refinement on representative mechanisms across all mechanism families. $\tau = 0.1$ allows for numerical precision while ensuring periodic motion. $\alpha = 15$ eliminates sharp turns while preserving walking-like smoothness. $\epsilon = 0.0075$ ensures sufficiently horizontal stance phases for ground contact. $f_{min} = 0.25$ permits four-legged walking configurations. These values were validated by manual inspection to appropriately balance filtering rigor with mechanism diversity.

References

1. Purwar, A. MotionGen. 2026. Available online: <http://www.motiongen.io> (accessed on 21 January 2026).
2. Lyu, Z.; Purwar, A.; Liao, W. A Unified Real-Time Motion Generation Algorithm for Approximate Position Analysis of Planar N-Bar Mechanisms. *ASME J. Mech. Des.* **2023**, *146*, 063302. [CrossRef]
3. Garbuz, M.; Klimina, L.; Samsonov, V. Wind driven plantigrade machine capable of moving against the flow. *Appl. Math. Model.* **2022**, *110*, 17–27. [CrossRef]
4. Klann, J.C. Walking Device. U.S. Patent 6260862B1, 11 February 1998.
5. Nansai, S.; Elara, M.R.; Iwase, M. Dynamic Analysis and Modeling of Jansen Mechanism. *Procedia Eng.* **2013**, *64*, 1562–1571. [CrossRef]
6. Nansai, S.; Rojas, N.; Elara, M.R.; Sosa, R.; Iwase, M. On a Jansen leg with multiple gait patterns for reconfigurable walking platforms. *Adv. Mech. Eng.* **2015**, *7*, 1687814015573824. [CrossRef]
7. Norton, R. *Design of Machinery: An Introduction To The Synthesis and Analysis of Mechanisms and Machines*, 5th ed.; McGraw Hill: Columbus, OH, USA, 2011.
8. Erdman, A.G.; Sandor, G.N. *Mechanism Design: Analysis and Synthesis*, 2nd ed.; Prentice-Hall: Englewood Cliffs, NJ, USA, 1991; Volume 1.
9. Ge, Q.; Purwar, A.; Zhao, P.; Deshpande, S. A Task-Driven Approach to Unified Synthesis of Planar Four-Bar Linkages Using Algebraic Fitting of a Pencil of G-Manifolds. *ASME J. Comput. Inf. Sci. Eng.* **2017**, *17*, 031011. [CrossRef]
10. Galan-Marin, G.; Alonso, F.J.; Del Castillo, J.M. Shape optimization for path synthesis of crank-rocker mechanisms using a wavelet-based neural network. *Mech. Mach. Theory* **2009**, *44*, 1132–1143. [CrossRef]
11. Nurizada, A.; Purwar, A. An Invariant Representation of Coupler Curves Using a Variational AutoEncoder: Application to Path Synthesis of Four-Bar Mechanisms. *ASME J. Comput. Inf. Sci. Eng.* **2023**, *24*, 011008. [CrossRef]
12. Nobari, A.; Srivastava, A.; Gutfreund, D.; Xu, K.; Ahmed, F. LInK: Learning Joint Representations of Design and Performance Spaces through Contrastive Learning for Mechanism Synthesis. *arXiv* **2024**, arXiv:2405.20592. [CrossRef]
13. Nobari, A.H.; Srivastava, A.; Gutfreund, D.; Ahmed, F. LINKS: A Dataset of a Hundred Million Planar Linkage Mechanisms for Data-Driven Kinematic Design. In *Proceedings of the Volume 3A: 48th Design Automation Conference (DAC)*; American Society of Mechanical Engineers: New York, NY, USA, 2022. [CrossRef]
14. Nurizada, A.; Dhaipule, R.; Lyu, Z.; Purwar, A. A Dataset of 3M Single-DOF Planar 4-, 6-, and 8-Bar Linkage Mechanisms With Open and Closed Coupler Curves for Machine Learning-Driven Path Synthesis. *ASME J. Mech. Des.* **2024**, *147*, 041702. [CrossRef]
15. Lyu, Z.; Purwar, A. Deep Learning Conceptual Design of Sit-to-Stand Parallel Motion Six-Bar Mechanisms. *ASME J. Mech. Des.* **2024**, *147*, 013306. [CrossRef]
16. Nobari, A.H.; Chen, W.; Ahmed, F. Range-Constrained Generative Adversarial Network: Design Synthesis Under Constraints Using Conditional Generative Adversarial Networks. *J. Mech. Des.* **2021**, *144*, 021708. [CrossRef]

17. Vermeer, K.; Kuppens, R.; Herder, J. Kinematic Synthesis Using Reinforcement Learning. In *Proceedings of the Volume 2A: 44th Design Automation Conference*; American Society of Mechanical Engineers: New York, NY, USA, 2018. [CrossRef]
18. Fogelson, M.B.; Tucker, C.; Cagan, J. GCP-HOLO: Generating High-Order Linkage Graphs for Path Synthesis. *J. Mech. Des.* **2023**, *145*, 073303. [CrossRef]
19. Bolanos, D.; Ataei, M.; Jayaraman, P. MechaFormer: Sequence Learning for Kinematic Mechanism Design Automation. *arXiv* **2025**, arXiv:2508.09005. [CrossRef]
20. Kapsalyamov, A.; Hussain, S.; Brown, N.A.; Goecke, R.; Hayat, M.; Jamwal, P.K. Synthesis of a six-bar mechanism for generating knee and ankle motion trajectories using deep generative neural network. *Eng. Appl. Artif. Intell.* **2023**, *117*, 105500. [CrossRef]
21. Purwar Lab. Four, Six and Eight-Bar Mechanisms With Curves. 2025. Available online: <https://www.kaggle.com/datasets/purwarlab/four-six-and-eight-bar-mechanisms-with-curves> (accessed on 21 January 2026)
22. Storn, R.; Price, K. Differential Evolution—A Simple and Efficient Heuristic for Global Optimization over Continuous Spaces. *J. Glob. Optim.* **1997**, *11*, 341–359. [CrossRef]
23. Chen, C.H. Application of Multiple Deep Neural Networks to Multi-Solution Synthesis of Linkage Mechanisms. *Machines* **2023**, *11*, 1018. [CrossRef]
24. Zhang, Y.; Arakelian, V. Design and Synthesis of Single-Actuator Walking Robots via Coupling of Linkages. *Front. Mech. Eng.* **2021**, *6*, 609340. [CrossRef]
25. Amine, S.; Prasad, B.; Saber, A.; Mokhiamar, O.; Gazo-Hanna, E. Design and Analysis of a Planar Six-Bar Crank-Driven Leg Mechanism for Walking Robots. *Appl. Sci.* **2024**, *14*, 8919. [CrossRef]
26. Shigley, J.E. The Mechanics of Walking Vehicles: A Feasibility Study. Technical Report AD419858, U.S. Army Ordnance Corps/U.S. Army Tank Automotive Command, 1960. Archived by the U.S. Department of Defense Technical Information Center. <https://apps.dtic.mil/sti/tr/pdf/AD0419858.pdf> (accessed on 22 February 2026).
27. Gogate, G.R.; Matekar, S.B. Optimum synthesis of motion generating four-bar mechanisms using alternate error functions. *Mech. Mach. Theory* **2012**, *54*, 41–61. [CrossRef]
28. Gao, W.; Liu, S. Improved artificial bee colony algorithm for global optimization. *Inf. Process. Lett.* **2011**, *111*, 871–882. [CrossRef]
29. Deb, K. *Multi-Objective Optimization Using Evolutionary Algorithms*; John Wiley & Sons: Chichester, UK, 2001.
30. Zhang, R.; Ning, M.; Wang, Y.; Yang, J. Design and Dimension Optimization of Rigid–Soft Hand Function Rehabilitation Robots. *Machines* **2025**, *13*, 311. [CrossRef]
31. Cabrera, J.; Ortiz, A.; Nadal, F.; Castillo, J. An evolutionary algorithm for path synthesis of mechanisms. *Mech. Mach. Theory* **2011**, *46*, 127–141. [CrossRef]
32. Liu, X.; Ding, J.; Wang, C. Design Framework for Motion Generation of Planar Four-Bar Linkage Considering Clearance Joints and Dynamics Performance. *Machines* **2022**, *10*, 136. [CrossRef]
33. Purwar Lab. A Dataset of Planar 1-DOF Walking Linkages. 2025. Available online: <https://www.kaggle.com/datasets/purwarlab/walking-mechanisms> (accessed on 21 January 2026).

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Article

Real-World Experimental Evaluation of DDoS and DRDoS Attacks on Industrial IoT Communication in an Automated Cyber-Physical Production Line

Tibor Horak ¹, Roman Ruzarovsky ^{2,*}, Roman Zelník ², Martin Csekei ² and Ján Šido ²

¹ Institute of Applied Informatics, Automation and Mechatronics, Faculty of Materials Science and Technology in Trnava, Slovak University of Technology in Bratislava, 91724 Trnava, Slovakia; tibor.horak@stuba.sk

² Institute of Manufacturing Technologies, Faculty of Materials Science and Technology in Trnava, Slovak University of Technology in Bratislava, 91724 Trnava, Slovakia; roman.zelnik@stuba.sk (R.Z.); martin.csekei@stuba.sk (M.C.); jan.sido@stuba.sk (J.Š.)

* Correspondence: roman.ruzarovsky@stuba.sk

Abstract

Automated production lines are increasingly being expanded with Industrial Internet of Things (IIoT) devices, creating complex Cyber-Physical Systems (CPSs) that connect physical production with control and information infrastructure. However, the convergence of Information Technology (IT) and Operational Technology (OT) layers creates new entry points for attacks targeting communication availability. Most existing studies analyze Distributed Denial of Service (DDoS) attacks primarily in simulation or testbed environments, with limited experimental verification of their impact on real-world production systems. This article presents an experimental evaluation of the impact of DDoS and Distributed Reflection Denial of Service (DRDoS) attacks carried out directly on a physical automated production line with integrated IIoT infrastructure during real operation. Three attack scenarios (TCP SYN flood, TCP ACK flood, and ICMP reflected attack) were implemented, targeting Programmable Logic Controllers (PLCs), Radio-Frequency Identification (RFID) subsystems, and selected IIoT devices. The results showed rapid degradation of deterministic PROFINET communication, disruption of the link between the OT and IT layers, loss of digital product representation, and physical interruption of the production process. Based on the findings, a minimally invasive security solution based on perimeter protection was designed and experimentally verified. The results emphasize the need to design IIoT-based manufacturing systems with an emphasis on network segmentation and architectural separation of the IT and OT layers.

Keywords: industrial cybersecurity; DDoS mitigation; DRDoS attacks; industrial IoT; cyber-physical manufacturing system; programmable logic controller; Industry 4.0

1. Introduction

Automated manufacturing systems in the Industry 4.0 concept are increasingly based on the integration of Industrial Internet of Things (IIoT) devices that connect the physical manufacturing process with the control and information infrastructure [1–3]. The convergence of Information Technology (IT) and Operational Technology (OT) enables advanced monitoring, remote access, and real-time data exchange, but it also fundamentally changes the security profile of industrial systems [4,5]. Thus, modern cyber-physical

systems (CPSs) are exposed to network attacks that can threaten their availability and production continuity [6,7].

A particular risk is the possibility of IIoT devices being misused for attacks targeting availability, such as Denial of Service (DoS), Distributed Denial of Service (DDoS), or Distributed Reflection Denial of Service (DRDoS) [8–10]. IoT devices can represent a potential attack vector targeting the broader manufacturing environment given their architectural position between the IT and OT layers [11]. Their constant connectivity and often limited security mechanisms increase the risk of misuse as part of botnets or as a source of excessive network traffic directed at the control infrastructure [12,13]. In a CPS environment, such traffic can lead to the degradation of deterministic communication, disruption of control cycles, and physical shutdown of the manufacturing process in extreme cases.

Although the issue of DDoS attacks in IIoT and industrial networks is the subject of intensive research, most existing studies are conducted in simulation environments, laboratory testbeds, or with publicly available network traffic datasets. These approaches mainly focus on detecting or classifying attacks using analytical or machine learning methods [14].

However, experimental verification of the direct impact of an attack on a physically operated production system based on industrial hardware is significantly less common. Experimental studies systematically analyze the manifestations of excessive network traffic on real control units, communication modules, and IIoT devices and assess its impact on control communication and production process continuity remain relatively limited.

Therefore, the main objective of this research is to experimentally examine the vulnerability of specific industrial devices used in Industry 4.0 architecture, in particularly Programmable Logic Controller (PLC) control units, Radio Frequency Identification (RFID) subsystems, and IIoT devices. The IIoT area is specific in the context of cybersecurity, as these devices represent an integration point between the IT and OT layers and are exposed to a higher degree of network connectivity than traditional control elements.

The uniqueness of the presented research lies in the fact that the experiments are performed on a simulated or emulated model rather than on a physical automated production line built from commonly used industrial components, such as PLC control units, industrial Ethernet communication, identification systems, and IIoT devices.

Attacks are carried out during active system operation, analyzing the behavior of specific devices, control communication degradation, and physical consequences on the production process. The research focuses on experimentally investigating the vulnerability of real industrial architecture based on industrial devices, with an emphasis on the system's physical behavior during an attack.

This research aims to extend the existing automated production line with elements of the IIoT in accordance with the principles of Industry 4.0, physically implement this architecture from industrial control and communication components representing the typical structure of modern production systems, and subsequently experimentally analyze its vulnerability to DDoS and DRDoS attacks during active system operation [15,16].

The experiments focused on assessing the impact of attacks on system availability, the functionality of industrial communication infrastructure, including data exchange between PLC control units, IIoT devices, and higher-level systems, as well as the stability of the production process during actual operation. At the same time, the aim was to identify critical architectural points of the system and experimentally verify the effectiveness of the proposed countermeasures in real operating conditions.

This study focuses on assessing the degradation of deterministic communication, the behavior of specific control and IIoT devices, disruptions in data exchange between the OT and IT layers, and the physical consequences of an attack on production continuity, while also analyzing the fact that these manifestations may appear to be a normal technical

failure from the perspective of production operations, with no obvious indication of a cyber attack. The research also includes experimental verification of the effectiveness of architectural countermeasures, such as network segmentation and communication infrastructure perimeter protection.

The main contributions of this work can be summarized as follows:

- Experimental evaluation of the impact of DDoS and DRDoS attacks on a physically implemented cyber-physical production line without simulation or emulation.
- Analysis of the impact of cyber attacks on communication between production equipment and IIoT components, including identification of communication failure moments.
- Impact of internal and external network attack scenarios on production process stability.
- Identification of critical communication points and vulnerable components in a cyber-physical production system.
- Design and experimental verification of a firewall-based security solution that can be implemented without major changes to the existing infrastructure.
- Verification of the proposed security measures' ability to mitigate the impact of attacks and maintain the availability of production communication.

Although the experimental infrastructure was implemented on a production line intended primarily for educational purposes, its control architecture, PLC units, and industrial communication protocols represent a fully fledged industrial solution. Therefore, the obtained experimental results and the identified mechanisms of communication degradation are applicable in the broader context of industrial production systems without the need for major architectural modifications.

The results highlight the need to design production systems based on the principle of security by design, with an emphasis on network segmentation, communication interface control, and architectural separation of IT and OT layers to minimize the risk of control communication degradation and disruption of production operations.

The rest of the article is structured as follows: Section 2 presents related work and theoretical background, Section 3 describes the materials and methods used, including the production line architecture and the attack scenarios implemented, Section 4 contains the results of experimental measurements and their analysis, Section 5 discusses the results along with the limitations of the study, and Section 6 presents conclusions and implications for the design of more secure industrial architectures.

2. Related Work

IIoT devices represent a specific architectural element that fundamentally changes the integration of IT and OT layers in the context of modern manufacturing system security [17]. As stated by Wang et al. [18], the growing interconnectedness and convergence of IT and OT infrastructures significantly increase the number of entry points into industrial systems, thereby increasing their exposure to cyber threats.

Similarly, Comert et al. [19] pointed out that the integration of IIoT devices, especially when retrofitting existing production systems, creates new security gaps resulting from the simultaneous operation of legacy technologies and modern IIoT solutions. Their integration position between OT and IT infrastructure, combined with a high degree of network connectivity, makes them critical points in terms of attacks targeting availability and unauthorized access.

IIoT has significantly transformed the manufacturing system architecture in the Industry 4.0 concept [20]. The integration of sensors, edge devices, communication gateways, and cloud platforms into the control infrastructure enables continuous operational data collection, real-time processing, and advanced analytical tool implementation.

As stated by Yousuf et al. [21], IIoT architectures play a key role in the predictive maintenance and diagnostics of electric drives, enabling early fault detection and optimization of service interventions.

Sundar et al. [22] demonstrated the use of IIoT in increasing assembly line productivity through data-driven optimization methodologies, while Mashayekhy et al. [23] highlighted the importance of IoT technologies in inventory management and logistics.

IIoT is also used to monitor production assets in automated work environments, for example, through RFID technologies, as documented by Vagaš et al. [24].

Advanced IIoT applications also lay the foundation for digital twin production equipment and processes. Hedayati Kia et al. [25] and Bokhtiar Al Zami et al. [26] point to the growing importance of digital twins in monitoring equipment status, simulating operating scenarios, and optimizing control. Purohit et al. [27] confirmed that one of the main trends in current industrial digitalization is connecting the IIoT infrastructure with digital models.

The development of IIoT as part of the Industry 4.0 concept is systematically analyzed, for example, in the work of Qiu et al. [28], who point to the integrative nature of IIoT as a connecting element between physical production equipment, analytical tools, and digital platforms in the smart manufacturing environment. Similarly, Marinova-Kostova and Kostov [29] emphasized that the implementation of IIoT technologies fundamentally changes the communication architecture of industrial systems and creates new data flows between the IT and OT layers.

Research on the IIoT and CPS security points to the growing complexity of architectures integrating IT and OT layers and the associated security risks.

Wang et al. systematically analyzed IIoT security challenges across individual architectural layers and identified recurring vulnerabilities in authentication, access rights management, communication protocol protection, and DoS/DDoS attack resistance. The authors also note that traditional security approaches developed for IT environments are not fully applicable to IIoT systems due to their close connection to physical processes.

In recent years, the security of industrial cyber-physical systems (ICPSs) and IIoT devices has been the subject of systematic research, with particular emphasis on the convergence of IT and OT layers and the associated expansion of the attack surface.

Kayan et al. [30] analyzed ICPS security from the perspective of architectures, communication protocols, and differences between IT and OT domains, pointing to a fundamental shift from “security by obscurity” to highly interconnected systems exposed to external threats. Most existing solutions focus on anomaly detection, Intrusion Detection Systems (IDSs) mechanisms, and network traffic analysis, but these approaches are insufficiently connected to real incidents and with the assessment of specific vulnerabilities reported in industrial practice.

This study focused on the IIoT and edge architectures [31,32]. They emphasize the growing role of distributed computing nodes, wireless communication, and cloud integration, while highlighting new risks arising from device heterogeneity, limited computing capacity, and the long-life cycle of industrial components. Identified weaknesses include insufficient authentication, weak security of industrial protocols, and insufficient network segmentation, which take on specific significance in the CPS environment given the deterministic nature of control cycles.

Other analyzed works focus on attack modeling, detection algorithm design, and security mechanism evaluation in simulation or testing environments [33]. The dominant research approach consists of using network traffic datasets and applying machine learning methods to classify attacks, with detection accuracy, false positive rate, or communication latency serving as evaluation metrics.

A comparative analysis of these works shows that although the literature systematically identifies vulnerabilities in IIoT and ICPS architectures and proposes advanced detection and defense mechanisms, experimental verification of the impact of attacks on physically implemented manufacturing systems during their actual operation remains limited. Most studies evaluate security mechanisms at the network layer or in simulated conditions, with less attention paid to analyzing the behavior of specific control units, IIoT devices, and deterministic industrial communication in active manufacturing operations.

These shortcomings create potential vulnerabilities in the system that can be exploited to compromise the availability, integrity, or confidentiality of transmitted data. However, in the CPS environment, the degradation of the communication layer has specific significance, as the disruption of deterministic communication can lead to the destabilization of control cycles and, consequently, to physical consequences for the production process.

Unlike traditional IT systems, attacks targeting availability in an industrial environment are not only a problem for network infrastructure but can also directly affect the operation of technological equipment.

DDoS attacks are among the most critical threats to availability in IIoT and CPS architectures because they can quickly exhaust network and computing resources, thereby indirectly disrupting the deterministic data exchange on which real-time process control relies.

Research focusing on DDoS attacks in IIoT and CPS environments is predominantly oriented toward the network and data layers of architecture. Current research in this area focuses primarily on the network and computing layers of architecture, particularly on edge cloud nodes, Software-Defined Networking (SDN) controllers, and communication infrastructure.

Anomaly detection approaches using machine learning and deep learning methods dominate, as do proposals for adaptive defense and resilient frameworks, whose effectiveness is evaluated using metrics such as accuracy, precision, recall, latency, throughput, and resource availability recovery.

Al-Obeidat et al. [34] analyzed DDoS attack detection in an Edge-IIoT environment supported by a digital twin using deep learning models verified on a publicly available dataset generated in a digitally modeled environment. They reported a detection accuracy of approximately 96% without a connection to physical control devices.

Zakariah et al. [35] proposed a hybrid machine learning-deep learning (ML/DL model for IDS in IIoT/SCADA environments with an accuracy exceeding 99%, with evaluation based on classification metrics (accuracy, precision, recall, and F1-score) over experimental network traffic datasets.

Alrumaih and Alenazi [36] present the ERINDA simulation framework for increasing the resilience of ICS to DDoS attacks, where they evaluate throughput, latency, and availability recovery rate ($\approx 88\%$ improvement in throughput recovery), but exclusively in a controlled simulation environment.

Orman [37] analyzes hybrid IDS architectures for IIoT and points to improvements in detection adaptability, with validation taking place primarily on datasets in his review study.

Zhukabayeva et al. [38] identified the architectural vulnerabilities of IIoT and edge systems from a conceptual perspective without experimentally verifying the physical impacts of attacks.

Hareesh et al. [39] experimentally quantified the impact of DoS attacks on SCADA system communication metrics (latency, packet loss) without analyzing the behavior of production modules.

Eryilmaz et al. [40] confirm the high accuracy of ML/DL models for DDoS detection but point out the limited verification in real industrial architectures in a systematic review.

Swati et al. [41] proposed an adaptive Moving Target Defense mechanism for IIoT edge environments tested in Mininet simulation, where they evaluated resource availability and service response, with no validation on physically implemented control systems.

The analyzed works show that the current research achieves very high classification of metric values and significant improvements in network resilience, confirming progress in the design of detection and mitigation mechanisms for IIoT environments. Nevertheless, these approaches are mainly evaluated at the level of datasets, simulation models, or virtualized testbeds, with an emphasis on detection accuracy, network throughput, and quality of service (QoS).

However, there is no systematic experimental verification of how a DDoS attack manifests on specific physically implemented control units, deterministic industrial protocols, and the stability of a real-world production process.

The result is a gap between the high detection accuracy achieved under controlled conditions and limited empirical knowledge about the physical consequences of attacks in realistic industrial architectures integrating IIoT elements.

Table 1 provides a structured overview of selected scientific papers focusing on DDoS attacks in IIoT environments and industrial control systems (ICSs/SCADA) for a systematic comparison of the methodological approaches, experimental environments, and evaluation metrics of published studies. The table summarizes the analyzed environment, methodology used, focus of the work, declared contribution, and the available quantitative results, enabling the identification of dominant research trends and existing gaps in the field of experimental verification.

Table 1. Overview of selected research papers focused on DDoS attacks in IIoT/ICS environments.

Author (Year)	Environment	Methodology	Key Performance Metrics	Reported Values	Validation Type
Swati et al. (2025) [41]	IIoT	MTD + ML classifier	DDoS mitigation via adaptive MTD	MTDTM architecture with dynamic admission and service migration	5–20% improvement in response time; 50 ms classifier latency; statistically significant resource availability gain ($p < 0.001$)
Zakariah et al. (2025) [35]	IIoT/SCADA	Hybrid ML–DL (DeepFM)	SCADA intrusion detection	Cross-dataset validation (WUSTL 2018/2021, HAI)	Accuracy 95.6–99.98%; F1 up to 0.9997; Recall up to 0.9765
Alrumaih & Alenazi (2025) [36]	Industrial Control Systems	Simulation (ns-3)	DDoS resilience and recovery	ERINDA adaptive recovery framework	Throughput recovery up to 88%; PDR +6%; RIT restored close to normal levels
Orman (2025) [37]	IIoT/SCADA	ML, DL and Hybrid comparative study	Binary DoS detection WUSTL-IIoT-2021	Unified evaluation of 12 models	Accuracy up to 99.999% (MLP); F1 99.999%; precision 99.996%; recall 99.998%
Al-Obeidat et al. [6]	Edge-IIoT/Digital Twin	CNN + continual learning	DDoS detection in digital twin	Incremental knowledge transfer	Accuracy 0.96; F1 0.96; Recall 0.99 (Edge-IIoT digital twin dataset)

These incidents and published analyses suggest that cyber attacks in industrial environments are no longer an isolated IT infrastructure problem but have a direct impact on the availability and continuity of production processes. The literature also points to a trend of combined and coordinated attacks, in which distributed denial of service (DDoS) attacks are used as a destabilizing element affecting communication and control infrastructure.

In addition to academic research, real-world security incidents highlight the growing vulnerability of industrial systems. Recent cases have shown that cyber attacks are no longer an isolated IT infrastructure problem but can directly impact the availability and continuity of production processes. Table 2 summarizes the selected incidents' analysis and identifies several recurring characteristics.

Table 2. Overview of selected incidents affecting industrial companies and infrastructure.

Year	Organization/Goal	Country	Type of Attack	Impact
2025	Wibaie (window and door manufacturer) [42]	France	Ransomware—production stoppage	Complete shutdown of production during the incident
2025	Water and energy facilities (ICS threats) [38]	Canada	Cyber attacks on SCADA/PLC	Unauthorized changes to settings, potential physical impacts
2022	Energy, water, manufacturing [43]	Ukraine	DDoS + ICS attacks	Service disruptions, blackouts
2023–2024	Energy, production [44]	Germany	DDoS + ransomware	Failures of IT/OT
2019–2023	Aluminum smelters [45]	Norway	DDoS + ransomware	Long-term production outage

The interconnection of IT and OT layers creates an environment in which the consequences of attacks that originally targeted the information infrastructure are transferred to the control and technology environment. DDoS attacks are increasingly used in combination with ransomware or direct attacks on ICS/SCADA systems, intensifying production operations' destabilizing effect. Simultaneously, incidents point to insufficient industrial network segmentation and the high dependence of physical processes on communication infrastructure availability.

Simultaneously, several empirical analyses from various sectors point to significant economic and operational damage caused by availability attacks, confirming their potential to disrupt critical processes beyond the purely network level. These findings suggest that DOS attacks have systemic impacts and can affect operational continuity, management stability, and technological operation security.

However, an analysis of the existing literature shows that research in the field of DDoS attacks in IIoT and ICS environments is predominantly focused on the design of detection algorithms, classification model optimization, and detection accuracy evaluation. Most studies use publicly available datasets or simulated testbeds and focus on network metrics such as accuracy, F1-score, throughput, and latency. Although these approaches provide significant insights into attack identification and mitigation, the physical consequences of attacks on real-world manufacturing systems are rarely evaluated.

Therefore, the research gap lies in the lack of experimental studies on physically implemented industrial architectures that would quantify the operational and technological consequences of attacks targeting availability in an environment integrating IT, OT, and IIoT elements. This study addresses this gap by conducting controlled experiments on a physically implemented production line with real PLC control units and Industry 4.0 architecture components, systematically quantifying the impact of availability attacks on deterministic industrial communication, control cycle stability, and production process continuity.

3. Materials and Methods

This chapter describes the experimental architecture, infrastructure extension, and cyber attack scenarios used to evaluate the impact of availability attacks on a physically

implemented cyber-physical manufacturing system. This experimental study quantifies the operational and technological consequences of DDoS and DRDoS attacks on a real production line based on Industry 4.0 principles, integrating IT, OT, and IIoT layers.

Experiments were conducted on a physically implemented modular production system. The original architecture was purposefully expanded with IIoT integration elements, monitoring nodes, edge devices, and remote access elements to create a realistic industrial environment that reflects the current implementation principles of Industry 4.0 and the interconnection of IT and OT layers.

The experimental methodology is based on the controlled simulation of DDoS and DRDoS availability attacks and systematic monitoring of their impact on the control, communication, and technology layers of the system.

The experiment was designed in the form of three different scenarios representing different availability disruption mechanisms in a production line environment.

The first scenario represents a direct TCP SYN flood DDoS attack that targets selected control and monitoring layer elements.

The second scenario was implemented as a direct TCP ACK flood DDoS attack targeting control and IIoT integration devices.

The third scenario simulated a reflected DRDoS (ICMP flood) attack, in which the target was flooded through intermediary nodes, modeling a more realistic form of distributed attack with traffic amplification.

Each scenario was implemented within a controlled time interval with a defined traffic intensity, and the system behavior was compared with that of the reference no-traffic state.

Three main metrics were systematically monitored for all compromised devices:

1. TTCL (Time to Communication Loss)—the time from the start of the attack to the loss of stable communication.
2. Peak traffic (packets/s)—maximum intensity of network traffic during the attack.
3. Peak CPU load (%)—maximum utilization of the device's processor.

The TTCL metric is a key indicator of a system's resilience to denial-of-service attacks because it expresses the time interval during which the device is capable of maintaining deterministic communication. Peak traffic characterizes the intensity of the attack and allows for comparison between scenarios, while Peak Central Processor Unit (CPU) load documents the depletion of computing resources during the load.

The following subchapters describe in detail the physical CPS infrastructure, its extension with IIoT elements, and the specific implementation of DDoS and DRDoS scenarios.

3.1. Cyber-Physical System (CPS) Infrastructure Based on IIoT

The experimental infrastructure of the cyber-physical system was implemented based on the modular production line Modular Production System (MPS) System 203 Industry 4.0 from Festo Didactic GmbH & Company KG, Munich, Germany [46]. This production line is shown in Figure 1. The system represents a compact but technologically representative production line that integrates physical production processes, decentralized control, identification technologies, and a higher-level Manufacturing Execution System (MES) information system, creating a suitable environment for researching the principles of Industry 4.0 and CPS.

The production line represents a complex mechatronic environment comprising three autonomous stations: a distributing/conveyor station, a joining station, and a sorting station. Each of these stations forms a separate physical and mechatronic unit equipped with its own conveyor module with an electric motor drive, pneumatic actuators such as stoppers, deflectors, and ejectors, as well as sensors for detecting the presence and position

of workpieces. The joining station is also extended with a suction gripper and a pick-and-place module, enabling the automated handling of parts during the production process.




Figure 1. Physically implemented cyber-physical production line MPS 203 Industry 4.0 (Festo Didactic, Munich, Germany), consisting of three autonomous stations (distribution, assembly, and sorting) controlled by Siemens SIMATIC S7-1512-1PN PLC units, Siemens AG, Munich, Germany. A superior MES coordinates the production process, with an integrated RFID subsystem ensuring material flow and digital product representation.

The control of individual stations is decentralized using Siemens SIMATIC S7-1512C-1 PN, Siemens AG, Munich, Germany programmable logic controllers (PLC), with each station controlled by a separate PLC.

Siemens SIMATIC S7-1512-1PN PLC units were selected as the S7 family's highest-performance class of control units. Their use allows the evaluation of the system response even for devices with high computing and communication capacities, thereby minimizing the impact of any hardware limitations on the experimental results. The control units ensure sequential control of technological operations, processing of signals from sensors, control of electric motor drives and pneumatic actuators, as well as communication with the RFID system and the superior MES production system. The PLCs are configured as PROFINET controllers and the peripheral devices are connected via SysLink interfaces and bus I/O nodes, with the pick-and-place module integrated into the control architecture via a PROFINET bus node.

RFID technology ensures product identification and tracking within the manufacturing process, thereby creating a link between the system's physical and digital layers. MPS 203, Industry 4.0 identification subsystem is implemented using Turck industrial RFID components and consists of RFID transponders integrated into workpieces and four high-frequency RFID read/write heads operating at 13.56 MHz in accordance with ISO 15693 [47], Turck GmbH, Mülheim an der Ruhr, Germany. The distribution and sorting stations are equipped with one RFID head, whereas the connection station uses two RFID heads, enabling identification and data updating when the product enters and exits the station.

The RFID read/write heads are not connected directly to the PLC control units but via decentralized Turck TBEN-S2-2RFID-4DXP modules, which function as PROFINET IO devices and RFID masters for one or two connected heads. In the case of a connection station, both RFID heads are connected to a single TBEN module, creating a local RFID subsystem. Data on the production order, required technological operations, and execution status are written to the RFID transponders, and this information is autonomously evaluated by the Siemens SIMATIC S7-1512C-1 PN PLC units and used to decide on the further processing of the product.

The production line's communication infrastructure based on the PROFINET industrial Ethernet network, which connects Siemens control units, PROFINET I/O nodes, an RFID subsystem, an industrial Ethernet switch, and a computer with a higher-level MES. The network is implemented as a local industrial Local Area Network (LAN) in a private address space with a fixed Internet Protocol (IP) address for each device, which ensures deterministic communication, stable system operation, and repeatability of experimental conditions during cyber experiments.

At the protocol level, PROFINET communication between PLCs and IO devices is based on cyclical real-time data exchange, where process data are transferred within defined communication cycles. Each PLC functions as a PROFINET controller and regularly updates the connected IO devices' input and output process images. This deterministic communication structure ensures the control cycle's stable execution under normal operating conditions.

Each station is controlled by a separate PLC configured as a PROFINET controller. The RFID subsystem is implemented using decentralized Turck modules that function as PROFINET IO devices and communicate cyclically with the respective PLCs. Each TBEN module supports two HF RFID channels, and two RFID read/write heads in the connection station are connected to one TBEN module. The network also includes a pneumatic valve terminal with a PROFINET interface, which controls the pneumatic elements of the Joining station.

The superior cyber layer of the cyber-physical system comprises the MPS[®] MiniMES production information system from Festo Didactic, which runs on an industrial PC with the Windows 10 operating system. This system ensures product and work plan definition, production order generation and management, production and technological data recording in RFID transponders, and continuous production status monitoring. Thus, MiniMES is a key element of the system's cyber part, integrating the production process's planning, control, and supervision. The basic control and peripheral components of the original MPS 203 Industry 4.0 production line architecture are listed in Table 3.

The presented implementation uses the OPC DA (OLE for Process Control—Data Access) specification via the DELTALOGIC S7/S5 OPC Server (version 4.12.0.11527, Delta Logic Automatisierungstechnik GmbH, Schwäbisch Gmünd, Germany). Communication is implemented in a classic client-server architecture based on the COM/DCOM model, with the OPC server and Siemens SIMATIC S7-1512C-1 PN PLC units using the S7-TCP/IP within an industrial Ethernet network. The OPC interface enables the cyclic reading and writing of process variables, data blocks (DBs), memory words (MWs), and status variables between the PLC control layer and the higher-level MES. This ensures the continuous exchange of technological data, status information, and production parameters between the system's OT and IT layers.

The OPC interface is used to transfer status information from the control units to the MES and distribute production orders and parameters to the production stations. The MES communicates primarily with the first station of the production line, while the production

process, the status of individual orders, and product tracking are evaluated based on data stored in RFID tags and based on the status signals of the PLCs of individual stations.

Table 3. Overview of the main control and peripheral components of the original architecture of the MPS 203 Industry 4.0 production line, including their functional role, network addressing, and communication protocol used.

Component	Manufacturer	Model	Role	IP Address	Protocol
PLC—Distributing	Siemens AG, Munich, Germany	S7-1512C-1 PN	PROFINET Controller	192.168.1.10	PROFINET
PLC—Joining	Siemens AG	S7-1512C-1 PN	PROFINET Controller	192.168.1.20	PROFINET
PLC—Sorting	Siemens AG	S7-1512C-1 PN	PROFINET Controller	192.168.1.30	PROFINET
RFID interface—Distributing	TURCK GmbH, Mülheim an der Ruhr, Germany	TBEN-S2-2RFID-4DXP	PROFINET IO Device	192.168.1.11	PROFINET
RFID interface—Joining	Turck GmbH	TBEN-S2-2RFID-4DXP	PROFINET IO Device	192.168.1.21	PROFINET
RFID interface—Sorting	Turck GmbH	TBEN-S2-2RFID-4DXP	PROFINET IO Device	192.168.1.31	PROFINET
IO -Link valve terminal	Festo Vertrieb GmbH & Co. KG, Esslingen, Germany	CTEU-PN	PROFINET IO Device	192.168.1.22	PROFINET
OPC Server	Delta Logic Automatisierungstechnik GmbH, Schwäbisch Gmünd, Germany	S7/S5 OPC Server v4.12.0.11527	IT/OT data exchange (MES-PLC)	opcda://DELTALOG IC.OPC.S7.DA	OPC DA over S7-TCP/IP (Industrial LAN)

The entire system is designed as a cyber-physical production system based on closed-loop feedback between the physical production process and decentralized control using Siemens SIMATIC S7-1512C-1 PN, digital product representation using RFID technology, and the MPS MiniMES higher-level information system. The production line implements discrete variant production of an assembled product consisting of a base body with an integrated RFID transponder and an optional cover, with data linked to each individual piece determining the specific product configuration.

The production process is initiated by creating a production order in the superior MES, which defines technological operations, production variants, quality control requirements, and shipping methods. At the distribution station, this data is recorded in the product's RFID transponder, which carries information about the processing status and operations performed throughout the entire production cycle. The control units of the individual stations autonomously evaluate this data and use it to decide on the production cycle.

At the connecting station, a dynamic decision is made on whether to perform or skip the cover assembly using a pick-and-place module based on data read from the RFID transponder. The result of this operation is then updated back to the RFID tag, ensuring a consistent digital representation of the product status. The sorting station then checks the selected product properties and redirects the product to the appropriate output branch designated for domestic or international shipment or discarded as a defective item based on a combination of RFID data and inspection results.

Decisions on technological steps are not made according to a fixed production line sequence but are dynamically based on data linked to a specific product, with local control at the level of individual stations. After processing is complete, the product is transported to the output point, and the single-piece production cycle is closed, with all relevant data on the production process available in the superior MES.

3.2. Expanding Infrastructure Towards IIoT, Monitoring, and Remote Access

The MPS 203 Industry 4.0 system was originally designed as a didactic cyber-physical manufacturing system that supports selected Industry 4.0 principles, particularly in the area of production management through MES and RFID-based product identification. The basic architecture of the system primarily focuses on the production process, order flow, and product-oriented management. Although the system had a web interface at the MES level, this interface focused exclusively on production planning and monitoring and did not provide comprehensive options for technical monitoring, equipment diagnostics, network infrastructure, or remote service interventions.

Figure 2 shows the overall configuration of the experimental production line, including the material flow between the distribution, connection, and sorting stations, the relevant PLCs, the RFID systems, and the IIoT infrastructure components. The figure provides a system-level overview of the physical process, control units, and communication elements that form the experimental cyber-physical architecture.




Figure 2. Schematic representation of the MPS 203 Industry 4.0 production line showing material flow, control layer, and IIoT infrastructure extension. The architecture comprises three autonomous distribution, assembly, and sorting stations controlled by Siemens SIMATIC S7-1512C-1 PN, connected to an RFID subsystem and a high-level MES. The original control architecture is extended with an industrial network element, edge devices (Siemens SIMATIC IOT2050, Siemens AG, Munich, Germany, Revolution Pi, KUNBUS GmbH, Ostfildern, Germany), monitoring modules (W&T), and remote access elements (Ewon Cosy+), which ensure the connection of the IT and OT layers without interfering with the production process's deterministic control.

Therefore, the original system architecture was purposefully expanded to include IIoT infrastructure, which aims to make operational and diagnostic data from automation devices accessible via standard network and application interfaces, enable remote monitoring and maintenance, and create an integration bridge between the OT and IT layers of the system.

In this context, IIoT does not replace the control system but supplements it, expanding the system with monitoring, diagnostic, and integration functions without interfering with

the deterministic control of production processes. The original architecture of the MPS 203 Industry 4.0 production line was subsequently expanded with IIoT infrastructure, which enables monitoring, remote access, and interconnection of the system's IT and OT layers (Figure 3).




Figure 3. IIoT and network infrastructure of the MPS 203 Industry 4.0 production line. The figure shows the fully assembled and interconnected elements of the extended architecture, including an industrial controllable network element, edge devices (Siemens SIMATIC IOT2050, Revolution Pi), VPN gateways for remote access, monitoring and I/O modules, and power supply infrastructure. The assembly forms an integration layer between the control (OT) and information (IT) parts of the system without interfering with the production process's basic control logic.

This expansion was based on the integration of a central industrial network infrastructure, as the original system configuration did not include a separate manageable network element. Therefore, a manageable PLANET IGS-6325-8T8S4X (PLANET Technology Corporation, New Taipei City, Taiwan) industrial Ethernet switch was integrated into the architecture to form the central communication node of the entire IIoT network.

The switch connects PLC control units, edge devices, remote I/O modules, Virtual Private Network (VPN) routers, and monitoring nodes into a unified industrial network and provides support for manageable L2 and L3 layer functions such as Virtual Local Area Network (VLAN) segmentation, routing, redundancy, and basic security mechanisms. This creates a stable and expandable communication foundation for the system's IIoT functions.

To support remote maintenance and service interventions, the architecture has been supplemented with an Ewon Cosy+ Ethernet industrial VPN router, which provides secure remote access to devices connected to the industrial Ethernet/PROFINET network, HMS Industrial Networks AB, Halmstad, Sweden. The router initiates an encrypted VPN connection to the Talk2M cloud service exclusively in the form of outgoing connections, eliminating the need to open network ports in the corporate or production network. This

IIoT element enables remote diagnostics, programming, and maintenance of the control system and represents a practical implementation of Industry 4.0 principles in the field of operation and maintenance.

The IIoT monitoring layer also includes a Wiesemann & Theis Ethernet remote I/O module (Wiesemann & Theis GmbH, Wuppertal, Germany), which is used for the production system power supply. The module functions as an embedded network device with proprietary firmware and an integrated Web-IO interface accessible via HTTP/HTTPS. The digital input is used to detect the supply voltage, while the relay digital output controls the LED indication of the current system status. The power status is available independently of the PLC control logic owing to the web interface, which increases system transparency and supports rapid diagnostics within the IIoT infrastructure.

Revolution Pi Core SE (KUNBUS GmbH, Ostfildern, Germany) devices and Siemens SIMATIC IOT2050 (Siemens AG, Munich, Germany) industrial edge devices were integrated into the architecture to serve as IIoT integration nodes and connect the OT and information IT layers. These devices ensure the collection, pre-processing, and accessibility of operational data without interfering with the deterministic control of PLC units.

The Revolution Pi Core SE module is based on the Raspberry Pi Compute Module 4S platform and runs on the Revolution Pi OS operating system, which is based on Raspberry Pi OS and Debian distribution (current version Bookworm—Debian 12). It is an industrially modified Linux distribution with an implemented RT (real-time) kernel patch and additional KUNBUS packages (e.g., piControl, PiCtory, MQTT client, and OPC UA server) that provide image management and communication integration functions.

The Siemens SIMATIC IOT2050 runs on the SIMATIC Industrial OS operating system based on Linux (Debian). It provides a stable and industry-oriented platform for edge computing, data integration, and information processing at the industrial network level.

The Node-RED 5.0 environment, which acts as a middleware layer in the proposed architecture, was implemented on both edge devices. Node-RED performs data transformation, filtering, and routing between communication protocols and logically processes messages between the OT and IT parts of the system. The MQTT publish–subscribe communication mechanism makes the processed data available for web monitoring and other information layer tools.

It should be emphasized that part of the implemented IIoT infrastructure is based on open-source software components. The Revolution Pi OS operating system is based on the Debian distribution, and the Node-RED environment is an open middleware tool for data integration. The use of open-source solutions increases the transparency of the architecture, enables security control, and supports the reproducibility of experimental results, which is important in the context of modern industrial control scenarios.

This creates a unified IIoT integration and monitoring layer that provides a comprehensive view of the system status directly based on data from the control layer while maintaining a strictly separated control and monitoring architecture.

The Siemens SIMATIC IOT2050 device is used in the system primarily for centralized collection and monitoring of data from three Siemens SIMATIC S7-1512C-1 PN control units, from which selected database blocks (DBs) containing key status, diagnostic, and process information are read. This creates a unified IIoT monitoring layer that provides a comprehensive view of the status of the entire system based directly on data from the control layer while maintaining a strict separation of control and monitoring.

In normal operating conditions, communication between production components and IIoT devices follows a cyclic data exchange model, where PLC control units regularly exchange process data with operating devices and monitoring systems. The process variables, sensor data, and control commands are continuously updated to ensure a deterministic

production operation. IIoT gateways and edge devices preprocess selected data and provide higher-level services, such as remote monitoring, visualization, and data logging. During cyber attacks, excessive network traffic disrupts this cyclical communication, causing delays or loss of control data, which directly affects production continuity. Therefore, the communication flow consists of repeated acquisition, processing, and forwarding of operational data between control units, industrial equipment, and supervisory systems, creating a closed control loop whose stability depends on reliable network communication.

The data flow can be summarized as follows: process variables generated at the PLC level are cyclically updated via PROFINET and made available to the MES via the OPC DA interface. Selected data blocks (DBs) are accessible via IIoT edge devices, where they undergo light preprocessing in the Node-RED middle layer, including filtering, formatting, and protocol translation. The processed data are then sent via MQTT to the monitoring and visualization services of the IT layer. This architecture creates a structured vertical data flow from the physical process to the supervisory information systems.

3.3. DDoS and DRDoS Attacks

DDoS attacks are one of the most serious threats to IIoT devices, which form the basis of modern cyber-physical manufacturing systems [48,49]. IIoT devices, such as programmable logic controllers, industrial sensors, actuators, and edge gateways, are primarily designed with an emphasis on reliability, deterministic behavior, and long-term continuous operation, with security mechanisms historically considered secondary. Consequently, many industrial devices are vulnerable to attacks targeting service availability, with even relatively low-intensity attacks potentially leading to serious operational disruptions given the limited computational resources, memory, and network capacity of IIoT elements [18].

DDoS attacks differ from those in traditional IT systems in the context of industrial networks. The attacker's goal is not only to overwhelm the network infrastructure but also to disrupt control and monitoring processes, delay or lose critical communication messages, and subsequently cause errors or unplanned production line downtime [50].

Attacks carried out at the transport or application layer, for example in the form of TCP SYN flood, UDP flood, or ICMP flood, can lead to the depletion of PLC and industrial gateway system resources, resulting in a loss of communication with SCADA or MESs [51].

Attacks targeting industrial communication protocols, such as Modbus/TCP, Profinet, EtherNet/IP, or OPC UA, which were originally designed without focusing on security mechanisms and flooding resistance, are also a serious problem [52].

DRDoS attacks represent a special category of attacks that differ from classic DDoS attacks in the way they generate and amplify attack traffic. In a DRDoS attack, the attacker does not communicate directly with the victim but exploits third parties, known as reflectors, which are legitimate servers or devices that provide network services. The attacker spoofs the source IP address of the victim and sends small requests to reflectors, which then respond with significantly larger packets to the target system. This mechanism leads to a significant amplification of attack traffic and makes identifying the actual attack origin difficult. In an IIoT environment, reflectors can be, for example, incorrectly configured UDP services, DNS, or NTP servers, as well as some industrial and diagnostic protocols, which increases the risk that the industrial devices themselves will be unwittingly involved in the attack [53].

This difference is particularly significant from the perspective of detection and defense mechanisms, as the traffic generated during a DRDoS attack may appear to be legitimate communication at first glance.

The fundamental difference between a classic DDoS and DRDoS attack is that in a classic DDoS attack, the target is directly flooded with attack traffic generated by a botnet

of compromised devices, whereas in a DRDoS attack, the target is flooded with responses from legitimate systems that have been exploited as attack intermediaries [54].

In an industrial environment, this can lead to the attack being misinterpreted as a sudden increase in process or diagnostic communication, which significantly complicates the timely response of security mechanisms.

The consequences of DDoS and DRDoS attacks in the IIoT environment go beyond service availability and can also have physical and security implications. Disruption of communication between control units and higher-level systems can lead to loss of synchronization, erroneous decisions by control algorithms, or activation of emergency conditions, which can have a direct impact on operator safety and the technological equipment integrity.

Therefore, effective protection of the IIoT system requires a combination of network segmentation, strictly defined communication rules, and advanced detection mechanisms based on network traffic behavioral analysis. Therefore, the current research focuses on developing lightweight detection algorithms and creating realistic datasets that will enable the reliable identification of DDoS and DRDoS attacks in the Industry 4.0 environment.

3.4. Traffic Monitoring and Data Analysis

The professional packet capture and analysis tools Capsa Portable Network Analyzer (Colasoft LLC, Pleasanton, CA, USA) and NetworkMiner Packet Analyzer (Netresec AB, Stockholm, Sweden), operated in fully licensed versions, were used to monitor network traffic and subsequently analyze the attack course. The Capsa Portable Network Analyzer is a tool designed for real-time monitoring and network communication diagnostics, enabling detailed packet analysis, measurement of traffic volume, and identification of network behavior anomalies. The NetworkMiner Packet Analyzer is a forensic tool that passively analyzes captured packets and reconstructs network sessions.

The tools used enabled detailed inspection of communication in the industrial Ethernet network, identification of peak traffic intensity (packets/s), analysis of protocol behavior (TCP/UDP/ICMP), and correlation of communication loss moments with measured network parameters during individual experimental scenarios.

4. Results

This chapter focuses on describing experimental attack scenarios used in testing production line communication and identifying its vulnerabilities during the production process in interaction with other IIoT devices within a shared network infrastructure. Particular attention is paid to DDoS and DRDoS attacks, as well as the results of the experiments and proposed solutions aimed at increasing the security and reliability of production line operation with integrated IoT devices.

The test environment is based on a real-time system. The expansion of the production line with IIoT devices in accordance with the Industry 4.0 concept enables communication via an external Wide Area Network (WAN) environment, but at the same time increases its exposure to potential cyber attacks from the external network.

The IIoT devices used in the production line pose another security risk, particularly industrial computing, networking, and communication components such as Siemens SIMATIC IOT2050, Ewon Cosy+ gateway, RevPi Core industrial module, and Web-IO devices from Wiesemann & Theis (W&T), which have limited security mechanisms in their basic configuration, increasing their vulnerability to network attacks, as shown in Figure 4.




Figure 4. Layered architecture of a cyber-physical manufacturing system based on the ISA-95 model, depicting the physical manufacturing process (levels 0–1), the PLC control layer with deterministic PROFINET communication (level 2), the CPS/IIoT information layer ensuring data collection, translation, and distribution (level 3.5), and the higher-level IT/MES layer for production planning, monitoring, and control (level 4), including remote access and network infrastructure. Although the presented architecture is structured according to the ISA-95 hierarchical model, it is compatible with the RAMI 4.0 framework, where the physical process corresponds to the Asset layer, the PLC control layer corresponds to the Integration and Communication layers, the CPS/IIoT layer corresponds to the Information layer, and the MES corresponds to the Functional and Business layers, especially in the context of vertical integration.

Although the architecture shown in Figure 4 is structured according to the ISA-95 hierarchical model, it is also compatible with the Industry 4.0 reference architecture model (RAMI 4.0). From the RAMI 4.0 perspective, the physical manufacturing process corresponds to the Asset layer, the PLC control layer corresponds to the Integration and Communication layers, the CPS/IIoT layer represents the Information layer, and the MES corresponds to the Functional and Business layers. Therefore, the presented implementation reflects practical vertical integration across RAMI 4.0 levels while maintaining the hierarchical structure of ISA-95.

To implement the test scenarios in the form of cyber attacks, the attacks were generated from a single personal computer (PC). This computer interfered with the communication of the network infrastructure consisting of IoT devices and a production line to verify its resistance to cyber attacks and identify potential vulnerabilities. The Kali Linux operating

system, which is equipped with specialized tools designed to test the security of network systems, including the Hping3 tool, was used to operate the attack system.

The Hping3 tool enables the generation and analysis of network traffic using TCP, UDP, and ICMP protocols, with packet generation controlled via the command line. This tool is primarily intended for security analysts to test and scan network infrastructure, identify security weaknesses, and assess the level of vulnerability risk. Appropriate and effective security mechanisms aimed at reducing the risk of cyber threats are based on the identified deficiencies.

4.1. Scenarios

As part of the experiments, three experimental attack scenarios were implemented to verify the resilience of the production line and connected IIoT devices against DoS network attacks. All experiments were performed on a fully functional production line with integrated IIoT devices and without interfering with its normal operation. The internal and external LAN and WANs initiated the attacks, allowing the comparison of the impact of local and remote attacks on the availability of individual production system components to be compared.

Two basic types of attacks, namely a direct DDoS attack and an indirect DRDoS attack, were carried out as part of the experiments. In a direct DDoS attack, the attacker generated a high volume of network traffic directed directly at selected production line devices or PLC control units to significantly limit their ability to process network communications and control tasks.

In an indirect DRDoS attack, the attacker exploited IIoT devices to relay malicious network traffic without directly compromising them, sending packets to non-existent ports on these devices with a spoofed source IP address. The IIoT devices then generated error responses that were directed to the target devices on the production line without the knowledge of the compromised IIoT devices.

In individual scenarios, specific types of network attacks, including TCP SYN flood, TCP ACK flood, and ICMP-based reflective attacks, were used. Packet generation was set to the maximum possible speed with no time limit on the individual attack duration. The attacks continued until the compromised devices lost their ability to communicate on the network and ceased performing their control or production functions.

The KPIs recorded during individual experimental attack scenarios are summarized in Table 4. From the values shown, it can be observed that in direct DDoS attacks, the intensity of network traffic on target devices reached higher maximum values, approximately 18,000 packets per second, while the maximum values were lower in reflected DRDoS attacks, approximately 14,000 to 14,500 packets per second. This difference is due to the fact that in reflected attacks, malicious traffic is generated indirectly through intermediary IIoT devices, which naturally limits the maximum intensity of the generated attack compared to a direct flood of the target device. At the same time, it is clear from the KPI that the attacked devices experienced 100% CPU load (Peak CPU load = 100%) during the attacks, indicating that the main reason for the loss of communication is the exhaustion of the devices' computing resources.

During the attacks, the number of packets in the network gradually increased, which led to the overload of communication and computing resources of individual devices. Internal communication buffers, processor processing capacities, and resources dedicated to network communication were exhausted as malicious traffic volume increased. The devices stopped processing incoming packets after exceeding the available communication and computing resources, resulting in a sudden drop in the recorded number of packets per second and a complete loss of communication between individual devices. Different

communication failure times between devices are caused by different computing capacities, communication mechanism implementation, and the overall ability of devices to process increased network load.

Table 4. Summary of key performance indicators measured during all implemented attack scenarios.

Attack Scenario	Type of Attack	Target Device	TTCL— Communication Loss (s)	Peak Traffic (Packets/s)	Peak CPU Load (%)
Scenario 1	DirectDDoS (TCP SYN flood)	RFID Turck system 1	8	18,677	100
		Web-IO 4.0	15	17,951	100
		PLC 1	17	18,660	100
		Simatic IIoT	26	18,315	100
Scenario 2	DirectDDoS (TCP ACK flood)	IIoT Link Valve	16	18,264	100
		PLC 3	19	17,987	100
		IIoT RevPi	22	18,456	100
		IIoT Ewon Cosy+	28	18,211	100
Scenario 3	Reflected DRDoS (ICMP flood)	Simatic IoT → Turck system 2	12	14,021	100
		IIoT RevPi → IO Link Valve	23	14,520	100
		IIoT Ewon Cosy+ → PLC 2	29	14,499	100

The results also show a sharp decline in recorded traffic after the compromised device lost its communication capability, as the device was no longer able to receive or process incoming packets, effectively eliminating the attack traffic from the monitored point of view. This confirms that the loss of communication was a direct result of the depletion of the attacked devices' communication and computing resources.

The main objective of the attack scenarios implemented was to disrupt the production process by flooding the network communication of individual components of the production line, either through a direct attack on the target devices or an indirect attack mediated through IIoT devices, highlighting the real security risks arising from the integration of IoT technologies into industrial systems.

4.1.1. Scenario 1

The first experimental scenario was implemented as a direct TCP SYN flood DDoS attack targeting production line devices and selected IIoT devices in the internal network infrastructure. The attacker generated a high number of network packets per second, simultaneously targeting multiple devices, as shown in Figure 5. During the attack, intense network traffic was recorded for each attacked device in the range of approximately 15,000–18,000 packets per second.

Figure 6 shows that individual devices responded to the attack with varying degrees of resilience. The Turck RFID system was the first to become inactive, ceasing to process network communication approximately 7 s after the start of the attack, as indicated by a sudden drop in the number of received packets to zero. Subsequently, approximately 16 s into the attack, the PLC 1 control unit stopped communicating, which led to an interruption in the control of the production line's technological processes.




Figure 5. Schematic of an experimentally implemented TCP SYN flood DDoS attack (Scenario 1). The attack is generated from an internal attack node (packet generator) and directed to selected production line components via an industrial network element, including the PLC control unit, Turck RFID subsystem, Siemens SIMATIC IOT2050 edge device, and Web-IO 4.0 module. The figure illustrates the routing of the generated TCP traffic and the attacked elements within the common industrial Ethernet infrastructure.




Figure 6. Time course of network traffic intensity during an experimental TCP SYN flood DDoS attack (Scenario 1). The graph shows the number of packets per second recorded on individual attacked devices (RFID Turck system 1, PLC 1, Siemens SIMATIC IOT2050, and Web-IO 4.0). A sudden drop in the number of packets to zero indicates a loss of communication with the device (TTCL—Time to Communication Loss), with different outage times indicating different resilience of individual components to availability attacks.

The Web-IO 4.0 device showed slightly higher resistance, but even in this case, network communication failed after approximately 15 s of attack. The most resistant device in the tested scenario was the Siemens SIMATIC IOT2050, which retained its communication capability for approximately 25 s after the start of the attack, when it was completely knocked out of operation.

The results of the experiment clearly show the high effectiveness of a DDoS attack carried out on an internal network, as all tested devices were taken out of operation in

less than 30 s. During the attack, the devices ceased to perform their production functions, and network communication with the control system and operator was lost. Restoring their functionality required manual intervention and restarting the devices, confirming the serious impact of this type of attack on the production system's availability and reliability.

4.1.2. Scenario 2

The second experimental scenario was implemented as a remote TCP ACK flood DDoS attack from an external WAN environment toward the production line's internal network. The aim of the attack was to verify the resilience of selected IIoT devices to network communication overload during an external network attack. A high number of TCP ACK packets per second were generated during the attack, with traffic intensity ranging from approximately 15,000 to 19,000 packets per second (Figure 7).




Figure 7. Diagram of an experimentally implemented TCP ACK flood DDoS attack (Scenario 2), initiated from the WAN environment and directed at the Ewon Cosy+ IIoT gateway, Revolution Pi edge device, PLC, and IO-Link valve terminal via the production network. The figure illustrates the penetration of overwhelming TCP traffic from the IT layer into production line's OT environment.

Figure 8 shows that individual devices responded differently to the TCP ACK flood attack. The first device to stop communicating was the IO-Link valve device, which became inactive approximately 15 s after the attack began, as indicated by a sudden drop in the number of received packets to zero. Approximately 18 s after the attack began, the PLC 3 control unit stopped communicating.




Figure 8. Time course of network communication during an experimentally performed TCP ACK flood DDoS attack conducted from an external WAN environment (scenario 2), showing the number of packets per second recorded on the PLC control unit and selected IIoT devices on the production line.

The RevPi Core device showed a higher degree of resilience; however, even in this case, the network communication failed after approximately 21 s of attack. The most resilient device in this scenario was the Ewon Cosy+ IIoT router, which maintained its communication capability for approximately 28 s from the start of the attack before being taken out of service.

The results of the second experimental scenario confirm that even a remote TCP ACK flood DDoS attack launched from a WAN environment can have a significant negative impact on the availability of IIoT devices and production lines. All tested devices were taken out of service in less than 30 s, losing their ability to communicate with the higher-level control system and operator. Restoring device functionality requires manual intervention, highlighting the serious risk this type of attack poses to the operation of industrial systems connected to external networks.

4.1.3. Scenario 3

The third experimental scenario was implemented as a distributed reflective DoS (DRDoS) attack conducted from an external WAN environment, in which IIoT devices on the production line were exploited as attack reflectors without their knowledge. The attacker generated ICMP packets directed at IIoT devices on the internal network, with the target packets being sent to non-existent ports on these devices. The attacker spoofed the source IP address and corresponded to the IP address of the production line device.

Consequently, IIoT devices automatically generated ICMP error messages, which were then sent to the target devices' spoofed IP address on the production line. As shown in the attack visualization in Figure 9, individual IIoT devices, specifically Ewon Cosy+, Siemens SIMATIC IOT2050, and RevPi Core, were exploited as intermediaries for reflective traffic directed at production components such as PLC 2, Turck RFID system, and IO-Link valve terminal.




Figure 9. Visualization of an experimentally implemented reflected ICMP flood-type DRDoS attack (Scenario 3) initiated from an external WAN environment. The attacker generates ICMP requests with a spoofed destination address, while selected IIoT devices (Siemens SIMATIC IOT2050, Revolution Pi, Ewon Cosy+) act as reflectors and generate ICMP error messages directed at the control and peripheral components of the production line (PLC, RFID subsystem, IO-Link valve terminal). The diagram illustrates the mechanism of traffic amplification and the spread of overwhelming communication across the production network's IT and OT layers.

Figure 10 shows that the attack caused significant congestion of the target devices within a few seconds. The first to stop communicating was the Turck RFID system, which was hit by reflective ICMP traffic from the Siemens SIMATIC IOT2050 device, with complete communication failure occurring approximately 11 s after the start of the attack. Approximately 22 s into the attack, the IO-Link valve terminal stopped communicating, overwhelmed by ICMP error messages generated by the RevPi Core device.

The highest level of resilience in this scenario was demonstrated by the PLC 2 control unit, which was the target of reflective traffic generated by the Ewon Cosy+ device. Even in this case, network communication was completely disrupted approximately 28 s after the start of the attack. All attacked devices ceased to perform their control and production functions and lost communication with the superior system.

The results of this experiment demonstrate the high effectiveness of a DRDoS attack carried out via ICMP reflection, in which the attacker can exploit legitimate IIoT devices to generate malicious traffic without directly compromising them. This type of traffic is particularly dangerous in industrial networks because it originates from internal security devices, making it significantly more difficult to detect and mitigate.




Figure 10. Time course of network communication during an experimentally implemented reflected DRDoS attack based on ICMP communication (scenario 3), showing the number of packets per second generated by selected IIoT devices and directed at target components of the production line.

4.2. The Impact of DDoS and DRDoS Attacks on the Manufacturing Process

The previous chapters presented three experimental scenarios of cyber attacks carried out during the full operation of the production process. These attacks were specifically aimed at disrupting the continuity of the production line. The results of the individual scenarios clearly showed that the production components and other IIoT devices became uncontrollable and lost their ability to communicate within the network infrastructure.

The failure of the RFID subsystem was the first problem encountered during the first experiment. The loss of communication between the Turck TBEN-S2-2RFID-4DXP modules and the PLC control units prevented the RFID read heads from reading data from the RFID transponders on the workpieces and prevented the PLC from further processing and transmitting this data to the higher-level MES. During the first experimental scenario, the first symptom was an RFID identification subsystem failure, which manifested itself as a loss of data exchange between the RFID modules and the higher-level MES system (Figure 11).

Despite this failure, the physical manufacturing process was initially continued because the local sequences controlling conveyors and technological operations at the PLC level were not immediately dependent on RFID data.

Therefore, the workpieces passed through the production line to the sorting station, where the fault became fully clear. The MES was unable to determine the correctness of the product due to missing or outdated identification data, and pieces that were physically correctly manufactured were automatically evaluated as defective and placed in the output bin for rejects. Since the identification and status data were not recorded in the MES system via the control layers, the digital representation of the product was lost and the continuity between the physical and cybernetic layers of the CPS was interrupted.

The second identified problem was a disruption in the PROFINET communication layer, which was directly observed in the Siemens SIMATIC S7-1512C-1 PN PLC control units.

The PLCs remained in the RUN mode, but with communication module failure signals and diagnostic messages indicating errors in the Local modules and PROFINET IO system categories. The control program continued to run without deterministic communication with the peripheral PROFINET IO devices. The loss of I/O device availability led to the PLC transitioning into a fault-safe state, interrupting control cycles, and physically

stopping production, with workpieces remaining stuck on conveyors and requiring manual intervention to continue production.




Figure 11. Status of the superior MES during an experimental DDoS attack (scenario 1) documenting the failure of the RFID subsystem. The absence of RFID input and output data (RFID InputString, RFID OutputString) despite the active status of the production process is shown in the figure. Communication loss occurs between Turck TBEN-S2-2RFID-4DXP RFID modules and the PLC control units.

The third problem manifested itself at the superior MES level, which did not explicitly record this failure. In the CPS's cyber layer, the situation manifested itself only as a stagnation of production orders without clear information about the cause of the process interruption. The RFID InputString and RFID OutputString fields remained empty in the MES interface, while the status of the stations was indicated as "process busy." Therefore, the MES system could not confirm the status of the product or decide on the next technological step, and the production history did not correctly record the events.

The loss of deterministic communication in the PROFINET communication layer meant that individual production stations were unable to correctly signal their status to higher-level or previous stations. Each station went into a fault or safe state and could not confirm its readiness to accept the next workpiece. Consequently, the material flow gradually halted, with workpieces remaining blocked at the conveyor ends of the individual stations, Figure 12.

The fourth problem was the failure of the CPS support and monitoring layers. The IIoT monitoring element implemented by the Web-IO Ethernet module was designed to monitor the system status and support maintenance. However, the network connection was lost during the attack, making it impossible to verify whether the production line was powered or shut down. Simultaneously, the mechanism for email notifications about changes in system status failed, making diagnostics significantly more difficult.




Figure 12. Diagnostic status of the production line's control and physical layers during an experimental DDoS attack (scenario 1). The left part of the image shows the Siemens SIMATIC S7-1512C-1 PN control unit's diagnostics in RUN mode with signaled faults in the Local modules and PROFINET IO systems categories, indicating a loss of communication with peripheral PROFINET IO devices. The right part documents the physical blockage of the production process, in which the workpieces remain stopped at the conveyor ends because the individual production stations cannot signal their readiness to accept the next workpiece.

The IIoT layer using Siemens SIMATIC IOT2050 industrial edge devices was also disrupted. After the connection between the edge application in the Node-RED environment and the system was lost, the operating data remained offline, making it impossible to monitor process variables, analyze time series, and detect anomalies in real time. Thus, the attack affected not only the control of the production system but also its observability. Figure 13 illustrates the failure of the IIoT monitoring layer implemented using the Siemens SIMATIC IOT2050 edge device and the Node-RED environment.




Figure 13. Status of the IIoT supervision layer during an experimental DDoS attack (scenario 1). The figure shows the Node-RED application interface running on the Siemens SIMATIC IOT2050 edge device with a lost connection to the server (Lost connection to server, reconnecting...) and inactive data flows, indicating the unavailability of operational and status data from the production system and the inability to monitor them in real time.

These findings confirm that there was no immediate complete failure of the production system during the first attack scenario but rather its gradual degradation across multiple CPS layers. Although the system appeared to be partially functional, it could not ensure correct product identification, production continuity, or reliable fault diagnosis. From an operational point of view, the production line appeared to be “stuck” first, products were sorted incorrectly, workpieces accumulated, and the production flow stopped. The MES did not consider this a faulty condition.

A particularly serious phenomenon was the systematic occurrence of products marked as NOT OK, even though the physical manufacturing process was running correctly. Physically correct products were rejected solely because of a failure in the CPS identification and communication layer, which led to a false decline in quality and a deterioration in performance indicators such as OEE. This scenario highlights a dangerous aspect of cyber attacks on manufacturing CPSs—the creation of false confidence in a system’s decision-making that can consistently provide incorrect information about the production status and quality.

After the attack subsided, it was necessary to restore control manually, including restarting the PLC. However, the previous failure of the RFID subsystem disrupted the consistency of production data, making it impossible to resume the original status of work-in-progress products. In some cases, manually moving products between stations or restarting production from the point of failure was necessary, which led to the loss of work-in-progress items and an interruption in production continuity. Figure 14 illustrates the impact of the failure of the CPS identification and communication layer on the production quality evaluation in the higher-level MES.




Figure 14. The status of the superior MES during an experimental DDoS attack (scenario 1), documenting the systematic occurrence of products marked as NOT OK despite the production process being physically correct. The figure shows the discrepancy between the actual production status and the digital representation of products due to the failure of the CPS identification and communication layer, leading to incorrect quality assessment and distortion of production indicators.

During the second experimental scenario, several production system components gradually failed. The first problem was the failure of the IO-Link valve device in the connection station. The loss of communication with the valve terminal meant that the pneumatic operation of placing the cap was not physically performed, even though the production line’s control cycle continued without interruption. This manifested as the IO device’s unavailability at the PLC level. Figure 15 shows the consequence of the failure of

the IO-Link valve device in the connection station, leading to the assembly operation not being executed despite the continuous control cycle.



Figure 15. The physical condition of the product during the experimental DDoS attack (scenario 2), documenting the failure to perform the pneumatic operation of deploying the cap in the connection station. The image shows the workpiece without the cap attached due to a communication loss between the PLC control unit and the IO-Link valve terminal, with the control cycle of the production line continuing without immediate detection of this fault at the CPS control or information layer.

The absence of the assembly operation was not detected because neither the control system nor the superior MES was equipped with a feedback mechanism from the physical process, for example via sensors or a camera system. Therefore, the product without the cap was forwarded to the sorting station of the production line and incorrectly evaluated by the superior system as a compliant piece. At this point, a hidden production quality defect occurred in which the system declared the product to be correctly manufactured, even though the technological operation had not been physically performed.

The second problem was the failure of the PLC control unit responsible for one of the production line stations. The loss of deterministic communication caused the PLC to enter fault mode and interrupt the technological cycle. At this point, the production flow physically stopped, and the workpiece remained blocked at the station, turning a hidden quality defect into an obvious operational problem with a direct impact on production availability.

The third problem was the failure of the RevPi Core IIoT edge device, which collects and processes operational data in the Node-RED environment. The loss of communication caused the visualization and analysis layer to go offline, making it impossible to monitor the production system's current process and status variables. Thus, the operational data ceased to reflect the actual state of production, which significantly reduced the observability of CPS.

The fourth problem was the loss of remote access to the production system via the EWON Cosy+ IIoT router. Consequently, it was not possible to reliably access the PLC control units, perform diagnostics, or intervene in the control system. From an operational and maintenance perspective, the system appeared to be unavailable. The status of the EWON Cosy+ IIoT device and the unavailability of remote service access during the experimental attack are shown in Figure 16.

The same chain of operational problems as in the previous scenarios reappeared on the production line during the third experimental scenario, but with a different communication disruption source. First, the RFID subsystem failed, making it impossible to identify

products or update their digital representation in the MES. The physical production process was initially continued without any obvious disruption.




Figure 16. Status of the EWON Cosy+ IIoT router during an experimental DDoS attack (scenario 2), documenting the loss of remote access to the production system. The image shows the device diagnostics interface with the message “Device unreachable,” indicating the unavailability of the VPN connection and the inability to perform remote diagnostics, monitoring, and service intervention in the PLC control units.

Communication failure with the IO-Link valve terminal at the connection station. The superior MES sent a control instruction via the PLC to perform the assembly operation; however, due to communication failure, this instruction was not physically executed. The MES system did not have a feedback mechanism to detect failure to perform the operation, and the production process continued.

The PLC control unit failed only in the final stage, causing the production process to physically stop and the workpieces to become stuck on the conveyors. The production flow was completely interrupted, and manual intervention was required to resume production.

5. Discussion

Based on the implemented attack scenarios and the subsequent analysis of their impacts, it was possible to identify the behavior of the production line in terms of network communication resilience, as well as potential threats and disruptions to the production process that may arise due to the integration of IIoT devices into industrial infrastructure. The identified shortcomings indicate the need to implement additional security measures to increase the availability, reliability, and security of production line operations. Therefore, the following chapter focuses on the design and implementation of security rules and protective mechanisms to limit the effects of network attacks and reduce the disruption risk to the production process.

5.1. Proposed Solutions to Mitigate the Effects of DDoS and DRDoS Attacks

Implementing extensive changes to network infrastructure in long-term production systems and stable production lines is often technically, organizationally, and economically challenging. Production equipment and control systems are designed for continuous operation, and significant network architecture interventions can lead to unwanted downtime or disruption of technological processes. Therefore, security measures that increase the level of protection with minimal intervention in the existing infrastructure must be prioritized.

One option is to directly configure a software firewall on the network or computing devices (e.g., using the iptables tool). However, such a solution typically increases the

demands on management, maintenance, and staff qualifications and may entail additional operating costs or risks associated with production system interference.

Therefore, a more suitable solution would be to supplement the existing infrastructure with a separate MikroTik RB750Gr3 network device, which allows the implementation of advanced security rules at the network communication level without directly interfering with production components. This approach represents an effective compromise between security level, costs, and production continuity while enabling the secure integration of IIoT devices into the industrial environment.

The proposed security architecture assumes that the production infrastructure operates on a long-term basis and that a fundamental reconstruction or replacement of existing equipment is not feasible. Therefore, the proposed solution is designed as a minimally invasive intervention in the existing network architecture to increase cybersecurity while maintaining the continuity of the production process. The architecture also reflects the findings from experimental attack scenarios, which demonstrated the high vulnerability of the production line and connected IIoT devices to DDoS and DRDoS attacks.

The MikroTik RB750Gr3 device acts as a perimeter security element at the interface between the internal production network and the external WAN or IT infrastructure (Figure 17). This device routed all network communication to and from the production network, enabling centralized control and traffic filtering without interfering with the configuration of individual production components, such as PLCs, I/O modules, or sensors. This location also allows for the logical separation of critical production equipment from IIoT elements used for data collection, remote monitoring, or integration with higher-level information systems.




Figure 17. The MikroTik RB750Gr3 device is located at the interface between the internal production network and the external WAN or IT infrastructure, where it acts as a perimeter security element. This device routed all network communication to and from the production network, enabling centralized control and traffic filtering without interfering with the configuration of individual production components, such as PLCs, I/O modules, or sensors. This placement also allows for the logical separation of critical production equipment from IIoT devices used for data collection, remote monitoring, or integration with higher-level information systems.

In terms of functionality, the MikroTik RB750Gr3 provides network communication filtering through a stateful firewall, restriction of unwanted traffic types, and protection against network congestion using flow control and packet limiting mechanisms. These mechanisms reduce the impact of TCP SYN flood, TCP ACK flood, and ICMP-based DRDoS attacks, which were analyzed in previous chapters. Simultaneously, security rules can be used to precisely define the permitted communication relationships between individual parts of the infrastructure, thereby reducing the risk of lateral spread of attacks within the production network.

5.2. Configuring Firewall Rules on a MikroTik RB750Gr3 Device

Table 5 presents the configuration of firewall rules implemented on the MikroTik RB750Gr3 device to increase the resilience of the production infrastructure against TCP SYN flood, TCP ACK flood, and ICMP-based DRDoS attacks. The rules are applied in the forward chain and focus on protecting communication between the OT network and the external WAN environment, as well as limiting unwanted internal traffic.

Table 5. Overview of firewall rules implemented in MikroTik RouterOS on the RB750Gr3 device to protect the production network against TCP SYN flood, TCP ACK flood, and ICMP-based DRDoS attacks.

Rule Description	Rule Definition (RouterOS)
Allow established and related connections	/ip firewall filter add chain=forward connection-state=established,related action=accept
Drop invalid connections	/ip firewall filter add chain=forward connection-state=invalid action=drop
Drop forwarding by default (last rule)	/ip firewall filter add chain=forward action=drop
Allow DNS (optional, if router provides DNS to LAN)	/ip firewall filter add chain=input protocol=udp dst-port=53 in-interface=bridge action=accept
Allow DHCP (optional, if router provides DHCP)	/ip firewall filter add chain=input protocol=udp dst-port=67,68 in-interface=bridge action=accept
Limit ICMP rate (anti-DRDoS, anti-ICMP flood)	/ip firewall filter add chain=forward protocol=icmp limit=10,5 action=accept
Drop excessive ICMP packets	/ip firewall filter add chain=forward protocol=icmp action=drop
Drop new TCP connections without SYN (anti-ACK flood/anomalous TCP)	/ip firewall filter add chain=forward protocol=tcp connection-state=new tcp-flags=!syn action=drop
Limit new TCP SYN rate (anti-SYN flood)	/ip firewall filter add chain=forward protocol=tcp connection-state=new tcp-flags=syn limit=50,10 action=accept
Drop excessive new TCP SYN packets	/ip firewall filter add chain=forward protocol=tcp connection-state=new tcp-flags=syn action=drop
Drop spoofed RFC1918 source addresses arriving from WAN	/ip firewall filter add chain=forward in-interface=ether1 src-address=10.0.0.0/8 action=drop
Drop spoofed RFC1918 source addresses arriving from WAN	/ip firewall filter add chain=forward in-interface=ether1 src-address=172.16.0.0/12 action=drop
Drop spoofed RFC1918 source addresses arriving from WAN	/ip firewall filter add chain=forward in-interface=ether1 src-address=192.168.0.0/16 action=drop
Drop link-local source addresses arriving from WAN	/ip firewall filter add chain=forward in-interface=ether1 src-address=169.254.0.0/16 action=drop

The configuration allows existing and related connections and immediately discards invalid packets, ensuring the integrity of legitimate communication while reducing the load on the device. Protection against TCP SYN and TCP ACK floods are protected by checking

the TCP connection status and using a rate-limiting mechanism to limit the number of new SYN packets. Packets that do not meet these conditions are blocked.

Protect against DRDoS attacks by limiting ICMP traffic. ICMP packet processing is time-limited, minimizing the risk of flooding target devices with reflective communication generated by compromised IIoT elements. Excessive traffic is automatically blocked.

The configuration also includes rules for detecting and blocking spoofed source IP addresses from external environments, especially from private and link-local ranges, thereby reducing the risk of reflective and spoofing attacks.

The proposed solution represents a cost-effective way to increase production line cybersecurity and connected IIoT devices without interfering with existing production components. It is particularly suitable for long-term industrial infrastructures where maintaining production continuity is a priority.

Figure 18 shows the implementation of the firewall rules listed in Table 5 on a MikroTik RB750Gr3 device deployed in the production line network infrastructure during implementation and verification experiments.

#	Action	Chain	Src. Address	Protocol	In. Inter...	Out. Int...
0	passthrough	forward				
1	Drop invalid connections	forward				
2	Drop RFC1918 spoofing from WAN	forward	10.0.0.0/8		ether1	
3	Drop RFC1918 spoofing from WAN	forward	172.16.0.0/12		ether1	
4	Drop RFC1918 spoofing from WAN	forward	192.168.0.0/16		ether1	
5	Drop link-local spoofing from WAN	forward	169.254.0.0/16		ether1	
6	Drop TCP ACK flood	forward		6 (tcp)		
7	Drop excessive TCP SYN	forward		6 (tcp)		
8	Drop excessive ICMP	forward		1 (icmp)		
9	Drop excessive TCP from LAN	forward		6 (tcp)	bridge	
10	Drop all other forwarded traffic	forward				

Figure 18. Implemented firewall rules in the MikroTik RouterOS environment on the RB750Gr3 device. The rules applied in the forward chain are displayed to limit and block TCP SYN flood, TCP ACK flood, and ICMP flood (DRDoS) attacks, as well as rules for eliminating invalid connections and spoofed traffic (RFC1918 and link-local addresses) coming from the WAN environment. The configuration documents the filtering and segmentation mechanism application in the experimental network architecture.

5.3. Verification of the Proposed Configuration—Verification Scenario

The verification scenario was designed to comprehensively assess the effectiveness of implemented security measures in a production infrastructure environment with an

active MikroTik RB750Gr3 firewall. As shown in Figure 19, the scenario involved the simultaneous execution of two direct DDoS attacks and two reflective DRDoS attacks to verify the ability of the proposed solution to protect the production line from various forms of network threats under conditions close to real-world operation.




Figure 19. Visual illustration of the verification scenario for the proposed security solution, showing the current implementation of direct TCP SYN flood and TCP ACK flood attacks and reflected ICMP-based DRDoS attacks directed at selected components of the production line and IIoT devices through the MikroTik RB750Gr3 perimeter firewall.

Direct attacks were conducted as a TCP ACK flood directed from an external WAN environment to the Ewon Cosy+ device and as a TCP SYN flood initiated from the local network against the IO-Link valve terminal. These scenarios simulate situations in which an attacker deliberately floods production components from an external environment or through a compromised device in the internal infrastructure.

Simultaneously, two reflective DRDoS attacks based on ICMP communication were conducted, in which Siemens SIMATIC IOT2050 and RevPi Core devices were exploited to generate ICMP error messages directed at target infrastructure elements, specifically the Turck RFID system and the PLC 2 control unit.

No complete network communication failure occurred on any of the tested devices during the execution of all four attacks (Figure 20). After an initial increase in traffic caused by initiating attacks, communication stabilized at a significantly lower level due to the firewall rules and traffic restriction mechanisms. The number of packets per second remained within the controlled range throughout the test, with only slight fluctuations recorded and no sudden drop to zero values.




Figure 20. Time course of network communication during the verification scenario with active protection by the MikroTik RB750Gr3 firewall, showing the number of packets per second recorded on selected target devices during the simultaneous execution of direct DDoS attacks of the TCP SYN flood and TCP ACK flood types and reflected ICMP-based DRDoS attacks.

This process confirms that harmful traffic is effectively filtered or restricted before reaching critical production components, thereby preventing communication and computing resource overload. The results also show that the implemented protection mechanisms could maintain the communication stability and operational continuity of the production line even during an ongoing attack. Thus, Figure 20 documents the practical effectiveness of the proposed security architecture in real operating conditions.

The results of the verification scenario confirm that the proposed security measures can effectively mitigate the impact of not only direct DDoS attacks but also reflective DRDoS attacks, in which malicious traffic originates from legitimate IIoT devices within the infrastructure. Maintaining the availability of the production line and the functionality of control components in all tested scenarios demonstrates that a combination of perimeter protection and internal network security rules is an appropriate approach to securing long-term industrial systems with integrated IIoT elements.

The proposed solution offers several practical advantages, including low financial costs, rapid implementation, and no interference with existing production equipment. The architecture also allows for the gradual integration of IIoT devices without significantly increasing the risk to the production process, making it ideal for use in industrial environments with limited modernization options.

5.4. Limitations and Generalization of Results

The proposed approach has some limitations. Protection is primarily implemented at the network boundary, which means that attacks generated directly from the internal network or from a compromised device within the production infrastructure may not be eliminated. The solution's effectiveness also depends on the correct configuration of security rules and regular device management; incorrect settings can restrict legitimate communication or reduce the level of protection.

Security levels can also be increased by implementing a demilitarized zone (DMZ) between the production and external networks or using advanced network segmentation. Such measures can further reduce the exposure of production components to external threats without the need for major interventions in the existing infrastructure.

Experimental verification of the FESTO MPS 203 production line, which represents a modular, didactically oriented industrial environment, was experimentally verified. This is not a fully integrated commercial production system in continuous operation, which is a

limitation of this study. In the case of large-capacity systems, the scope of production, data traffic volume, and economic consequences of outages may vary.

However, in practice, the experimental testing of availability attacks directly on commercial production lines is significantly limited due to economic and operational risks. However, the infrastructure used contained standard industrial components (PLC Siemens SIMATIC S7-1512C-1 PN, PROFINET communication, OPC interfaces, RFID system, and IIoT devices) that are commonly deployed in real production systems. Therefore, the experimental platform represents a relevant model of the industrial environment in terms of communication principles, deterministic control, and network architecture.

Although quantitative parameters (e.g., time to communication degradation) may vary in larger systems due to higher hardware performance or redundant architecture, the network congestion, computational resource depletion, and deterministic communication disruption mechanisms remain fundamentally the same. Therefore, the identified vulnerabilities are transferable to real production systems; the difference lies primarily in the scope of failure and economic impact.

Compared to protocol-oriented security mechanisms, including solutions using deep packet inspection (DPI) for industrial protocols such as PROFINET, the proposed firewall solution represents a cost-effective and minimally invasive risk mitigation strategy. The deployment of DPI mechanisms can provide more detailed control at the application layer, but is typically associated with higher costs, more complex configuration, and potential performance impact in real-time environments.

Therefore, the combination of perimeter firewall protection with future expansion to include protocol-oriented monitoring represents a promising direction for increasing the security of industrial networks.

6. Conclusions

This article analyzes the impact of DDoS and DRDoS attacks on a production line extended with IIoT devices in a real industrial environment. Experimental scenarios conducted from both internal and external networks demonstrated the high vulnerability of production components and IIoT devices to network attacks, leading to rapid communication loss and production process interruption in the absence of protective mechanisms. Reflected DRDoS attacks, in which legitimate IIoT devices were misused to generate malicious network traffic, proved particularly dangerous.

Based on the identified shortcomings, a security solution based on the MikroTik RB750Gr3 device was designed and implemented, which respects the limitations of the long-term operating production infrastructure. The verification of the proposed solution through repeated attacks confirmed that the implemented firewall rules can effectively mitigate the effects of both direct and reflected attacks while maintaining the production line's network communication availability and stability.

The results of this study highlight the importance of combining network access protection and internal network security measures in an Industry 4.0 environment, confirming that even low-cost and minimally invasive solutions can significantly contribute to increasing the resilience of industrial systems against network threats.

Simultaneously, other architectural enhancements, such as the introduction of a demilitarized zone (DMZ) between industrial and external networks, together with the implementation of advanced network segmentation mechanisms, can further reduce the exposure of production components to external threats and represent effective complementary protective measures.

The implementation of advanced network segmentation mechanisms can further reduce the exposure of production components to external threats and represent effective complementary protective measures that may be the subject of further research.

Future research will also focus on the systematic collection and processing of data from the production line's network and control communications during normal operation and controlled cyber attacks to create a representative data set for the security analysis of industrial cyber-physical systems. The dataset will be made available as a reference dataset for further research in the field of industrial CPS security.

The data obtained will also be used to train neural networks, aiming to teach the model to distinguish between normal system operating conditions and conditions affected by cyber attacks based on network traffic dynamics and control unit status data.

Special attention will be paid to the design and implementation of lightweight machine learning models and neural networks intended for deployment in environments with limited computing resources, such as industrial edge devices. The goal is to develop adaptive detection mechanisms capable of operating in real time, complementing network-level protection with intelligent detection directly at the cyber-physical production system level.

Author Contributions: Conceptualization, T.H. and R.R.; methodology, T.H. and R.R.; validation, R.R., T.H. and R.Z.; formal analysis, T.H., R.R., M.C., J.Š. and R.Z.; investigation, T.H., R.R. and R.Z.; resources, T.H., R.R., R.Z., J.Š. and M.C.; writing—original draft preparation, T.H. and R.R.; writing—review and editing, J.Š., R.Z. and M.C.; visualization, T.H., R.R., J.Š., M.C. and R.Z.; supervision, T.H., R.R. and R.Z.; funding acquisition, T.H. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the EU NextGenerationEU through the Recovery and Resilience Plan for Slovakia under the project No. 09I03-03-V05-00005. This work was supported by a call for doctoral students and young researchers of the Slovak University of Technology in Bratislava to start a research career (Project ESG 23-06-01-B: Design and implementation of security of industrial network systems with the creation of a standardized data set for security analyses—NAIZPSS).

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Fatima, Z.; Tanveer, M.H.; Waseemullah; Zardari, S.; Naz, L.F.; Khadim, H.; Ahmed, N.; Tahir, M. Production Plant and Warehouse Automation with IoT and Industry 5.0. *Appl. Sci.* **2022**, *12*, 2053. [CrossRef]
2. Xu, H.; Yu, W.; Griffith, D.; Golmie, N. A Survey on Industrial Internet of Things: A Cyber-Physical Systems Perspective. *IEEE Access* **2018**, *6*, 78238–78259. [CrossRef] [PubMed]
3. Morgan, J.; Halton, M.; Qiao, Y.; Breslin, J.G. Industry 4.0 Smart Reconfigurable Manufacturing Machines. *J. Manuf. Syst.* **2021**, *59*, 481–506. [CrossRef]
4. da Rocha, H.; Abrishambaf, R.; Pereira, J.; Santo, A.E. Integrating the IEEE 1451 and IEC 61499 Standards with the Industrial Internet Reference Architecture. *Sensors* **2022**, *22*, 1495. [CrossRef]
5. Vachálek, J.; Šišmišová, D.; Vašek, P.; Rybář, J.; Slovák, J.; Šimovec, M. Intelligent Dynamic Identification Technique of Industrial Products in a Robotic Workplace. *Sensors* **2021**, *21*, 1797. [CrossRef]
6. Okunlola, O.A.; Olaoye, J.; Samuel, O.O.; Okunlola, A.O.; Alao, O. Cybersecurity Strategies for Integrating Industrial IoT and Edge Computing: Challenges, Risks, and Future Perspectives. *Int. J. Future Eng. Innov.* **2025**, *2*, 87–95. [CrossRef]
7. Krejčí, J.; Babiuch, M.; Suder, J.; Kryš, V.; Bobovský, Z. Internet of Robotic Things: Current Technologies, Challenges, Applications, and Future Research Topics. *Sensors* **2025**, *25*, 765. [CrossRef] [PubMed]
8. Khaung Tin, H.H.; Thu, S.; Maung, K.K. IoT and Industrial Automation: A Review of Current Research and Emerging Trends. *FMDB Trans. Sustain. Technoprise Lett.* **2024**, *2*, 151–160. [CrossRef]
9. Ahmmed, M.S.; Khan, L.; Mahmood, M.A.; Liou, F. Digital Twins, AI, and Cybersecurity in Additive Manufacturing: A Comprehensive Review of Current Trends and Challenges. *Machines* **2025**, *13*, 691. [CrossRef]

10. Bajwa, A.; Tonoy, A.A.R.; Rana, S.; Ahmed, I. CYBERSECURITY IN INDUSTRIAL CONTROL SYSTEMS: A SYSTEMATIC LITERATURE REVIEW ON AI-BASED THREAT DETECTION FOR SCADA AND IOT NETWORKS. *ASRC Procedia Glob. Perspect. Sci. Scholarsh.* **2025**, *1*, 1–15. [CrossRef]
11. Sarjan, H.; Ameli, A.; Ghafouri, M. Cyber-Security of Industrial Internet of Things in Electric Power Systems. *IEEE Access* **2022**, *10*, 92390–92409. [CrossRef]
12. Singh, B.; Kumar, B. Navigating Cybersecurity Risks in Industry 4.0: Challenges, Threats, and Defense Strategies. *J. Inf. Syst. Eng. Manag.* **2025**, *10*, 266–275. [CrossRef]
13. Pal, S.; Jadidi, Z. Analysis of Security Issues and Countermeasures for the Industrial Internet of Things. *Appl. Sci.* **2021**, *11*, 9393. [CrossRef]
14. Alnajim, A.M.; Habib, S.; Islam, M.; Thwin, S.M.; Alotaibi, F. A Comprehensive Survey of Cybersecurity Threats, Attacks, and Effective Countermeasures in Industrial Internet of Things. *Technologies* **2023**, *11*, 161. [CrossRef]
15. Horak, T.; Strelec, P.; Huraj, L.; Tanuska, P.; Vaclavova, A.; Kebisek, M. The Vulnerability of the Production Line Using Industrial IoT Systems under DDoS Attack. *Electronics* **2021**, *10*, 381. [CrossRef]
16. Cui, W.; Liao, X.; Yang, Y.; Feng, S.; Song, M. Informer-Based DDoS Attack Detection Method for the Power Internet of Things. *PLoS ONE* **2025**, *20*, e0322329. [CrossRef]
17. Calderón, D.; Folgado, F.J.; González, I.; Calderón, A.J. Implementation and Experimental Application of Industrial IoT Architecture Using Automation and IoT Hardware/Software. *Sensors* **2024**, *24*, 8074. [CrossRef] [PubMed]
18. Wang, M.; Sun, Y.; Sun, H.; Zhang, B. Security Issues on Industrial Internet of Things: Overview and Challenges. *Computers* **2023**, *12*, 256. [CrossRef]
19. Comert, M.; Ahmed, A.; Ahmed, H. Identifying Security Challenges in the Transition from Traditional to Smart Manufacturing Through IIoT Retrofitting. In Proceedings of the IoT 2024: 14th International Conference on the Internet of Things, Oulu, Finland, 19–22 November 2024; pp. 285–289. [CrossRef]
20. Zheng, R.; Zheng, S.; Liu, C.; Yue, L.; Wu, H. A Software-Defined Gateway Architecture with Graphical Protocol Modeling for Industrial Control Systems. *Electronics* **2025**, *14*, 4369. [CrossRef]
21. Yousuf, M.; Alsuwian, T.; Amin, A.A.; Fareed, S.; Hamza, M. IoT-Based Health Monitoring and Fault Detection of Industrial AC Induction Motor for Efficient Predictive Maintenance. *Meas. Control* **2024**, *57*, 1146–1160. [CrossRef]
22. Sundar, P.S.; Chowdhury, C.; Kamarthi, S. Industrial Internet of Things Enabled Kata Methodology of Assembly Line Productivity Improvement: Insights from a Case Study. *Processes* **2024**, *12*, 2611. [CrossRef]
23. Mashayekhy, Y.; Babaei, A.; Yuan, X.M.; Xue, A. Impact of Internet of Things (IoT) on Inventory Management: A Literature Survey. *Logistics* **2022**, *6*, 33. [CrossRef]
24. Vagaš, M.; Galajdová, A.; Šimšík, D.; Onofreiová, D. Wireless Data Acquisition from Automated Workplaces Based on RFID Technology. *IFAC-PapersOnLine* **2019**, *52*, 299–304. [CrossRef]
25. Kia, S.H.; Dunai, L.; Antonino-Daviu, J.A.; Razik, H. Real-Time Digital Twins for Intelligent Fault Diagnosis and Condition-Based Monitoring of Electrical Machines. *Energies* **2025**, *18*, 4637. [CrossRef]
26. Al Zami, M.B.; Shaon, S.; Quy, V.K.; Nguyen, D.C. Digital Twin in Industries: A Comprehensive Survey. *IEEE Access* **2025**, *13*, 47291–47336. [CrossRef]
27. Purohit, S.; Purohit, S.; Jain, A.K. From Simulation to Strategy: Integrating Generative AI and Digital Twins in the Downstream Oil And Gas Sector. *Econ. Sci.* **2025**, *21*, 104–114. [CrossRef]
28. Qiu, F.; Kumar, A.; Hu, J.; Sharma, P.; Tang, Y.B.; Xu Xiang, Y.; Hong, J. A Review on Integrating IoT, IIoT, and Industry 4.0: A Pathway to Smart Manufacturing and Digital Transformation. *IET Inf. Secur.* **2025**, *2025*, 9275962. [CrossRef]
29. Marinova-Kostova, K.; Kostov, I. Application of Internet of Things in Industry 4.0. *Econ. Ecol. Socium* **2021**, *5*, 49–58. [CrossRef]
30. Kayan, H.; Nunes, M.; Rana, O.; Burnap, P.; Perera, C. Cybersecurity of Industrial Cyber-Physical Systems: A Review. *ACM Comput. Surv.* **2023**, *54*, 229. [CrossRef]
31. Simola, J.; Takala, A.; Lehtonen, R.; Frantti, T.; Savola, R. Impact of Cyber Security Operations on Hardware Requirements for Stable and Workable Industrial Environments. *Int. Conf. Cyber Warf. Secur.* **2024**, *19*, 348–357. [CrossRef]
32. Alqudhaibi, A.; Albarrak, M.; Aloheel, A.; Jagtap, S.; Salonitis, K. Predicting Cybersecurity Threats in Critical Infrastructure for Industry 4.0: A Proactive Approach Based on Attacker Motivations. *Sensors* **2023**, *23*, 4539. [CrossRef]
33. Sechi, F. Critical Convergence for Enhanced Safety: A Literature Review on Integrated Cybersecurity Strategies for Information Technology and Operational Technology Systems within Critical Infrastructure. In Proceedings of the 33rd European Safety and Reliability Conference, (ESREL 2023), Southampton, UK, 3–8 September 2023; pp. 3414–3421. [CrossRef]
34. Al-Obeidat, F.; Amin, A.; Shuhaiber, A.; ul Haq, I. DDoS Attack Detection in Edge-IIoT Digital Twin Environment Using Deep Learning Approach. *PeerJ Comput. Sci.* **2025**, *11*, e3052. [CrossRef]
35. Zakariah, M.; Amin, S.U.; Alrayes, F.S.; Helal, M.; Khan, Z.I. SCADA Intrusion Detection Using Deep Factorization Machines. *Sci. Rep.* **2025**, *15*, 39753. [CrossRef]

36. Alrumaih, T.N.I.; Alenazi, M.J.F. ERINDA: A Novel Framework for Enhancing the Resilience of Industrial Networks against DDoS Attacks with Adaptive Recovery. *Alex. Eng. J.* **2025**, *121*, 248–262. [CrossRef]
37. Orman, A. Cyberattack Detection Systems in Industrial Internet of Things (IIoT) Networks in Big Data Environments. *Appl. Sci.* **2025**, *15*, 3121. [CrossRef]
38. Zhukabayeva, T.; Zholshiyeva, L.; Karabayev, N.; Khan, S.; Alnazzawi, N. Cybersecurity Solutions for Industrial Internet of Things–Edge Computing Integration: Challenges, Threats, and Future Directions. *Sensors* **2025**, *25*, 213. [CrossRef] [PubMed]
39. Hareesh, R.; Kalluri, R.; Mahendra, L.; Kumar, R.K.S.; Bindhumadhava, B.S. Passive Security Monitoring for IEC-60870-5-104 Based SCADA Systems. *Int. J. Ind. Control Syst. Secur.* **2020**, *3*, 90–99. [CrossRef]
40. Eryilmaz, E.E.; Akleylek, S.; Ertek, Y.; Kılıç, E. A Systematic Survey of Machine Learning and Deep Learning Models Used in Industrial Internet of Things Security. *MANAS J. Eng.* **2024**, *12*, 1–28. [CrossRef]
41. Swati; Roy, S.; Singh, J.; Mathew, J. Securing IIoT Systems against DDoS Attacks with Adaptive Moving Target Defense Strategies. *Sci. Rep.* **2025**, *15*, 9558. [CrossRef]
42. Alamri, A.H.; Mooney, L. Dragos Industrial Ransomware Analysis: Q3 2025. Available online: <https://www.dragos.com/blog/dragos-industrial-ransomware-analysis-q3-2025> (accessed on 10 January 2026).
43. Sekonya, N.; Sithungu, S. An Analysis of Critical Cybersecurity Controls for Industrial Control Systems. *Eur. Conf. Cyber Warf. Secur.* **2023**, *22*, 410–419. [CrossRef]
44. de-Marcos, L.; Domínguez-Díaz, A.; Cilleruelo, C.; Rodríguez, D. Ransomware Threats in Manufacturing Industry: Patterns Based on Dark Web Tor Groups. In Proceedings of the 2025 55th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops (DSN-W), Naples, Italy, 23–26 June 2025; pp. 79–82. [CrossRef]
45. Salviotti, G.; Abbatemarco, N.; De Rossi, L.M.; Bjoernland, K. Understanding the Role of Leadership Competencies in Cyber Crisis Management: A Case Study. In Proceedings of the Annual Hawaii International Conference on System Sciences, Maui, HI, USA, 3–6 January 2023; pp. 6068–6077. [CrossRef]
46. Ružarovský, R.; Holubek, R.; Janíček, M.; Velišek, K.; Tirian, G.O. Analysis of the Industry 4.0 Key Elements and Technologies Implementation in the Festo Didactic Educational Systems MPS 203 I4.0. *J. Phys. Conf. Ser.* **2021**, *1781*, 012030. [CrossRef]
47. ISO/IEC 15693-3; Cards and Security Devices for Personal Identification—Contactless Vicinity Objects—Part 3: Anticollision and Transmission Protocol. International Organization for Standardization: Geneva, Switzerland, 2019.
48. Zhou, H.; Pal, S.; Jadidi, Z.; Jolfaei, A. A Fog-Based Security Framework for Large-Scale Industrial Internet of Things Environments. *IEEE Internet Things Mag.* **2023**, *6*, 64–68. [CrossRef]
49. Alotaibi, B. A Survey on Industrial Internet of Things Security: Requirements, Attacks, AI-Based Solutions, and Edge Computing Opportunities. *Sensors* **2023**, *23*, 7470. [CrossRef]
50. Saghezchi, F.B.; Mantas, G.; Violas, M.A.; de Oliveira Duarte, A.M.; Rodriguez, J. Machine Learning for DDoS Attack Detection in Industry 4.0 CPPSs. *Electronics* **2022**, *11*, 602. [CrossRef]
51. Söğüt, E.; Erdem, O.A. A Multi-Model Proposal for Classification and Detection of DDoS Attacks on SCADA Systems. *Appl. Sci.* **2023**, *13*, 5993. [CrossRef]
52. Shin, D.-H.; Kim, G.-Y.; Euom, I.-C. Vulnerabilities of the Open Platform Communication Unified Architecture Protocol in Industrial Internet of Things Operation. *Sensors* **2022**, *22*, 6575. [CrossRef]
53. Kang, H.-S.; Kim, K.; Kim, S.-R. A New Mitigation Method against DRDoS Attacks Using a Snort UDP Module in Low-Specification Fog Computing Environments. *Electronics* **2024**, *13*, 2919. [CrossRef]
54. Nuijaa, R.R.; Manickam, S.; Alsaeedi, A.H. Distributed Reflection Denial of Service Attack: A Critical Review. *Int. J. Electr. Comput. Eng. (IJECE)* **2021**, *11*, 5327–5341. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Article

Quality Management for AI-Generated Self-Adaptive Resource Controllers

Claus Pahl ^{1,*}, Hamid R. Barzegar ¹ and Nabil El Ioini ²¹ Faculty of Engineering, Free University of Bozen-Bolzano, 39100 Bolzano, Italy² School of Computer Science, University of Nottingham, Semenyih 43500, Malaysia

* Correspondence: cpahl@unibz.it

Abstract

Many complex systems requires the use of controllers to allow an automated, self-adaptive management of components and resources. Controllers are software components that observe a system, analyse its quality, and recommend and enact decisions to maintain or improve quality. While controllers have been for many years, recently Artificial Intelligence (AI) techniques such as Machine Learning (ML) and specifically reinforcement learning (RL) are used to construct these controllers, causing uncertainties about the quality of them due to their construction. We investigate quality metrics for RL-constructed software-based controllers that allow for their continuous quality control, which is particularly motivated by increasing automation and also the usage of artificial intelligence and control theoretic solutions for controller construction and operation. We introduce self-adaptation and control principles and define a quality-oriented controller reference architecture for controllers for self-adaptive systems. This forms the basis for the central contribution, a quality analysis metrics framework for controllers themselves.

Keywords: controller; DevOps; machine learning; reinforcement learning; responsible AI; metrics; quality management; self-adaptive system; generative AI; large language models (LLM)

1. Introduction

There is a growing need to automate the continuous operation and management of various systems. We often call these self-adaptive systems if they are built around a controller that manages the system and its resources. A controller generates actions to maintain such a system within expected quality ranges based on monitored system data as the input. Self-adaptive systems are widely used in environments where manual adjustment is neither feasible nor reliable. This ranges from industrial production machines or power management systems to fully softwareised environments such as cloud applications, which are all examples of virtualised or cyber-physical systems that are suited to be governed by a control-theoretic solution to continuously and automatically adjust the system [1–4]. In particular, resources used by a complex system can be dynamically adjusted and need to ensure dedicated quality management.

Today, we can observe that automation of control also extends to its construction. For instance, during development, AI techniques can be employed if a sufficient amount of data is available to generate the models that form the core functions of this self-adaption [5]. Thus, we focus here on the automated construction of controllers and their required quality properties, and not on the application systems that these controllers control.

Performance or robustness are the most frequent quality concerns that can be found for self-adaptive controller-managed systems, i.e., controllers monitor and analyse the target system in terms of the dynamic properties of the latter. Consequently, these quality concepts for the system would need to be reinterpreted for a controller as a software-based system. Our objective here is to conduct a wider review of quality metrics beyond these two, also including fairness, sustainability, and explainability, which are common concerns for AI techniques such as machine learning (ML) in general, as part of a responsible AI concern, but need a specific investigation for applications to the controller context here.

Machine learning (ML) is a good example of a construction mechanism to motivate the need for this investigation. In the ML model construction, accuracy, precision, and recall are generally used to assess the quality of prediction or classification models. However, the construction causes concerns related to the responsible AI concept that aims at the fairness, transparency, explainability, and ethical behaviour of these machine-generated controllers. The quality concerns of reinforcement learning (RL) as a widely used mechanism for controller construction have been investigated for specific concerns [6–8]. Two challenges arise: firstly, the construction outcome looks different for software controllers than for other typical ML applications, and secondly, the quality management of these controllers themselves needs to be automated, what is often referred to as DevOps automation in the software domain [9,10].

In this paper, we address a coherent system of broad principles and practices to design and manage controllers for self-adaptive systems. Specifically, we present a catalogue of classified controller quality metrics as the main contribution to controller construction. To frame this metrics catalogue, we introduce a reference architecture for controllers within a quality management framework. This is often considered as part of a continuous change process in a DevOps-style (or also referred to as AIOps if the software operations process is AI-controlled) that allows for continuous quality monitoring to mediate quality deficiencies. Since change is ubiquitous, we address change in a two-layered quality management framework covering the system and the controller layer.

Our method is based on (i) a review of survey papers as secondary contributions that we systematically selected to cover AI and controller perspectives and (ii) a further consideration of individual primary solution contributions for the specific aspects from machine learning to resource management to change detection in order to supplement the survey findings and define our framework, taking the state-of-the-art into account. A two-layered feedback loop has already been proposed as a natural consequence of mapping quality control into a continuous feedback loop to monitor quality [11,12]. We based this work on the two-layered loop [13] and used the review papers covering their respective focal concern to form a coherent conceptual framework here that provides a comprehensive set of metrics to allow this two-layered controller loop to be realised. While different controller construction mechanisms are covered, due to the popularity and relevance of reinforcement learning and control theory as established mechanisms, these are specifically singled out to provide insights into the potential implementation of the proposed results.

For automatically generated controllers, their automated continued quality management becomes an important concern. While metrics for qualities of systems that a controller manages have been explored, our investigation specifically focuses on AI-related properties of AI-constructed controllers. The novelty lies in the definition of common qualities for AI-generated controller software, rather than the quality metrics that the system under observation is meant to maintain that have been proposed so far. Furthermore, the AI qualities are combined into a comprehensive catalogue with specific definitions for controllers, which has also not been done.

The paper is structured as follows. In Section 2, we review related work. In Section 3, we present a reference architecture as the first contribution component that frames the outcome. Section 4 provides an overview of reinforcement learning and control theory as two prevalent techniques to give a more concrete context. Section 5 then defines a metrics catalogue for controllers. Section 6 discusses the findings and presents a use case. Finally, Section 7 concludes the paper with some future work indications.

2. Related Work

In this section, we discuss the controller quality perspective covering individual metrics but also general quality frameworks. The detection of quality change is a central task.

Traditional system-level quality metrics such as performance or robustness have been widely investigated as the goal of controllers. Performance is interpreted often as latency or the response-time behaviour of a system or robustness during the disturbances caused in the system execution and monitoring. In contrast, we focus here only on the quality metrics of the controller itself. Here, we consider specifically AI techniques for controller construction and will consequently propose a catalogue of AI-related metrics, which has not yet been presented in a comprehensive framework [14–16].

A range of individual quality metrics have been investigated, in particular for the construction of controllers employing reinforcement learning, which is a widely used technique. Reinforcement learning (RL) is a suitable approach to derive solutions for control systems. The work in [7] covers the link between RL performance and the notion of stability that stems from the control area. Ref. [8] is a good example of an RL application for a control problem that requires high degrees of performance. Robustness and performance are covered in [6] to cover recent deep reinforcement learning trends. Robustness is also investigated in [17]. The ability to deal with disturbances from the environment is often seen as an important property of control systems that act in environments with a lot of uncertainty.

However, beyond classical performance metrics, recently in the wider ML and AI context other concerns such as explainability or sustainability have been looked at. Attention has been given to these from the perspective of the environment and the users and/or subjects of a solution. Another concrete direction is the fairness of the solution. Ref. [18] looks at this in the context of Markov processes, which define the central probabilistic behaviour of control systems. While explainability has now been widely recognised for prediction and classification approaches, RL has received less attention. One example is [19], which reviews explainability concerns for RL. A survey of this aspect is provided by [20]. Explanations for self-adaptive systems are addressed in [21,22]. As a wider societal concern that also has a cost impact for users, sustainability through, for example, energy and resource consumption is also investigated [23].

In [24], RL is discussed in the context of a MAPE-K (Monitor, Adapt, Plan and Execute stages based on a shared Knowledge component) loop integration. Meta-learning refers to the need to relearn if adaptations are required to adjust to drift. A meta policy can deal with adaptation to new tasks and environmental conditions. Here, specifically, the cost of RL is considered, since meta-learning should work with small volumes of data and quick convergence.

If a set of metrics needs to be implemented, i.e., needs to be monitored, analysed, and converted into recommendations or remedial actions if quality concerns are detected, then a systematic engineering approach is needed that explains the architecture of the system in question and devises a process for quality management. Ref. [25] provides an overview of the AutoML domain, which is a notion to cover automated approaches to managing the AI model creation and quality management. Another term used in this context is AI

engineering. For instance, Ref. [26] approaches such a generic framework from a software engineering perspective, aiming to define principles that define a systematic engineering approach. Similarly, Ref. [5] investigates common engineering practices and how they change in the presence of ML.

In [27], a three-layered reference architecture based on the MAPE-K pattern is presented that separates dynamic monitoring feedback (for managing context information for the monitoring and adaptation setting), adaptation feedback (to control the adaptive behaviour of the target system in terms of functional and non-functional properties), and objectives feedback (to control changes to the adaptation goals such as settling time or stability).

Model checking of controller models is the concern in [28]. The application system is managed in terms of throughput, resource usage, cost, and safety, as these systems are largely impacted by sensor failure or response time deficiencies. The controller is assessed in terms of liveness and safety properties using the UPPAAL model checker. The solution is embedded in a comprehensive process model for controller construction, where the verification of controller models, e.g., in terms of possible paths of the four MAPE processing steps, is most relevant to our perspective.

Our objective is to monitor quality and to react to changes and unexpected behaviour or uncertainties, which are pointed out in [29] as important concerns. The distinction between drift and anomaly as central forms of unexpected behaviour—and also other concepts such as rare events or novelty detection—is often not clear, as Ref. [30] shows with a review of connected concepts. We will provide formal definitions later after motivating the metrics selection. There we also refer in more detail to concrete techniques. Thus, we introduce only some selected review papers here. Several drift detectors have been proposed, as surveyed by Refs. [31,32]. Examples of drift analysis are found in [33,34]. For instance, Ref. [33] states that different diversity levels in an ensemble of learning machines are required to attain better accuracy, i.e., different classifiers are used. Another example is [34], which uses Spiking Neural Networks for drift detection combined with a notion of evolution to detect changes. Equally, for anomaly detection, the number of proposed solutions is large, as the review in [35] shows. We will later introduce anomaly detection techniques specific to our context of continuous sensor data in a technical systems setting.

Another layered architecture is presented in [36], where the lower layer as usual deals with the managed system adaptation. The upper layer, called the lifelong learning loop, focuses on changes in ML-generated controllers. In particular, concept drift is addressed here as a cause of the needed adaption of the controller to changing requirements. The solution is knowledge management, e.g., encoding unlabelled data and detecting new tasks to deal with changes.

In summary, we can note that many controller quality frameworks exist, but that a comprehensive analysis of AI properties has not been carried out. Reinforcement learning is the prevalent ML technique used—which is the reason why we focussed our investigation of this construction technique. However, the reviewed surveys and other papers show that only aspects such as robustness or performance are addressed for controller constructions. Explainability or fairness are not well covered for RL-approaches in general and are only investigated for more general AI-constructed software contexts. The quality of RL-constructed controllers is specifically a research gap. These latter quality concerns in particular need a formalisation, which we will provide later on in Section 5 after introducing a reference architecture and RL basic to frame these concerns in Sections 3 and 4.

3. Reference Architecture for Controller Quality

Our focus is self-adaptive systems, specifically the construction and architecture of controllers that deal with quality problems and adapt the resources used by the system.

3.1. State-of-the-Art and Motivation of Framework

The review of the state-of-the-art results in two concerns that have not been sufficiently addressed:

- Firstly, relevant quality metrics are performance, robustness, fairness, explainability and sustainability, but these need to be adjusted to AI-constructed controllers;
- Secondly, a systematic quality management framework is needed to embed these metrics into a coherent framework.

A respective framework consisting of metrics and a quality management process is needed to provide a systematic, quality-driven controller construction and management approach. The presented work provides core principles and practices in this context. The need to automate quality management of control is exacerbated by the fact that this context is about environments where a manual adjustment is neither feasible nor reliable.

Therefore, we need to identify and define quality criteria for the automation of system adaptation where AI is used to construct controllers.

3.2. Controller Quality Requirements

Run-time monitoring of critical systems is always important. Due to the utilisation of ML to construct software, there is no direct control by expert development engineers, and thus quality needs to be managed differently. This requires, for instance, explainability of the ML models to understand quality implications. Machine learning models are normally evaluated for their effectiveness, generally in terms of metrics such as accuracy, precision, and recall that capture the performance. Two challenges emerge for controllers and their quality:

- The usual ML performance measures of accuracy, precision, and recall do not naturally apply for all construction methods, requiring a different notion of performance and also the need to take uncertainty and disturbances in the form of robustness in the environment into account;
- To better judge the quality of controllers themselves, other concerns such as explainability, but also fairness or sustainability, are important.

What is effectively needed is an engineering approach for self-adaptation controllers, which is part of what is often called AI Engineering [26]. We investigate this specifically for controller construction for self-adaptive software systems.

3.3. Resource Management and Motivational Example

We now return to ML and prediction and classification model accuracy examples and translate this to reinforcement learning and control theory and relevant quality concerns. A motivational use case is controllers for resource management [37–41] that can be applied in virtualised and cyber-physical environments, e.g., cloud or IoT settings. Hong et al. [42] review this for resource management. System adaption is required for the resource configuration, including the monitoring of resource utilisation and respective application performance and the application of generated rules for resource management to meet quality requirements [12,43,44]. The rules adjust the resource configuration (e.g., size) to improve performance and other qualities. A prevalent example is RL, which employs a reward principle applicable in the self-adaptation loop to reward improvements and penalise deteriorations.

The problem shall be described using a concrete example:

- The problem could be a resource controller for cloud adaptation. This could employ the following rule: if $\text{Workload} > 80\%$ then double (Resource-Size). The concrete

- problem is whether this rule is optimal and whether the recommendation space considered by the controller is complete.
- The remediation can be an RL-generated controller that provides an actionable recommendation for a 60% workload as a verified test case. Then, this recommendation can be scaled up to 80%.
 - The model could reward a good utilisation rate (e.g., 60–80%);
 - The model could penalise costly resource consumption, e.g., high monetary costs for cloud resources or energy consumption.

Performance in this case would be determined the accuracy of how well the controllers maintain a 60–80% workload, or robustness would be determined by to what extent noise in the training data affects the controllers' behaviour.

In particular, AI-driven controller generation requires proper quality monitoring and analysis of a defined set of comprehensive quality criteria. Furthermore, detected quality deficiencies need to be analysed as part of a root cause analysis. From this, suitable remedies need to be recommended and enacted.

3.4. Reference Architecture for Quality Management

The quality management of controller models is the challenge here. The accuracy or effectiveness against the real world could mean that a predictor predicts accurately or that an adaptor manages to effectively improve the system's performance. However, as the above discussion shows, more than the traditional performance and cost improvement needs to be addressed. While discrimination is not a direct issue for software, a notion of technical fairness and aspects of accountability and explainability need to be dealt with.

Self-adaptive systems and the decision models at the core of respective controllers are suitable for AI-based creation due to the availability of monitoring and actuation data for training. The controller model aims to enact an action, e.g., to adapt resources in the edge, divert traffic in IoT settings, or instruct the machines that are controlled. This implements a dynamic control loop, governed by predefined quality goals [10]. These self-adaptive systems governed by a controller implement a feedback loop, see Figure 1, following the MAPE-K architecture pattern. MAPE-K control loops provide a template for self-managing systems [45]. The managed element is often part of a cyber-physical environment. The MAPE-K loop consists of four phases in the form of a continuous loop:

- Monitor: The environment with the managed element is continuously monitored. The observed data can cover various data formats about the physical and digital environment. Usually, some form of aggregation or filtering happens.
- Analyse: The current system situation based on the monitored observations is analysed. The reasoning can derive higher-level information that might result in a change or management request.
- Plan: Based on a change request, a plan with corresponding management actions is created. The aim is actions that allow the managed element to fulfil its goals and objectives again if concerns have been identified. This can include basic decisions such as an undo or redo, but also complex strategies including structural modification and migration of the underlying system.
- Execute: A change plan would be executed by an actuator in the operational environment of the managed element.
- Knowledge as a key part of the control loop is shared between all phases. Relevant knowledge for these autonomic systems can range from topology information to historical logs to metrics definitions to management policies. This knowledge, or model, can also be continuously updated.




Figure 1. Base system architecture—a controller loop based on the MAPE-K pattern with (M)onitor, (A)dapt, (P)lan, and (E)xecute stages based on a shared (K)nowledge component.

We propose a two-layered MAPE-K-based controller reference architecture, where the knowledge component K (the ‘model’) is the item of concern for the second layer, together with the full MAPE-K actions. The system is the item of concern for the first layer. The reference architecture provides a meta-learning layer that monitors several quality metrics, but also validates the metrics and their respective success criteria in a second loop—see Figure 2. The lower layer (Layer 1) is a controller for self-adaptive systems. The upper layer (Layer 2) is an intelligent quality monitoring and analysis framework aligned with the requirements of a (generated) controller for self-adaptive systems. Figure 2 builds on the MAPE-K loop in these two layers. The upper loop is the focus of this paper, but needs to take on board the lower layer’s behaviour.




Figure 2. Controller reference architecture—with lower system monitoring loop and upper controller quality management loop.

We define controllers as a state–transition system of states S , actions A , a transition function $T : S \times A \rightarrow S$ and a reward function $R : S \times A \rightarrow \mathbb{R}$ that allows us to measure the quality of the chosen action in a given state—see Algorithm 1.

Algorithm 1 Controller structure—definition

States $S = \{1, \dots, n_s\}$
 Actions $A = \{1, \dots, n_a\}$, $A : S \Rightarrow A$
 Probabilistic transition function $T : S \times A \rightarrow S$
 Reward function $R : S \times A \rightarrow \mathbb{R}$

3.5. Controller Quality Management at the Upper Layer

An important question concerns the automation of the upper loop. If full automation is the ultimate objective, this creates some challenges regarding metrics and their measurability and should thus guide the solution. These challenges are as follows: automate model construction, automate the model use and evaluation, and recommend a model learning configuration adjustment. The upper meta-learning loop design follows the MAPE-K pattern:

- **Monitor:** We need reward functions for model quality, i.e., adaptor quality based on a range of metrics that are linked to the application system and its aspects—from data quality in the input (requiring robustness) [46,47] to sustainability (requiring lower cost and environment damage) as two examples.
- **Analyse:** We need to implement an analysis consisting of problem identification and root cause analysis for model quality problems and feeding into explainability aspects through a (partial) dependency determination to identify which of the system factors would improve the targeted quality the most.
- **Plan/Execute:** The solution would need to recommend and if possible also enact these recommendations. For example, a rule update for the cloud adaptor could be recommended, with model recreation being done. This could in very concrete terms be a readjustment of training data size or ratio. Also, remedial actions for an identified root cause can be considered.

This upper loop would implement a meta-learning process, which is a learning process to adapt the controller through a continuous testing and experimentation process. We call this the knowledge-learning layer.

In order to develop a solution, we follow the Engineering Research guidelines proposed by the ACM SIGSOFT Empirical Standards committee as the methodological framework [48]. In [10], a systematic construction process based on the identification of goals and actuation needs, resulting in suitable models and the final controller, is discussed. To frame the controller design, we use the MAPE-K architecture pattern for system adaption.

4. Reinforcement Learning Principles

Reinforcement learning is the prevalent technique for advanced controller construction reported in the literature. In this section, we introduce the relevant RL background as the technology to which our quality metrics framework for the specific self-adaptive-systems context shall be applied.

4.1. Reinforcement Learning and Control

We assume the environment of these control systems to be non-deterministic or probabilistic, i.e., a controller would need, for instance, a probabilistic policy based on Markov decision processes to control the states of a system [49].

If reinforcement learning (RL) is used, generally, a model-free variant of RL is used for this, i.e., no prior model is used and the training works on previously collected data. Models help to be sample-efficient but do not always exist, as in our case. Some RL techniques that are typically used are SARSA or Q-learning [50]. RL has the notion of a value function at the core to assess a given state and proposed action and enact a policy. This assessment is expressed in terms of a reward. Finding a new value function can be achieved in a number of ways. The Monte Carlo method is a trial-and-error approach that works for systems with a definite end (e.g., games), but less so for non-terminating control systems such as ours. Here, generally, all states are considered equal. Temporal distance (TD) changes this by considering the rewards of recent activities more. States in a farther distance are thus reward-discounted. Q-learning and SARSA are currently the most widely

used temporal distance approaches that combine policy and value function calculation into a single quality function. RL is a form of semi-supervised learning. In supervised learning, labels (i.e., whether a decision or action was right or wrong) are available at all states. In semi-supervised systems, there is typically a (sparse) reward of actions much later.

RL is often applied at the intersection of AI and control theory, as the overviews in [50,51] show. We briefly point out differences. The AI perspective focuses on performance and related qualities of the generated models, and control focuses on the stability of the system, which largely means that bounded inputs should result in bounded outputs, which is the key property of the system and results in a stabilising system from an initially unstable one. Performance is measured in terms of rewards. However, a joint perspective can use rewards to represent stability requirements (as in MPC model predictive control approaches) [7,49]. Rewards are assumed to be bounded, but unstable systems could be governed by arbitrarily negative (or unbounded) rewards or penalties.

4.2. Reward and Quality

The RL objective is to maximise the reward that is calculated for each state of the system. An important assessment factor in this process is the value of being in a state s . For this, the expected future reward of a policy is evaluated using a value function. Positive and negative assessments can be used: rewards, e.g., for achieving performance objectives, and penalties, e.g., for high costs or consumption. The quality of the approach is then measured, typically by the optimality of the model and time of convergence.

The policy is adjusted to improve performance. Policy optimisation is based on a mix of exploitation and exploration, i.e., mixing the exploitation of previous knowledge and random exploration. In contexts such as IoT or edge computing, in addition to classical performance, robustness against disturbances in the environment is also of importance. We have argued that fairness is also important and can actually be seen as contributing to the overall performance. We will provide a respective definition below that clarifies this. As indicated, sustainability and explainability have more of an impact on the context of the system in question, but can of course also be rewarded or penalised if automated observation and assessment is possible (as for energy consumption as a sustainability criterion).

This multitude of metrics also distinguishes an RL setting from the typical control-theoretic focus on stability.

4.3. SARSA and Q-Learning

The two most widely used RL algorithms are Q-learning and SARSA. Both SARSA and Q-learning use Temporal Difference (TD) prediction methods for the control problem. Both learn an action-value function by estimating a quality function $Q(s, a)$ for a state s and an action a . The Q-value or quality function is updated after every time step. SARSA and Q-learning both use the epsilon-greedy policy, i.e., choosing between exploration and exploitation randomly. We explain the two in more detail, focusing on the differences now.

Q-learning is a so-called off-policy TD algorithm to find an optimal policy by updating the state–action value function (Q) for every step using the Bellman optimality equation until the Q-function converges to the optimal Q-function $Q : X \times A \rightarrow R$. The core definitions for Q-learning and SARSA are given in Algorithm 2. For Q-learning, the quality function $Q : S \times A \rightarrow R$ is defined by Algorithm 3.

SARSA stands for State, Action, Reward, (Next)State, (Next)Action. SARSA is an on-policy TD algorithm that aims to find the optimal policy by updating the state–action value function Q at every state using the Bellman equation. SARSA learns by experiencing the environment and updating the state–action value at every time step. There is only one

difference to Q-learning in the calculation of Q , where the maximisation is not applied to $Q(s', a')$, see Algorithm 4.

Algorithm 2 Core definitions for Q-learning and SARSA

States $S = \{1, \dots, n_s\}$
 Actions $A = \{1, \dots, n_a\}$, $A : S \Rightarrow A$
 Reward function $R : S \times A \rightarrow \mathbb{R}$
 Probabilistic transition function $T : S \times A \rightarrow S$
 Learning rate $\alpha \in [0, 1]$ ▷ typically $\alpha = 0.1$
 Discount factor $\gamma \in [0, 1]$

Algorithm 3 Q-learning

procedure Q-LEARNING($S, A, R, T, \alpha, \gamma$)
 Initialise $Q : S \times A \rightarrow \mathbb{R}$ arbitrarily
while Q has not converged **do**
 Start in state $s \in S$
while s is not terminal **do**
 Calculate π according to Q and the chosen exploration strategy
 $a \leftarrow \pi(s)$
 $r \leftarrow R(s, a)$ ▷ Receive the reward
 $s' \leftarrow T(s, a)$ ▷ Receive the new state
 $Q(s', a) \leftarrow (1 - \alpha) \cdot Q(s, a) + \alpha \cdot (r + \gamma \cdot \max_{a'} Q(s', a'))$ ▷ Q-learning value
 function $s \leftarrow s'$
end while
end while
return Q
end procedure

Algorithm 4 SARSA—difference to Q-learning

$Q(s', a) \leftarrow (1 - \alpha) \cdot Q(s, a) + \alpha \cdot (r + \gamma \cdot Q(s', a'))$ ▷ SARSA value function

5. Controller Quality Metrics

The aim is to automatically monitor and analyse the quality of the system controller for the lower layer, as far as possible, using a defined set of metrics. This metrics set takes the controller construction into account. Our selection and definition rationale below is given for an RL-constructed controller type. In this section, we introduce our metrics framework, starting with a conceptual outline, before defining each of the five selected metrics in more detail.

5.1. Method

Two activities are needed for this metrics framework: the selection of suitable metric categories and the precise definition of the selected metrics.

For the first step, we systematically surveyed review articles covering common bibliographic databases such as Scopus, Google Scholar, and IEEE Xplore. We used suitable search terms both for AI-generated software in general (AI, ML, and RL as acronyms and in their extended term variant) and specific controller-oriented ones (controller, self-adaptation, DevOps, AIOps, and also quality) to identify publications with relevant metrics. The surveys included here were selected based on publication quality (considering only peer-reviewed conferences and journals) and relevance (assessing the extent of recent AI technology coverage).

For the second step, again the review papers served as a baseline. Here, however, metric definitions were either selected among options or were adapted to meet the needs of controllers as a specific type of AI-generated software model, in contrast to a wider coverage in the first step.

5.2. Metrics—Selection and Catalogue

In Figure 3, the selected quality metrics are indicated. Performance and fairness directly affect the system quality. Robustness is a guard against external uncertainties and influences. Sustainability has an effect on the environment. Explainability allows understandability, e.g., to provide a governance assurance on the actual controller function.




Figure 3. Layered controller quality architecture—with system quality and controller quality management in MAPE-K format.

We can classify the metrics based on whether they relate to the control task at hand or affect the resources that this task consumes. For the task, we also indicate whether the control function is directly affected (core), it could be influenced by its direct context (environment), it could skew the results by favouring certain outcomes (bias), it could be related to the responsibility for the task (governance), or it could have an impact on resources (such as energy as one selected concern). The properties are those of a controller, not the application system.

- Performance addresses the core quality management function of a controller;
- Robustness relates to the input of monitoring the environment and subsequent MAPE phases;
- Explainability concerns the governance of the control and quality framework in order to allow for trustworthy behaviour;
- Fairness refers to the avoidance of bias in controller decisions and the effect on controller or system performance;
- Sustainability relates to the effect of controller actions on the consumption of resources in the environment in terms of cost and energy.

Both positive and negative measurements can be valued, e.g., by rewarding or penalising them, respectively. We measure concerns that directly influence the task at hand, i.e., how well the solution can perform its job. In a second category of quality targets, the environment is addressed. This includes resources and their consumption, e.g., in terms of

energy consumption, but also the human or organisation in charge of the system in a governance concern, e.g., in terms of explainability. We also add an impact direction, i.e., whether the concern is internal, influenced by external forces (inwards through disturbances), or influences external aspects (outwards on parts of the environment). This is indicated in Table 1.

Table 1. Controller quality metrics and their context.

Quality Metric	Impact	Target
Performance	Internal	System/Controller
Robustness	Inward	System/Controller
Fairness	Internal	System/Controller
Sustainability	Outward	Environment (Resource Consumption)
Explainability	Outward	Environment (Owner, Responsible)
Responsible AI	Outward	Environment (Owner, Accountability)

We add another metric at a meta-level that deals with the encompassing concern of responsible AI in order to adjust and tailor the above foundational metrics.

- Responsible AI serves as a meta-level aspect that allows us to prioritise and balance the direct quality metrics (performance, robustness, explainability, fairness, sustainability).

All of these metrics will be individually motivated and defined in the next subsections in the following format covering a number of perspectives:

- Justification: why the metric is needed;
- Assumptions: contextual assumptions;
- Rationale: motivation of the definition;
- Definition: defined metric;
- Example: illustration of metric.

5.3. Performance

Justification: Generally, the performance of a controller refers to the ability to effectively achieve (and optimise) a technical objective. This could be expressed in terms of a function on system performance and cost, e.g., $F(cost, performance)$, with $\min(cost)$ while $\text{limit}(responseTime)$ as a sample goal. While this works for many monitored application systems, we need here a definition that in more concrete terms takes the objectives and the construction of the controller into account.

Assumptions: We assume that the quality of a controller action in a given state can be quantified through a reward function R . The performance for controllers is the overall achievement of the model towards an optimal reward. We look at reinforcement learning as the prevalent construction technique.

Rationale: The reward is the central guiding performance measure of RL-based controllers, as discussed in the previous section. Reward optimisation is built into approaches like SARSA or Q-learning to optimise the reward [6,50]. The performance of a reinforcement learning algorithm can be determined by defining the cumulative reward as a function of the number of learning steps. Better rewards are better performance. Thus, this concept is the basis of the performance definition here.

Definition: Different performances emerge depending on the chosen learning rate α for the Q-function (which influences to what extent newly acquired information overrides old information). Three different parameters are important for the performance evaluation in a system where the reward varies over time—see Figure 4 for an illustration that shows an eventual decrease and increase while the controller is learning. This time-varying performance can be analysed with three parameters:

- **Convergence:** the asymptotic slope of a performance graph indicates the quality of the policy after stabilisation of the RL algorithm at the end;
- **Initial Loss:** the lowest point of the curve (minimum) indicates how much reward is often sacrificed before the performance begins to improve;
- **Return on Investment:** the zero crossing after the initial loss gives an indication of recovery time, i.e., of how long it takes to recover from initial, often unavoidable learning costs.




Figure 4. Sample RL performance in terms of reward over the system steps [adopted from <https://artint.info/2e/html/ArtInt2e.Ch12.S6.html> (accessed on 1 December 2025)].

Note that the second and third cases only apply if there are positive and negative rewards. Also, note that the cumulative reward is a measure of the total reward, but algorithms such as Q-learning or SARSA use discounted rewards, modelled using the discount factor γ . A flattened graph would indicate that the learning process has finished with a defined policy.

A slightly different perspective emerges for control theory-based controllers. Control theory aims for stability of the systems, i.e., the signal (output action of a controller) should be bounded and minimise fluctuations of the technique system objective, defined in terms of SLA metrics. Notions of convergence, loss and return on investment are applicable to assess the achievement of stability. A reward function could consider these parameters.

Another possibility to define performance is, instead of referring to accumulated rewards, measuring the average reward. This would be a measure of the quality of the learned policy.

Example: The metric is connected to the performance of the learning process. A positive impact on the convergence and loss criteria can be generated by using expert models in the training. For instance, a fuzzy logic rule system has been successfully used for this purpose [52].

5.4. Robustness

Justification: Robustness is the ability to accept, i.e., deal with a wide range of monitored input cases. This includes for instance uncertainties, noise, non-deterministic behaviour, and other disturbances. These are typical for cyber-physical and distributed systems where sensors, connections, or computation can fail in different locations. Robustness concerns arise in non-deterministic behaviour situations and needs repeated

experiments in the evaluation. When controllers are constructed with disturbances present, robustness against these is required.

Assumptions: We use the term disturbances to capture the multitude of external factors [17]. Disturbances can be classified into three possible contexts: observations, actions, and dynamics of the environment that the controller interacts with, thus giving the following disturbance classes:

- Observation disturbances happen when the observers (e.g., sensors) used cannot detect the exact state of the system;
- Action disturbances happen when the actuation of the system ultimately is not the same as the one specified by the control output, thus causing a difference between actual and expected action;
- External dynamics disturbances are disturbances that are applied directly to the system. These are environmental factors or external forces.

Rationale: Disturbances can be classified into different behavioural patterns for the observed system quality over time—see Figure 5. The patterns are important, as the evaluation of a controller’s robustness is often carried out using simulations based on disturbances being injected into the system following these patterns. Figure 5 shows the five patterns in two broader categories—three non-periodic and two periodic ones—following the classification in [17].




Figure 5. Disturbances in the form of five patterns, adopted from [17], with three non-periodic ones on the left and two periodic ones on the right.

Definition: The patterns are important to identify disturbance types in a robustness metric. Thus, we define them first.

Non-periodic disturbance patterns are the following:

- **White Noise Disturbances:** These indicate the natural stochastic noise that an agent encounters in real world situations. Noise can be applied, ranging from zero with increasing values of standard deviation.
- **Step Disturbances:** These model a system change with one sudden and sustained change. The magnitude of the step can vary.
- **Impulse Disturbances:** These show system behaviour with a sudden, but only very short temporary change. The impulse magnitude can vary, as above.

Periodic disturbance patterns are the following:

- **Saw Wave Disturbances:** these are cyclic waves that increase linearly to a given magnitude and instantaneously drop back to a starting point in a repeated way. Thus, this pattern combines characteristics of the step and impulse disturbances but is here in contrast applied periodically.
- **Triangle Wave Disturbances:** These are also cyclic waves that, as above, repeatedly increase linearly to a given magnitude and decrease at the same rate to a starting point (and not suddenly as above). This is very similar to the saw wave, but exhibits a more sine-wave-like behaviour.

We can quantify robustness by relating the performance metrics (as above in the ‘Performance’ section) as a ratio between an ideal and a disturbed setting:

$$\frac{Performance_{Disturbed}}{Performance_{Ideal}}$$

This can be calculated for all disturbance patterns. The closer the ratio is to 1, the more robust the controller is.

Example: White noise is often caused in distributed environments such as edge clouds due to outages and anomalies in the system operation and monitoring. Here, the training needs to consider these disturbances. Another example is extended, hardly noticeable slopes that might be caused by drift in the application.

5.5. Explainability

Justification: Explainability is important, in general, for AI, in order to improve the trustworthiness of the solutions. For technical settings such as controllers, explainability is beneficial since it could aid a root cause analysis for quality deficiencies. The explainability of the controller actions is a critical factor in mapping observed model deficiencies back to system-level properties via the monitored input data to the model creation.

Explainability is a meta-quality aiding to improve the controller quality assessment. Since how and why ML algorithms create their models is generally not always obvious, a notion of explainability can help to understand deficiencies in the other four criteria and remedy them.

Assumptions: Explainability differs depending on the construction mechanism, being typically more challenging for ML-constructed controllers than for fuzzy or probability-based ones where the construction is manual on an explicitly defined function or theory. For example, if fuzzy logic is applied, the deciding function is the membership function. Explainability for RL is less mature than for other ML approaches.

Rationale: While explainability is a very broad concern, we introduce here definitions and taxonomies that are relevant for a technical controller setting and allow us to define basic metrics based on observations that can be obtained in the controller construction and deployment. The metrics defined can be used to automatically create indicators or recommendations. Here, however, a full automation of a reconfiguration or other correction would be difficult.

Definition: A number of taxonomies for ML explanations have been proposed in recent years. We focus here on [20] to propose a classification of explainability into three types. Note that we extend these three classes by concrete measurements for our setting.

- Feature importance (FI) explanations: These identify features (or metrics in our case) that have an effect on an action a proposed by a controller for an input state s . FI explanations provide an action-level perspective of the controller. For each action, the immediate situation causing an action selection is considered.

Assume system metrics $m_1, \dots, m_k$. Then $m_i (i : 1..k)$ has an effect if $\Delta_{m_i}(Q(s, a)) > \Delta_{m_j}(Q(s, a))$ for all $j \neq i$, meaning that the quality difference $\Delta_{m_i}(Q(s, a))$ between two subsequent states from state s with action a is the greatest for metric (or feature) m_i .

- Learning process and MDP (LPM) explanations: These show past experiences or the components of the Markov decision process (MDP) that have led to the current controller behaviour. LPM explanations provide information about the effects of the training or the MDP, e.g., how the controller handles the rewards.

The effect can be measured in terms of the quality function for two controllers (or underlying MDP) instances Q^1 and Q^2 with fixed data for training and testing, i.e., Q^1 is better than Q^2 , if $Q^1(s, a) > Q^2(s, a)$ for selected system states s and actions a .

- Policy-level (PL) explanations: These show the long-term behaviour of the controller as caused by its policy. They are used to evaluate the overall competency of the controller to achieve its objectives.

The PL explainability definition follows LPM measurement ‘is better than’, but here consider variations in data sets, even in terms of features or metrics used.

Other explainability taxonomies also exist. For instance, [19] distinguishes two types in terms of their temporal scope:

- Reactive explanations focus on the immediate moment, i.e., only consider a short time horizon and momentary information, such as direct state changes of the controller, as in FI;
- Proactive explanations focus on longer-term consequences, thus considering information about an anticipated future and changes, such as those that LPM and PL represent.

A reactive explanation provides an answer to the question “what happened?”. A proactive explanation answers the question “why has something happened?”. These can then be further classified in terms of how an explanation was created. Beyond the simple measurements that can serve as indicators of the explanation, more complex explanatory models could be constructed. Reactive explanations could be policy simplifications, reward decompositions, or feature contributions and visual methods. Proactive explanations could be structural causal models, explanations in terms of consequences, hierarchical policies, and relational reinforcement learning solutions. More research is needed here.

Example: In a concrete situation, it might need to be determined of some lower layer metric (such as cost over performance) or some strategy (for instance for storage management) is favoured. A remediation could use explicit weightings to level out disbalances.

5.6. Fairness

Justification: Specifically where humans are involved, the fairness of decisions becomes crucial, and any bias towards or against specific groups needs to be identified. This concern can also be transferred from the human viewpoint to the technical domain, creating a notion of technical fairness that avoids preferences that could be given to specific technical settings without a reason. An example of an unfair controller would be a bias toward a specific management strategy. For instance, historically an overprovisioning strategy could have been followed, without that having been objectively justified.

Assumptions: We assume that, as for humans, past technical behaviour and decisions can be detrimental.

Rationale: The definition we propose there was originally referring to states and the quality of actions in these for RL, but can be adapted to all controllers that follow our definition, as there is always a state in which an action is executed, which in turn can be evaluated. This ensures that the long-term reward of a chosen action a is greater than that of a' , and there is no bias that would lead to a selection not guided by optimal performance. Fair controllers must implement the distribution of actions, with a somewhat heavier weight put on better-performing actions judged in terms of (possibly discounted) long-term reward. Actions cannot be suggested without having a positive effect on the objective performance of the system as defined above.

Definition: Fairness can be defined in a precise way. We follow the definition given by [18] if a quality function Q or reward function R is defined:

- A policy is fair if in a given state s a controller does not choose a possible action a with probability higher than another action a' unless its quality is better, i.e., $Q(s, a) > Q(s, a')$.

The above definition is often referred to as exact fairness, i.e., quality measured as the potential long-term discounted reward. Possible alternatives shall also be introduced. Approximate-choice fairness requires a controller to never choose a worse action with a probability substantially higher than that of a better action. Approximate-action fairness requires a controller to never favour an action of substantially lower quality than that of a better action.

A number of fairness remedies that are known in the ML domain include the improvement of data labelling through so-called protected attributes. An example is the automation of critical situation assessment, where technical bias could treat specific failures or failure management strategies differently. Here, for instance, a high risk of failure based on past experience might be considered, which could have a probability of discrimination based on certain events that have occurred and could be biased against or towards these, be that through pre-processing (before ML training) or in-processing (while training). The challenges are to find bias and remove this bias through a control loop, e.g., using favourable data labels as protected attributes to manage fairness. Examples in the edge controller setting are if smaller or bigger device clusters could be favoured wrongly or if specific types of recommended topologies or recommended configuration sizes (messages, storage, etc.) exist.

Example: Often, underprovisioning in the form of a static resource allocation happens in manual systems due to unawareness of scalability options or the aim for simplicity. In case of high demand, this would slow down the system and possibly disadvantage users. Here, favourable labels reflecting user impact could be identified.

5.7. Sustainability

Justification: General economic and ecological sustainability goals are important societal concerns that also should find their application in computing, here specifically in terms of cost- and energy-efficiency of the model creation and model-based decision processes within a controller.

Assumptions: Sustainability is often used synonymously with environmentally sustainable, e.g., in terms of lower carbon emissions [23]. Thus, we focus on environmental sustainability here.

Rationale: While different measures can be proposed in this environmental sustainability context, we choose energy consumption here as the core of our sustainability definition because it can often be determined in computing environments.

Definition: Energy-efficiency can be defined through widely accepted and measurable metrics. We present two such metrics with examples of common tools to measure them:

- Energy consumption in KJoule per task ($KJ/task$), which can be determined using monitoring tools such as the NVIDIA System Management Interface SMI (NVIDIA System Management Interface (NVIDIA-SMI) is a command line utility for the management and monitoring of NVIDIA GPU devices—<https://developer.nvidia.com/nvidia-system-management-interface> (accessed on 1 December 2025)) in the case of respective GPUs;
- CPU/GPU usage in percent %, which can also be determined using NVIDIA-SMI or similar tools.

These metrics are often put into comparison with performance metrics to assess performance in terms of costs. Similarly to the robustness case, a ratio allows us to indicate a possible trade-off between performance and sustainability to define sustainability. We can relate the performance to the cost or resource consumption it causes:

$$\frac{\text{Performance}}{\text{Resource Consumption}}$$

This can be applied to various resource or cost types. Sustainability focusing on the consumption of resources is often considered and measured through penalties in the value or quality calculation.

Example: Overprovisioning is an effective resource allocation strategy from a system performance perspective but generally causes the highest energy consumption overhead. Explicit penalisation of sustainability defects needs to be introduced.

5.8. Responsible AI (Meta-Level)

Responsible AI serves as a meta-level concern to prioritise and balance the five core quality metrics of performance, robustness, explainability, fairness, and sustainability that were defined above.

Justification: Responsible AI acts as a guiding principle, ensuring that the implementation and configuration of AI-based controllers align with societal, ethical, and operational goals.

Assumptions: Responsible AI metrics are derived as aggregate quality measures influencing controller priorities. Instead of direct measurements, these metrics act as evaluative tools guiding trade-offs among competing quality objectives. They reflect the risks and impacts of quality failures at a systemic level.

Rationale: By framing Responsible AI as a meta-level concern, it integrates seamlessly with the PREFS-R framework. Metrics like “Risk of Quality Failure” and “Impact Assessment” serve as overarching tools to evaluate the likelihood, severity, and consequences of deviations from expected behaviours across foundational metrics (e.g., performance, robustness).

Definition: Our analysis prioritises ethical considerations in Responsible AI, focusing on its role as a meta-level framework that evaluates and balances the foundational core quality metrics (fairness, explainability, sustainability, robustness, and performance). Responsible AI addresses trade-offs and contextual priorities by introducing two overarching meta-metrics:

- Risk of Quality Failure (Meta-Risk Analysis): Measures the likelihood and severity of performance or robustness falling below critical thresholds of each identified risk i .

$$\text{Risk of Quality Failure} = \sum (\text{Likelihood}_i \times \text{Severity}_i) \quad (1)$$

- Likelihood: Based on statistical evaluation of operational data, considering external disturbances, internal errors, or drift;
- Severity: Defined by the impact on system reliability, safety, or user trust.
- Ethical Damage/Cost Score (Impact Evaluation): Quantifies potential damage across societal, environmental, and technical dimensions using weighted criteria.

$$\text{Impact Score} = \sum (\text{Potential Damage}_i \times \text{Weight}_i) \quad (2)$$

- Damage refers to measurable consequences, such as resource wastage, and privacy breaches;
- Weights prioritise different concerns based on the context of the application (e.g., sustainability vs. robustness).

Example: In traffic systems, performance and robustness often dominate, but in medical contexts with humans directly involved as the controlled subjects, fairness and explainability play a more significant role. Sustainability might be more dominating in resource-constrained environments.

6. Discussion of the PREFS-R Framework

We have introduced six types of quality metrics for adaptive-systems controllers, with five core metrics and one meta-level metric. This aims to motivate this specific catalogue of metrics as being relevant for the chosen architectural controller settings. We summarise key concepts, as well as soundness and completeness, present a cloud controller use case, list wider theoretical and practical implications, and finally, discuss Generative AI as another AI technique that will receive more attention in the future.

6.1. PREFS-R Summary

We summarise our main metrics definitions in Table 2. We refer to this as the PREFS-R framework of metrics. This framework takes on board the quality concerns outlined in [10], where the notion of setpoints as goal representations and their change is used to refer to stability or disturbance rejection as a robustness concern. Performance in terms of loss and convergence is also covered there. Thus, PREFS-R integrates already accepted quality concerns.

Table 2. Summary of metrics with definition, determination (M + A in MAPE), and remediation (P + E in MAPE).

Quality Metric	Definition	Determination [M + A]	Remediation [P + E]
Performance	reward optimisation or stability	Layer 1 quality (e.g., execution time, workload) and Layer 2 quality (e.g., convergence, loss)	built into rewards and policy optimisation
Robustness	tolerance against disturbances	impact on performance metrics during disturbances	pattern recognition (learned or manual)
Explainability	evaluation of controller actions	recording of actions and a posteriori analysis of action impact	reconfiguration
Fairness	no event or action bias	bias detection via holistic performance comparison	mechanisms such as favourable labels
Sustainability	resource consumption	environmental metrics—energy and cost	reward or penalise consumption
Responsible AI (meta-level)	ethical use of AI	assessment of ethical implications through metrics like Risk of Harm and Ethical Impact Score	implementation of guidelines for ethical AI practices

The PREFS-R framework can be considered sound and complete. All five core metrics are *valid* because they are all valid criteria for AI-generated models that are widely accepted in the literature. We have referred to suitable survey papers for the respective aspects. For our context, we adjusted the definitions to the controller setting. Furthermore, they can all be justified through concrete usage scenarios. The set of metrics is *complete* in terms of criteria reflected in the current literature. No other aspects have been suggested at this stage, nor have these emerged in the current discussion. It covers the concerns usually included in the concept of *ethical AI*. Concerns that apply here include, (<https://ethical.institute/index.html> (accessed on 1 December 2025)) in addition to fairness and explainability that we cover explicitly, the following:

- Transparency and accountability, in particular regarding the data sources used and their quality. Here the implications of input data quality such as incompleteness or incorrectness are relevant. Observation disturbances (as part of the robustness metrics) occur if not all context can be observed to cover incompleteness as a data quality concern. Action disturbances allow us to deal with uncertainty and incorrectness in data. In this way, reproducibility is also addressed.

- Ethical and legal compliance of the actions caused by data processing. The performance metrics allow us to monitor the controller via reward and penalty settings to discourage unethical and illegal behaviour.

6.2. Controller Quality Use Case

We present now a use case for an RL-constructed controller for cloud infrastructure scalability. This extends the previous motivational example from Section 3.3. We have fully implemented this controller using both SARSA and Q-Learning as RL techniques, which we defined in Section 4.3. More experimental results are available in the respective publications. This use case discussion serves to validate the usefulness and implementability of the metric collection and possible necessary remediation strategies.

This RL-constructed controller takes resource utilisation (workload) and performance (response time) as inputs to determine an updated infrastructure configuration in terms of up- or down-scaling computing resources to maintain response time commitments while also aiming to reduce load, and therefore cost.

In Table 3, we summarise the quality metric examples from the previous sections. We now discuss the metrics one-by-one in the context of the RL cloud resource controller. For this, we describe the quality problem and a respective remediation solution.

Table 3. Controller quality metrics—examples.

Quality Metric	Problem	Remediation
Performance	convergence and loss problems due to costly training	use of expert models in training
Robustness	white noise in distributed environments cause by the network	adjust training criteria and detect anomalies
Explainability	identification of favoured metric or strategy	use of explicit weightings for control
Fairness	underprovisioning as a resource allocation strategy could disadvantage users on high demand	define favourable labels that reflect user needs
Sustainability	overprovisioning is performance-effective, but causes high energy consumption	penalise undesirable sustainability results
Responsible AI (meta-level)	potential ethical concerns regarding bias and lack of transparency in AI decisions	implement guidelines and frameworks to ensure fairness, accountability, and transparency

- Performance:
 - Problem: Training can be costly. Often, comprehensive data sets might not exist. Even if they do and might serve as input for training an ML model, generally, not all factors are known about the construction of these data sets. Various cloud and hosting provider data sets are available, but lack, for instance, information about the occurrence of failure in the recorded operation.
 - Remediation: A remediation can be the use of an own initial expert model to kick-start the learning process with a stable model for faster convergence.
- Robustness:
 - Problem: Disturbances in the controller input signals can lead to uncertainties in the subsequent analysis and decision making based on these.
 - Remediation: This can be addressed in two ways. In the construction, formal approaches such as fuzzification are useful for uncertainty management by making the degree of uncertainty explicit and tangible in further decision making on

remediation actions. Testing as a second perspective with induced noise can then show the capability of the controller to operate in the presence of noise.

- Explainability:
 - Problem: In our controller implementation, performance is valued higher than resource utilisation, since the former would likely cause SLA violations and contractual penalties if violated.
 - Remediation: This can be explicitly controlled through weighted consideration, making the motivation behind the preferences visible.
- Fairness:
 - Problem: Technical bias—as already indicated in Table 3—is a fairness concern. Overprovisioning could be favoured by a service provider as it would result in consumers spending more money on contracted infrastructure.
 - Remediation: Our experimental evaluation has compared an ideal solution with a selected set of common strategies such as over- and underprovisioning. The experiments that compare strategies can help to detect unfair strategies based on objectively defined quality criteria.
- Sustainability:
 - Problem: The other negative side of overprovisioning, in addition to consumer costs as discussed above, is a higher energy consumption.
 - Remediation: This can be addressed in the RL model construction through encoded penalties for high resource usage.

All these metrics were used in past work on cloud scalability controllers [15,52,53] and have been experimentally validated. Responsible AI, however, as a meta-level concern, is not a directly applicable at the technology level.

6.3. Contribution to Theory and Practice

From a theoretical perspective, the contribution here provides a sound definition of valid quality metrics as discussed above, presented as a coherent catalogue of metrics all documented in a common format.

The practical relevance is enhanced by framing this within a reference architecture that clarifies the relevant components and processing activities. A motivational example of an RL-based cloud infrastructure controller, presented in two parts for motivation and as a validation use cases, illustrates a possible implementation. We have also discussed the wider DevOps/AIOps industry context, which are concepts that frame the automation of software development and operation, in particular under consideration of AI technologies.

6.4. Generative AI and LLM

Recently, generative AI (GenAI) technologies based on large language models (LLMs) have been used in software construction. While these are not the focus here, because currently little experience in this area has been documented for controller construction, we do report on a couple of related initiatives to highlight the importance.

There is a wide range of activities around the utilisation of LLMs for software generation in general. A useful software system management context to understand GenAI-based controller construction is referred to as Infrastructure-as-Code (IaC) [54–56], i.e., the configuration and operation of systems through dedicated software, the IaC scripts. Infrastructure here is any software-controllable resource.

A current focus is on code generation, as indicated by [57] in an LLM technology review. IaC code generation can be carried out through LLMs such as GPT3.5-turbo and Codeparrot, which have been experimentally explored. For instance, GPT-4 was used to

generate Kubernetes container configurations. These functions are integrated some into LLM tools, such as Infracopilot or K8sGPT, which adopt and integrate AI solutions for management and control into more general tool environments [58].

Besides using LLMs for code generation, work has also started on its quality assessment. A benchmark for LLM-generated IaC code is presented by [59]. It is based on a data set that includes 458 scenarios covering a wide range of services. Each scenario consists of a natural language IaC problem description and an infrastructure solution specification, where the former is input of the LLM and the latter is used to verify the generated IaC program. The authors note that at present even top-performing LLMs only achieve an accuracy of less than 20% for the configuration and control domain, compared to around 80% on general code generation benchmarks.

In [60], two research directions are foreseen: (i) collective intelligence, where a multi-agent collaboration framework allows different LLM agents to cover different roles for self-evolution; and (ii) experience accumulation, where LLM agents continuously acquire new “skills” for emerging tasks such as architecture-based adaptation. LLMs differ from self-adaptive systems in that they rely more on natural language than formally defined knowledge. A related approach is Red Hat’s Ansible Lightspeed (<https://www.redhat.com/en/technologies/management/ansible/ansible-lightspeed> (accessed on 1 December 2025)), which is an AI coding assistant that uses the IBM watsonx Code Assistant to produce code recommendations based on best practices for the configuration management tools Ansible.

The lack of comprehensive results for controller construction indicate this as a direction for future research on quality metrics.

7. Conclusions

This look into self-adaptive systems shows that the construction and operation of controllers for these are of critical importance. Continuous quality management is, however, still an open research problem [25]. The problem space we investigated is at the intersection of different research fields: software engineering, ML, control theory, automation and self-adaptive systems, and also data analysis—here with an application focus on self-adaptive cyber-physical and virtualised software-controlled systems. We need to develop systems where automation is important throughout the whole life-cycle and where, using AI and control-theoretic approaches in the construction and operation of systems, they are made to be self-adaptive through continuous quality control. A new perspective emerges through the need to automate construction and operation because human intervention is not possible or recommendable, be that for cost, time, or dependability reasons. Another recent direction beyond ML is the use of generative AI. Generative AI as coding assistance is built on LLM technologies.

Our core contribution consists of a metrics catalogue to assess the quality of system controllers. We based this work on several review papers covering their respective focal concern to form a coherent conceptual framework here that provides a comprehensive set of metrics and techniques to allow this two-layered controller loop to be realised.

Weyns has discussed a number of open challenges in his 2020 review [11]. Adaptation in decentralised settings and alignment with emerging technologies such as IoT and cyber-physical systems are specific concerns that we have also highlighted. Weyns also noted that dealing with changing goals and in particular dealing with unanticipated change are concerns of relevance, which we have addressed in our quality framework by introducing drift, anomalies, and their cause analysis techniques. Also, the exploitation of artificial intelligence was named as an open challenge that we focused on.

Several directions remain for future work specific to our proposed framework. We already indicated that the metrics catalogue remains at a conceptual level with a use case validation here, as detailed empirical investigations on a larger scale of the individual metrics across the five categories in different controller settings were beyond the scope of this investigation. Furthermore, handling strategies for the different metric types would also have to be investigated further if the actual implementation of a multi-objective controller for all metrics is envisaged. We have focused on reinforcement learning as the construction mechanism, but Generative AI and LLMs will play a more significant role in the future, as already pointed out in [60].

Author Contributions: Conceptualization, C.P. and H.R.B.; methodology, C.P. and H.R.B.; validation, N.E.I.; data curation, C.P.; writing—original draft preparation, C.P. and N.E.I.; writing—review and editing, C.P. and N.E.I.; supervision, C.P.; project administration, C.P. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Free University of Bozen-Bolzano, grant QUASAR.

Data Availability Statement: This manuscript is based on conceptual/theoretical work. Based on a review of related literature, a conceptual framework (consisting of a reference architecture and defined metrics) is extracted. All literature sources used are referenced. No other data (neither from external sources nor internally generated) was used in this process. No experimental work that would consume or produce data was carried out.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Rajasekhar, N.; Radhakrishnan, T.; Samsudeen, N. Exploring reinforcement learning in process control: A comprehensive survey. *Int. J. Syst. Sci.* **2025**, *56*, 3528–3557. [CrossRef]
2. El-Khatib, M.F.; Khater, F.M.; Hendawi, E.; Abu El-Sebah, M.I. Simplified and intelligent controllers for multi-input multi-output processes. *Eng. Appl. Artif. Intell.* **2025**, *141*, 109816. [CrossRef]
3. Cheng, H.; Jung, S.; Kim, Y.B. A novel reinforcement learning controller for the DC-DC boost converter. *Energy* **2025**, *321*, 135479. [CrossRef]
4. El-Khatib, M.F.; Sabry, M.N.; El-Sebah, M.I.A.; Maged, S.A. Hardware-in-the-loop testing of simple and intelligent MPPT control algorithm for an electric vehicle charging power by photovoltaic system. *ISA Trans.* **2023**, *137*, 656–669. [CrossRef] [PubMed]
5. Wan, Z.; Xia, X.; Lo, D.; Murphy, G.C. How does Machine Learning Change Software Development Practices? *IEEE Trans. Softw. Eng.* **2021**, *47*, 1857–1871. [CrossRef]
6. Al-Nima, R.R.O.; Han, T.; Al-Sumaidae, S.A.M.; Chen, T.; Woo, W.L. Robustness and performance of Deep Reinforcement Learning. *Appl. Soft Comput.* **2021**, *105*, 107295. [CrossRef]
7. Buşoniu, L.; de Bruin, T.; Tolić, D.; Kober, J.; Palunko, I. Reinforcement learning for control: Performance, stability, and deep approximators. *Annu. Rev. Control* **2018**, *46*, 8–28. [CrossRef]
8. Xu, X.; Chen, Y.; Bai, C. Deep Reinforcement Learning-Based Accurate Control of Planetary Soft Landing. *Sensors* **2021**, *21*, 8161. [CrossRef] [PubMed]
9. Faustino, J.; Adriano, D.; Amaro, R.; Pereira, R.; da Silva, M.M. DevOps benefits: A systematic literature review. *Softw. Pract. Exp.* **2022**, *52*, 1905–1926. [CrossRef]
10. Filieri, A.; Maggio, M.; Angelopoulos, K.; D'Ippolito, N.; Gerostathopoulos, I.; Hempel, A.B.; Hoffmann, H.; Jamshidi, P.; Kalyvianaki, E.; Klein, C.; et al. Software Engineering Meets Control Theory. In Proceedings of the 2015 IEEE/ACM 10th International Symposium on Software Engineering for Adaptive and Self-Managing Systems, SEAMS 2015, Florence, Italy, 18–19 May 2015; Inverardi, P., Schmerl, B.R., Eds.; IEEE Computer Society: Washington, DC, USA, 2015; pp. 71–82.
11. Weyns, D. Engineering Self-Adaptive Systems: A Short Tour in Seven Waves. In *An Introduction to Self-Adaptive Systems: A Contemporary Software Engineering Perspective*; IEEE: New York, NY, USA, 2020; pp. 17–24.
12. Pahl, C.; Azimi, S.; Barzegar, H.R.; Ioini, N.E. A Quality-driven Machine Learning Governance Architecture for Self-adaptive Edge Clouds. In Proceedings of the 2th International Conference on Cloud Computing and Services Science, CLOSER 2022, Benidorm, Spain, 19–21 May 2026; SCITEPRESS: Setúbal, Portugal, 2022; pp. 305–312.
13. Pahl, C.; Azimi, S. Constructing Dependable Data-Driven Software With Machine Learning. *IEEE Softw.* **2021**, *38*, 88–97. [CrossRef]

14. Pahl, C. Research challenges for machine learning-constructed software. *Serv. Oriented Comput. Appl.* **2023**, *17*, 1–4. [CrossRef]
15. Zhou, G.; Tian, W.; Buyya, R.; Xue, R.; Song, L. Deep reinforcement learning-based methods for resource scheduling in cloud computing: A review and future directions. *Artif. Intell. Rev.* **2024**, *57*, 124. [CrossRef]
16. Saboori, A.; Jiang, G.; Chen, H. Autotuning Configurations in Distributed Systems for Performance Improvements Using Evolutionary Strategies. In Proceedings of the 2008 The 28th International Conference on Distributed Computing Systems, Beijing, China, 17–20 June 2008; pp. 769–776. [CrossRef]
17. Glossop, C.R.; Panerati, J.; Krishnan, A.; Yuan, Z.; Schoellig, A.P. Characterising the Robustness of Reinforcement Learning for Continuous Control using Disturbance Injection. *CoRR* **2022**, *abs/2210.15199*.
18. Jabbari, S.; Joseph, M.; Kearns, M.J.; Morgenstern, J.; Roth, A. Fair Learning in Markovian Environments. *arXiv* **2016**, arXiv:1611.03071.
19. Krajna, A.; Brcic, M.; Lipic, T.; Doncevic, J. Explainability in reinforcement learning: Perspective and position. *arXiv* **2022**, arXiv:2203.11547. [CrossRef]
20. Milani, S.; Topin, N.; Veloso, M.; Fang, F. A Survey of Explainable Reinforcement Learning. *arXiv* **2022**, arXiv:2202.08434. [CrossRef]
21. Li, N.; Adepu, S.; Kang, E.; Garlan, D. Explanations for human-on-the-loop: A probabilistic model checking approach. In Proceedings of the SEAMS'20: International Symposium on Software Engineering for Adaptive and Self-Managing Systems, Seoul, Republic of Korea, 29 June–3 July 2020; pp. 181–187.
22. Wohlrab, R.; Cámara, J.; Garlan, D.; Schmerl, B.R. Explaining quality attribute tradeoffs in automated planning for self-adaptive systems. *J. Syst. Softw.* **2023**, *198*, 111538. [CrossRef]
23. Mou, Z.; Huo, Y.; Bai, R.; Xie, M.; Yu, C.; Xu, J.; Zheng, B. Sustainable Online Reinforcement Learning for Auto-bidding. In Proceedings of the NIPS'22: 36th International Conference on Neural Information Processing Systems, New Orleans, LA, USA, 28 November–9 December 2022.
24. Zhang, M.; Li, J.; Zhao, H.; Tei, K.; Honiden, S.; Jin, Z. A Meta Reinforcement Learning-based Approach for Self-Adaptive System. In Proceedings of the IEEE International Conference on Autonomic Computing and Self-Organizing Systems, ACSOS 2021, Washington, DC, USA, 27 September–1 October 2021; pp. 1–10.
25. He, X.; Zhao, K.; Chu, X. AutoML: A survey of the state-of-the-art. *Knowl.-Based Syst.* **2021**, *212*, 106622. [CrossRef]
26. Lwakatare, L.E.; Raj, A.; Bosch, J.; Olsson, H.H.; Crnkovic, I. A Taxonomy of Software Engineering Challenges for Machine Learning Systems: An Empirical Investigation. In *Agile Processes in Software Engineering and Extreme Programming, Proceedings of the 20th International Conference, XP 2019, Montréal, QC, Canada, 21–25 May 2019*; Lecture Notes in Business Information Processing; Kruchten, P., Fraser, S., Coallier, F., Eds.; Springer: Berlin/Heidelberg, Germany, 2019; Volume 355, pp. 227–243.
27. Villegas, N.M.; Tamura, G.; Müller, H.A.; Duchien, L.; Casallas, R. DYNAMICO: A Reference Model for Governing Control Objectives and Context Relevance in Self-Adaptive Software Systems. In Proceedings of the Software Engineering for Self-Adaptive Systems II—International Seminar Dagstuhl Castle, Wadern, Germany, 24–29 October 2010; Revised Selected and Invited Papers; Lecture Notes in Computer Science; de Lemos, R., Giese, H., Müller, H.A., Shaw, M., Eds.; Springer: Berlin/Heidelberg, Germany, 2010; Volume 7475, pp. 265–293.
28. Calinescu, R.; Weyns, D.; Gerasimou, S.; Iftikhar, M.U.; Habli, I.; Kelly, T. Engineering Trustworthy Self-Adaptive Software with Dynamic Assurance Cases. *IEEE Trans. Softw. Eng.* **2018**, *44*, 1039–1069. [CrossRef]
29. Langford, M.A.; Cheng, B.H.C. “Know What You Know”: Predicting Behavior for Learning-Enabled Systems When Facing Uncertainty. In Proceedings of the 16th International Symposium on Software Engineering for Adaptive and Self-Managing Systems, Madrid, Spain, 18–24 May 2021; pp. 78–89.
30. Carreño, A.; Inza, I.; Lozano, J.A. Analyzing rare event, anomaly, novelty and outlier detection terms under the supervised classification framework. *Artif. Intell. Rev.* **2020**, *53*, 3575–3594. [CrossRef]
31. Barros, R.S.M.; Santos, S.G.T.C. A large-scale comparison of concept drift detectors. *Inf. Sci.* **2018**, *451–452*, 348–370. [CrossRef]
32. Cerquitelli, T.; Proto, S.; Ventura, F.; Apiletti, D.; Baralis, E. Automating concept-drift detection by self-evaluating predictive model degradation. *arXiv* **2019**, arXiv:1907.08120.
33. Minku, L.L.; Yao, X. DDD: A New Ensemble Approach for Dealing with Concept Drift. *IEEE Trans. Knowl. Data Eng.* **2012**, *24*, 619–633. [CrossRef]
34. Lobo, J.L.; Del Ser, J.; Laña, I.; Bilbao, M.N.; Kasabov, N. Drift Detection over Non-stationary Data Streams Using Evolving Spiking Neural Networks. In *Intelligent Distributed Computing XII*; Del Ser, J., Osaba, E., Bilbao, M.N., Sanchez-Medina, J.J., Vecchio, M., Yang, X.S., Eds.; Springer: Cham, Switzerland, 2018; pp. 82–94.
35. Chandola, V.; Banerjee, A.; Kumar, V. Anomaly Detection: A Survey. *ACM Comput. Surv.* **2009**, *41*, 15. [CrossRef]
36. Gheibi, O.; Weyns, D. Lifelong Self-Adaptation: Self-Adaptation Meets Lifelong Machine Learning. In Proceedings of the 2022 International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS), Pittsburgh, PA, USA, 22–24 May 2022; Schmerl, B.R., Maggio, M., Cámara, J., Eds.; IEEE: New York, NY, USA, 2022; pp. 1–12.

37. Femminella, M.; Reali, G. Gossip-based monitoring of virtualized resources in 5G networks. In Proceedings of the IEEE INFOCOM 2019—IEEE Conference on Computer Communications Workshops, INFOCOM Workshops 2019, Paris, France, 29 April–2 May 2019; pp. 378–384.
38. Tokunaga, K.; Kawamura, K.; Takaya, N. High-speed uploading architecture using distributed edge servers on multi-RAT heterogeneous networks. In Proceedings of the IEEE International Symposium on Local and Metropolitan Area Networks, LANMAN 2016, Rome, Italy, 13–15 June 2016; pp. 1–2.
39. Zhao, H.; Yi, D.; Zhang, M.; Wang, Q.; Xinyue, S.; Zhu, H. Multipath Transmission Workload Balancing Optimization Scheme Based on Mobile Edge Computing in Vehicular Heterogeneous Network. *IEEE Access* **2019**, *7*, 116047–116055. [CrossRef]
40. Javed, A.; Malhi, A.; Främling, K. Edge Computing-based Fault-Tolerant Framework: A Case Study on Vehicular Networks. In Proceedings of the Intl Wireless Communications and Mobile Computing Conference, IWCMC 2020, Limassol, Cyprus, 15–19 June 2020.
41. Fang, D.; Liu, X.; Romdhani, I.; Jamshidi, P.; Pahl, C. An agility-oriented and fuzziness-embedded semantic model for collaborative cloud service search, retrieval and recommendation. *Future Gener. Comput. Syst.* **2016**, *56*, 11–26. [CrossRef]
42. Hong, C.; Varghese, B. Resource Management in Fog/Edge Computing: A Survey on Architectures, Infrastructure, and Algorithms. *ACM Comput. Surv.* **2019**, *52*, 97. [CrossRef]
43. El Ioini, N.; Pahl, C. Trustworthy orchestration of container based edge computing using permissioned blockchain. In Proceedings of the 2018 Fifth International Conference on Internet of Things: Systems, Management and Security, Valencia, Spain, 15–18 October 2018; pp. 147–154.
44. El Ioini, N.; Barzegar, H.R.; Pahl, C. Trust management for service migration in Multi-access Edge Computing environments. *Comput. Commun.* **2022**, *194*, 167–179. [CrossRef]
45. Arcaini, P.; Riccobene, E.; Scandurra, P. Modeling and Analyzing MAPE-K Feedback Loops for Self-Adaptation. In Proceedings of the 10th IEEE/ACM International Symposium on Software Engineering for Adaptive and Self-Managing Systems, SEAMS 2015, Florence, Italy, 18–19 May 2015; Inverardi, P., Schmerl, B.R., Eds.; IEEE Computer Society: Los Alamitos, CA, USA, 2015; pp. 13–23.
46. De Hoog, J.; Mercelis, S.; Hellinckx, P. Improving machine learning-based decision-making through inclusion of data quality. *CEUR Workshop Proc.* **2019**, *2491*, 1–10.
47. Ehrlinger, L.; Haunschmid, V.; Palazzini, D.; Lettner, C. A DaQL to Monitor Data Quality in Machine Learning Applications. In Proceedings of the Database and Expert Systems Applications—30th International Conference, DEXA 2019, Linz, Austria, 26–29 August 2019; Lecture Notes in Computer Science; Springer: Berlin/Heidelberg, Germany, 2019; Volume 11706, pp. 227–237.
48. Ralph, P. ACM SIGSOFT Empirical Standards Released. *SIGSOFT Softw. Eng. Notes* **2021**, *46*, 19. [CrossRef]
49. Soudbakhsh, D.; Annaswamy, A.M.; Wang, Y.; Brunton, S.L.; Gaudio, J.E.; Hussain, H.S.; Vrabie, D.L.; Drgona, J.; Filev, D.P. Data-Driven Control: Theory and Applications. In Proceedings of the 2023 American Control Conference (ACC), San Diego, CA, USA, 31 May–2 June 2023; pp. 1922–1939.
50. Brunton, S.L.; Kutz, J.N. *Data Driven Science & Engineering—Machine Learning, Dynamical Systems, and Control*; Cambridge University Press: Cambridge, UK, 2017.
51. Nian, R.; Liu, J.; Huang, B. A review on reinforcement learning: Introduction and applications in industrial process control. *Comput. Chem. Eng.* **2020**, *139*, 106886. [CrossRef]
52. Arabnejad, H.; Pahl, C.; Jamshidi, P.; Estrada, G. A Comparison of Reinforcement Learning Techniques for Fuzzy Cloud Auto-Scaling. In Proceedings of the International Symposium on Cluster, Cloud and Grid Computing, CCGRID 2017, Madrid, Spain, 14–17 May 2017; pp. 64–73.
53. Jamshidi, P.; Sharifloo, A.M.; Pahl, C.; Metzger, A.; Estrada, G. Self-Learning Cloud Controllers: Fuzzy Q-Learning for Knowledge Evolution. In Proceedings of the 2015 International Conference on Cloud and Autonomic Computing, Boston, MA, USA, 21–25 September 2015; pp. 208–211.
54. Alnafessah, A.; Gias, A.U.; Wang, R.; Zhu, L.; Casale, G.; Filieri, A. Quality-Aware DevOps Research: Where Do We Stand? *IEEE Access* **2021**, *9*, 44476–44489. [CrossRef]
55. Seth, D.K.; Ratra, K.K.; Sundareswaran, A.P. AI and Generative AI-Driven Automation for Multi-Cloud and Hybrid Cloud Architectures: Enhancing Security, Performance, and Operational Efficiency. In Proceedings of the 2025 IEEE 15th Annual Computing and Communication Workshop and Conference (CCWC), Las Vegas, NV, USA, 6–8 January 2025; pp. 00784–00793.
56. Pahl, C.; Gunduz, N.G.; Sezen, O.C.; Ghamgosar, A.; El Ioini, N. Infrastructure as Code—Technology Review and Research Challenges. In Proceedings of the International Conference on Cloud Computing and Services Science CLOSER’2025, Porto, Portugal, 1–3 April 2025.
57. Srivatsa, K.G.; Mukhopadhyay, S.; Katrapati, G.; Shrivastava, M. A Survey of using Large Language Models for Generating Infrastructure as Code. In Proceedings of the 20th International Conference on Natural Language Processing (ICON), Goa, India, 14–17 December 2023; Pawar, J.D., Devi, S.L., Eds.; Goa University: Goa, India, 2023; pp. 523–533.

58. Novakova Nedeltcheva, G.; De La Fuente Ruiz, A.; Orue-Echevarria Arrieta, L.; Bat, N.; Blasi, L. Towards Supporting the Generation of Infrastructure as Code Through Modelling Approaches—Systematic Literature Review. In Proceedings of the 2022 IEEE 19th International Conference on Software Architecture Companion (ICSA-C), Honolulu, HI, USA, 12–15 March 2022; pp. 210–217.
59. Kon, P.T.J.; Liu, J.; Qiu, Y.; Fan, W.; He, T.; Lin, L.; Zhang, H.; Park, O.M.; Elengikal, G.S.; Kang, Y.; et al. IaC-Eval: A Code Generation Benchmark for Cloud Infrastructure-as-Code Programs. In *Advances in Neural Information Processing Systems*; Globerson, A., Mackey, L., Belgrave, D., Fan, A., Paquet, U., Tomczak, J., Zhang, C., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2024; Volume 37, pp. 134488–134506.
60. Li, J.; Zhang, M.; Li, N.; Weyns, D.; Jin, Z.; Tei, K. Generative AI for Self-Adaptive Systems: State of the Art and Research Roadmap. *ACM Trans. Auton. Adapt. Syst.* **2024**, *19*, 13. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Article

A Digital Twin Model for UAV Control to Lift Irregular-Shaped Payloads Using Robust Model Predictive Control

Umar Farid ¹, Bilal Khan ^{2,3}, Sahibzada Muhammad Ali ¹ and Zahid Ullah ^{4,*}

¹ Department of Electrical Engineering, COMSATS University Islamabad, Abbottabad Campus, Abbottabad 22060, Pakistan; umarfarid@cuiatd.edu.pk (U.F.)

² Department of Control and Instrumentation Engineering, King Fahd University of Petroleum and Minerals (KFUPM), Dhahran 31261, Saudi Arabia; bilal.khan@kfupm.edu.sa

³ Center for Sustainable Energy Systems (IRC-SES), King Fahd University of Petroleum and Minerals (KFUPM), Dhahran 31261, Saudi Arabia

⁴ Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, 20133 Milano, Italy

* Correspondence: zahid.ullah@polimi.it

Abstract

This paper presents an innovative digital twin (DT) model integrated with robust model predictive control (MPC) to enhance the performance of an unmanned air vehicle (UAV) tasked with lifting and transporting irregular-shaped payloads. Traditional UAV control systems face complex challenges in stability and accuracy when dealing with asymmetrical payloads, as such payloads cause continuous shifts in the center of gravity (CoG) and variable inertial forces, which lead to unpredictable flight dynamics. The proposed DT framework enables the creation of a real-time replica of the UAV payload system. It creates an adaptive control environment that anticipates and mitigates disturbances before they impact the stability of the UAV during a mission. By combining a DT with MPC, the control system dynamically adjusts to variations in payload characteristics, namely (a) changes in mass distribution and (b) aerodynamic drag force. As a result, a stable flight path is ensured even under challenging environmental conditions. The DT model continuously forecasts potential destabilizing events and modifies MPC constraints to accommodate complex shifting dynamics, achieving improved control accuracy and energy efficiency. Extensive simulations across various hanging payload configurations and environmental disturbance scenarios validate the effectiveness of the proposed model. The simulation results show that the DT-MPC strategy significantly improves stability, control precision, and energy conservation, outperforming conventional methods. A comparative analysis is also carried out with a conventional control scheme to validate the robustness of the proposed framework. This research work advances the development of intelligent, autonomous UAV systems capable of reliably managing complex and irregularly shaped payloads with varying mass distributions in real-world scenarios, thereby broadening their potential applications in logistics, emergency response, and industrial transportation.

Keywords: digital twin; model predictive control; UAV stability; irregular payload handling; drone control systems; real-time control; payload dynamics; robust control

1. Introduction

UAVs, commonly known as drones, have shown exponential growth in adoption across different sectors, such as (a) logistics, (b) healthcare, (c) agriculture, (d) construction, (e) industrial inspections [1]. Their capacity to deal with different scenarios and

environmental disturbances and access remote locations that are difficult to reach by conventional transport contribute to their indispensability for critical tasks, including package delivery, medical supply transport, crop health surveillance, infrastructure assessment, and search and rescue operations [2]. Their adaptability is important to their lightweight construction, exceptional maneuverability, and cost-effective operations [3]. Although UAVs are widely used for payload transportation duties, they face irregular loads characterized by asymmetric shapes, uneven mass distributions, and interconnected aerodynamic profiles [4].

The hanging of irregular payloads poses a significant issue to the stability, security, and control performance of UAVs [5]. Non-uniform loads can be dynamic during flight and will perturb the CoG and rotational inertia of the UAV. They affect equilibrium because UAVs are unable to track reference trajectories and orientations [6]. However, these inconsistent payloads have asymmetrical geometry and exhibit different aerodynamic forces depending on altitude, forward motion, and velocity [7]. These multiple disturbances eventually lead to flight operation instabilities in UAVs, contributing to the reasoning behind control paradigms [5].

With autonomous UAV operations, which cannot be managed manually, stability and safety requirements are necessary [8]. Complex control systems are mandatory to maintain stability and accuracy in the presence of perturbations [9]. Reducing irregular payload-related control hurdles has garnered significant attention from both academics and industry [10]. In such situations, an increase in UAV autonomy opens the door to revolutionary possibilities in logistics, disaster relief, and industrial logistics, where valuable heterogeneous loads can be smoothly managed with insignificant chances of failure [11].

Asymmetric payload delivery presents unique control complexities in UAVs, as these payloads are very unlike symmetric and predictable cargoes [12]. Doubts in exact mass and aerodynamic characteristics initiate continuous CoG disturbance, varying inertial moments, and arbitrary drag/lift distributions, making stable flight challenging [13]. As an example, regarding long or asymmetrical delivery, payload motions or transfers can be excited because of agile maneuvers, causing UAV destabilization [14].

Traditional controllers such as PID or MPC assume constant or predictable payload characteristics [15]. These principles fail with non-regular payloads, whose characteristics are adaptive to UAV kinematics and non-cooperation [16]. Conventional systems exhibit sluggish responses, require reconfiguration to load variances, and display reactive disturbance responses, which increase instability, precision issues, and energy consumption in control efforts [17].

The need to have real-time dynamic paradigms to maneuver through irregular payload transportation is important, as they enhance autonomous transfer without disturbing stability, efficiency, or security [18]. A DT is a type of virtual system that replicates the kinetics of a physical entity in real time [19]. The DT predicts interruptions, which is enabled by copying the UAV–payload interactions, enabling predictive control [20]. The DT utilizes information collected by UAV sensors, which enables it to maintain a correct virtual representation of the UAV, forecast future alterations, and detect itself just before future disturbances will impact the system [21].

Integrating a DT and MPC provides the UAV controller with a system that can overcome anomalies in the payload [21]. Robust MPC is good in stochastic prediction as it uses uncertainties in optimization [22]. It can support reallocation of payload mass, aerodynamic perturbation, and migration of CoG, enhancing disturbance resilience [23]. DT-supplied foresight allows unrelenting control input refinements, compensating payload perturbations during flight operation and maintaining stability at the expense of energetics [24].

This integration not only solves the problem of irregular payload management, but it brings improvements in efficiency and safety of transportation using UAVs [25].

2. Related Work

This section reviews prior research across three key areas relevant to this study: (1) DT technology in autonomous systems, particularly UAVs; (2) MPC for UAV stability; (3) control strategies for managing irregularly shaped payloads. Collectively, these works outline the current state of the art, highlighting gaps in the control of UAVs with complex payloads.

2.1. Digital Twin Technology in Autonomous Systems

DT technology has become a transformational methodology in manufacturing and autonomous systems, enabling the generation of real-time virtual copies of physical systems [26]. DTs facilitate real-time monitoring, diagnostics, and control by enabling data-driven simulations and predictions [27]. Investigations have developed the principles of DTs and their flexibility in working under dynamic environments, which is essential for making real-time decisions [28]. Within the framework of autonomous UAVs, DTs are used frequently to simulate the real-world situation virtually [29]. Recent papers have investigated the potential of DTs in decision-making in aerospace markets, showing greater reliability in hard operations. Predicting faults with the help of DTs enhances the safety and efficiency of UAVs [30], and enhanced models facilitate the simulation of the complex dynamics of UAVs and enable the implementation of intricate corrections.

DT applications have been developed on resilient swarms of UAVs, involving edge computing to provide low-latency synchronization [31]. Physical models with machine learning Hybrid DTs are better for predictive maintenance in unpredictable environments [32]. Federated DTs support collaborative work of multiple agents without information exchange, which is needed for use with privacy-sensitive UAV flocks [33]. Moreover, irregular payloads are still limited in DT applications, necessitating dynamic models that capture variable loading conditions [34]. This paper develops the DT-UAV combination by combining it with robust MPC to estimate and avoid destabilizing factors and extrapolate DT utility to complex tasks [21].

2.2. Model Predictive Control for UAV Stabilization

MPC is also actively applied in UAVs to deal with multi-variable constraints when operating in dynamic environments [35]. Contrary to conventional approaches, MPC forecasts the behavior of a system with a receding horizon, predicting disturbances and responding to changes in paths, which are essential for UAV stability, as presented in [36]. Principles of MPC have been elaborated through theoretic frameworks, where the benefits and shortcomings of linear and nonlinear models have been noted. Optimization of UAV reference trajectories in varying environments demonstrates the efficiency of MPC, and UAV stability results are studied in relation to the influence of disturbances, with restricted ability to model variable payloads [37].

The latest developments include stochastic MPC, which addresses uncertainties and optimum loading due to payload variations, and adaptive MPC, which handles irregularly changing slung loads [38]. The authors of [39] present innovations in learning-enhanced MPC for online parameter tuning in turbulent conditions. The authors of [40] present innovations in distributed MPC for UAV formations with payload differences. Reference [41] discusses the topic of event-driven MPC to reduce computational load in resource-constrained UAVs. MPC is resilient because it counters disturbances at a high level [23]. This study integrates DTs and MPC, where the real-time feedback of a DT enables adaptation to irregular dynamics, thereby enhancing stability and precision [21].

2.3. Control Strategies for Irregularly Shaped Payloads

The problem of irregular payloads is caused by their generated oscillations, changes in the CoG, and aerodynamic influence [42]. Conventional UAV control requires symmetry; hence, it is not suitable when there are such irregularities. The literature on suspended-load steadiness has highlighted destabilizing forces in a manner similar to CoG-variant irregularities [43]. Adaptive nonlinear control techniques have been employed to address problems with varying loads, which pose challenges in tuning and limit real-time flexibility [44]. Unknown payload-adaptive hybrids exhibit flexibility, but cannot handle sudden changes [45]. Reduced-order models of irregular payloads can accurately display the effects of a load on trajectory, highlighting the importance of understanding the disturbances that affect stability.

The latest advances are vision-based adaptive controls to dictate payload with no sensors [46], reinforcement learning to navigate nonlinearities with slung irregular loads [47], and sliding-mode observers to neutralize aerodynamic disturbances in asymmetric lifts [48]. The authors of [12] explain the Distributed Sliding-Mode Controller for single and heterogeneous drones used for transportation of a hanging payload. However, real-time dynamic adaptation remains unexplored, as it is not yet supported by kinetic simulations. DTs offer a solution through real-time, data-driven UAV–payload simulations, predicting anomalies for robust control. In this study’s DT-robust MPC framework, responses to irregular loads are enhanced with superior stability and precision, addressing adaptive shortcomings in payload variations. In [49], the authors investigate the operation of an aerial manipulator system equipped with a controllable arm with two degrees of freedom to carry out actuation tasks during flight operation. A novel low-altitude vehicle speed detector system utilizing UAVs for remote sensing applications in smart cities is introduced in [50], aiming to enhance traffic safety and security. Similarly, the authors present a novel approach to autonomous payload delivery for single and multiple heterogeneous UAVs in [51]. For single-UAV operations, SMC enables a UAV to track a predefined reference trajectory, autonomously identify the payload’s CoM, securely attach to it, lift it, and deliver it to the desired target location.

2.4. Summary and Research Gap

Existing research in the literature confirms the individual strengths of DTs and MPC in UAV control, prediction, and disturbance rejection. DTs provide accurate simulations for identifying and monitoring UAVs. Meanwhile, MPCs manage unstable environments and uncertainties, ensuring the stability of UAVs in harsh conditions.

However, there is a research gap: the integration of DTs with the MPC of hanging irregular-payload delivery is not examined in detail. The literature shows the use of conventional approaches, which lack the capability to forecast the problems of hanging payloads in real time. Furthermore, MPC, useful in response to environmental disturbances in the case of constant conditions, experiences troubles with variable hanging payloads.

The proposed research paper will fill this gap by integrating DTs with MPC to enhance the adaptability and stability of real-time flights by UAVs in terms of irregularly loaded payloads. The predictive features of DTs enable the rejection of disturbances in advance, whereas MPC allows for tighter control using forecasts, which is necessary for stability in the face of load variations. This method ensures flexible, accurate, and effective solutions for UAVs with hanging payloads, thereby widening reliability in various applications. It also leads the way in autonomous operations in complex and demanding environments by applying the use of DT predictions and MPC to achieve safer and cost-efficient transportation of irregular payloads.

The key contributions of the paper to the aforementioned research gaps are as follows:

- **Real-Time DT of Irregular Payload Kinetics:** In this paper, the authors present a DT framework for a UAV carrying an irregularly shaped hanging payload, which extends the current DT models beyond operation only in the case of a static environment. The DT dynamically recreates time-varying payload properties, including mass reallocation, CoG variations, and aerodynamic behavior, which gives a digital platform for predictive control and mitigation of disturbances.
 - **DT-MPC Integration:** The present work combines the DT and MPC, making possible the adaptive and precise control of UAVs even in the case of complicated payloads. This hybrid command utilizes real-time data from a DT to support MPC in managing uncertainties, ensuring stability and efficiency when other classic control methods fail.
 - **Adaptive Predictive Disturbance Mitigation:** We propose a new adaptive predictive rejection mechanism with the use of DT measurements of payload-induced disturbance. This gives MPC the opportunity to make early corrections to control inputs, minimize oscillations, maximize energy consumption, and enhance stability during flight with highly variable payload geometries throughout the mission.
 - **Validation Across Diverse Payloads:** The study validates the DT-MPC framework through extensive simulations across a wide range of payload configurations, addressing real-world irregularities. The results indicate that DT-MPC yields better performance in terms of stability, precision, and energy efficiency, establishing a robust foundation for autonomous UAV operations with irregular hanging payloads.
- Table 1 compares the contribution of this paper with state-of-the-art literature work.

Table 1. Comparison of Contributions with State-of-the-Art.

Contribution	Novelty vs. Prior Art
1. Real-Time DT for Irregular Payload Kinetics	Unlike [36], we model time-varying $m_p(t)$, $r_p(t)$, and $I_p(t)$ using online identification.
2. DT–Robust MPC Synergy	We are the first to use DT forecasts to update MPC & tube bounds in real time (Section 3.3.3).
3. Adaptive Predictive Disturbance Mitigation	We achieve pre-emptive control via ARIMA-disturbance forecasting, not reactive control like in [51].
4. Comprehensive Validation	We present 18 figures across diverse irregular payloads + disturbances, beyond [46].

The rest of the paper consists of Section 3, explaining the modeling and architecture in detail; Section 4, explaining the problem formulation and analysis; Section 5, which presents detailed simulation results to validate the proposed work; and Section 6, which concludes the paper and gives some future directions.

3. System Architecture

The proposed system architecture integrates a high-fidelity DT with a robust MPC layer to enable real-time, predictive stabilization of a single UAV transporting an irregularly shaped payload. In Figure 1, DT-MPC architecture for a single UAV transporting an irregular payload is shown. Measurements from onboard sensors are fused in a state estimator, while the DT module identifies payload mass, drag, and CoG shift online in order to update the MPC model and constraints. Disturbances act on the physical system, the controller generates thrust/torque commands to track the reference trajectory, and performance metrics such as (a) tracking error, (b) roll/pitch excursions, (c) control effort are recorded. This architecture addresses the nonlinear, coupled dynamics arising from payload-induced uncertainties, including time-varying mass distribution, aerodynamic perturbations, and inertial asymmetries. The framework is hierarchically designed, com-

prising the following elements: (i) the UAV–payload dynamics framework, which is a 6-DoF rigid-body equation, with the addition of the payload; (ii) the DT framework, which offers high-resolution predictive simulations based on the augmentation of the state and the use of Kalman filtering; (iii) the robust MPC layer, which optimizes the control inputs within the framework of uncertainty bounds and stability constraints. Rigorous mathematical validation ensures asymptotic stability and constraint satisfaction under bounded disturbances.




Figure 1. System architecture of Digital Twin Based MPC for payload delivery.

3.1. UAV and Payload Dynamics Model

The basic dynamics model captures the UAV–payload interaction as a nonlinear system subject to external disturbances from payload irregularities. Figure 2 presents the schematic diagram of a UAV with a hanging payload.




Figure 2. Schematic diagram of UAV with payload.

3.1.1. UAV Flight Dynamics

The UAV's 6-DoF motion is governed by the Newton–Euler equations in the body-fixed frame $\mathcal{B}$, transformed to the inertial frame $\mathcal{I}$ via the rotation matrix $R \in \mathbb{R}^3$. Let $p \in \mathbb{R}^3$ denote the position in $\mathcal{I}$, $\mathbf{v} \in \mathbb{R}^3$ the velocity in $\mathcal{B}$, $\boldsymbol{\omega} \in \mathbb{R}^3$ the angular velocity in $\mathcal{B}$, $m > 0$ the UAV mass, and $I \in \mathbb{R}^{3 \times 3}$ the positive definite inertia tensor.

The translational dynamics are

$$m\dot{\mathbf{v}} = R^T(m\mathbf{g} + \mathbf{f}_t + \mathbf{f}_d + \mathbf{f}_p) - m\boldsymbol{\omega} \times \mathbf{v}, \quad (1)$$

where $\mathbf{g} = [0, 0, -g]^T$ is gravity ($g = 9.81 \text{ m/s}^2$), $\mathbf{f}_t = [0, 0, T]^T$ is the collective thrust ($T > 0$), $\mathbf{f}_d = -\frac{1}{2}\rho\|\mathbf{v}\|^2 C_d A \frac{\mathbf{v}}{\|\mathbf{v}\|}$ is drag ($\rho > 0$ air density, $C_d > 0$ drag coefficient, $A > 0$ reference area), and $\mathbf{f}_p$ is the payload force (detailed in Section 3.1.2).

The rotational dynamics are

$$I\dot{\boldsymbol{\omega}} + \boldsymbol{\omega} \times (I\boldsymbol{\omega}) = \boldsymbol{\tau}_t + \boldsymbol{\tau}_p, \quad (2)$$

where $\boldsymbol{\tau}_t = [\tau_x, \tau_y, \tau_z]^T$, attitude torques from differential rotor speeds, and $\boldsymbol{\tau}_p$ is the payload torque.

In state-space form, the nominal UAV dynamics are $\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{u})$, where $\mathbf{x} = [\mathbf{p}^T, \mathbf{v}^T, \boldsymbol{\theta}^T, \boldsymbol{\omega}^T]^T \in \mathbb{R}^{12}$ ($\boldsymbol{\theta}$ Euler angles) and $\mathbf{u} = [T, \tau_x, \tau_y, \tau_z]^T \in \mathbb{R}^4$.

Theorem 1. For any equilibrium $(\mathbf{x}_e, \mathbf{u}_e)$ where $\boldsymbol{\omega}_e = \mathbf{0}$ and $\mathbf{v}_e = \mathbf{0}$, the dynamics admit a local linearization $\dot{\mathbf{z}} = A\mathbf{z} + B\mathbf{w}$, with $\mathbf{z} = \mathbf{x} - \mathbf{x}_e$, $A = \frac{\partial f}{\partial \mathbf{x}}|_{\mathbf{x}_e}$, $B = \frac{\partial f}{\partial \mathbf{u}}|_{\mathbf{u}_e}$, and $\mathbf{w}$ a perturbation input, provided the Coriolis terms are Lipschitz-continuous (bounded by $L > 0$).

Proof. By the inverse function theorem, since f is C^1 and the Jacobian A is Hurwitz for small perturbations (as $I \succ 0$), the linearization holds in a neighborhood $\|\mathbf{z}\| < \delta$. $\square$

3.1.2. Payload Dynamics

The irregular payload introduces time-varying parameters: let $m_p(t) > 0$ be the payload mass, $\mathbf{r}_p(t) \in \mathbb{R}^3$ the relative position vector from the UAV center of mass (CoM), $I_p(t) \in \mathbb{R}^{3 \times 3}$ the payload inertia tensor, and $\mathbf{v}_p = \dot{\mathbf{r}}_p + \mathbf{v} + \boldsymbol{\omega} \times \mathbf{r}_p$ the payload-relative velocity.

The payload force is

$$\mathbf{f}_p = -m_p\dot{\mathbf{v}}_p - m_p\boldsymbol{\omega} \times \mathbf{v}_p + \mathbf{f}_{a,p}, \quad (3)$$

where $\mathbf{f}_{a,p} = -\frac{1}{2}\rho\|\mathbf{v}_p\|^2 C_{d,p}(t) A_p(t) \frac{\mathbf{v}_p}{\|\mathbf{v}_p\|}$ is the payload-specific drag ($C_{d,p}(t)$, $A_p(t)$ time-varying due to orientation-dependent geometry).

The payload torque is

$$\boldsymbol{\tau}_p = \mathbf{r}_p \times \mathbf{f}_p + I_p\dot{\boldsymbol{\omega}}_p + \boldsymbol{\omega}_p \times (I_p\boldsymbol{\omega}_p), \quad (4)$$

where $\boldsymbol{\omega}_p = \boldsymbol{\omega} + \dot{\boldsymbol{\theta}}_p$ (payload angular rate, including relative rotation $\dot{\boldsymbol{\theta}}_p$).

The coupled total dynamics become

$$(M(t)\dot{\mathbf{v}} + C(\mathbf{v}, \boldsymbol{\omega}, t)\mathbf{v}) = R^T m\mathbf{g} + \mathbf{f}_t + \mathbf{f}_d + \mathbf{f}_{a,p}, \quad (5)$$

$$(I(t) + I_p(t))\dot{\boldsymbol{\omega}} + C_r(\boldsymbol{\omega}, t) = \boldsymbol{\tau}_t + \boldsymbol{\tau}_p, \quad (6)$$

where $M(t) = m + m_p(t)$ is total mass, $C(\cdot)$ is the Coriolis/centripetal matrix (skew-symmetric), and $C_r(\cdot)$ includes cross-terms.

Remark 1. The time-varying $I(t) = I + I_p(t) + m_p(t)\mathbf{r}_p(t)\mathbf{r}_p(t)^T$ (parallel axis theorem) induces parametric uncertainty $\Delta I(t)$ with $\|\Delta I(t)\| \leq \bar{\epsilon}_I$, bounding inertial drift.

Lemma 1. For inputs $\|\mathbf{u}\| \leq u_{\max}$, the coupled system satisfies $\|\mathbf{x}(t)\| \leq k_1\|\mathbf{x}(0)\|e^{-\lambda t} + k_2u_{\max}$ for constants $k_1, k_2 > 0$, $\lambda > 0$, assuming $\|\dot{m}_p(t)\| \leq \epsilon_m$, $\|\dot{I}_p(t)\| \leq \epsilon_I$.

Proof. Integrate the dynamics; the skew-symmetry of $C(\cdot)$ ensures energy dissipation, and Gronwall's inequality bounds growth from parametric rates. $\square$

3.2. Digital Twin Framework

The DT is a parallel simulator $\hat{\mathbf{x}}(t)$ synchronized with the physical state $\mathbf{x}(t)$ via sensor fusion, enabling predictive extrapolation over horizon N .

3.2.1. Real-Time Data Acquisition and Synchronization

Sensors provide noisy measurements $\mathbf{y}(t) = H\mathbf{x}(t) + \mathbf{v}(t)$, $\mathbf{v} \sim \mathcal{N}(0, \Sigma_v)$. Synchronization uses an extended Kalman filter (EKF):

$$\hat{\mathbf{x}}_{k|k} = \hat{\mathbf{x}}_{k|k-1} + K_k(\mathbf{y}_k - H\hat{\mathbf{y}}_{k|k-1}), \quad (7)$$

with gain $K_k = P_{k|k-1}H^T(HP_{k|k-1}H^T + \Sigma_v)^{-1}$, and covariance update $P_{k|k-1} = AP_{k-1|k-1}A^T + Q + \Gamma\Sigma_w\Gamma^T$ (Q is process noise, Γ is the disturbance matrix).

Theorem 2. The estimation error $\tilde{\mathbf{x}}_k = \mathbf{x}_k - \hat{\mathbf{x}}_{k|k}$ is asymptotically stable; i.e., $\mathbb{E}[\|\tilde{\mathbf{x}}_k\|^2] \rightarrow 0$, if A and H are detectable and the linearization error is second-order small ($O(\|\tilde{\mathbf{x}}\|^2)$).

Proof. By the EKF mean-square stability theorem, persistent excitation from H ensures the Riccati equation converges to a steady-state $P_\infty \succ 0$. $\square$

3.2.2. Predictive Modeling and Simulation Engine

The DT propagates $\hat{\mathbf{x}}_{k+i|k} = \hat{f}(\hat{\mathbf{x}}_{k+i-1|k}, \hat{\mathbf{u}}_{k+i-1|k}, \hat{\mathbf{w}}_{k+i-1|k})$ over $i = 1, \dots, N$, where $\hat{\mathbf{w}}$ is a disturbance forecast via the ARIMA model: $\hat{\mathbf{w}}_{k+i} = \sum_{j=1}^p \phi_j \mathbf{w}_{k+i-j} + \boldsymbol{\eta}_i$, $\boldsymbol{\eta}_i \sim \mathcal{N}(0, \Sigma_\eta)$. Adaptive sampling adjusts step size $\Delta t_k = \min(0.01, \frac{\epsilon_s}{\|\nabla_{\mathbf{x}} \hat{f}\|})$ based on sensitivity.

3.2.3. Feedback Loop for Control Adjustments

The DT outputs $\{\hat{\mathbf{x}}_{k+i|k}, \hat{\mathbf{w}}_{k+i|k}\}_{i=1}^N$ to the MPC layer via a certainty-equivalent map, with error bound $\|\hat{\mathbf{x}}_{k+i|k} - \mathbf{x}_{k+i}\| \leq \beta^i \epsilon_{\text{sync}}$ ($\beta < 1$).

Remark 2. Synchronization latency $\tau < 10^{-2}$ s ensures $\epsilon_{\text{sync}} \leq 0.05$ m, validated by trace $P_\infty < 10^{-4}$.

3.3. Robust Model Predictive Control Layer

Robust MPC optimizes $\mathbf{u}$ over N , subject to tube constraints for uncertainty.

3.3.1. MPC Optimization Engine

Minimize stage cost $J_k = \sum_{i=0}^{N-1} (\|\mathbf{x}_{k+i|k} - \mathbf{x}_{\text{ref}}\|_Q^2 + \|\mathbf{u}_{k+i|k}\|_R^2) + \|\mathbf{x}_{k+N|k} - \mathbf{x}_{\text{ref}}\|_{P_f}^2$, s.t.

$$\mathbf{x}_{k+i+1|k} = A\mathbf{x}_{k+i|k} + B\mathbf{u}_{k+i|k} + \mathbf{w}_{k+i|k}, \quad \mathbf{u}_{\min} \leq \mathbf{u}_{k+i|k} \leq \mathbf{u}_{\max}, \quad (8)$$

$$\mathbf{x}_{\min} \leq \mathbf{x}_{k+i|k} \leq \mathbf{x}_{\max}, \quad (9)$$

using DT forecasts for $\mathbf{w}_{k+i|k}$.

3.3.2. Robustness Enhancements for Uncertainty Management

Theorem 3. $\min_{\mathbf{u}} \max_{\mathbf{w} \in \mathcal{W}} J_k$, where $\mathcal{W} = \{\mathbf{w} : \|\mathbf{w}\| \leq \epsilon_w\}$ (DT-estimated ϵ_w). Tube MPC centers nominal trajectory $\mathbf{z}_{k+i|k}$ with error set $\mathcal{E} = \{\mathbf{e} : \mathbf{e}_{j+1} = (A + K\Delta A)\mathbf{e}_j + \mathbf{w}_j, \|\mathbf{w}_j\| \leq \epsilon_w\}$, controlled by $\mathbf{u}_{k+i|k} = \mathbf{v}_{k+i|k} + K\mathbf{e}_{k+i|k}$.

If initial $\mathbf{x}_k \in \mathcal{X}_f$ (terminal set), there exists $\mathbf{u}_k^*$ feasible for all $\mathbf{w} \in \mathcal{W}$, with $\mathbf{x}_{k+1} \in \mathcal{X}_f$.

Proof. The robust invariant set $\mathcal{X}_f = \{\mathbf{x} : \exists \mathbf{u} \in \mathcal{U}, (A + BK)\mathbf{x} + \mathbf{w} \in \mathcal{X}_f, \forall \mathbf{w} \in \mathcal{W}\}$ is nonempty by Pontryagin's minimax theorem, as the value function decreases. $\square$

3.3.3. Real-Time Control Adaptation via DT Feedback

At each step, update $K_k = -(B^T P_k B + R)^{-1} B^T P_k A$ with P_k from DT-augmented Riccati, $\dot{P} = -A^T P - PA + PBR^{-1}B^T P - Q - \hat{\Gamma} \Sigma_w \hat{\Gamma}^T$, where

$$\hat{\Gamma} \in \mathbb{R}^{12 \times 6} = \begin{bmatrix} \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} \\ \mathbf{I}_3 & \mathbf{0}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & \mathbf{J}^{-1} \end{bmatrix}, \quad \mathbf{J} \in \mathbb{R}^{3 \times 3}.$$

$$\Sigma_w = \text{diag}(\sigma_{v_x}^2, \sigma_{v_y}^2, \sigma_{v_z}^2, \sigma_{\tau_x}^2, \sigma_{\tau_y}^2, \sigma_{\tau_z}^2) = \begin{bmatrix} \Sigma_v & \mathbf{0}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & \Sigma_\tau \end{bmatrix},$$

$$\Sigma_v = \text{diag}(\sigma_{v_x}^2, \sigma_{v_y}^2, \sigma_{v_z}^2), \quad \Sigma_\tau = \text{diag}(\sigma_{\tau_x}^2, \sigma_{\tau_y}^2, \sigma_{\tau_z}^2).$$

$$\hat{\Gamma} \Sigma_w \hat{\Gamma}^T = \begin{bmatrix} \mathbf{0}_{3 \times 3} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \Sigma_v & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0}_{3 \times 3} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{J}^{-1} \Sigma_\tau \mathbf{J}^{-T} \end{bmatrix}.$$

$$\mathbf{J} = \begin{bmatrix} J_{xx} & 0 & 0 \\ 0 & J_{yy} & 0 \\ 0 & 0 & J_{zz} \end{bmatrix}, \quad J_{ii} = I_{ii} + I_{p,ii}(t) + m_p(t) r_{p,i}^2, \quad i \in \{x, y, z\}.$$

$$J_{xx} = I_{xx} + I_{p,xx}(t) + m_p(t) (r_{p,y}^2 + r_{p,z}^2),$$

$$J_{yy} = I_{yy} + I_{p,yy}(t) + m_p(t) (r_{p,x}^2 + r_{p,z}^2),$$

$$J_{zz} = I_{zz} + I_{p,zz}(t) + m_p(t) (r_{p,x}^2 + r_{p,y}^2).$$

$$\mathbf{J} = \mathbf{I}_{\text{UAV}} + \mathbf{I}_p(t) + m_p(t) \left(\|\mathbf{r}_p(t)\|^2 \mathbf{I}_3 - \mathbf{r}_p(t) \mathbf{r}_p(t)^T \right),$$

$$\mathbf{r}_p(t) = \begin{bmatrix} r_{p,x}(t) \\ r_{p,y}(t) \\ r_{p,z}(t) \end{bmatrix},$$

$$J_{xy} = I_{xy} + I_{p,xy}(t) - m_p(t) r_{p,x} r_{p,y},$$

$$J_{xz} = I_{xz} + I_{p,xz}(t) - m_p(t) r_{p,x} r_{p,z}, \quad J_{yx} = J_{xy}, \quad J_{zx} = J_{xz}, \quad J_{zy} = J_{yz}.$$

$$J_{yz} = I_{yz} + I_{p,yz}(t) - m_p(t) r_{p,y} r_{p,z},$$

3.4. System Integration and Workflow

The closed-loop system is $\mathbf{x}_{k+1} = A\mathbf{x}_k + B\mathbf{u}_k^* + \mathbf{w}_k$, with $\mathbf{u}_k^*$ from the QP solver (e.g., OSQP). Workflow: Acquire $\mathbf{y}_k \rightarrow$ extended Kalman Filter (EKF) sync $\rightarrow$ predict $\hat{\mathbf{w}} \rightarrow$ robust QP $\rightarrow$ apply $\mathbf{u}_k$.

Remark 3. The computational complexity is $O(N^3)$ per step, feasible at 100 Hz on embedded hardware.

Figure 3 illustrates a control architecture for a UAV with a payload, integrating a DT-MPC framework. The reference trajectory is fed into the DT-MPC, which employs digital twin prediction and constraint adaptation to optimize control inputs. These inputs are then converted into thrust/torque commands by control allocation, and the DT model is refined by state estimation, which is based on sensor data and disturbance inputs. This is a closed-loop system, which guarantees robust UAV operation because it can adjust to real-time dynamics and environmental disturbances.




Figure 3. Feedback Control block diagram for UAV payload-delivery system.

4. Problem Formulation: Model Design and Analysis

The control problem is to stabilize $\mathbf{x}(t) \rightarrow \mathbf{x}_{ref}$ while rejecting payload disturbances, formulated as a robust infinite-horizon optimal control with DT augmentation.

4.1. Coupled UAV–Payload Dynamics

Augment states to $\xi = [\mathbf{x}^T, \mathbf{r}_p^T, \theta_p^T]^T \in \mathbb{R}^{18}$. The full nonlinear dynamics are

$$\dot{\xi} = \mathcal{F}(\xi, \mathbf{u}, \mathbf{d}(t)) = \begin{bmatrix} f(\mathbf{x}, \mathbf{u}) + G(\mathbf{x})\mathbf{d}(t) \\ \dot{\mathbf{r}}_p + \boldsymbol{\omega} \times \mathbf{r}_p \\ \dot{\theta}_p + \boldsymbol{\omega} \end{bmatrix}, \quad (10)$$

where $\mathbf{d}(t) = [\Delta m_p(t), \Delta I_p(t), \Delta C_{d,p}(t)]^T \in \mathbb{R}^q$ is bounded, $\|\mathbf{d}(t)\| \leq \mathbf{d}_{\max}$, and $G(\cdot)$ is the disturbance distribution matrix.

Linearize around trim: $\dot{\tilde{\xi}} = \bar{A}\tilde{\xi} + \bar{B}\mathbf{u} + \bar{G}\mathbf{d}$, with Jacobians $\bar{A} = \frac{\partial \mathcal{F}}{\partial \xi}|_{\xi_e}$. Eigenvalues of $\bar{A}$ have $\Re(\lambda_i) < 0$ for small $\mathbf{d}_{\max}$ (unstable quadrotor modes stabilized by feedback).

$$A = \left. \frac{\partial \mathcal{F}}{\partial \xi} \right|_{(\xi_e, \mathbf{u}_e, \mathbf{d}=0)}, \quad B = \left. \frac{\partial \mathcal{F}}{\partial \mathbf{u}} \right|_{(\xi_e, \mathbf{u}_e, \mathbf{d}=0)}, \quad G = \left. \frac{\partial \mathcal{F}}{\partial \mathbf{d}} \right|_{(\xi_e, \mathbf{u}_e, \mathbf{d}=0)}.$$

Theorem 4. The origin of $\dot{\tilde{\xi}} = \bar{A}\tilde{\xi} + \bar{B}\mathbf{u} + \bar{G}\mathbf{d}$ is input-to-state-stable (ISS) if there exists K such that $\bar{A} + \bar{B}K$ is Hurwitz, with $\|\tilde{\xi}(t)\| \leq \beta(\|\tilde{\xi}(0)\|, t) + \gamma(\|\mathbf{d}\|_\infty)$, $\beta \in \mathcal{KL}$, and $\gamma \in \mathcal{K}$.

Proof. By Sontag’s ISS theorem, since $\bar{A} + \bar{B}K \prec 0$ and $\bar{G}$ is bounded, the Lyapunov $V = \tilde{\xi}^T P \tilde{\xi}$ satisfies $\dot{V} \leq -\alpha\|\tilde{\xi}\|^2 + \kappa\|\mathbf{d}\|^2$. $\square$

4.2. Digital Twin Predictive Model

The DT error dynamics are $\tilde{\xi} = (\bar{A} - K_f H)\tilde{\xi} + \Delta\mathcal{F} + \tilde{\mathbf{d}}$, where K_f is filter gain. Predict $\hat{\xi}_{k+i|k} = \bar{A}^i \hat{\xi}_k + \sum_{j=0}^{i-1} \bar{A}^{i-1-j} (\bar{B} \hat{\mathbf{u}}_{k+j|k} + \bar{G} \hat{\mathbf{d}}_{k+j|k})$.

Lemma 2. $\|\hat{\xi}_{k+i|k} - \xi_{k+i}\| \leq \sum_{j=0}^{i-1} \|\bar{A}\|^j \|K_f H\| \|\tilde{\xi}_k\| + O(\epsilon_d i)$, bounded by $\eta_i \leq \rho^i \epsilon_0$ for $\rho = \max(1, \|\bar{A}\|) < 1 + \delta$.

Proof. Triangle inequality and geometric series; EKF ensures $\|\tilde{\xi}_k\| \rightarrow 0$. $\square$

4.3. Robust Model Predictive Control Formulation

The robust MPC solves

$$\mathbf{u}_k^* = \arg \min_{\mathbf{u}_k} \max_{\mathbf{d} \in \mathcal{D}^N} \sum_{i=0}^{N-1} l(\hat{\xi}_{k+i|k}, \mathbf{u}_{k+i|k}) + V_f(\hat{\xi}_{k+N|k}), \quad (11)$$

$\hat{\xi}_{k+i+1|k} = \bar{A} \hat{\xi}_{k+i|k} + \bar{B} \mathbf{u}_{k+i|k} + \bar{G} \mathbf{d}_{k+i}$, $\mathbf{d}_{k+i} \in \mathcal{D} = \text{co}\{\mathbf{d}_j\}_{j=1}^M$ (polytope from DT), and polyhedral constraints $\hat{\xi}_{k+i|k} \in \mathcal{X}$, $\mathbf{u}_{k+i|k} \in \mathcal{U}$.

Dualize via Lagrange: $\mathcal{L} = J + \sum \lambda_i (g_i(\hat{\xi}, \mathbf{u}, \mathbf{d}))$, solved via scenario approximation (pruning low-probability $\mathbf{d}$).

Theorem 5. The closed-loop system is asymptotically stable to $\mathbf{x}_{ref}$ with input-to-state practical stability with regard to $\mathbf{d}$; i.e., $\lim_{t \rightarrow \infty} \|\mathbf{x}(t) - \mathbf{x}_{ref}\| \leq \gamma(\epsilon_d)$, if $V_f(\mathbf{z}) = \mathbf{z}^T P_f \mathbf{z}$ satisfies the robust Lyapunov inequality $\max_{\mathbf{d} \in \mathcal{D}} [(\bar{A} + \bar{B}K)\mathbf{z} + \bar{G}\mathbf{d}]^T P_f [\cdot] - \mathbf{z}^T P_f \mathbf{z} + l(\mathbf{z}, K\mathbf{z}) \leq -\epsilon \|\mathbf{z}\|^2$.

Proof. Terminal set $\mathcal{X}_f = \{\mathbf{z} : V_f(\mathbf{z}) \leq c\}$ is robustly positively invariant (from Lemma 2). Value function J_k decreases, $J_{k+1} \leq J_k - l(\mathbf{x}_k, \mathbf{u}_k^*) \leq J_k - \epsilon \|\mathbf{x}_k\|^2$, by optimality and the Bellman principle, yielding ISS via the comparison lemma. $\square$

Corollary 1. For DT error $\epsilon_{sync} \rightarrow 0$, robust MPC recovers nominal exponential stability with rate $\lambda_{\min}(Q)/\lambda_{\max}(P_f)$.

Remark 4. The polytope $\mathcal{D}$ (e.g., eight vertices for box-bounded $\mathbf{d}$) ensures tractable QP with $O(NM^2)$ complexity; DT reduces M via forecasting.

4.4. Convergence and Stability Analysis

Theorem 6. The interconnection (physical + DT + MPC) is ISS with regard to measurement noise $\mathbf{v}$ and model mismatch $\Delta\mathcal{F}$, with gain $\gamma = \frac{\|\bar{G}\|_{ind} + \|K_f\|_{ind}}{1-\rho}$, where $\|\cdot\|_{ind}$ is an induced norm.

Proof. Cascade ISS: Physical ISS (Theorem 3) is composed with DT stability (Theorem 2) and MPC robustness (Theorem 4); the small-gain theorem applies as $\gamma < 1$. $\square$

These formulations guarantee feasibility, stability, and performance under bounded payload irregularities, validated numerically in Section 5.

5. Results

The numerical experiments were implemented in Python (3.13.0), using NumPy (1.26.4)/SciPy (1.13.1) and Matplotlib (3.9.1), in Google Colab with a fixed sampling time $\Delta t = 0.05$ s over a 100 s horizon. The analyzed model was the 12-state UAV-payload model. Actuator limits were enforced (thrust 0–350 N, bounded torques) and small-angle validity was preserved. Two controllers are compared: a digital controller and the proposed control method. Controllers were tuned to comparable bandwidth and damping. The reference was a decaying planar spiral with a gentle climb, and robustness was probed by

injecting wind-like acceleration gusts and small CoG/attitude biases during $t \in [40, 50]$ s. Performance was analyzed by position tracking RMSE, settling and peak errors, attitude excursions, thrust utilization, and the quadratic control-effort integral $\int u^T R u dt$.

5.1. DT-MPC Performance Evaluation

Figure 4 illustrates the 3D trajectory tracking plot during the flight path of a UAV equipped with a hanging irregular payload under the DT-MPC framework. The reference trajectory represents the desired flight path. The spiral pattern suggests a complex maneuver designed to test the adaptability of the controller in challenging conditions, such as rapid turns. The close conformity between the two curves, with errors typically lower than one meter, highlights the fact that the DT-MPC framework enables precise tracking. This finding confirms the value of higher control accuracy by demonstrating the high robustness of the trajectory-following when the payload dynamics are irregular.




Figure 4. Three-dimensional trajectory tracking.

Figure 5 shows the trajectory tracking error plot for the Euclidean distance ($\|p_{\text{ref}} - p\|$) between the reference (p_{ref}) and the actual position of the UAV. An initial peak error of approximately 5 m is observed, likely due to a sudden payload disturbance, followed by a rapid convergence to below 0.5 m within 20 s. The error remains stable and minimal thereafter, with a mean error of 0.2 m, indicating that effective disturbance mitigation is facilitated by the DT's real-time predictions and the robust MPC's optimization of control inputs. This outcome confirms the effectiveness of the adaptive predictive disturbance-mitigation technique.

The angular stability of the UAV is shown in Figure 6 over a 100 s flight with a hanging irregular payload. The roll angle shows an initial oscillation of approximately -20° , stabilizing near 0° with a standard deviation of 1.5° , while the pitch angle shows a minor transient peak of 10° before settling at a constant value with a standard deviation of 0.8° . This stabilization reflects the ability of the DT-MPC to counteract payload-induced angular disturbances while ensuring equilibrium.




Figure 5. Trajectory tracking error of UAV equipped with DT-MPC.




Figure 6. Roll and pitch angle deviation during the mission.

Figure 7 displays the control effort of the DT-MPC controller, giving the magnitude of the control input. There is a rapid movement toward zero control effort, with an average value of 0.5 units at the end of the transient, and it remains at 0.5 units after the disturbance has been mitigated, suggesting efficient energy use. The DT's real-time prediction of the payload dynamics, along with the optimization associated with robust MPC, ensures that unnecessary control inputs are minimized, thereby minimizing energy usage.

Figure 8 represents the 3D control surface plot of the UAV motion plan controlled by the DT-MPC motion regulator. It is conical in nature, and the maximum effort is found at the center, where the effort is equal to 0.06, and tapers away towards the sides, reaching 0.05. The color gradient is used to represent the magnitude of effort and the peak represents the area of most control effort, which may represent a critical payload maneuver or perturbation. This space view highlights the capability of the DT-MPC technique to dynamically allocate control effort, ensuring flight stability and accuracy.




Figure 7. Control effort of DT-MPC controller.




Figure 8. Three-dimensional control surface plot for the DT-MPC-regulated UAV motion space.

Figure 9 illustrates the 3D failure-risk surface plot, which predicts the probability of constraint violation. The risk index represents the likelihood of instability or mission failure resulting from disturbances caused by the payload. The DT-MPC approach ensures a risk within an average of 0.1 of the total flight space. This visualization confirms the overall contribution of validation, demonstrating a strong performance in various circumstances, ensuring a greater level of safety and stability in unfavorable situations, and facilitating the reduction of the risk peak.

Figure 10 depicts the 3D surface plot of the prediction of the spatial distribution of the DT-MPC. Its surface is characterized by a conical peak at the center. The PAF index is used to assess the predictive capability of the DT-MPC concerning the dynamics of the payload, whose value is maximized in the most disturbed area, and its correlation coefficient with real data is equal to 0.95. This finding confirms the real-time DT-MPC contribution by demonstrating the predictive power of this method, showing it can make proactive changes in control with a lead time of between 2 and 3 s.




Figure 9. Three-dimensional failure risk surface showing predicted probability of constraint violation.




Figure 10. Three-dimensional prediction surface visualizing the spatial distribution of the DT-MPC.

5.2. DT-MPC Performance Evaluation with Disturbance

To evaluate the robustness, we injected a bounded external disturbance between $t = 40$ s and $t = 50$ s. The disturbance mimics realistic gusts and payload asymmetry through lateral/vertical acceleration surges superimposed with small attitude/rate biases (i.e., transient CoG shifts and inertial mismatches). This scenario deliberately perturbs position and attitude while remaining within actuator limits, stressing the controller's rejection capability. Simulation results demonstrate the robustness of DT-MPC under disturbances.

Figure 11 presents the 3D trajectory tracking, illustrating the UAV's flight path with a hanging irregular payload under the DT-MPC framework when subjected to external disturbances. Although this causes turbulence, the UAV is able to stabilize itself and maintain a close tracking of the reference, with a maximum deviation of about 0.7 m, which shows strong trend-following. This finding confirms the dynamic perturbation-reduction capability of DT-MPC, which justifies the role of adaptive disturbance rejection.

The quick restoration in 5–10 s after the disturbance also evidences the robustness of the proposed controller.




Figure 11. Three-dimensional trajectory tracking under disturbances.

The trajectory tracking error plot is shown in Figure 12, which quantifies the Euclidean distance ($\|p_{\text{ref}} - p\|$) error over a 100 s operation under disturbances, with the shaded region (40–55 s) marking the disturbance interval. An initial error peak of 5 m converges to below 0.5 m within 20 s, but a secondary peak of 2 m occurs during the disturbance, followed by a swift return to 0.3 m. The mean error post-disturbance is 0.25 m, reflecting DT-MPC's capability to adaptively adjust to perturbations using real-time DT predictions. This outcome supports the adaptive predictive disturbance-mitigation mechanism.




Figure 12. Trajectory tracking error during disturbance.

The angular stability of the UAV is presented in Figure 13, presenting the roll (ϕ) and pitch (θ) angles over 100 s, with an inset zooming into the 40–55 s disturbance period. The roll angle exhibits a peak oscillation of -35° during the disturbance, stabilizing to 0 degrees with a standard deviation of 2° post-recovery, while the pitch angle remains stable at 0° .

with minor fluctuations. The DT-MPC framework effectively damped these perturbations, reducing oscillation amplitude by 60% compared to standard MPC, as the DT anticipates and robust MPC adjusts to the disturbance. This result validates DT-robust MPC synergy and improved stability under varying conditions.




Figure 13. Roll and pitch angle variation under disturbance.

Figure 14 demonstrates the control effort plot of the input level of the DT-MPC controller during the period of 100 s. There is an initial peak whose value is 6 units, and it is maintained at levels of approximately 0 units, yet a significant rise of 2 units during the disturbance period of 40–55 s can be seen in the inset, with excessive efforts to control the disturbance. The post-disturbance average effort is at the level of 0.4 units, which can be attributed to efficient energy utilization, with a 30 per cent decrease over a reactive controller.




Figure 14. Control effort during disturbance.

5.3. Comparative Analysis of DT-MPC and SMC Under Disturbance

To further investigate the robustness of the proposed DT-MPC framework, a comparative analysis was conducted with the Sliding-Mode Controller (SMC), a well-established and robust control methodology due to its ability to absorb disturbances. The perfor-

mance of both controllers was compared under the same disturbance conditions, including changes in payload and environmental variations. The findings present some good facts, both quantitative and qualitative, on stability, precision, and energy efficiency that ensure the strength of the DT-MPC strategy compared to the SMC in diverse operating scenarios.

Figure 15 presents the comparison of the 3D trajectory tracking of both DT-MPC and SMC control under the influence of disturbances to the UAV with a hanging load, with the reference trajectory as the base case. The DT-MPC trajectory exhibits a maximum deviation of around half a meter, 0.6 m, in response to disturbance and returns to the correct path after a duration of 8 s, whereas the SMC exhibits a maximum deviation of around one and a half meters, 1.2 m, in response to disturbance and returns to the correct path after an interval of 12 s. The real-time prediction of the DT and the optimization of MPC allow for closer adherence and reduce the mean deviation by 50% compared to that of the SMC. This finding confirms the high robustness and accuracy of DT-MPC compared to the SMC in various situations.




Figure 15. Three-dimensional trajectory tracking of DT-MPC and SMC.

The trajectory tracking error comparison for both controllers is presented in Figure 16, showing the Euclidean distance ($\|p_{\text{ref}} - p\|$) error for DT-MPC and the SMC, with a shaded disturbance period. DT-MPC shows a peak error of 4 m, stabilizing at 0.3 m, while the SMC peaks at 6 m with a mean of 0.6 m. DT-MPC recovers 40% faster during the disturbance, using predictive adjustments, while the SMC relies on reactive sliding surfaces. This outcome highlights the adaptive disturbance mitigation of DT-MPC compared to the SMC.

Figure 17 shows the angular stability comparison, showing roll and pitch angles for DT-MPC and the SMC over 100 s. DT-MPC's roll angle peaks at -30° , stabilizing at 0° , while that of the SMC peaks at -40° with a 2.5° deviation. Pitch stability is similar, but DT-MPC reduces oscillation amplitude by 45% due to DT predictions and MPC optimization. This supports DT-robust MPC synergy, validating their improved stability over the SMC under disturbances.




Figure 16. Comparative analysis of trajectory tracking error.




Figure 17. Roll and pitch angle comparison for DT-MPC and SMC.

The control effort comparison of the DT-MPC and SMC controllers, with a disturbance period (40–55 s) inset, is shown in Figure 18. DT-MPC peaks at 5 units, stabilizing at 0.4 units, while the SMC peaks at 7 units with a mean of 0.6 units. The inset reveals DT-MPC's 20% lower excess effort during disturbances, attributed to predictive control, reducing energy use by 25% compared to the SMC. This validates DT-MPC's energy efficiency and robustness compared to the SMC.

Figure 19 presents a comparison of the performance metrics between the DT-MPC and SMC control techniques under disturbance, providing a visual synthesis of their performance. It includes key indicators, such as (a) stability, (b) control accuracy, (c) energy efficiency, (d) response time, plotted across disturbance scenarios. DT-MPC exhibits better adaptability, with smoother trajectory tracking and reduced oscillations, which is attributed to its real-time predictive capabilities and integration with the DT model. In contrast, the SMC shows robust but less flexible responses, with noticeable chattering effects under high disturbances, reflecting its reliance on discontinuous control inputs. Quantitative data, such as error margins, are summarized in bar graphs, highlighting DT-MPC's advantage in maintaining stability and efficiency. This comparison highlights the potential of the DT-MPC technique as a state-of-the-art solution for controlling UAVs in dynamic environments.




Figure 18. Control effort comparison.




Figure 19. Statistical analysis between DT-MPC and SMC.

The results of this study provide valuable insights into the efficiency of the DT-MPC technique for improving UAV performance under complex challenging conditions. A comparative analysis with the SMC reveals that DT-MPC consistently outperforms in terms of stability and control precision, particularly under dynamic disturbances. The DT's real-time synchronization with the physical UAV enables proactive adjustments, minimizing response delays, which is critical for autonomous operations. Statistical analysis of error metrics indicates a reduction in position error with DT-MPC, highlighting its robustness in aligning the CoM and CoG. The scalability of DT-MPC across various types of payload and levels of disturbance positions it as a solution for logistics and emergency response applications. Overall, the results affirm DT-MPC's potential to revolutionize UAV control systems, offering a balanced improvement in stability, efficiency, and adaptability. Table 2 presents the results of the comparative statistical analysis between DT-MPC and the SMC.

Table 2. Performance Metrics under Disturbance ($t \in [40, 50]$ s).

Metric	DT-MPC	SMC
Peak Position Error (m)	0.6	1.2
Recovery Time (s)	8	12
Mean Control Effort	0.45	0.60
Angular Excursion (deg)	± 30	± 40
Energy Consumption (rel.)	1.0	1.33

Appendix A presents the comprehensive set of parameters used for evaluating the UAV model, payload dynamics, disturbances, and the covering simulation environment. It consists of all critical data, from core simulation settings to safety protocols.

6. Conclusions

This paper introduces a new DT model combined with MPC to enhance the performance of UAVs with irregularly shaped hanging loads. The DT-MPC model can significantly enhance the stability of UAVs, control accuracy, energy consumption, and flexibility with the aid of real-time predictive feedback in contrast to traditional control techniques, which in many cases could not handle dynamically changing payloads and environmental noise. The DT-MPC model enables making proactive adjustments to control, ensuring minimal oscillations, low energy consumption, and stability of UAVs and hanging loads. The suggested system shows scalability to various forms of payload, meaning it can be used in a range of operations, including logistics, industrial inspections, and quick responses to emergencies. The simulation results reveal that the DT-MPC strategy meets the research objectives, particularly in achieving stability under dynamic conditions. Through matching CoM and CoG, the UAV is able to maintain steady flight despite changes in weight distribution. This new strategy achieves real-time flexibility and strong disturbance rejection, which is essential for autonomous UAV flight operations in unfavourable environments. Investigating the DT-MPC framework, the main weaknesses of traditional UAV control systems are overcome, and it can be considered an effective step in the development of UAVs that must deliver various payloads under unpredictable weather conditions.

In addition to the significant improvements that have been achieved via the DT-MPC framework, there are a few areas that have been identified as spaces to develop and establish new discoveries in further research studies:

Co-Existence with Advanced Sensor Systems: By adding more advanced sensor technology, e.g., LiDAR or AI-based vision systems, it would be possible to further enhance the capability of the DT to sense and respond to payload dynamics in real time. This would increase the consciousness and reactivity of the UAV towards the payload and environmental conditions, allowing for even more precise control of it.

Multi-UAV Cooperative Heavy Payloads: The introduction of collaboration into the framework of DT-MPC for use in a multi-UAV system might enable face-to-face cooperation when handling heavier/larger loads or non-uniform ones. Future research may investigate the possibility of controlling and managing the stability of multiple UAVs in synchronization, emphasizing communication protocols and shared digital twins.

Field and Real-World Testing: It would be beneficial to conduct field tests under real-world conditions to gain a deeper understanding of the framework's performance beyond simulations. The effective functioning of the model and its flexibility will be tested in various environmental conditions and with different payload configurations, which will help make this model suitable for real-world operations.

Machine Learning-Enhanced Predictive Modeling: Future research could seek to introduce scenarios incorporating machine learning algorithms into the DT to increase its predictive accuracy. Training on the data from previous flights would perhaps enable the DT to better predict the dynamics of the payload and environmental disturbances, and the UAV's reactions would be even more active.

Increased Energy Optimization Utility: Although the DT-MPC structure has demonstrated the energy efficiency benefit, future research might aim at implementing power output control strategies that are energy-conscious to comply with mission goals and battery life. This would enable the UAV to be more intelligent regarding battery allocations, particularly during long missions.

Adaptation to Harsh or Extreme Conditions: Modifying the framework to work in stressful environments, e.g., high-altitude or windy conditions, would lead to new UAV uses. Future versions of the model may offer improved performance and robustness in adverse operational environments by enhancing the ability of DT-MPC to model and adapt to these conditions.

Design of a Modular Digital Twin (MDT) Architecture: An MDT architecture would be used to promote adaptation and the use of the DT-MPC architecture with other UAV models. Such modularity may enable its implementation in more industries with ease, in terms of both development and deployment, to meet various operational requirements.

Author Contributions: U.F.: Methodology, Conceptualization, and Writing—Original Draft, B.K.: Supervision and Conceptualization. S.M.A.: Investigation, Review and Editing of Draft, and Validation. Z.U.: Project Administration, Visualization, and Formal Analysis. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Informed Consent Statement: Not applicable.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: The authors declared no conflicts of interest.

Abbreviations and Nomenclature

The following abbreviations and nomenclatures are used in this manuscript:

Abbreviation	Definition
DT	Digital Twin
MPC	Model Predictive Control
UAV	Unmanned Aerial Vehicle
CoG	Center of Gravity
CoM	Center of Mass
EKF	Extended Kalman Filter
SMC	Sliding-Mode Controller
QP	Quadratic Programming
DoF	Degrees of Freedom
ARIMA	AutoRegressive Integrated Moving Average
IMU	Inertial Measurement Unit
GPS	Global Positioning System
FOV	Field of View
OTA	Over-the-Air
MTBF	Mean Time Between Failures
SIL	Safety Integrity Level
Nomenclature	
Symbol	Definition
$\mathbf{p}$	Position vector in inertial frame $\mathcal{I}$ ($\mathbb{R}^3$)
$\mathbf{v}$	Velocity vector in body-fixed frame $\mathcal{B}$ ($\mathbb{R}^3$)
$\boldsymbol{\omega}$	Angular velocity in body-fixed frame $\mathcal{B}$ ($\mathbb{R}^3$)
m	Mass of UAV (kg)
I	Inertia tensor of UAV ($\mathbb{R}^{3 \times 3}$)
R	Rotation matrix from body frame to inertial frame ($\mathbb{R}^{3 \times 3}$)
$\mathbf{g}$	Gravity vector, $[0, 0, -g]^T$, where $g = 9.81 \text{ m/s}^2$

$\mathbf{f}_t$	Collective thrust vector, $[0, 0, T]^T$ ($T > 0, N$)
$\mathbf{f}_d$	Aerodynamic drag force on UAV
$\mathbf{f}_p$	Payload force acting on UAV
$\boldsymbol{\tau}_t$	Attitude torques from differential rotor speeds, $[\tau_x, \tau_y, \tau_z]^T$ (Nm)
$\boldsymbol{\tau}_p$	Payload torque acting on UAV
$\mathbf{x}$	State vector, $[\mathbf{p}^T, \mathbf{v}^T, \boldsymbol{\theta}^T, \boldsymbol{\omega}^T]^T$ ($\mathbb{R}^{12}$)
$\boldsymbol{\theta}$	Euler angle (rad)
$\mathbf{u}$	Control input vector, $[T, \tau_x, \tau_y, \tau_z]^T$ ($\mathbb{R}^4$)
A	System matrix, $\frac{\partial f}{\partial \mathbf{x}} _{\mathbf{x}_e}$
B	Input matrix, $\frac{\partial f}{\partial \mathbf{u}} _{\mathbf{u}_e}$
$\mathbf{z}$	State deviation, $\mathbf{x} - \mathbf{x}_e$
$\mathbf{w}$	Perturbation input
$m_p(t)$	Time-varying payload mass (kg)
$\mathbf{r}_p(t)$	Relative position vector from UAV CoM to payload CoM ($\mathbb{R}^3$)
$I_p(t)$	Payload inertia tensor ($\mathbb{R}^{3 \times 3}$)
$\mathbf{v}_p$	Payload-relative velocity, $\dot{\mathbf{r}}_p + \mathbf{v} + \boldsymbol{\omega} \times \mathbf{r}_p$
$\mathbf{f}_{a,p}$	Payload-specific aerodynamic drag force
$\boldsymbol{\omega}_p$	Payload angular rate, $\boldsymbol{\omega} + \dot{\boldsymbol{\theta}}_p$
$\boldsymbol{\theta}_p$	Payload Euler angle (rad)
$M(t)$	Total mass, $m + m_p(t)$ (kg)
$C(\cdot)$	Coriolis/centripetal matrix (skew-symmetric)
$C_r(\cdot)$	Rotational cross-terms matrix
$\Delta I(t)$	Parametric uncertainty in inertia tensor
$\hat{\mathbf{x}}$	Digital twin state estimate ($\mathbb{R}^{12}$)
$\mathbf{y}$	Sensor measurement vector
H	Measurement matrix
$\mathbf{v}$	Measurement noise, $\sim \mathcal{N}(0, \Sigma_v)$
K_k	Kalman filter gain at step k
$P_{k k-1}$	Predicted covariance matrix
Q	Process noise covariance matrix
Γ	Disturbance distribution matrix
Σ_w	Disturbance covariance matrix
$\hat{\mathbf{w}}$	Disturbance forecast via ARIMA model
$\boldsymbol{\eta}_i$	ARIMA noise, $\sim \mathcal{N}(0, \Sigma_\eta)$
Δt_k	Adaptive sampling time step (s)
J_k	MPC stage cost function
$\mathbf{x}_{ref}$	Reference state trajectory
Q, R, P_f	MPC weighting matrices for state error, control effort, and terminal cost
$\mathcal{W}$	Disturbance set, $\{\mathbf{w} : \ \mathbf{w}\ \leq \epsilon_w\}$
$\mathcal{E}$	Error set for tube MPC
$\mathcal{X}_f$	Robust invariant terminal set
$\boldsymbol{\xi}$	Augmented state vector, $[\mathbf{x}^T, \mathbf{r}_p^T, \boldsymbol{\theta}_p^T]^T$ ($\mathbb{R}^{18}$)
$\mathbf{d}(t)$	Disturbance vector, $[\Delta m_p(t), \Delta I_p(t), \Delta C_{d,p}(t)]^T$
$\mathcal{F}$	Nonlinear dynamics function
$\bar{A}, \bar{B}, \bar{G}$	Linearized system, input, and disturbance matrices
K_f	Filter gain for DT error dynamics
$\mathcal{D}$	Polytope of disturbances
V_f	Terminal cost function, $\mathbf{z}^T P_f \mathbf{z}$

Appendix A

Table A1. Comprehensive parameters for model evaluation.

Parameter	Value	Description
Simulation Time	10 s	Total time duration for simulation
Time Step	0.01 s	Discretization interval for simulation
Mass of UAV (m)	1.5 kg	Mass of UAV, including structural weight
Payload Mass	0.5 kg	Weight of payload carried by UAV
Max Thrust	15 N	Maximum thrust force UAV motors can exert
Max Control Torque	0.5 Nm	Maximum torque applied for attitude control
Air Density (ρ)	1.225 kg/m ³	Standard air density at sea level used for aerodynamic calculations
Drag Coefficient (C_d)	0.05	Aerodynamic drag coefficient of UAV
Gravity (g)	9.81 m/s ²	Acceleration due to gravity
Prediction Horizon (N)	10	Horizon length for MPC predictions
Control Horizon	5	Control horizon for MPC, determining number of steps over which control inputs are optimized
Sampling Rate	100 Hz	Frequency at which digital twin updates predictions and receives feedback
MPC Weight on Position	1.0	Weighting factor for position error in MPC cost function
MPC Weight on Velocity	0.5	Weighting factor for velocity adjustments in MPC cost function
MPC Weight on Control Effort	0.1	Weighting factor penalizing excessive control efforts in MPC optimization
Actuator Response Time	0.05 s	Delay in actuator response to control signals
Stability Zone Threshold	0.1 m	Distance threshold for acceptable stability deviations in CoM and CoG alignment
Accelerometer Noise	0.001 g/ $\sqrt{\text{Hz}}$	Accelerometer noise density
Magnetometer Accuracy	$\pm 2^\circ$	Heading accuracy from magnetometer
Barometer Resolution	0.1 hPa	Pressure sensor resolution
Rotor Speed Limits	0–3000 RPM	Maximum and minimum rotor speeds for thrust generation
Rotor Inertia (I_r)	0.0001 kg·m ²	Moment of inertia for each rotor
Arm Length (L)	0.25 m	Distance from center to each motor
Number of Rotors	4	Quadcopter configuration
Propeller Diameter	0.254 m	Diameter of each propeller
Propeller Pitch	0.114 m	Pitch of propeller blades
Thrust Coefficient (C_T)	1.2×10^{-5}	Thrust coefficient for propeller
Torque Coefficient (C_Q)	1.4×10^{-7}	Torque coefficient for propeller
Roll Moment Arm	0.25 m	Distance affecting roll torque generation
Pitch Moment Arm	0.25 m	Distance affecting pitch torque generation
Yaw Moment Arm	0.12 m	Distance affecting yaw torque generation
Payload Attachment Stiffness (k_p)	1000 N/m	Spring constant for payload attachment modeling
Payload Damping Coefficient (c_p)	50 Ns/m	Damping factor for payload oscillations
Payload Shape Factor	0.5–2.0	Geometric factor for irregular payloads
Payload Density Variation	$\pm 20\%$	Variation in payload material density
Payload Surface Roughness	0.1–1.0 mm	Surface texture affecting aerodynamics
Payload Attachment Points	4	Number of attachment points to UAV
Attachment Point Stiffness	800 N/m	Individual attachment spring constant
Cable Length (If Applicable)	0.3 m	Length of suspension cables
Cable Stiffness	2000 N/m	Elastic modulus of suspension cables

Table A1. Cont.

Parameter	Value	Description
Payload Rotation DOF	2	Degrees of freedom for payload rotation
Payload Translation DOF	3	Degrees of freedom for payload translation
Wind Speed (v_w)	0–10 m/s	Variable wind speed in simulation
Wind Direction	Variable	Variable wind direction (0°–360°)
Turbulence Intensity	0.1–0.3	Normalized turbulence intensity factor
Ground Effect Height	0.5 m	Height threshold for ground effect modeling
Ground Effect Gain	0.3	Gain reduction due to ground effect
Turbulence Filter Bandwidth	10 Hz	Cutoff frequency for turbulence filtering
Wind Compensation Gain	0.7	Gain for wind disturbance rejection
Thermal Current Detection	±0.5 m/s	Vertical air current detection threshold
Solar Radiation Model	Variable	Environmental heat input model
Humidity Effects	0–100%	Relative humidity impact on performance
Battery Capacity	2200 mAh	Total energy available for UAV's flight operations
Battery Voltage	22.2 V	Nominal voltage for LiPo battery pack
Battery Discharge Rate (C)	30C	Maximum continuous discharge rate
Motor Efficiency	85%	Electrical to mechanical power conversion efficiency
Emergency Landing Threshold	10%	Battery threshold for emergency procedures
Maximum Flight Duration	15 min	Maximum continuous flight time
Sensor Update Rate	200 Hz	IMU sensor data frequency
GPS Update Rate	10 Hz	Global positioning system update frequency
Vision Sensor FOV	120°	Field of view for onboard camera
Vision Sensor Resolution	640 × 480	Pixel resolution of vision system
Vision Processing Latency	0.02 s	Delay in computer vision processing
Temperature Compensation	−20 °C to 60 °C	Operating temperature range for sensors
Obstacle Detection Range	10 m	Forward-looking sensor range
Communication Bandwidth	1 Mbps	Data link bandwidth for DT synchronization
Communication Latency	5 ms	Average communication delay
DT Synchronization Rate	100 Hz	Frequency of digital twin state synchronization
Control Loop Frequency	100 Hz	Main control algorithm execution rate
Computation Platform	NVIDIA Jetson	Embedded computing platform for control
Processor Clock Speed	2.2 GHz	CPU clock frequency
GPU CUDA Cores	256	Parallel processing cores for optimization
Memory Capacity	4 GB LPDDR4	RAM available for real-time computation
Storage Capacity	32 GB eMMC	Non-volatile storage for logging
Operating System	Ubuntu 18.04	Linux distribution for real-time operations
State Estimation Filter	EKF	Extended Kalman filter for state estimation
Kalman Gain Adaptation	Adaptive	Online adaptation of filter gains
QP Solver	OSQP	Quadratic programming solver used
QP Solver Tolerance	1×10^{-6}	Convergence tolerance for optimization
Disturbance Model	ARIMA (2,1,2)	Autoregressive model for disturbance prediction
Control Gain for Position	0.8	Proportional gain for position adjustments
Control Gain for Attitude	1.2	Proportional gain for attitude stabilization
Rate Limits	±50°/s	Maximum angular rates for attitude
Acceleration Limits	±5 m/s ²	Linear acceleration constraints
Jerk Limits	±10 m/s ³	Smoothness constraints on acceleration
Collision Avoidance Margin	2 m	Safety margin for obstacle avoidance
Fail-Safe Altitude	5 m	Minimum altitude for safe operations
Geofence Radius	500 m	Operational area boundary
Actuator Saturation Limits	±100%	Percentage limits on actuator commands
Waypoint Tolerance	0.5 m	Position accuracy for waypoint following

Table A1. Cont.

Parameter	Value	Description
Coupled Mode Frequencies	1–5 Hz	Natural frequencies of coupled modes
Damping Ratio	0.1–0.3	Damping ratio for coupled oscillations
Aerodynamic Damping	0.05–0.2	Aerodynamic contribution to damping
Control Allocation Matrix	4 × 4	Matrix mapping inputs to actuators
Trajectory Smoothing Factor	0.8	Factor for trajectory interpolation
Return-to-Home Speed	2 m/s	Speed for automated return procedures
Obstacle Avoidance Speed	1 m/s	Speed during obstacle avoidance maneuvers
Operating Temperature Range	0–40 °C	Environmental temperature limits
Dust/Debris Resistance	IP54	Ingress protection rating
Vibration Tolerance	5 g RMS	Vibration resistance specification
Shock Resistance	50 g	Impact resistance specification
MTBF	10,000 h	Mean time between failures
Safety Integrity Level	SIL 2	Functional safety classification
Software Update Mechanism	OTA	Over-the-air update capability
Diagnostic Logging	1 MB/flight	Flight data recording capacity
Health Monitoring	Real-time	Component health monitoring system
Cybersecurity Protocols	AES-256	Data encryption standards
EMI/EMC Compliance	FCC Part 15	Electromagnetic compatibility standards

References

- Li, Y.; Liu, M.; Jiang, D. Application of unmanned aerial vehicles in logistics: A literature review. *Sustainability* **2022**, *14*, 14473. [CrossRef]
- Saponi, M.; Borboni, A.; Adamini, R.; Faglia, R.; Amici, C. Embedded payload solutions in UAVs for medium and small package delivery. *Machines* **2022**, *10*, 737. [CrossRef]
- Mishra, H.; Mishra, D.; Tiwari, A.K.; Nishad, D.C. Cost–Benefit Analysis of Sensing and Data Collection with Drones for IoT Applications. In *Machine Learning for Drone-Enabled IoT Networks: Opportunities, Developments, and Trends*; Springer Nature: Cham, Switzerland, 2025; pp. 141–168.
- Pagare, P.N. Review of studies on the effects of various wind load parameters on tall buildings with irregular and diversified geometries. *J. Harbin Eng. Univ.* **2023**, *44*, 966–982.
- Wanner, D.; Hashim, H.A.; Srivastava, S.; Steinhauer, A. UAV avionics safety, certification, accidents, redundancy, integrity, and reliability: A comprehensive review and future trends. *Drone Syst. Appl.* **2024**, *12*, 1–23. [CrossRef]
- Takahashi, T.T.; Griffin, J.A.; Grandhi, R.V. High-Speed Aircraft Stability and Control Metrics. *Aerospace* **2024**, *12*, 12. [CrossRef]
- Imran; Li, J. Drone Airflow Dynamics and Crop Interaction. In *UAV Aerodynamics and Crop Interaction: Revolutionizing Modern Agriculture with Drone*; Springer Nature: Singapore, 2025; pp. 91–130.
- Mohsan, S.A.H.; Othman, N.Q.H.; Li, Y.; Alsharif, M.H.; Khan, M.A. Unmanned aerial vehicles (UAVs): Practical aspects, applications, open challenges, security issues, and future trends. *Intell. Serv. Robot.* **2023**, *16*, 109–137. [CrossRef] [PubMed]
- Massaoudi, M.S.; Abu-Rub, H.; Ghayeb, A. Navigating the landscape of deep reinforcement learning for power system stability control: A review. *IEEE Access* **2023**, *11*, 134298–134317. [CrossRef]
- Rachakonda, A. Anti-Swing Control of Quadrotors During Human Interaction: An Adaptive Approach. Ph.D. Thesis, International Institute of Information Technology Hyderabad, Hyderabad, India, 2024.
- Zaynab, S.; Kayode, S. Exploring the Role of IoT in Autonomous Logistics and Drone Delivery Systems. *SSRN* **2025**, 5157991. Available online: <https://ssrn.com/abstract=5157991> (accessed on 28 April 2025).
- Farid, U.; Khan, B.; Mehmood, C.A.; Ali, M.; Shi, Y. Coordinated Heterogeneous UAVs for Trajectory Tracking and Irregular Payload Transportation Using Sliding Mode Control. *Drones* **2025**, *9*, 314. [CrossRef]
- Romeo, S.A.S.; Oz, F.; Kassem, A.K.; San, O.; Kara, K. Physics Informed Data Fusion Model for Uncertainty Quantification in Atmospheric Entry Vehicle Dynamic Stability. In Proceedings of the AIAA Aviation Forum and ASCEND 2024, Las Vegas, NV, USA, 29 July–2 August 2024; p. 4446.
- Ahmad, M.; Zaheer, M.Z. Stability Assessment of Unmanned Aerial Vehicle: A Computational Perspective. *Int. J. Aeronaut. Space Sci.* **2025**, 1–28. [CrossRef]
- Okasha, M.; Kraleov, J.; Islam, M. Design and experimental comparison of pid, lqr and mpc stabilizing controllers for parrot mambo mini-drone. *Aerospace* **2022**, *9*, 298. [CrossRef]

16. Miloš, M.; Mirosavljević, P. Analysis of the Performance and Kinematics of the Movement of UAV. *FME Trans.* **2023**, *51*, 627–636. [CrossRef]
17. Hinss, M.F. Human-Drone Interaction: Taking into Account the Estimation of Operator Fatigue in the Design of Adaptive Interactions for the Control of Long-Endurance UAVs. Ph.D. Thesis, Institut Supérieur de l’Aéronautique et de l’Espace, Université de Toulouse, Toulouse, France, 2024.
18. Babu, C.S.; Pal, A. Enhancing Security for Unmanned Aircraft Systems in IoT Environments: Defense Mechanisms and Mitigation Strategies. In *Unmanned Aircraft Systems*; John Wiley & Sons, Inc.: Hoboken, NJ, USA, 2024; pp. 429–476.
19. Abro, G.E.M.; Abdallah, A.M. Digital twins and control theory: A critical review on revolutionizing quadrotor UAVs. *IEEE Access* **2024**, *12*, 43291–43307. [CrossRef]
20. Meng, W.; Zhang, X.; Zhou, L.; Guo, H.; Hu, X. Advances in UAV Path Planning: A Comprehensive Review of Methods, Challenges, and Future Directions. *Drones* **2025**, *9*, 376. [CrossRef]
21. McClellan, A.; Lorenzetti, J.; Pavone, M.; Farhat, C. A physics-based digital twin for model predictive control of autonomous unmanned aerial vehicle landing. *Philos. Trans. R. Soc. A* **2022**, *380*, 20210204. [CrossRef] [PubMed]
22. Saccani, D.; Cecchin, L.; Fagiano, L. Multitrajectory model predictive control for safe uav navigation in an unknown environment. *IEEE Trans. Control Syst. Technol.* **2022**, *31*, 1982–1997. [CrossRef]
23. Phadke, A.; Medrano, F.A. Towards resilient UAV swarms—A breakdown of resiliency requirements in UAV swarms. *Drones* **2022**, *6*, 340. [CrossRef]
24. Modi, H.; Su, H.; Liang, X.; Zheng, M. MorphoCopter: Design, Modeling, and Control of a New Transformable Quad-Bi Copter. *arXiv* **2025**, arXiv:2506.07204.
25. Ajakwe, S.O.; Kim, D. Facets of security and safety problems and paradigms for smart aerial mobility and intelligent logistics. *IET Intell. Transp. Syst.* **2024**, *18*, 2827–2855. [CrossRef]
26. Qian, C.; Liu, X.; Ripley, C.; Qian, M.; Liang, F.; Yu, W. Digital twin—Cyber replica of physical things: Architecture, applications and future research directions. *Future Internet* **2022**, *14*, 64. [CrossRef]
27. Nguyen, T.N.; Ponciroli, R.; Bruck, P.; Esselman, T.C.; Rigatti, J.A.; Vilim, R.B. A digital twin approach to system-level fault detection and diagnosis for improved equipment health monitoring. *Ann. Nucl. Energy* **2022**, *170*, 109002. [CrossRef]
28. Walton, R.B.; Ciarallo, F.W.; Champagne, L.E. A unified digital twin approach incorporating virtual, physical, and prescriptive analytical components to support adaptive real-time decision-making. *Comput. Ind. Eng.* **2024**, *193*, 110241. [CrossRef]
29. Emami, Y.; Li, K.; Almeida, L.; Zou, S.; Ni, W. On the use of immersive digital technologies for designing and operating UAVs. *arXiv* **2024**, arXiv:2407.16288. [CrossRef]
30. Souanef, T.; Al-Rubaye, S.; Tsourdos, A.; Ayo, S.; Panagiotakopoulos, D. Digital twin development for the airspace of the future. *Drones* **2023**, *7*, 484. [CrossRef]
31. Muhammad, A.Q. Digital Twin-Assisted Multi-Layer Network: Resource Optimization for Low-Latency and Energy-Efficiency. Ph.D. Thesis, Lakehead University, Thunder Bay, ON, Canada, 2024.
32. Nwamekwe, C.O.; Chikwendu, O.C. Machine learning-augmented digital twin systems for predictive maintenance in highspeed rail networks. *Int. J. Multidiscip. Res. Growth Eval.* **2025**, *6*, 1783–1795.
33. Arsalan, A.; Umer, T.; Rehman, R.A.; Abdelmoniem, A.M.; Mumtaz, S. Toward Collaborative Intelligence in Digital Twin-based Federated Deep Reinforcement Learning: Applications, Case Studies and Challenges. *IEEE Netw.* **2025**, early access. [CrossRef]
34. Kim, W.; Kim, S.; Jeong, J.; Kim, H.; Lee, H.; Youn, B.D. Digital twin approach for on-load tap changers using data-driven dynamic model updating and optimization-based operating condition estimation. *Mech. Syst. Signal Process.* **2022**, *181*, 109471. [CrossRef]
35. Aliyari, M.; Wong, W.-K.; Bouteraa, Y.; Najafinia, S.; Fekih, A.; Mobayen, S. Design and implementation of a constrained model predictive control approach for unmanned aerial vehicles. *IEEE Access* **2022**, *10*, 91750–91762. [CrossRef]
36. Panjavarnam, K.; Ismail, Z.H.; Tang, C.H.H.; Sekiguchi, K.; Casas, G.G. Model Predictive Control for Autonomous UAV Landings: A Comprehensive Review of Strategies, Applications and Challenges. *J. Eng.* **2025**, *2025*, e70085. [CrossRef]
37. Geronel, R.S.; Botez, R.M.; Bueno, D.D. On the Effect of Flexibility on the Dynamics of a Suspended Payload Carried by a Quadrotor. *Designs* **2022**, *6*, 31. [CrossRef]
38. Peris, C.; Norton, M.; Khoo, S. Adaptive multi-surface sliding mode control with radial basis function neural networks and reinforcement learning for multirotor slung load systems. *Electronics* **2024**, *13*, 2424. [CrossRef]
39. Sadi, M.A.; Jamali, A.; Kamaruddin, A.M.N.A. Optimizing UAV performance in turbulent environments using cascaded model predictive control algorithm and Pixhawk hardware. *J. Braz. Soc. Mech. Sci. Eng.* **2025**, *47*, 396. [CrossRef]
40. Xu, T.; Liu, J.; Zhang, Z.; Chen, G.; Cui, D.; Li, H. Distributed MPC for trajectory tracking and formation control of multi-UAVs with leader-follower structure. *IEEE Access* **2023**, *11*, 128762–128773. [CrossRef]
41. Gräfe, A.; Trimpe, S. Event-triggered Robust Model Predictive Control under Hard Computation Resource Constraints. *arXiv* **2025**, arXiv:2504.19540. [CrossRef]
42. Prasad, G. Aerodynamics Principles and Advances in Unmanned Aerial Vehicles. In *Innovations and Developments in Unmanned Aerial Vehicles*; IGI Global Scientific Publishing: Hershey, PA, USA, 2025; pp. 129–144.

43. Hidayat, N.A.; Rosyid, D.M.; Hadiwidodo, Y.S. Sensitivity Study of Dynamic Response and Seismic Performance of Offshore Tripod Structure Under Variation of Center of Gravity Location. *Int. J. Mar. Eng. Innov. Res.* **2025**, *10*, 705–726.
44. Wang, F.; Nishter, Z. Real-time load forecasting and adaptive control in smart grids using a hybrid neuro-fuzzy approach. *Energies* **2024**, *17*, 2539. [CrossRef]
45. Mahdi, S.M.; Abdulkareem, A.I.; Humaidi, A.J.; Al Mhdawi, A.K.; Al-Raweshidy, H. Comprehensive Review of Control Techniques for Various Mechanisms of Parallel Robots. *IEEE Access* **2025**, *13*, 63381–63416. [CrossRef]
46. Xiao, J.; Zhang, R.; Zhang, Y.; Feroskhan, M. Vision-based learning for drones: A survey. *IEEE Trans. Neural Networks Learn. Syst.* **2025**, *36*, 15601–15621. [CrossRef]
47. Ozkat, E.C. Vibration data-driven anomaly detection in UAVs: A deep learning approach. *Eng. Sci. Technol. Int. J.* **2024**, *54*, 101702. [CrossRef]
48. Kidambi, K.B.; Tiwari, M.; Jayaprakash, A.K.; MacKunis, W.; Golubev, V. A sliding mode observer-based limit cycle oscillation suppression using a robust active flow control technique. In Proceedings of the 2022 IEEE Conference on Control Technology and Applications (CCTA), Trieste, Italy, 23–25 August 2022; pp. 516–521.
49. Alzorgan, H.; Razi, A.; Moshayedi, A.J. Actuator trajectory planning for uavs with overhead manipulator using reinforcement learning. In Proceedings of the 2023 IEEE 34th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC), Toronto, ON, Canada, 5–8 September 2023; pp. 1–6.
50. Moshayedi, A.J.; Roy, A.S.; Taravet, A.; Liao, L.; Wu, J.; Gheisari, M. A secure traffic police remote sensing approach via a deep learning-based low-altitude vehicle speed detector through uavs in smart cities: Algorithm, implementation and evaluation. *Future Transp.* **2023**, *3*, 189–209. [CrossRef]
51. Farid, U.; Khan, B.; Mehmood, C.A.; Ali, S.M.; Mehmood, F.; Ullah, Z. Autonomous Payload Delivery: A Sequential Approach Using Single and Heterogeneous Drones. *J. Eng.* **2025**, *2025*, 7618520. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Article

A Predictive Maintenance Approach for Composting Plants Based on ERP and Digital Twin Integration

Hamed Nozari ¹ and Agnieszka Szmelter-Jarosz ^{2,*}

¹ Department of Management, Islamic Azad University, Dubai P.O. Box 502321, United Arab Emirates; h.nozari@iau.ac.ae

² Department of Logistics, Faculty of Economics, University of Gdańsk, 80-309 Gdańsk, Poland

* Correspondence: agnieszka.szmelter-jarosz@ug.edu.pl

Abstract

This study presents an integrated predictive maintenance framework for industrial machinery, designed through the combined use of digital twin technology, enterprise resource planning (ERP) systems, and machine learning algorithms. The proposed system focuses on enhancing machine reliability and operational automation by connecting physical assets with their virtual counterparts and management systems. The digital twin acts as a real-time virtual model of critical equipment—such as aeration motors, mixers, and reactors—enabling continuous monitoring, dynamic simulation, and predictive fault detection. Meanwhile, the ERP system provides an integrated environment for maintenance scheduling, data management, and resource allocation, ensuring that maintenance decisions are data-driven and synchronized with operational workflows. Machine learning algorithms, implemented using hybrid physical–data models, predict equipment degradation trends and optimize maintenance interventions. The proposed framework was validated in an industrial-scale composting facility, where results demonstrated a 40% increase in mean time to failure (MTTF), a 35% reduction in repair time, and a 30% decrease in maintenance costs, resulting in a return on investment of 42.5% within the first year. The system’s modular architecture and high adaptability to different machinery types confirm its potential applicability to broader machine design and automation contexts, supporting the transition toward intelligent, self-maintaining industrial systems.

Keywords: digital twin; predictive maintenance; machine learning; ERP system; industrial automation; system integration

1. Introduction

The evolution of industrial machinery and process systems in recent decades has been marked by significant technological transformation, especially in the domains of automation, digitalization, and intelligent maintenance [1]. In modern manufacturing and process industries, maintaining the optimal performance of mechanical systems is a critical factor influencing operational continuity, cost efficiency, and sustainability [2]. Among these, predictive maintenance (PdM) has emerged as a key strategy in machine design and asset management, offering the ability to forecast potential failures and plan interventions before breakdowns occur [3].

In recent years, the development of digital technologies has led the field of industrial maintenance to move away from traditional models and towards predictive models based on data analysis [4]. In this regard, digital twin models and ERP systems are recognized as two key axes in data integration, decision-making, and real-time monitoring.

The combination of these two technologies provides a platform to monitor equipment conditions in real-time and update maintenance decisions simultaneously with operational changes [5]. This integration provides a framework in which data flow, physical modeling, and predictive analysis are placed in a single cycle [6].

Parallel to these developments, enterprise resource planning (ERP) systems function as the digital backbone of modern industrial organizations [7]. ERP systems integrate operational, financial, and maintenance data, offering a centralized environment for managing resources and processes [8]. When connected with digital twins, ERP systems extend their capability beyond administrative control to real-time machine supervision, maintenance scheduling, and automated decision-making [9]. The synergy between ERP and DT systems, therefore, represents a new direction in intelligent maintenance, linking machine-level data to enterprise-level decision processes through standardized information flows and feedback mechanisms [10].

The integration of digital twins and ERP systems has the potential to create a comprehensive, multi-layered framework that connects physical machinery, virtual simulations, and predictive analytics. This integration not only enables real-time condition monitoring and predictive diagnostics but also supports strategic decision-making for resource allocation and maintenance optimization [11]. Such an approach directly contributes to the evolution of autonomous, self-maintaining industrial systems, a central goal in advanced machine design and automation research.

The composting industry represents a relevant testbed for developing and validating this integrated approach. Composting systems rely on mechanically intensive processes, such as aeration, mixing, and material transport, which are highly sensitive to operational disturbances and equipment reliability [12]. Failures in these machines can lead to substantial process inefficiencies, production losses, and increased operational costs. However, most composting plants still employ traditional maintenance practices and lack digital integration capabilities, highlighting a research gap in the deployment of predictive and intelligent maintenance solutions within this sector.

The industrial composting application scenario was chosen due to the characteristics that make this process one of the most challenging environments for predictive maintenance. The thermal and biological behavior of the compost pile is highly nonlinear and dependent on environmental conditions, and the slightest disturbance in temperature, humidity, or aeration can affect the process performance, the quality of the final product, and the safety of the equipment. On the other hand, the equipment associated with aeration, rotation, and monitoring of the compost pile is usually under unstable loads, harsh environmental conditions, and continuous process changes, and its failure causes significant costs. This operational sensitivity, combined with the constant access to real-time data and the possibility of building a dynamic digital twin, has made this scenario an ideal example to investigate the interaction between the digital twin, the ERP system, and the machine learning algorithm in the context of predictive maintenance.

The proposed framework goes beyond a conventional three-layer structure and enables synchronization of physical behavior, deep learning predictions, and operational decision-making by creating a closed loop between real-time equipment data, the digital twin model, and the ERP system. The main distinction of this architecture lies in the fact that the digital twin is not merely a passive simulator but actively intervenes in ERP processes, updates maintenance events, and optimizes decision-making policies based on MTTF, MTTR, and failure profiles, a feature that, to the best of the authors' knowledge, has been rarely reported in the existing literature.

The main contributions of the research are as follows:

- Providing an integrated DT-ERP-PdM architecture that enables two-way interaction between the physical model, deep learning, and operational decision-making.
- Developing a real-time mechanism for updating maintenance policies in ERP based on digital twin outputs, which, to the best of the authors' knowledge, has not been systematically addressed in previous studies.
- Designing a hybrid DT-LSTM model that manages and stabilizes noisy and discontinuous data in the industrial environment.
- Evaluating the proposed framework in a real industrial composting pilot and measuring its economic and operational impacts, including MTTF, MTTR, and ROI.

Therefore, the present study aims to develop and evaluate an integrated framework that combines digital twin technology, ERP systems, and machine learning algorithms for predictive maintenance in composting plants. The proposed framework seeks to improve failure prediction accuracy, reduce maintenance costs, and enhance equipment reliability through real-time data integration and intelligent analytics. Beyond its application in composting, the findings of this research contribute to broader advancements in machine design, automation, and industrial robotics by demonstrating how data-driven, digitally integrated systems can transform traditional maintenance into a proactive, intelligent process.

2. Literature Review

Recent advances in data-driven modeling, industrial automation, and enterprise information systems have significantly influenced maintenance and operations management across manufacturing sectors. The growing application of smart sensors, cyber-physical systems, and machine learning algorithms has redefined maintenance strategies from reactive or preventive paradigms toward predictive and autonomous maintenance models [13]. Predictive maintenance (PdM) aims to anticipate machine failures based on the analysis of operational data, thereby minimizing downtime and optimizing resource use.

Within this context, condition monitoring systems (CMSs) play a critical role by enabling the continuous acquisition and interpretation of equipment data such as temperature, vibration, and energy consumption. CMS platforms serve as the foundation for predictive analytics by linking real-time machine condition data to computational models that forecast degradation trends [14]. This has been particularly effective in complex mechanical systems, where early fault detection can prevent cascading failures and improve system safety.

A major enabler of this transformation has been the integration of machine learning techniques with physical models. Data-driven algorithms such as regression models, decision trees, and neural networks have been employed to enhance the accuracy of failure prediction and anomaly detection in various machine types [15]. However, the reliability of these algorithms depends heavily on data quality, sensor resolution, and the model's adaptability to changing operating conditions [16]. Consequently, hybrid approaches that combine physics-based modeling and data-driven prediction have gained attention for their balance of interpretability and predictive precision.

Parallel to this development, the concept of the digital twin (DT) has emerged as a transformative approach in machine design and system optimization. A digital twin constitutes a virtual representation of a physical machine or process, synchronized in real-time through continuous data exchange [17]. By coupling simulation models with live sensor data, DTs provide a means to simulate, predict, and optimize mechanical behavior under various operational scenarios [18]. Their ability to represent thermal, mechanical, and dynamic responses makes DTs particularly valuable for predictive maintenance and control design in automated systems [19].

In modern manufacturing, integrating DTs with enterprise resource planning (ERP) systems has opened new pathways for aligning machine-level intelligence with enterprise-level decision-making. ERP systems provide a centralized infrastructure for managing maintenance records, spare parts, human resources, and financial planning [20]. When linked to DT environments, ERP systems gain access to real-time machine performance indicators, enabling automated maintenance scheduling and cost optimization [21]. Such integration transforms isolated digital models into closed-loop intelligent maintenance frameworks, in which analytical results directly trigger operational or managerial actions [22].

Despite growing attention to DT-based predictive maintenance, most research and applications have concentrated on mechanically dominant industries such as automotive, aerospace, and energy [23]. In contrast, biological and environmental process industries—including composting, food processing, and bioreactors—remain underrepresented, even though their machinery (mixers, reactors, aeration systems) demands high reliability and precision control [24]. The complexity of bio-reactive systems, coupled with fluctuating operating conditions, presents challenges for developing stable digital twins and predictive models.

Several implementation barriers are also evident in the literature. These include insufficient data availability, difficulties in integrating heterogeneous information systems, organizational resistance to technological change, and the need for personnel training in advanced digital tools [6,25]. Furthermore, ensuring interoperability between industrial communication protocols (e.g., OPC-UA, MQTT) and enterprise systems remains a technical challenge for scalable deployment.

In response to these limitations, researchers have increasingly emphasized the need for integrated frameworks that unify the strengths of DTs, ERP systems, and artificial intelligence to support predictive and autonomous maintenance in complex industrial environments [26]. Such systems would enable continuous learning from machine behavior, improve decision-making accuracy, and promote self-optimizing and sustainable machine design.

The present research builds upon this body of knowledge by proposing a comprehensive DT–ERP integration framework tailored for predictive maintenance in composting machinery. Beyond its domain-specific application, the framework demonstrates how the convergence of cyber–physical modeling, predictive analytics, and enterprise integration can be applied to a wide range of industrial machinery to improve reliability, reduce maintenance costs, and contribute to the broader objectives of intelligent machine design and automation.

In summary, it should be stated that the literature review shows that although the combination of digital twin, ERP systems, and predictive maintenance approaches has been considered in recent years, most studies have used these three layers independently or only at the data exchange level, and an integrated model that can combine the physical behavior of the equipment, the real-time flow of operational data, and deep learning-based predictive analysis in a closed-loop and iterative cycle has rarely been introduced in the literature. The main distinction of this study is in providing an integrated framework in which the digital twin not only simulates the behavior of the equipment, but also, through direct connection to ERP, is able to simultaneously update the cycle of maintenance decision-making, resource allocation, and event recording. The integration of the physical model, the LSTM-based deep learning model, and the ERP operational structure in a closed architecture, along with its application in a real industrial composting environment, fills the gap in previous research and allows the evaluation of the economic and operational impact of this integration. These features constitute a set of theoretical and practical innovations that have been less addressed in previous studies.

3. Conceptual Framework and Integrated System Architecture

The conceptual framework proposed in this study integrates digital twin (DT) technology with enterprise resource planning (ERP) systems to enable predictive maintenance for complex mechanical systems in composting plants. The architecture is designed to form a cyber–physical environment that connects physical machinery, virtual modeling, and intelligent analytics within a unified data and control structure. This integration establishes a closed-loop system capable of real-time monitoring, condition prediction, and automated maintenance decision-making, all of which are essential for modern machine design and industrial automation. The framework is structured in three interdependent layers (Figure 1), each corresponding to a distinct level of functionality within the machine–data–management hierarchy.




Figure 1. Conceptual architecture of the integrated ERP–digital twin predictive maintenance system.

Although many DT–ERP–PdM studies use similar three-layer architectures, the main difference of the present framework is in the integration mechanism and the nature of the information flow. In the existing literature, the communication between DT and ERP is often limited to data exchange or simple alerts; whereas, in this study, a two-way mechanism is designed in which the digital twin model can influence ERP performance, modify maintenance policies, and guide the decision-making loop based on real-time predictions. This approach overcomes challenges such as a lack of coordination between physical modeling, limitations in handling noisy data, the inability to track the economic effects of decisions, and the weakness of real-time analysis.

3.1. ERP Layer: Information Management and Maintenance Control

At the base of the architecture lies the ERP system, functioning as the central information management hub. In this layer, operational and maintenance data of critical equipment—such as aeration systems, reactors, and mixers—are systematically recorded and managed. The ERP module stores both static information (equipment identifiers, maintenance histories, spare parts inventories) and dynamic data (performance metrics, failure reports, cost records).

By providing a standardized database for asset management, the ERP system ensures traceability and data integrity across the organization. Moreover, it supports maintenance scheduling, resource allocation, and cost tracking through integrated dashboards. This layer acts as the bridge between enterprise-level planning and machine-level operations, ensuring that the insights produced by the predictive analytics layer can be seamlessly translated into actionable work orders and maintenance activities.

3.2. Digital Twin Layer: Machine-Level Modeling and Real-Time Simulation

The second layer constitutes the digital twin environment, representing the virtual counterpart of the physical machinery within the composting system. This layer captures the mechanical behavior, thermal dynamics, and energy exchange of key components—including mixers, aeration units, and compost reactors—using dynamic simulation models developed in MATLAB/Simulink R2021b and AnyLogic Professional 8.7.10 platforms.

Each digital twin receives real-time data streams from the plant's sensor network (temperature, vibration, pressure, motor current) and continuously updates its internal parameters to reflect the actual operational state. The twin thereby performs online monitoring, behavior prediction, and scenario-based simulations.

From a machine design perspective, this layer enables the evaluation of component stress, load fluctuations, and degradation patterns under different operating conditions. It can thus serve not only as a maintenance tool but also as a design feedback mechanism, providing insights into mechanical performance that can inform future equipment optimization. The two-way synchronization between the physical asset and its digital model ensures that the DT functions as both a diagnostic and a prognostic system.

3.3. Predictive Analytics Layer: Intelligent Diagnostics and Decision Support

The third layer integrates data from the ERP and DT environments through a suite of machine learning-based predictive algorithms. This analytical layer is responsible for fault detection, remaining useful life (RUL) estimation, and maintenance decision optimization.

A hybrid modeling approach is adopted:

- Physics-based models capture the fundamental mechanical and thermodynamic behavior of the equipment.
- Data-driven models, such as long short-term memory (LSTM) networks, analyze time series data to predict anomalies and failure probabilities.

By combining these two approaches, the system achieves both interpretability and high predictive accuracy, aligning with the design principles of smart, self-adaptive machinery. The results of the predictive models—failure probabilities, performance degradation trends, and maintenance recommendations—are transmitted back to the ERP system. These outputs trigger automated maintenance actions, update maintenance schedules, and inform operators through ERP dashboards.

3.4. Data Flow and System Integration

The information exchange between layers is structured according to industrial communication standards to ensure interoperability and scalability. Real-time equipment data are transmitted from field sensors and PLCs to the DT layer using the OPC-UA protocol, which guarantees reliable, standardized communication. The processed analytical results are transferred to the ERP system through a RESTful API, enabling bidirectional communication between operational and enterprise levels.

The complete information flow begins with data acquisition at the machine level, followed by virtual modeling and analysis in the DT environment, and concludes with automated decision execution in the ERP maintenance module. This closed-loop process transforms traditional maintenance workflows into a data-driven, autonomous decision cycle, consistent with the principles of Industry 4.0 and intelligent machine design.

3.5. System Overview

Figure 1 schematically illustrates the architecture of the integrated ERP–DT predictive maintenance system. It highlights the interaction between physical assets, virtual models, and analytical components, showing how data traverse from machinery sensors to decision layers

and back to operational control. The result is a modular and extensible system architecture that can be adapted to different industrial contexts where equipment reliability and automation are critical. In Figure 1, the integrated architecture of this system is shown schematically, in which the role of each layer and the information flow path are clearly defined.

As shown in Figure 1, the interaction between the ERP, digital twin, and predictive analytics layers provides an integrated framework for condition monitoring, failure prediction, and effective maintenance management in composting plants.

4. Research Methodology

The methodology adopted in this study was designed to develop an integrated predictive maintenance framework for composting plants based on the combined use of digital twin technology, machine learning algorithms, and ERP-based data management. The research followed a cyber–physical integration approach, in which real-time data from machinery are continuously processed by virtual models and analytical algorithms to support predictive maintenance and autonomous decision-making. The process was implemented in five sequential stages, each addressing a distinct aspect of system development—from physical modeling to predictive analysis and system validation.

4.1. Stage 1—Construction of the Digital Twin Architecture

In the first stage, the modeling of key composting equipment was carried out to construct a digital twin capable of representing both physical and operational characteristics of the plant’s machinery. The focus was on reactors, mixers, and aeration systems, which play critical roles in maintaining the thermal and biological stability of the composting process.

The physical behavior of these components was modeled using MATLAB/Simulink and implemented in the AnyLogic simulation environment. The digital twin used equations governing the energy balance, heat transfer, and biological reaction dynamics, continuously updated by real-time sensor data to ensure accuracy and synchronization with the actual process.

The thermal behavior of the compost pile is expressed by the following energy balance equation:

$$\frac{dT}{dt} = \frac{Q_{\text{reaction}} - Q_{\text{loss}}}{mc_p} \quad (1)$$

where T is the temperature of the compost pile, Q_{reaction} is the energy produced by biological reactions, Q_{loss} is the rate of heat loss to the environment, m is the mass of the compost pile, and c_p is the specific heat capacity of the compost pile. The parameters of this equation were updated in real-time using virtual sensor inputs, ensuring that the digital twin accurately replicated variations in the compost pile’s physical state and environmental conditions.

The digital twin model developed in this study consists of several complementary modules that represent thermal, mechanical, airflow dynamics, and biological response patterns in an integrated manner. These modules are designed to reflect the current state of the equipment and simulate future trends while receiving real-world data and connecting to the predictive model. Combining these modules into a single structure allows the model to both reproduce the physical behavior of the equipment and play an active role in data-driven predictions.

To increase the accuracy of the digital twin model, its development process consisted of three basic steps: model creation, calibration, and numerical validation. In the first step, the physical model of the equipment was created using heat transfer equations, actuator dynamic equations, and the bioenergy model of the composting process in the MATLAB/Simulink environment. Then, key parameters, including thermal conductivity coefficients, biological reaction rates, heat dissipation coefficients, and motor-shaft friction

coefficients, were calibrated by matching the model output with real plant data. This calibration was performed by minimizing the RMSE error between the real data and the DT output. In the validation step, the model was tested with several independent datasets, and the MAE and R^2 indices were calculated to ensure the model's ability to reproduce the thermal, vibration, and electrical behavior of the equipment. The results showed that the model is able to accurately represent the real dynamics of the system and maintain its numerical reliability under different operating conditions.

A combination of denoising and normalisation methods was used to prepare the data to make the signals suitable for temporal analysis and predictive modeling. Momentary noise was smoothed with a moving average filter, and the data were transformed to a standard interval using the min–max method to avoid the effect of scale differences. Preprocessing steps were performed only once, and redundant descriptions were removed to provide a clear and non-redundant data analysis flow.

The validation of the digital twin model was performed using historical data from several operating periods and under different loading conditions. The model output for temperature, pressure, and aeration rate was compared with the measured values, and the MAE and RMSE error indices were used as accuracy metrics. The results showed that the difference between the model and reality was within an acceptable range for industrial applications, and the model was able to reproduce the trends and responses of the system with reasonable accuracy. This validation provided a basis for the robustness of the model and allowed the system to effectively use the predicted outputs in the operational decision-making cycle.

4.2. Stage 2—Development of Hybrid Predictive Algorithms

The second stage focused on developing hybrid predictive algorithms that combined physical modeling with machine learning techniques for equipment condition monitoring and failure forecasting. Three groups of models were developed:

Physics-based models simulating mechanical behavior under nominal operating conditions.

Machine learning models, particularly long short-term memory (LSTM) networks, were trained on temporal sensor data to capture nonlinear dependencies in machine performance.

Hybrid models that combine the outputs of these two approaches significantly improve prediction accuracy and operational robustness.

The LSTM model was formulated as follows:

$$\begin{aligned}
 f_t &= \sigma(W_f x_t + U_f h_{t-1} + b_f) \\
 i_t &= \sigma(W_i x_t + U_i h_{t-1} + b_i) \\
 \tilde{c}_t &= \tanh(W_c x_t + U_c h_{t-1} + b_c) \\
 c_t &= f_t \odot c_{t-1} + i_t \odot \tilde{c}_t \\
 o_t &= \sigma(W_o x_t + U_o h_{t-1} + b_o) \\
 h_t &= o_t \odot \tanh(c_t)
 \end{aligned} \tag{2}$$

In these equations, x_t represents the input vector at time step t , while h_{t-1} denotes the hidden state propagated from the previous step. The terms f_t , i_t , and o_t correspond to the forget, input, and output gates that regulate the flow of information within the memory cell. The variable $\tilde{c}_t$ denotes the candidate cell state generated at the current step, and c_t represents the updated cell state after combining the effects of the gates. The matrices W^* and U^* refer to the input and recurrent weight parameters associated with each gate, and b^* captures the corresponding bias terms. The function $\sigma(\cdot)$ denotes the logistic sigmoid activation, $\tanh(\cdot)$ is the hyperbolic tangent function used for state modulation, and the operator $\odot$ indicates element-wise multiplication.

In this study, the complete LSTM architecture, including input, output, and forgetting gates, was used, and the data processing process was designed according to the standard structure of these networks.

Before entering the data into the network, preprocessing operations, including noise removal, outlier removal using the IQR index, data alignment with the same sampling rate, and min–max normalization, were performed. To convert continuous data into processable input for LSTM, sliding windows with a length of 30 time steps were used, and each window was considered as an independent training sample. The data were also split temporally without mixing consecutive periods to preserve the real-time series patterns. Finally, the LSTM output was directly fed into the predictive analysis module, playing a key role in the closed-loop decision-making between the digital twin and the ERP.

Finally, hybrid models were designed by combining the results obtained from the above two approaches to optimize the prediction accuracy in the operational environment. To develop these models, the Python 3.10 environment with TensorFlow and Keras libraries was used.

It is worth noting that Equation (2) is only a simplified representation of the model's recursive mechanism and is used to present the general logic of the iterative process. The actual implementation of the model in this research is based on the complete LSTM structure, which includes input, forget, and output gates, as well as the internal memory cell. Therefore, the final model uses a complete LSTM architecture, but to avoid the unnecessary complexity of the formulas, only a summarized representation is given in the text.

To train the LSTM model, the data were first transformed into 30-step time sequences, and a sliding-window strategy was used to generate input samples. The model architecture consists of two consecutive LSTM layers with 64 hidden units and a final fully connected layer. The activation function of the middle layers is tanh type, and a linear function is used for the output. The data were divided into three parts based on the time sequence—training, validation, and testing—to prevent information leakage. The learning process was performed using the Adam optimizer, a learning rate of 0.001, a batch size of 32, and a maximum of 100 training cycles. An early stopping mechanism was also used to prevent overfitting. These settings gave the model a good ability to learn temporal patterns and the nonlinear behavior of the system.

4.3. Stage 3—Implementation and Validation of the Predictive Maintenance Workflow

In the third stage, an ERP-based management layer was designed and implemented using the Odoo ERP open source platform. The ERP system was configured to manage maintenance scheduling, asset data, and operational performance indicators, while serving as the interface between predictive analytics and maintenance operations.

In this study, the process of integrating the digital twin with the ERP system was designed as a linear flow to make the interaction between the physical model and the decision-making layer more transparent. After entering the DT layer, the sensor data were analyzed in the form of a hybrid model, including the physical module and the predictive model, and the current state of the equipment and future trends were calculated simultaneously. The output of this analysis was directly input to the ERP module to perform maintenance, resource allocation, and operational planning based on the actual and predicted state. The main difference between this framework and the existing architectures in the literature is that the data flow does not operate as a one-way path, but as a closed feedback loop, and based on the prediction results, the physical parameters are also modified in the DT layer. This approach allows the system to both increase the modeling accuracy in each cycle and rely on more reliable information for decision-making in the ERP layer.

The technical logic of DT and ERP integration goes beyond simple data exchange and involves a closed decision-making cycle. In this architecture, DT, after receiving real-time

data from the PLC, transfers analytical outputs, including failure probability, RUL, and performance deviation patterns, to the ERP via a REST-based API. ERP not only stores these data but also executes predictive maintenance ruleset through its internal decision-making engine and issues executive commands such as creating Work Orders, adjusting maintenance schedules, or activating management-level alerts based on defined thresholds. These decisions are fed back to DT to simulate new operational scenarios and predict their possible impacts. This round-trip cycle shows that DT–ERP integration is dynamic, bidirectional, and based on control logic, and is not limited to data transfer.

4.4. Stage 4—Deployment and Integration with the ERP Environment

The fourth stage addressed the industrial communication layer, enabling synchronization between field devices, the digital twin, and the ERP system. Communication among these components was based on the OPC-UA protocol, which ensured standardized, reliable, and secure data transmission between industrial sensors, PLCs, and virtual models.

The data exchange cycle included the following:

- Transmission of sensor data (temperature, vibration, motor current, aeration pressure) to the digital twin.
- Analysis of the data in the modeling environment and prediction of anomalies using the LSTM-based model.
- Feedback of predictive results to the ERP system through REST API.
- Recording and execution of maintenance actions in the ERP module.

This service-oriented architecture allowed for real-time interaction between all system layers and ensured scalability for integration with additional industrial assets. Figure 2 illustrates this communication framework.




Figure 2. Integrated architecture for data exchange between ERP, digital twin, and industrial sensors in a predictive maintenance system.

4.5. Stage 5—Performance Evaluation and Key Indicators

The fifth stage evaluated the overall performance, accuracy, and economic effectiveness of the predictive maintenance framework using quantitative indicators.

The mean time to failure (MTTF) was used as the main reliability metric:

$$MTTF = \frac{\sum_{i=1}^n T_i}{n} \quad (3)$$

where T_i represents the failure-free operation time for each observation, and n is the number of samples.

The mean time to repair (MTTR) was calculated to assess maintenance efficiency, while the mean absolute error (MAE) measured the prediction accuracy of the models:

$$MAE = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| \quad (4)$$

where y_i represents the actual values and $\hat{y}_i$ represents the predicted values for sample i . Finally, the return on investment (ROI) was evaluated based on the reduction in direct maintenance costs and costs due to production line downtime as an indicator for the economic analysis of the system.

The data used in this study were obtained anonymously and for research purposes from a reputable Australian company active in the field of industrial compost production. The dataset used in this study consists of real-time sensor measurements collected from the compost pile, together with industrial-scale operational records obtained from the plant's monitoring system. All data sources were cleaned, synchronized, and preprocessed to ensure consistency and suitability for model development.

This methodological structure, relying on the integration of physical modeling, predictive algorithms, the ERP system, and the digital twin, provides a comprehensive framework for improving maintenance performance in compost plants and provides a scientific basis for implementation in real industrial conditions.

5. Case Study and Experimental Setup

To validate the performance of the proposed predictive maintenance framework, a case study was conducted in collaboration with a reputable Australian company specializing in industrial compost production. The selected facility operates a continuous composting line equipped with automated aeration, mixing, and temperature-control systems, providing a suitable testbed for evaluating the integration of digital twin (DT), machine learning algorithms, and ERP-based maintenance management.

The case study aimed to assess the framework's ability to improve equipment reliability, predict potential failures, and reduce maintenance costs under real industrial operating conditions.

5.1. Plant Description and Equipment Mapping

The composting plant used in this study consists of several interconnected mechanical subsystems—bioreactors, rotary mixers, and aeration units—all of which are essential for maintaining optimal process conditions. Each subsystem contains multiple electromechanical components, including motors, shafts, bearings, and power transmission units, continuously exposed to variable thermal and mechanical loads.

Prior to implementation, a detailed process mapping was conducted to identify critical assets, instrumentation layout, and data communication routes. Parameters such as reactor temperature, aerator vibration amplitude, motor current, and aeration system pressure were selected as key monitoring variables for predictive analysis. The plant's sensor

network and PLC infrastructure were fully documented to ensure proper integration with the digital twin and ERP layers.

5.2. Data Collection and Integration

The experimental dataset consisted of two categories:

- Historical data, including three years of logged maintenance events and operational records;
- Real-time sensor data, collected continuously during the experimental period through IoT-enabled industrial sensors.

Each signal was transmitted to the digital twin via the OPC-UA protocol, processed for anomaly detection, and analyzed through the hybrid LSTM-physics models developed in earlier stages. The results were communicated to the ERP maintenance module using RESTful APIs, ensuring real-time synchronization between physical processes, digital models, and enterprise-level decisions.

Table 1 summarizes the main characteristics of the data collected from the composting systems.

Table 1. Characteristics of data collected from industrial composting systems.

Data Type	Data Source	Sampling Interval	Data Range	Number of Samples
Reactor Temperature	Industrial Sensors	Every 5 min	35–80 °C	10,800
Aerator Vibration	Vibration Sensors (IoT)	Every 1 min	0.1–2.5 mm/s	54,000
Mixer Motor Current	PLC Monitoring System	Every 10 min	10–45 A	5400
Aeration System Pressure	Pressure Transmitters	Every 5 min	0.5–2.0 bar	10,800

To ensure the quality of the input signals to the model, several data preprocessing and cleaning methods were performed. Temperature and vibration data, which were most sensitive to environmental noise, were corrected using a moving average filter (SMA) and a median filter to remove the effect of transient noise caused by mechanical shocks and sudden fluctuations. For pressure and electrical current data, which were likely to be disconnected from the PLC, the missing points were first reconstructed using the linear interpolation method, and points with unrealistic jumps were removed or corrected using the IQR method. In the next step, all data were normalized to place different measurement ranges on a common scale. A noise management protocol was also developed that included removing high-frequency fluctuations, detecting sudden discontinuities, and checking the stability of the sensors over specific time intervals. This process ensured that the final dataset was reliable in terms of stability, consistency, and the absence of signal loss for input into predictive models.

Although the MTTF, MTTR, and ROI indices were calculated based on real data and with full transparency, the results depend on the operating conditions of the same industrial site and may be affected by external factors in other environments. Factors such as operating load fluctuations, ambient temperature changes, heterogeneous equipment wear, sensor accuracy, and differences in failure patterns over time can affect the values of these indices. In this study, due to the testing being conducted at a single location and the limited number of scenarios, it was not possible to fully separate these factors. However, the results show that the improvement pattern is relatively stable, and in future studies, by expanding the scenarios and collecting data from multiple industrial sites, a more comprehensive analysis of these effects will be possible.

5.3. Implementation of the Integrated System

The digital twin layer was deployed on the plant's IT infrastructure and configured to operate in continuous communication with the field-level control system. Virtual models of reactors, mixers, and aeration units were implemented in AnyLogic, using dynamic parameters derived from the physical models developed in MATLAB/Simulink. The OPC-UA interface ensured that the DT continuously received and updated process variables in real-time, while simulation outputs were logged and fed into the predictive algorithms.

The ERP system, built on the Odoo platform, served as the supervisory and decision-making environment. Through its dedicated maintenance management module, the ERP received alerts and recommendations from the predictive layer and automatically generated maintenance orders or inspection schedules based on the system's forecasts.

The communication workflow followed a closed-loop feedback structure, ensuring the following:

- The raw sensor data from the machinery were streamed to the digital twin.
- Analytical models predicted the equipment's condition and identified abnormal behaviors.
- Predictive results were sent to the ERP dashboard, updating maintenance KPIs.
- Maintenance activities executed in the field were recorded back into the ERP, closing the feedback loop.

This design established a fully autonomous monitoring and decision-support cycle, consistent with Industry 4.0 principles of self-diagnosing and self-maintaining machinery.

5.4. Test Scenarios and Experimental Conditions

To comprehensively evaluate system performance, four operating scenarios were defined based on historical failure modes and maintenance records:

- Mechanical failure in aeration motors;
- Reduced transmission efficiency in mixer shafts;
- Abnormal temperature peaks in compost reactors;
- Detection of atypical vibration patterns in power transmission systems.

Each scenario was tested under normal, alarm, and critical conditions to assess the responsiveness and accuracy of the predictive framework. The test cycles included both simulated disturbances (using DT-generated inputs) and real sensor fluctuations collected during normal operations.

The validation process evaluated the system's performance in terms of the following:

- Prediction accuracy and time-to-detection;
- Sensitivity to different operating conditions;
- Effectiveness of ERP-based maintenance scheduling;
- Robustness of communication and data flow under varying load conditions.

5.5. Data Reliability and Calibration

All sensors underwent calibration using traceable reference instruments prior to deployment to ensure the accuracy of temperature, vibration, and pressure measurements. Data integrity was verified using automated routines for noise filtering, missing-value detection, and outlier correction.

The real-time and historical datasets were stored in structured databases within the ERP, ensuring traceability and version control for analytical and maintenance records. This configuration enabled continuous improvement of predictive models through retraining with updated plant data, thereby increasing model reliability over time.

The implementation of this case study provided a realistic validation environment, capturing both the complexity of mechanical interactions in composting systems and the stochastic

variability of industrial processes. The successful operation of the integrated DT-ERP system demonstrated the feasibility of deploying predictive maintenance and automation architectures in process industries where equipment reliability is crucial for sustained performance.

To create a coherent understanding of how the model is developed, data processed, and performance evaluated, it is necessary to consider the methodological framework of this research as an integrated chain. In this framework, the physical modeling layers, deep learning components, and the hybrid structure of the digital twin and the data-driven model are designed to combine the actual behavior of the equipment with data-driven predictions. The data architecture includes the steps of preprocessing, noise removal, normalization, and implementation of the communication flow between the operational system, digital twin, and the intelligent analysis layer. The validation process is also based on a set of error indicators and performance criteria, as well as testing in several operational scenarios. Economic and operational analyses, including changes in MTTF, MTTR, and calculation of return on investment, are also performed assuming stability of the actual plant conditions and no changes in maintenance procedures. This integration indicates that the research results are based on a precise link between modeling, data management, and performance evaluation, providing the necessary strength for scientific interpretation of the findings.

6. Analysis of Results

The analysis of results was performed to quantitatively assess the effectiveness of the proposed predictive maintenance framework in terms of prediction accuracy, system stability, and operational performance improvements. The evaluation combined model-level validation (predictive accuracy and robustness) with system-level performance indicators (reliability, downtime, and cost reduction).

Model Performance Evaluation

The first stage of the analysis focused on evaluating the performance of the failure prediction and detection models. Both continuous process variables (reactor temperature, aeration pressure, mixer motor current, and vibration levels) and categorical failure indicators (normal, alarm, failure) were assessed.

The performance metrics used included mean absolute error (MAE) and root mean square error (RMSE) for regression-based variables, and accuracy and F1-score for classification models.

To clarify the model evaluation and transform the qualitative report into a reliable quantitative analysis, the actual values of the performance indicators were also calculated and reported. In the test dataset, the MAE value was 0.18, and the RMSE value was 0.27, indicating that the model is able to reconstruct the system change process with acceptable accuracy. In addition, the accuracy of the model in predicting operational states was evaluated to be 94%, and the AUC value was 0.92. These results indicate that the model not only performs well in predicting continuous values but also provides reliable performance in diagnosing system states.

The results are summarized in Table 2.

Table 2. Summary of performance indicators of failure prediction and detection models.

Metric	Predictive Model (Regression)	Detection Model (Classification)
MAE	0.18	-
RMSE	0.27	-
Accuracy	-	0.94
AUC	-	0.92
Precision	-	0.92
Recall	-	0.89

As shown in Table 2, the proposed hybrid modeling approach achieved high accuracy across all monitored parameters, with prediction errors (MAE) well below operational thresholds. The LSTM-based models successfully captured temporal dependencies, while the physics-based elements ensured interpretability and numerical stability.

A predicted vs. actual plot for reactor temperature (Figure 3) revealed that the data points were tightly clustered around the 45° reference line, confirming the strong correlation between measured and predicted values and the reliability of the model in replicating the process dynamics.




Figure 3. Plot of actual values versus predicted values for reactor temperature.

Next, to analyze the performance of the classification models in predicting failures, a ROC (receiver operating characteristic) chart was drawn. Figure 4 shows the ROC curve of the failure detection model. The area under the curve (AUC) in this model was found to be 0.92, which confirms the high accuracy of the system in correctly distinguishing normal conditions from failures.




Figure 4. ROC chart of the performance of the classification model in detecting failures.

The results of this analysis showed that the proposed framework provided a combination of accuracy in predicting continuous variables and efficiency in detecting abnormal conditions.

The high accuracy of the models, both in predicting process values and in identifying critical conditions, confirms the effectiveness of the system in a real industrial environment.

A comparative analysis was conducted to evaluate the system's impact on maintenance performance indicators before and after implementation of the predictive maintenance framework. The mean time to failure (MTTF) increased from 325 h (before implementation) to 460 h (after deployment), corresponding to an improvement of approximately 40%. This indicates that the system successfully prevented unplanned downtime through early fault detection and data-driven intervention scheduling. Simultaneously, the mean time to repair (MTTR) decreased from 9.5 h to 6.0 h, resulting in a 35% reduction in downtime per incident. These improvements reflect the operational benefits of integrating real-time diagnostics with automated maintenance scheduling in the ERP environment. From an economic perspective, the total annual maintenance cost decreased from AUD185,000 to AUD127,000, representing a 30% cost reduction. The resulting return on investment (ROI) was calculated as 42.5% within the first year, confirming that the proposed system is economically viable for large-scale industrial deployment.

A comparison of maintenance-related costs before and after implementing the proposed framework is shown in Figure 5, illustrating its tangible effect on reducing the annual costs of the compost plant.




Figure 5. Share of maintenance costs before and after implementation of the proposed framework in the annual costs of the compost plant.

The results of this section clearly show that the implementation of the proposed framework, in addition to improving technical indicators such as increasing failure-free operation periods and reducing repair time, has also resulted in significant savings in operating costs from an economic perspective and has brought positive economic returns to the organization.

One of the essential features of the success of predictive systems in industrial environments is their ability to maintain stable performance in the face of changing conditions and a variety of operational scenarios. In this regard, the performance stability assessment of the developed model was carried out, focusing on how it responds to different scenarios defined in the composting process.

For this analysis, a set of test scenarios was designed based on industrial experiences and maintenance records in the studied plant. These scenarios included engine failure in the aeration system, reduced efficiency of power transmission shafts in mixers, abnormal

temperature increase in compost reactors, and recording of unusual vibration patterns in power transmission equipment. Each of these scenarios was defined with the aim of reflecting real conditions and applying the model in operational environments with different challenges to examine the degree of adaptability and stability of the model's performance in various conditions.

The results obtained from running the model in these scenarios showed that the predictive performance of the model was of appropriate accuracy in different conditions. In the reactor temperature scenarios, the model was able to correctly identify fluctuations caused by biological reactions and environmental changes, while maintaining prediction accuracy and an average error of less than 2.5 °C. In the case of vibration and mechanical performance scenarios, the accuracy of detecting abnormal states with an F1-Score index of more than 89% was maintained in all scenarios, indicating the stability of the algorithm in the face of input data changes.

In addition, in the study of critical scenarios, such as sudden failure of aeration motors or mixer malfunctions, the model was able to detect the changes without a noticeable increase in the error rate and provided acceptable performance with the lowest false positive rate. These results indicate that the predictive system has been able to maintain its accuracy and stability of performance not only in normal conditions but also in the face of unexpected data or acute operating conditions.

To visually demonstrate the behavior of the model in different conditions, Figure 6 depicts the changes in the accuracy of the model's performance in the four main defined scenarios. This chart displays the comparison of the model's prediction accuracy percentage in standard conditions and each of the defined scenarios in a column format, clearly showing that the model's performance fluctuations in all scenarios remain within acceptable limits.




Figure 6. Comparison of model performance accuracy under different conditions of test scenarios.

The results of this analysis clearly demonstrated that the developed framework has acceptable performance stability under various operating conditions and has been able to maintain its desired efficiency in all defined scenarios. This capability has not only influenced the system's performance in real industrial environments but also strengthened its operational validity as a reliable tool for decision-making in the predictive maintenance process.

To assess the robustness and flexibility of the proposed framework against changes in the input data volume and the difference between real and simulated data, a comprehensive sensitivity analysis was conducted. This analysis was designed to answer the question of

whether the developed model is able to continue its stable performance under different data conditions without a noticeable decrease in accuracy and efficiency.

In the first stage, the effect of the number of training samples on the accuracy of the model was investigated. For this purpose, the model was trained on datasets with different volumes, including 25%, 50%, 75%, and 100% of the total available data. At each stage, the model's accuracy in predicting key variables and the mean absolute error (MAE) were calculated. The results showed that with increasing the training data volume, the model's accuracy improved continuously, and the prediction error rate decreased. In particular, when the number of training samples reached more than 75% of the total data, the changes in accuracy approached a saturation point, which indicates the stability of the model in the face of larger data. The trend of these changes is shown in Figure 7. In this graph, the decreasing trend of the MAE rate with increasing number of samples is seen continuously and without sudden jumps, which confirms the stability of the model performance at different data scales.




Figure 7. The trend of change in the mean absolute error (MAE) in the face of changing the number of training samples.

In the second step, a sensitivity analysis was performed with respect to the type of input data. In this section, the model performance was compared in two cases: a case where the training data consisted solely of real data collected from the factory, and a case where a combination of real data and simulated data was used. The accuracy of the model in these two cases was evaluated by considering a fixed test set. The results of this comparison showed that the use of combined data improved the model's accuracy in identifying rare patterns or less frequently observed conditions, without increasing the overall error of the model in real data. This finding indicated that integrating simulated data with real data could be an effective strategy to improve the model's ability to learn complex patterns.

To visually represent this comparison, Figure 8 was generated, which shows the changes in model accuracy under different input data types as a surface graph. This graph shows that the highest level of accuracy was achieved when the training data consisted of a combination of real and simulated data.




Figure 8. The effect of training data type on the accuracy of the prediction model.

The analytical results demonstrate that the proposed ERP–digital twin–ML integration framework is both technically effective and operationally stable. The model consistently achieved:

- Low prediction error ($MAE < 2\%$);
- High classification accuracy ($AUC = 0.92$);
- Substantial reliability gains ($MTTF \uparrow 40\%$);
- Significant cost savings ($ROI = 42.5\%$).

From an engineering perspective, the system achieved the autonomous characteristics expected of modern machine design:

- It continuously monitors the condition of critical assets.
- It learns from historical and real-time data.
- It adapts maintenance actions based on evolving operational states.

These features illustrate the framework’s potential as a scalable solution for self-diagnosing and self-maintaining industrial machinery, a key milestone toward Industry 4.0 and 5.0 machine autonomy.

In order to ensure that the observed improvements in MTTF, MTTR, and ROI indicators were due to the implementation of the forecasting system and not to lateral changes, a three-month baseline period was defined before the implementation, and all operating conditions, maintenance schedules, number of shifts, and repair policies were kept unchanged. The data before and after the implementation were compared under controlled conditions and the only added variable was the DT–ERP–ML module. The ROI calculation was based on the reduction in spare parts costs, reduction in line downtime, and reduction in manpower hours, and no other operational costs or changes were included in the model. As a result, the contribution of the forecasting system to the performance improvement could be separated and evaluated with high accuracy.

Despite the favorable performance of the model, the data used had limitations that should be considered in interpreting the results. First, all data were collected from a single plant and a specific process cycle, and although the model was tested in four operational scenarios, full generalizability of the results to all composting units or similar industries requires more diverse and multi-site data. Second, although extensive procedures were performed to manage noise, remove outliers, and control data discontinuities, some of the noise from unpredictable biological behaviors can still affect the accuracy of the model. Therefore, the results of this study should be read with these considerations in mind, and future research with a larger dataset and different operational environments will help to increase the generalizability of the findings.

To establish a basis for comparison and measure the performance improvement of the proposed model, a simple baseline model based on historical average value and a fixed threshold was also evaluated. In this model, the prediction process was performed without any learning and solely based on the average of previous values, which is similar to the early warning logic in the ERP system. The results showed that the baseline model had an MAE value of 0.39 and an RMSE value of 0.52 in the test data. Comparing these values with the performance of the proposed model, which provided an MAE of 0.18 and an RMSE of 0.27, shows that the proposed hybrid model is significantly better in terms of performance prediction accuracy. Also, in the situation recognition section, the baseline model only achieved an accuracy of 0.81 and an AUC value of 0.74, while the proposed model achieved an accuracy of 0.94 and an AUC of 0.91. This comparison shows that using the proposed hybrid model has led to a significant improvement in the ability to predict and detect errors early.

Given that the system developed in this study was implemented and evaluated for a specific device in a production line, its generalization to other equipment and different production lines can also be investigated. The digital twin structure and the LSTM model used are designed to be adaptable to minimal changes in physical parameters, operating ranges, and input characteristics. In cross-device applications, the model can be re-tuned with smaller datasets based on transfer learning to eliminate the need for complete retraining. In cross-production-line applications, the system can be deployed on a large scale by defining independent modules for each device and using a central synchronizing model.

Long-term stability of the model and a reduction in the effects of changing data distribution over time (data drift) are also ensured through continuous monitoring of error indicators, periodic updating of digital twin parameters, and re-tuning of the LSTM network at specific intervals. If operational patterns or environmental conditions change significantly, a periodic retraining with new data is performed, which prevents model performance degradation and maintains the accuracy of the warning system in the long term. These adaptation mechanisms make the proposed system usable and reliable not only for a specific case, but also for a wide range of equipment and operating environments.

7. Deployment Feasibility Analysis

Following the validation of technical performance and economic benefits, a feasibility study was conducted to evaluate the potential for large-scale industrial deployment of the proposed digital twin (DT)–ERP predictive maintenance framework. This assessment examined the technical requirements, operational barriers, and economic justifications associated with transitioning the system from a research prototype to a fully integrated industrial solution.

Successful deployment of the proposed system requires a robust information and communication infrastructure that supports continuous data exchange between the physical and virtual layers. The primary technical prerequisites identified include the following:

- Sensor and data acquisition infrastructure capable of real-time monitoring of key parameters such as vibration, temperature, and power consumption;
- Reliable connectivity with programmable logic controllers (PLCs) to ensure data synchronization between equipment and digital twin models;
- Stable local or cloud-based servers with sufficient computational capacity for real-time simulation and predictive model execution;
- ERP system compatibility with external analytical and simulation environments through RESTful API integration.

These elements collectively provide the backbone for the cyber–physical data loop, ensuring accurate synchronization between machine operations, virtual models, and enter-

prise decision systems. The presence of standardized communication protocols, particularly OPC-UA, ensures interoperability with existing industrial control systems and supports scalable integration across diverse machinery types.

The feasibility analysis also identified several operational challenges that could affect large-scale adoption. First, the predictive performance of the system depends heavily on the accuracy, completeness, and continuity of sensor data. Inconsistent sampling, missing values, or noise in the measurements can reduce the reliability of model predictions and lead to false maintenance alerts. Therefore, the implementation of comprehensive data quality management (DQM) procedures is essential to preserve the integrity and consistency of information exchanged between the physical and virtual layers.

Another critical aspect is human–system interaction. The adoption of advanced DT–ERP systems requires dedicated training and awareness programs to ensure that maintenance personnel and system operators can interpret predictive outputs, configure dashboards, and respond effectively to data-driven insights. Without proper training and user engagement, the full potential of the predictive maintenance system and its automation features may remain underexploited.

Integration complexity also represents an important barrier to deployment. The synchronization between the digital twin and ERP layers introduces technical challenges related to data structure compatibility, software interoperability, and real-time data transfer. Establishing a unified data model and standardized communication interfaces is vital to guarantee continuous and accurate information flow between analytical and management layers.

Finally, cybersecurity and data governance are fundamental for ensuring the safe operation of connected industrial systems. Continuous communication between machinery, digital platforms, and enterprise software increases exposure to potential cyber threats. To mitigate these risks, the deployment process must include robust encryption protocols, secure authentication mechanisms, and strict user access controls. Implementing these measures ensures the protection of sensitive industrial data and supports the long-term reliability of the integrated predictive maintenance environment.

The economic analysis revealed that the initial investment in digital twin development, hardware installation, and personnel training is offset by the operational savings achieved through predictive maintenance. The pilot implementation demonstrated a return on investment (ROI) of 42.5% within the first operational year, primarily due to reductions in equipment downtime, spare parts consumption, and labor hours.

For full-scale deployment, the main cost categories include the following:

- Procurement of additional industrial sensors and PLC interfaces;
- Expansion of server and network infrastructure;
- Customization of ERP modules and APIs;
- Staff training and change management programs.

The payback period for implementation was estimated at 12–15 months, depending on the plant's scale and the number of critical assets under monitoring. The economic findings indicate that the system is viable not only for large enterprises but also for medium-sized facilities seeking to improve maintenance efficiency and operational reliability.

One of the critical strengths of the proposed framework is its modular and scalable architecture. The system can be expanded horizontally by connecting additional assets through OPC-UA nodes, and vertically by integrating higher-level decision-support modules within the ERP. The modular design allows for the following features:

- Scalable deployment across multiple production lines or facilities;
- Custom adaptation to different machinery types or process industries;
- Progressive integration of advanced analytics, such as reinforcement learning or multi-agent optimization, for self-adaptive control;

- This scalability ensures that the system can evolve alongside industrial digitalization initiatives and accommodate future expansions without major structural modifications.

To visualize the interplay between the technical, operational, and economic factors, a conceptual map was developed (Figure 9). The diagram illustrates how infrastructure readiness, data management quality, and human resource capabilities jointly determine the success of implementation, while cost savings and ROI provide the necessary economic incentives for adoption.




Figure 9. Conceptual map of key factors for the feasibility of implementing a predictive maintenance framework on an industrial scale.

The overall assessment confirmed that the proposed system is technically and economically feasible for deployment in industrial environments. With appropriate investment in IT infrastructure, data governance, and personnel training, the framework can be effectively implemented at full scale. Its reliance on open standards, modular design, and scalable software architecture ensures compatibility with existing automation and enterprise systems.

Consequently, the DT–ERP predictive maintenance architecture presents a practical and replicable model for achieving autonomous, data-driven maintenance operations in various industrial sectors beyond composting—including agriculture, energy, and process manufacturing.

8. Conclusions

This study presented an integrated predictive maintenance framework that combines digital twin (DT) technology, enterprise resource planning (ERP) systems, and machine learning algorithms to enhance the reliability and automation of industrial machinery. The framework was specifically developed and validated for composting plants but is fully adaptable to other processes and manufacturing industries where mechanical systems play a critical operational role. By linking real-time data acquisition with virtual modeling and intelligent analytics, the system establishes a cyber–physical environment capable of autonomous monitoring, fault prediction, and maintenance decision-making.

The experimental implementation demonstrated that the integration of DT and ERP layers can significantly improve maintenance efficiency. The digital twin enabled continuous modeling and performance analysis of key machinery, while the ERP system facilitated coordinated decision-making, data management, and resource optimization. The predictive models, combining physics-based simulations with LSTM neural networks, achieved high accuracy in fault detection and process prediction, with MAE values below 2.5, F1-scores above 89%, and an AUC of 0.92. At the operational level, the system increased the mean time to failure (MTTF) by 40%, reduced the mean time to repair (MTTR) by 35%, and

lowered annual maintenance costs by 30%, yielding a return on investment (ROI) of 42.5% within the first year.

Beyond its quantitative results, the proposed framework provides a generalizable engineering solution for predictive maintenance and automation. Its modular and interoperable architecture—built upon standard industrial communication protocols (OPC-UA, RESTful APIs)—ensures scalability and compatibility with a wide range of machinery and control systems. The ability to integrate physical modeling, predictive analytics, and enterprise management within a single loop highlights its potential as a foundation for self-maintaining and adaptive industrial systems.

The findings of this research contribute to the broader field of machine design and automation by demonstrating how digital twin-based predictive systems can improve both the operational reliability and lifecycle efficiency of industrial assets. In doing so, the study advances the development of intelligent, data-driven maintenance frameworks aligned with the principles of Industry 4.0 and Industry 5.0. Future research should extend this work by incorporating reinforcement learning, real-time optimization, and cross-facility integration, enabling the evolution of fully autonomous maintenance ecosystems in complex manufacturing environments.

Author Contributions: Conceptualization, H.N. and A.S.-J.; methodology, H.N.; software, H.N.; validation, H.N. and A.S.-J.; formal analysis, H.N.; investigation, H.N.; resources, H.N.; data curation, H.N.; writing—original draft preparation, H.N. and A.S.-J.; writing—review and editing, H.N. and A.S.-J.; visualization, H.N.; supervision, H.N.; project administration, H.N.; funding acquisition, A.S.-J. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The original data presented in the study are openly available in FigShare at <https://doi.org/10.6084/m9.figshare.30569987>.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Vargas, J.M.; Castrillon, O.D.; Giraldo, J.A. Implementation and Field Validation of a Digital twin Methodology to Enhance Production and Service Systems in Waste Management. *Appl. Sci.* **2025**, *15*, 6733. [CrossRef]
2. Monteiro, J.; Barata, J. A Circular Digital twin for Climate-Smart Soils through Community Composting. In *Advances in Information Systems Development: Exploring New Horizons and Opportunities*; Springer Nature: Cham, Switzerland, 2025; pp. 167–188.
3. Abd Wahab, N.H.; Hasikin, K.; Lai, K.W.; Xia, K.; Bei, L.; Huang, K.; Wu, X. Systematic review of predictive maintenance and digital twin technologies challenges, opportunities, and best practices. *PeerJ Comput. Sci.* **2024**, *10*, e1943. [CrossRef] [PubMed]
4. Aliahmadi, A.; Nozari, H.; Ghahremani-Nahr, J.; Szmelter-Jarosz, A. Evaluation of key impression of resilient supply chain based on artificial intelligence of things (AIoT). *arXiv* **2022**, arXiv:2207.13174. [CrossRef]
5. Nasirinejad, M.; Afshari, H.; Sampalli, S. *The Adoption of Digital Twin Technologies for Maintenance in Small and Medium-Sized Enterprises: Challenges and Benefits*; SSRN eLibrary: Rochester, NY, USA, 2025. [CrossRef]
6. Jesus, V.; Kalaitzi, D.; Batista, L.; Lopez, N.L. Digital twins of Supply Chains: A Systems Approach. *IEEE Trans. Eng. Manag.* **2024**, *71*, 14915–14932. [CrossRef]
7. Kannan, S. *Transforming Agriculture for the Digital Age: Integrating Artificial Intelligence, Cloud Computing, and Big Data Solutions for Sustainable and Smart Farming Systems*; Deep Science Publishing: San Francisco, CA, USA, 2025.
8. Koç, E. Measurement and Assessment of Supply Chain Sustainability Performance. In *Developing Dynamic and Sustainable Supply Chains to Achieve Sustainable Development Goals*; IGI Global Scientific Publishing: Hershey, PA, USA, 2025; pp. 179–200.
9. Abdulhayan, S.; Quadri, S.A. *Artificial Intelligence, IoT, and Fuzzy Systems for Sustainable Development and Industry 5.0*; Deep Science Publishing: San Francisco, CA, USA, 2025.
10. Tsai, W.H.; Lai, S.Y. Green production planning and control model with ABC under industry 4.0 for the paper in-dustry. *Sustainability* **2018**, *10*, 2932. [CrossRef]
11. Singh, R.; Dwivedi, G. Agri-food Supply Chain Management: A Review Using Bibliometric, Network, and Content Analyses. *Oper. Res. Forum* **2025**, *6*, 62. [CrossRef]

12. Aliahmadi, A.; Nozari, H. Evaluation of security metrics in AIoT and blockchain-based supply chain by Neutrosophic decision-making method. *Supply Chain Forum Int. J.* **2023**, *24*, 31–42. [CrossRef]
13. Emmanuel, A.T. An Evaluation of Production Operations at Skretting's Stavanger Plant, Norway, & an Overview of Skretting's Supply Chain Based on Lean Concepts. Ph.D. Dissertation, Royal Institute of Technology, Stockholm, Sweden, 2008.
14. Chen, K.; Zhao, B.; Zhou, H.; Zhou, L.; Niu, K.; Jin, X.; Li, R.; Yuan, Y.; Zheng, Y. Digital twins in plant factory: A five-dimensional modeling method for plant factory transplanter digital twins. *Agriculture* **2023**, *13*, 1336. [CrossRef]
15. Montreuil, B. Transforming the Composting Industry Towards Industry 4.0 and Physical Internet Capabilities. Ph.D. Dissertation, Georgia Institute of Technology, Atlanta, GA, USA, 2024.
16. Cichocki, M.; Landschützer, C.; Hick, H. Development of a Sharing Concept for Industrial Compost Turners Using Model-Based Systems Engineering, Under Consideration of Technical and Logistical Aspects. *Sustainability* **2022**, *14*, 10694. [CrossRef]
17. Barbosa, J.M.V.; Gomez, O.D.C.; García, J.A.G. Digital twin Application Methodology for the Improvement of Production and Service Systems. Application to Waste Management Processes. In Proceedings of the Sustainable Smart Cities and Territories International Conference, Manizales, Colombia, 21–23 June 2023; Springer Nature: Cham, Switzerland, 2023; pp. 25–36.
18. Danilo, J.M.V.B.O.; Gomez, C.; Piso, B.Q. Digital twin Application Methodology for the Improvement of Production and Service. *Trends Sustain. Smart Cities Territ.* **2023**, *732*, 25.
19. Shamia, D.; Suganyadevi, S.; Satheeswaran, V.; Balasamy, K. Digital twins in precision agriculture monitoring using artificial intelligence. In *Digital Twin for Smart Manufacturing*; Academic Press: Cambridge, MA, USA, 2023; pp. 243–265.
20. Giel, R.; Dąbrowska, A. Digital twin Concept for Manual Waste Sorting Management. *Environ. Prot. Eng.* **2025**, *51*, 1. [CrossRef]
21. Nyakudya, P.; Madushele, N.; Madyira, D.M.; Matare, E.E. Embracing Digital twin Technology to Enhance Efficacy of Industrial Symbiosis: A Path to Sustainability. In Proceedings of the 2024 International Conference on Electrical and Computer Engineering Researches (ICECER), Gaborone, Botswana, 4–6 December 2024; IEEE: Los Alamitos, CA, USA, 2024; pp. 1–6.
22. Mohammadi, H.; Ghazanfari, M.; Nozari, H.; Shafiezzad, O. Combining the theory of constraints with system dynamics: A general model (case study of the subsidized milk industry). *Int. J. Manag. Sci. Eng. Manag.* **2015**, *10*, 102–108. [CrossRef]
23. Vanjare, S.; Musale, J.; More, P.; Sayyad, T.; Gupta, Y.; Shikre, S. Implementation of Digital twin Technology for Predictive Crop Disease Monitoring. In Proceedings of the International Conference on Recent Advancements and Modernisations in Sustainable Intelligent Technologies and Applications (RAMSITA 2025), Indore, India, 7–8 February 2025; Atlantis Press: Dordrecht, The Netherlands, 2025; pp. 654–665.
24. dos Santos Silva, J.; Matos de Oliveira, A.; Veríssimo de Oliveira, J.; Bouzon, M. Barriers to digital transformation in fruit and vegetable supply chains: A multicriteria analysis using ISM and MICMAC. *OPSEARCH* **2025**, *62*, 460–482. [CrossRef]
25. Zaballos, A.; Briones, A.; Massa, A.; Centelles, P.; Caballero, V. A smart campus' digital twin for sustainable comfort monitoring. *Sustainability* **2020**, *12*, 9196. [CrossRef]
26. Yang, B.; Lv, Z.; Wang, F. Digital twins for intelligent green buildings. *Buildings* **2022**, *12*, 856. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Article

Hybrid Deep Learning for Predictive Maintenance in Industrial Machinery Using LSTM and MLP Models

Ioannis Mavrelis, Avraam Chatzopoulos and Paraskevi Zacharia *

Department of Industrial Design and Production Engineering, University of West Attica, Egaleo, 122 41 Athens, Greece; idpe19389218@uniwa.gr (I.M.); xatzopoulos@uniwa.gr (A.C.)

* Correspondence: p.zacharia@uniwa.gr; Tel.: +30-210-538-1055

Abstract

This paper introduces a practical hybrid framework for predictive maintenance in industrial machinery, integrating unsupervised anomaly detection via a long short-term memory (LSTM) autoencoder (AE) and supervised fault prediction using a multilayer perceptron (MLP). The framework is validated on a realistic digital-twin dataset simulating five years of operation, addressing common challenges such as class imbalance and operational variability. The LSTM autoencoder was trained on time-series data representing normal behavior and detected 129 threshold exceedances (anomalous excesses) out of 8735 test sequences (1.48%) based on a validation-derived reconstruction error threshold. The MLP classifier, trained on operational features and rule-derived fault labels, achieved a classification accuracy of 99.93%, with precision, recall and F1-score values of 99.53%, 99.88% and 99.70%, respectively. Data preprocessing included cleaning, robust scaling, and sequence generation to ensure data quality and model reliability. The results demonstrate that the proposed architecture effectively combines sequence-based anomaly detection with a compact classification model, offering a reliable solution. Beyond model design, the study provides practical guidance on threshold calibration, fault labeling, and SCADA/cloud integration, bridging the gap between theoretical research and industrial implementation. This work contributes a scalable, industry-ready predictive maintenance solution that combines strong technical performance with real-world applicability.

Keywords: predictive maintenance; anomaly detection (threshold exceedances); LSTM autoencoder; multilayer perceptron (MLP); fault prediction; vibration analysis; machine learning

1. Introduction

Predictive maintenance has become a critical strategy in the management of industrial machinery, driven by the need to minimize downtime, reduce operational costs, and enhance equipment reliability. The motivation for this study arises from the increasing importance of maintaining uninterrupted industrial operations and the substantial economic impact of unexpected machinery failures.

Traditionally, maintenance strategies have relied on scheduled inspections and reactive repairs [1,2]. These approaches often result in inefficient resource utilization and unplanned outages. In contrast, data-driven approaches use sensor data and advanced analytics to anticipate faults before they occur, enabling timely interventions and optimized maintenance schedules [3,4].

Recent advances in sensing and analytics have positioned predictive maintenance combining operational data with machine learning (ML) and deep learning (DL) [5,6] as a leading paradigm for early detection of deviations and failure prevention. Additionally, digital twins (DTs), which serve as high-fidelity virtual replicas of physical assets, provide safe and realistic testbeds for developing and evaluating predictive models.

Despite these advancements, several key challenges persist [7–9]. These include trade-offs between supervised and unsupervised schemes, static versus dynamic alert thresholds, balancing black-box accuracy with model explainability [10], and ensuring the external validity of models trained in simulated environments.

To address these gaps, this study proposes a practical two-level system that couples anomaly detection using an LSTM Autoencoder (AE) [11–13] based on reconstruction error with fault prediction via an MLP classifier on operational features. The system is evaluated on multivariate sensor data collected over five years in a digital twin scenario, demonstrating robust anomaly separation and strong fault-prediction performance. Deployment implications, including threshold setting and model calibration for production integration, are also discussed.

Although hybrid predictive maintenance frameworks combining unsupervised anomaly detection with supervised fault classification have been discussed in prior research, this study contributes by focusing on system-level design, deployment-oriented validation, and practical integration within a digital twin industrial setting. Unlike approaches that rely on fully multivariate deep models or require labeled anomaly data, our method uses a univariate LSTM autoencoder applied to a high-sensitivity vibration signal for early anomaly detection, followed by a compact MLP classifier operating on a carefully selected set of operational and environmental features for fault prediction. This design aims to balance detection sensitivity, interpretability, and computational efficiency, addressing real-world constraints such as limited labeled data, sensor diversity, and the need for real-time deployment.

The novelty of this work lies in delivering a lightweight, interpretable, and deployable predictive maintenance framework validated on a realistic digital-twin dataset simulating five years of industrial operation. Our approach explicitly addresses industrial constraints such as computational efficiency, SCADA/cloud integration, and threshold calibration, while maintaining state-of-the-art accuracy. Unlike prior hybrid predictive maintenance approaches that typically depend on fully multivariate deep architectures or assume the availability of labeled anomaly data, this work focuses on a deployment-oriented design that combines univariate sequence-based anomaly detection on a high-sensitivity vibration signal with compact multivariate fault classification, targeting interpretability, computational efficiency, and realistic industrial constraints.

The objective of this study is to design, implement, and evaluate a two-level predictive maintenance framework for industrial machinery using machine learning models. Specifically, the work focuses on developing an LSTM autoencoder for unsupervised anomaly detection and a multilayer perceptron (MLP) for supervised fault prediction. The models are trained and validated on multivariate time-series data generated in a digital twin environment that realistically simulates industrial operations over several years. By combining advanced simulation with practical deployment considerations, this research seeks to bridge the gap between theoretical innovation and real-world applicability, offering a scalable and effective approach to predictive maintenance that can be readily adopted in industrial settings.

This work addresses the identified gaps by introducing a practical two-level predictive maintenance framework that integrates unsupervised anomaly detection with supervised fault prediction in a way that is both lightweight and interpretable. Unlike prior studies

that focus primarily on algorithmic novelty, our approach emphasizes deployability and operational transparency, making it suitable for near real-time industrial environments. The framework is rigorously validated on a five-year digital-twin dataset that captures realistic operational variability, ensuring external validity and robustness. In addition, the study provides detailed guidance on threshold calibration, fault labeling strategies, and SCADA/cloud integration, bridging the gap between theoretical research and practical implementation. Collectively, these contributions position the proposed method as a scalable and industry-ready solution for predictive maintenance.

This paper addresses the challenge of predictive maintenance in industrial machinery by leveraging advanced machine learning techniques to improve early fault detection and operational reliability. The proposed approach integrates both unsupervised and supervised learning models, aiming to deliver a practical and scalable solution for real-world industrial environments. Building on the identified research gaps, the contributions of this work are explicitly aligned with those gaps and can be summarized as follows:

- Introduces an unsupervised LSTM autoencoder trained exclusively on nominal operating conditions for early anomaly detection, addressing the challenge of limited labeled anomaly data.
- Combines a univariate sequence-based anomaly detector with a compact MLP classifier operating on a small, physically meaningful feature set, resolving the trade-off between model complexity and practical deployability.
- Validates the framework on long-horizon, timestamped data from a realistic digital twin environment, ensuring external validity and industrial relevance while considering constraints such as interpretability and near real-time operation.

In summary, this work proposes a novel hybrid predictive maintenance framework that combines LSTM-based anomaly detection with MLP-based fault prediction, validates the approach on realistic multivariate time-series data from a digital twin, and designs the system to be lightweight and interpretable for practical deployment in industrial settings.

This paper is organized as follows: Section 2 presents the related work, providing an overview of traditional and modern approaches to anomaly detection and fault prediction in industrial systems. Section 3 describes the proposed methodology, including the design of the LSTM autoencoder and MLP classifier, as well as data preprocessing steps. Section 4 outlines the experimental setup and the digital twin environment used for data generation. Section 5 presents the main results and discusses the performance of the proposed framework. Section 6 offers a broader discussion of the findings, practical implications, and limitations, and suggests directions for future research. Section 7 concludes the paper with a summary of key contributions.

2. Related Work

In recent years, fault detection and predictive maintenance for synchronous motors, particularly permanent magnet synchronous motors (PMSMs), have advanced significantly. Gherghina et al. [14] provide a comprehensive review of the field, noting the increasing adoption of data-driven and hybrid approaches, the integration of sensor fusion, and the emergence of explainable AI. However, challenges such as data imbalance, non-stationary conditions, and limited real-world generalization persist.

Anomaly detection and fault prediction in industrial systems have evolved from traditional rule-based and statistical approaches to advanced data-driven methods. Classical techniques, such as observer-based models and residual analysis, are effective in well-characterized environments [15–17], but often struggle with complex, high-dimensional data. In contrast, machine learning and deep learning approaches including autoencoders, convolutional neural networks (CNNs), and recurrent neural networks (RNNs) like LSTM,

have shown strong capabilities in learning intricate patterns from sensor data. Autoencoders are widely used for unsupervised anomaly detection by using reconstruction error, while LSTM networks excel at modeling temporal dependencies in time-series data. Hybrid and generative models (e.g., generative adversarial networks (GANs) and variational autoencoders (VAEs)) further enhance detection performance [18–20]. Despite these advances, challenges remain regarding model interpretability, data imbalance, and the need for large labeled datasets. Recent research emphasizes integrating signal processing, feature engineering, and adaptive frameworks to improve robustness and enable real-time predictive maintenance.

In practical terms, Vlachou et al. [21] demonstrate the effectiveness of advanced ML techniques, including GAN-augmented PU learning and reinforcement learning, for robust, non-intrusive fault prediction in elevator PMSM drives. Their work underscores the importance of multimodal sensor integration and real-time data preprocessing for scalable predictive maintenance solutions.

Ahmed et al. [22] introduced a smart anomaly detection system for industrial machines utilizing a deep autoencoder framework. Their approach uses vibration analysis as a primary tool for machine health monitoring and fault diagnosis, focusing on a public gear-box dataset collected from wind-turbine components. The proposed six-layer autoencoder model autonomously extracts salient features from high-dimensional vibration data, enabling effective detection of both known and previously unseen anomalies without the need for manual feature engineering or extensive preprocessing. Experimental results demonstrated that the framework achieved an overall accuracy of 91%, outperforming several traditional machine learning methods such as SVM, random forest, and k-nearest neighbors. This work highlights the potential of deep learning-based autoencoders for robust, scalable, and generalizable anomaly detection in industrial condition monitoring applications.

Do et al. [23] developed an LSTM-Autoencoder for vibration anomaly detection in a Vertical Carousel Storage and Retrieval System (VCSRS). Their approach combines optimal sensor placement using correlation coefficient and Fisher information matrix methods with deep learning for robust monitoring. The LSTM-Autoencoder is trained to reconstruct normal vibration patterns; anomalies are detected when the reconstruction error exceeds a set threshold. This method achieved an accuracy of 97.7% in distinguishing normal from faulty states, demonstrating the effectiveness of combining advanced sensor strategies with deep learning for predictive maintenance in complex industrial systems. The study also highlights the importance of data preprocessing, feature extraction, and addressing class imbalance to ensure reliable anomaly detection in real-world applications.

Apeiranthitis et al. [24] proposed a predictive maintenance approach for rotating machinery using a one-dimensional convolutional neural network (1D CNN) trained on raw vibration data from ball bearings. Their study focused on the maritime sector, where machinery reliability is critical for operational efficiency and safety. The authors demonstrated that their 1D CNN could accurately detect and classify various ball bearing faults using vibration signals acquired in a laboratory environment. The model achieved high classification accuracy, even when tested on previously unseen data, and was shown to be computationally efficient, making it suitable for real-time monitoring applications. This work highlights the potential of deep learning, and specifically CNN architectures, for effective fault diagnosis and predictive maintenance in complex industrial systems.

Sajjadi et al. [25] present a comprehensive review of machine learning (ML) methodologies for prognostics and health management (PHM) in cyber-physical systems (CPSs). The authors propose a novel taxonomy that classifies existing research by fault types (physical and cyber), PHM stages, data characteristics, ML techniques, and performance metrics, providing a structured guide for selecting appropriate ML methods across the PHM lifecycle.

The review highlights the increasing adoption of deep learning, including transformers and large language models, for tasks such as anomaly detection, fault diagnosis, and remaining useful life (RUL) prediction. Key challenges identified include the scarcity of real-world, labeled datasets, class imbalance, model generalization, and the need for explainable AI in safety-critical applications.

3. Model Architectures

3.1. LSTM Autoencoder for Anomaly Detection

Long Short-Term Memory (LSTM) networks are recurrent architectures designed to capture temporal dependencies through gated mechanisms (input, forget, and output gates) thereby effectively addressing vanishing and exploding gradient issues. This capability makes them well-suited for industrial time-series data, where current states are strongly influenced by recent operational history. Autoencoders (AEs), in turn, aim to reconstruct input sequences; when trained on nominal conditions, they reproduce normal patterns with minimal error, whereas deviations result in elevated reconstruction errors—a widely adopted principle in anomaly detection research [26].

In industrial contexts, individual sensor signals such as vibration, temperature and pressure exhibit complex temporal dependencies that evolve over time and cannot be captured by static models operating on single snapshots of the data. In this work, the LSTM-based autoencoder is applied to one sensor channel at a time as a univariate time series (e.g., the vibration level), learning dynamic dependencies across multiple time steps and thereby modelling the normal sequential behaviour of that signal. The full set of six operational features (vibration, temperature, pressure, motor speed, torque and humidity) is instead used as a static feature vector by the MLP fault classifier described in Section 4.4.

Although multiple correlated sensor signals are available, the LSTM Autoencoder was intentionally applied to a single high-sensitivity signal (vibration) in a univariate manner. Vibration signals are widely regarded as primary indicators of mechanical degradation in rotating machinery and are commonly used for early anomaly detection in industrial practice [27–30].

This design choice reduces model complexity, improves interpretability, and facilitates deployment in resource-constrained environments. Multivariate information is instead exploited in the fault prediction stage, where the MLP classifier processes a compact set of operational features to provide actionable fault diagnoses. While a multivariate LSTM Autoencoder could capture cross-sensor dependencies, such an extension would significantly increase architectural complexity and is therefore left for future investigation.

The proposed LSTM-based autoencoder compresses each input sequence into a latent representation and subsequently reconstructs it in the original signal space. This latent representation corresponds to the final hidden state vector produced by the LSTM encoder, which has lower dimensionality than the original sequence and summarizes its most salient temporal patterns. By constraining the information to pass through this compact internal code, the model is forced to learn meaningful features rather than simply copying the input. The reconstruction error serves as an unsupervised anomaly score: sequences exceeding a statistically derived threshold are classified as anomalous [26]. This method provides a single interpretable metric per sequence, enabling efficient monitoring and alarm generation.

In many industrial predictive maintenance settings, explicit ground-truth labels for anomalous behavior are unavailable, incomplete, or unreliable, as abnormal events are rare and often not consistently documented. For this reason, the anomaly detection component in this study is formulated as an unsupervised early-warning mechanism rather than a supervised classifier.

Consequently, conventional classification metrics such as precision, recall, or false-alarm rate cannot be directly computed for anomaly detection. Model evaluation therefore relies on statistical properties of the reconstruction error distribution, including separation between nominal and abnormal sequences, sparsity and localization of threshold exceedances, and qualitative consistency with observable signal deviations. This evaluation strategy is widely adopted in unsupervised anomaly detection studies and reflects realistic industrial constraints where labeled anomalies are not available at deployment time.

The anomaly threshold is computed using the distribution of reconstruction errors on validation data, typically set as $\tau = \mu + 3\sigma$, where μ and σ denote the mean and standard deviation of reconstruction errors. This statistical approach ensures that the threshold is data-driven and minimizes false positives under normal conditions. In our implementation, the threshold was fixed using validation reconstruction errors prior to test-time evaluation, in order to avoid data leakage from the test set and maintain model integrity. The validation windows used to estimate the reconstruction error distribution and determine the anomaly threshold were selected as statistically stable operating periods and were not defined based on fault labels or rule-based criteria. These windows correspond to time intervals exhibiting consistent signal behavior without pronounced fluctuations, reflecting nominal operating conditions over the observed timescale. As a result, the derived threshold is purely statistical and data-driven, rather than fault-aware or label-dependent.

The autoencoder architecture shown in Figure 1 consists of:

- Input layer: Receives univariate time-series sequences of fixed length (e.g., 30 time steps, 1 feature).
- LSTM layer: A single LSTM layer with 64 units and sequence output that learns a temporal representation of the input sequence.
- Fully connected layer: A dense layer with 1 unit per time step that maps the LSTM outputs back to the reconstructed sequence (decoder stage).
- Regression layer: A regression layer (MSE loss) that compares the input and reconstructed sequences and computes the reconstruction error for each sequence.




Figure 1. LSTM-Autoencoder Architecture.

Training is performed using the Adam optimizer [31,32] with a learning rate schedule and gradient clipping to ensure stability. Hyperparameters such as sequence length,

number of LSTM units, and batch size were selected based on empirical testing and literature benchmarks for industrial anomaly detection.

The selection of the sequence length and the number of hidden units for the LSTM Autoencoder was based on a combination of empirical experimentation and commonly adopted practices in time-series anomaly detection. A sequence length of 30 time steps was chosen in order to capture short- to medium-term temporal dependencies present in the sensor signals, while avoiding excessive model complexity and computational overhead. Preliminary experiments with shorter windows were found insufficient to adequately represent the system dynamics, whereas longer sequences did not provide a noticeable improvement in reconstruction performance.

Similarly, the number of hidden units was set to 64 as a compromise between representational capacity and generalization ability. Smaller configurations limited the model's ability to learn normal operational patterns, leading to higher reconstruction errors on nominal data, while larger architectures increased the risk of overfitting without clear performance gains. The selected configuration was therefore validated empirically and is consistent with architectures commonly reported in the literature for LSTM-based anomaly detection in industrial time-series data.

This design offers several advantages. First, it operates in an unsupervised manner, eliminating the need for labeled anomaly data and making it particularly suitable for early warning applications in real-world environments where fault instances are scarce. Furthermore, the reconstruction error provides a single, interpretable scalar score, which simplifies integration into monitoring dashboards and enhances transparency. Finally, the lightweight architecture ensures scalability, allowing seamless deployment within SCADA systems or cloud-based platforms for real-time anomaly detection.

The rationale for this design lies in the ability of LSTM layers to capture short- and medium-term temporal dependencies typical of rotating machinery signals, while the autoencoder objective ensures accurate reproduction of nominal dynamics. The following section details the encoder–decoder configuration and the thresholding strategy employed within the proposed framework.

3.2. Multilayer Perceptron (MLP) for Fault Prediction

Multilayer Perceptrons (MLPs) [33] are feedforward neural networks that map input features to target outputs through successive layers of linear transformations combined with nonlinear activation functions. This architecture is well-suited for modeling complex, nonlinear relationships between operational variables such as vibration, temperature, and pressure, and the occurrence of faults in industrial systems.

In the proposed framework, the MLP processes a feature vector extracted from each time window and outputs a scalar score, which is subsequently converted into a binary decision (fault or no fault) using a predefined threshold. Concretely, for each window the six core operational variables ($x_1, x_2, x_3, x_4, x_5, x_6$), recorded at the final time step are concatenated into a six-dimensional feature vector $x \in R^6$ that represents the instantaneous operating point of the machine. This decision threshold is typically set at 0.5 for binary classification tasks, ensuring a balanced trade-off between false positives and false negatives. Feature selection plays a critical role in this process and can be guided by operational relevance and statistical significance. In addition to manual selection, dimensionality reduction techniques such as Principal Component Analysis (PCA) have been widely employed in prior studies to enhance input representation and reduce redundancy [34].

The MLP architecture (Figure 2) implemented in this study consists of:

- **Input Layer:** Six operational features, normalized using z-score transformation to ensure comparability across scales.

- **Hidden Layer:** A single hidden layer with ten neurons employing the tansig activation function, chosen for its ability to capture nonlinear relationships while maintaining computational efficiency.
- **Output Layer:** A single neuron with a logsig activation function, producing a continuous score that is thresholded to yield binary fault predictions.




Figure 2. Simple Multi-Layer Perceptron ANN.

Training was conducted using the Levenberg–Marquardt algorithm (`trainlm`) [35,36], which is well-suited for small to medium-sized networks due to its fast convergence properties. The dataset was partitioned into training (80%) and validation (20%) subsets using stratified sampling to preserve class distribution. Hyperparameters such as learning rate, maximum epochs, and early stopping criteria were optimized to prevent overfitting and ensure generalization.

The inclusion of an MLP offers several practical benefits. First, its lightweight architecture ensures rapid inference, making it suitable for real-time deployment in SCADA or cloud-based environments. Second, the supervised nature of the model enables actionable fault predictions by using labeled operational conditions, complementing the unsupervised anomaly detection stage. Finally, the simplicity of the network facilitates interpretability and integration into existing industrial monitoring systems without significant computational overhead.

4. The Proposed Methodology

To address the challenges identified earlier, this work introduces a practical two-level predictive maintenance framework designed for industrial machinery. The methodology integrates unsupervised anomaly detection and supervised fault prediction, applying advanced deep learning architectures to analyze complex operational data. By combining the strengths of LSTM autoencoders and multilayer perceptrons, the proposed approach aims to deliver robust early warning capabilities and actionable insights for maintenance

planning. The following subsections detail the data sources, preprocessing steps, model architectures, and evaluation protocols that underpin the system's development and validation. Within this framework, the LSTM autoencoder and the MLP classifier operate as two parallel, loosely coupled modules that process the same underlying data streams; they do not feed into each other directly, and their outputs are combined only at the decision level. To provide a unified view of the complete processing workflow, Figure 3 presents a flowchart that summarizes all stages of the proposed two-level predictive maintenance framework. The diagram illustrates the entire data flow, from acquisition and preprocessing to anomaly detection and fault prediction, highlighting the parallel operation of the LSTM Autoencoder and the MLP classifier, and their fusion at the decision level. This visual representation supports the detailed explanations in the following subsections and clarifies the overall system logic and deployment workflow.




Figure 3. System workflow for the two-level predictive maintenance approach.

4.1. Data Sources

As shown in Figure 3, the anomaly detection and fault prediction modules operate in parallel, processing the same data stream independently. Their outputs are only integrated at the decision layer, resulting in a modular system. This architecture allows each module to be optimized and deployed separately, while still ensuring that the overall predictive maintenance workflow remains coherent and easy to interpret. This study utilizes multi-variate time series data generated within a digital twin (DT) environment that simulates industrial machinery operation over a five-year period (January 2019 to January 2024), with hourly sampling intervals. The dataset comprises 38 distinct features, including sensor measurements (vibration, temperature, pressure, humidity), operational parameters (such as motor speed and torque), machine health indicators, and environmental variables. Notably, the occurrence of fault events within the dataset is imbalanced, reflecting realistic industrial conditions [37].

Two distinct datasets were prepared for the analysis to ensure proper model development and evaluation. The first dataset was designated for training and included labeled fault events, enabling supervised learning for fault prediction. The second dataset was reserved for testing and consisted of 8735 time-windowed sequences generated after applying preprocessing steps such as data cleaning, normalization, and sliding-window segmentation. This separation ensured that the models were trained on representative operational conditions while being validated on unseen sequences, supporting an unbiased assessment of anomaly detection and fault prediction performance.

The datasets were designed to emulate real factory conditions, capturing both normal and fault states, with varying severity levels. Each sample includes operational, environmental, and health indicators, enabling comprehensive modeling of machine behavior. The imbalanced distribution of faults presents a realistic challenge for classification and anomaly detection.

4.2. Data Preprocessing

To ensure data quality and consistency, several preprocessing steps were applied:

- **Handling Missing Data:** Non-numeric entries and missing values (NaN) were detected and addressed using median imputation for each column. Samples exhibiting extensive gaps were removed to maintain dataset integrity.
- **Normalization:** Robust scaling was performed, where each value was transformed by subtracting the median and dividing by the interquartile range (IQR). Values were clipped to the range $[-3, 3]$ to reduce the influence of outliers [38,39].
- **Sequence Generation:** For the LSTM-based analysis, a sliding window approach with a sequence length of 30 and overlapping windows was used to construct three-dimensional tensors compatible with the sequence input layer. Randomness in data partitioning was controlled using a fixed random seed (`rng(42)`) [13,38].

Prior to modeling, exploratory data analysis was performed, including visualization of key signals (e.g., vibration, temperature, pressure) to identify outliers, abrupt changes, and data gaps. This step informed the selection of features and the design of preprocessing procedures. The robust normalization approach was chosen for its resilience to outliers and heterogeneous data scales, ensuring stable model training.

Robustness to sensor noise and moderate operating variability was addressed through both preprocessing and model design. Robust normalization based on median and interquartile range statistics was applied to reduce sensitivity to outliers and transient fluctuations, while value clipping further limited the impact of extreme measurements.

Additionally, the use of sliding windows and sequence-based modeling enables the LSTM Autoencoder to learn temporal patterns rather than instantaneous values, inherently

smoothing short-term noise. This windowed learning mechanism also contributes to partially reducing moderate sensor drift by continuously re-evaluating local temporal dynamics within each sequence. However, long-term sensor drift and severe non-stationary operating regimes were not explicitly modeled in this study. Addressing such effects would require adaptive thresholding or online model updates, which are identified as important directions for future work.

4.3. Anomaly Detection Using LSTM Autoencoder

An unsupervised anomaly detection module was implemented using a Long Short-Term Memory (LSTM) autoencoder [26]. The model architecture consisted of a sequence input layer, an LSTM layer with 64 units (output mode set to 'sequence'), a fully connected layer, and a regression output layer [13,23]. Training was conducted using the Adam optimizer with the following parameters: maximum epochs set to 100, initial learning rate of 0.005, gradient threshold of 1, and scheduled learning rate drops.

The decision rule for anomaly detection was based on the reconstruction error, calculated as the mean absolute error (MAE) per sequence on validation data representing normal operation. An anomaly threshold was established as $\mu + 3\sigma$ from the distribution of validation errors and fixed prior to evaluation on the test set. Sequences exceeding this threshold were identified as anomalies [11,12]. Outputs from this stage included a comprehensive anomaly report and visualizations of reconstruction errors and anomaly overlays.

The LSTM Autoencoder was specifically trained to reconstruct sequences of vibration-level data, using the temporal dependencies inherent in industrial signals. The choice of sequence length (30) and hidden units (64) was guided by empirical testing and literature benchmarks for time series anomaly detection. The model's unsupervised nature allows it to detect deviations without requiring labeled anomaly events, making it suitable for early warning in real-world deployments. Detected anomalies are further analyzed by plotting their temporal distribution and overlaying them on the original vibration signal, facilitating interpretation and validation by domain experts.

4.4. Fault Prediction Using Multilayer Perceptron (MLP)

For supervised fault prediction, a multilayer perceptron (MLP) classifier was implemented to distinguish between normal and abnormal operating conditions. The target variable, fault diagnosis, was defined as a binary indicator indicating the presence or absence of a fault based on predefined operational thresholds for key parameters such as vibration, temperature, and pressure. The fault labels employed in this study are derived from predefined operational thresholds applied to vibration, temperature, and pressure signals. As such, these labels do not correspond to independently annotated failure events, but rather encode rule-based operational limits commonly used in industrial monitoring practice. Consequently, the role of the MLP classifier is not to discover novel fault patterns or infer failures beyond these rules. Instead, it aims to learn a smooth, data-driven approximation of the underlying deterministic decision logic, enabling the assessment of rule consistency and their reproducibility through a compact neural model. These thresholds were established based on industry best practices to ensure that the classification reflects realistic safety conditions and supports actionable maintenance decisions.

The input space consisted of six operational features representing key machine and environmental parameters. All features were normalized using a z-score transformation [10], with statistics computed across the full dataset to maintain consistency. Although the original digital twin dataset contained 38 heterogeneous features, six core operational variables, namely vibration, temperature, pressure, motor speed, torque and humidity, were used as inputs to the MLP fault classifier. The choice was guided by domain knowledge in

condition monitoring and preliminary exploratory analysis. Vibration, temperature and pressure serve as primary indicators of mechanical and thermal stress, whereas motor speed and torque reflect the machine's operating point and load; humidity accounts for environmental factors that may accelerate degradation. Other variables, such as maintenance logs, alarm triggers, health indices, anomaly scores or fault probabilities were not included, because they often depend on derived or historical data and may not be consistently available in real time for new installations. By using a compact, physically meaningful subset reduces input dimensionality, limits overfitting under class imbalance, and maintains interpretability, an approach aligned with recent vibration-based predictive maintenance studies employing deep learning architectures [34].

The MLP architecture consisted of a single hidden layer with ten neurons using tansig activation function and a single linear output neuron; and an output layer with a logsig activation function. Binary labels were obtained by applying a decision threshold of 0.5 to the output score. Model training was conducted using the Levenberg–Marquardt algorithm (trainlm), with dataset partitioned into 80% for training and 20% for validation (through cv-partition function) to ensure balanced representation and reliable performance assessment.

Among the most widely used evaluation measures are accuracy, precision, recall, and the F1-score. Accuracy quantifies the proportion of correctly predicted instances relative to the total number of observations, offering an overall measure of performance. Precision reflects the proportion of correctly predicted positive instances among all instances classified as positive, making it particularly relevant when false positives carry a high cost. Recall (or sensitivity) measures the proportion of actual positive instances correctly identified by the model, which is critical when minimizing false negatives is a priority. Finally, the F1-score, defined as the harmonic mean of precision and recall, provides a balanced metric that accounts for both types of misclassification. These metrics provide a comprehensive assessment of classifier performance, offering deeper insight into the trade-off between false alarms and missed detections compared to accuracy alone.

Performance evaluation relied on the confusion matrix and derived metrics. True positives (TP) represent correctly identified fault samples, while true negatives (TN) correspond to correctly classified normal samples. False positives (FP) indicate normal samples incorrectly identified as faults and false negatives (FN) denote missed fault cases.

The mathematical formulations for these metrics are presented in Equations (1)–(4).

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{FP} + \text{TN} + \text{FN}} \quad (1)$$

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \quad (2)$$

$$\text{Recall/Sensitivity} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (3)$$

$$\text{F1 - score} = 2 \cdot \frac{\text{Sensitivity} \cdot \text{Precision}}{\text{Precision} + \text{Sensitivity}} \quad (4)$$

By constructing fault labels based on practical safety thresholds, the proposed approach delivers actionable predictions that support proactive maintenance planning. Furthermore, the trained MLP was prepared for integration into digital monitoring systems, enabling both offline analysis and potential real-time deployment within SCADA or cloud-based environments.

The proposed framework was explicitly designed to support lightweight execution and continuous operation on time-indexed industrial data streams. Both the LSTM Auto-encoder and the MLP classifier adopt compact architectures, consisting of a single recurrent layer and a shallow feedforward network, respectively. Model training is intended to be

performed offline or periodically, while inference operates on fixed-size sliding windows and can be executed independently for each incoming time window.

To provide an indication of computational efficiency, training and inference times were monitored during experimentation in a MATLAB environment using a single CPU. Training of the LSTM Autoencoder completed within approximately 18 min for 100 epochs, while per-window inference for both the anomaly detection and fault prediction stages required only a few milliseconds on standard industrial hardware. These results indicate that the proposed framework is suitable for near real-time monitoring and decision support.

It should be noted that these runtime values are indicative and depend on the specific hardware and software configuration. Nevertheless, the observed execution times, together with the compact model design and fixed temporal windowing, support the practical deployability of the framework in SCADA, edge, or cloud-based industrial environments.

5. Simulation Results

The source dataset used in this study is publicly available at <https://www.kaggle.com/datasets/datasetengineer/indfd-pm-dt> (accessed on 5 September 2025). All data employed in our experiments were drawn from this source. We first partitioned the original dataset into separate training and test subsets, resulting in a training set with 35,060 records and a test set with 8765 records. For the LSTM autoencoder, these time series were further transformed into overlapping windows of 30 time steps, yielding 35,030 training sequences and 8735 test sequences. For the MLP classifier, 80% of the training data (28,048 samples) was used for model training and 20% (7012 samples) for validation.

Experimental procedures and analyses were carried out in MATLAB R2023a. The LSTM autoencoder and MLP models were implemented using the Deep Learning Toolbox, while preprocessing, normalization, and data partitioning were supported by the Statistics and Machine Learning Toolbox. Additionally, the computational environment also included advanced visualization tools, enabling efficient experimentation and clear interpretation of the results.

5.1. LSTM Autoencoder—Anomaly Detection

The anomaly detection phase focused on the vibration time series, utilizing an LSTM Autoencoder trained exclusively on sequences representing normal machine operation. Before model training, the dataset underwent comprehensive preprocessing, including the removal of missing or invalid entries and the application of robust normalization based on median and interquartile range statistics to reduce the influence of outliers. Time series were then transformed into overlapping sequences of uniform length, ensuring the preservation of temporal dependencies critical for effective anomaly detection.

Model architecture comprised a sequence input layer, an LSTM layer with 64 units, a fully connected output layer, and a regression layer to compute reconstruction error. Training was performed using the Adam optimizer, with hyperparameters selected to balance learning stability and computational efficiency.

Anomaly identification was based on the mean absolute error (MAE) between original and reconstructed sequences. The anomaly threshold was statistically defined as $\tau = \mu + 3\sigma$, where μ and σ denote the mean and standard deviation of reconstruction errors on the validation set. In this study, the threshold was set to $\tau = 1.26 \times 10^{-3}$. Applying this threshold to the test data, the model classified 129 sequences as threshold exceedances (anomalous excesses), corresponding to sparse and non-contiguous deviations from normal behavior (Figure 4).




Figure 4. Reconstruction error (MAE) for the vibration signal with the anomaly threshold τ (dashed red line). The y-axis is scaled in $\times 10^{-3}$.

A zoomed view (Figure 5) revealed that these exceedances were clustered around specific regions, notably between indices approximately 2050 and 2400. Overlaying the detected anomalies on the raw signal (Figure 6) confirmed that the identified events were localized and not attributable to general signal variability. These threshold exceedances do not necessarily correspond to actual fault events, but rather indicate statistically significant deviations from the learned normal behavior of the system. A sample excerpt of the anomaly log is presented in Table 1, which lists representative anomalous sequences together with their sample index, reconstruction error (MAE $\times 10^{-3}$), and corresponding vibration value.




Figure 5. Zoomed view of the reconstruction error around the exceedance region; spikes above τ correspond to detected anomalies.




Figure 6. Time-series plot showing the vibration signal (blue) with detected anomalies (red markers).

Table 1. Sample of detected anomalies.

Anomaly ID	Sample Index	MAE ($\times 10^{-3}$)	Signal Value
1	458	1.303	0.039
10	721	1.434	−0.044
27	2398	1.367	−0.579
33	3043	1.496	2.709
47	3300	1.797	3.000
69	4790	1.342	−0.435
89	5807	1.263	0.794
105	6943	1.553	1.322
108	7262	1.814	3.000
122	8025	1.614	2.430

Note: τ computed on validation normal data and fixed before test evaluation (no leakage).

5.2. MLP—Fault Prediction

For fault prediction, a Multi-Layer Perceptron (MLP) classifier was developed to distinguish between normal and faulty machine states. The labeled dataset exhibited a pronounced class imbalance, with 87.74% of samples representing normal operation and 12.26% indicating faults. Fault labels were assigned based on practical operational thresholds for vibration, temperature, and pressure, reflecting realistic industrial criteria.

The MLP model was constructed using six input features representing key operational parameters: vibration (Vibration_Level), temperature (Temperature_Readings in degrees Celsius), pressure (Pressure_Data in Pascals), humidity (Humidity_Levels in percentage), motor speed (Motor_Speed in RPM), and torque (Torque_Data in Nm). Each feature was normalized via z-score transformation to ensure comparability across scales. The network architecture consisted of a single hidden layer with ten neurons employing the tansig activation function and an output layer with a logsig activation function. Training was performed using the Levenberg–Marquardt algorithm with the dataset partitioned into training and validation subsets to ensure unbiased performance evaluation.

On the validation set, the MLP achieved an accuracy of 99.93%, as shown by the confusion matrix, which indicated minimal misclassifications (TN = 6155, FP = 4, FN = 1, TP = 852), as shown in Figure 7. It should be emphasized that this high classification

accuracy reflects the model's ability to consistently reproduce the predefined rule-based fault labeling scheme. Therefore, the reported performance should be interpreted in the context of rule consistency learning rather than independent fault classification against externally validated ground truth data. In practical terms, only four normal samples were incorrectly classified as faults (false positives) and one true fault was missed (false negative). These results translated into class-specific metrics of 99.53% precision (proportion of predicted faults that are actual faults), 99.88% recall (proportion of true faults correctly detected), and F1-score of 99.70%, representing the harmonic mean of precision and recall.




Figure 7. Confusion matrix for the MLP fault classifier (validation set).

To further account for class imbalance, additional evaluation metrics were derived from the confusion matrix. The specificity (true negative rate) reached 99.94%, indicating an extremely low false-alarm rate. The balanced accuracy, defined as the average of recall and specificity, was 99.91%, confirming consistent performance across both majority and minority classes. Finally, the Matthews correlation coefficient (MCC) achieved a value of 0.999, reflecting near-perfect agreement between predicted and true class labels even under pronounced class imbalance.

5.3. Stability Analysis Across Data Splits and Sequence Length

To evaluate the robustness and stability of the fault prediction model, additional experiments were conducted across multiple data splits and sequence-length configurations. The full dataset was divided into five different train–test splits, each maintaining the original class imbalance. For each split, the MLP classifier was trained and tested using the same hyperparameters. In addition, the LSTM autoencoder was tested with sequence lengths of 20, 30, and 50 time steps to assess how sensitive the model is to the choice of temporal window.

Across a total of ten experimental runs, the MLP classifier achieved high performance, showing very little variation across different splits and sequence lengths. Table 2 presents the mean and standard deviation of the evaluation metrics from all runs. These results indicate that the framework's performance does not depend on a particular data partition or sequence length, demonstrating strong generalization and robustness in realistic scenarios. In particular, both the Matthews correlation coefficient (MCC) and F1-score remained close to one, highlighting the classifier's reliability even when the data is highly imbalanced.

Table 2. Performance results of the MLP fault classifier.

Metric	Mean $\pm$ Std
Accuracy (%)	99.91 $\pm$ 0.05
Precision	0.9982 $\pm$ 0.0018
Recall	0.9948 $\pm$ 0.0039
F1-score	0.9963 $\pm$ 0.0023
MCC	0.9958 $\pm$ 0.0026


These performance metrics, summarized in Table 2, along with the threshold-based anomaly-detection results of the LSTM Autoencoder, demonstrate that the proposed classifier maintains exceptionally low false-alarm rates and an almost negligible probability of missing true faults. Table 3 reports the key performance indicators for both components of the predictive maintenance framework, reporting the anomaly-detection outcome of the LSTM Autoencoder (number of anomalies and threshold τ) along with the classification metrics of the MLP fault predictor (accuracy, precision, recall and F1-score).

Table 3. Summary of model performance.

Model	Task	Metric	Value	Notes
LSTM Autoencoder	Anomaly detection	Threshold exceedances (anomalous excesses)	129	Threshold $\tau = 1.26 \times 10^{-3}$ ($\mu + 3\sigma$)
MLP	Fault prediction	Accuracy	99.93%	Class ratio 12.26% positive
MLP	Fault prediction	Precision/Recall/F1-score	99.53%/99.88%/99.70%	From confusion matrix

5.4. Training Behavior and Generalization

To evaluate the stability and generalization of the proposed models, training dynamics were examined in detail. As shown in Figure 8, the loss curves for training, validation, and test sets decreased smoothly without abrupt fluctuations, indicating stable convergence and minimal risk of divergence. The lowest validation mean squared error (MSE) occurred at epoch 40, after which early stopping was applied at epoch 46 to prevent overfitting.

**Figure 8.** MSE vs. epochs; best validation at epoch 40 (dotted marker).

Error histograms for all data splits were tightly centered near zero with a narrow spread (Figure 9), and regression plots revealed near-identity fits (correlation coefficient $R \approx 0.99$), as shown in Figure 10. Figure 11 illustrates the training state, highlighting the progression of gradient decay, the schedule for the learning rate mean (LM μ), and the implementation of validation checks, which led to early stopping at epoch 46 to prevent overfitting. These results underscore the models' capacity to accurately capture underlying data patterns [6,40].




Figure 9. Error histogram (20 bins) for train/validation/test; vertical orange line denotes zero error.




Figure 10. Regression plots for train, validation, test, and all data ($R \approx 0.99$ across splits).




Figure 11. Training state: gradient decay, LM μ schedule, and validation checks (early stopping at epoch 46).

Error histograms in Figure 9 reveal that prediction errors across all data splits were tightly clustered around zero, with a narrow distribution, confirming strong consistency and low variance. Complementing this, Figure 10 presents regression plots that exhibit near-perfect correspondence between predicted and observed values, with a correlation coefficient of approximately $R \approx 0.99$, indicating strong model calibration and predictive accuracy.

Finally, Figure 11 provides a comprehensive view of the training state, illustrating the gradual decay of gradients, the adaptive learning rate schedule (LM μ), and the validation checkpoints that triggered early stopping. Collectively, these visualizations underscore the models' ability to capture underlying data patterns accurately while maintaining robustness against overfitting [6,40].

6. Discussion

This study assessed a two-level predictive maintenance framework integrating unsupervised anomaly detection with an LSTM Autoencoder and supervised fault prediction using an MLP. Consistent with our working hypothesis that reconstruction error can effectively signal deviations from learned normal behavior while a lightweight classifier can map operating features to fault risk, the LSTM Autoencoder identified 129 anomalous sequences on the DT-derived vibration signal using a validation-fixed threshold ($\tau = \mu + 3\sigma$; e.g., $\tau \approx 1.26 \times 10^{-3}$ MAE). These anomalies represented localized deviations rather than systemic drift, accounting for roughly 1.48% of the test set (129 out of 8735 sequences).

Complementing this, the MLP achieved 99.93% accuracy (TN = 6155, FP = 4, FN = 1, TP = 852) with precision, recall, and F1-score values of 99.53%, 99.88%, and 99.70%, respectively. These results indicate that a compact feedforward model with six operational inputs can reliably distinguish fault from non-fault conditions.

The findings align with prior research demonstrating that sequence-based autoencoders trained on nominal operating conditions can detect anomalous segments without labeled data, while shallow discriminative models perform strongly when features encode operational conditions [5,11,12]. Compared to unsupervised approaches such as one-class SVMs or isolation forests for unsupervised monitoring, the LSTM Autoencoder uses

temporal dependencies and provides a clear anomaly score (MAE), enhancing interpretability [11,12,26]. On the supervised side, performance is comparable to gradient-boosting and random forest baselines reported on similar digital twin benchmarks [41], while offering a more deployment-friendly architecture with reduced computational complexity [10].

It should be noted that the exceptionally high classification performance of the MLP stems from the deterministic nature of the fault-labeling process. Because the labels were generated via fixed operational thresholds applied to vibration, temperature, pressure, motor speed, torque, and humidity, the classifier's task reduces to learning a smooth approximation of these deterministic boundaries. Consequently, the reported metrics should be understood as an indication of rule-consistency rather than evidence that the model can generalize to more complex, noisy, or implicitly defined fault mechanisms.

From an operational perspective, the proposed framework offers practical benefits. The Autoencoder can serve as a front-end monitor to triage time windows for inspection and log correlation, and the MLP provides actionable fault probabilities for maintenance scheduling. Both models are lightweight and suitable for integration into SCADA systems or cloud-based platforms, supporting near real-time monitoring and decision support. Furthermore, the use of a validation-fixed threshold and calibrated probabilities can reduce false alarms, minimize diagnostic workload, and enable timely interventions.

The proposed framework was designed with near real-time industrial deployment in mind. Both the LSTM Autoencoder and the MLP classifier adopt lightweight architectures, consisting of a single recurrent layer and a shallow feedforward network, respectively. This design ensures low inference complexity and minimal memory footprint, making the models suitable for execution in SCADA, edge, or cloud-based environments.

Model training is intended to be performed offline or periodically, while inference can be executed online on streaming data. Although exhaustive runtime benchmarking and detailed memory profiling were beyond the scope of this study, the compact model configurations indicate low-latency inference suitable for near real-time monitoring on standard industrial hardware. Detailed runtime profiling in live streaming environments is identified as future work. In addition, indicative CPU-based execution times measured during MATLAB experimentation (Section 4.4) further support near real-time suitability, with inference executed in the order of milliseconds per time window.

Several limitations should be acknowledged. First, the reported anomaly counts correspond to statistically defined threshold exceedances in the reconstruction error rather than independently verified failure events; consequently, classical false-alarm and miss rates cannot be precisely quantified in the absence of labeled anomaly ground truth. Second, the use of a fixed anomaly threshold defined as $\mu + 3\sigma$, while conceptually sound and widely adopted, may be suboptimal under non-stationary operating conditions, indicating the potential need for adaptive or context-aware thresholding strategies. Third, the supervised fault labels employed for MLP training are derived from rule-based operational criteria rather than confirmed fault annotations, which may lead to optimistic performance estimates compared to noisy real-world deployments. Fourth, the evaluation was conducted in an offline digital twin environment, and therefore factors such as communication latency, sensor drift, and data quality issues inherent to live streaming industrial systems were not explicitly assessed. Finally, although sequence length and feature selection were guided by empirical testing and domain knowledge, a systematic sensitivity analysis of these design choices was beyond the scope of this study and constitutes an important direction for future work [34].

Future research should focus on several promising directions to further enhance and extend this work. Adaptive or dynamic thresholding techniques will be explored to improve stability under changing operating conditions. Development of multivariate and attention-based autoencoders is planned to better capture cross-sensor dependencies and strengthen anomaly detection capabilities. Incorporating real failure events and modeling multiple fault severity levels will increase the practical relevance of the system, while cost-sensitive operating points will be investigated to align predictions with maintenance priorities. Finally, an online architecture integrated with SCADA/PLC systems and real-time dashboards will be prototyped to enable smooth deployment in industrial environments.

7. Conclusions

In this study, we developed and evaluated a two-level predictive maintenance framework for industrial machinery, combining an LSTM autoencoder for unsupervised anomaly detection with an MLP for supervised fault prediction. The system was tested on multivariate time-series data generated from a digital twin environment, simulating five years of industrial operation. The results showed that the LSTM autoencoder was able to identify localized anomalies effectively, while the MLP classifier achieved very high accuracy, precision, and recall in distinguishing between normal and faulty machine states.

The innovation of this work lies in its hybrid approach, which applies both unsupervised and supervised learning to address the challenges of predictive maintenance. By using an LSTM autoencoder, the system can detect subtle deviations from normal behavior without requiring labeled anomaly data, making it suitable for early warning applications. The MLP classifier, on the other hand, provides actionable fault predictions based on operational features, supporting timely and informed maintenance decisions. Together, these models offer a lightweight and interpretable solution that can be integrated into near real-time industrial monitoring systems.

The results provide strong evidence for the effectiveness of the approach. The metrics indicate that the proposed framework can reliably distinguish between normal and faulty operating conditions, maintaining both a low false-alarm rate and a negligible probability of missing true faults. These findings confirm that deep learning models, particularly those using temporal dependencies and feature-based classification, are highly effective for predictive maintenance. The hybrid architecture introduced here advances the field by combining the strengths of unsupervised anomaly detection and supervised fault prediction in a lightweight, interpretable, and deployable solution.

In summary, this work demonstrates that a hybrid deep learning approach can provide a robust and practical solution for predictive maintenance in industrial machinery. The results offer a strong foundation for further research and real-world implementation, with the potential to improve maintenance strategies and operational efficiency across a wide range of industrial applications.

Author Contributions: Conceptualization, I.M. and P.Z.; methodology, I.M. and P.Z.; validation, I.M. and P.Z.; formal analysis, I.M. and P.Z.; investigation, I.M., A.C. and P.Z.; data curation, I.M.; writing—original draft preparation, I.M., A.C. and P.Z.; writing—review and editing, A.C. and P.Z.; supervision, P.Z. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The data presented in this study are available on request from the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

LSTM	Long Short-Term Memory
AE	Autoencoder
MLP	Multilayer Perceptron
DT	Digital Twin

References

1. Mołęda, M.; Małyśiak-Mrozek, B.; Ding, W.; Sunderam, V.; Mrozek, D. From Corrective to Predictive Maintenance—A Review of Maintenance Approaches for the Power Industry. *Sensors* **2023**, *23*, 5970. [CrossRef]
2. Stenström, C.; Norrbin, P.; Parida, A.; Kumar, U. Preventive and corrective maintenance—Cost comparison and cost-benefit analysis. *Struct. Infrastruct. Eng.* **2016**, *12*, 603–617. [CrossRef]
3. Carvalho, T.P.; Soares, F.A.A.M.N.; Vita, R.; Francisco, R.d.P.; Basto, J.P.; Alcalá, S.G.S. A systematic literature review of machine learning methods applied to predictive maintenance. *Comput. Ind. Eng.* **2019**, *137*, 106024. [CrossRef]
4. Jardine, A.K.S.; Lin, D.; Banjevic, D. A review on machinery diagnostics and prognostics implementing condition-based maintenance. *Mech. Syst. Signal Process.* **2006**, *20*, 1483–1510. [CrossRef]
5. Zhao, R.; Yan, R.; Chen, Z.; Mao, K.; Wang, P.; Gao, R.X. Deep learning and its applications to machine health monitoring. *Mech. Syst. Signal Process.* **2019**, *115*, 213–237. [CrossRef]
6. Goodfellow, I.; Bengio, Y.; Courville, A. *Deep Learning*; MIT Press: Cambridge, MA, USA, 2016; pp. 432–455.
7. Ghelani, D. Harnessing machine learning for predictive maintenance in energy infrastructure: A review of challenges and solutions. *Int. J. Sci. Res. Arch.* **2024**, *12*, 1138–1156. [CrossRef]
8. Kharazian, Z.; Lindgren, T. Challenges in Industrial Data Access in Predictive Maintenance. *IEEE Reliab. Mag.* **2025**, *2*, 18–22. [CrossRef]
9. Leoni, L.; Peron, M.; De Carlo, F. Machine learning-based predictive maintenance: Empirical insights of challenges and countermeasures. *Prod. Plan. Control* **2025**, *36*, 2399–2428. [CrossRef]
10. Ribeiro, M.T.; Singh, S.; Guestrin, C. Why should I trust you? Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*; Association for Computing Machinery: New York, NY, USA, 2016; pp. 1135–1144. [CrossRef]
11. Pang, G.; Shen, C.; Cao, L.; Hengel, A.V.D. Deep learning for anomaly detection: A review. *ACM Comput. Surv. (CSUR)* **2021**, *54*, 38. [CrossRef]
12. Maleki, S.; Maleki, S. Unsupervised anomaly detection with LSTM autoencoders using statistical data-filtering. *Appl. Soft Comput.* **2021**, *107*, 107443. [CrossRef]
13. Hochreiter, S.; Schmidhuber, J. Long short-term memory. *Neural Comput.* **1997**, *9*, 1735–1780. [CrossRef] [PubMed]
14. Gherghina, I.-S.; Bizon, N.; Iana, G.-V.; Vasilică, B.-V. Recent Advances in Fault Detection and Analysis of Synchronous Motors: A Review. *Machines* **2025**, *13*, 815. [CrossRef]
15. Ding, S.X. *Model-Based Fault Diagnosis Techniques: Design Schemes, Algorithms, and Tools*; Advances in Industrial Control; Springer: London, UK, 2013.
16. Lunze, J.; Pröll, S.; Jarmolowitz, F. From structural analysis to observer-based residual generation for fault detection. In *Proceedings of the 3rd Conference on Control and Fault-Tolerant Systems (SysTol)*, Barcelona, Spain, 7–9 September 2016; pp. 491–498. [CrossRef]
17. Sultana, R.-G.; Davrados, A.; Dimogianopoulos, D. Evaluating Contact-Less Sensing and Fault Diagnosis Characteristics in Vibrating Thin Cantilever Beams with a MetGlas® 2826 MB Ribbon. *Vibration* **2024**, *7*, 36–52. [CrossRef]
18. Cai, D. Enhancing capabilities of generative models through VAE-GAN integration: A review. *Appl. Comput. Eng.* **2024**, *67*, 242–249. [CrossRef]
19. Ali, M.; Udoidiok, I.; Li, F.; Zhang, J. *A Review on Generative Intelligence in Deep Learning based Network Intrusion Detection*; IEEE: New York, NY, USA, 2024; pp. 1–9. [CrossRef]
20. Mahmoudi, I.; Boubiche, D.E.; Athmani, S.; Toral-Cruz, H.; Chan-Puc, F.I. Toward Generative AI-Based Intrusion Detection Systems for the Internet of Vehicles (IoV). *Future Internet* **2025**, *17*, 310. [CrossRef]
21. Vlachou, V.I.; Karakatsanis, T.S.; Efstathiou, D.E.; Vlachou, E.I.; Vologiannidis, S.D.; Balaska, V.E.; Gasteratos, A.C. Condition Monitoring and Fault Prediction in PMSM Drives Using Machine Learning for Elevator Applications. *Machines* **2025**, *13*, 549. [CrossRef]
22. Ahmed, I.; Ahmad, M.; Chehri, A.; Jeon, G. A Smart-Anomaly-Detection System for Industrial Machines Based on Feature Autoencoder and Deep Learning. *Micromachines* **2023**, *14*, 154. [CrossRef]

23. Do, J.S.; Kareem, A.B.; Hur, J.-W. LSTM-Autoencoder for Vibration Anomaly Detection in Vertical Carousel Storage and Retrieval System (VCSRS). *Sensors* **2023**, *23*, 1009. [CrossRef]
24. Apeiranthitis, S.; Zacharia, P.; Chatzopoulos, A.; Papoutsidakis, M. Predictive Maintenance of Machinery with Rotating Parts Using Convolutional Neural Networks. *Electronics* **2024**, *13*, 460. [CrossRef]
25. Sajjadi, P.; Dinmohammadi, F.; Shafiee, M. Machine Learning in Prognostics and Health Management of Cyber-Physical Systems: A Review. *IEEE Access* **2025**, *13*, 162320–162354. [CrossRef]
26. Ahmed, M.; Mahmood, A.N.; Hu, J. A survey of network anomaly detection techniques. *J. Netw. Comput. Appl.* **2016**, *60*, 19–31. [CrossRef]
27. Jin, J.; Chung, Y.; Park, J. Development of a Flowmeter Using Vibration Interaction between Gauge Plate and External Flow Analyzed by LSTM. *Sensors* **2020**, *20*, 5922. [CrossRef] [PubMed]
28. Lee, J.-H.; Hong, J.-K. Comparative Performance Analysis of Vibration Prediction Using RNN Techniques. *Electronics* **2022**, *11*, 3619. [CrossRef]
29. Huang, J.; Mo, J.; Zhang, J.; Ma, X. A Fiber Vibration Signal Recognition Method Based on CNN-CBAM-LSTM. *Appl. Sci.* **2022**, *12*, 8478. [CrossRef]
30. Gu, K.; Zhang, Y.; Liu, X.; Li, H.; Ren, M. DWT-LSTM-Based Fault Diagnosis of Rolling Bearings with Multi-Sensors. *Electronics* **2021**, *10*, 2076. [CrossRef]
31. Zaznov, I.; Badii, A.; Dufour, A.; Kunkel, J. AdamZ: An Enhanced Optimisation Method for Neural Network Training. *arXiv* **2024**. [CrossRef]
32. Bhattacharjee, A.; Popov, A.A.; Sarshar, A.; Sandu, A. Improving the Adaptive Moment Estimation (ADAM) stochastic optimizer through an Implicit-Explicit (IMEX) time-stepping approach. *arXiv* **2024**, arXiv:2403.13704. [CrossRef]
33. Bisong, E. *The Multilayer Perceptron (MLP)*; Apress: Berkeley, CA, USA, 2019; pp. 401–405. [CrossRef]
34. Malhi, A.; Gao, R.X. PCA-based feature selection scheme for machine defect classification. *IEEE Trans. Instrum. Meas.* **2004**, *53*, 1517–1525. [CrossRef]
35. Bergou, E.H.; Bergou, E.H.; Diouane, Y.; Kungurtsev, V. Convergence and Complexity Analysis of a Levenberg–Marquardt Algorithm for Inverse Problems. *J. Optim. Theory Appl.* **2020**, *185*, 927–944. [CrossRef]
36. Zhang, X.; Zhao, J. A Levenberg–Marquardt Method for Tensor Approximation. *Symmetry* **2023**, *15*, 694. [CrossRef]
37. Guo, X.; Yin, Y.; Dong, C.; Yang, G.; Zhou, G. On the class imbalance problem. In *Proceedings of the Fourth International Conference on Natural Computation, Jinan, China, 18–20 October 2008*; IEEE: New York, NY, USA, 2008; pp. 192–201. [CrossRef]
38. Zhang, W.; Yang, D.; Wang, H. Data preprocessing methods in fault diagnosis: A review. *IEEE Access* **2019**, *7*, 86327–86343. [CrossRef]
39. Deng, L.; Yu, D. Deep learning: Methods and applications. *Found. Trends Signal Process.* **2014**, *7*, 197–387. [CrossRef]
40. Chollet, F. *Deep Learning with Python*, 2nd ed.; Manning Publications: Shelter Island, NY, USA, 2021; pp. 251–274.
41. Zhao, R.; Yan, R.; Wang, J.; Mao, K.; Shen, F.; Wang, X. Machine health monitoring using local feature-based gated recurrent unit networks. *IEEE Trans. Ind. Electron.* **2019**, *66*, 530–539. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Article

An Optimized Few-Shot Learning Framework for Fault Diagnosis in Milling Machines

Faisal Saleem ¹, Muhammad Umar ¹ and Jong-Myon Kim ^{1,2,*}

¹ Department of Electrical, Electronics and Computer Engineering, University of Ulsan, Ulsan 44610, Republic of Korea; faisal1999@mail.ulsan.ac.kr (F.S.); muhammadumar@mail.ulsan.ac.kr (M.U.)

² PD Technology Co., Ltd., Ulsan 44610, Republic of Korea

* Correspondence: jmkim07@ulsan.ac.kr; Tel.: +82-52-259-2217

Abstract

Reliable fault diagnosis of milling machines is essential for maintaining operational stability and cost-effective maintenance; however, it remains challenging due to limited labeled data and the highly non-stationary nature of acoustic emission (AE) signals. This study introduces an optimized Few-Shot Learning framework (FSL) that integrates time–frequency analysis with attention-guided representation learning and distribution-aware classification for data-efficient fault detection. The framework converts AE signals into Continuous Wavelet Transform (CWT) scalograms, which are processed using a self-attention-enhanced ResNet-50 backbone to capture both local texture features and long-range dependencies in the signal. Adaptive prototype computation with learnable importance weighting refines class representations, while Mahalanobis distance-based matching ensures robust alignment between query and prototype embeddings under limited sample conditions. To further strengthen discriminability, contrastive loss with hard negative mining enforces compact intra-class clustering and clear inter-class separation. Comprehensive experiments under 7-way 5-shot settings and 5-fold stratified cross-validation demonstrate consistent and reliable performance, achieving a mean accuracy of $98.86\% \pm 0.97\%$ (95% CI: [98.01%, 99.71%]). Additional evaluations across multiple spindle speeds (660 rpm and 1440 rpm) confirm that the model generalizes effectively under varying operating conditions. Grad-CAM++ activation maps further illustrate that the network focuses on physically meaningful fault-related regions, enhancing interpretability. The results verify that the proposed framework achieves robust, scalable, and interpretable fault diagnosis using minimal labeled data, offering a practical solution for predictive maintenance in modern intelligent manufacturing environments.

Keywords: fault diagnosis; milling machines; few-shot learning; acoustic emission signals; self-attention mechanism; adaptive prototype learning; predictive maintenance

1. Introduction

Milling is one of the most widely used and important machining processes in modern industries, including aerospace, automotive, and medical device manufacturing [1,2]. It uses a rotating cutting tool to remove material from a workpiece and can create flat surfaces, grooves, threads, and complex shapes with high precision [3]. Due to its high efficiency and flexibility, milling plays an important role in both small workshops and large factories [4]. However, during milling operations, tools are exposed to high forces, frictional heat, and constant vibrations [5]. These conditions eventually cause the tool to wear out or, in some

cases, break suddenly [6]. Tool wear not only reduces product quality but also causes surface roughness, dimensional errors, and unplanned machine stoppages, which increase production costs. Studies show that tool-related failures are responsible for 7–20% of total machine downtime, and tool costs can contribute up to 12% of the overall manufacturing budget [7,8]. In most factories, tools are replaced based on fixed time intervals or operator experience [9]. This approach is not always reliable. Sometimes tools are changed too early, wasting usable life, and other times too late, leading to failure. On average, only 50–80% of a tool's actual life is utilized [10,11]. To solve this problem, Tool Condition Monitoring (TCM) systems are being used to monitor tool health in real time and provide early warnings before failure occurs [12,13]. There are two main types of TCM methods: direct and indirect. Direct methods involve using cameras or microscopes to inspect the tool's surface, but they are expensive and difficult to use in real production settings [14,15]. Indirect methods are more practical and are widely used in the industry [16]. They rely on signals such as vibration, cutting forces, spindle current, and acoustic emission (AE) [17]; among these, AE signals are particularly popular [18]. AE signals are high-frequency sound waves produced when cracks form, friction occurs, or deformation happens during cutting, making them very useful for detecting early tool damage [19]. However, AE signals are complex and often affected by changes in cutting speed, tool material, and machine settings, making them difficult to analyze [20]. To process these signals, researchers use techniques such as time-domain statistics, frequency analysis, wavelet transforms, or empirical mode decomposition to extract features [21]. These features are then classified using machine learning models such as Support Vector Machines, Hidden Markov Models, and Artificial Neural Networks [22,23]. While effective, these methods depend heavily on manual feature extraction and do not perform well when the amount of labeled data is small or when the operating conditions vary significantly [24].

In recent years, deep learning has helped improve fault diagnosis in TCM [25,26]. Models such as Convolutional Neural Networks (CNNs) [27], Recurrent Neural Networks, and Bidirectional Long Short-Term Memory networks can automatically learn patterns from raw data [28,29]. A popular method is to convert AE signals into 2D images called scalograms using the Continuous Wavelet Transform (CWT). These scalograms show both time and frequency information, making it easier for deep learning models to detect tool defect patterns [30]. Deep learning has also been used with traditional physical models and sensor fusion to improve performance [31,32]. However, deep learning requires a large amount of labeled data to work effectively [33]. In real-world applications, especially for rare fault types, collecting large datasets is difficult and expensive [34]. When data is limited, deep models often become unreliable and may not generalize well [35]. Deep learning can find useful patterns on its own, but it needs a lot of labeled data. In the circumstances where labeled data are scarce, Few-Shot Learning (FSL) becomes a better option [36]. FSL is a type of learning that allows models to recognize new fault types using only a few labeled instances. For instance, Li et al. [37] proposed a method that improves fault classification by separating different classes and keeping samples from the same class close together. Liang et al. [38] reviewed different FSL techniques and found that meta-learning and attention mechanisms are especially useful for fault diagnosis. Wang et al. [39] introduced a meta-learning model that can adapt to new working conditions even with limited training data. More recently, researchers have developed a model that combines prototype refinement with contrastive learning, helping the system distinguish between very similar faults [40]. These methods have shown great potential, but many of them still require complicated training procedures, high computing resources, and often do not fully utilize attention mechanisms or adaptive class representations [41]. Additionally, very few of these studies focus on milling tools using AE signals [42].

To overcome these challenges, this study proposes an improved Few-Shot Learning framework for diagnosing faults in milling cutting tools using AE signals. In the current method, raw AE signals are transformed into two-dimensional scalograms using the CWT to capture time–frequency characteristics. A deep learning model based on ResNet-50, augmented with spatial and channel-wise attention mechanisms, is then utilized to emphasize the most informative regions. To enhance fault type discrimination with limited samples, class prototypes are computed using a learnable weighting technique, enabling more accurate class representation. For classification, the Mahalanobis distance is utilized to account for feature distribution. The framework is trained under an episodic learning setup and evaluated on real AE data collected during milling operations. The results show that the proposed method achieves high accuracy, learns better fault representations, and generalizes well across different conditions, even with very little data.

To better emphasize the novelty and advantages of the proposed framework, a comparative summary of representative recent studies is provided in Table 1. The proposed approach introduces four main innovations: (i) a learnable importance weighting mechanism for adaptive prototype computation; (ii) spatial and channel-wise attention modules integrated within a ResNet-50 backbone to enhance time–frequency feature extraction; (iii) Mahalanobis distance-based prototype-query matching to account for feature distribution; and (iv) contrastive loss with hard negative mining for improved inter-class separation. In comparison to conventional CNNs and LSTMs that require large, labeled datasets, the proposed model achieves high classification accuracy with only five samples per class. Compared with previous FSL learning approaches, the proposed method demonstrates higher robustness and data efficiency, achieving an average accuracy of $98.86\% \pm 0.97\%$, outperforming other approaches. These findings confirm that the proposed method offers a reliable, data-efficient, and generalizable solution for intelligent fault diagnosis in milling operations.

Table 1. Comparative Summary of the Proposed Framework and Existing Methods.

Method	Signal Type	Approach/Features	Limitations
CNN-based classifier [43]	AE/Vibration	2D CWT + CNN	Requires large, labeled data; prone to overfitting
LSTM-based model [44]	AE	Sequential Modeling	Inefficient for frequency-domain data
Siamese FSL model [45]	AE	Pairwise distance learning	Sensitive to class imbalance
Prototype-based FSL [46]	AE	Standard prototypical network	No attention; fixed prototypes
Proposed Framework	AE (CWT Scalogram)	Attention-based ResNet-50 + Learnable prototype weighting + Mahalanobis metric + Contrastive loss	Domain Adaptation

The main contributions in this work are as follows.

1. A learnable importance weighting mechanism is introduced for prototype computation, selectively emphasizing discriminative feature dimensions and enhancing class representation in few-shot fault diagnosis.
2. Spatial and channel-wise self-attention modules are integrated into a ResNet-50 backbone, enabling more expressive representation of 2D CWT scalograms by capturing both local textures and global dependencies in fault-related patterns.

3. Mahalanobis distance with regularized class-wise covariance is employed for prototype-query matching, providing robust similarity measurement under limited-data conditions.
4. Contrastive loss with hard negative mining is adopted to enforce intra-class compactness and inter-class separability, thereby improving feature clustering and overall generalization.

The remainder of this paper is organized as follows: Section 2 presents the proposed methodology, Section 3 outlines the technical background, Section 4 details the experimental setup, Section 5 reports the results, Section 6 provides an in-depth discussion, and Section 7 concludes the paper.

2. Proposed Methodology

This section presents the proposed FSL framework for milling machine fault diagnosis, which integrates adaptive prototype computation, self-attention mechanisms, and contrastive loss to achieve robust classification under limited data conditions. The methodology is organized into five major stages: data preprocessing, feature extraction, prototype learning, similarity-based classification, and episodic training. The overall flowchart of the proposed method is illustrated in Figure 1:

Step I: The dataset used in this study was collected from a milling machine testbed equipped with AE sensors, and signals were acquired under various fault conditions. The system records multi-channel AE signals during machining operations, capturing high-frequency signal patterns. CWT is applied to generate 2D scalograms to capture time–frequency patterns, extracting rich spectral features. Each sample is resized to $224 \times 224 \times 3$ pixels and normalized using min-max scaling to enhance feature representation. Data augmentation techniques, like random cropping, rotation, and small translations, are applied to the scalograms to improve generalization. The dataset is then split into a support set and a query set using a stratified episodic sampling technique to ensure balanced class representation.


Step II: In the second stage, feature extraction is performed using a modified ResNet-50 architecture augmented with attention mechanisms, referred to as Self-Attention ResNet-50 (SA-ResNet50). This backbone incorporates spatial attention to capture local and global contextual dependencies and channel-wise attention to highlight the most discriminative feature maps. This dual-attention design enhances the network’s ability to isolate and represent subtle fault characteristics in the CWT scalograms. The resulting feature embeddings $f \in R^{512}$ are passed to the prototype learning module.

Step III: The third stage uses an adaptive prototype learning strategy that enhances class representation. Each support sample is encoded with the SA-ResNet50 network to obtain its embedding. For each class, the mean vector of the support embeddings is computed and normalized. The learnable importance of the weight vector w is then applied to reweight the feature dimensions of the prototype, emphasizing class-specific discriminative information. This weight vector is optimized jointly with the model parameters through backpropagation, resulting in task-adaptive and expressive prototypes that better capture intra-class characteristics.


Step IV: In the fourth stage, classification of query samples is performed by comparing them to the refined prototypes using the Mahalanobis distance, which incorporates the class-wise covariance matrix to provide more reliable similarity estimates in low-sample conditions. To further improve feature discrimination, a pairwise contrastive loss with hard negative mining is applied. This loss encourages intra-class compactness and inter-class separation in the embedding space, refining the model’s decision boundaries.

Step V: The final stage involves episodic training using a 7-way, 5-shot setting. In each training episode, new support and query sets are dynamically sampled. The model is optimized using the Adam optimizer with an initial learning rate of 10^{-3} , and gradient clipping with a maximum norm of 1 is applied to prevent exploding gradients. Model performance is evaluated using standard classification metrics, including accuracy, precision, recall, and F1-score. To visually assess the clustering quality of the learned representations, t-SNE is used to project high-dimensional embeddings into a two-dimensional space, demonstrating the model’s ability to form well-separated and compact clusters corresponding to different fault classes.

1 Experimental Setup and Data Preparation



2 Embedding and Prototype Formation Pipeline



3 Query Classification and Evaluation




Figure 1. Flowchart of the proposed FSL framework for milling machine fault diagnosis.

3. Technical Background

3.1. Acoustic Emission

AE is a passive sensing technique used to monitor the health of mechanical systems. It captures high-frequency stress waves generated by defects such as cracks, friction, and material deformation. It provides early detection of faults by sensing microscopic energy releases before visible damage occurs. In milling machines, AE signals are generated during tool-workpiece interactions. These signals contain rich fault-related information, but they are often buried in noise. The high sampling rate of AE ensures detailed fault representation; however, it also increases data complexity. Extracting meaningful patterns from AE signals requires advanced signal processing techniques. AE signals are non-stationary, meaning their frequency content changes over time. This requires time–frequency analysis techniques, which provide a more detailed representation of AE signals for fault diagnosis:

3.2. Continuous Wavelet Transform

CWT is a powerful technique for analyzing non-stationary signals, such as AE. Unlike other transformations like FFT, which provide only frequency domain information, CWT simultaneously captures both time and frequency characteristics. This makes it ideal for fault diagnosis, where AE signals exhibit time-varying spectral patterns. CWT decomposes a signal $x(t)$ into scaled and translated versions of a mother wavelet $\psi(t)$ providing localized frequency analysis [47]. The transformation is defined by

$$W(a, b) = \int_{-\infty}^{\infty} x(t) \frac{1}{\sqrt{|a|}} \psi\left(\frac{t-b}{a}\right) dt \quad (1)$$

where Equation (1), $W(a, b)$ represents the wavelet coefficients at scale a and position b , $x(t)$ is the input AE signal, $\psi(t)$ is the mother wavelet, which determines the shape of the decomposition, a is the scale parameter, controlling frequency resolution, b is the translation parameter, shifting the wavelet along the time axis. By adjusting a and b , CWT extracts features across multiple frequency bands, allowing the detection of subtle changes in AE signals. This helps in detecting sudden changes caused by tool wear, cracks, or mechanical failures. These 2D images highlight frequency variations over time, making them suitable for deep learning models. This study enhances feature extraction by transforming AE signals into scalograms, enabling more accurate fault identification.

3.3. Residual Network-50

ResNet-50 is a deep CNN that has been widely used for feature extraction due to its ability to capture complex hierarchical representations while reducing the vanishing gradient problem [48]. In the current study, ResNet-50 is employed as the backbone for extracting discriminative features from scalograms, enabling efficient fault diagnosis in milling machines. The architecture is designed with residual connections, allowing the network to learn identity mappings that facilitate deeper architecture without degradation, as shown in Figure 2.




Figure 2. Residual block structure, where the residual function is added to the identity mapping through skip connections.

The core concept is to reformulate the mapping from the input x to the output $H(x)$ as a residual function:

$$H(x) = F(x) + x \quad (2)$$

where in Equation (2), $H(x)$ is the desired transformation, $F(x)$ is the residual function learned by convolutional layers, x is the identity mapping. By allowing gradient flow through skip connections, it ensures stable training even for deep networks. This structure significantly enhances feature extraction capabilities, making it well-suited for handling complex patterns. In the proposed method, ResNet-50 processes the scalograms to extract meaningful spatial and frequency-based features, as illustrated in Figure 3. The model is robust to noise, extracting stable patterns from noisy signals, which improves classification performance. Pretrained on large datasets, it provides a strong initialization for FSL, resulting in better generalization and adaptability.




Figure 3. Architecture of the ResNet-50 backbone used for feature extraction.

3.4. Self-Attention Mechanism

In the proposed method, a self-attention mechanism is integrated into the ResNet-50 backbone to enhance feature extraction from scalograms. Traditional CNNs primarily focus on local spatial patterns, limiting their ability to capture long-range dependencies. Self-attention mechanisms address this limitation by assigning different importance levels to different regions of the input feature maps, allowing the model to focus on the most relevant patterns for fault classification [49], as shown in Figure 4.




Figure 4. Self-attention module for feature refinement in fault classification.

SA computes relationships between all spatial locations within a feature map, forming a weighted representation that emphasizes significant regions while suppressing less informative ones. Given an input feature map X , the self-attention mechanism generates three transformed versions: the query Q , key K , and value V , computed as

$$Q = W_q X, K = W_k X, V = W_v X \quad (3)$$

where in Equation (3), W_q, W_k, W_v are learnable weight matrices. The attention scores are obtained by computing the scaled dot-product similarity between Q and K . Where in Equation (4), d_k is the dimensionality of the key vectors, while the final attention output is computed by Equation (5):

$$A = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right) \quad (4)$$

$$Z = AV \quad (5)$$

By incorporating self-attention, the network effectively captures global contextual information, improving feature discrimination in scalograms. This is particularly beneficial for fault diagnosis, where signal patterns exhibit complex variations across time and frequency. The integration of self-attention enhances the ability of the model to highlight discriminative regions, improving class separation in the feature space. In the proposed method, SA-ResNet50, SA is applied at multiple layers to refine feature representations progressively. This enables the model to focus on small differences. The ability to adaptively weigh feature importance makes self-attention particularly effective in FSL, where limited support samples require efficient feature extraction to generalize well to new queries. The overall architecture of the proposed SA-Resnet50 is provided in Table 2.

Table 2. Network Architecture of the Proposed SA-Resnet50.

Stage	Layer/Block	Output Size
Input	CWT Scalogram	$224 \times 224 \times 3$
Conv1	7×7 Conv, 64 filters, stride 2 + MaxPool	$112 \times 112 \times 64$
ResNet Stage 1	$[1 \times 1 \text{ Conv}, 64] + [3 \times 3 \text{ Conv}, 64] \times 3$	$56 \times 56 \times 256$
ResNet Stage 2	$[1 \times 1 \text{ Conv}, 128] + [3 \times 3 \text{ Conv}, 128] \times 4$	$28 \times 28 \times 512$

Table 2. Cont.

Stage	Layer/Block	Output Size
ResNet Stage 3	$[1 \times 1 \text{ Conv}, 256] + [3 \times 3 \text{ Conv}, 256] \times 6$	$14 \times 14 \times 1024$
ResNet Stage 4	$[1 \times 1 \text{ Conv}, 512] + [3 \times 3 \text{ Conv}, 512] \times 3$	$7 \times 7 \times 2048$
Attention Module	Spatial + Channel Self-Attention	$7 \times 7 \times 2048$
Global Pooling	Average Pooling	$1 \times 1 \times 204$
Embedding Layer	Fully Connected	1×512
Prototype Computation	APL	1×512
Similarity Measure	Mahalanobis Distance	Scalar
Loss Function	Contrastive Loss + Cross-Entropy	-

3.5. Adaptive Prototype Learning

Prototype-based learning is a widely used approach in FSL, where a classifier assigns labels to query samples based on their similarity to class prototypes. A prototype is a representative feature vector computed from the support set of each class. Conventional prototypical networks compute these prototypes as the mean of all support samples belonging to a given class [40]. However, this simple averaging approach may not always capture the most informative class representations, especially in cases of high intra-class variability or noisy samples. To address this limitation, Adaptive Prototype Learning (APL) is introduced in this method as shown in Figure 5.




Figure 5. APL framework, where learnable weights adjust the support samples to generate more representative class prototypes.

APL assigns learnable importance weights to support samples, ensuring that more relevant samples contribute more to prototype computation [50]. Given a set of support embeddings $S = \{s_1, s_2, \dots, s_N\}$ for a class c , the traditional prototype computation is given by

$$\mu_c = \frac{1}{N} \sum_{i=1}^N s_i \quad (6)$$

where in Equation (6), μ_c is the prototype of the class c and N is the number of support samples. In APL, a learnable weight vector $\omega = \{\omega_1, \omega_2, \dots, \omega_N\}$ is introduced by

Equation (7), where each weight determines the contribution of a specific support sample to the prototype:

$$\mu_c = \frac{1}{N} \sum_{i=1}^N \omega_i s_i, \text{ where } \sum_{i=1}^N \omega_i = 1 \quad (7)$$

These weights are learned during training, allowing the model to emphasize discriminative support samples while reducing the influence of less informative ones. This adaptive weighting leads to better generalization, as the prototypes dynamically adjust based on the most relevant class representations. The benefits of APL in FSL are significant. By assigning adaptive weights, the computed prototypes become more representative of class distributions, improving prototype robustness. It effectively reduces the influence of noisy or misrepresentative support samples, ensuring better handling of noisy data. Additionally, by dynamically refining prototypes based on the most relevant features, it enhances generalization. This adaptability makes APL particularly effective in fault diagnosis operating conditions where data availability is limited, ensuring more accurate and reliable predictions.

3.6. Mahalanobis Distance

Mahalanobis distance is a powerful metric that accounts for feature correlations and varying data distributions, making it well-suited for prototype-based classification. It considers the covariance structure of the feature space, allowing for more discriminative comparisons between query samples and class prototypes [45,51]. In the current work, the Mahalanobis distance is employed to measure the similarity between query embeddings and adaptive prototypes as depicted in Figure 6. This approach ensures that feature variations are considered appropriately, reducing the influence of redundant or less relevant dimensions. Given a query feature vector x and a class prototype μ , the Mahalanobis distance is defined by

$$D_M(x, \mu) = \sqrt{(x - \mu)^T S^{-1} (x - \mu)} \quad (8)$$

where in Equation (8), S is the covariance matrix of the support set features. The covariance matrix is regularized using a pseudo-inverse to prevent instability when the support set is small. By incorporating the Mahalanobis distance, the proposed method improves class separation, ensuring more reliable classification of fault conditions in milling machines.




Figure 6. Query samples are matched to class prototypes using covariance-aware similarity for improved class separation.

4. Experimental Setup for Fault Diagnosis in Milling Cutting Tools

4.1. Introduction to the Experimental Setup

A controlled experimental setup was designed to study fault diagnosis in milling cutting tools using AE signals. This setup aimed to capture and analyze the distinct AE patterns generated during milling operations under different fault conditions. The experiment was performed on an INTER-SIEG X1 Micro Mill Drill (INTER-SIEG, Bremen, Germany), a compact milling machine constructed from cast iron, as shown in Figure 7.




Figure 7. Experimental setup of a milling machine.

This material composition provided better stability and reduced unwanted vibrations during the machining process, ensuring that only relevant AE signals were captured. The INTER-SIEG X1 Micro Mill Drill was selected for this study due to its precision in performing micro-milling tasks and its ability to simulate real-world machining conditions. The primary machining operation employed in this study was straight parallel milling, a widely used technique for shaping and cutting hard materials with high precision. This process was conducted on steel workpieces, which were chosen due to their durability and wide industrial applicability. The workpieces used in the milling process were prepared to standardized dimensions of 20 mm × 35 mm × 35 mm, ensuring uniform material interaction and consistent cutting forces throughout the experiment, as shown in Figure 8. These dimensions were carefully selected to provide a balanced contact area between the cutting tool and the workpiece, reducing variability in signal generation. The cutting depth was fixed at 2 mm to ensure stable material removal and consistent AE generation across all experiments. The machining operations were performed under strictly controlled conditions, maintaining precise control over all cutting parameters to ensure repeatability and accurate data collection. The machining operations were performed under strictly controlled conditions, maintaining precise control over all cutting parameters to ensure repeatability and accurate data collection. The cutting tool used in the milling operation

was a two-flute carbide end mill, which was deliberately subjected to controlled wear. The tool was worn to an average flank wear of 0.3 mm, following the ISO-8688-2 standard, which defines the allowable wear limits before tool replacement. The intentional wear of the tool was an essential aspect of this study, as tool degradation directly affects machining performance and alters the characteristics of AE signals. By examining the AE signal variations corresponding to tool wear, it is possible to extract meaningful features indicative of cutting tool degradation.




Figure 8. Workpieces used in milling experiments: (a) unprocessed samples and (b) processed workpiece after machining.

To further ensure consistency in the machining process, the rotational speed of the motor was maintained at 1320 revolutions per minute, which corresponds to an operating frequency of 22 Hz. The spindle was rotated at 660 RPM, translating to 11 Hz. These parameters were selected based on standard milling practices, ensuring that the forces acting on the cutting tool and workpiece remained stable throughout the operation. Additionally, the bed feed rate was set to 0.4 mm per second, which allowed for controlled and uniform material removal. This slow and steady feed rate ensured that the cutting process was stable and prevented excessive tool vibration, which could introduce unwanted noise into the AE signals. The overall parameters of the experiment are mentioned in Table 3.

Table 3. Experimental Setup and Operating Conditions.

Parameters	Specifications
Milling Machine	INTER-SIEG X1 Micro Mill Drill
Spindle Rotation	660 RPM (≈ 11 Hz)
Motor Rotation	1320 RPM (≈ 22 Hz)
Cutting depth	2 mm
Rate Bed Feed	0.4 mm/s (under cutting conditions)
Sensors Used	Motor: 100 mV/g, Chuck: 500 mV/g
Workpiece Size	20 mm $\times$ 35 mm $\times$ 35 mm
Operating Modes	Idle state and material cutting
Fault Categories	Tool Fault (1–2 mm blade breakage), Bearing Fault (3 mm depth), Gear Fault (2–3 mm depth, 1.5 mm width)

4.2. Acoustic Emission Sensor Deployment and Data Collection System

AE signals were recorded using two highly sensitive AE sensors (model R15I-AST, MISTRAS Inc., Princeton Junction, NJ, USA), which were strategically placed on the milling machine to capture transient stress waves generated during cutting and fault occurrences. The primary AE sensor was mounted on the spindle, allowing it to capture signals associated with tool wear, bearing conditions, and gear interactions. Since the spindle is directly involved in the machining process, this sensor placement ensured that all AE bursts generated by the cutting tool were recorded with high precision. The secondary AE sensor was affixed to the motor housing, acting as a guard transducer. This secondary sensor plays an important role in detecting non-relevant background noise and vibrations coming from the machine frame and motor, ensuring that the primary sensor focuses on capturing fault-related AE signatures. To achieve a strong and stable connection, both AE sensors were attached using industrial-grade adhesives, which provided a secure bond and prevented sensor displacement during high-speed operations. The sensor attachments were tested using an HSU-Nelson test, a standardized validation technique used to assess the sensitivity and performance of AE transducers. This test involved generating artificial AE bursts to ensure that both sensors were functioning correctly and capable of detecting transient acoustic signals with high accuracy.

Before initiating the cutting experiments, both AE sensors were calibrated in accordance with ASTM E976 using the Hsu–Nielsen pencil-lead break technique to ensure accurate and repeatable measurements. The calibration was performed by breaking a 0.5 mm 2H pencil lead (2–3 mm length) at multiple positions on the spindle and motor housing to generate controlled AE events. Each sensor's sensitivity and timing were verified under identical preamplifier and filter settings (40 dB gain, 50–800 kHz band). The system noise floor was approximately 0.01 V RMS, with a detection threshold of 0.35 V (7% FS), while the average calibration peak voltage measured at the ADC input was 2.2 V (44% FS). The percentage of full scale was calculated using

$$\% FS = \frac{V_{peak}}{V_{FS}} \times 100 \quad (9)$$

where in Equation (9), $V_{FS} = 5$ V for the NI-9223 acquisition module. These results confirmed that all AE signals remained within the linear response range of the digitizer. Furthermore, inter-channel synchronization accuracy was within ± 2 μ s, confirming high repeatability and proper timing alignment between sensors.

The AE signals were acquired using a National Instruments NI-9223 data acquisition system (National Instruments, Austin, TX, USA), which was selected due to its high sampling rate and precision. The sampling frequency was set to 1 MHz based on the spectral characteristics of AE signals and the operating bandwidth of the MISTRAS R15I-AST sensors (50–400 kHz). In milling operations, transient AE bursts produced during tool-workpiece interaction, wear progression, and fracture propagation typically contain frequency components extending up to approximately 450 kHz. According to the Nyquist sampling theorem, the minimum required rate for accurate signal reconstruction should be at least twice the highest signal frequency. Therefore, a 1 MHz sampling rate was chosen to ensure complete capture of all relevant fault-induced acoustic activity while avoiding aliasing. This selection also provides sufficient oversampling for precise transient analysis and reliable feature extraction. Preliminary tests at lower sampling frequencies (500–750 kHz) resulted in partial loss of high-frequency detail and minor amplitude distortion, whereas higher rates offered no additional diagnostic advantage but increased data size and computational cost. Consequently, the 1 MHz rate represents an optimal balance between signal fidelity, data efficiency, and diagnostic performance [52]. A bandpass filter was applied to

the recorded AE signals to isolate the frequency components relevant to cutting-tool fault dynamics. The cutoff frequencies were determined based on both the operating bandwidth of the R15I-AST AE sensors and the observed spectral distribution of the acquired data. The sensors exhibit maximum sensitivity in the 50–400 kHz range, which encompasses the dominant energy of AE bursts produced by tool wear, bearing impacts, and gear defects. Frequencies below 50 kHz were largely attributed to mechanical vibrations of the machine frame and spindle, while frequencies above 450 kHz contained random electromagnetic interference. Accordingly, a 50–450 kHz bandpass filter was adopted to retain all relevant acoustic emission activity while suppressing low-frequency structural noise and high-frequency disturbances. This range provides an optimal balance between noise reduction and information preservation, ensuring that only fault-related AE signatures were used for feature extraction and model training. By isolating the relevant AE frequency range, the recorded data was optimized for fault diagnosis. To facilitate multi-sensor data synchronization, the data acquisition system was operated in dual-channel mode, ensuring that AE signals from both the spindle-mounted and motor-mounted sensors were recorded simultaneously. This synchronized data collection allowed for comparative analysis between normal and faulty operating conditions, making it possible to differentiate between tool-related faults and background noise. The overall setup for the AE data acquisition is shown in Figure 9.




Figure 9. Schematic diagram of the milling machine experimental setup.

4.3. Fault Introduction and Simulated Defect Conditions

To evaluate the fault diagnosis system, three types of defects were deliberately introduced into the milling machine components. These faults were designed to mimic real-world wear and failure mechanisms commonly observed in industrial machining environments. The first fault type was tool wear, which was introduced by subjecting the cutting tool to progressive flank wear until it reached an average of 0.3 mm, as shown in Figure 10a. As the tool deteriorated, the AE signals exhibited distinct changes in energy levels, burst frequencies, and transient peak amplitudes, reflecting the increased resistance experienced during machining. The second fault involved bearing defects, which were artificially created by introducing localized damage to the outer race of the bearing, as shown in Figure 10b. This defect was designed to simulate rolling element fatigue, a common issue in high-speed machining. When the bearing defect interacted with the rotating spindle, it produced repetitive impact forces, generating periodic bursts in the AE signal. These bursts were characterized by high-amplitude oscillations and increased energy levels,

making them distinguishable from normal bearing operation. The third fault concerned gear damage, introduced by removing a small metal fragment from one of the gear teeth responsible for torque transmission. The defect was created in the main spindle gearbox, which transfers torque from the motor to the spindle. A fragment approximately 2–3 mm deep and 1.5 mm wide was removed from a single gear tooth to simulate localized wear. Although this modification produced measurable AE variations, it represents a secondary diagnostic condition compared to the tool and bearing faults emphasized in this study, as shown in Figure 10c. This fault altered the load distribution in the gear system, leading to intermittent fluctuations in torque and creating irregular shock waves. These gear-related AE signals exhibited unique patterns that differentiated them from tool wear and bearing faults. All fault conditions were tested under both idle and cutting conditions, ensuring that the dataset captured AE variations across multiple operating scenarios. By recording AE signals in these conditions, it became possible to construct a robust classification model capable of distinguishing between normal operation and faulty conditions. The faults introduced in this study were designed to emulate realistic deterioration levels typically encountered in industrial milling operations rather than artificially exaggerated damage. The tool fault was produced by progressive edge wear up to an average flank wear of 0.3 mm, in compliance with ISO 8688-2 standards for allowable wear prior to tool replacement. The bearing defect involved a 3 mm localized outer race indentation, simulating rolling element fatigue and replicating natural pitting effects observed in high-speed spindles. The gear fault was introduced by removing a 2–3 mm section (1.5 mm width) from one gear tooth to create a mild imbalance in load transmission without disrupting mechanical alignment. These controlled defect magnitudes represent early to moderate fault stages that generate measurable yet non-trivial acoustic emission variations. All faults were tested under both idle and cutting conditions to capture AE signatures across varying load conditions. The resulting signals exhibited gradual changes in energy and frequency distribution, confirming that the proposed framework effectively recognizes both minor and advanced fault states. This demonstrates the system’s high sensitivity and practical applicability for early fault detection and predictive maintenance in real milling environments.




Figure 10. Faulty machine components considered in this study: (a) end-mill tool with edge damage, (b) gear with surface defect, and (c) bearing with localized fault.

4.4. Dataset Composition and Validation

For each operating condition, including normal, tool wear, bearing fault, and gear fault, a total of 280 AE signal samples were recorded. Each sample was recorded for a duration of 1 s, resulting in 1 million data points per sample due to the 1 MHz sampling rate, as shown in Table 4. The collected dataset provided a comprehensive representation of AE patterns, ensuring that the classification models could be trained with diverse signal variations.

Table 4. Data Acquisition Under Different Operating Conditions.

Operational Conditions	Samples Collected	Acquisition Rate	Duration
Normal (N)	40	1 MHz	2 min
Normal—Idle (NI)	40	1 MHz	2 min
Bearing Fault (BF)	40	1 MHz	2 min
Bearing Fault—Idle (BFI)	40	1 MHz	2 min
Gear Fault (GF)	40	1 MHz	2 min
Gear Fault—Idle (GFI)	40	1 MHz	2 min
Tool Fault (TF)	40	1 MHz	2 min

To validate the dataset, both time-domain and frequency-domain analyses were conducted. The time-domain analysis focused on measuring burst amplitude, signal energy, and transient wave patterns, as shown in Figure 11, while the frequency-domain analysis examined the spectral distribution of AE signals, detecting frequency shifts caused by different fault conditions. Additional post-processing techniques such as signal normalization, noise suppression, and feature extraction were applied to refine the dataset, improving its suitability for machine learning applications.

Synchronization of AE signals across both sensors ensured accurate comparisons between normal and faulty states, enhancing the diagnostic capability of fault classification models. The resulting dataset served as a foundation for developing advanced deep learning-based approaches for automated fault detection and predictive maintenance in milling operations. The experimental setup established in this study provides a highly controlled and repeatable framework for investigating AE-based fault diagnosis in milling cutting tools. By integrating precision machining, high-resolution AE sensing, synchronized data acquisition, and systematic fault introduction, the collected dataset is well-structured for developing intelligent fault detection models. The comprehensive data collection process ensures that real-world faults can be accurately identified, classified, and predicted, contributing to improved reliability in manufacturing operations.




Figure 11. AE signals in the time domain under various operating conditions: (a) N; (b) NI; (c) BF; (d) BFI; (e) GF; (f) GFI; and (g) TF.

5. Results

This section presents a detailed experimental evaluation of the proposed FSL framework for fault diagnosis in milling machines. The experiments were designed to replicate challenging real-world scenarios where only a limited number of labeled samples are available per fault category. All evaluations were conducted using CWT scalogram images generated from AE signals. To ensure fairness and reproducibility, the same preprocessing pipeline, training strategy, and computational environment were applied to both the proposed model and all comparison architectures. The analysis is structured into three parts: (a) overall classification performance, (b) few-shot generalization, and (c) visualization of learned representations.

The proposed approach was compared against widely used CNN architectures, including ResNet-18, ResNet-50, ShuffleNetV2, MobileNetV3 Large, DenseNet-201, and SqueezeNet. All models were trained using the same preprocessing strategy: images were resized to 224×224 , normalized to $[-1, 1]$, and augmented using random horizontal flips and slight rotations. Training and evaluation were conducted on a single NVIDIA RTX 3060 GPU using PyTorch 2.8.0. Classification performance was evaluated using accuracy, precision, recall, F1-score, specificity, Matthew's correlation coefficient (MCC), and training time, as provided in Table 5. The proposed model achieved 99.32% accuracy, reflecting perfect generalization and better classification performance under data-constrained conditions. In contrast, ResNet-18 achieved 91.43% accuracy, ResNet-50 85.71%, ShuffleNetV2 70.00%, MobileNetV3 84.29%, DenseNet-201 89.29, and SqueezeNet 78.57%. Despite DenseNet-201 requiring over 11 min for training, it still underperformed compared to the proposed lightweight framework, which required only one minute. This highlights the proposed method's better classification performance and computational efficiency, making it suit-

able for real-time and resource-limited industrial environments. The high precision and recall values achieved by the proposed framework demonstrate its ability to correctly classify fault conditions while minimizing both false positives and false negatives. These results confirm the model's strong regularization ability and effective optimization under low-data conditions, with no signs of overfitting. The visual evidence aligns with these quantitative results.

Table 5. Classification Performance Comparison.

Model	Acc (%)	Prec (%)	Rec (%)	F1 (%)	MCC (%)	Spec (%)	Time (min)
ResNet-18	91.43	91.20	91.40	91.30	90.34	93.10	1:09
ResNet-50	87.14	87.00	87.10	87.00	83.57	92.00	1:15
ShuffleNetV2	70.00	73.80	70.00	71.20	58.84	88.30	0:27
MobileNetV3	84.29	84.60	84.30	84.20	80.38	91.00	0:30
DenseNet-201	89.29	89.40	89.30	89.20	80.77	93.40	11:21
SqueezeNet	78.57	78.90	78.60	78.50	71.02	90.10	1:10
Proposed	99.32	99.30	99.40	99.30	98.90	99.40	1:00

This reliability is further illustrated by the confusion matrix, which shows perfect classification across all operating conditions as shown in Figure 12a. The training curves of the proposed method further support these findings. The model exhibited a smooth ascent in accuracy and a consistent decline in loss, converging to 100% accuracy with near-zero training loss over 500 epochs, as shown in Figure 12b,c. Occasional small fluctuations in the loss were attributed to the dynamics of contrastive loss and the episodic sampling strategy, but did not impact convergence stability.

To better understand the distribution and discriminability of the learned features, we plotted t-distributed Stochastic Neighbor Embedding (t-SNE) to visualize high-dimensional embeddings in a two-dimensional space. The resulting plots showed tight intra-class clusters and well-separated inter-class margins as shown in Figure 12d, indicating the model's capability to learn highly discriminative feature representations. This improvement can be attributed to the self-attention mechanism, which allows the model to focus on the most informative time–frequency regions in the scalograms, and the contrastive loss function, which reinforces both intra-class compactness and inter-class separation. These insights are supported by the better feature learning capacity of the model's SA-ResNet50 backbone, which enhances the extraction of meaningful and contextually rich representations.

To further quantify the model's decision confidence and class discrimination, Receiver Operating Characteristic (ROC) curves were plotted for each class as shown in Figure 12e. The proposed framework achieved AUC scores close to 1.0 across all fault categories, confirming its high confidence and robustness even in low-data scenarios. These high AUC values validate that the model not only distinguishes between classes with great accuracy but also maintains strong generalization across unseen query samples.

One of the core challenges in industrial fault diagnosis is adapting to conditions with extremely limited labeled samples. Table 6 presents the classification performance under 1-shot, 3-shot, and 5-shot learning setups. As expected, performance improves steadily with more labeled samples. The 1-shot configuration yielded 87.46% accuracy, sufficient for coarse classification, while the 3-shot setting improved accuracy to 93.18%. With 5 shots, the model achieved 99.32% accuracy, highlighting its ability to generalize even under highly constrained supervision. The proposed framework consistently outperformed conventional CNNs across all evaluations. Its ability to deliver high accuracy under 5-shot settings,

adverse feature clustering, and computational efficiency establishes it as a highly practical solution for fault diagnosis in modern manufacturing environments.




Figure 12. Result visualization of proposed method: (a) Confusion matrix, (b) Accuracy curve, (c) Loss curve, (d) t-SNE, and (e) ROC curve.

Table 6. Few-Shot Generalization Performance.

Setting	Acc (%)	Prec (%)	Rec (%)	F1 (%)
1-shot	87.46	87.50	87.40	87.40
3-shot	93.18	93.20	93.20	93.10
5-shot	99.32	99.30	99.40	99.30

To further assess the generalization capability of the proposed method, additional experiments were conducted using datasets collected under two distinct spindle speeds: 1440 rpm and 660 rpm. These conditions were chosen to simulate different cutting environments and evaluate the robustness of the model under variable AE characteristics. The quantitative results for both spindle speeds are summarized in Table 7. At 1440 rpm, the model achieved an average precision of 97.4%, a recall of 97.2%, and an F1-score of 97.3%, indicating reliable discrimination among tool-wear classes under high-speed conditions. At 660 rpm, all performance metrics exceeded 99.8%, demonstrating the framework's ability to maintain nearly perfect accuracy under low-speed operation. The model achieved high and balanced classification performance for both the updated and newly collected data. The confusion matrix in Figure 13a for 660 rpm displays strong diagonal dominance, and the t-SNE visualization in Figure 13b shows clear separation between all states, indicating that the learned embeddings remain discriminative under high-speed operation. The confusion matrix, as in Figure 14a for the 1440 rpm, confirms better predictions, and the t-SNE projection in Figure 14b demonstrates complete separability among the feature clusters. The model achieved high and balanced classification performance for both the updated and newly collected data. The confusion matrix in Figure 13a for 660 rpm displays strong

diagonal dominance, and the t-SNE visualization shows clear separation between all states, indicating that the learned embeddings remain discriminative under high-speed operation. The confusion matrix, as in Figure 14a for the 1440 rpm, confirms better predictions, and the t-SNE projection in Figure 14b demonstrates complete separability among the feature clusters. These results clearly indicate that the proposed framework maintains consistent discriminative ability across different spindle speeds. The consistent results across both speed operations validate the robustness and stability of the proposed network and confirm its ability to generalize effectively across diverse machining conditions.

Table 7. Comparative Performance of the Proposed Model Under Different Spindle Speeds.

Spindle Speed (rpm)	Class	Precision	Recall	F1-Score	Support
660	BF	0.998	0.999	0.998	15
	GF	0.999	0.997	0.998	15
	N	0.998	0.998	0.998	15
	TF	0.999	0.999	0.999	15
	Macro Avg	0.999	0.998	0.998	60
1440	BF	0.976	0.973	0.974	14
	GF	0.972	0.974	0.973	14
	N	0.974	0.970	0.972	14
	TF	0.975	0.972	0.974	14
	Macro Avg	0.974	0.972	0.973	56




Figure 13. Results for the 660 rpm dataset: (a) confusion matrix, and (b) t-SNE visualization.




Figure 14. Results for the 1440 rpm dataset: (a) confusion matrix, and (b) t-SNE visualization.

6. Discussion

The integration of Mahalanobis distance into the prototype matching process played a significant role in enhancing classification robustness. Unlike other distances, Mahalanobis distance considers the covariance structure of class-wise features, allowing for more accurate similarity matching, especially in cases involving subtle signal differences. Furthermore, the adaptive prototype learning mechanism dynamically refined class representations during training, improving intra-class consistency and inter-class distinctiveness. These components, working together with contrastive loss, contributed to robust class separability and minimized misclassifications. The proposed model's strong few-shot learning performance is thus attributed to this combination of adaptive prototype learning, attention-guided representation, and covariance-aware matching, which enables it to generalize effectively even under data scarcity.

To address concerns regarding overfitting in the proposed method, a comprehensive 5-fold stratified cross-validation was conducted to rigorously evaluate the robustness and generalization. The detailed fold-wise results confirm that the model maintains better consistency across all partitions, with validation accuracy ranging from 97.14% to 100% and a mean accuracy of $98.86\% \pm 0.97\%$ (95% CI = [98.01%, 99.71%]). The low standard deviation ($\sigma = 0.97\%$) across folds indicates that the model's performance is not sensitive to specific train-test splits. The training dynamics illustrated in Figure 15a show uniform and stable convergence behavior across all five folds, with each fold achieving rapid convergence within the first 200 episodes. The training-validation comparison shown in Figure 15b further supports this, as the performance gap remains consistently narrow ($\approx 1\%$ on average) between the training accuracy ($99.99\% \pm 0.02\%$) and validation accuracy ($98.86\% \pm 0.97\%$). Such a minimal gap demonstrates effective generalization and indicates that the model learns meaningful prototypical representations. Moreover, the mean confusion matrix in Figure 15c exhibits pronounced diagonal dominance across all seven conditions, underscoring the model's discriminative ability and class-level stability. The strong consistency across folds substantiates that the proposed network achieves reliable few-shot learning performance even under constrained data conditions. Although environmental noise and sensor calibration can affect AE data acquisition, the use of the Hsu–Nielsen test and 50–450 kHz band-pass filtering ensured consistent sensor sensitivity and minimized external interference during all experiments.

To further interpret the network's decision process, Gradient-weighted Class Activation Mapping++ (Grad-CAM++) was applied to visualize the most influential regions within the CWT scalograms. Figure 16 presents the class-wise average activation maps for all seven fault conditions (BF, BFI, GF, GFI, N, NI, and TF). The red regions indicate higher activation intensity, representing the areas where the model focused during classification. The visualization results show that each fault type activates distinct time–frequency regions, revealing the network's ability to capture fault-specific acoustic characteristics. For example, severe breakage conditions (BF, TF) show concentrated activations in high-frequency regions associated with impact-type emissions, whereas gradual wear faults (GF, GFI) produce distributed mid-frequency responses. The normal states (N, NI) exhibit low and stable activation areas, corresponding to the absence of sudden AE bursts. These findings confirm that the attention-augmented ResNet backbone effectively emphasizes meaningful discriminative patterns rather than noise. Overall, the Grad-CAM++ analysis provides valuable interpretability by visualizing the relationship between learned representations and physical fault behaviors, thereby reinforcing the transparency and reliability of the proposed few-shot learning framework for fault diagnosis.




Figure 15. 5-fold cross-validation results of the proposed model: (a) training convergence across folds, (b) training vs. validation accuracy comparison, and (c) mean confusion matrix showing consistent classification performance.




Figure 16. Class-wise Grad-CAM++ activation maps showing discriminative time–frequency regions.

These findings confirm the proposed framework’s effectiveness in overcoming the core challenges of few-shot fault diagnosis in milling machine applications. Its perfect classification performance, interpretability through visualization tools, strong generalization under minimal data, and computational efficiency establish it as a strong candidate for

real-world deployment. The framework is highly scalable and adaptable, and through its integration of adaptive prototype learning, Mahalanobis distance, and attention-driven feature extraction, consistently outperforms baseline few-shot learning approaches. It holds strong potential for extension to other industrial domains such as predictive maintenance, real-time monitoring, and anomaly detection. In summary, the proposed framework consistently demonstrated near-perfect performance under the 5-shot settings while maintaining strong generalization under more restrictive conditions. Its scalability, interpretability, and computational efficiency make it a practical and robust solution for real-world milling machine fault diagnosis.

Although the proposed framework demonstrates strong diagnostic performance and robustness across varying spindle speeds, certain limitations remain that present valuable opportunities for future improvement. The present study was conducted using AE data collected under controlled laboratory conditions, which may not fully capture the range of variability present in real industrial environments. Future work will therefore focus on extending the dataset to include multiple machines, tool materials, and cutting parameters to further validate domain generalization. In addition, while the current model performs effectively in offline analysis, deploying it in real-time monitoring systems will require optimization for computational efficiency and latency reduction. Moreover, since this study relied solely on AE signals, future extensions will explore the integration of multi-sensor information, such as vibration, cutting force, and spindle current, to improve robustness under noisy conditions. Finally, further evaluation under extremely low-shot or cross-domain conditions will be pursued to enhance model stability and adaptability. Addressing these aspects will help transition the proposed framework from laboratory validation to a fully scalable and explainable deployment in industrial TCM.

7. Conclusions

This paper presented an optimized Few-Shot Learning framework for fault diagnosis in milling machines using acoustic emission (AE) signals. The approach effectively addresses the challenges of data scarcity and non-stationary signal behavior by integrating CWT-based time–frequency feature extraction, self-attention ResNet-50 encoding, and adaptive prototype learning with Mahalanobis distance matching. The incorporation of contrastive loss with hard negative mining further enhances class separability, leading to compact and well-defined embedding spaces. Experimental validations conducted under multiple spindle speeds (660 rpm and 1440 rpm) and 5-fold stratified cross-validation confirm the robustness and reproducibility of the model, with consistently high accuracy, precision, and F1-scores across all partitions. The narrow performance variance ($\sigma = 0.97\%$) demonstrates that the model generalizes effectively beyond specific train–test splits. Moreover, interpretability analysis using Grad-CAM++ visualizations revealed that the network attends to fault-relevant signal regions, providing transparency in the decision-making process. The findings highlight the framework’s strong diagnostic sensitivity across both minor and severe fault stages, reflecting its real-world applicability for early-stage fault detection and predictive maintenance. While the model shows outstanding performance on controlled laboratory data, its reliance on AE signals from a single machine setup represents a limitation that future work will address through multi-machine validation, domain adaptation, and multi-modal sensor integration. Future research will also explore online learning strategies and adaptive episodic training to enhance real-time deployment in dynamic industrial environments. Overall, the proposed framework provides a reliable, data-efficient, and interpretable AI solution for intelligent fault diagnosis in modern manufacturing systems.

Author Contributions: Conceptualization, F.S., M.U. and J.-M.K.; methodology, F.S., M.U. and J.-M.K.; validation, F.S., M.U. and J.-M.K.; formal analysis, F.S., M.U. and J.-M.K.; resources, F.S., M.U. and J.-M.K.; writing—original draft preparation, F.S., M.U. and J.-M.K.; writing—review and editing, J.-M.K.; visualization, F.S., M.U. and J.-M.K.; project administration, J.-M.K.; funding acquisition, J.-M.K. All authors have read and agreed to the published version of the manuscript.

Funding: This result was supported by the “Regional Innovation System & Education (RISE)” through the Ulsan RISE Center, funded by the Ministry of Education (MOE) and the Ulsan Metropolitan City, Republic of Korea (2025-RISE-07-001). This work was also supported by the Technology Innovation Program (20023566, Development and Demonstration of Industrial IoT and AI Based Process Facility Intelligence Support System in Small and Medium Manufacturing Sites) funded By the Ministry of Trade, Industry&Energy (MOTIE, Republic of Korea).

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The raw data supporting the conclusions of this article will be made available by the authors on request.

Conflicts of Interest: Author Jong-Myon Kim was employed by the company Prognosis and Diagnostic Technologies Co., Ltd. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

References

1. Soori, M.; Arezoo, B.; Dastres, R. Machine learning and artificial intelligence in CNC machine tools, A review. *Sustain. Manuf. Serv. Econ.* **2023**, *2*, 100009. [CrossRef]
2. Przybyś-Małaczek, A.; Antoniuk, I.; Szymanowski, K.; Kruk, M.; Kurek, J. Application of Machine Learning Algorithms for Tool Condition Monitoring in Milling Chipboard Process. *Sensors* **2023**, *23*, 5850. [CrossRef]
3. Umar, M.; Ahmad, Z.; Ullah, S.; Saleem, F.; Siddique, M.F.; Kim, J.-M. Advanced Fault Diagnosis in Milling Machines Using Acoustic Emission and Transfer Learning. *IEEE Access* **2025**, *13*, 100776–100790. [CrossRef]
4. Zhu, Z.; Tang, X.; Chen, C.; Peng, F.; Yan, R.; Zhou, L.; Li, Z.; Wu, J. High precision and efficiency robotic milling of complex parts: Challenges, approaches and trends. *Chin. J. Aeronaut.* **2022**, *35*, 22–46. [CrossRef]
5. Wang, W.; Guo, Q.; Yang, Z.; Jiang, Y.; Xu, J. A state-of-the-art review on robotic milling of complex parts with high efficiency and precision. *Robot. Comput. Integr. Manuf.* **2023**, *79*, 102436. [CrossRef]
6. Kaliyannan, D.; Thangamuthu, M.; Pradeep, P.; Gnansekaran, S.; Rakkiyannan, J.; Pramanik, A. Tool Condition Monitoring in the Milling Process Using Deep Learning and Reinforcement Learning. *J. Sens. Actuator Netw.* **2024**, *13*, 42. [CrossRef]
7. Bourassa, D.; Gauthier, F.; Abdul-Nour, G. Equipment failures and their contribution to industrial incidents and accidents in the manufacturing industry. *Int. J. Occup. Saf. Ergon.* **2016**, *22*, 131–141. [CrossRef] [PubMed]
8. Ahmad, Z.; Ullah, S.; Maliuk, A.S.; Kim, J.-M. Milling machine fault detection and identification based on a novel vitality index and temporal-residual network. *Appl. Acoust.* **2025**, *239*, 110861. [CrossRef]
9. Cavalcante, C.A.V.; Lopes, R.S.; Scarf, A. Inspection and replacement policy with a fixed periodic schedule. *Reliab. Eng. Syst. Saf.* **2021**, *208*, 107402. [CrossRef]
10. Twardowski, P.; Tabaszewski, M.; Wiciak-Pikuła, M.; Felusiak-Czyryca, A. Identification of tool wear using acoustic emission signal and machine learning methods. *Precis. Eng.* **2021**, *72*, 738–744. [CrossRef]
11. Zhou, J.-H.; Pang, C.K.; Zhong, Z.-W.; Lewis, F.L. Tool Wear Monitoring Using Acoustic Emissions by Dominant-Feature Identification. *IEEE Trans. Instrum. Meas.* **2011**, *60*, 547–559. [CrossRef]
12. Zhou, Y.; Liu, C.; Yu, X.; Liu, B.; Quan, Y. Tool wear mechanism, monitoring and remaining useful life (RUL) technology based on big data: A review. *SN Appl. Sci.* **2022**, *4*, 232. [CrossRef]
13. Cheng, Y.; Gai, X.; Guan, R.; Jin, Y.; Lu, M.; Ding, Y. Tool wear intelligent monitoring techniques in cutting: A review. *J. Mech. Sci. Technol.* **2023**, *37*, 289–303. [CrossRef]
14. Pimenov, D.Y.; Gupta, M.K.; da Silva, L.R.R.; Kiran, M.; Khanna, N.; Krolczyk, G.M. Application of measurement systems in tool condition monitoring of Milling: A review of measurement science approach. *Measurement* **2022**, *199*, 111503. [CrossRef]
15. Dimla, D.E. Sensor signals for tool-wear monitoring in metal cutting operations—A review of methods. *Int. J. Mach. Tools Manuf.* **2000**, *40*, 1073–1098. [CrossRef]

16. Zaman, W.; Siddique, M.F.; Ullah, S.; Saleem, F.; Kim, J.-M. Hybrid Deep Learning Model for Fault Diagnosis in Centrifugal Pumps: A Comparative Study of VGG16, ResNet50, and Wavelet Coherence Analysis. *Machines* **2024**, *12*, 905. [CrossRef]
17. Sick, B. ON-LINE AND INDIRECT TOOL WEAR MONITORING IN TURNING WITH ARTIFICIAL NEURAL NETWORKS: A REVIEW OF MORE THAN A DECADE OF RESEARCH. *Mech. Syst. Signal Process* **2002**, *16*, 487–546. [CrossRef]
18. Umar, M.; Siddique, M.F.; Ullah, N.; Kim, J.-M. Milling Machine Fault Diagnosis Using Acoustic Emission and Hybrid Deep Learning with Feature Optimization. *Appl. Sci.* **2024**, *14*, 10404. [CrossRef]
19. Nasir, V.; Dibaji, S.; Alaswad, K.; Cool, J. Tool wear monitoring by ensemble learning and sensor fusion using power, sound, vibration, and AE signals. *Manuf. Lett.* **2021**, *30*, 32–38. [CrossRef]
20. Prickett, W.; Johns, C. An overview of approaches to end milling tool monitoring. *Int. J. Mach. Tools Manuf.* **1999**, *39*, 105–122. [CrossRef]
21. Marinescu, I.; Axinte, D.A. A critical analysis of effectiveness of acoustic emission signals to detect tool and workpiece malfunctions in milling operations. *Int. J. Mach. Tools Manuf.* **2008**, *48*, 1148–1160. [CrossRef]
22. Turšič, N.; Klančnik, S. Tool Condition Monitoring Using Machine Tool Spindle Current and Long Short-Term Memory Neural Network Model Analysis. *Sensors* **2024**, *24*, 2490. [CrossRef]
23. Yan, S.; Sui, L.; Wang, S.; Sun, Y. On-line tool wear monitoring under variable milling conditions based on a condition-adaptive hidden semi-Markov model (CAHSM). *Mech. Syst. Signal Process* **2023**, *200*, 110644. [CrossRef]
24. Zhou, Y.; Zhi, G.; Chen, W.; Qian, Q.; He, D.; Sun, B.; Sun, W. A new tool wear condition monitoring method based on deep learning under small samples. *Measurement* **2022**, *189*, 110622. [CrossRef]
25. Xiang, G.; Chen, W.; Peng, Y.; Wang, Y.; Qu, C. Deep Transfer Learning Based on Convolutional Neural Networks for Intelligent Fault Diagnosis of Spacecraft. In Proceedings of the 2020 Chinese Automation Congress (CAC), Shanghai, China, 6–8 November 2020; pp. 5522–5526. [CrossRef]
26. Siddique, M.F.; Saleem, F.; Umar, M.; Kim, C.H.; Kim, J.-M. A Hybrid Deep Learning Approach for Bearing Fault Diagnosis Using Continuous Wavelet Transform and Attention-Enhanced Spatiotemporal Feature Extraction. *Sensors* **2025**, *25*, 2712. [CrossRef] [PubMed]
27. Yan, J.; Cheng, Y.; Wang, Q.; Liu, L.; Zhang, W.; Jin, B. Transformer and Graph Convolution-Based Unsupervised Detection of Machine Anomalous Sound Under Domain Shifts. *IEEE Trans. Emerg. Top. Comput. Intell.* **2024**, *8*, 2827–2842. [CrossRef]
28. Zhu, Z.; Lei, Y.; Qi, G.; Chai, Y.; Mazur, N.; An, Y.; Huang, X. A review of the application of deep learning in intelligent fault diagnosis of rotating machinery. *Measurement* **2023**, *206*, 112346. [CrossRef]
29. Saleem, F.; Ahmad, Z.; Kim, J.-M. Real-Time Pipeline Leak Detection: A Hybrid Deep Learning Approach Using Acoustic Emission Signals. *Appl. Sci.* **2024**, *15*, 185. [CrossRef]
30. Lupea, I.; Lupea, M. Continuous Wavelet Transform and CNN for Fault Detection in a Helical Gearbox. *Appl. Sci.* **2025**, *15*, 950. [CrossRef]
31. Barbosh, M.; Dunphy, K.; Sadhu, A. Acoustic emission-based damage localization using wavelet-assisted deep learning. *J. Infrastruct. Preserv. Resil.* **2022**, *3*, 6. [CrossRef]
32. Saleem, F.; Ahmad, Z.; Siddique, M.F.; Umar, M.; Kim, J.-M. Acoustic Emission-Based Pipeline Leak Detection and Size Identification Using a Customized One-Dimensional DenseNet. *Sensors* **2025**, *25*, 1112. [CrossRef] [PubMed]
33. Zhang, B.; Wang, W.; He, Y. A hybrid approach combining deep learning and signal processing for bearing fault diagnosis under imbalanced samples and multiple operating conditions. *Sci. Rep.* **2025**, *15*, 13606. [CrossRef] [PubMed]
34. Neupane, D.; Bouadjene, M.R.; Dazeley, R.; Aryal, S. Data-driven machinery fault diagnosis: A comprehensive review. *Neurocomputing* **2025**, *627*, 129588. [CrossRef]
35. LeCun, Y.; Bengio, Y.; Hinton, G. Deep learning. *Nature* **2015**, *521*, 436–444. [CrossRef]
36. Gao, Y.; Li, H.; Fu, W. Few-shot learning for image-based bridge damage detection. *Eng. Appl. Artif. Intell.* **2023**, *126*, 107078. [CrossRef]
37. Li, K.; Ye, H.; Gao, X.; Zhang, L. Meta-Learning With Intraclass and Interclass Optimization for Few-Shot Fault Diagnosis. *IEEE Trans. Industr. Inform.* **2025**, *21*, 713–722. [CrossRef]
38. Liang, X.; Zhang, M.; Feng, G.; Wang, D.; Xu, Y.; Gu, F. Few-Shot Learning Approaches for Fault Diagnosis Using Vibration Data: A Comprehensive Review. *Sustainability* **2023**, *15*, 14975. [CrossRef]
39. Wang, D.; Zhang, M.; Xu, Y.; Lu, W.; Yang, J.; Zhang, T. Metric-based meta-learning model for few-shot fault diagnosis under multiple limited data conditions. *Mech. Syst. Signal Process* **2021**, *155*, 107510. [CrossRef]
40. Li, K.; Shang, C.; Ye, H. Reweighted Regularized Prototypical Network for Few-Shot Fault Diagnosis. *IEEE Trans. Neural Netw. Learn. Syst.* **2024**, *35*, 6206–6217. [CrossRef]
41. ALataiqeh, M.; Shi, H.; Qu, Q.; Mei, X.; Wang, H. A Novel Approach to a Few Shot Learning Techniques Based on Thermal Error Modeling for Slant Bed CNC Lathe Machine. *Int. J. Precis. Eng. Manuf.* **2025**, *26*, 1431–1448. [CrossRef]
42. Gong, Z.; Huo, D. Tool condition monitoring in micro milling of brittle materials. *Precis. Eng.* **2024**, *87*, 11–22. [CrossRef]

43. Wen, L.; Li, X.; Gao, L.; Zhang, Y. A New Convolutional Neural Network-Based Data-Driven Fault Diagnosis Method. *IEEE Trans. Ind. Electron.* **2018**, *65*, 5990–5998. [CrossRef]
44. Zhao, R.; Yan, R.; Chen, Z.; Mao, K.; Wang, P.; Gao, R.X. Deep learning and its applications to machine health monitoring. *Mech. Syst. Signal Process* **2019**, *115*, 213–237. [CrossRef]
45. Vu, M.-H.; Nguyen, V.-Q.; Tran, T.-T.; Pham, V.-T.; Lo, M.-T. Few-Shot Bearing Fault Diagnosis Via Ensembling Transformer-Based Model With Mahalanobis Distance Metric Learning From Multiscale Features. *IEEE Trans. Instrum. Meas.* **2024**, *73*, 1–18. [CrossRef]
46. Chen, Y.; Yue, J.; Liu, Z.; Chen, J. A semi-supervised wise-attention weighted prototype network for rolling bearing fault diagnosis under noisy and limited labeled data conditions. *Neurocomputing* **2025**, *647*, 130563. [CrossRef]
47. Ahmed, Y.S.; Arif, A.F.M.; Veldhuis, S.C. Application of the wavelet transform to acoustic emission signals for built-up edge monitoring in stainless steel machining. *Measurement* **2020**, *154*, 107478. [CrossRef]
48. Xie, S.; Girshick, R.; Dollar, P.; Tu, Z.; He, K. Aggregated Residual Transformations for Deep Neural Networks. In Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA, 21–26 July 2017; pp. 5987–5995. [CrossRef]
49. Yu, M.; Liu, H.; Wang, R.; Kong, X.; Hu, Z.; Li, X. An improved CNN based on attention mechanism with multi-domain feature fusion for bearing fault diagnosis. In Proceedings of the 2021 IEEE International Conference on Prognostics and Health Management (ICPHM), Detroit, MI, USA, 7–9 June 2021; pp. 1–7. [CrossRef]
50. Li, X.; Li, Y.; Zheng, Y.; Zhu, R.; Ma, Z.; Xue, J.; Cao, J. ReNAP: Relation network with adaptive prototypical learning for few-shot classification. *Neurocomputing* **2023**, *520*, 356–364. [CrossRef]
51. Zhao, X.; Zhao, Z.; Cui, X.; Yin, J.; Cao, J.; Wang, H. Few-shot Bearing Fault Diagnosis using Adaptive Detail Convolution and Global KAN-Transformer with Mahalanobis Distance. *J. Vib. Eng. Technol.* **2025**, *13*, 296. [CrossRef]
52. Magadán, L.; Ruiz-Cárcel, C.; Granda, J.C.; Suárez, F.J.; Starr, A. Explainable and interpretable bearing fault classification and diagnosis under limited data. *Adv. Eng. Inform.* **2024**, *62*, 102909. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Article

Defective Photovoltaic Module Detection Using EfficientNet-B0 in the Machine Vision Environment

Minseop Shin, Junyoung Seo, In-Bae Lee and Sojung Kim *

Department of Industrial and Systems Engineering, Dongguk University-Seoul, Seoul 04620, Republic of Korea; 2025126785@dgu.ac.kr (M.S.); seojun9909@dgu.ac.kr (J.S.); leeinbae@dgu.ac.kr (I.-B.L.)

* Correspondence: sojungkim@dgu.ac.kr; Tel.: +82-2-2260-3375

Abstract

Machine vision based on artificial intelligence technology is being actively utilized to reduce defect rates in the photovoltaic module production process. This study aims to propose a machine vision approach using EfficientNet-B0 for defective photovoltaic module detection. In particular, the proposed approach is applied to the electroluminescence (EL) operation, which identifies microcracks in PV modules by using polarization current. The proposed approach extracts low-level structures and local brightness variations, such as busbars, fingers, and cell boundaries, from a single convolutional block. Furthermore, the mobile inverted bottleneck convolution (MBConv) block progressively transforms defect patterns—such as microcracks and dark spots—that appear at various shooting angles into high-level feature representations. The converted image is then processed using global average pooling (GAP), Dropout, and a final fully connected layer (Dense) to calculate the probability of a defective module. A sigmoid activation function is then used to determine whether a PV module is defective. Experiments show that the proposed Efficient-B0-based methodology can stably achieve defect detection accuracy comparable to AlexNet and GoogLeNet, despite its relatively small number of parameters and fast processing speed. Therefore, this study will contribute to increasing the efficiency of EL operation in industrial fields and improving the productivity of PV modules.

Keywords: renewable energy; machine vision; artificial intelligence; photovoltaic module; photovoltaic

1. Introduction

Machine vision is a technology that uses camera sensors installed in manufacturing facilities and uses artificial intelligence (AI) to assess product quality based on captured images [1]. It automates the product quality assessment process previously performed by humans. By automating human quality assessment through machine vision, machine vision can reduce errors in quality judgments caused by human perception and worker fatigue [2]. With recent advancements in AI technology, demand for machine vision has steadily increased, and its application fields are expanding to include semiconductors, electric vehicles, and healthcare [3]. The machine vision market was estimated at approximately \$13.95 billion in 2025 and is projected to reach \$19.81 billion by 2030, growing at a compound annual growth rate (CAGR) of 7.27% during the forecast period (2025–2030) [4].

The production of photovoltaic (PV) modules—devices that directly convert solar energy into electrical energy using semiconductors—also actively utilizes machine vision technology, as in semiconductor production. Because solar cells, the core components of PV

modules, are made primarily of semiconductor materials like silicon, quality inspection of PV modules is crucial for ensuring excellent photovoltaic performance (the phenomenon of sunlight striking a semiconductor and generating electricity) [5]. Machine vision is utilized to accurately verify the alignment and position of PV cells, minimizing errors during assembly [6]. In 2024, global new solar power generation capacity was expected to reach approximately 599 GW, and it was projected to increase 17% year-on-year to approximately 700 GW by 2025. Consequently, demand for photovoltaic modules is expected to continue to grow. PV module manufacturing capacity is projected to steadily increase from 1155 GW in 2023 to 1546 GW by 2035 [7], implying continued growth in machine vision applications.

According to [6], the PV module manufacturing process typically involves sequential production operations, including a tapping operation to connect solar cells, a framing operation to assemble frames onto individual modules, and a junction box mounting operation to connect busbars and junction box terminals. Because connecting solar cells and individual PV modules is a key manufacturing step, defects in solar cells and modules can occur in various ways during the connection process [8]. These defects lead to additional tasks, such as product disassembly, reassembly, and disposal, which increase production costs [9]. Therefore, efforts to reduce defects in the PV module manufacturing process are actively being pursued in major PV module-producing countries (e.g., China [10], Taiwan [11], and South Korea [12]), including the development of early defect detection technologies using machine vision.

In particular, because electroluminescence (EL) is an important operation in the PV module manufacturing process and can be used to obtain images of microcracks, automating the detection of defective PV modules using machine vision can improve manufacturing efficiency and reduce defective products. This study aims to improve the machine vision approach by utilizing EfficientNet-B0 to enhance the production quality of PV modules. EfficientNet-B0 employs a hybrid model scaling strategy that simultaneously expands network depth, channel width, and input resolution. It is designed with a lightweight architecture that includes mobile inverted bottleneck convolution (MBConv) and squeeze-and-excitation (SE) modules [13]. It is designed to have excellent scalability across depth, width, and resolution, effectively limiting computational load and model size while maintaining high-resolution input, which makes it suitable for defective PV module inspection environments during the EL operation. To address the class imbalance problem between good and bad PV module images, this study uses ImageNet pretrained weights as initial values and applies data augmentation and hierarchical segmentation strategies in parallel to stabilize representation learning for the bad class and mitigate overfitting. In experiments, we compare the performance of the proposed methodology with existing convolutional neural network (CNN)-based algorithms (e.g., AlexNet, ResNet50, DenseNet121, and GoogLeNet). We also perform a sensitivity analysis of how performance changes with the optimal image capture angle and bad detection threshold for the proposed methodology to achieve optimal performance. As a result, the proposed EfficientNet-B0-based methodology is a suitable solution for integration with machine vision due to its small model size and easy deployment on industrial PCs or GPUs.

The contributions of this study are as follows. First, the proposed Efficient-B0-based methodology demonstrates superior performance in PV module defect detection compared to existing CNN-based algorithms (e.g., AlexNet, ResNet50, DenseNet121, and GoogLeNet) in terms of accuracy, precision, recall, and F1-score. Compared to existing CNN-based algorithms, EfficientNet-B0 achieves high expressive power with fewer parameters and fewer floating-point operations (FLOPs). Second, this study analyzes the performance sensitivity to the image capture angle, confirming that the proposed Efficient-B0-based methodology achieves optimal performance. Unlike previous studies that focused solely

on performance comparisons of AI algorithms, we analyzed performance sensitivity to the image capture angle, considering that defective component detection performance can vary depending on the angle of the image sensor used to capture PV modules during the EL process. Finally, we performed a sensitivity analysis on the performance of the proposed methodology with respect to the classification threshold settings used to detect defective PV modules. Since the performance of defect detection varies not only with the given algorithm but also with the threshold value that serves as the classification criterion, this analysis reveals which threshold value should be set for more accurate detection of defective PV modules in the EL process.

This study is structured as follows. Section 2 describes the PV manufacturing process and a PV module defect detection technique utilizing a machine vision framework and Efficient-B0 technology. Section 3 validates the performance of the proposed Efficient-B0-based PV module defect detection method using data collected from a PV module manufacturing facility in Korea. Section 4 presents the results of this study and proposes future research directions.

2. Materials and Methods

2.1. Machine Vision in Photovoltaic Module Manufacturing

A photovoltaic (PV) system consists of a frame, front and back covers, an encapsulant, solar cells, and a junction box that converts solar radiant energy into direct current (DC) through multiple PV modules [14]. It is crucial to manufacture PV modules in which solar cells can efficiently convert photon energy associated with a specific wavelength into electrical energy through semiconductors (i.e., p-type and n-type silicon layers) at high efficiency [15]. That is, when a solar cell receives sunlight, electrons within the silicon layer are emitted from the n-type silicon layer and move to an external wire through the p-type silicon layer, generating direct current. Since the final output is determined by the physical conversion efficiency, research is being conducted on the crystal structure of the solar cell itself and the utilization of materials such as cadmium telluride (CdTe) and CuInGaSe₂ (CIGS) [5].

In terms of the manufacturing process, the focus is on the efficient production of PV modules, and the production process typically involves seven sequential operations, including tapping, automatic bussing, EL, laminating, a framing station, and frame and junction box mounting [14]. The tapping process is the first step in which multiple solar cells are soldered and connected in series to form string cells, while automatic bussing is the process of connecting panel strings consisting of multiple solar cells to a board to manufacture PV modules [16,17]. EL testing utilizes a polarization current to emit EL radiation from solar cells, which can be used to obtain high-resolution images that can identify microcracks [18]. Laminating is the process of covering the panel strings with encapsulating material under high pressure and high temperature conditions; the framing station manufactures frames for the encapsulated panels, and frame and junction box mounting is the operation of attaching the encapsulated panel strings to the PV frame and using sealant or adhesive, attaching the junction box to the back of the frame, which is connected to the busbar of the solar string and transmits the generated electricity to external electrical devices [19].

Among these seven operations, the EL operation is the one in which machine vision can be utilized to detect defective manufactured parts, as it can provide images of microcracks. Figure 1 shows the machine vision framework for PV manufacturing [6], which monitors the appearance of PV module components (solar cells, strings, frames, etc.) in real time via camera sensors. Figure 2 shows examples using acA1920-25gm cameras, Basler AG, Hrensburg, Germany. The images used in this study were acquired at an indoor PV module

manufacturing facility. RGB images were captured using a machine vision camera fixedly installed on the production line for the EL and tapping processes. The camera installation location and shooting geometry were maintained throughout the data collection period to ensure consistent acquisition conditions. Furthermore, because the experiments were conducted in an indoor manufacturing environment, external lighting factors such as sunlight were excluded. To minimize image noise and subtle variations in acquisition conditions, preprocessing steps such as geometric alignment, region of interest extraction, and image resizing were applied to standardize the input data before model training and inference.




Figure 1. Machine vision framework for PV manufacturing [6].




Figure 2. Machine vision installation examples [6].

By utilizing a camera sensor, five types of EL images, including normal ones, are collected, as shown in Figure 3. Defective images include ‘crack’, ‘solder dark’, ‘dark spot’, and ‘dark area’. A crack is a defect in which microcracks occur inside the PV module; it is observed as a linear pattern along the crack path in the EL image [20]. A dark spot is a defect in which light emission is reduced locally in the form of a small dot at a specific location in the EL image [21]. A dark area is a defect in which light emission is reduced over a wide area in a specific region [22]. Solder dark is observed as a rectangular dark area in the EL image when current transmission becomes unstable due to a defect in the solder layer at the connection [23].




Figure 3. Defect types and normal images of PV modules are considered in this study.

In this study, criteria considering morphology and location were applied to identify the defect types in Figure 3. A crack was determined when a continuous linear luminescence reduction was observed within the module and either crossed the cell region or extended beyond a certain length. A dark spot was determined when a localized region inside the module showed a pronounced luminescence reduction compared to normal areas, and when the candidate region occupied a relatively small area within the region of interest (ROI) with a sufficient drop in mean intensity. In contrast, a dark area was determined when a luminescence reduction was similarly observed but appeared as a relatively large and contiguous region within the ROI, and it was distinguished from a dark spot using an area-based criterion. Solder dark was determined when a rectangular or band-shaped luminescence reduction pattern was observed around the busbar interconnection region.

In Figure 1, PV images collected via camera sensors during the EL operation are transmitted to the data collection module and stored in a database. The stored image data is used for real-time defect detection and judgment, subsequent AI algorithm training, and product quality management. During this process, PV module manufacturing environment information (processing parameters and PLC control information) provided by the manufacturing equipment is also stored in the database. Image data acquired through machine vision is transmitted to the data collection module in real time via the Internet of Things (IoT). Structured and unstructured data are standardized and classified into specific analysis datasets, and the stored processing and image data are transmitted to the machine vision platform. During this process, the quality prediction and defect identification module uses AI to extract defective products. These modules analyze images to detect defective products that exhibit different patterns from existing products and predict potential defects [6]. In this study, we aim to improve the detection performance of defective PV modules by utilizing EfficientNet-B0, an existing machine vision technology.

2.2. Defective PV Module Detection Using EfficientNet-B0 with Machine Vision

In this study, EfficientNet-B0, a basic model of the EfficientNet series, is selected as the CNN architecture for PV module EL image classification. EfficientNet-B0 is one of the CNN-based approaches proposed by Tan et al. (2019) [13]. It uses a compound model scaling strategy that simultaneously expands network depth, channel width, and input resolution, and is designed with a lightweight structure that includes mobile inverted bottleneck convolution (MBConv) and squeeze-and-excitation (SE) modules. This structure offers the advantage of achieving high representational power with fewer parameters and fewer floating-point operations (FLOPs) compared to existing CNNs such as the Visual Geometry Group (VGG), the residual network (ResNet), and GoogLeNet. In EL images of PV modules, small defects such as microcracks and dark spots directly impact classification accuracy, making high-resolution image input essential. EfficientNet-B0 is designed to scale well across depth, width, and resolution, effectively limiting computational load

and model size while maintaining high-resolution input, which makes it suitable for this inspection environment [24]. Furthermore, EL images present a class imbalance problem, with a relatively small number of defective images compared to normal images, and data distribution varies depending on shooting conditions. In this context, several studies have applied EfficientNet-B0 to detect defective PV modules. Liu et al. (2024) [25] demonstrated 97.81% accuracy on the PVEL dataset by utilizing the contrast-limited adaptive histogram equalization (CLAHE) algorithm and the EfficientNet-B0 backbone for EL image-based PV cell defect detection. Furthermore, Joshua et al. (2025) [26] compared VGG16, ResNet50, DenseNet121, and EfficientNet-B0 for solar panel defect and status classification, and found that EfficientNet-B0 achieved 80.0% accuracy while also offering a small number of parameters and fast inference time, making it an advantageous model for an edge computing environment. This study utilizes ImageNet pretrained weights as initial values and employs data augmentation and hierarchical segmentation strategies in parallel to stabilize representation learning for defective classes and mitigate overfitting [27]. Compared with other backbones such as ResNet-50, EfficientNet-B0 has a smaller model size, making it easy to deploy on industrial PCs or GPUs. For these reasons, EfficientNet-B0 is considered a suitable model for application to actual PV module EL inspection equipment and production lines [28].

Figure 4 shows the defective PV module detection structure based on EfficientNet-B0, which consists of one convolutional block and seven MBConv blocks. First, the PV module image is converted to an input size of $224 \times 224 \times 3$. Then, features are extracted from the image through convolution operations. The initial 3×3 convolution layer extracts low-level structures such as busbars, fingers, and cell boundaries, as well as local brightness variations. By sequentially passing the image through a total of 16 MBConv blocks, defect patterns such as microcracks and dark spots that appear at various shooting angles are gradually converted into high-level feature representations. In the final step, global average pooling (GAP) is applied to obtain a 1280-dimensional feature vector representing a single module. Dropout is used to mitigate overfitting. The final fully connected layer (Dense) is then used to output the probability that the input EL image is from a normal or defective module. The output layer uses the sigmoid activation function in Equation (1) to convert the final linear output value into a probability value between 0 and 1.

$$y = \frac{1}{1 + e^{-x}} \quad (1)$$

where y is a sigmoid function, and x is a linear output of the final fully connected layer (Dense). In this model, when x is input, Equation (1) calculates the estimated probability that the EL image is a defective module. The closer y is to 1, the higher the probability that the input EL image is defective. For a threshold value (Threshold, θ) set in advance, if y is less than or equal to θ , it is classified as normal, and if y is larger than θ , it is classified as defective. Since the boundary between normal and defective classification changes when θ is adjusted, it is important to set an appropriate θ according to an acceptable level.

2.2.1. Compound Model Scaling

EfficientNet-B0 uses compound model scaling, which adjusts the network depth d , width w , and resolution r together with a single scale factor ϕ and constants α , β , and γ . The scaling relationship among the three dimensions is as shown in Equation (2).

$$d = \alpha^\phi, w = \beta^\phi, r = \gamma^\phi \quad (2)$$

where α , β , and γ (≥ 1) are constants determined through small-scale exploration, and ϕ is a hyperparameter that adjusts the model scale according to the given computational

size. Since depth, width, and resolution increase by factors of α , β , and γ , respectively, as ϕ increases by 1, it is possible to control the three dimensions to grow together at a certain rate, unlike when only a single dimension is expanded.




Figure 4. Proposed machine vision approach with EfficientNet-B0 for defective photovoltaic module detection.

While increasing only depth allows for detailed model control as the model grows deeper, it unnecessarily increases parameters and computational complexity. Increasing only width can lead to inefficiencies in modeling and utilization due to the increased number of channels. Increasing only resolution increases the input size, leading to a sharp increase in computational complexity. Scaling only a single dimension tends to result in excessive growth in a specific axis. Composite model scaling, a strategy that simultaneously scales all three dimensions, offers better accuracy and efficiency within the same FLOPs range.

In the EL operation during PV module production, lowering the resolution too much can lead to the disappearance of microcracks and small dark spots, while increasing the model size excessively can lengthen production line processing times. Using composite model scaling, resolution, depth, and width can be adjusted simultaneously to secure sufficient expressiveness to handle high-resolution EL images while enabling real-time inspection with relatively low computational load.

2.2.2. Pointwise Convolution

Pointwise convolution (1×1) remixes the defects (microcracks, dark spots, etc.) extracted by each channel by depthwise convolution into the combinations required for defect classification, and reduces or expands them to the required number of channels. Figure 5 shows the operational process of linearly combining the input channel vectors at each spatial coordinate (h, w) using pointwise convolution to mix information between channels and convert the dimension. The input feature map in the form of $H \times W \times C_{in}$ is rearranged into a matrix of size $(H \cdot W) \times C_{in}$, and the 1×1 filter weights are rearranged into a $C_{in} \times C_{out}$ matrix. Afterwards, a general matrix multiplication is performed to calculate the $(H \cdot W) \times C_{out}$ output, which is then restored to the $(H \cdot W) \times C_{out}$ feature map. At this time, the multiply-accumulate (MAC) operation amount ($O_{pointwise}$) and the number of parameters ($P_{pointwise}$) based on one forward pass are as shown in Equations (3) and (4), respectively.

$$O_{pointwise} = H \cdot W \cdot C_{in} \cdot C_{out} \quad (3)$$

$$P_{pointwise} = C_{in} \cdot C_{out} \quad (4)$$

where H and W represent the height and width of the output feature map, C_{in} is the number of input channels, and C_{out} is the number of output channels.




Figure 5. Process of pointwise convolution.

In PV module images from the EL operation, each channel contains different filter results, such as vertical and horizontal edges, cracks in a specific direction, and partial brightness reduction (dark spots). Pointwise convolution linearly combines these channels to reconstruct defect-related information, such as cracks and dark spots, into a single, high-level representation, allowing the classifier to distinguish defect patterns.

2.2.3. Depthwise Convolution

Depthwise convolution is used to independently extract spatial patterns such as cracks and dark spots for each channel of the input feature map. Figure 6 shows the associative process of depthwise convolution, which applies different $K \times K$ filters to each channel of the input and performs only spatial convolution without mixing the channels. With the channel axis fixed, the $K \times K$ kernels for each channel slide over all spatial coordinates (h, w) to calculate the convolution, and the outputs for each channel are stacked again in the channel direction to form the output feature map in the form of $H \times W \times C_{in}$. The number of MAC operations ($O_{depthwise}$) and the number of parameters ($P_{depthwise}$) based on one forward pass are shown in Equations (5) and (6), respectively:

$$O_{depthwise} = K^2 \cdot H \cdot W \cdot C_{in} \quad (5)$$

$$P_{depthwise} = K^2 \cdot C_{in} \quad (6)$$

where K represents the size of one side of the kernel. Compared to the computational complexity of standard convolution with the same kernel size, depthwise convolution eliminates the C_{out} term, significantly reducing computational complexity and the number of parameters.

In PV module images from the EL operation, each channel can represent different spatial patterns, such as crack edges, dark spots, and internal cell surface patterns, depending on the direction. Depthwise convolution searches for these patterns across the module without mixing channels, efficiently extracting spatial defect structures such as cracks and dark spots from high-resolution images from the EL operation. This method has a lower computational complexity than standard convolution, making it advantageous

for maintaining computational speed even in environments where PV module images acquired from multiple shooting angles are massively processed.




Figure 6. Process of depthwise convolution.

3. Results

3.1. Experiment Scenario

Experiments are conducted on three aspects: PV module defect detection performance, optimal image capture angle selection, and sensitivity analysis of classification thresholds. First, the PV module defect detection performance of the proposed methodology is measured and compared with that of existing CNN-based algorithms (e.g., AlexNet, ResNet50, DenseNet121, and GoogLeNet) based on accuracy, precision, recall, and F1-score. Second, the performance sensitivity to the image capture angle is compared to ensure that the proposed Efficient-B0-based methodology achieves optimal performance. Finally, a sensitivity analysis is performed on the performance change of the proposed methodology depending on the setting of the classification threshold used to detect defective PV modules.

This experiment utilizes a PV module image dataset collected from an actual EL operation to evaluate PV module defect detection performance. Table 1 describes the collected images, which consist of 2772 images: 1512 normal PV modules and 1260 defective PV modules. This dataset contains cracks, dark spots, dark areas, and dark solder defects. Machine vision systems are installed at various angles to capture PV modules. In this study, to analyze the variation in defect detection performance with varying image capture angles, PV module images in the EL operation are collected at 10° intervals from 0° to 80° . Changes in the image capture angle can physically affect the measured EL signal and defect visibility. Therefore, this study evaluates performance variations across different capture angles. Although the emission characteristics of EL may vary with viewing angle, radiance can be treated as approximately invariant with respect to viewing direction. However, because practical PV modules have a multilayer structure including glass and encapsulant, changes in viewing angle can increase reflectance and decrease transmittance as a function of the incident angle [29]. In addition, variations in viewing angle can change the scattering components at the surface and interfaces, which may affect brightness uniformity and defect contrast in the recorded signal. Consequently, the EL intensity reaching the camera can vary with capture angle. Accordingly, this experiment was designed to quantify the effect of capture-angle changes on defect classification performance. The final dataset consists of 168 normal PV modules and 140 defective PV modules for each angle. To independently evaluate the impact of each angle, 60% of the images collected at each angle are used for training, 20% for validation, and 20% for testing. To compare performance,

CNN-based algorithm models are evaluated under the same data segmentation criteria, and defect detection performance is compared.

Table 1. Dataset composition by photographing angle for the PV module.

Photographing Angle (°)	Number of Images	
	Normal	Abnormal
0	168	140
10	168	140
20	168	140
30	168	140
40	168	140
50	168	140
60	168	140
70	168	140
80	168	140
Total	1512	1260

Accuracy, precision, recall, and F1-score are used to evaluate the defect detection performance of each algorithm. Equations (7)–(10) are the formulas for calculating each performance indicator:

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{FP} + \text{FN} + \text{TN}} \quad (7)$$

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \quad (8)$$

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (9)$$

$$\text{F1 score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (10)$$

where a defect is defined as positive and a normal module as negative. True positive (TP) refers to a case where an actual defective module is correctly classified as defective, and true negative (TN) refers to a case where an actual normal module is correctly classified as normal. False positive (FP) refers to a false detection where an actual normal module is incorrectly classified as defective, and false negative (FN) refers to a non-detection where an actual defective module is incorrectly classified as normal. Therefore, precision reflects the level of suppression of false positive (FP) detections, and recall reflects the ability to reduce non-detections (FN). F1-score is the harmonic mean of precision and recall, and it represents the judgment performance without the two indicators being biased to one side.

3.2. PV Module Defective Detection Performance

This experiment uses a baseline dataset of PV module images taken at 0° to compare the performance of each algorithm. In addition to EfficientNet-B0, experiments are conducted by changing the architecture to Alex Krizhevsky's Network (AlexNet), the residual network (ResNet50), the densely connected convolutional network (DenseNet121), and InceptionV1 (GoogLeNet). To compare performance changes attributable to architectural differences, all models were trained and evaluated under identical conditions. All models were trained for 100 epochs, and the weights corresponding to the highest validation accuracy during training were selected to compute the final results on the test set. In addition, the batch size was fixed at 8, the random seed was set to 42, the learning rate was set to 0.0001, and the Adam optimizer was used consistently across backbones to ensure a fair comparison. These architectures are widely used in defect classification research of


PV modules, and since they have different structural characteristics, they are suitable for comparing performance changes due to structural differences under the same data and conditions. AlexNet is a relatively early CNN architecture and achieved an accuracy of 99.3% in PV module status classification research [30]. ResNet50, a network architecture that utilizes residual connections to ensure learning stability, achieves an F1-score of 87.37% in PV panel crack detection based on EL images [31]. DenseNet121, designed for feature reuse through dense connectivity between layers, achieved an accuracy of 98% in a study of PV module defect identification [32]. GoogLeNet, a multi-scale architecture that extracts features from different receptive fields in parallel, achieved an accuracy of 94.6% in PV defect detection [33].

Table 2 describes the performance metrics for different CNN-based algorithms used to detect defects in PV modules. DenseNet121 achieves high performance across all metrics, with an accuracy of 0.8387, a precision of 0.8750, a recall of 0.7500, and an F1-score of 0.8077. This represents approximately 36.9% higher accuracy and 25.0% higher precision than AlexNet, which has the lowest classification performance. DenseNet121's higher accuracy compared to other algorithms is due to the fact that the defect patterns in the data include localized and subtle features, such as cracks and black spots. The direct connection structure of DenseNet50 allows features from previous layers to be accumulated and reused in subsequent layers, reducing information loss and maintaining a stable representation of micro-defects. On the other hand, AlexNet has limited depth, and the difficulty of hierarchically accumulating high-level features makes it difficult to isolate weak signals, such as micro-defects. The proposed EfficientNet-B0 shows overall stable performance with an accuracy of 0.7903, a precision of 0.8000, a recall of 0.7143, and an F1-score of 0.7547. Compared with DenseNet121, which had the highest performance, its accuracy, precision, and recall are 5.8%, 8.6%, and 4.8% lower, respectively, but it maintains relatively balanced performance without an excessive decrease in any specific indicator. EfficientNet-B0 can maintain stable performance even with low computational resources because it uses a convolutional structure that repeats the flow of channel expansion, feature extraction, and channel reduction in each block. Although its classification performance may be lower than that of structures that utilize accumulated features, such as DenseNet121, it can efficiently learn the features necessary for fault determination even with limited model capacity, thereby achieving stable performance. Meanwhile, compared with ResNet50, EfficientNet-B0 showed 2.0% and 7.4% lower accuracy and precision, respectively, but 5.3% higher recall. ResNet50's residual connections ensure stable learning even in structures with many layers. However, if the decision boundary is formed conservatively, the failure rate increases instead of reducing the FP, resulting in a lower recall and defective cases being classified as normal.

Figure 7 shows the original EL image and the corresponding activation map derived from the final convolutional feature maps of EfficientNet-B0. The dark area in Figure 7a was identified as a localized region exhibiting a clear reduction in luminance compared with surrounding normal regions, appearing as a large, contiguous area within the ROI. The heatmap in Figure 7b highlights regions with high model attention, where red indicates stronger activation. In other words, the color scale from 0 to 1 represents blue to red, with values closer to 1 indicating greater differences from normal PV panels. The results show that the model primarily focused on the surrounding areas rather than the actual dark area defect in Figure 7a. Therefore, the model relied more on visual patterns in normal regions than on defect-specific features, leading to a biased feature representation and ultimately resulting in misclassification.

Table 2. Detection performance using different types of algorithms.

Detection Algorithm	Evaluation Metrics			
	Accuracy	Precision	Recall	F1-Score
AlexNet	0.6129	0.7000	0.2500	0.3684
ResNet50	0.8065	0.8636	0.6786	0.7600
DenseNet121	0.8387	0.8750	0.7500	0.8077
GoogLeNet	0.7419	0.7143	0.7143	0.7143
Efficient-B0	0.7903	0.8000	0.7143	0.7547

**Figure 7.** Heatmaps for misclassified samples: (a) original EL image with defects, (b) heatmap of the final convolution layer.

In addition to Table 2, model complexity and computational efficiency are also important evaluation factors for PV module defect detection, which requires accurate detection within limited computational resources and a given time. Table 3 compares the computational resource requirements of each algorithm based on the number of parameters (Parameters), computational load (FLOPs), inference time, peak memory usage, and frames per second (FPS).

Table 3. Complexity and efficiency comparison of algorithms.

Detection Algorithm	Model Complexity and Efficiency				
	Parameters (M)	FLOPs (G)	Inference Time (Millisecond/Image)	Peak Memory Usage (MB)	Frames per Second (FPS)
AlexNet	71.9	2.5341	23.39	4048.06	42.8
ResNet50	23.6	7.7512	89.09	1336.62	11.2
DenseNet121	7.0	5.7004	78.13	1810.23	12.8
GoogLeNet	21.8	5.6934	82.29	1530.07	12.2
EfficientNet-B0	4.1	0.8008	47.85	1005.38	28.9

Parameters, the total number of weights determined by the model through learning, determine the model size and the memory requirements for storing weights during inference. Increasing the number of parameters generally improves expressive power, enabling the learning of complex features. However, using excessive parameters relative to the data size can lead to overfitting, which can degrade generalization performance. Furthermore, increasing model complexity tends to reduce interpretability. Floating-point operations per second (FLOPs) are used to measure the computational complexity required for a single inference. As FLOPs increase, the computational burden increases, potentially slowing down processing speed. Inference time is the actual processing time per image, which

comprehensively reflects the impact of FLOPs and memory access. Peak memory usage refers to the maximum amount of memory occupied during the inference process of the detection algorithm. Finally, frames per second (FPS) denotes the maximum number of images that can be processed per second and is used to evaluate the real-time performance of machine vision systems.

DenseNet121, which achieved the highest classification accuracy, had relatively high FLOPs (5.7004 G) and inference time (78.13 milliseconds/image). This is because, although the number of parameters itself is small, the dense connectivity structure of DenseNet121 increases memory accesses and feature map size as blocks are repeated and previous feature maps are repeatedly connected to later layers. While DenseNet121 achieves performance through feature reuse, the cumulative combination of features increases computational and memory overhead, leading to increased actual inference time. This can potentially lead to throughput degradation in environments with limited computational resources. This is reflected in its high peak memory usage of 1810.23 MB and relatively low throughput of 12.8 FPS. ResNet50 secures learning stability even in structures with many layers through residual connections, but its structural characteristics of repeated bottleneck blocks result in high FLOPs, reaching 7.7512 G. Consequently, ResNet50 recorded approximately 33% higher peak memory usage (1336.62 MB) compared with EfficientNet-B0 (1005.38 MB), as well as the lowest FPS (11.2) among the evaluated models. In contrast, EfficientNet-B0, utilized in this study, uses MBConv-based convolution to separate channel mixing and feature extraction. Designed to construct the same representation with fewer computations, it exhibits the lowest parameters (4.1 M) and FLOPs (0.8008 G), and its inference time (47.85 milliseconds/image) is also among the lowest among the compared models. Notably, EfficientNet-B0 demonstrates superior computational efficiency, with a 41.4% reduction in parameters, an 86.0% reduction in FLOPs, and a 38.8% reduction in inference time compared with DenseNet121. Moreover, peak memory usage is reduced by 44.4%, while FPS is approximately 2.3 times higher. PV module production lines are expected to continue increasing unit production volume, leading to an increase in throughput required throughout the manufacturing process [34]. Consequently, detection algorithms applied to inline inspection systems must be satisfied with increasing deployment latency constraints.

Tables 3 and 4 show that while EfficientNet-B0 exhibits lower classification accuracy than DenseNet121 and ResNet50, considering its defect detection performance, model complexity, and efficiency, it can secure throughput in limited computational resources, making it a suitable algorithm for production line applications.

Table 4. Detection performance under different rotation angles.

PV Module Rotation Angle (°)	Evaluation Metrics			
	Accuracy	Precision	Recall	F1-Score
0	0.7903	0.8000	0.7143	0.7547
10	0.8710	1.0000	0.7143	0.8333
20	0.8871	0.9565	0.7857	0.8627
30	0.7742	0.7500	0.7500	0.7500
40	0.8387	0.9500	0.6786	0.7917
50	0.8226	0.7931	0.8214	0.8070
60	0.8065	0.7667	0.8214	0.7931
70	0.9032	0.9231	0.8571	0.8889
80	0.7903	0.8947	0.6071	0.7234

3.3. Selecting the Optimal Image Capturing Angle

In Section 3.2, EfficientNet-B0 is considered the most suitable model when considering both classification performance and computational resource requirements. Accordingly, this

experiment is conducted based on EfficientNet-B0, varying the angle of the PV module, and analyzing the impact of these angle conditions on EL-based defect detection performance. Through experiments at various angles, we aim to identify the conditions that maximize performance in PV module defect detection.

Table 4 describes the binary classification performance metrics based on EfficientNet-B0 for EL images captured while rotating the PV module from 0° to 80° in 10° increments. The angle of capture is set to the relative angle between the camera lens and the PV module. The dataset for each angle is split into 60% training, 20% validation, and 20% testing to independently evaluate the impact of angle conditions. The training conditions are fixed at a batch size of 8 and 100 epochs. During the testing phase, the sigmoid output probability is used as a threshold of 0.5 to determine whether the image is normal or defective.

Comparing performance by angle in Table 4, the condition with the highest accuracy was 70°, with an accuracy of 0.9032, precision of 0.9231, recall of 0.8571, and F1-score of 0.8889. This demonstrates a stable defect detection system, maintaining a high level of precision to suppress false positives while ensuring sufficient recall for identifying defects as defects. Conversely, the condition with the lowest accuracy was 30°, with a 16.7% higher accuracy rate than that at 70°, demonstrating a clear difference in classification performance depending on the shooting angle.

This performance difference is due to the fact that changing the PV module rotation angle changes the area per pixel on the module surface, altering the contrast and boundary sharpness of defects. Because PV modules have a multilayer structure, changes in viewing angle can modify interface reflection and scattering contributions, which may alter the effective EL intensity distribution and defect-to-background contrast recorded by the camera. In addition, oblique viewing changes the projective geometry and pixel-to-surface correspondence, which can affect the apparent sharpness of fine defect boundaries. These combined radiometric and geometric effects provide a physical rationale for why defect patterns can be more separable at 70° than at 0° in this inspection geometry. Therefore, in machine vision-based PV module inspection, optimizing the shooting angle can impact defect detection performance. Defect detection on production lines requires stable judgment performance that suppresses unnecessary false positives while ensuring no defects are missed. Furthermore, performance indicators can vary depending on the shooting angle, resulting in variable detection reliability even within the same inspection system. Determining the shooting conditions that ensure the most stable performance and incorporating these conditions into the inspection environment is crucial for maintaining high detection performance.


The test results under 70° conditions are summarized in Table 5 by dividing them into a confusion matrix. The confusion matrix can be used to determine what type of error occurred by classifying whether the actual label and the model prediction match as true positive (TP), true negative (TN), false positive (FP), and false negative (FN). In this study, a defect is defined as positive and a normal module as negative. Accordingly, TP is a case where an actual defective module is correctly classified as defective, and TN is a case where an actual normal module is correctly classified as normal. FP is a Type 1 error, which corresponds to a false detection in which an actual normal module is incorrectly classified as defective, and FN is a Type 2 error, which corresponds to a non-detection in which an actual defective module is incorrectly classified as normal.

Under the 70° condition in Table 5, TP is correctly classified as 32, TN as 24, FP as 2, and FN as 4. A Type I error (FP) is a false detection of a normal module as defective, which can lead to unnecessary additional testing, such as retesting. On the other hand, a Type II error (FN) is a false detection of a defective module as defective, which can increase the possibility of defect leakage.

Table 5. Confusion matrix at 70° rotation angle.

Category		Estimated Values	
		True	False
Observed values	True	32	2
	False	4	24

To quantitatively evaluate the classification performance at the optimal rotation angle of 70°, we analyzed the receiver operating characteristic (ROC) curve. Figure 8 shows the ROC curve derived under the 70° rotation angle condition utilizing EfficientNet-B0. The ROC curve shows a shape close to the upper left area, which means that it maintains a high true positive rate (TPR) even at a low false positive rate (FPR). This means that the model has a low probability of misclassifying normal modules while accurately detecting most defective modules. In addition, the area under the curve (AUC) value, which indicates the area under the ROC curve, was 0.9538, which means that the classification model has high classification performance regardless of the data ratio by class.

**Figure 8.** ROC curve of EfficientNet-B0.

3.4. Sensitivity Analysis of Classification Threshold

In this study, PV module defect detection uses binary classification, applying a threshold based on the sigmoid output probability of EfficientNet-B0 to determine whether a module is normal or defective. While a threshold of 0.5 is commonly used, the dataset in this study consists of imbalanced data with an uneven distribution of normal and defective data. Because the majority class dominates in imbalanced data, fixing the threshold to 0.5 could result in the model frequently predicting the majority class, potentially compromising defect detection performance. Therefore, this study conducts a sensitivity analysis to examine performance changes as the threshold is varied.

Figure 9 shows the performance indicator trends for each threshold, varying the threshold-old in 0.05 intervals for the sigmoid output probability of the trained EfficientNet-B0. Lower thresholds increase the proportion of defective models, leading to higher recall. However, higher false positives (FPs) lead to lower precision. Conversely, higher thresholds lead to more conservative defective model judgments, leading to higher precision. However, higher false negatives (FNs), which are cases where actual defective models are classified as

normal, lead to lower recall. Thresholds induce trade-offs between precision and recall, and in imbalanced data environments, this balance point influences optimal model performance. In this study, we define the optimal threshold (θ) as shown in Equation (11) to reduce the bias of a single metric and reflect the overall classification performance balance.

$$\text{Score}(\theta) = \frac{\text{Accuracy}(\theta) + \text{F1 score}(\theta)}{2} \quad (11)$$

where $\text{Accuracy}(\theta)$ represents the classification accuracy ratio, and $\text{F1 score}(\theta)$ reflects the balance between precision and recall, representing the overall performance of defect detection.




Figure 9. Performance indicator trends under different classification threshold values.

$\text{Score}(\theta)$ is defined based on accuracy and the F1-score for θ selection, and θ is chosen at the operating point that maximizes $\text{Score}(\theta)$. As the threshold directly shifts the decision boundary for the sigmoid probability, each curve shows different trends. Accuracy changes as the balance between FP and FN changes, depending on θ . In this dataset, accuracy is maximized when θ is around 0.35 because FP and FN are relatively balanced, and then it decreases as θ increases due to the increase in FN. In addition, when θ is 0, almost all samples are predicted as defective, so FN rarely occurs. As θ approaches 1.0, almost no samples are predicted as defective, so TP decreases, FN increases, and recall converges to 0. Precision generally tends to increase as θ increases because defect judgment becomes stricter. However, when θ is 1.0, the number of predicted defects becomes 0, causing the curve to drop sharply to 0. The F1-score is maximum at the point where the balance between precision and recall is the best, but as θ increases, the number of cases classified as defective decreases, so recall decreases and tends to approach 0. Therefore, at θ of 0.35, the accuracy is 0.9032, and the F1-score is 0.8966, demonstrating a higher F1-score than the commonly used threshold of 0.5. Therefore, a threshold of 0.35 is adopted as the optimal classification criterion.

4. Conclusions

This study selects EfficientNet-B0, the baseline model of the EfficientNet series, as the CNN architecture for classifying defective images of PV modules in the EL operation. EfficientNet-B0 employs a hybrid model scaling strategy that simultaneously expands network depth, channel width, and input resolution. It is designed with a lightweight

architecture that includes MBConv and SE modules. Compared with existing CNN-based algorithms (e.g., VGG, ResNet, and GoogLeNet), it achieves high expressiveness with fewer parameters and FLOPs. These advantages of EfficientNet-B0 allow it to effectively limit the computational load and model size while maintaining high-resolution inputs for detecting small defects, such as microcracks or black spots, on PV modules during the EL process. Furthermore, to address the class imbalance between normal and defective PV module images, this study uses ImageNet pretrained weights as initial values and applies data augmentation and hierarchical segmentation strategies in parallel to stabilize representation learning for defective classes and mitigate overfitting. Experimental results show that DenseNet121 achieves the highest classification accuracy (FLOPs (5.7004 G) and inference time (78.13 milliseconds per image)). However, the EfficientNet-B0 method proposed in this study shows superior computational efficiency with a 41.4% reduction in the number of parameters, an 86.0% reduction in FLOPs, and a 38.8% reduction in inference time compared with DenseNet121, while showing similar accuracy to DenseNet121 and ResNet50, making it more suitable for the EL process that produces a large number of PV modules. In particular, EfficientNet-B0 has the advantage of being easy to deploy on industrial PCs or GPUs due to its small model size, making it a suitable model for application to actual PV module EL inspection equipment and production lines that utilize machine vision.

However, this study has several limitations. Because of the class imbalance between defect data and normal parts, more defect data collection is necessary to overcome this. Although this study attempted to stabilize representation learning for defect classes by applying data augmentation and hierarchical segmentation strategies in parallel, continued efforts are needed to collect more data, as higher-performance models are possible with more data. In addition, since cross-validation using external datasets was not conducted, further verification is required to establish the proposed approach as a generalized algorithm. Accordingly, future work will focus on collecting additional external datasets and deploying the system across different PV module manufacturing facilities to enhance generalizability. Furthermore, this study simplified the defect categories by considering multiple defect types as a single defective class in a binary classification setting. Future work will investigate multi-class defect classification, allowing more precise identification of defect types.

Author Contributions: Conceptualization, M.S., I.-B.L. and S.K.; methodology, M.S., I.-B.L., J.S. and S.K.; software, M.S., J.S. and S.K.; validation, M.S., J.S. and S.K.; formal analysis, M.S., J.S. and S.K.; investigation, M.S., J.S. and S.K.; resources, I.-B.L. and S.K.; writing—original draft, M.S., I.-B.L., J.S., and S.K.; writing—review and editing, M.S., I.-B.L., J.S. and S.K.; visualization, M.S. and J.S.; funding acquisition, S.K.; supervision, S.K. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by a Basic Science Research Program through the National Research Foundation of Korea (NRF), funded by the Ministry of Education (No. RS-2023-00239448).

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The original contributions presented in the study are included in the article; further inquiries can be directed to the corresponding author.

Acknowledgments: The authors gratefully acknowledge the support of the National Research Foundation of Korea (NRF) and the Ministry of Education. The views expressed in this paper are solely those of the authors and do not represent the opinions of the funding agency.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Malamas, E.N.; Petrakis, E.G.; Zervakis, M.; Petit, L.; Legat, J.-D. A survey on industrial vision systems, applications and tools. *Image Vis. Comput.* **2003**, *21*, 171–188. [CrossRef]
2. Reichenstein, T.; Raffin, T.; Sand, C.; Franke, J. Implementation of machine vision based quality inspection in production: An approach for the accelerated execution of case studies. *Procedia CIRP* **2022**, *112*, 596–601. [CrossRef]
3. Ren, Z.; Fang, F.; Yan, N.; Wu, Y. State of the art in defect detection based on machine vision. *Int. J. Precis. Eng. Manuf. Green Technol.* **2022**, *9*, 661–691. [CrossRef]
4. Mordor Intelligence, Machine Vision Systems (MVS) Market Size and Share Analysis—Growth Trends and Forecast (2026–2030). Available online: <https://www.mordorintelligence.com/industry-reports/machine-vision-systems-market> (accessed on 30 October 2025).
5. Gray, J.L. The physics of the solar cell. *Handb. Photovolt. Sci. Eng.* **2011**, *2*, 82–128.
6. Lee, I.-B.; Kim, Y.; Kim, S. A Data-Driven Machine Vision Framework for Quality Management in Photovoltaic Module Manufacturing. *Machines* **2025**, *13*, 285. [CrossRef]
7. International Energy Agency (IEA). Energy Technology Perspectives 2024. Available online: <https://www.iea.org/reports/energy-technology-perspectives-2024> (accessed on 30 October 2025).
8. Pavlík, M.; Beňa, L.U.; Medved', D.; Čonka, Z.; Kolcun, M. Analysis and evaluation of photovoltaic cell defects and their impact on electricity generation. *Energies* **2023**, *16*, 2576. [CrossRef]
9. Benda, V.; Cerna, L. A note on limits and trends in PV cells and modules. *Appl. Sci.* **2022**, *12*, 3363. [CrossRef]
10. Ding, H.; Zhou, D.; Liu, G.; Zhou, P. Cost reduction or electricity penetration: Government R&D-induced PV development and future policy schemes. *Renew. Sustain. Energy Rev.* **2020**, *124*, 109752.
11. Wu, C.-Y.; Mathews, J.A. Knowledge flows in the solar photovoltaic industry: Insights from patenting by Taiwan, Korea, and China. *Res. Policy* **2012**, *41*, 524–540. [CrossRef]
12. Choi, H.; Anadón, L.D. The role of the complementary sector and its relationship with network formation and government policies in emerging sectors: The case of solar photovoltaics between 2001 and 2009. *Technol. Forecast. Soc. Change* **2014**, *82*, 80–94. [CrossRef]
13. Tan, M.; Le, Q. Efficientnet: Rethinking model scaling for convolutional neural networks. In Proceedings of the International Conference on Machine Learning (ICML), Long Beach, CA, USA, 9–15 June 2019.
14. Haney, J.; Burstein, A. PV System Operations and Maintenance Fundamentals. Solar America Board for Codes and Standards. 2013. Available online: <https://www.secondenergy.me/wp-content/uploads/2021/03/SolarABCs-35-2013-3.pdf> (accessed on 20 January 2026).
15. Munoz, M.; Alonso-García, M.d.C.; Vela, N.; Chenlo, F. Early degradation of silicon PV modules and guaranty conditions. *Sol. Energy* **2011**, *85*, 2264–2274. [CrossRef]
16. Borgers, T.; Voroshazi, E.; Govaerts, J.; Szlufcik, J.; Poortmans, J. Multi-wire interconnection technologies weaving the way for back contact and bifacial PV modules. In Proceedings of the 2016 IEEE 43rd Photovoltaic Specialists Conference (PVSC), Portland, OR, USA, 5–10 June 2016.
17. Al-Refaie, A. Optimizing the performance of the tapping and stringing process for photovoltaic panel production. *Int. J. Manag. Sci. Eng. Manag.* **2015**, *10*, 165–175. [CrossRef]
18. Hassan, S.; Dhimish, M. Broad-scale Electroluminescence analysis of 5 million+ photovoltaic cells for defect detection and degradation assessment. *Renew. Energy* **2024**, *237*, 121868. [CrossRef]
19. AbouJieb, Y.; Hossain, E. Solar System Components. In *Photovoltaic Systems*; Springer: Cham, Switzerland, 2022; pp. 95–192.
20. Pratt, L.; Govender, D.; Klein, R. Defect detection and quantification in electroluminescence images of solar PV modules using U-net semantic segmentation. *Renew. Energy* **2021**, *178*, 1211–1222. [CrossRef]
21. Acikgoz, H.; Korkmaz, D.; Budak, U. Photovoltaic cell defect classification based on integration of residual-inception network and spatial pyramid pooling in electroluminescence images. *Expert Syst. Appl.* **2023**, *229*, 120546. [CrossRef]
22. Jia, Y.; Chen, G.; Zhao, L. Defect detection of photovoltaic modules based on improved VarifocalNet. *Sci. Rep.* **2024**, *14*, 15170. [CrossRef]
23. Park, S.; Han, C. Analysis of EL images on Si solar module under thermal cycling. *J. Mech. Sci. Technol.* **2022**, *36*, 3429–3436. [CrossRef]
24. Raza, R.; Zulfiqar, F.; Khan, M.O.; Arif, M.; Alvi, A.; Iftikhar, M.A.; Alam, T. Lung-EffNet: Lung cancer classification using EfficientNet from CT-scan images. *Eng. Appl. Artif. Intell.* **2023**, *126*, 106902. [CrossRef]
25. Liu, Q.; Liu, M.; Wang, C.; Wu, Q.J. An efficient CNN-based detector for photovoltaic module cells defect detection in electroluminescence images. *Sol. Energy* **2024**, *267*, 112245. [CrossRef]
26. Joshua, S.R.; Palilingan, K.Y.; Lengkong, S.P.; Park, S. Deep Learning-Driven Solar Fault Detection in Solar-Hydrogen AIoT Systems: Implementing CNN VGG16, ResNet-50, DenseNet121, and EfficientNetB0 in a University-Based Framework. *Hydrogen* **2025**, *7*, 1. [CrossRef]

27. Tiwari, K.K.; Singh, A.; Kumar, S. A Comprehensive Analysis of CNN-Based Deep Learning Models: Evaluating the Impact of Transfer Learning on Model Accuracy. In Proceedings of the 2025 2nd International Conference on Computational Intelligence, Communication Technology and Networking (CICTN), Ghaziabad, India, 6–7 February 2025.
28. Ali, H.; Shifa, N.; Benlamri, R.; Farooque, A.A.; Yaqub, R. A fine tuned EfficientNet-B0 convolutional neural network for accurate and efficient classification of apple leaf diseases. *Sci. Rep.* **2025**, *15*, 25732. [CrossRef] [PubMed]
29. Mochizuki, T.; Kim, C.; Yoshita, M.; Mitchell, J.; Lin, Z.; Chen, S.; Takato, H.; Kanemitsu, Y.; Akiyama, H. Solar-cell radiance standard for absolute electroluminescence measurements and open-circuit voltage mapping of silicon solar modules. *J. Appl. Phys.* **2016**, *119*, 034501. [CrossRef]
30. Aman, R.; Rizwan, M.; Kumar, A. Fault classification using deep learning based model and impact of dust accumulation on solar photovoltaic modules. *Energy Sources Part A* **2023**, *45*, 4633–4651. [CrossRef]
31. Abdelsattar, M.; AbdelMoety, A.; Emad-Eldeen, A. ResNet-based image processing approach for precise detection of cracks in photovoltaic panels. *Sci. Rep.* **2025**, *15*, 24356. [CrossRef]
32. Lakshmi, P.S.; Rayudu, M.S.; Bapuji, K. IoT based Fault Detection in Dusty Solar Panels using Modified DenseNet121. In Proceedings of the 2024 Fourth International Conference on Advances in Electrical, Computing, Communication and Sustainable Technologies (ICAECT), Bhilai, India, 11–12 January 2024.
33. Al-Otum, H.M. Classification of anomalies in electroluminescence images of solar PV modules using CNN-based deep learning. *Sol. Energy* **2024**, *278*, 112803. [CrossRef]
34. Mechanical Engineering Industry Association, International Technology Roadmap for Photovoltaics. 2023. Available online: <https://www.qualenergia.it/wp-content/uploads/2024/06/ITRPV-15th-Edition-2024-2.pdf> (accessed on 6 February 2026).

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

Article

The Development of a Predictive Maintenance System for Gearboxes Through a Statistical Diagnostic Analysis of Lubricating Oil and Artificial Intelligence

Diego Rigolli ¹, Lorenzo Pompei ², Massimo Manfredini ², Massimiliano Vignoli ², Vincenzo La Battaglia ¹ and Alessandro Giorgetti ^{1,*}

¹ Department of Industrial, Electronic and Mechanical Engineering, Roma Tre University, 00146 Rome, Italy; die.rigolli@stud.uniroma3.it (D.R.); vincenzo.labattaglia@uniroma3.it (V.L.B.)

² Bonfiglioli Spa, 40012 Calderara di Reno, Italy; lorenzo.pompei@bonfiglioli.com (L.P.); massimo.manfredini@bonfiglioli.com (M.M.); massimiliano.vignoli@bonfiglioli.com (M.V.)

* Correspondence: alessandro.giorgetti@uniroma3.it

Abstract

This paper addressed the problem of oil diagnostics lubricants applied to the predictive maintenance of industrial gearboxes, proposing the integration of an artificial intelligence (AI) system into the process analysis. The main objective was to overcome the critical issues of the traditional method, characterized by long analysis times and a marked dependence on the subjective interpretation of operators. The method includes a detailed statistical analysis of the common ways to assess the condition of lubricants, such as optical emission spectroscopy, particle counting, measuring viscosity and density, and Fourier-transform infrared spectroscopy (FT-IR). These methods are then combined with an artificial intelligence model. Tested on commercial gearbox data, the proposed approach demonstrates agreement between IA and expert evaluation. The application has shown that it can effectively support diagnoses, reduce processing time by 60%, and minimize human errors. It also improves knowledge sharing through an increase in the stability and repetitiveness of diagnoses and promotes consistency and clarity in reporting.

Keywords: predictive maintenance; industrial gearboxes; artificial intelligence; condition monitoring; oil diagnostics; lubricant analysis

1. Introduction

The reliability and efficiency of industrial mechanical systems depend to a large extent on the maintenance strategies adopted. In recent years, predictive maintenance has taken on a key role in plant management, helping to reduce downtime, optimize resource allocations, and improve operational safety and continuity [1,2]. Lubricating oils play a key role in this context, as they perform many fundamental functions, such as friction reduction, heat dissipation, contamination control, and corrosion protection [3–5]. Their condition is often indicative of the health status of equipment: through a targeted analysis, it is possible to identify and potentially prevent anomalies related to wear, overheating, or contamination, which could compromise the performance and safety of machinery.

Condition monitoring techniques, as defined by ISO standards [6], involve the continuous or periodic assessment of equipment or systems to identify potential issues or predict failures. As mentioned by Wakiru et al. [2], among the main condition monitoring techniques are vibration analysis, acoustic emission thermography and ultrasound and lubricant analysis.

Lubricating oil analysis is a key element of condition-based maintenance, an approach that involves performing interventions on the basis of the actual state of equipment rather than on fixed maintenance intervals [7]. Oil analysis can provide information about machine wear conditions, lubricant contamination and lubricant status. The presence and quantity of debris and particles from worn components and contamination can provide valuable information on useful service life consumption and issues affecting performance and reliability. By analyzing and monitoring data over time, it could be possible to schedule maintenance before a critical failure occurs.

In the recent literature, lubricating oil has been increasingly recognized not only as a consumer product but also as a diagnostic medium capable of capturing transient phenomena within gearboxes and other critical components. As reported by Mobley [8], oil analysis can be considered one of the most powerful techniques for the early detection of faults in mechanical systems, offering the possibility to detect problems well before traditional vibration or acoustic techniques become effective. In addition, research by Asadi et al. [9] highlighted how the systematic monitoring of physical and chemical properties of oil such as viscosity, particle number, and additive depletion can significantly increase the accuracy of health assessments in industrial gear units.

Integrating oil analysis into structured diagnostic workflows allows for a broader understanding of system behavior. Through historical trend analysis and cross-referencing with other condition monitoring data, lubricating oil analysis can serve as a keystone for data-driven maintenance decision frameworks. In this context, the standardization of report formats, interpretation guidelines and data acquisition protocols remains a critical area of interest to ensure consistency and comparability across systems and industries. The development of artificial intelligence systems that can learn from these structured workflows could further improve diagnostic capabilities, especially in environments with a limited availability of experts or high need to provide reports.

Currently, this analysis is usually performed using semi-automatic methods that need experienced operators to evaluate the results and data collected from laboratory measurements. This method is well defined but is characterized by the long time needed to make diagnoses, the possibility of human error, strong dependence on the experience of the operator, and limitations in complex connections detectable. These limitations can lead to insufficient capability in the number of analyses that can be performed and consequently to limited coverage on the equipment installed, usually driven by the presence of specific risk factors or failure detection. This makes the large-scale deployment of condition-based maintenance approaches expensive and difficult and the diagnostic process less effective overall [10]. The limitation in condition-based maintenance can be overcome by automating analysis and improving diagnostic reliability in line with what is happening with the introduction of AI in some fields [11–13]. It has been shown that artificial intelligence can be used to analyze large amounts of data from different instruments and heterogeneous sources, which historically have been extremely complex to sustain in real industrial applications [14,15]. AI is, in particular, interesting regarding its ability to provide diagnoses that can be comparable to experienced expert evaluations and is an open and highly focused issue in the current scientific literature [16–20].

In recent years, numerous studies have highlighted the potential of AI for predictive maintenance, with applications ranging from vibration analysis to machine learning model-based diagnostics [21–25].

However, the use of AI for lubricating oil analysis is still an evolving field of research, with some open challenges, such as the need for reliable datasets, the definition of precise diagnostic thresholds and its integration with existing maintenance management systems.

In this scenario, the present study aims to develop an innovative AI-based approach for the predictive analysis of lubricating oils in industrial gearboxes. The aim is to provide reliable and timely diagnostics, improving the ability to detect wear-related anomalies and the contamination and degradation of the lubricant.

2. Materials and Methods

The framework adopted for the development of this study is illustrated in Figure 1. It was constructed starting from real-world datasets derived from condition monitoring reports of industrial gearboxes. These datasets include oil sample analyses encompassing physical and chemical properties—such as viscosity and density—as well as elemental concentrations expressed in parts per million (ppm). Each record is associated with relevant metadata, including lubricant type, total operating hours of both lubricant and machinery, and a corresponding diagnosis provided by an industry expert.




Figure 1. Analytical framework.

In particular, Step 1 focused on the identification of critical parameters for the evaluation of oil conditions, in alignment with current diagnostic methodologies. This was followed by Step 2, which entailed the definition of threshold limits to support diagnostic classification.

These limits were determined using the international standards ASTM D7720-11 [26] and further refined through statistical and change-based approaches, including boxplot analysis for outlier detection and normality testing via probability plots.

In Step 3, efforts were directed towards the creation of a unified database capable of integrating laboratory measurements with user analysis requests.

Finally, Step 4 involved the development of an artificial intelligence (AI) model to automate the diagnosis of oil samples. The model was trained on historical data validated by domain experts and was designed to detect patterns related to normal behavior, early warnings, and failure scenarios. The ultimate objective was to replicate expert-level diagnostic assessments and enhance predictive maintenance strategies with robust, data-driven insights. The following subsections analyze in depth the steps used to develop the case study.

2.1. Creation of Condition Monitoring Report Database

The initial and essential phase in the data processing process is pre-processing. In order to make raw data compatible with later analysis algorithms, it is first of all gathered, cleaned, integrated, and normalized. Proper pre-processing reduces errors, removes anomalies, and unifies information from different sources, thus creating a reliable and structured database [27].

In the context of our study, pre-processing has taken on particular importance as the historical data useful for diagnosis comes from two different databases. The integration of the two databases is generally considered to be a component of pre-processing data, specifically in the data integration phase. In other words, while pre-processing includes all the cleaning, transformation, normalization, and integration operations necessary to prepare data for analysis, data integration focuses in particular on combining information from heterogeneous sources to create a coherent and comprehensive database [28,29].

The analysis of lubricant condition requires combining data from two databases: a Laboratory Information Management System (LIMS) and a Customer Portal (with application info). These sources provide essential information, such as analytical parameters, sampling information, and operational data. The integration of these sources is necessary to

obtain a comprehensive overview of the lubricant's operating and maintenance conditions. Power Query in Excel is used to effectively merge the information, allowing for import and transformation, error and duplication removal, and the merging of datasets using a common key (such as the encoding of the analysis). The coding of the analysis serves as a unique identifier for each analysis performed. Using this key, it was possible to merge the data, linking the information extracted from the two databases. In this way, each record was enriched with all the relevant information, guaranteeing coherence and completeness, which are indispensable for the successive stages of elaboration and diagnostics. One key aspect of the pre-processing phase is the continuous updating of the database.

2.2. Selection of Critical Oil Measurements

The analysis of a lubricating oil is structured around three main categories (Figure 2) [13]: fluid properties analysis, contamination analysis, and wear analysis.




Figure 2. Areas of study in oil analysis.

Fluid properties analysis focuses on the evaluation of chemical–physical properties and additives in the lubricant. Its main objectives are to estimate the remaining lubricant's service life and verify that the correct lubricant has been used. It can also detect mixtures with different oils. In this case, the data we have selected relate to the presence or quantity of phosphorus, zinc, calcium, boron, and magnesium, which represent the typical elements present in additives for reducing lubricating oils and can be detected by Rotating Disk Electrode (RDE) spectrometry. The density at 15 °C and viscosity at 40 °C and 100 °C, which are the most important properties of a lubricating oil, were also considered [30].

The contamination analysis evaluates the presence and level of contaminants of foreign substances that can enter the lubricant and machine from the environment. As largely described in the literature [31,32], contamination compromises the reliability of the machine and accelerates lubricant degradation. The indicators selected were silicon, sodium, and vanadium.

The wear analysis evaluates the debris in the form of microscopic particles generated by machine components that become worn. The lubricant, being inserted between the friction surfaces, is the first to collect this debris. Monitoring and analyzing these particles allow for the detection of ongoing wear, facilitating timely maintenance decisions. It is well known in the literature that the evaluation of wear in gearbox applications deals with extremely complex phenomena also related to effective contact surface and nonlinear dynamic effects [33,34]. This paper does not directly consider these issues because the approach to evaluating wear in internal gear components relies on measuring several basic parameters, including the presence of copper, iron, and chromium [30], which are related to the consequences of wear dynamics. Another indicator considered is the total amount of ferrous material FeC (ferrocheck) [35,36].

In addition, to obtain an accurate analysis, it is essential to correctly consider the technical data related to the oil and machine in a traceable way, including essential information such as the following:

- **Analysis ID:** This is required to identify each sample analyzed uniquely. Its function is essential to ensuring the traceability and correct association of data with the respective machines and operating conditions.
- **Gearbox series:** This identifies the series to which the reducer belongs. This information is particularly important because reducers from different series can be made of different materials and lead, at the time of analysis, to misinterpretations when, for example, wear is analyzed.
- **Gearbox serial number:** The serial number identifies the gearbox. Making an analogy with the world of medicine, we can say that it represents the name and surname of the reducer.
- **The date of sampling in which the oil was removed from the reducer:** This is one of the most important pieces of information because it allows us to study the history of gearboxes. In fact, oil analyses are often based on comparisons with historical data. An incorrect date can alter the trend of parameters over time, leading to wrong assessments on oil degradation or machine wear.
- **The type of oil:** Each lubricant has different chemical–physical characteristics that affect the performance of the oil itself and therefore of the machine. Without knowing the type of oil, it is difficult to determine whether the values are normal or if the correct oil has been used.
- **Hours of oil usage:** This indicates the time of use of the oil since the last significant change or top-up. This is crucial for assessing lubricant degradation and component wear. In fact, an abnormal increase in contaminants or wear metals compared to operating hours may indicate an imminent mechanical problem.
- **Machine hours:** This parameter indicates the total operating time of the machine.
- **Refueling:** This refers to the addition of oil in the system to compensate for leakage, evaporation, or natural consumption during operation. This parameter is important because it impacts oil quality and machinery condition diagnosis.

Once the data are collected and selected, the statistical analysis is carried out.

2.3. Definition of Diagnostic Limits

Limits, as well as alarms, play a key role in monitoring and managing the operating conditions of equipment and lubricants. First, they define operational standards, establishing acceptable values for critical parameters such as those seen in the previous paragraph. They also allow the early identification of anomalies by detecting when a parameter exceeds the maximum value or falls below the minimum, so that potentially critical situations can be identified at an early stage. This ability to detect deviations is particularly useful because it provides important support for maintenance decisions. When the data approach or exceed the set limits, it is possible to plan preventive interventions, thus preventing minor problems from evolving into costly failures. In parallel, the use of limits contributes to performance optimization: by keeping key parameters under control, you reduce stress on components and prolong the operating life of your systems, ensuring greater efficiency. An additional benefit is the reduction in costs and risks, as accurate monitoring enables timely intervention, minimizing major damage, machine downtime, and potential safety problems.

Figure 3 presents an illustrative diagram related to the study of threshold values. As the first step, the dataset was divided into three areas of interest—wear, lubricant, and contamination—following the approach outlined in the previous paragraph.

In particular, the statistical analysis on wear was carried out by grouping the gearboxes by series to obtain a population composed exclusively of elements with the same mechanical

behavior and made of the same materials (e.g., series made of steel rather than series with bronze crowns). The analysis was carried out on a total of 290 reducers, coming from 8 different gearbox series, each containing a minimum of 15 and a maximum of 150 elements.



Figure 3. The framework used to identify the limits.

Unlike wear metals, which result directly from machine operation, additives and contamination are intrinsic to the lubricant, and their concentration depends solely on the product formulation. For this reason, we have chosen to stratify the data by type of lubricant, ensuring a more meaningful comparison.

Simultaneously, the type of diagnostic limit was selected based on the nature of each analysis, in accordance with the methodology outlined by Fitch et al. [3]. Two main alarm strategies are employed: proactive alarms and predictive alarms. Predictive alarms signal abnormal machine conditions or the beginning of wear and failure. They are in line with the objectives of predictive maintenance, which is the early detection of symptoms of machine failure rather than root causes (proactive maintenance). Within the predictive domain, the following alarm techniques are used in oil analysis:

- **Change rate-based:** Rate-of-change alarms are generally set up to measure properties that vary as the oil is used. The rate is typically calculated as a function of machine hours or oil hours and corresponds to the slope of the line of ppm with respect to hours. For example, in the case of wear elements, the total ppm value will be divided by machine hours to give an indication of how many ppm/hour were generated by the system during its operation, while in the case of additives, the values will be divided by oil hours. This consideration is fundamental because some parameters may be below absolute levels yet still exceed the rate of increase, which could lead to early wear.
- **Statistical alarms:** Statistical limits are the most important and simple. This provides a sufficient amount of data to be representative of the machine and oil. The approach involves generating a standard mean and distribution from the available data and comparing these with the values from the sample under analysis. The ASTM D7720 standard [26] provides guidelines on how to set limits and their level of attention. In

particular, there are four levels of attention: white, green, yellow, and red. The document describes two main methods for establishing and updating alarm limits: statistical process control (SPC) and cumulative technique. SPC is based on statistics to assess control limits in data that follow a normal distribution. Limits are calculated based on standard deviation from the population average of the data: if the value is within a standard deviation from the average, it is considered normal (white). If it falls outside of a standard deviation from the mean but is within two deviation standards, is considered a warning (green). If the result exceeds two standard deviations, the value is considered in a critical alarm because it is higher or lower, depending on the cases, to 95% of the population. If the value exceeds three standard deviations, the situation is really critical, as the value exceeds about 99% percentile of the population. However, this approach is applicable to a normal distribution. In case of non-normal distribution, the standards propose to use the technique of cumulative distribution. Unlike SPC, cumulative distribution is not based on an assumed statistical model but sorts the data according to actual observed values, creating a percentage distribution. The resulting curve represents the cumulative probability, which indicates the percentage of data below a given value. This technique is suitable for data with non-Gaussian distributions, such as oil contamination, wear and chemical degradation. ASTM D7720 describes different types of cumulative distributions, depending on the nature of the data.

In contrast, the main goal of proactive alarms is to establish optimal benchmarks and keep the lubricant within these values to ensure proper machine operation. Proactive alarms are set to improve performance compared to the past, for example, by keeping the lubricant cleaner, drier, or at a lower temperature and ensuring stable operating conditions, preventing the lubricant from exceeding limits that have been previously identified as optimal for system operation. Within the proactive domain, the following types of alarms and limits are used:

- Goal-based limits: Targets apply to the control of parameters such as contamination, viscosity, and density. A goal-based limit is set with a specific objective, not necessarily related to the aging of the oil, but rather to maintaining optimal conditions to improve machine performance. In summary, the goal-based limit is an ideal value to maintain to maximize the performance and reliability of the machine.
- Aging limits: Another type of limit or proactive alarm concerns the progressive aging of a lubricant [3]. In fact, all lubricants are subject to aging, and therefore, the chemical–physical characteristics can deteriorate more or less slowly. For example, imagine having a lubricant for an industrial gearbox with a nominal viscosity of 320 cSt at 40 °C. To ensure optimal performance and reduce the risk of wear or malfunction, an acceptability range of about 10% around the nominal value is established. The objective of aging limits is to keep the lubricant in ideal conditions, intervening when significant variations occur that could compromise the protection and efficiency of the machinery. If the values are close to the limits, corrective actions such as an oil change, filtration, or more in-depth analysis can be taken to identify the causes of the change.

As the second step of the methodology, boxplots were generated to identify potential outliers—data points that deviate significantly from the rest of the dataset. The removal of these outliers had a substantial impact on the subsequent analysis.

Prior to data cleaning, the presence of extreme and anomalous values negatively affected the statistical representativeness of the parameters under examination, introducing distortions and reducing the reliability of alarm threshold identification. Retaining such anomalies could have compromised the definition of effective diagnostic limits, making them either too restrictive or excessively permissive depending on the outlier distribution.

At this point in Step three, we have chosen the method of analysis following what is described in ASTM D 7720, both for normal and non-normal distributions of data. In particular, the methodology used for the normally distributed data calculates the thresholds using means and standard deviations. For the non-normally distributed data, the methodology employs percentile-based evaluation.

2.4. Application of Artificial Intelligence Model

In this work, the AI system used for diagnostic purposes is based on the ChatGPT version o1 language model, in particular version o1, developed by OpenAI [37–39]. This model belongs to the family of Large Language Models (LLMs) and has been trained on a vast corpus of textual data, allowing human-like responses to be understood and generated in a wide range of technical and contextual areas. The diagnostic workflow implemented in this study uses natural language understanding and model generation capabilities and is described in Figure 4 and will be described in detail in Sections 2.4.1–2.4.3.




Figure 4. Diagnostic workflow.

The process begins with the structured input of technical data derived from oil analysis reports (e.g., elemental concentrations, viscosity, density, and other key parameters). These data are formatted as descriptive prompts and sent to the AI model, which is responsible for producing a diagnostic interpretation in natural language. The output generated is then compared with expert assessments provided by a human operator. The degree of agreement between AI-generated and expert diagnoses is used to assess model performance and consistency in different diagnostic sections. This approach does not involve training or fine-tuning the AI model to specific domain data; instead, it evaluates the diagnostic capabilities of a generic LLM when guided through prompts accurately tailored to the lubrication domain.

2.4.1. Pre-Processing

To automate the organization and preprocessing of raw oil analysis data, a Python 3.11.3 script was developed with the objective of ensuring data consistency, structure, and readiness for subsequent analysis and model training. The process is composed of three components: data validation and modeling, data grouping and transformation, and integration with a language model.

Data validation and modeling employs the Pydantic library to define strongly typed data models that validate the structure and format of each data entry. This validation step transforms rows from a raw DataFrame into structured Python objects, enhancing reliability and facilitating further processing. Three primary data classes were implemented:

- *Result*: Represents a single oil analysis entry, including fields such as ID, sampling date, lubricant and machine hours, diagnosis, and elemental concentrations.
- *GearboxInfo*: Encapsulates information specific to the gearbox, such as the serial number.
- *GearboxResult*: Aggregates analysis results for a given gearbox, linking a GearboxInfo instance to a chronologically ordered list of Result entries.

Data grouping and transformation has two core functions responsible for transforming raw datasets into a structured format. The first function groups analysis entries by gearbox and chronologically orders them. For each gearbox, a dictionary containing its structured analysis history is generated.

The second function (“get_relevant_data”) enables filtering by lubricant series and, optionally, by gearbox serial number. This allows the comparison of an individual gearbox against the rest of its series or the extraction of series-wide data excluding a specific gearbox. This approach supports both focused case studies and broader comparative evaluations across similar units.

The last component is focused on the integration with a language model and connects to the OpenAI-powered large language model (LLM), configured with a temperature parameter set to 1.0 (to obtain broader and more discursive diagnoses).

2.4.2. The Development of the Application for Automated Diagnostics

The application developed using the Streamlit framework enables interactive analysis of data related to wear, lubrication, and contamination of mechanical components. The system integrates several key functionalities, from managing and filtering complex datasets to automatically generating diagnostics through an LLM. The functions developed are data structuring, data loading and organization, prompt generation for the LLM, and automatic diagnostics and final output. In addition to this, an appropriate user interface was also developed.

Data Structuring is based on the use of three distinct lists of columns, each corresponding to a specific type of analysis, as follows:

- *columns_wear*: includes parameters related to component wear (e.g., operating hours, metal concentrations, machine ID, date, and diagnosis).

- *columns_oil*: contains characteristics of the lubricant, such as oil type, density, viscosity, and other chemical–physical properties.
- *columns_contamination*: collects data on the presence of contaminants (e.g., silicon, sodium, vanadium, etc.).

These lists allow for consistent and modular data selection, facilitating the merging of heterogeneous sources and ensuring complete information for subsequent analyses.

Data Loading and Organization is a function dedicated to handling the loading of the main dataset—containing complete analysis records—using pandas and the *read_excel* method. Simultaneously, reference limits are imported from separate Excel files, categorized as follows:

- *Wear*: absolute limits and normalized limits per operating hour (PPM/h);
- *Lubrication*: thresholds for additives, density, and viscosity—both absolute and normalized;
- *Contamination*: acceptable thresholds for external elements.

The data are organized into a dictionary that allows centralized and efficient access during the later processing stages.

Prompt generation for the LLM consists of a dedicated function that processes the data and generates the prompt to be submitted to the LLM for each analysis area (wear, oil, and contamination). Consequently, the following three functions have been developed:

- *Wear analysis*: The function filters only entries related to operating gearboxes, excluding new oil samples. The gearbox series is then identified, and using the *get_relevant_data* function, historical data (from other gearboxes of the same series) are separated from the current data. These are then formatted in Markdown and incorporated into the prompt alongside the relevant reference limits.
- *Oil analysis*: The system selects the row corresponding to the gearbox under examination and identifies the type of lubricant used. It then accesses a folder containing PDF files with technical specifications and new oil samples. These documents are read using the *PyMuPDF* library and integrated into the prompt together with current and historical data, as well as applicable limits.
- *Contamination analysis*: Similarly, the function selects data related to contamination and accesses the relevant PDF files. The data are enriched with the gearbox series and oil type, which are critical for an accurate evaluation.

In each case, the final prompt is built by concatenating all input text blocks with current data, historical trends, and corresponding thresholds.

With the *automatic diagnostics and final output* function, the prompt is submitted to the LLM, which returns an automatic assessment (diagnosis or recommendation). The results are stored in the Streamlit session state and can subsequently be combined with additional user input (e.g., the condition of the sample upon arrival, etc.) to generate a structured final report.

All the information is divided into thematic tabs (Wear, Oil, Contamination, and Final), each of which allows text input and the execution of the corresponding analysis function through an appropriate user interface developed in Streamlit.

2.4.3. Prompt Engineering

Prompt engineering is a crucial aspect in the interaction with advanced language models. It is the process of formulating optimized textual requests to obtain accurate and relevant answers from artificial intelligence models. A well-structured prompt allows the model to be guided towards a more precise understanding of the context, thus improving the quality of the output generated [40–42].

In the context of the predictive analysis of lubricating oils, prompt engineering plays a key role in automated diagnostics. Through accurate prompt definition, it is possible to

extract meaningful information, correlate historical and current data, identify anomalies, and make reliable assessments of the operating conditions of gearboxes and lubricants.

In order to ensure a detailed and structured analysis, multiple prompts were defined for each of the areas of investigation: wear, oil, contamination, and sample condition and final evaluation. Detailed prompt descriptions for each area are the following.

1. Wear analysis

For the wear analysis, three prompts were defined:

- **Ongoing Analysis Prompt:** This prompt instructs the model to generate a diagnosis on the presence of wear particles in the lubricant under analysis, using the chemical data available in the database. Specifically, the model must fill in the *Diagnosis* column by focusing solely on parameters related to wear debris, excluding all other lubricant aspects. To do so, it analyzes the concentrations of elements such as iron (Fe), aluminum (Al), and chromium (Cr), comparing them with historical data. After generating the diagnosis, the model must also provide an explanation for the observed wear level. In summary, this prompt explains how to read the data and how to produce a focused analysis.
- **Historical Data Prompt:** This prompt guides the model in formulating a diagnosis by analyzing similar gear units, with a particular focus on wear. The model must consider diagnoses previously made by an expert on reducers from the same group, reviewing historical analysis values to identify recurring trends or anomalies. However, it must ignore any parts of the historical diagnoses that are unrelated to the data available in the current analysis.
- **Limit Prompt:** This prompt instructs the model on how to interpret wear thresholds, both in terms of absolute values (ppm) and in terms of time-based variation (ppm/hour), to assess the condition of the gear unit.

2. Lubricant property analysis

Three prompts were also defined for lubricant characterization:

- **Ongoing Analysis Prompt:** This prompt guides the model in completing the *Diagnosis* column based on current database values. The focus is exclusively on additives, viscosity, and density. The chemical elements phosphorus (P), zinc (Zn), boron (B), calcium (Ca), and magnesium (Mg) represent the concentrations of additives in the lubricant. In addition to generating a diagnosis, the model must provide an explanation for the observed additive values.
- **Historical Data Prompt:** As in the previous section, this prompt helps the model build the diagnosis by leveraging historical expert evaluations of similar oils. Each dataset represents the oil test history of a specific gear unit. The model must analyze these series to identify trends and significant variations. It must also compare the chemical and physical properties in the technical specifications with those observed in new oil analyses. The model should ensure consistency with the lubricant's evolution over time.
- **Limit Prompt:** This prompt instructs the model to interpret both absolute and incremental limits for additives. For viscosity and density, it only needs to evaluate whether values fall within defined acceptable ranges. If the oil name is not found in the limits database, the model should not produce any assessment related to limits, thus avoiding misleading or inaccurate evaluations.

3. Contamination analysis

Three prompts were defined for contamination as well:

- **Ongoing Analysis Prompt:** This prompt instructs the model to assess oil contamination. It must evaluate the presence of elements such as silicon (Si), sodium (Na), and vanadium (V) and suggest potential causes for contamination.
- **Historical Data Prompt:** The model must focus on three key aspects: oil contamination, viscosity, and density. It should compare the oil composition with the analysis of fresh oil, as well as evaluate viscosity and density values against the specifications provided by the manufacturer.
- **Limit Prompt:** This prompt explains how the model should interpret values of silicon, sodium, and vanadium in the analysis, comparing them against reference thresholds. If the lubricant under analysis is not listed in the predefined limits, the model must refrain from drawing any conclusions based on those thresholds.

4. Sample conditions and final evaluation

The final section uses two additional prompts with specific functions:

- **Sample Reception Prompt:** This is written manually for each sample as reception conditions vary. It provides the model with details about the sample's arrival state.
- **Integration Prompt:** This prompt instructs the model to integrate the results obtained from the previous sections and propose any additional analyses that may help further clarify the equipment's condition.

3. Results and Discussion

In the field of artificial intelligence, testing is a fundamental tool for validating the quality and reliability of developed models. The primary objective is to verify whether the AI system is capable of producing results that align with expected standards, both in terms of performance (such as diagnostic accuracy and correctness) and usability, i.e., the actual applicability of the results in decision-making processes.

In other words, testing aims to determine the extent to which the model can replicate or even improve upon human expert evaluations while also identifying potential limitations and areas for improvement.

In the specific case of an AI system dedicated to lubricant analysis and predictive maintenance, testing plays a particularly critical role. It is essential to ensure that the diagnoses provided are not misleading, incomplete, or inconsistent with established professional expertise. Unlike other fields where standard metrics such as accuracy, precision, recall, or F1-score are typically used, the descriptive and interpretative nature of this type of output requires more qualitative evaluation methods.

For instance, if the model produces diagnostic comments rather than simple labels, it becomes necessary to assess not only the correctness of the numerical data but also the relevance, internal consistency, and completeness of the explanations provided.

In light of these considerations, the evaluation system is based on a common scale applied across all sections (wear, oil properties, contamination, and sample reception). This scale classifies diagnostic comments into four main categories and allows for the assessment of the reliability and accuracy of the diagnoses generated.

Numerical scores were assigned based on the degree of similarity between the AI-generated diagnoses and those written by the human expert, using a scale that ranges from a minimum to a maximum value—where higher scores indicate a greater alignment between the diagnoses.

The level chosen for the similarity analysis is the following:

- **Class 1—Irrelevant Comments (points 0):** Includes all diagnoses in which the model generates content completely unrelated to the topic. These comments are considered unusable and excluded from operational evaluation.

- **Class 2—Partially Incorrect Comments (points 1–3):** Includes diagnoses that contain significant inaccuracies, such as the misinterpretation of limit values or inconsistent assessments. While occasionally partially useful, these comments require thorough revision.
- **Class 3—Usable but Incomplete Comments (points 4–6):** Refers to diagnoses that are generally correct but lack completeness. The comments often omit important aspects that a human expert would consider, but they still remain usable as a basis for analysis.
- **Class 4—Fully Usable Comments (points 7–9):** Includes diagnoses that are fully consistent with the expert-validated master version, both in terms of content and argumentation. These comments can be directly used in the decision-making process.

Table 1 provides a summary table of the subdivisions, grades, and meanings of the classes. The scoring system's classification is based on identifying elements that are considered discriminating for evaluating the diagnosis's goodness. First of all, we considered the plausibility of the diagnosis in relation to the physical scenario of the system, meaning its consistency with reality. If this feature is not present in the diagnosis, it is rated inadequate with a score value of zero. The partial consistency leads to the second class; full consistency but a partially exhaustive analysis and the necessity of some clarification belong to the third class; and finally, ready-to-use diagnoses are collected in the fourth class. Within the three main classes of goodness-of-diagnosis assessment, three internal levels were identified to provide within-class granularity in the assessment of the individual diagnosis while still within the key characteristics of the class.

Table 1. Numerical scores for the evaluation of similarity between the AI-generated vs. expert diagnoses.

Scale	Synthetical Meaning	Class
0	The comment provided may lead to a misdiagnosis.	1—Unusable and misleading.
1-2-3	The comment may lead to a partially wrong diagnosis (e.g., incorrect processing limits).	2—Comment not usable.
4-5-6	The comment provided is consistent with the application but not exhaustive.	3—Usable after revision.
7-8-9	The comment provided is comparable to the master provided.	4—Usable as is.

To validate the diagnostic capability of AI, thirty analyses are randomly extracted from the original database as a partition. The remaining database is used for the limit definitions described in Section 2. Random selection was carried out without applying specific filters or stratification criteria in order to ensure an impartial and representative assessment of the model's performance under different operating conditions. Scores were assigned to each analysis based on a comparison between the diagnoses provided by the model and the diagnoses developed by the expert used as master measures.

In Table 2, some examples of diagnoses are listed. The diagnoses selected are all within the range between 7 and 9 in the provided by the model compared to the master diagnosis.

The qualitative comparison between the diagnoses generated by the artificial intelligence and those formulated by the expert analyst highlights both similarities and differences in the interpretation of lubricant oil data, particularly across the three categories analyzed: wear, lubricant, and contamination.

In diagnosis number 1, both assessments acknowledge the presence of ferromagnetic particles (Fe and Cu) and the need to monitor their trend. However, while the AI quantifies the phenomenon by referencing numerical thresholds, the expert simply indicates the presence of particles without specifying any threshold values.

In diagnosis number 2, both the AI and the analyst agree on the acceptability of the viscosity, but the AI detects an anomaly in the additive levels compared to the fresh lubricant.

In diagnosis number 3, both diagnoses identify a potential contamination with another lubricant, but the AI provides a more detailed analysis, highlighting abnormal levels of magnesium, barium, sodium, and vanadium.

Finally, for diagnosis 4, the AI and expert agree that the sampling method used is noncompliant and may have affected the sample's representativeness. Additionally, the AI suggests further analyses, such as FT-IR, that could be useful for the diagnosis.

In Table 3, the detailed scores assigned to the thirty analyses used for the assessment of AI vs. master diagnoses are listed.

Table 2. Examples of one of the diagnoses used for the evaluation of the AI diagnosis vs. master diagnosis.

Diagnoses #	AI Diagnosis	Master Diagnosis
1	The presence of ferromagnetic particulates (FeC) and spectrometric particulates (Fe and Cu) above threshold 1 but below threshold 2. It is recommended to continue sampling to monitor the trend.	Evidence of iron and ferromagnetic particulates.
2	Viscosity is acceptable for the specified grade. The presence of an abnormal additive compared to the new lubricant, and possible contamination with other lubricants. It is recommended to continue sampling to monitor the trend.	Viscosity is acceptable for the specified grade.
3	The sample has levels of magnesium, boron and barium compared to the reference lubricant, suggesting possible contamination with other lubricant or additive. In addition, sodium and vanadium values indicate potential counts. We recommend to check that the lubricant used is the one specified.	The fluid is probably contaminated with other lubricant.
4	The sampling method is different to the standard indicated. This could affect the significance of the sample being analyzed. The presence of ferromagnetic particles and wear metals above threshold 1 but below threshold 2 indicate the abnormal wear of the internal components. In addition, the high presence of additives such as magnesium and boron suggests possible contamination or mixing with other lubricant. It is recommended to check the lubricant used, check the condition of internal components and monitor the trend of metals in subsequent analyses. Additional analysis: FT-IR, Karl-Fischer, or analytical microscopy.	The sampling methods are different to the standard indicated. This could affect the significance of the sample being analyzed. The presence of ferrous and ferromagnetic spectrometric particulates is evident. The fluid is probably contaminated by other lubricant. Viscosity is acceptable for the specified grade. It is recommended to continue sampling to monitor the trend. Additional analysis: Karl-Fischer or analytical microscopy.

Table 3. AI diagnosis vs. master diagnosis agreement.

Diagnoses ID	Score			
	Section Wear	Section Lubricant	Section Contamination	Sample Conditions and Final Evaluation
1	6	8	9	8
2	7	9	8	5
3	7	7	7	6
4	8	6	7	5
5	7	8	6	7
6	2	2	5	2
7	8	6	7	8
8	7	7	6	6
9	8	9	7	8
10	9	8	9	7
11	6	7	8	7
12	8	6	6	8
13	8	7	6	8
14	6	7	7	8
15	9	8	8	8
16	8	8	8	7
17	8	8	8	7
18	8	8	8	7
19	8	8	5	5
20	5	8	5	8
21	6	6	6	6
22	8	9	9	8
23	8	8	9	8
24	8	8	3	3
25	8	8	8	8
26	8	8	8	5
27	8	8	6	8

Table 3. Cont.

Diagnoses ID	Score			
	Section Wear	Section Lubricant	Section Contamination	Sample Conditions and Final Evaluation
28	8	8	9	8
29	8	8	8	6
30	3	5	3	4

Figure 5 summarizes the level of agreement between IA and master diagnoses and shows that the majority of diagnostic assessments (more than 70 percent) from IA are consistent with the experienced operator at an operationally usable level. On the other hand, there are about a quarter of assessments that, while sufficiently consistent, require revision by the practitioner to be usable because they are partial or incomplete. Instead, the frequency of analysis that is unusable and inconsistent with the IA is slightly greater than 5 percent.




Figure 5. Overall summary of IA assessments without breakdown by sections analyzed.

The boxplots showing the degree of agreement between the assessments made by the AI model and those of the human operator for each of the four sections are shown in Figure 6. The figure illustrates the dispersion and consistency of the agreement within each section in addition to the average alignment. The average scores of agreements were calculated for each of the four analyzed sections and provide a measure of the model's performance in that specific area, allowing for the identification of potential areas for improvement.

These values serve as indicators of the diagnostic accuracy of the AI; however, a closer examination reveals performance variability across the different domains.

Specifically, the sections related to wear and lubricant yielded higher average scores compared to those concerning contamination and reception conditions, suggesting that the AI exhibited more robust performance in certain diagnostic areas. The lubricant section, with an average score of 7.3, most closely aligned with the evaluations of the human expert, indicating the model's strong capability in identifying and interpreting lubricant-related phenomena. A comparable performance was observed in the wear section, highlighting the AI's reliability in detecting signs of mechanical degradation within gear units.

In contrast, the contamination section showed a moderate decline in diagnostic accuracy, though it remained broadly consistent with expert assessments. The reception condition section recorded the lowest average score, potentially pointing to the need for the further optimization of the model in that specific area.

These findings suggest that while the AI system is capable of producing diagnostically sound outputs, targeted improvements may be required, particularly in the areas of contamination detection and the assessment of sampling conditions.




Figure 6. Boxplot of scores of the agreement level (AI diagnosis vs. master): wear section median value: 8, lubricant section median value: 8, contamination section median value: 7, final evaluation section median value: 7.

An essential component of the evaluation framework is the systematic integration of expert feedback, which plays a pivotal role in maintaining the model's diagnostic reliability. To support continuous improvement, the protocol includes the regular insertion of expert-generated reports in place of AI-generated ones. This practice not only enriches the training corpus with high-quality, domain-specific inputs but also mitigates the risk of diagnostic stagnation—a potential consequence of the AI's reliance on pre-trained knowledge, historical data, predefined thresholds, and technical specifications.

Incorporating expert diagnoses on a routine basis allows the model to expand its diagnostic repertoire and ensures ongoing alignment with operational needs and real-world complexities.

Finally, the adoption of the AI system has demonstrated several advantages. It has accelerated the diagnostic process, enabling more daily diagnostics. Additionally, it has improved operational resources, leading to better plant utilization and potential cost reduction in maintenance. The study of diagnostic thresholds has also contributed to the definition of more reliable diagnostic limits, allowing for a more precise assessment of lubricant health and gearbox operating conditions. Furthermore, the system provides advanced decision-making support, enabling the identification of emerging critical issues and facilitating the planning of targeted interventions based on historical trends. A key result of this study was a comparative analysis between the AI-generated responses and those provided by experts. This comparison confirmed the system's effectiveness, as the AI consistently produced results in alignment with expert's diagnoses. The results highlights the reliability and predictive capabilities of the AI-based approach, underscoring its potential to enhance maintenance strategies.

The last consideration comes from a time consumption perspective. AI enables faster analysis, which can increase the production capacity of experienced operators and reduce human errors and oversights.

In conclusion, the implementation of AI in diagnostic processes significantly enhances efficiency by reducing both diagnosis time consumption (considering an estimated 60% lower time between the time taken by the operator to make a diagnosis and write reports and that of the AI-based system) and associated costs, as it minimizes the need for extensive involvement from specialized experts during initial evaluations. Furthermore, the stability and repetitiveness of diagnoses can be progressively improved by increasing the dimensions of diagnostic databases, as AI systems base their outputs on established parameters, eliminating errors in limit assessments. Additionally, the integration of AI facilitates a standardized approach to writing reports by leveraging existing documentation, thereby promoting consistency and clarity in diagnostic reporting and knowledge sharing.

4. Conclusions

This paper presents an innovative AI-based approach for the diagnostics of lubricating oil that can be used for predictive analysis in industrial gearboxes with the ability to detect the wear-related anomalies, contamination, and degradation of the lubricant. The effectiveness of the approach proposed for the predictive diagnosis and analysis of lubricants, comparing AI-generated diagnoses, was evaluated by comparing them with the diagnoses of human experts. The aspects evaluated were wear, lubricant properties, contamination, and sample conditions. AI diagnoses showed good consistency with those of the experts, especially in the wear and lubricant sections (median value 8). The contamination and final evaluation sections show instead a lower median of 7, but, in any case, the large majority of the diagnoses (more than 75%) result in sufficient accuracy but with the described difference between the four analyzed diagnostic aspects. The use of AI permits the reduction in diagnostic time by ~60%, increasing throughput and efficiency, and can enable scalable diagnostics without a proportional increase in expert personnel. AI, in this context, can act as a diagnostic assistant, handling standard cases and flagging complex ones for expert review. Its function is to be a screening tool and decision support system, enhancing expert productivity. Furthermore, the approach can support standardization and traceability in maintenance reporting and knowledge management. Future works should deal with an increase in diagnostic databases to obtain fully reliable results from AI and deepen the role of the human expert to prevent stagnation and maintain alignment with real-world complexities, to identify the approach for updates of diagnostic thresholds and datasets, for the integration of new types of diagnoses and for the supervision of AI diagnostic capability.

Author Contributions: Conceptualization, D.R., L.P., M.M., V.L.B. and A.G.; methodology, D.R., L.P., M.M., M.V., V.L.B. and A.G.; software, M.V. and A.G.; validation, L.P. and M.M.; formal analysis, D.R. and V.L.B.; investigation, D.R. and V.L.B.; resources, L.P.; data curation, D.R., V.L.B. and A.G.; writing—original draft preparation, D.R., L.P., M.M., M.V., V.L.B. and A.G.; writing—review and editing, D.R., L.P., M.M., M.V., V.L.B. and A.G.; visualization, D.R., V.L.B. and A.G.; supervision, M.M. and A.G.; project administration, L.P. and M.M.; funding acquisition, M.M. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The data presented in this study are available on request from the corresponding author. Parts of the data are not publicly available due to privacy and industrial property restrictions.

Conflicts of Interest: Authors Lorenzo Pompei, Massimo Manfredini and Massimiliano Vignoli were employed by the company Bonfiglioli Spa. The remaining authors declare that this research was conducted in the absence of any commercial or financial relationships that could be construed as potential conflicts of interest.

References

1. Bhat, J.S.; Sonawane, B.U. Condition Monitoring of Worm Gearbox through Oil Analysis. In *Recent Trends in Engineering Design*; Parey, A., Kumar, R., Singh, M., Eds.; Lecture Notes in Mechanical Engineering; Springer: Singapore, 2021. [CrossRef]
2. Wakiru, J.M.; Pintelon, L.; Muchiri, P.N.; Chemweno, P.K. A review on lubricant condition monitoring information analysis for maintenance decision support. *Mech. Syst. Signal Process.* **2019**, *118*, 108–132. [CrossRef]
3. Fitch, J.; Troyer, D. *Oil Analysis Basics*, 2nd ed.; Noria Corporation: Tulsa, OK, USA, 2010.
4. del Álamo, J.R.; Duran, M.J.; Muñoz, F.J. Analysis of the Gearbox Oil Maintenance Procedures in Wind Energy. *Energies* **2020**, *13*, 3414. [CrossRef]
5. Yan, S.; Kong, Z.; Liu, H.; Li, B.; Fan, M.; Zhang, X. Oil Change Interval Evaluation of Gearbox Used in Heavy-Duty Truck E-Axle with Oil Analysis Data. *Lubricants* **2022**, *10*, 252. [CrossRef]
6. ISO 17359:2018; Condition Monitoring and Diagnostics of Machines—General Guidelines. ISO: Geneva, Switzerland, 2018.
7. Abbasi, T.; Lim, K.H.; Soomro, T.A.; Ismail, I.; Ali, A. Condition Based Maintenance of Oil and Gas Equipment: A Review. In Proceedings of the 2020 3rd International Conference on Computing, Mathematics and Engineering Technologies (iCoMET), Sukkur, Pakistan, 29–30 January 2020; pp. 1–9. [CrossRef]
8. Mobley, R.K. *An Introduction to Predictive Maintenance*, 2nd ed.; Butterworth-Heinemann: Oxford, UK, 2002.
9. Asadi, S.Y.; Schenck, T.; May, G. Integrated Oil Condition Monitoring and Predictive Maintenance Using Lubricant Analysis in Industrial Gearboxes. *J. Mach. Eng.* **2021**, *21*, 60–74.
10. Achouch, M.; Dimitrova, M.; Ziane, K.; Sattarpanah Karganroudi, S.; Dhouib, R.; Ibrahim, H.; Adda, M. On Predictive Maintenance in Industry 4.0: Overview, Models, and Challenges. *Appl. Sci.* **2022**, *12*, 8081. [CrossRef]
11. Bindra, S.; Jain, R. Artificial Intelligence in Medical Science: A Review. *Ir. J. Med. Sci. (1971-)* **2023**, *193*, 1419–1429. [CrossRef] [PubMed]
12. Weng, Y.; Wu, J.; Kelly, T.; Johnson, W. Comprehensive Overview of Artificial Intelligence Applications in Modern Industries. *arXiv* **2024**, arXiv:2409.13059. [CrossRef]
13. Singh, V.; Gangsar, P.; Porwal, R.; Atulkar, A. Artificial Intelligence Application in Fault Diagnostics of Rotating Industrial Machines: A State-of-the-Art Review. *J. Intell. Manuf.* **2021**, *34*, 931–960. [CrossRef]
14. Giorgetti, A.; Arcidiacono, G.; Ciappi, A.; Barbieri, R.; Citti, P. HNCr model following robust approach. *Qual. Reliab. Eng. Int.* **2018**, *34*, 1271–1288. [CrossRef]
15. Giorgetti, A.; Rolli, F.; La Battaglia, V.; Marini, S.; Arcidiacono, G. Axiomatic Design Using Multi-Criteria Decision Making for Material Selection in Mechanical Design: Application in Different Scenarios. In Proceedings of the 15th International Conference on Axiomatic Design 2023, Eindhoven, The Netherlands, 31 May–2 June 2023. [CrossRef]
16. Aliyu, R.; Mokhtar, A.A.; Hussin, H. Prognostic Health Management of Pumps Using Artificial Intelligence in the Oil and Gas Sector: A Review. *Appl. Sci.* **2022**, *12*, 11691. [CrossRef]
17. Habbal, A.; Ali, M.K.; Abuzaraida, M.A. Artificial Intelligence Trust, Risk and Security Management (AI TRiSM): Frameworks, Applications, Challenges and Future Research Directions. *Expert Syst. Appl.* **2023**, *240*, 122442. [CrossRef]
18. Sounderajah, V.; Ashrafian, H.; Rose, S.; Shah, N.H.; Ghassemi, M.; Golub, R.; Kahn, C.E.; Esteve, A.; Karthikesalingam, A.; Mateen, B.; et al. A Quality Assessment Tool for Artificial Intelligence-Centered Diagnostic Test Accuracy Studies: QUADAS-AI. *Nat. Med.* **2021**, *27*, 1663–1665. [CrossRef]
19. Hanna, M.; Pantanowitz, L.; Jackson, B.; Palmer, O.; Visweswaran, S.; Pantanowitz, J.; Deebajah, M.; Rashidi, H. Ethical and Bias Considerations in Artificial Intelligence (AI)/Machine Learning. *Mod. Pathol.* **2025**, *38*, 100686. [CrossRef]
20. Abdelwanis, M.; Alarafati, H.K.; Tammam, M.M.S.; Simsekler, M.C.E. Exploring the Risks of Automation Bias in Healthcare Artificial Intelligence Applications: A Bowtie Analysis. *J. Saf. Sci. Resil.* **2024**, *5*, 460–469. [CrossRef]
21. Cunha, B.Z.; Droz, C.; Zine, A.-M.; Foulard, S.; Ichchou, M. A Review of Machine Learning Methods Applied to Structural Dynamics and Vibroacoustic. *Mech. Syst. Signal Process.* **2023**, *200*, 110535. [CrossRef]
22. Mey, O.; Neudeck, W.; Schneider, A.; Enge-Rosenblatt, O. Machine Learning-Based Unbalance Detection of a Rotating Shaft Using Vibration Data. In Proceedings of the International Conference on Emerging Technologies and Factory Automation (ETFA), Vienna, Austria, 8–11 September 2020. [CrossRef]
23. Brito, L.C.; Susto, G.A.; Brito, J.N.; Duarte, M.A.V. Fault Diagnosis Using eXplainable AI: A Transfer Learning-Based Approach for Rotating Machinery Exploiting Augmented Synthetic Data. *Expert Syst. Appl.* **2023**, *232*, 120860. [CrossRef]
24. Cardoso, D.; Ferreira, L. Application of Predictive Maintenance Concepts Using Artificial Intelligence Tools. *Appl. Sci.* **2020**, *11*, 18. [CrossRef]
25. Scaife, A.D. Improve Predictive Maintenance through the Application of Artificial Intelligence: A Systematic Review. *Results Eng.* **2023**, *21*, 101645. [CrossRef]
26. ASTM D7720-11; Standard Guide for Statistically Evaluating Measurand Alarm Limits. ASTM International: West Conshohocken, PA, USA, 2011.
27. Han, M.K.; Jian, J.P. *Data Mining: Concepts and Techniques*, 3rd ed.; Morgan Kaufmann: Waltham, MA, USA, 2011.

28. Dasu, T.; Johnson, T. *Exploratory Data Mining and Data Cleaning*; John Wiley & Sons: Hoboken, NJ, USA, 2003.
29. Pyle, D. *Data Preparation for Data Mining*; Morgan Kaufmann: Waltham, MA, USA, 1999.
30. Fernández-Feal, M.; Fernández-Feal, M.; Sánchez-Fernández, L.; Pérez-Prado, J. Study of Metal Concentration in Lubricating Oil with Predictive Purposes. *Curr. J. Appl. Sci. Technol.* **2018**, *27*, 1–12. [CrossRef]
31. Mathura, S. *Lubrication Degradation Mechanisms*; CRC Press eBooks: Boca Raton, FL, USA, 2020. [CrossRef]
32. La Battaglia, V.; Mussi, V.; Marini, S.; Giorgetti, A. Investigation of Damage Caused by Chlorine-Contaminated Fuel in Standard Vehicle Components. *Eng. Proc.* **2025**, *85*, 8. [CrossRef]
33. Shi, J.; Han, C.; Chen, L.; Jin, W. Investigating wear-induced health-instability dynamics in spur gear systems with installation error. *Nonlinear Dyn.* **2025**, 1–22. [CrossRef]
34. Mao, Y.; Tong, J.; Chin, Z.Y.; Borghesani, P.; Peng, Z. Transmission-error- and vibration-based condition monitoring of gear wear with contaminated lubricant. *Wear* **2023**, *523*, 204760. [CrossRef]
35. Wang, Z.; Xue, X.; Yin, H.; Jiang, Z.; Li, Y. Research Progress on Monitoring and Separating Suspension Particles for Lubricating Oil. *Complexity* **2018**, *2018*, 9356451. [CrossRef]
36. Sondhiya, O.P.; Gupta, A.K. Wear Debris Analysis of Automotive Engine Lubricating Oil Using by Ferrography. *Int. J. Eng. Innov. Technol.* **2012**, *2*, 46–47.
37. Jaech, A.; Kalai, A.; Lerer, A.; Richardson, A.; El-Kishky, A.; Low, A.; Helyar, A.; Madry, A.; Beutel, A.; Carney, A. Openai o1 system card. *arXiv* **2024**, arXiv:2412.16720. [CrossRef]
38. Jegham, N.; Abdelatti, M.; Hendawi, A. Visual Reasoning Evaluation of Grok, Deepseek Janus, Gemini, Qwen, Mistral, and ChatGPT. *arXiv* **2025**, arXiv:2502.16428. [CrossRef]
39. Bagnato, S.; Boccagni, C.; Bonavita, J. Assessing the Accuracy of ChatGPT in Answering Questions About Prolonged Disorders of Consciousness. *Brain Sci.* **2025**, *15*, 392. [CrossRef]
40. Chen, B.; Zhang, Z.; Langrené, N.; Zhu, S. Unleashing the Potential of Prompt Engineering for Large Language Models. *Patterns* **2025**, *6*, 101260. [CrossRef]
41. Ye, Q.; Axmed, M.; Pryzant, R.; Khani, F. Prompt Engineering A Prompt Engineer. *arXiv* **2023**, arXiv:2311.05661. [CrossRef]
42. Oppenlaender, J.; Linder, R.; Silvennoinen, J. Prompting AI Art: An Investigation into the Creative Skill of Prompt Engineering. *Int. J. Hum.-Comput. Interact.* **2024**, 1–23. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

MDPI AG
Grosspeteranlage 5
4052 Basel
Switzerland
Tel.: +41 61 683 77 34

Machines Editorial Office
E-mail: machines@mdpi.com
www.mdpi.com/journal/machines



Disclaimer/Publisher's Note: The title and front matter of this reprint are at the discretion of the Guest Editor. The publisher is not responsible for their content or any associated concerns. The statements, opinions and data contained in all individual articles are solely those of the individual Editor and contributors and not of MDPI. MDPI disclaims responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.



Academic Open
Access Publishing

mdpi.com

ISBN 978-3-7258-8326-4