

Application of Artificial Intelligence in Decision Making

Muhammad Syafrudin  and Norma Latif Fitriyani * 

Department of Artificial Intelligence and Data Science, Sejong University, Seoul 05006, Republic of Korea

* Correspondence: norma@sejong.ac.kr

1. Introduction

Artificial intelligence (AI) is increasingly influencing decision-making processes in medicine, communications, cybersecurity, manufacturing, finance, software engineering, and social computing. The Special Issue “Application of Artificial Intelligence in Decision Making” in Mathematics serves as an interdisciplinary platform for methods and applications that prioritize practical decision support under real-world constraints (Contributions 1–13). The contributions in this Special Issue demonstrate methodological diversity, with methods including ensemble learning, attention mechanisms, self-supervised learning, sentiment analysis, secure network architectures, and fault detection pipelines.

A distinguishing feature of this Special Issue is its balanced integration of application-driven studies and foundational modeling perspectives. In addition to problem-specific systems for healthcare, 5G/6G security, IoT optimization, and industrial quality assurance (Contributions 1–7, 12), this Special Issue presents a comprehensive review of self-supervised learning that situates emerging representation-learning paradigms (Contribution 13). This range of methodologies underscores the central role of AI as both a methodological foundation and a decision-support framework across disciplines. This collection also emphasizes the importance of rigorous empirical evaluation, with most contributions reporting performance metrics on real-world datasets and comparing these against relevant baselines. This empirical grounding enhances the practical relevance of the research and provides a roadmap for future work that seeks to bridge methodological innovation with real-world impact.

2. Special Issue Coverage and Summary

The Special Issue comprises 13 papers, including 12 research articles and 1 review, published between March 2024 and July 2025. Together, they address core questions in AI-driven decision support across a wide range of application domains, demonstrating methodological breadth alongside rigorous empirical evaluation.

From a topical perspective, the Special Issue is naturally organized into three major research clusters, each addressing a distinct domain of AI-enabled decision support, followed by a cross-cutting review that provides the representational foundations shared across all clusters.

2.1. Cluster I: AI for Secure and Intelligent Communication Infrastructures

The first cluster addresses the challenge of designing and securing next-generation wireless communication networks against both conventional cyber threats and emerging risks from quantum computing. Four contributions collectively advance this agenda across the complementary dimensions of network intelligence, access control, cyber-physical integration, and edge deployment.



Received: 19 May 2026

Accepted: 22 May 2026

Published: 30 June 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Contribution 2 develops a neuro-driven, agent-based security framework specifically designed for quantum-safe 6G networks, in which intelligent software agents coordinate dynamic defense strategies in response to evolving adversarial patterns. The framework embeds decision logic parametrized by neural network components within a distributed agent architecture, enabling adaptive, quantum-resistant policies to be enforced at network scale without centralized coordination. This design reflects the broader shift in 6G security thinking from static, perimeter-based protections toward continuously learning, autonomous defense architectures.

Building on the premise that 6G connectivity must simultaneously deliver security and ultra-low-latency computation, Contribution 4 investigates the synergistic integration of edge computing with 6G network infrastructure to meet the stringent latency, throughput, and reliability requirements of real-time IoT applications. By offloading computation to edge nodes co-located with 6G access points, the proposed architecture substantially reduces end-to-end response times while preserving bandwidth efficiency at the network core. The results demonstrate that this architectural coupling is a viable and scalable pathway toward latency-sensitive industrial and consumer IoT deployments.

Addressing securing at the access layer of 5G networks, Contribution 5 proposes an adaptive Zero Trust policy management framework tailored to the dynamic and heterogeneous topology of 5G environments, where traditional perimeter-based security models are structurally insufficient. The framework continuously evaluates trust levels for connected devices and services based on behavioral signals and contextual attributes, enforcing fine-grained access controls that evolve in response to detected anomalies. This policy-driven approach provides principled security guarantees in high-mobility, multi-tenant 5G deployments without incurring prohibitive computational overhead.

Extending Zero Trust principles from the network access layer to the domain of cyber-physical systems (CPSs), Contribution 6 examines how these strategies can be systematically adapted to the real-time operational and safety constraints of CPSs operating within 6G network environments, where security failures can propagate into physical infrastructure and threaten human safety. The study addresses the challenge of reconciling strict timing guarantees and deterministic actuation requirements with the continuous re-authentication demanded by Zero Trust architectures. The resulting framework provides a unified security posture spanning both the cyber and physical domains, addressing a vital gap in protecting critical connected infrastructure.

2.2. Cluster II: AI for Health and Biomedical Decision Support

The second cluster addresses clinical and epidemiological decision-making, applying diverse AI methodologies to tasks in which early, accurate, and reliable predictions have direct consequences for patient outcomes and public health policy. Three contributions span preventive cardiology, population-level disease surveillance, and diagnostic imaging.

Contribution 1 targets the early prediction of cardiovascular disease (CVD) onset, which remains among the most consequential problems in preventive medicine given CVD's status as a leading cause of preventable mortality worldwide. The study proposes an ensemble pipeline combining Bagging with Histogram Gradient Boosting and an advanced feature selection strategy that mitigates the curse of dimensionality inherent in high-dimensional clinical datasets while retaining the most diagnostically informative variables. The resulting classifier achieves competitive performance on standard clinical benchmarks, offering a practical and interpretable decision-support tool suitable for deployment in preventive cardiology settings.

Extending AI-driven health intelligence from the individual patient to the population level, Contribution 3 presents a real-time monitoring system that tracks infectious disease

dynamics in urban environments. The system integrates heterogeneous urban data streams and applies machine learning models that detect early outbreak signals, enabling rapid public health responses before incidence rates can reach epidemic thresholds. By situating disease surveillance within a smart city computational framework, this work connects epidemiological modeling and applied AI in a deployable, infrastructure-compatible monitoring architecture.

As part of the biomedical cluster, Contribution 8 introduces the Multi-Feature Attention Network (MFAN) for automated breast cancer classification from histopathological imaging data. This architecture employs multi-scale feature extraction combined with attention mechanisms that selectively weight diagnostically discriminative image regions, thereby improving classification robustness under intra-class morphological variation and challenging conventional convolutional approaches. The reported experimental results establish MFAN as a clinically viable second-opinion system that reduces the cognitive burden on radiologists and pathologists in high-throughput diagnostic workflows.

2.3. Cluster III: AI for Data-Intensive Decision Analytics

The third cluster demonstrates the applicability of advanced AI methods to decision-making problems in domains characterized by high data volume, adversarial patterns, semantic ambiguity, or structural complexity. Five contributions address financial fraud detection, social media sentiment analysis, software quality assurance, institutional communication, and industrial manufacturing inspection.

Contribution 7 focuses on detecting tax fraud, a domain characterized by severe class imbalance and adversarially manipulated data patterns that systematically degrade conventional supervised classifiers. The study constructs a soft-voting ensemble that combines Generative Adversarial Network (GAN)-augmented training data with an encoder-based feature representation, leveraging GANs' generative capacity to address class imbalance, while the ensemble voting mechanism stabilizes decision boundaries in highly skewed distributions. Empirical evaluations confirm that this architecture improves detection sensitivity while maintaining acceptable false-positive rates, directly supporting the development of more reliable automated audit systems.

Addressing the challenge of extracting structured opinion signals from noisy social media data, Contribution 9 presents a stacking ensemble approach for the sentiment analysis of Twitter posts. The ensemble integrates lexicon-based sentiment scoring with multiple machine learning classifiers, combining the domain-agnostic coverage of rule-driven lexicons with the adaptive inductive capacity of learned models to achieve greater robustness than either paradigm in isolation. The resulting system demonstrates measurable accuracy gains over individual baseline methods, providing a reliable tool for monitoring public opinion dynamics in real time.

Contribution 10 addresses software quality assurance through the lens of effort-aware fault-proneness prediction, targeting the identification of defect-prone software modules to enable development teams to allocate testing resources more efficiently. The study introduces a set of non-API-based, package-modularization metrics that capture structural and organizational properties of software components without relying on API-level coupling measurements, thereby improving generalizability across heterogeneous codebases with varying API usage conventions. Crucially, the effort-aware evaluation protocol adopted in the study assesses predictive models not only on detection accuracy but also on the practical cost of applying them in real project contexts, providing a more practically meaningful evaluation framework.

Contribution 11 examines the intersection of natural language processing and institutional strategic communication, presenting a case study on how a major research university manages and amplifies its organizational identity on X (formerly Twitter). By applying

NLP and sentiment analysis to a corpus of institutional posts and corresponding audience responses, the study quantifies the relationship between content strategy, sentiment polarity, and engagement metrics in a controlled real-world setting. The findings translate directly into practical guidelines for institutional communicators who seek to use social media analytics to strengthen organizational brand presence in higher education and similar professional contexts.

Completing the cluster, Contribution 12 addresses quality control in wiring harness manufacturing, a safety-critical industrial process in which undetected defects can cause cascading failures in automotive and aerospace systems. The study develops an AI-based inspection pipeline that integrates computer vision and automated classification models to detect faults along the production line in real time, replacing or augmenting manual inspection with high-throughput automated analysis. The reported detection performance demonstrates that the system satisfies industrial tolerance requirements, establishing a reproducible methodological template for AI-driven quality assurance in safety-critical production environments.

2.4. Cross-Cutting Foundation: Advances in Self-Supervised Learning

Providing the collection with a unifying theoretical foundation, Contribution 13 presents a comprehensive review of self-supervised learning (SSL) and its consequential advancements within deep learning research. The review systematically classifies SSL paradigms, including contrastive learning, masked autoencoders, and generative pretraining, and evaluates their applicability across vision, language, and multimodal domains, synthesizing a rapidly evolving literature into a coherent conceptual framework. By establishing a shared theoretical vocabulary for representation learning, Contribution 13 contextualizes the cross-domain transferability of methods exemplified throughout the Special Issue, underscoring that advances in unsupervised pretraining constitute a unifying methodological thread that extends across healthcare, security, analytics, and manufacturing applications alike.

Table 1 summarizes the 13 contributions by core topic, primary method, and application domain.

Table 1. Summary of published papers in the Special Issue.

Contribution	Core Topic	Primary Method	Application Domain
1	CVD onset prediction	Bagging, Histogram Gradient Boosting, feature selection	Healthcare
2	Quantum-safe 6G security	Neuro-driven agent-based architecture	6G networks
3	Infectious disease surveillance	AI-assisted real-time monitoring	Public health
4	IoT real-time processing	Edge computing + 6G integration	IoT/communications
5	Zero Trust policy management	Adaptive framework for 5G	5G network security
6	Cyber-physical system security	Zero Trust strategies	6G/CPS
7	Tax fraud detection	Soft-voting ensemble (GAN + encoder)	Finance
8	Breast cancer classification	Multi-Feature Attention Network (MFAN)	Biomedical imaging
9	Sentiment analysis of tweets	Stacking ensemble (lexicon + ML)	Social media
10	Fault-proneness prediction	Effort-aware, non-API metrics	Software engineering
11	Institutional identity on social media	NLP and sentiment analysis	Institutional communication
12	Wiring harness fault detection	AI-based inspection pipeline	Manufacturing
13	Self-supervised learning (review)	Comprehensive SSL taxonomy	Cross-domain AI foundations

3. Potential Future Open Research

The breadth and depth of the 13 contributions collected in this Special Issue illuminate both the substantial progress made in AI-driven decision-making and the research challenges that remain open. The following directions are grounded in limitations and extensions identified across the papers published in this Special Issue and, in the editors' assessment, represent the most consequential avenues for future investigation.

3.1. Explainability, Calibration, and Trustworthy Decision Intelligence

A persistent gap across application-driven AI systems is the disconnect between high predictive accuracy and the interpretability required for deployment in high-stakes settings. The ensemble-based CVD predictor (Contribution 1) and the MFAN breast cancer classifier (Contribution 8) both achieve strong discriminative performance, yet clinical adoption ultimately depends on physicians' ability to understand, audit, and override model recommendations. Future work in this area should move beyond post hoc feature attribution methods such as SHAP or LIME, which explain individual predictions but do not expose the model's internal decision logic, toward inherently interpretable architectures, including attention-regularized networks, concept-based models, and rule-extraction pipelines designed for clinical deployment.

Closely related is the problem of calibration: a model that is 90% accurate but systematically overconfident in its wrong predictions poses greater risk in a clinical or security context than a slightly less accurate but well-calibrated alternative. The security frameworks in this Special Issue (Contributions 2, 5, and 6) make access-control decisions continuously and at scale; integrating conformal prediction or Bayesian uncertainty quantification into these architectures would allow the system to escalate ambiguous decisions to human operators rather than acting unilaterally. As a concrete research target, future studies might evaluate whether calibrated uncertainty estimates improve operator response times and decision quality in simulated 5G/6G threat scenarios.

3.2. Federated Learning and Privacy-Preserving AI for Distributed Deployments

Several contributions in this Special Issue operate on sensitive data, patient clinical records (Contributions 1, 3, and 8), financial transaction logs (Contribution 7), and social media content associated with identifiable individuals (Contributions 9 and 11), yet all adopt centralized training paradigms. This architectural choice, while appropriate for controlled research benchmarks, creates substantial barriers to real-world deployment in jurisdictions subject to data-protection regulations such as GDPR or HIPAA and in organizations whose data cannot be legally or ethically pooled.

Federated learning (FL) offers a principled alternative in which model updates, rather than raw data, are shared across participating institutions, enabling collaborative model training without centralizing sensitive records. For the healthcare applications in this Special Issue, a federated variant of the CVD prediction pipeline (Contribution 1) trained across multiple hospital networks could simultaneously improve model generalization across demographic subgroups and eliminate the need for data-sharing agreements. Similarly, the infectious disease surveillance system (Contribution 3) could be extended under a federated architecture in which municipal health authorities contribute locally computed gradient updates, preserving patient-level privacy while enabling population-scale outbreak detection. Future research should also address the specific vulnerabilities of FL, including gradient inversion attacks and model poisoning, which are especially acute in adversarial contexts such as the fraud-detection setting in Contribution 7.

3.3. Cross-Domain Transfer and Self-Supervised Pretraining for Data-Scarce Applications

The comprehensive SSL review (Contribution 13) establishes a theoretical foundation for representation learning that is directly relevant to several practical limitations encountered by the application papers published in this Special Issue. Many supervised classifiers developed for specialized domains, including histopathological cancer imaging (Contribution 8), wiring harness defect detection (Contribution 12), and software fault prediction (Contribution 10), operate under conditions of limited labeled data, where the

cost and expertise required for expert annotation impose a fundamental constraint on the scale of labeled datasets.

Self-supervised pretraining on unlabeled domain data, followed by fine-tuning on small labeled sets, represents a promising pathway to overcome this constraint. For example, a masked autoencoder pretrained on large, unlabeled corpora of histopathological slides could substantially reduce the number of expert-annotated cases needed to train a high-performing cancer classifier, thereby directly extending the approach of Contribution 8. In the industrial domain, contrastive pretraining on normal wiring harness images, which are abundant in production environments, could enhance the system developed in Contribution 12's few-shot fault-detection capability, enabling robust detection of novel defect types with minimal additional annotation. Future research should systematically benchmark the trade-off between pretraining corpus size, fine-tuning data volume, and downstream classification performance across these application settings.

3.4. Real-Time and Resource-Constrained AI for Edge and Embedded Deployments

The edge computing architecture of Contribution 4 and the manufacturing inspection pipeline of Contribution 12 both operate under resource constraints, limited memory, computational throughput, and energy budgets that are characteristic of embedded and IoT deployment environments. While the results reported in both studies confirm that AI-based decision support is feasible under these constraints, the broader challenge of systematically co-designing model architecture and deployment hardware remains underexplored.

Model compression techniques, including structured pruning, quantization-aware training, and knowledge distillation, offer pathways for deploying high-capacity models on resource-constrained edge nodes without a proportional degradation in predictive performance. A concrete future direction would be the development of a distilled variant of the MFAN architecture (Contribution 8) optimized for deployment on embedded medical devices, enabling point-of-care breast cancer screening in low-resource clinical settings where access to centralized computing infrastructure is limited. Similarly, the Zero Trust policy enforcement frameworks of Contributions 5 and 6 could benefit from hardware–software co-design studies that quantify the latency and energy cost of continuous trust re-evaluation on 5G/6G network edge nodes, providing empirical guidance for practical deployment at scale.

3.5. Adaptive and Adversarially Robust AI Under Distribution Shift

A shared assumption underlying many classifiers in this context is that the statistical distribution of the deployment data closely resembles that of the training data. In practice, this assumption is violated in several of the target domains. Infectious disease outbreak patterns shift as pathogens evolve (Contribution 3); financial fraud tactics adapt to the detection systems deployed against them (Contribution 7); and the sentiment landscape of social media evolves in response to news cycles, geopolitical events, and platform policy changes (Contributions 9 and 11). Models trained under static distributional assumptions may degrade imperceptibly in production, producing confident but incorrect decisions precisely in the high-stakes scenarios for which they were designed.

Future research should explicitly evaluate AI decision-support systems under controlled distribution shift protocols, for example, training on pre-COVID-19 epidemiological data and evaluating pandemic-period outbreak signals, or training fraud detectors on historical tax records and testing on data generated after regulatory changes. Online and continual learning frameworks, which adapt model parameters incrementally as new data arrives, represent a natural response to this challenge but require careful design to prevent catastrophic forgetting of previously learned patterns. In adversarial domains such as fraud

detection and network security (Contributions 2 and 7), adaptive AI systems must also be evaluated against dynamic adversaries who actively probe and circumvent the detection model, a setting that can be formalized using game-theoretic frameworks and adversarial machine learning benchmarks.

3.6. Ethical Governance, Fairness Auditing, and Human-in-the-Loop Design

Several contributions in this Special Issue develop systems that influence decisions with direct consequences for individuals: fraud determinations with direct financial consequences for individuals (Contribution 7), clinical disease predictions (Contributions 1, 3, and 8), and automated assessments of software module quality (Contribution 10). The ethical dimensions of deploying such systems, including algorithmic fairness, accountability, and the appropriate scope of human oversight, have not yet been systematically addressed in existing contributions and represent a significant gap in the current literature.

Future research should incorporate fairness auditing as a first-class evaluation criterion alongside accuracy metrics, examining whether models exhibit differential performance across demographic subgroups defined by age, sex, ethnicity, or socioeconomic status. For the CVD prediction system (Contribution 1), this would involve stratified performance analysis across patient subgroups, with explicit reporting of any disparities and an investigation of their causes, whether in feature selection, training data composition, or model architecture. More broadly, the design of human-in-the-loop interfaces, in which AI systems refer uncertain or high-stakes cases to human reviewers rather than issuing fully automated decisions, should be studied as an integrated system rather than as a secondary consideration. Empirical studies measuring how AI confidence scores, explanation interfaces, and decision thresholds affect human operator accuracy and response time in security (Contributions 2, 5, and 6) and clinical settings (Contributions 1 and 8) would provide evidence-based design guidance for responsible AI deployment.

4. Conclusions

The Special Issue demonstrates that AI-enabled decision making is progressing toward broader applicability and methodological maturity. This 13-paper collection captures substantial thematic breadth, from secure communication infrastructures and cyber-physical resilience to healthcare intelligence, financial analytics, and manufacturing optimization. Overall, this collection offers an integrated perspective: AI is most effective for decision-making when algorithmic innovation is coupled with domain constraints, interpretability requirements, and deployment-aware evaluation.

Author Contributions: Conceptualization, M.S. and N.L.F.; formal analysis, M.S. and N.L.F.; writing—original draft preparation, M.S. and N.L.F.; writing—review and editing, M.S. and N.L.F. All authors have read and agreed to the published version of the manuscript.

Funding: This study was supported by the Sejong University Industry-Academic Cooperation Foundation (Grant No. 2026-0019-01).

Acknowledgments: The authors sincerely thank all reviewers for their rigorous and constructive evaluations, all authors and contributors for their high-quality submissions, the Assistant Editors for their professional editorial support throughout the process, and the readers for their continued interest in and engagement with this Special Issue.

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

List of Contributions

1. Fitriyani, N.L.; Syafrudin, M.; Chamidah, N.; Rifada, M.; Susilo, H.; Aydin, D.; Qolbiyani, S.L.; Lee, S.W. A Novel Approach Utilizing Bagging, Histogram Gradient Boosting, and Advanced Feature Selection for Predicting the Onset of Cardiovascular Diseases. *Mathematics* **2025**, *13*, 2194. <https://doi.org/10.3390/math13132194>.
2. Alwakeel, M. Neuro-Driven Agent-Based Security for Quantum-Safe 6G Networks. *Mathematics* **2025**, *13*, 2074. <https://doi.org/10.3390/math13132074>.
3. Alwakeel, M.M. AI-Assisted Real-Time Monitoring of Infectious Diseases in Urban Areas. *Mathematics* **2025**, *13*, 1911. <https://doi.org/10.3390/math13121911>.
4. Alwakeel, A.M. Synergistic Integration of Edge Computing and 6G Networks for Real-Time IoT Applications. *Mathematics* **2025**, *13*, 1540. <https://doi.org/10.3390/math13091540>.
5. Alnaim, A.K. Adaptive Zero Trust Policy Management Framework in 5G Networks. *Mathematics* **2025**, *13*, 1501. <https://doi.org/10.3390/math13091501>.
6. Alnaim, A.K.; Alwakeel, A.M. Zero Trust Strategies for Cyber-Physical Systems in 6G Networks. *Mathematics* **2025**, *13*, 1108. <https://doi.org/10.3390/math13071108>.
7. Alrasheedi, M.A.; Ijaz, S.; Alrashdi, A.M.; Lee, S.-W. Advanced Tax Fraud Detection: A Soft-Voting Ensemble Based on GAN and Encoder Architecture. *Mathematics* **2025**, *13*, 642. <https://doi.org/10.3390/math13040642>.
8. Nasir, I.M.; Alrasheedi, M.A.; Alreshidi, N.A. MFAN: Multi-Feature Attention Network for Breast Cancer Classification. *Mathematics* **2024**, *12*, 3639. <https://doi.org/10.3390/math12233639>.
9. Malebary, S.J.; Abulfaraj, A.W. A Stacking Ensemble Based on Lexicon and Machine Learning Methods for the Sentiment Analysis of Tweets. *Mathematics* **2024**, *12*, 3405. <https://doi.org/10.3390/math12213405>.
10. Shaikh, M.; Tunio, I.; Khan, J.; Jung, Y. Effort-Aware Fault-Proneness Prediction Using Non-API-Based Package-Modularization Metrics. *Mathematics* **2024**, *12*, 2201. <https://doi.org/10.3390/math12142201>.
11. Albarrak, K.M.; Sorour, S.E. Boosting Institutional Identity on X Using NLP and Sentiment Analysis: King Faisal University as a Case Study. *Mathematics* **2024**, *12*, 1806. <https://doi.org/10.3390/math12121806>.
12. Song, J.; Kumar, P.; Kim, Y.; Kim, H.S. A Fault Detection System for Wiring Harness Manufacturing Using Artificial Intelligence. *Mathematics* **2024**, *12*, 537. <https://doi.org/10.3390/math12040537>.
13. Abdulrazzaq, M.M.; Ramaha, N.T.A.; Hameed, A.A.; Salman, M.; Yon, D.K.; Fitriyani, N.L.; Syafrudin, M.; Lee, S.W. Consequential Advancements of Self-Supervised Learning (SSL) in Deep Learning Contexts. *Mathematics* **2024**, *12*, 758. <https://doi.org/10.3390/math12050758>.

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.