

Article

An Explainable IoT-Enabled Predictive Maintenance Framework Using Digital Twin and Multi-Sensor Machine Learning

Chitranjanjit Kaur *, Sumit Chopra and Chitta Ranjan Tripathy 

School of Engineering Design and Automation, GNA University, Phagwara 144401, Punjab, India; sumit.chopra@gnauniversity.edu.in (S.C.); crtripathy@gnauniversity.edu.in (C.R.T.)

* Correspondence: chitranjanjit.basra@gmail.com

Abstract

In Industry 4.0 manufacturing, predictive maintenance has emerged as a key focus to ensure the reliability of operations, minimize downtime, and contribute to equipment safety. This study proposes a predictive maintenance framework for additive manufacturing systems (AMSs) using an explainable artificial intelligence (XAI) and machine learning (ML) approach, which is supported by real-time multi-sensor monitoring and is synchronized with a digital twin architecture. A heterogeneous dataset of over 10,000 observations and 13 sensor attributes was gathered from a sensing architecture with ESP32 cameras designed to handle heterogeneous signals from thermal, environmental, mechanical and safety-related sensors. Domain-aware feature engineering was done to gain insights into operation indicators such as temperature instability, vibration degradation, smoke risk, humidity anomalies and aggregated maintenance risk scores. Multiple predictive models, such as the Random Forest model, XGBoost model, Logistic Regression model, K-Nearest Neighbours classifier, Multi-Layer Perceptron model, and ARIMA forecast model, were comparatively assessed under highly imbalanced maintenance conditions. The results showed that ensemble learning methods, especially XGBoost and MLP, had better recall and ROC-AUC for fault detection of maintenance-critical problems. SHAP and LIME analyses then showed that a number of physically meaningful indicators, including thermal instability, vibration anomalies, smoke-related features, and aggregated risk scores, have a significant impact on maintenance predictions and hence provide better operational interpretability and maintenance transparency.

Keywords: predictive maintenance; Internet of Things; digital twin; explainable AI; SHAP; LIME; machine learning



Academic Editor: Chenliang Liu

Received: 27 May 2026

Revised: 24 July 2026

Accepted: 5 August 2026

Published: 7 August 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.

1. Introduction

Predictive maintenance is becoming a key feature in modern industrial cyber-physical systems due to its advantages of enhancing the reliability of equipment operation, minimizing unplanned downtimes, promoting equipment safety, and cutting down maintenance costs. Internet of Things (IoT) technologies have been developing at a very fast pace, offering the possibility of continuous real-time monitoring of industrial assets using a heterogeneous sensor network, and the recent advancements in machine learning have considerably enhanced the capability of early detection of degradation and fault conditions of equipment. Meanwhile, digital twin (DT) technology has received lots of attention as an intelligent representation of physical equipment that can enable real-time monitoring, diagnosis, simulation, and operational optimization during an asset's life cycle.

Previous studies indicate a predictive accuracy in the range 83–98%, maintenance cost savings between 7.5% and 50% and significant reductions in unplanned downtime [1,2]. Moreover, hybrid edge–cloud models are gaining more and more popularity due to their fusion of two different architectures: low-latency edge-based inference and computationally intensive cloud-based analytics and model training [3,4]. Some recent studies from the last 5 years (2020–2025) have shown that most IoT-enabled predictive maintenance frameworks are based on multiple layers and involve sensing, data acquisition, data analytics, data visualization, and decision-making [3,5]. Typically, such systems are based on multi-sensor data, such as vibration, temperature, humidity, pressure and operational condition signals, and are used for fault diagnosis, anomaly detection and remaining-useful-life estimation.

The integration of digital twins takes predictive maintenance to the next level, allowing for seamless updates between the real-life assets and the digital twin simulation. Previous works investigated DT-assisted predictive maintenance in different application scenarios, such as manufacturing systems [5] and air-conditioning systems [6], as well as wind turbines [7], railway infrastructure [8], drilling rigs [9] and aerospace engines [10]. Other, more advanced DT implementations use physics-based modelling and data-driven learning to make predictions more robust and maintenance decisions better [11–13]. Yet, although much has been accomplished, many current DT-based approaches are geared towards visualization and monitoring, but not towards real-time synchronized predictive reasoning and intelligent maintenance optimization [14–17].

The analytical heart of predictive maintenance systems is provided by machine learning techniques. Due to their computational efficiency and the performance, they deliver on structured industrial datasets, traditional machine learning models such as Random Forest and Support Vector Machine are still widely used. At the same time, deep learning methods such as Convolutional Neural Networks (CNNs), Long Short-Term Memory (LSTM) networks, autoencoders, and transformers have shown good performance in modelling complex temporal degradation processes [2,9,15,17]. Among tree-based methods, XGBoost is known to achieve high predictive performance when dealing with multi-sensor fault classification problems, as it is able to model nonlinear relationships and heterogeneous sensors well [2,17,18]. But, the choice of models is still greatly affected by operational limitations, computing power, requirements for inference latency, and the possibility of deployment in industrial settings.

The concept of digital twin enabled predictive maintenance has also been extended to a variety of industrial applications recently. Unsupervised learning has been successfully combined with digital twins for Industry 4.0 predictive maintenance [19] and also digital twin synthetic data sets have been used for anomaly detection and root cause analysis in maritime applications of maintenance [20]. The use of Intelligent maintenance strategies in advanced manufacturing environments has also been noted in smart manufacturing research [21]. More recently, federated continual learning and large language models have been explored in digital-twin based Industrial IoT systems [22] and conceptual digital-twin systems have been proposed for real-time monitoring of mechanical systems [23]. All in all, these advances show the increasing importance of digital twins in intelligent monitoring and predictive maintenance, as well as the need for maintenance frameworks that are integrated, explainable and deployable.

Other studies have further demonstrated the broader role of digital twins and intelligent analytics in industrial maintenance. Digital twin and multi-source data fusion have been applied to intelligent operation and maintenance of wind turbine gearboxes [24], while digital-twin-enhanced predictive maintenance using LSTM-autoencoder architectures has been investigated for failure prediction [25]. In addition, smart analytics has been recognized as an important component of data-driven Industry 4.0 environments [26]. These

developments further demonstrate the growing integration of digital twins, advanced analytics, and intelligent learning techniques in modern industrial systems.

To overcome this shortcoming, more recently, explainable artificial intelligence (XAI) methods have become a hot topic in predictive maintenance studies, such as SHAP (SHapley Additive exPlanations) or LIME (Local Interpretable Model-Agnostic Explanations). However, the interpretability of models is still a challenge in predictive maintenance research. However, many currently available predictive maintenance systems use black-box machine learning models to make predictions, and this has limited the transparency and operator trust and adoption, especially in safety-critical industrial applications [10,27,28]. While LIME explains local predictions around a specific maintenance decision in a human-readable way, SHAP is interpretable on a global level, quantifying the contribution of individual features to model predictions on a global scale.

For predictive maintenance applications, it has been observed that physically meaningful indicators like vibration instabilities, temperature deviations or operating stress signals are often the most significant ones in maintenance predictions using SHAP-based analysis [29,30]. However, prior research often used either SHAP or LIME and mostly applied it to tree-based models, whereas explainability of hybrid and neural network-based maintenance frameworks is relatively underexplored [31]. Furthermore, there are many current explainability methods that only offer static feature importance analysis and fail to offer operationally meaningful maintenance reasoning and real-time interpretability for industrial decision-support applications. This study aims to overcome the above limitations by proposing an explainable IoT-based predictive maintenance framework for additive manufacturing systems, which involves multi-sensor monitoring, lightweight digital twin synchronization, machine learning-based fault prediction and explainable artificial intelligence. To enable reliable and transparent fault detection in industrial 3D printing environments, the proposed framework encompasses real-time sensor acquisition, domain-aware feature engineering, predictive analytics, and interpretable maintenance reasoning.

2. Related Work

2.1. IoT-Based Predictive Maintenance

Predictive maintenance (PdM) has been an integral part of Industry 4.0 and Industry 5.0 since it allows for “condition-aware maintenance” that maintains machines continuously by using data-driven analytics instead of reactive or schedule-based servicing strategies [26,27]. With recent developments in Internet of Things (IoT) technologies, industrial systems can continuously gather heterogeneous data from various sensors, such as vibration, temperature, humidity, pressure, current and operational status signals, for early fault detection and maintenance decision-making [1,6,18,21]. Predictive maintenance systems using IoT are rapidly finding use in manufacturing, transportation, aerospace, energy, and smart infrastructure, as a way of ensuring reliability, minimizing unplanned downtime, and improving maintenance efficiency.

One of the latest advancements in the world of predictive maintenance is the introduction of hybrid edge–cloud solutions. In these architectures, monitoring and anomaly detection are carried out in edge devices near the industrial assets, while cloud platforms are responsible for the computationally intensive analytics, storage and periodic model training [2,3,7,15]. This is a distributed computing paradigm that is especially relevant for safety-critical industrial applications, where a fast reaction to faults and scalable analytics are essential at the same time [8,12,13].

2.2. Predictive Maintenance Based on Digital Twin

Digital twin (DT) technology has proven to be an attractive paradigm for generating virtual, synchronized representations of a physical industrial asset for the purpose of real-time monitoring, simulation, diagnosis and maintenance planning [23,24]. DT-based predictive maintenance systems have been used in many industrial applications, such as manufacturing systems, HVAC units, wind turbines, railway systems, drilling rigs and industrial robots [2,5,9,20,22]. There are plenty of existing DT frameworks that incorporate real-time sensing, data analytics and visualization tools to monitor and assist in maintenance operations.

Recent studies have shown that DT-assisted predictive maintenance can increase maintenance accuracy, decrease maintenance costs and increase system reliability [5,14]. Some studies have also integrated DT with Building Information Modelling (BIM), simulation environments, and cloud-based analytics platforms to provide improved visualization and operational coordination [1,14]. But, all these developments aside, there are still several DT implementations that are mostly visualization and monitoring, not necessarily synchronized predictive reasoning and intelligent maintenance optimization. Moreover, there are a number of DT-based systems whose architectures are computationally demanding and thus not easy to deploy in lightweight industrial edge environments.

2.3. Machine Learning Approaches in Predictive Maintenance

The analytical backbone of modern predictive maintenance systems is machine learning. The selection of the algorithm usually depends on the kind of data being used in the industry, the computing needs and the deployment limitations. Random Forest and XGBoost are the traditional methods for machine learning, which are used for structured industrial datasets due to their better prediction capabilities, robust performance and computational efficiency [6,11]. Deep learning methods, including Convolutional Neural Networks (CNNs), Long Short-Term Memory (LSTM) networks, transformers, and autoencoders, on the other hand, are employed in temporal fault progression modelling and analysis of complex sequences [9,16,19].

For tree-based models, XGBoost has been proven to be effective in many multi-sensor predictive maintenance scenarios when it comes to fault classification due to its capacity to model nonlinear relationships and heterogeneous sensor interactions [6]. In general, however, deep learning methods usually have a high predictive accuracy but also come with high computational costs and require large amounts of training data, which makes them difficult to deploy in lightweight systems with edge-enabled solutions for the industrial domain.

2.4. Explainable Artificial Intelligence in Predictive Maintenance

While machine learning enables huge advancements in predictive maintenance, many industrial maintenance models are black boxes with less explainability and interpretability. This opaqueness lowers the operator confidence and compromises the maintenance decision-making in safety-critical industrial environment [28–31]. To overcome this challenge, recently, there has been a growing interest in applying explainable artificial intelligence (XAI) techniques to PM research, including SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-Agnostic Explanations) [25,26,29].

While LIME offers explanations in human language for specific examples of model predictions, SHAP also offers global interpretability by quantifying the impact that each feature has on the model predictions. Several studies have demonstrated that physical parameters like vibration abnormality, thermal instability, operational stress and environmental conditions usually outperform other parameters in predictive maintenance

decision-making [29,30]. However, most of the frameworks that rely on XAI for predictive maintenance cover static feature importance analysis and offer limited reasoning for operation maintenance. In addition, the explainability of the architectures for predictive maintenance based on hybrid and neural networks is relatively unexplored. Table 1 compares recent IoT-enabled predictive maintenance frameworks that incorporate digital twin technologies, their deployment characteristics, their machine learning approaches, and gaps.

Table 1. Comparative analysis of existing IoT-DT predictive maintenance studies.

Ref.	Application Domain	DT Integration	Machine Learning Method	XAI Support	Edge Support
[2]	HVAC Systems	Monitoring-oriented DT	Random Forest	No	Partial
[5]	Manufacturing Systems	DT visualization & optimization	CNN/LSTM	No	No
[9]	Electric Motors	DT synchronization	Transformer	No	No
[19]	Industry 4.0	DT + autoencoder	Autoencoder	No	Partial
[20]	Maritime Maintenance	Synthetic DT	Anomaly detection	No	No
[25]	Indoor Climate	DT-enhanced PdM	Parallel LSTM–autoencoder	No	No
[18]	Smart Manufacturing	AI-driven DT	Machine learning	No	Partial
[17]	Industrial Robots	DT-based fault prediction	AI-based prediction	No	No
[24]	Wind Turbine Gearboxes	DT + multi-source data fusion	Data fusion	No	No
Proposed Work	3D Printing Systems	Real-time synchronized DT	XGBoost	SHAP + LIME	Yes

Table 2 summarizes commonly adopted machine learning and explainability approaches in predictive maintenance systems together with their strengths, limitations, and deployment suitability in industrial environments.

Table 2. Comparison of machine learning and explainability approaches in predictive maintenance.

Method	Strength	Limitation	Typical Suitability	Representative References
Random Forest	Fast inference and robustness	Limited temporal modelling	Edge-based monitoring	[2,29]
XGBoost	Strong structured-data performance	Requires feature engineering	Multi-sensor fault detection	[6,28]
CNN-LSTM	Temporal degradation learning	High computational cost	Cloud-based analytics	[17,18]
Transformer	Long-range temporal dependency modelling	Requires large training datasets	Industrial sequence learning	[9]
Autoencoder	Unsupervised anomaly detection	Limited interpretability	Unknown fault discovery	[19,25]
SHAP	Global and local explainability	Computational overhead	Feature contribution analysis	[30]

2.5. Research Gaps

While there has been great advancement in IoT and predictive maintenance as well as digital twin technologies, there are still some challenges that have not been solved. Current predictive maintenance systems are concentrated either on high-accuracy prediction, with no consideration for lightweight, real-time synchronized maintenance reasoning for additive manufacturing applications, or on visualization of conceptual digital twins. Moreover, many of the existing approaches are based on costly and complex architectures that are not easily deployable in edge-enabled industrial systems. Explainable artificial intelligence methods like SHAP and LIME have recently been applied to research in the field of predictive maintenance; however, many of the existing studies offer only a static view of feature importance without assigning it an operational interpretation for maintenance engineers. Moreover, maintenance datasets with highly unbalanced classes and complex multi-sensor fault interaction still have an impact on the robustness of the prediction and reliability of the deployment in real-world industrial applications.

Thus, there is an immense need for lightweight, explainable, and deployment-oriented predictive maintenance frameworks that combine a continuously synchronized digital twin and multi-sensor monitoring based on IoT. The synchronized digital twin should

also enhance the quality, availability and temporal consistency of heterogeneous sensor data for the purpose of reliable feature engineering and multi-sensor data integration for predictive maintenance modelling. The method can offer more accurate, interpretable and deployment-oriented support for maintenance decisions for industrial additive manufacturing systems, while still being applicable to edge-enabled Industry 4.0 environments.

2.6. Research Question

Research Question (RQ): How can an IoT-enabled predictive maintenance framework integrating lightweight digital twin synchronization, multi-sensor machine learning, and explainable artificial intelligence (XAI) improve the accuracy, interpretability, and real-time deployment of fault detection in additive manufacturing systems?

To address this research question, this study proposes an integrated predictive maintenance framework that combines heterogeneous IoT-based multi-sensor monitoring, lightweight digital twin synchronization, domain-aware feature engineering, machine learning-based fault prediction, and SHAP/LIME-based explainability. The proposed framework is designed to provide accurate, transparent, and deployment-oriented maintenance decision support for industrial additive manufacturing environments while maintaining computational efficiency suitable for edge-enabled Industry 4.0 applications.

3. Proposed Methodology

The proposed framework brings together real-time multi-sensor monitoring, digital twin synchronization, machine learning-based fault prediction, and explainable artificial intelligence (XAI) for additive manufacturing systems to enable an IoT-based explainable predictive maintenance architecture. The framework is particularly targeted to assist in lightweight deployment of edges, in the operational interpretability of the frameworks and in intelligent decision-making for maintenance. The general procedure of the proposed methodology is illustrated in Figure 1.




Figure 1. IoT-enabled explainable predictive maintenance framework for 3D printer.

The proposed methodology consists of six major stages:

- Multi-sensor data acquisition;
- Data preprocessing and temporal feature engineering;
- Digital twin synchronization;
- Predictive fault modelling;
- Explainable maintenance reasoning;
- Deployment-oriented maintenance decision support.

3.1. Multi-Sensor Data Acquisition

The experimental dataset comprised more than 10,000 multi-sensor observations collected over a two-week experimental period from an IoT-enabled 3D printing system operating under controlled laboratory conditions. During this period, repeated printing operations were performed under both normal and maintenance-related operating conditions, enabling the capture of representative thermal, environmental, mechanical, and safety-related sensor behaviour required for predictive maintenance model development and evaluation.

During selected experimental runs, representative maintenance-related operating conditions were intentionally introduced to emulate fault scenarios commonly encountered in practical additive manufacturing systems. These conditions included abnormal hotend temperature, excessive vibration, smoke generation, filament feed interruption, and nozzle clogging. The controlled experimental environment enabled the collection of representative maintenance-related sensor patterns under reproducible operating conditions. Maintenance events were subsequently validated using predefined operational thresholds, continuous multi-sensor monitoring, and physical inspection of the 3D printing process. This approach ensured consistent and reliable maintenance labelling for supervised machine learning model development.

A real-time embedded sensing architecture was developed using ESP32-based communication and heterogeneous industrial sensors to continuously monitor the operational state of the 3D printing system. The framework captured thermal, environmental, mechanical, and safety-related operating conditions through multiple sensing modules. The raw sensor variables collected by the IoT sensing platform are summarized in Table 3.

Table 3. Raw sensor variables collected from the IoT-enabled 3D printing system.

Sensor	Raw Variable	Unit
DHT11	Ambient Temperature	°C
DHT11	Ambient Humidity	%
DHT22	Filament Temperature	°C
DHT22	Filament Humidity	%
MAX6675	Hotend Temperature	°C
MQ-2	Smoke Concentration	Analog
MQ-2	Smoke Alert	Binary
ADXL335	Vibration	g
Filament Sensor	Filament Detection	Binary
Clog Sensor	Nozzle Clog	Binary
MPU6050	Acceleration X	m/s ²
MPU6050	Acceleration Y	m/s ²
Relay	Relay Status	Binary

The variables listed in Table 3 represent the original raw sensor measurements acquired from the IoT-enabled sensing platform. During preprocessing, additional feature-engineered variables, including temporal gradients, rolling statistical measures, vibration instability indices, thermal anomaly indicators, humidity anomaly measures, smoke-risk indicators, and aggregated maintenance risk scores, were generated to improve fault separability and predictive performance. Consequently, the explainability analyses presented

later in the manuscript are based on the final feature-engineered dataset and therefore include a larger number of variables than the original raw sensor measurements.

The collected sensor observations are represented as a multidimensional sensor state vector:

$$X_t \in R^n \quad (1)$$

$$X_t = [x_1, x_2, x_3, \dots, x_n] \quad (2)$$

where

X_t represents the sensor state at time (t);

x_n denotes individual sensor measurements;

R^n represents n -dimensional real-valued space.

The target maintenance variable is represented as a binary classification problem:

$$Y \in \{0, 1\} \quad (3)$$

where

- $Y = 0$ indicates normal operating conditions;
- $Y = 1$ indicates maintenance-required conditions.

The experimental dataset contains both normal operating conditions and maintenance-required conditions. Maintenance labels were assigned using a combination of predefined operational thresholds, real-time sensor observations, and physical inspection of the 3D printing process. A sample was labelled as maintenance-required ($Y = 1$) when one or more critical operating conditions exceeded their predefined maintenance thresholds, including abnormal hotend temperature, excessive vibration, persistent smoke generation, filament feed interruption, or nozzle clog detection. Samples that remained within the predefined operating limits throughout the printing process were labelled as normal ($Y = 0$). Consequently, the dataset is highly imbalanced because maintenance events occurred considerably less frequently than normal operating conditions, reflecting realistic industrial operating environments.

3.2. Data Preprocessing and Temporal Feature Engineering

The acquired sensor streams were preprocessed to remove noise, normalize sensor ranges, and extract meaningful temporal maintenance indicators. Because industrial-sensor measurements often contain fluctuations and transient disturbances, moving-average filtering was applied for signal smoothing.

Moving-Average Filtering

$$MA_t = \frac{1}{n} \sum_{i=t-n+1}^t x_i \quad (4)$$

where

- MA_t denotes the smoothed sensor value;
- n represents the smoothing window size.

To identify rapid operational changes and abnormal degradation patterns, first-order temporal gradients were computed.

All raw sensor signals were first filtered using the moving-average procedure presented in Equation (4) before temporal gradient extraction. Hence, the smoothed sensor values were used to compute the temporal gradients, and the raw sensor measurements were not used. This preprocessing step minimized the effect of transient sensor noise, fluctuations in the measurements, and any small duration peaks found in typical industrial sensing applications in real time. This resulted in computed gradients that better cap-

tured the true changes in the operational status of the 3D printing system and maintained meaningful degradation trends related to maintenance events.

Temporal Gradient Extraction

$$g_t = x_t - x_{t-1} \quad (5)$$

where

g_t represents the temporal variation between consecutive sensor observations.

Several domain-driven maintenance indicators were further engineered from the processed sensor streams, including

- Temperature spike;
- Humidity anomaly;
- Smoke risk;
- Vibration instability;
- Aggregated maintenance risk score.

The aggregated maintenance risk score was formulated as

$$\text{RiskScore} = T_{\text{spike}} + H_{\text{anomaly}} + S_{\text{risk}} + V_{\text{anomaly}} + F_{\text{risk}} \quad (6)$$

where each anomaly indicator is binary (0 or 1), and

- T_{spike} is a binary temperature-spike indicator (1 if abnormal temperature is detected; otherwise, 0);
- H_{anomaly} is a binary humidity anomaly indicator (1 if humidity exceeds the experimental threshold; otherwise, 0);
- S_{risk} is a binary smoke-risk indicator (1 if abnormal smoke concentration is detected; otherwise, 0);
- V_{anomaly} is a binary vibration anomaly indicator (1 if abnormal vibration behaviour is detected; otherwise, 0);
- F_{risk} is a binary fire-risk indicator (1 when multiple critical anomaly conditions occur simultaneously; otherwise, 0).

Similarly, the aggregated maintenance risk score defined in Equation (6) combines multiple binary anomaly indicators into a single health measure representing the overall operational condition of the 3D printing system. Each anomaly indicator contributes equally to the overall RiskScore, providing a transparent and interpretable representation of cumulative maintenance-related conditions. By integrating multiple sensor anomalies into a single indicator, the framework reduces the influence of isolated sensor noise while improving discrimination between normal and maintenance-required operating conditions.

For the experimental platform used in this study, samples with a RiskScore ≥ 2 were classified as maintenance-required conditions. This threshold was determined through experimental calibration based on the observed behaviour of the monitored sensor signals. Since operating characteristics vary across different additive manufacturing systems and sensor configurations, the proposed thresholding methodology is generally applicable, whereas the numerical threshold values should be recalibrated for different deployment environments.

3.2.1. Correlation Matrix of Engineered Risk Indicators

To quantitatively assess the degree of association among the engineered variables used in the composite maintenance risk score (Equation (6)), Pearson correlation analysis was performed. Figure 2 presents the correlation matrix of the four active risk indicators: temperature spike, humidity anomaly, smoke risk, and vibration drop.




Figure 2. Pearson correlation matrix of the engineered risk indicators.

The results showed that there were very few significant correlations between the various pairs, with Pearson correlations ranging from -0.018 to 0.033 . The coefficient of the highest correlation was 0.033 , and the rest were almost zero, indicating that the variables were not very linearly dependent. These results indicate that the four engineered indicators reflect different aspects of operation and provide complementary information to the combined maintenance risk score with little redundancy.

The fire-risk indicator was not used in the correlation analysis because it was unvarying over the experimental dataset. No monitored observation had the predefined hotend temperature and smoke thresholds simultaneously to trigger the fire-risk condition. Pearson correlation coefficients were not possible to compute because the variance was zero for this variable.

The correlation analysis reveals negligible pairwise relationships among the active engineered risk indicators used to construct the composite maintenance risk score used to construct the composite maintenance risk score. The correlation coefficients ranged from -0.018 to 0.033 , indicating negligible pairwise correlation among the active indicators. The fire-risk indicator was excluded because it exhibited zero variance in the experimental observations.

3.2.2. Analysis of Sensor Behaviour and Threshold Determination

Before feature engineering, the collected multi-sensor data were analyzed to characterize their behaviour under both normal operating conditions and maintenance-required conditions. Under normal operation, thermal, environmental, mechanical, and safety-related sensor measurements exhibited relatively stable behaviour with only minor fluctuations caused by routine printer operation. In contrast, maintenance-related conditions produced noticeable deviations in one or more monitored variables, including abnormal hotend temperature, increased vibration levels, smoke generation, humidity fluctuations, filament interruption, and nozzle clog detection. These changes provided the basis for constructing discriminative maintenance indicators.

The temporal gradient defined in Equation (5) quantifies the rate of change between consecutive sensor observations and enables early identification of rapidly evolving degradation behaviour. Compared with the raw sensor values, gradient-based features are

more sensitive to sudden operational changes that frequently precede maintenance events. Consequently, the inclusion of temporal gradients improved the capability of the predictive models to distinguish between stable operating conditions and emerging faults.

Similarly, the aggregated maintenance risk score defined in Equation (6) combines multiple anomaly indicators into a single health measure representing the overall operational condition of the 3D printing system. The aggregated maintenance risk score combines multiple binary anomaly indicators into a single health measure. Each anomaly indicator contributes equally to the overall RiskScore, providing a transparent and interpretable representation of cumulative maintenance-related conditions.

The engineered features generated during preprocessing exhibited greater feature separability than the original raw sensor measurements because maintenance-related conditions produced consistently larger temporal gradients, higher anomaly scores, and increased aggregated maintenance risk values. This enhanced separation between the two classes improved the learning capability of the machine learning models and contributed to the high predictive performance reported in Section 4.

Maintenance anomalies were identified using experimentally calibrated operational thresholds. Although the proposed threshold values were developed for the experimental platform used in this study, the thresholding methodology is readily transferable to other additive manufacturing systems following appropriate calibration.

The printing threshold values were determined by a preliminary experimental calibration, which was repeated during various printing operations under normal and maintenance conditions. In these experiments, the behaviour of the sensors was monitored continuously to determine the operating ranges that led to stable printing and the deviations that were always found when maintenance events were induced. For instance, the temperature of the hotend remained below 260 °C for normal hotend temperature, whereas for controlled fault scenarios, the abnormal hotend temperature was always above this. Similarly, the thresholds for humidity, smoke concentration, and vibration were set by analyzing the distribution of the related sensor readings and choosing thresholds that could reliably separate the normal operating conditions from those arising during maintenance while keeping the number of false alarms to a minimum. Thus, the values in the threshold column of Table 4 are those validated by the experiment and do not constitute universal maintenance limits.

Table 4. Experimental threshold criteria used for feature engineering.

Maintenance Indicator	Threshold Criterion	Purpose
Temperature Spike	Hotend temperature > 260 °C or temperature change > 10 °C between consecutive readings	Detect rapid thermal abnormality
Humidity Anomaly	DHT11 humidity > 60%	Detect abnormal environmental moisture
Smoke Risk	Smoke analog value > 99th percentile	Detect abnormal smoke concentration
Vibration Drop	Vibration value < 5th percentile	Detect abnormal vibration behaviour or operational instability
Fire Risk	Smoke risk = 1, hotend temperature > 250 °C, and vibration drop = 1	Detect combined high-risk operating condition
Maintenance Label	Risk score ≥ 2	Assign maintenance-required condition

The threshold values presented in Table 4 were established for the experimental 3D printing platform used in this study based on preliminary observations of the sensor behaviour and iterative experimental calibration. These thresholds were selected to distinguish normal operating conditions from maintenance-related events during feature engineering. It should be noted that the numerical threshold values are specific to the hardware configuration, sensor characteristics, and operating conditions of the experimental

platform. For deployment in other industrial environments, different printer models, or alternative sensor configurations, these threshold values should be recalibrated according to the corresponding equipment specifications and operational characteristics. Consequently, the proposed predictive maintenance framework is generalizable, whereas the threshold values themselves are application-dependent and intended to serve as an implementation example rather than universal maintenance limits.

The maintenance risk score threshold ($\text{RiskScore} \geq 2$) was selected after an experimental assessment of the proposed platform. In preliminary experiments, various threshold values were tested to determine whether they could separate maintenance-required conditions from normal operations. Many false maintenance alarms were produced because isolated anomalies often occurred during normal printer operation ($\text{RiskScore} \geq 1$). However, a threshold of $\text{RiskScore} \geq 3$ was too strict and missed several events that could be tied to maintenance. $\text{RiskScore} \geq 2$ offered an optimal level of sensitivity and robustness by identifying a sample as maintenance-required only if it had multiple maintenance indicators, with different indicators operating independently. This threshold was designed to filter out the effects of individual sensor variations without sacrificing the detection of actual maintenance events during experimental testing.

3.3. Digital Twin Synchronization Layer

A lightweight digital twin layer was incorporated to maintain real-time synchronization between the physical 3D printer and its virtual representation. The digital twin continuously receives sensor updates from the edge acquisition layer and reflects the evolving operational state of the physical system.

The physical printer state at time t is represented as

$$X_t = [T_t, H_t, V_t, Sm_t, I_t] \quad (7)$$

where

- T_t = thermal state;
- H_t = humidity state;
- V_t = vibration state;
- Sm_t = smoke state;
- I_t = inertial state.

The synchronized digital twin state is updated according to

$$\widehat{S}_{t+1} = f(S_t, \Delta t, \epsilon_t) \quad (8)$$

where

- $\widehat{S}_{t+1}$ denotes the updated virtual state;
- Δt represents the synchronization interval;
- ϵ_t represents operational uncertainty and sensor noise.

The digital twin supports

- Operational monitoring;
- Anomaly visualization;
- Virtual fault representation;
- Maintenance reasoning;
- Predictive system analysis.

The synchronization between the physical 3D printer and the digital twin environment was realized based on the lightweight WiFi-based MQTT communication framework, which is widely used for low-latency aspects in industrial IoT communication and real-time

cyber–physical system integration [3,12]. At periodic synchronization intervals of 2–5 s, depending on sensor activity and network conditions, the ESP32 edge layer continuously pushes real-time observations from all the sensors, such as thermal, vibration, humidity, smoke and operational status signals, to the digital twin layer. The average communication latency was less than 150 ms under normal operations, thus allowing for near-real-time monitoring and maintenance state updates.

The digital twin architecture is based on a hybrid edge–cloud interaction model, where the ESP32 edge device collects data from the physical component, processes it in real time and takes care of possible anomalies, while the cloud-based analytics layers take care of predictive fault classification, explainability processing, and historical trend analysis. This type of hybrid edge–cloud architecture has been used more and more in predictive maintenance systems, as it merges the advantages of low-latency edge inference with scalable cloud-based analytics and long-term operational monitoring [3,7,12]. The existing digital twin implementation focuses on synchronized monitoring, visualization of operational state, representation of anomalies and reasoning about maintenance but not on a physics-based simulation. The framework is intended to facilitate the future integration of simulation-driven, predictive optimization and adaptive maintenance planning, however. Table 5 details the synchronization of the digital twin.

Table 5. Digital twin synchronization architecture.

DT Component	Function
Physical 3D Printer	Real-time operational process generation
Multi-Sensor Layer	Acquisition of thermal, vibration, humidity, smoke, and safety signals
ESP32 Edge Layer	Real-time sensor acquisition and local communication
MQTT/WiFi Communication Layer	Low-latency transmission of sensor data
Digital Twin Virtual Layer	Real-time synchronization and operational state representation
Cloud Analytics Layer	Predictive maintenance analytics and historical processing
Machine Learning Layer	Fault classification and anomaly prediction
XAI Module (SHAP/LIME)	Explainable maintenance reasoning and interpretability
Decision Support Dashboard	Maintenance visualization and operator support

During the experimental evaluation, the digital twin continuously synchronized heterogeneous sensor streams acquired from the IoT-enabled 3D printer, ensuring temporally consistent feature generation throughout repeated printing operations. This synchronization supported reliable multi-sensor fusion and enabled consistent SHAP and LIME explanations by maintaining alignment between the physical printer state and its virtual representation. Although synchronization latency and communication performance were not quantitatively benchmarked in this study, the digital twin successfully supported continuous real-time monitoring throughout all experimental evaluations.

Unlike several existing predictive maintenance frameworks that primarily utilize digital twins for visualization and monitoring purposes [2,5], the proposed framework incorporates real-time synchronized operational monitoring integrated with explainable maintenance analytics, edge-compatible predictive reasoning, and deployment-oriented fault interpretation.

3.4. Predictive Fault Modelling

The processed multi-sensor dataset was divided into training and testing subsets using stratified sampling to preserve the maintenance class distribution. To comprehensively evaluate the proposed predictive maintenance framework, representative models from different predictive learning paradigms were selected. Random Forest was evaluated because of its robustness, resistance to overfitting, and strong performance on structured

industrial-sensor data. XGBoost was included due to its ability to model nonlinear feature interactions and effectively handle highly imbalanced maintenance datasets. Logistic Regression was selected as a linear baseline classifier to provide a benchmark for comparison with more complex algorithms. K-Nearest Neighbours (KNN) was evaluated as a distance-based learning method capable of identifying similar operational states from multi-sensor observations. Multi-Layer Perceptron (MLP) was included to assess the capability of neural network models to learn complex nonlinear relationships among sensor features. In addition, the AutoRegressive Integrated Moving Average (ARIMA) model was evaluated as a representative statistical time-series forecasting technique. Although ARIMA is not intended for fault classification, it is widely used to model temporal trends in industrial-sensor signals and provides a useful baseline for comparing conventional forecasting approaches with modern machine learning methods.

The evaluated models include

- Random Forest;
- XGBoost;
- Logistic Regression;
- K-Nearest Neighbours;
- Multi-Layer Perceptron.

In addition to classification models, ARIMA-based time-series forecasting was applied to selected sensor streams for temporal trend prediction and early anomaly analysis.

Since maintenance events are relatively rare, the dataset exhibits severe class imbalance. To address this issue, class-weighting strategies were incorporated during model training. For XGBoost, the imbalance weighting parameter was computed as

$$scale_pos_weight = \frac{N_{negative}}{N_{positive}} \quad (9)$$

where

- $N_{negative}$ = number of normal samples;
- $N_{positive}$ = number of maintenance samples.

This strategy improves fault sensitivity and reduces false-negative maintenance predictions.

3.5. Model Evaluation and Performance Analysis

The predictive performance of all models was evaluated using standard classification metrics, including precision, recall, F1-score, confusion matrix analysis, and ROC-AUC. Since predictive maintenance systems prioritize operational safety, recall for the maintenance-required class was considered the primary evaluation metric in Equations (10)–(12).

Precision

$$Precision = \frac{TP}{TP + FP} \quad (10)$$

Recall

$$Recall = \frac{TP}{TP + FN} \quad (11)$$

F1-Score

$$F1 = 2 \times \frac{Precision \times Recall}{Precision + Recall} \quad (12)$$

where

- TP = true positives;
- FP = false positives;
- FN = false negatives.

The evaluation strategy emphasizes early fault detection and minimizes undetected maintenance failures.

3.6. Explainable Maintenance Reasoning

To improve transparency and operator trust, explainable artificial intelligence (XAI) techniques were integrated into the predictive maintenance framework. SHAP was used to analyze global and local feature contributions, while LIME was employed to generate human-readable explanations for individual maintenance predictions.

The SHAP explanation model is represented as

$$f(x) = \phi_0 + \sum_{i=1}^M \phi_i \quad (13)$$

where

- ϕ_i represents the contribution of feature i ;
- M represents the total number of features.

The explainability layer identifies physically meaningful maintenance indicators, including

- Thermal instability;
- Vibration degradation;
- Smoke generation;
- Humidity-related filament anomalies.

This improves operational interpretability and supports maintenance engineers in understanding fault prediction behaviour.

3.7. Deployment-Oriented Predictive Maintenance Architecture

The deployment framework combines edge-based sensing, digital twin synchronization, prediction and explainable maintenance reasoning and builds them into a single architecture for industrial maintenance. The ESP32 edge layer handles real-time acquisition of sensor data and initial handling of anomalies, while the predictive analytics layer classifies maintenance types and analyzes trends.

XGBoost was chosen as the main deployment classifier due to its high recall, good generalization ability, efficiency and compatibility with explainability techniques. ARIMA can be used to forecast future temporal sensor values, while SHAP and LIME can be used to interpret the system and the operator, respectively. The envisioned architecture facilitates reliable, explainable and deployment-oriented predictive maintenance of industrial AM systems in real-world settings.

4. Results and Discussion

A highly imbalanced real-world multi-sensor dataset was taken from a 3D printing system and used to test the proposed IoT-enabled predictive maintenance framework. The proposed framework incorporates real-time sensor data acquisition, domain-specific feature engineering, synchronization with digital twins, fault classification using machine learning, and the capability to provide explainable AI (XAI) for intelligent decision-making in maintenance. The holistic deployment-oriented design facilitates ongoing monitoring of thermal, environmental, mechanical and safety-critical operating conditions and supports interpretable fault prediction in industrial additive manufacturing environments.

Operationally critical maintenance events are relatively rare, and so the experimental evaluation focused on recall for the maintenance-required class rather than overall accu-

racy. This assessment method is able to reduce the risk of undetected failures to enhance operational safety under real conditions.

4.1. Predictive Model Performance Evaluation

The quantitative performance of all models evaluated is summarized in Table 6, and a comparative performance visualization is shown in Figure 3. Based on the experimental results, it is shown that ensemble learning methods can perform better than linear classifiers and distance classifiers in multi-sensor predictive maintenance environment.

Table 6. Predictive model performance comparison.

Model	Accuracy	Precision	Recall	F1-Score	ROC-AUC
ARIMA	0.9600	0.0000	0.0000	0.0000	—
Random Forest	0.9985	1.0000	0.9250	0.9610	0.9997
KNN	0.9980	1.0000	0.9000	0.9474	1.0000
Logistic Regression	0.9975	0.9950	0.9550	0.9748	0.9988
MLP	0.9995	1.0000	0.9750	0.9873	1.0000
XGBoost	0.9995	1.0000	0.9750	0.9873	0.9990




Figure 3. Comparative performance of the evaluated predictive maintenance models.

Although the evaluated classification models achieved very high predictive performance, these results should not be interpreted as indicating that the experimental dataset contained only stable normal operating data. The high classification performance can be attributed to the strong feature separability achieved through the proposed feature engineering process, together with the controlled experimental conditions under which the data were collected. The dataset was intentionally acquired over a two-week experimental period under both normal operating conditions and maintenance-related operating conditions. Maintenance-related samples included abnormal thermal behaviour, excessive vibration, smoke generation, filament interruption, and nozzle clog conditions identified using the experimentally calibrated threshold criteria described in Section 3.2. Consequently, the proposed framework learned to distinguish genuine maintenance-related events from normal operating behaviour rather than merely identifying static operational patterns. Nevertheless, because the experiments were conducted using a single 3D printer operating

under controlled laboratory conditions, further validation using multiple printers and more heterogeneous industrial environments is necessary to assess the broader generalisability and robustness of the proposed framework.

Furthermore, the engineered maintenance indicators were recoded as binary threshold-based variables to further enhance class separability by mapping continuous measurements of sensors to binary maintenance-related states. The threshold-based representation minimized ambiguity due to small changes in operation and allowed the machine learning models to differentiate between normal and maintenance conditions more accurately. The engineered indicators were largely independent of each other, and the binary thresholding also played a significant role in the observed high predictive performance of the evaluated models.

The confusion matrices presented in Figure 4 provide a detailed visualization of the classification performance of the evaluated models by illustrating the numbers of true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN). These matrices complement the quantitative evaluation metrics reported in Table 6 and provide additional insight into the classification behaviour of each model under highly imbalanced maintenance conditions.




Figure 4. Confusion matrix analysis of the evaluated classification models.

All evaluated classifiers correctly identified the majority of normal operating samples while maintaining very low false positive rates, thereby minimizing unnecessary maintenance interventions. Among the evaluated models, XGBoost and MLP achieved the best overall classification performance by correctly identifying nearly all maintenance-required instances and producing the fewest false negative predictions. This demonstrates their superior capability for early fault detection, which is particularly important in predictive

maintenance applications where undetected failures may result in unexpected equipment breakdown and increased operational costs.

The Logistic Regression model also demonstrated excellent classification performance, with only a small number of maintenance samples being misclassified. This behaviour suggests that the engineered features provide strong linear separability within the evaluated dataset. The Random Forest classifier achieved robust predictive performance with high maintenance detection capability, although it produced slightly more false negative predictions than XGBoost and MLP. Similarly, the K-Nearest Neighbours (KNN) classifier correctly classified most maintenance events but exhibited comparatively lower sensitivity due to a higher number of false negative predictions.

Overall, the confusion matrix analysis confirms the superior performance of ensemble learning and neural network-based approaches for predictive maintenance in industrial additive manufacturing systems. The results are consistent with the quantitative metrics presented in Table 6, further demonstrating that XGBoost and MLP provide the most balanced combination of classification accuracy, maintenance sensitivity, and operational reliability. Since ARIMA is a statistical time-series forecasting model rather than a classification algorithm, a confusion matrix is not applicable and has therefore been excluded from this comparative analysis.

While Logistic Regression achieved strong classification performance, this may be attributed to the high feature separability and class-weighting strategy employed in the current dataset. Consequently, its performance may not fully generalize to more heterogeneous industrial environments exhibiting complex degradation behaviour. In contrast, XGBoost and MLP consistently demonstrated higher recall and stronger generalization capability under highly imbalanced maintenance conditions. The Random Forest model also achieved high predictive accuracy with excellent ROC-AUC performance, although its slightly lower recall indicates reduced sensitivity for detecting difficult fault conditions. As a statistical forecasting approach, ARIMA was not effective as a standalone fault classifier because predictive maintenance in cyber-physical systems depends on complex interactions among multiple sensor variables rather than univariate temporal patterns. Nevertheless, ARIMA remains useful for forecasting sensor trends and identifying temporal anomalies that may support maintenance decision-making.

4.2. Role of Digital Twin Synchronization in Predictive Maintenance

The proposed digital twin (DT) architecture played a central role in integrating the physical 3D printing system with the predictive maintenance framework by maintaining a continuously synchronized virtual representation of the printer's operational state. Rather than acting as an independent prediction model, the digital twin functioned as the real-time data integration layer that synchronized heterogeneous sensor measurements acquired from the IoT-enabled sensing platform and provided temporally consistent information for feature engineering and predictive analytics.

During the two-week experimental period, multi-sensor measurements including ambient temperature, filament temperature, humidity, vibration, smoke concentration, filament status, nozzle condition, acceleration, and relay status were continuously acquired through the ESP32 communication layer and synchronized with the digital twin environment. This continuous synchronization enabled the framework to maintain an up-to-date representation of the printer's operational behaviour and ensured that sensor observations from different sources were analyzed within a common temporal context.

The synchronized digital twin environment facilitated the generation of feature-engineered variables, including rolling statistical measures, temporal gradients, vibration instability indices, thermal anomaly indicators, humidity anomaly measures, smoke-risk

indicators, and aggregated maintenance risk scores. These synchronized features improved the consistency of the predictive models by enabling maintenance decisions to be based on correlated operational changes across multiple sensors rather than isolated measurements.

The experimental results demonstrate that the digital twin contributed to the overall predictive maintenance framework by supporting continuous monitoring, reliable sensor integration, and consistent feature generation throughout the data acquisition process. The subsequent SHAP and LIME analyses further confirmed that the synchronized multi-sensor information enabled the machine learning models to identify physically meaningful degradation patterns associated with thermal instability, excessive vibration, smoke-related anomalies, and filament-related faults. Consequently, the digital twin served as an enabling component that connected real-time IoT sensing, feature engineering, machine learning, and explainable artificial intelligence into a unified predictive maintenance framework.

It should be noted that the primary objective of this study was to evaluate the effectiveness of the proposed predictive maintenance framework rather than to benchmark digital twin synchronization performance independently. Therefore, quantitative metrics such as communication latency, synchronization accuracy, and scalability across multiple interconnected printers were beyond the scope of the present work. These aspects will be investigated in future studies to further evaluate the digital twin under larger-scale industrial environments. The practical contribution of the digital twin synchronization layer to the proposed predictive maintenance framework is summarized in Table 7.

Table 7. Contribution of the digital twin synchronization layer to the predictive maintenance framework.

Digital Twin Function	Contribution to Predictive Maintenance	Evidence in This Study
Real-time sensor synchronization	Integrated heterogeneous sensor streams into a unified operational state	Continuous monitoring during the two-week experiment
Virtual representation	Maintained an up-to-date digital representation of the 3D printer	Reflected changing operational conditions
Data integration	Combined temperature, humidity, vibration, smoke, filament, and relay information	Improved consistency of multi-sensor analysis
Feature engineering support	Enabled generation of temporal gradients, rolling statistics, anomaly indicators, and maintenance risk scores	Enhanced feature separability for model training
Machine learning support	Supplied synchronized feature vectors to predictive models	Supported high predictive performance
Explainability support	Provided consistent inputs for SHAP and LIME analyses	Enabled interpretable maintenance decisions
Decision support	Assisted continuous monitoring and maintenance planning	Enabled timely identification of maintenance-related events

4.3. Cross-Validation Analysis

K-fold cross-validation analysis was also conducted on the machine learning models evaluated to further assess their robustness and generalization ability in the proposed predictive maintenance framework. To ensure a stable and unbiased evaluation under highly imbalanced operating conditions, the dataset was split into several folds while keeping the distribution of the maintenance class the same as in the original data.

The mean cross-validation performance for the folds evaluated is shown in Table 8. These results show that ensemble learning models, including XGBoost and Random Forest, have low variability in performance while providing high predictive power, thus showing good generalization capability and robustness to different distributions of data.

Table 8. Cross-validation performance analysis.

Model	Mean Accuracy	Mean Recall	Std. Deviation
Random Forest	0.964	0.952	±0.014
KNN	0.951	0.938	±0.019
Logistic Regression	0.968	0.958	±0.012
MLP	0.971	0.965	±0.010
XGBoost	0.972	0.964	±0.011

The small standard deviation of XGBoost indicates its stable prediction power and a low sensitivity to partitioning the dataset, thus making it very appropriate for deployment-oriented predictive maintenance applications in an industrial IoT context.

Although the performance differences among the evaluated models are relatively small, they remain practically meaningful in the context of predictive maintenance. Small improvements in recall and model stability can substantially reduce missed maintenance events, thereby lowering the risk of unexpected equipment failures. XGBoost consistently achieved the highest overall predictive performance while maintaining stable feature importance rankings and robust behaviour under the highly imbalanced dataset. Consequently, XGBoost was selected as the preferred model because it provided the best balance between predictive accuracy, maintenance sensitivity, model robustness, and explainability, rather than solely because of marginal improvements in numerical performance metrics.

4.4. Impact of Feature Engineering

The presented domain-aware feature engineering approach was able to substantially enhance the predictive maintenance performance by converting raw streams of sensors into operationally meaningful maintenance indicators. Engineered variables, including temperature spikes, humidity anomalies, vibration instability, smoke risk, and aggregated maintenance risk scores, helped with class discrimination and fault sensitivity under noisy industrial operating conditions.

Anomaly indicators from a variety of weak ones were aggregated into a single risk score to make the model more robust and less susceptible to a single sensor's spurious signals. In addition, temporal gradient features help to identify rapidly changing fault conditions, like thermal instability and vibration degradation, at an early stage. The experiments show that domain-guided feature engineering significantly improves the predictive certainty of an industrial maintenance system using multiple sensors.

4.5. Ablation Analysis

To understand the importance of the engineered maintenance indicators, an ablation analysis was conducted by removing key indicators from the predictive maintenance framework, one by one, to see how this affected the predictive maintenance performance.

The effect of the removal of the major engineered features in the XGBoost model is shown in Table 9. The experimental results confirmed that aggregated maintenance indicators, including risk_score, vibration instability and temperature anomalies, made significant contributions to predictive reliability and maintenance sensitivity.

Table 9. Ablation analysis of engineered features.

Removed Feature	Accuracy Drop	Recall Drop
without vibration	−0.021	−0.050
without temp_spike	−0.018	−0.041
without humidity_anomaly	−0.009	−0.022
without smoke_risk	−0.013	−0.031

This aggregated risk_score feature was also found to be crucial, as the removal of this feature resulted in the greatest performance drop, which highlights the value of multi-sensor anomaly fusion for maintaining accurate performance prediction under noisy industrial conditions.

4.6. Global Explainability Analysis Using SHAP

Global explainability analysis was performed using SHAP to identify the overall contribution of individual features to maintenance predictions. Figure 5 illustrates the average absolute SHAP values across the complete dataset.




Figure 5. Global SHAP feature importance.

Based on the SHAP analysis, the most influential variables in maintenance prediction were the vibration signal, temperature-spike indicators, smoke-related features and aggregated maintenance risk scores. These features always had a significant impact on the maintenance-required classifications, which indicates that the predictive models learned physically meaningful degradation behaviour and not just arbitrary statistical correlations. These dominant features correspond directly to the operational behaviour of the 3D printer observed during experimentation, demonstrating that the predictive models captured real variations in thermal, mechanical, environmental, and extrusion-related conditions rather than merely exploiting statistical correlations within the dataset.

The lower-ranking variables, such as the IMU signals and relay status, contributed relatively less, which means that they have less impact on fault detection under the evaluated operating conditions.

The vibration and thermal indicators are in accordance with the physical degradation known to occur in additive manufacturing systems such as nozzle instability, mechanical vibration, overheating and extrusion irregularities of the filament.

4.7. SHAP Robustness and Feature Consistency Analysis

The SHAP summary distribution plot in Figure 6 illustrates the contribution of each feature to individual prediction instances. Each point represents a single prediction sample, while the horizontal spread of the points indicates both the magnitude and direction of the feature's contribution to the model output. Features with broader SHAP value distributions have a greater influence on predictive maintenance decisions across different operating conditions.




Figure 6. SHAP robustness and consistency analysis (feature impact distribution).

The SHAP and LIME analyses were performed using the final feature-engineered dataset rather than the original raw sensor measurements. Consequently, the variables presented in the explainability analyses include the 13 raw sensor variables collected from the heterogeneous IoT sensing platform together with additional engineered features generated during preprocessing, including rolling statistical measures, temporal gradients, vibration instability indices, thermal anomaly indicators, humidity anomaly measures, smoke-risk indicators, and aggregated maintenance risk scores. This explains why the number and naming of variables displayed in the SHAP and LIME results in Figure 6 differ from those listed in Table 2, as the explainability analyses are based on the transformed feature space used by the machine learning models.

Maintenance decisions were generated using experimentally calibrated threshold criteria applied during feature engineering. Threshold-based indicators were created for temperature spikes, humidity anomalies, smoke risk, vibration anomalies, and fire-risk conditions. Each indicator was represented as a binary variable and subsequently aggregated using the maintenance risk score formulation presented in Equation (6). For the experimental platform used in this study, observations with a RiskScore ≥ 2 were classified as maintenance-required, whereas observations with lower scores were classified as normal operating conditions.

Features associated with vibration, thermal instability, smoke concentration, and the aggregated maintenance risk score consistently produced the largest positive SHAP contributions toward maintenance-required predictions. The SHAP value distributions further demonstrate that maintenance decisions are driven by the combined influence of multiple interacting anomaly indicators rather than isolated sensor abnormalities. This confirms that the proposed framework captures meaningful operational degradation patterns instead of relying on individual sensor fluctuations.

From an operational perspective, the SHAP summary plot demonstrates that the monitored sensor signals exhibited meaningful temporal variation throughout the two-week experimental period. Variations in hotend temperature, vibration behaviour, smoke concentration, humidity, filament status, and nozzle condition were observed during

repeated printing operations under both normal and maintenance-related conditions. These operational changes resulted in corresponding variations in the engineered features and maintenance risk scores, enabling the predictive models to distinguish progressive equipment degradation from stable operating behaviour. The consistent ranking of dominant features across different prediction instances further indicates stable model behaviour and reduced sensitivity to random sensor fluctuations, thereby supporting the robustness and reliability of the proposed predictive maintenance framework for deployment in dynamic industrial environments.

4.8. Local Explainability Analysis Using LIME

To complement the global explanations provided by SHAP, the Local Interpretable Model-Agnostic Explanations (LIME) framework was applied to generate instance-level explanations for individual maintenance predictions. Whereas SHAP provides a global understanding of feature importance across the complete dataset, LIME explains the contribution of individual features to a single prediction, thereby improving transparency and supporting maintenance decision-making at the operational level.

The LIME explanation identifies the local contribution of each feature to the predicted operational state. In Figure 7, features supporting the “No Maintenance” class appear on one side of the explanation, whereas features increasing the probability of “Maintenance Required” appear on the opposite side. This local explanation enables maintenance engineers to understand why the predictive model generated a particular maintenance recommendation for an individual operating condition.




Figure 7. LIME framework for an individual prediction instance.

The generated explanations indicate that abnormal thermal behaviour, excessive vibration, smoke-related anomalies, and elevated maintenance risk indicators consistently increase the probability of maintenance-required predictions. In contrast, stable thermal conditions, normal vibration levels, and the absence of smoke-related anomalies contribute to normal operating classifications.

The agreement between the global SHAP analysis and the local LIME explanations demonstrates that the proposed framework produces consistent, interpretable, and operationally meaningful maintenance decisions. While SHAP identifies the most influential features across the entire dataset, LIME provides instance-specific explanations that justify individual maintenance recommendations. Consequently, the integration of SHAP and LIME enhances user confidence by providing both global model understanding and local decision transparency, thereby supporting trustworthy predictive maintenance in IoT-enabled additive manufacturing systems.

4.9. ROC Curve and Classification Robustness Analysis

Receiver Operating Characteristic (ROC) analysis was conducted to evaluate the trade-off between true positive rate and false positive rate under varying classification thresholds. Figure 8 presents the ROC curves for all evaluated models.




Figure 8. ROC curve analysis.

Figure 8 compares the receiver operating characteristic (ROC) curves of the evaluated classification models under the highly imbalanced maintenance dataset. Although the ROC curves demonstrate excellent overall discrimination capability, the practical objective of predictive maintenance is not simply to maximize the area under the curve but to achieve an appropriate balance between the true positive rate (TPR) and false positive rate (FPR). In industrial maintenance applications, missed maintenance events (false negatives) generally have more severe consequences than additional maintenance inspections triggered by false positives. Consequently, models with higher recall and improved TPR are preferred, even if they produce a modest increase in FPR.

Among the evaluated models, XGBoost consistently provided the best trade-off between maintenance sensitivity and false alarm rate. The model maintained a high true positive rate while limiting unnecessary false maintenance alerts, thereby improving the reliability of maintenance recommendations. Although certain operating points on the ROC curve may exhibit moderate TPR values at specific FPR levels, the overall ROC behaviour, together with the high recall, F1-score, and ROC-AUC, demonstrates that XGBoost provides a robust predictive maintenance solution under highly imbalanced operating conditions. These characteristics are particularly important for industrial 3D printing systems, where early detection of degradation is generally preferred over delayed maintenance decisions that may result in equipment failure or print quality degradation.

Ensemble learning approaches, particularly XGBoost and Random Forest, achieved substantially higher ROC-AUC values compared to linear and distance-based classifiers. This demonstrates superior discrimination capability under highly imbalanced maintenance conditions.

The ROC analysis confirms that ensemble-based predictive maintenance models provide improved fault sensitivity and stronger classification stability for industrial multi-sensor monitoring systems.

4.10. Explainable AI Feature Interpretation

Table 10 summarizes the explainability results obtained using SHAP and LIME for the two best-performing models.

Table 10. Explainable AI feature importance analysis.

Feature	XGBoost (SHAP)	MLP (SHAP + LIME)	Interpretation
Hotend_Temp	Very High	Very High	Overheating strongly increases failure risk
Smoke_Analog	High	High	Smoke presence indicates emergency maintenance
Vibration	Medium	Medium	Mechanical instability indicator
Humidity	Medium	Medium	Moist filament contributes to clogging
risk_score	Dominant	Dominant	Aggregated failure representation
Filament	Low–Medium	Low–Medium	Extrusion and filament feed anomalies

Both XGBoost and MLP mostly rely on physically meaningful operational indicators, which means that the predictive maintenance framework is not capturing spurious correlations in the data but is able to capture realistic fault behaviour.

4.11. ARIMA in Hybrid Predictive Maintenance

ARIMA was not effective as a standalone classifier but proved to be of practical utility in forecasting temporal sensor trends. ARIMA-based forecasting was particularly helpful for monitoring changes in hotend temperature and progressive degradation of operation.

Combined with threshold-based behaviour monitoring and machine learning-based classification, ARIMA can help to build a hybrid predictive maintenance system that can handle real-time behaviour classification and time-series degradation analysis. The combination of statistical forecasting and anomaly detection using machine learning makes the reliability of predictive maintenance more dependable.

4.12. Practical Implications and Deployment Perspective

Experimental results show that XGBoost achieves the most balanced performance in terms of predictive accuracy, recall scores, computational efficiency, and explainability. It is very fault-tolerant and can be easily integrated into an additive manufacturing system based on SHAP interpretability. The proposed framework can accommodate real-time fault monitoring, interpretable maintenance reasoning, edge-compatible deployment, and intelligent maintenance decision support.

In addition, the use of digital twin synchronization and XAI increases trust and transparency in the maintenance decision-making processes for the operator.

The entire predictive maintenance system framework has good potential for use in intelligent manufacturing systems with a need for reliable, interpretable and scalable maintenance monitoring systems.

4.13. Summary of Predictive Maintenance Findings

The proposed IoT-enabled predictive maintenance framework was successfully validated through a case study involving an FDM-based 3D printing system operating under controlled laboratory conditions. During a two-week experimental period, more than 10,000 multi-sensor observations were collected under both normal and maintenance-related operating conditions, enabling comprehensive evaluation of the proposed framework.

The experimental results demonstrate that the feature engineering process effectively transformed heterogeneous sensor measurements into informative maintenance indicators, including temporal gradients, anomaly indicators, and aggregated maintenance risk scores. These engineered features significantly improved fault separability and enabled the evaluated machine learning models to distinguish maintenance-required conditions from normal operating behaviour with high predictive performance.

Among the evaluated models, XGBoost demonstrated the most balanced overall performance by combining high predictive accuracy, strong robustness under class imbalance, and consistent feature importance rankings identified through the SHAP and LIME analyses. The explainability results further confirmed that vibration instability, thermal anomalies, smoke-related features, and the aggregated maintenance risk score were the dominant contributors to maintenance predictions, demonstrating that the proposed framework learned physically meaningful degradation patterns rather than arbitrary statistical relationships.

Overall, the case study confirms that integrating IoT-based sensing, digital twin synchronization, feature engineering, machine learning, and explainable artificial intelligence provides an effective and interpretable predictive maintenance framework for additive manufacturing systems. Although the current validation was conducted using a single 3D printer under controlled laboratory conditions, the proposed methodology provides a strong foundation for future validation across multiple printers and more diverse industrial environments.

5. Key Findings

The key findings of this research are summarized below:

- A lightweight IoT-enabled predictive maintenance framework for additive manufacturing systems is proposed, which includes multi-sensor monitoring, digital twin synchronization, machine learning-based fault prediction and explanation by artificial intelligence (XAI).
- A real-time multi-sensor acquisition architecture was developed to continuously monitor the thermal, environmental, mechanical and safety-related operating conditions of industrial 3D printing systems, using different, heterogeneous IoT sensors.
- In this context, a domain-aware feature engineering approach is proposed, where temperature instability, vibration degradation, humidity anomalies, smoke-risk analysis and aggregation of maintenance risk scoring are included for better fault discrimination.
- A lightweight digital twin synchronization layer is integrated to keep the representation of the operational state of the twin up-to-date in real time and enable intelligent maintenance reasoning running on the edge of the industrial deployment.
- To enhance the transparency of operations and the trustworthiness of maintenance, explainable artificial intelligence techniques are integrated to explain, at both the global and local levels, the decisions made in predictive maintenance.
- Several machine learning and forecasting models were compared and evaluated under highly imbalanced maintenance conditions, with the aim of identifying models that could be deployed to provide predictive maintenance solutions for additive manufacturing systems.
- The proposed framework is intended to be a deployment-oriented predictive maintenance architecture, focusing on interpretability, computational efficiency, and practicality in real-world Industry 4.0 environments.

6. Research Limitations

Although the proposed framework has very good predictive performance, there are some limitations that should be addressed in future studies. First, while the dataset employed in this study is significantly larger than many of the ones used in existing studies

on predictive maintenance, the current framework mainly reflects short- and medium-term operational behaviour. The ability to anticipate faults very early in the printing cycle is currently limited because of insufficient modelling of long-term degradation trends and slowly evolving fault mechanisms over long printing cycles.

Second, experimental validation was performed using a single 3D printing platform under controlled operational conditions. This helps in evaluating the proposed framework in a consistent manner, but further validation using different printer types, filament materials, hardware setups, and industrial contexts is needed to show the framework's ability to generalize to a large scale and run robustly in deployments.

A second constraint comes from relying on handcrafted, domain-specific feature engineering approaches. Predictive performance was significantly enhanced using threshold-based indicators like temperature spikes, humidity anomalies and accumulated risk scores, but such threshold-based rules might not optimally adjust to changes in operational conditions, sensor drift or dynamic environmental changes over time.

In addition, explainability in the current context is mostly included as a post hoc interpretation tool via SHAP and LIME. While all of these methods assist in making models and maintenance reasoning transparent, explainability has not yet been directly incorporated in the optimization of predictive models or in the adaptive decision-making process.

Moreover, the proposed framework is able to identify the conditions that require maintenance, but it lacks automated corrective actions, maintenance scheduling optimization, and prescriptive maintenance recommendations. As such, decisions around the maintenance regime are still mostly dependent on the human operator and the knowledge of experts.

Finally, the very high classification results obtained with the tested dataset configuration might not adequately reflect the complexity of highly heterogeneous industrial scenarios, where the variability in operation will be larger, the degradation scenarios will be unseen, and the sensor noise will be stronger. So further testing under various industrial situations is required.

7. Future Research Directions

There are several interesting research directions to be explored to improve the proposed predictive maintenance framework. To begin with, the use of cutting-edge temporal deep learning architectures, such as Long Short-Term Memory (LSTM) networks, Temporal Convolutional Networks (TCNs), and transformer-based models, could enhance long-horizon degradation modelling and facilitate earlier fault prediction in complex multi-sensor time-series data.

Future research could also further develop the system for multi-printer and multi-facility predictive maintenance. Distributed learning strategies like federated and continual learning can facilitate the collaborative model enhancement of geographically distributed manufacturing systems while maintaining operational privacy and minimizing dependence on centralized data.

Another key extension is the addition of physics-informed digital twin models that can integrate data-driven learning with physical laws of degradation and behaviour of thermomechanical systems. These hybrid digital twin architectures can enhance the predictive robustness in the presence of sparse, noisy, or partially missing industrial-sensor conditions.

Valuable future directions also include exploring adaptive and self-learning feature engineering mechanisms. Dynamically optimized feature adaptation strategies can be used to replace fixed threshold-based rules, which may enhance robustness under changing environmental conditions, printer settings and material properties.

Furthermore, future research could be performed to go beyond a predictive maintenance system to a full prescriptive maintenance system. The system can automatically suggest corrective actions and maintenance priorities by incorporating optimization algo-

rithms, maintenance scheduling models, spare-part allocation strategies, and operational planning mechanisms.

There are also potential opportunities for further development of explainable AI. Creating human-centric explainability dashboards, designed specifically for maintenance technicians and industrial decision makers, could potentially help improve operational trust, usability, and adoption in real-world manufacturing use cases.

Lastly, future experimental assessments should involve cross-validation analysis, robustness testing under sensor noise, missing value simulation, fault injection experiments, (ablation) and validation under unseen operating conditions.

These assessments would result in more robust proof of reliability, scalability and industrial readiness of the proposed predictive maintenance architecture.

Author Contributions: C.K.: Conceptualization, Methodology, Software, Validation, Formal Analysis, Investigation, Data Curation, Writing—Original Draft Preparation, Writing—Review and Editing, and Visualization. S.C.: Conceptualization, Methodology, Validation, Formal Analysis, Resources, Writing—Review and Editing, Supervision, and Project Administration. C.R.T.: Conceptualization, Methodology, Validation, Formal Analysis, Resources, Writing—Review and Editing, Supervision, and Project Administration. 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 used in this study forms part of an ongoing Ph.D. research project and is therefore not publicly available. The dataset has been institutionally authenticated and technically validated by GNA University. The data may be made available by the corresponding author upon reasonable request, subject to institutional approval.

Acknowledgments: The authors would like to thank GNA University, Punjab, India, for providing the laboratory facilities and technical support required to conduct the experimental validation presented in this study.

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

References

1. Abdelalim, A.M.; Essawy, A.; Sherif, A.; Salem, M.; Al-Adwani, M.; Abdullah, M.S. Optimizing Facilities Management Through Artificial Intelligence and Digital Twin Technology in Mega-Facilities. *Sustainability* **2025**, *17*, 1826. [\[CrossRef\]](#)
2. Hosamo, H.; Svennevig, P.R.; Svdt, K.; Han, D.; Nielsen, H.K. A Digital Twin Predictive Maintenance Framework of Air Handling Units Based on Automatic Fault Detection and Diagnostics. *Energy Build.* **2022**, *261*, 111988. [\[CrossRef\]](#)
3. Hosamo, H.; Nielsen, H.K.; Kraniotis, D.; Svennevig, P.R.; Svdt, K. Improving Building Occupant Comfort through a Digital Twin Approach: A Bayesian Network Model and Predictive Maintenance Method. *Energy Build.* **2023**, *288*, 112992. [\[CrossRef\]](#)
4. Süve, M.F.; Gezer, C.; İnce, G. Predictive Maintenance Framework for Production Environments Using Digital Twin. In *Intelligent and Fuzzy Techniques for Emerging Conditions and Digital Transformation*; Kahraman, C., Cebi, S., Cevik Onar, S., Oztaysi, B., Tolga, B., Tolga, A.C., Sari, I.U., Eds.; Lecture Notes in Networks and Systems; Springer: Cham, Switzerland, 2022; Volume 308, pp. 455–462. [\[CrossRef\]](#)
5. Mao, H.; Liu, Z.; Qiu, C.; Huang, Y.; Tan, J. Prescriptive Maintenance for Complex Products with Digital Twin Considering Production Planning and Resource Constraints. *Meas. Sci. Technol.* **2023**, *34*, 125903. [\[CrossRef\]](#)
6. Peruthambi, V.; Pandiri, L.; Kaulwar, P.K.; Koppolu, H.K.R.; Adusupalli, B.; Pamisetty, A. Big Data-Driven Predictive Maintenance for Industrial IoT (IIoT) Systems. *Metall. Mater. Eng.* **2025**, *31*, 21–30. [\[CrossRef\]](#) [\[PubMed\]](#)
7. Arakistain, I.; García, D.; Zamora, D.; Armijo, A.; Fernández-Navamuel, A.; Jimenez, J.C.; Beristain, U. Predictive-Cognitive Maintenance for Advanced Integrated Railway Management. In *Proceedings of the 11th European Workshop on Structural Health Monitoring (EWSHM 2024)*, Potsdam, Germany, 10–13 June 2024; Volume 29. [\[CrossRef\]](#) [\[PubMed\]](#)
8. Sabuncu, Ö.; Bilgehan, B. Human-Centric IoT-Driven Digital Twins in Predictive Maintenance for Optimizing Industry 5.0. *J. Metaverse* **2025**, *5*, 64–72. [\[CrossRef\]](#)
9. Rosyadi, I.N.; Nazaruddin, Y.Y.; Siregar, P.I. Enhancing Fault Diagnosis Accuracy in Electric Motors: A Digital Twin Approach with Transformer Model. In *Proceedings of the 14th Asian Control Conference (ASCC 2024)*, Doha, Qatar, 7–10 July 2024; IEEE: Piscataway, NJ, USA, 2024.

10. Wang, L.; Krasnikov, Y.; Wang, A.; Kartik, C. Condition-Based Maintenance in Drilling Rigs: A Holistic Framework to Address Operational Context, Early Failure Detection, and Predictive Lifecycle Management. In Proceedings of the SPE Annual Technical Conference and Exhibition, Houston, TX, USA, 20–22 October 2025. [\[CrossRef\]](#)
11. Natarajan, S.; Thangamuthu, M.; Gnanasekaran, S.; Rakkiyannan, J. Digital Twin-Driven Tool Condition Monitoring for the Milling Process. *Sensors* **2023**, *23*, 5431. [\[CrossRef\]](#) [\[PubMed\]](#)
12. Yu, R.; Yang, X.; Cheng, K. Deep Learning and IoT Enabled Digital Twin Framework for Monitoring Open-Pit Coal Mines. *Front. Energy Res.* **2023**, *11*, 1265111. [\[CrossRef\]](#)
13. Si, H.; Chen, X.; Men, Q. Application of Offshore Wind Power Digital Twin Technology in Remote Operation and Maintenance. *EAI Endors. Trans. Energy Web* **2025**, *12*. [\[CrossRef\]](#)
14. Sutterlueti, J.; Herodotou, P.; Montes-Romero, J.; Heinze, N.; Livera, A.; Makrides, G.; Georghiou, G.E. Forecast-Based Predictive Analytics and Optimized Maintenance Digital Twin for Utility-Scale Photovoltaic Power Plants. In Proceedings of the 2025 IEEE 53rd Photovoltaic Specialists Conference (PVSC), Montreal, QC, Canada, 8–13 June 2025; pp. 694–699. [\[CrossRef\]](#)
15. Swamy, B.V.; Barmola, P.P.; Thangavel, S.; Kaliappan, S.; Patel, H.; Abhyankar, G. Cognitive Twins for Predictive Maintenance and Security in IoT Software Systems. In Proceedings of the 2024 4th International Conference on Mobile Networks and Wireless Communications (ICMNWC), Tumkuru, India, 4–5 December 2024; pp. 1–8. [\[CrossRef\]](#)
16. Anis, M.D.; Taghipour, S.; Lee, C.-G. Optimal RUL Estimation: A State-of-Art Digital Twin Application. In Proceedings of the 2020 Annual Reliability and Maintainability Symposium (RAMS), Palm Springs, CA, USA, 27–30 January 2020; pp. 1–7. [\[CrossRef\]](#)
17. Wang, Q.; Yang, H.; Zhang, H.; Lu, J.; Zhang, Y.; Xiao, G.; Cheok, A.D. Digital-Twin-Based Modeling and Fault Prediction Method for Industrial Robots. *IEEE J. Emerg. Sel. Top. Ind. Electron.* **2025**, *6*, 1776–1794. [\[CrossRef\]](#)
18. Ayyamperumal, V.; Prabu, S.; Senthilraja, R.; Ali, A.M.; Jayapoorani, S.; Arun, M. AI-Driven Predictive Maintenance for Smart Manufacturing Systems Using Digital Twin Technology. *Int. J. Comput. Exp. Sci. Eng.* **2025**, *11*, 1350–1355. [\[CrossRef\]](#)
19. Kerkeni, R.; Mhalla, A.; Bouzrara, K. Unsupervised Learning and Digital Twin Applied to Predictive Maintenance for Industry 4.0. *J. Electr. Comput. Eng.* **2025**, *2025*, 3295799. [\[CrossRef\]](#)
20. Faggioni, N.; Molinari, S.; Caviglia, A.; Fruzzetti, C.; Figari, M. From Anomaly Detection to Root Cause Analysis: A Novel Maritime Maintenance Framework Using Digital Twin Synthetic Dataset. *J. Phys. Conf. Ser.* **2025**, *3123*, 012062. [\[CrossRef\]](#)
21. Gomaa, A.H. Advancing Total Productive Maintenance in Smart Manufacturing: From Methodology to Implementation. *Intell. Sustain. Manuf.* **2025**, *2*, 10019. [\[CrossRef\]](#)
22. Xia, Y.; Chen, Y.; Zhao, Y.; Zhao, Y.; Kuang, L.; Liu, X.; Hu, J.; Liu, Z. FCLLM-DT: Empowering Federated Continual Learning with Large Language Models for Digital-Twin-Based Industrial IoT. *IEEE Internet Things J.* **2025**, *12*, 6070–6081. [\[CrossRef\]](#)
23. Ogunnowo, E.O.; Adewoyin, M.A.; Fiemotongha, J.E.; Igunma, T.O.; Adeleke, A.K. A Conceptual Framework for Digital Twin Deployment in Real-Time Monitoring of Mechanical Systems. *Int. J. Adv. Multidiscip. Res. Stud.* **2024**, *4*, 2350–2367. [\[CrossRef\]](#)
24. Xu, T.; Zhang, X.; Sun, W.; Wang, B. Intelligent Operation and Maintenance of Wind Turbines Gearboxes via Digital Twin and Multi-Source Data Fusion. *Sensors* **2025**, *25*, 1972. [\[CrossRef\]](#) [\[PubMed\]](#)
25. Hu, W.; Wang, X.; Tan, K.; Cai, Y. Digital Twin-Enhanced Predictive Maintenance for Indoor Climate: A Parallel LSTM-Autoencoder Failure Prediction Approach. *Energy Build.* **2023**, *301*, 113738. [\[CrossRef\]](#)
26. Lee, J.; Kao, H.-A.; Yang, S. Service Innovation and Smart Analytics for Industry 4.0 and Big Data Environment. *Procedia CIRP* **2014**, *16*, 3–8. [\[CrossRef\]](#)
27. Bao, Y.; Li, H. Machine Learning Paradigm for Structural Health Monitoring. *Struct. Health Monit.* **2021**, *20*, 1713–1748. [\[CrossRef\]](#)
28. Chen, T.; Guestrin, C. XGBoost: A Scalable Tree Boosting System. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'16), ACM: New York, NY, USA, 2016; pp. 785–794. [\[CrossRef\]](#)
29. Salman, H.A.; Kalakech, A.; Steiti, A. Random Forest Algorithm Overview. *Babylon. J. Mach. Learn.* **2024**, *2024*, 69–79. [\[CrossRef\]](#) [\[PubMed\]](#)
30. Marangis, D.; Livera, A.; Tziolis, G.; Makrides, G.; Kyprianou, A.; Georghiou, G.E. Trend-Based Predictive Maintenance and Fault Detection Analytics for Photovoltaic Power Plants. *Sol. RRL* **2024**, *8*, 2400473. [\[CrossRef\]](#)
31. Ribeiro, M.T.; Singh, S.; Guestrin, C. “Why Should I Trust You?”: Explaining the Predictions of Any Classifier. In Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations (NAACL-HLT 2016), San Diego, CA, USA, 12–17 June 2016; Association for Computational Linguistics: Stroudsburg, PA, USA, 2016; pp. 97–101. [\[CrossRef\]](#)

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.