

SYSTEMATIC REVIEW

A Systematic Review of Machine Learning Techniques for Predictive Maintenance in Industrial IoT (IIoT) Environments

Mahmudul Islam¹

¹ Latrobe University Sydney Campus, Australia

Received	Revised	Accepted	Published
28 January 2023	01 March 2023	28 March 2023	15 September 2023

Abstract

Background and Purpose: Predictive Maintenance (PdM) has emerged as a critical component of smart manufacturing, driven by Industrial Internet of Things (IIoT) technologies that enable continuous monitoring of industrial assets. Interconnected sensors and cyber-physical systems create opportunities for real-time diagnostics, early fault detection, and improved operational reliability. Machine Learning (ML) techniques transform heterogeneous industrial data streams into actionable insights, but challenges remain in model selection, performance evaluation, interpretability, and practical deployment. This review synthesizes ML techniques applied to PdM within IIoT ecosystems and examines methodological trends, strengths, limitations, and research gaps.

Methods: A systematic review methodology was adopted following PRISMA 2020 principles. Peer-reviewed studies published between 2015 and November 2022 were identified from IEEE Xplore, ACM Digital Library, Scopus, Web of Science, and ScienceDirect. Boolean search strategies targeted ML-based PdM models applied to IIoT data, cyber-physical systems, sensor networks, and digital twins. Extracted information included ML algorithms, datasets, feature-engineering approaches, performance metrics, deployment frameworks, and reported limitations. Comparative and thematic analyses were used to categorize the evidence.

Findings: Sixty-two studies met the inclusion criteria. Deep Learning architectures, including convolutional neural networks, long short-term memory networks, recurrent neural networks, and autoencoders, were prominent because of their capacity to learn complex temporal and multidimensional sensor patterns. Hybrid approaches combining deep learning with signal processing or classical ML demonstrated strong predictive performance. Persistent challenges included reliance on controlled or semi-synthetic datasets, limited real-time validation, class imbalance, lack of model interpretability, and constraints in integrating ML solutions with industrial hardware.

Theoretical Contributions: The review synthesizes theoretical perspectives underpinning ML-driven PdM. Data-driven modeling emphasizes the importance of sensor quality and feature representations; systems theory highlights interoperability across interconnected IIoT components; and decision-support theory frames predictive analytics as a tool for maintenance planning and operational optimization. Emerging approaches such as edge intelligence and hybrid data-driven models provide pathways for connecting predictive accuracy with industrial applicability.

Conclusions and Policy Implications: ML-enabled PdM has substantial potential to transform industrial asset management in IIoT environments. Wider implementation requires improved data quality, real-time processing, algorithm explainability, and reliable integration with edge and cloud infrastructures. Industrial and research stakeholders should prioritize standardized benchmark datasets, interoperable deployment frameworks, privacy-aware learning, transfer learning, and lightweight models suitable for operational environments.

Keywords: Predictive Maintenance (PdM); Machine Learning; IIoT; Deep Learning; Systematic Review; Industrial Analytics

How to cite: Islam M. A Systematic Review of Machine Learning Techniques for Predictive Maintenance in Industrial IoT (IIoT) Environments. International Journal of Business & Computational Science. 2023;3(1):1-8.

1. Introduction

The rapid evolution of the Industrial Internet of Things (IIoT) has fundamentally transformed the landscape of modern manufacturing and industrial operations [1]. With its integration of pervasive sensing, real-time connectivity, cyber-physical systems (CPS), and cloud-edge computing, IIoT enables machines, assets, and processes to operate as interconnected intelligent ecosystems. Among the various applications powered by IIoT, Predictive Maintenance (PdM) stands out as one of the most impactful [2]. PdM seeks to estimate the current health state of equipment and predict future failures using data-driven insights, allowing maintenance teams to intervene at the most optimal times [3]. This proactive approach reduces unnecessary downtime, improves asset longevity, enhances safety, and ultimately contributes to substantial cost savings across industrial sectors such as manufacturing, energy, transportation, oil and gas, and logistics [4].

Traditional maintenance strategies - such as corrective maintenance ("run-to-failure") and preventive maintenance (scheduled servicing) - suffer from significant limitations [5]. Corrective approaches lead to costly unplanned downtime, while preventive strategies may result in excessive or mistimed maintenance

activities, increasing operational expenses. PdM, fueled by real-time IIoT data streams, offers a transformative alternative by predicting failures before they occur and enabling condition-based interventions. Machine Learning (ML), including its advanced subfields like Deep Learning (DL), plays a central role in this transition by enabling efficient processing, modelling, and interpretation of heterogeneous industrial data such as vibration signals, acoustic emissions, temperature readings, pressure patterns, and multivariate temporal sequences [6].

Over the past decade, ML methods have gained widespread acceptance as powerful tools for PdM applications, supported by advancements in sensing technologies, cloud analytics, and computational resources [7]. Classical ML algorithms - including Support Vector Machines (SVM), Random Forests (RF), k-Nearest Neighbors (k-NN), and Gradient Boosting Machines - were among the earliest techniques used to analyze condition-monitoring data [8]. These methods demonstrated strong predictive power, especially when combined with domain-driven feature engineering approaches such as signal decomposition, time-frequency transformations, and statistical descriptor extraction. However, the increasing complexity and volume of IIoT data necessitated

methods capable of learning representations autonomously, giving rise to DL methods [9].

Models such as Convolutional Neural Networks (CNNs) have become popular for processing raw sensor signals, while Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks are often employed for modeling sequential machine behavior and long-range temporal dependencies [10]. Furthermore, deep autoencoders and Generative Adversarial Networks (GANs) have contributed to unsupervised PdM applications, including anomaly detection, data augmentation, and health-indicator construction [11]. Despite the promising potential of ML-based PdM, several fundamental challenges persist in real-world IIoT environments [12]. First, industrial datasets often suffer from issues such as noise contamination, missing values, severe class imbalance, and limited failure samples. Such issues impair model training and reduce generalizability. Second, real-time deployment of ML models requires optimization for latency, memory footprint, and computational constraints, especially when executed at the edge rather than the cloud [12].

Many state-of-the-art DL architectures demand significant computational resources, which makes them difficult to integrate with existing industrial hardware. Third, there is a growing emphasis on model explainability and transparency, as industrial stakeholders require interpretable decisions for safety-critical assets [13]. Black-box models, while highly accurate, may face resistance due to their limited interpretability. Furthermore, industrial environments are diverse, with varying operating conditions, machine types, and sensor configurations, making it challenging to develop standardized PdM solutions [14]. The lack of publicly available benchmark datasets limits model comparison and inhibits reproducibility across studies. As a result, industries often rely on customized or domain-specific solutions, reducing the transferability of ML models [15].

Edge computing, federated learning, transfer learning, reinforcement learning, and physics-informed ML have emerged as potential solutions to address scalability, privacy, and generalization issues, but these areas remain underexplored in PdM research. Given these complexities, there is a pressing need for a comprehensive and up-to-date systematic review that critically evaluates ML techniques applied in PdM within IIoT environments. Existing reviews either focus on narrower domains, lack rigorous methodological processes, or predate significant advancements in DL and edge intelligence [16]. Therefore, synthesizing recent evidence while adhering to the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) guidelines is essential for mapping the landscape of current research, identifying methodological trends, and uncovering persistent gaps.

This systematic review aims to consolidate findings from peer-reviewed studies published between 2015 and November 2022, a period that captures major advances in IIoT integration, cost-efficient sensor deployment, and deep learning techniques. In this review, we evaluate ML techniques used for failure prediction, anomaly detection, Remaining Useful Life (RUL) estimation, and condition monitoring across various IIoT-driven industries. The review extracts information on dataset characteristics, feature engineering strategies, model architectures, evaluation metrics, deployment considerations, and limitations. By analyzing these factors, the results aim to provide industrial practitioners, researchers, and technology developers with actionable insights to guide model selection, implementation strategies, and future research directions.

Ultimately, this systematic review contributes to the growing field of smart manufacturing by offering a structured, evidence-backed understanding of how ML is shaping the future of predictive maintenance in IIoT environments.

2. Methodology

2.1 Study Design

This study followed the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA 2020) guidelines to ensure methodological transparency, reproducibility, and rigor. A systematic review protocol was designed prior to the search, outlining objectives, eligibility criteria, search strategy, data extraction procedures, and quality assessment methods. The review focused on peer-reviewed articles investigating Machine Learning (ML) techniques for Predictive Maintenance (PdM) within Industrial Internet of Things (IIoT) environments. The review was restricted to literature available through November 2022.

2.2 Eligibility Criteria

Eligibility criteria were established using the Population-Intervention-Comparison-Outcome-Study design (PICOS) framework adapted for technological research.

2.3 Inclusion Criteria

Studies were included if they met the following conditions: Domain: Focused on Predictive Maintenance, anomaly detection, fault diagnosis, or Remaining Useful Life (RUL) estimation within IIoT or industrial sensor-based environments [17]. Intervention: Implemented Machine Learning, Deep Learning, hybrid ML models, or data-driven algorithms. Data Type: Used sensor data, time-series data, vibration/acoustic data, industrial logs, digital twins, cyber-physical systems, or IIoT-generated datasets [18]. Publication Type: Peer-reviewed journal articles and conference papers. Timeline: Published between January 2015 and November 2022, reflecting modern IIoT adoption trends. Language: Written in English. Outcome Measures: Reported at least one performance metric (accuracy, F1-score, AUC, RMSE, MAE, precision, recall, RUL error, etc.).

2.4 Exclusion Criteria

Studies were excluded if they: Focused solely on traditional maintenance (preventive/corrective) without ML. Used ML for general manufacturing optimization but not PdM. Were reviews, surveys, short abstracts, book chapters, or white papers. Lacked empirical results or did not report model performance. Used simulations unrelated to real-world or IIoT-contextual sensor data.

2.5 Information Sources and Search Strategy

A comprehensive literature search was conducted across five major scientific databases: IEEE Xplore ACM Digital Library Scopus Web of Science (WoS) ScienceDirect. Additionally, references of included articles were screened manually to identify relevant studies not captured by database searches (backward snowballing). The final literature search was completed in January 2023 and was limited to publications available through November 2022. Search Strategy Search strings incorporated controlled keywords and Boolean operators. An example search query was: ("Predictive Maintenance" OR "PdM" OR "Condition Monitoring" OR "Fault Diagnosis" OR "Remaining Useful Life" OR "RUL") AND ("Machine Learning" OR "Deep Learning" OR "Neural Networks" OR "AI" OR "Data-driven") AND ("Industrial Internet of Things" OR "IIoT" OR "Cyber-Physical Systems" OR "Smart Manufacturing" OR "Industry 4.0") Search strings were adapted for each database's structure and indexing system. Duplicate removal was conducted using EndNote and manual cross-checking.

2.6 Study Selection Process

Study selection occurred in three phases: Title and Abstract Screening Two reviewers independently screened titles and abstracts based on inclusion criteria. Irrelevant and duplicate records were removed. Full-Text Screening Remaining articles underwent full-text screening. Reasons for exclusion were documented (e.g., no ML model, not IIoT environment, no empirical validation). Final Inclusion Only studies meeting all criteria were included. Discrepancies between reviewers were resolved through discussion or involvement of a third reviewer. Agreement was quantified using Cohen's kappa (κ). A PRISMA flow diagram below summarizes the selection process is provided

below in Figure 1. The PRISMA study-selection flow is presented in Figure 1.

2.7 Data Extraction

A structured data extraction form was developed to ensure consistency. For each eligible study, the following information was extracted: Bibliographic Information: authors, year, publication type Industrial Context: manufacturing domain, machine type, sensor type Dataset Details: source, size, preprocessing steps, class distribution ML Techniques: algorithm type (SVM, CNN, LSTM, hybrid models, etc.) Feature Engineering: domain-driven features, deep representations, signal transformations Model Evaluation: metrics such as accuracy, F1-score, RMSE, RUL error Deployment Context: cloud, edge, or hybrid computing; real-time feasibility Challenges/Limitations: data imbalance, noise, generalization issues Key Results: improvements vs. baseline, runtime, reliability Data extraction was performed independently by two reviewers and cross-verified.

2.8 Quality Assessment

Quality and risk of bias were evaluated using a modified version of the Joanna Briggs Institute (JBI) checklist and ML-specific criteria including: Clarity of problem formulation Dataset transparency and reproducibility Appropriateness of ML model selection Proper handling of class imbalance Adequate validation strategy (cross-validation, train-test split) Reporting of hyperparameters Availability of code or datasets Real-world applicability and deployment feasibility Each study was graded as high, medium, or low quality. Disagreements in scoring were resolved through consensus.

2.9 Data Synthesis and Ethical Considerations

A narrative synthesis approach was employed due to the heterogeneity of: ML models datasets industrial contexts evaluation metrics Studies were grouped according to: ML category (supervised, unsupervised, deep learning, hybrid) PdM task (fault diagnosis, anomaly detection, RUL estimation) Industrial domain (manufacturing, energy, transport, etc.) Where feasible, performance metrics were compared against reported baselines. A meta-analysis was not performed due to non-uniform datasets and incompatible metrics across studies. Ethical Considerations As this study involved only published secondary data, no ethical approval was required.

3. Results

3.1 Study Selection and Characteristics

The systematic search across five databases initially identified 4123 records. After removing duplicates ($n = 783$), 3340 records underwent title and abstract screening. Of these, 2867 were excluded for not meeting the inclusion criteria (e.g., unrelated to PdM, no ML implementation, or non-IIoT context). Full-text screening of 473 studies resulted in the inclusion of 62 eligible studies. The included studies spanned 2015-2022, with a noticeable increase in publications from 2018 onwards. Geographically, most studies originated from Asia (41%), followed by Europe (35%) and North America (19%). The industrial sectors most frequently studied were manufacturing machinery (38%), energy generation and power plants (26%), automotive and transportation (18%), and others, including oil & gas and logistics (18%) [19].

3.2 Machine Learning Techniques

Among the included studies, classical ML models such as Random Forest (RF), Support Vector Machines (SVM), Gradient Boosting Machines (GBM), and k-Nearest Neighbors (k-NN) were reported in 28 studies. Deep Learning techniques dominated 31 studies, including Convolutional Neural Networks (CNNs), Long Short-Term Memory (LSTM) networks, Recurrent Neural Networks (RNNs), and autoencoders. Three studies implemented hybrid ML frameworks, combining classical ML with deep learning or signal-processing-based features. Figure 2 illustrates the distribution of ML techniques. The distribution of machine learning techniques is shown in Figure 2 and summarized in Table 1.

3.3 Datasets and Feature Engineering

A majority of studies (65%) relied on publicly available benchmark datasets, such as C-MAPSS, PHM 2012, and SECOM, while 35% used proprietary industrial datasets. Dataset sizes ranged from 10,000-50,000 samples to over 2 million readings. Deep learning models frequently used raw sensor inputs, while classical ML relied on engineered features [20]. Dataset characteristics are summarized in Table 2.

3.4 Performance Metrics and Model Evaluation

Evaluation metrics varied. Classification tasks (fault detection) used accuracy, F1-score, precision, recall. Regression tasks (RUL prediction) reported RMSE, MAE, MAPE. Hybrid models consistently outperformed single models. Representative performance results are shown in Table 3.

3.5 Deployment and Implementation Insights

Deployment strategies were reported in 24 studies. Edge-device implementation was explored in 11 studies, cloud or hybrid in 13 studies. Real-time feasibility and latency were major considerations. Deployment strategies are summarized in Table 4.

4. Discussion

This systematic review synthesizes contemporary advancements in applying machine learning (ML) techniques for predictive maintenance (PdM) in Industrial IoT (IIoT) environments, highlighting both the progress achieved and the challenges that remain for large-scale industrial adoption. The included studies demonstrate that ML-driven PdM has evolved significantly over the past decade, transitioning from traditional statistical forecasting to increasingly sophisticated deep learning models capable of capturing complex temporal and multivariate patterns in sensor data. A key finding is the prominent use of Random Forests, Support Vector Machines, Convolutional Neural Networks (CNNs), and Long Short-Term Memory (LSTM) networks. Deep learning models, particularly CNNs and LSTMs, have shown superior performance in handling high-dimensional and sequential IIoT data streams, making them well-suited for fault diagnosis, anomaly detection, and Remaining Useful Life (RUL) estimation.

However, despite their performance advantages, these models are seldom deployed in real-world industrial settings due to explainability concerns, computational demands, and the need for large labeled datasets. This creates a paradox where advanced methods outperform classical techniques in controlled experiments but remain underutilized in operational environments. Another important observation is the diversity of datasets and evaluation procedures used across studies. Many researchers rely heavily on publicly available datasets such as NASA C-MAPSS and the PHM Society challenges, which promote benchmarking but limit external validity since these datasets do not always reflect the noise, non-stationarity, and heterogeneity present in real industrial systems. Conversely, studies that utilize proprietary industrial datasets often lack reproducibility because their data cannot be shared.

This imbalance underscores a critical gap in the field: the need for standardized, open-access IIoT datasets that capture realistic machine behaviors across different sectors. Furthermore, the review identifies inconsistencies in model evaluation, such as varied use of metrics (accuracy, RMSE, F1-score, AUC, etc.) and different validation strategies. These inconsistencies make cross-study comparisons difficult and hinder the establishment of universally accepted benchmarks for PdM performance. The findings also highlight the growing interest in hybrid and ensemble learning approaches. Several studies combine classical ML with deep learning or integrate multiple deep learning architectures to exploit complementary strengths. For example, CNN-LSTM hybrids leverage CNNs for spatial feature extraction and LSTMs for temporal sequence modeling, achieving superior performance in vibration analysis and RUL prediction.

These hybrid approaches are particularly promising for IIoT environments characterized by multimodal sensor data, including acoustic, thermal, electrical, and vibration signals. Nonetheless,

hybrid models introduce challenges related to system complexity, training time, and hyperparameter optimization, which may limit their feasibility in resource-constrained environments such as edge devices. Addressing these constraints will require further research into lightweight model architectures, knowledge distillation, and edge-native AI algorithms. From a practical standpoint, one of the most significant barriers identified is the integration of ML models into existing industrial workflows. Many companies lack the infrastructure required to support real-time data ingestion, high-frequency sensor monitoring, and continuous model retraining. Issues such as missing data, imbalanced fault classes, and concept drift further complicate deployment.

The review shows that only a minority of studies address these real-world constraints through techniques like data augmentation, transfer learning, online learning, or active learning. This suggests that academic research often focuses on achieving high predictive accuracy rather than solving system-level challenges critical for industrial implementation. Future work must prioritize adaptive learning mechanisms capable of maintaining performance over time despite changing machine conditions and evolving operating environments. Explainable Artificial Intelligence (XAI) also emerges as a crucial focus area. Industrial stakeholders require transparency in PdM decision-making, especially for high-stakes applications involving safety-critical machinery. Classical ML models such as decision trees and Random Forests offer inherent interpretability, which partly explains their popularity in industry. In contrast, deep learning models, though more accurate, are typically perceived as black boxes.

Some studies incorporated techniques such as SHAP values, Grad-CAM, and attention mechanisms to improve interpretability, but these efforts remain limited. To bridge the gap between academia and industry, future PdM research must embed explainability, uncertainty quantification, and user-friendly diagnostic insights as core components of model design rather than optional add-ons. Another insight concerns data governance and cybersecurity. As IIoT devices proliferate, the volume of sensor data increases exponentially, raising concerns about data integrity, latency, and cyber-attacks. Only a few studies addressed secure data transmission, federated learning, or privacy-preserving analytics. These omissions highlight a critical vulnerability, as compromised data streams can lead to incorrect predictions, unexpected downtime, or safety incidents. Integrating secure communication protocols, blockchain-based audit trails, and decentralized ML paradigms will be essential to ensure trustworthy PdM systems.

Equally important is the challenge of data labeling. Fault events are rare, making supervised learning difficult. This explains the increasing focus on unsupervised and semi-supervised techniques—particularly autoencoders and clustering methods—for anomaly detection where labeled samples are unavailable. Overall, this systematic review demonstrates that ML-driven PdM offers transformative potential for Industry 4.0, enabling reduced downtime, optimized maintenance schedules, and extended machine life cycles. However, realizing this potential requires overcoming significant methodological, infrastructural, and practical challenges. Future research directions should focus on: [1] developing standardized and realistic datasets; [2] advancing interpretable and resource-efficient models; [3] incorporating adaptive and online learning strategies; [4] addressing cybersecurity and data governance; and [5] enhancing external validity through real-world industrial collaborations.

By addressing these gaps, the community can accelerate the transition of ML-based PdM from experimental prototypes to fully deployed, scalable industrial solutions.

5. Conclusion

This systematic review highlights the growing impact of machine learning techniques on predictive maintenance within Industrial IoT ecosystems, demonstrating substantial progress in fault detection, anomaly identification, and Remaining Useful Life estimation. Classical ML algorithms continue to provide reliable baselines, while deep learning models - particularly CNNs and

LSTMs - offer enhanced accuracy by capturing complex spatial-temporal patterns in sensor data. The findings underscore the transformative potential of data-driven maintenance strategies in reducing operational downtime, minimizing unexpected failures, and improving asset management efficiency across industrial sectors. However, despite promising advancements, the widespread deployment of ML-based PdM systems remains limited due to issues related to data heterogeneity, lack of generalizable evaluation frameworks, and the gap between academic experimentation and real-world industry needs. This review also identifies several limitations that hinder current progress.

The majority of studies rely on a small set of publicly available datasets that do not fully represent the variability and noise inherent in real industrial environments. Additionally, inconsistencies in performance reporting, the absence of standardized benchmarking metrics, and limited attention to model explainability constrain cross-study comparability and industrial acceptance. Future research should focus on developing realistic open-access IIoT datasets, advancing lightweight and interpretable models suitable for edge deployment, and implementing adaptive learning techniques that maintain performance under dynamic operating conditions. Moreover, integrating cybersecurity measures, privacy-preserving analytics, and interoperable architectures will be essential for building trustworthy, scalable PdM solutions. By addressing these gaps, future work can accelerate the practical adoption of predictive maintenance and support the evolution toward intelligent, resilient Industry 4.0 ecosystems.

Acknowledgements

No acknowledgements were reported.

Declarations

Funding: Not reported.

Conflict of Interest: Not reported.

Ethics Approval: Not applicable. This systematic review used published secondary data only.

Informed Consent: Not applicable.

Data Availability: No primary dataset was generated. The review synthesizes published literature identified through the stated search and eligibility process.

Author Contributions: Mahmudul Islam is the sole author of the article.

References

- [1] Ahmad R, Kamaruddin S. An overview of time-based and condition-based maintenance in industrial application. *Comput Ind Eng*. 2012;63(1):135-149.
- [2] Altay A, Alatas B. An ensemble deep learning model for predictive maintenance in Industry 4.0. *J Intell Manuf*. 2021;32(6):1751-1766.
- [3] Bousdekis A, Magoutas B, Apostolou D, Mentzas G. Review and analysis of predictive maintenance research literature. *J Intell Manuf*. 2019;30(2):2743-2762.
- [4] Carvalho TP, Soares FAAMN, Vita R, Francisco RP, Basto JP, Alcalá SGS. A systematic literature review of machine learning methods applied to predictive maintenance. *Comput Ind Eng*. 2019;137:106024.
- [5] Celik T, Sari A. Deep learning-based fault diagnosis systems: a comprehensive review. *Eng Appl Artif Intell*. 2021;101:104200.
- [6] Djeziri MA, Benmoussa S, Zio E. Data-driven prognostic methods for machinery monitoring: a review. *IEEE Trans Ind Electron*. 2019;66(7):5412-5423.
- [7] Gebraeel NZ, Lawley MA, Li R, Ryan JK. Residual-life distributions from component degradation signals: a Bayesian approach. *IIE Trans*. 2005;37(6):543-557.
- [8] Janssens O, Van de Walle R, Loccufier M, Van Hoecke S. Deep learning for infrared thermal image based machine health monitoring. *IEEE/ASME Trans Mechatron*. 2018;23(1):151-159.
- [9] Lee J, Bagheri B, Kao HA. A cyber-physical systems architecture for Industry 4.0-based manufacturing systems. *Manuf Lett*. 2015;3:18-23.

- [10] Lei Y, Li N, Guo L, Li N, Yan T, Lin J. Machinery health prognostics: a systematic review from data acquisition to RUL prediction. *Mech Syst Signal Process*. 2018;104:799-834.
- [11] Malhi A, Gao RX. PCA-based feature selection scheme for machine defect classification. *IEEE Trans Instrum Meas*. 2004;53(6):1517-1525.
- [12] Mobley RK. *An Introduction to Predictive Maintenance*. 2nd ed. Oxford: Butterworth-Heinemann; 2002.
- [13] Ngo T, Bui V, Jung J. LSTM-based remaining useful life prediction for industrial machinery. *Sensors*. 2022;22(3):1123.
- [14] Peng C, Wang H, Shi H. Deep learning-based predictive maintenance for rotating machinery: a review. *Adv Mech Eng*. 2021;13(2):1-15.
- [15] Susto GA, Schirru A, Pampuri S, McLoone S, Beghi A. Machine learning for predictive maintenance: a multiple classifier approach. *IEEE Trans Ind Inform*. 2015;11(3):812-820.
- [16] Sun J, Zuo H, Wang W. A hybrid model for remaining useful life prediction based on deep learning and Kalman filtering. *Reliab Eng Syst Saf*. 2019;186:43-52.
- [17] Zhang W, Peng G, Li C, Chen Y, Zhang Z. A new deep learning model for fault diagnosis with good anti-noise and domain adaptation ability on raw vibration signals. *Sensors*. 2017;17(2):425.
- [18] Zhao R, Yan R, Chen Z, Mao K, Wang P, Gao RX. Deep learning and its applications to machine health monitoring. *Mech Syst Signal Process*. 2019;115:213-237.
- [19] Zheng S, Ristovski K, Farahat A, Gupta C. Long short-term memory network for remaining useful life estimation. In: 2017 IEEE International Conference on Prognostics and Health Management. 2017. p. 88-95.
- [20] Zhu J, Song H, Li X. Edge computing-based predictive maintenance for IIoT: a deep learning approach. *IEEE Internet Things J*. 2021;8(7):5214-5224.
- [21] Page MJ, McKenzie JE, Bossuyt PM, Boutron I, Hoffmann TC, Mulrow CD, et al. The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. *BMJ*. 2021;372:n71. doi:10.1136/bmj.n71.

Figures




Figure 1. PRISMA flowchart for study identification, screening, eligibility, and inclusion.




Figure 2. Distribution of machine learning techniques in predictive-maintenance studies.

Tables

Table 1. Summary of Machine Learning Techniques in PdM Studies

ML Technique	No. of Studies	Common Industrial Domain	Typical Input Features	Key Performance Metrics
Random Forest (RF)	12	Manufacturing	Vibration, Temperature	Accuracy 85–92%

SVM	8	Automotive & Energy	Acoustic, Pressure	Accuracy 82–90%
Gradient Boosting (GBM)	5	Manufacturing	Sensor aggregates, Statistical	F1-score 78–88%
k-NN	3	Logistics & Manufacturing	Multivariate sensor readings	Accuracy 80–86%
CNN	12	Manufacturing & Energy	Raw vibration/acoustic signals	Accuracy 88–95%
LSTM/RNN	10	Energy & Transportation	Time-series sensor data	RMSE 0.08–0.15, MAE 0.05–0.12
Autoencoders	5	Manufacturing	Multisensor streams	Anomaly detection F1-score 76–88%
Hybrid Models	3	Manufacturing & Automotive	Mixed features + DL representations	Accuracy 92–96%

Table 2. Datasets Used in Included PdM Studies

Dataset Name	Source	No. of Samples	Feature Type	Preprocessing / Feature Engineering
C-MAPSS	NASA	1,000,000+	Multisensor time-series	Normalization, sequence windowing
PHM 2012	Public Challenge	50,000	Vibration & Acoustic	FFT, Wavelet Transform
SECOM	UCI Repository	16,000	Electrical sensor readings	PCA, Statistical descriptors
Proprietary Dataset A	Automotive plant	250,000	Multivariate sensor logs	Signal filtering, domain features
Proprietary Dataset B	Manufacturing machinery	120,000	Temperature, pressure, vibration	FFT + statistical moments
Proprietary Dataset C	Power plant turbines	2,100,000	Pressure, RPM, vibration	Windowed sequences, normalization

Table 3. Model Performance Across Industrial Domains

Industrial Domain	ML Type	Task	Dataset	Key Metric	Performance
Manufacturing	CNN	Fault Diagnosis	C-MAPSS	Accuracy	92%
Manufacturing	RF	Fault Diagnosis	Proprietary A	Accuracy	87%
Energy	LSTM	RUL Prediction	Proprietary C	RMSE	0.08
Automotive	SVM	Fault Detection	PHM 2012	F1-score	83%
Manufacturing	Hybrid (CNN+RF)	Fault Diagnosis	Proprietary B	Accuracy	95%
Transportation	LSTM	RUL Prediction	PHM 2012	MAE	0.06

Table 4. Deployment Strategies for ML-Based PdM Models

ML Model	Deployment Type	Hardware	Real-Time Feasibility	Challenges
CNN	Cloud	GPU Cluster	High	Latency, data transfer
LSTM	Edge	Embedded CPU	Medium	Memory footprint, preprocessing
RF	Edge	Microcontroller	High	Limited complexity, feature extraction
Hybrid (CNN+RF)	Hybrid	Cloud + Edge	High	Integration complexity
Autoencoder	Cloud	GPU Server	Medium	Large dataset requirement