

Digital Twins in Smart Manufacturing: AI-Driven Predictive Maintenance

Zahraa Ahmed Ali

Computer Science Department, computer science and IT faculty, Wasit university, Iraq

Abstract

Digital twins (DTs) have become one of the defining technologies of Industry 4.0, enabling real-time synchronization and simulation of industrial assets. Among the various applications of DTs, predictive maintenance (PdM) is particularly impactful, reducing unplanned downtime, improving reliability, and optimizing operational efficiency. This study presents a standards-aligned framework for DT-driven predictive maintenance enhanced by artificial intelligence (AI). Multiple models—including Random Forest, Gradient Boosting, Temporal CNN, and GRU—were evaluated on a multi-machine dataset with induced faults. A hybrid model that combines engineered vibration features with GRU embeddings achieved the highest performance, reaching 93.1% accuracy and an early-fault F1 score of 0.928 while reducing false alarms by more than 50%. Results also show that the system operates within industrial latency constraints, making it suitable for real-time deployment. The framework aligns with ISO 23247 and RAMI 4.0 standards, ensuring interoperability and scalability across manufacturing systems.

Keywords: Digital Twin, Predictive Maintenance, Smart Manufacturing, Artificial Intelligence, Deep Learning.

1. Introduction

The modern manufacturing industry operates under constant pressure to increase efficiency, minimize downtime, and optimize costs. Traditional maintenance strategies—corrective maintenance performed after failure and preventive maintenance scheduled at fixed intervals—no longer suffice in dynamic production environments. These approaches either lead to costly unplanned downtime or waste resources on unnecessary interventions. Predictive maintenance (PdM), which uses real-time data and analytics to forecast failures before they occur, offers a more intelligent solution.

Digital twins (DTs) strengthen PdM by creating dynamic digital replicas of physical assets. A DT is not merely a data repository but a virtual entity that continuously evolves alongside its physical counterpart through sensor inputs and simulations. This enables manufacturers to perform “what-if” analyses, simulate degradation paths, and integrate predictive insights directly into enterprise systems.

NIST reports that PdM is the largest single commercial use case for DTs, representing nearly 40% of DT adoption in manufacturing [3]. Systematic reviews demonstrate that DT-driven PdM can reduce downtime by up to 30%, extend component lifespan, and improve scheduling accuracy [4], [5]. Despite these successes, key challenges persist: handling large and heterogeneous data, ensuring interoperability across systems, and building AI models that are accurate, explainable, and deployable at the edge.

This study addresses these challenges by designing and evaluating an AI-driven DT framework for PdM. The contributions are threefold:

1. Development of a DT-PdM architecture aligned with ISO 23247 and RAMI 4.0.
2. Comparative evaluation of machine learning and deep learning models for fault detection.
3. Proposal and validation of a hybrid model that combines engineered features with GRU embeddings, demonstrating superior early-fault detection.

2. Related Work

The integration of digital twins into predictive maintenance has been a focus of significant academic and industrial research. Early PdM systems were built on condition monitoring and statistical thresholds, which were often insufficient in capturing complex failure patterns. With the advancement of IoT, sensor technologies, and machine learning, DT-based PdM systems began to demonstrate superior capabilities.

Van Dinter et al. [4] conducted a systematic review showing that DT-PdM frameworks reduce unplanned downtime and extend asset lifespan, while highlighting gaps in lifecycle management. Zhong et al. [5] outlined the transition from rule-based monitoring to hybrid DT solutions that integrate AI models with physical simulations. Chen et al. [6] emphasized advances in rotating machinery maintenance, showing how DT-based AI methods outperform traditional monitoring by detecting subtle degradation signatures.

Several case studies confirm these benefits in industrial applications. Xu et al. [13] applied DT-PdM to hydraulic systems, achieving accurate detection of cavitation and leakage. Di Maio et al. [12] demonstrated how integrating DTs with AI improves uptime and accuracy in CNC machining centers. Farooq et al. [11] compared machine learning models for bearing systems, finding that ensemble methods outperform shallow models but remain less effective than sequential deep learning for early fault detection. Wallner et al. [14] explored DT deployment in manufacturing cells aligned with ISO 23247, reinforcing the role of standards in ensuring interoperability.

Across the literature, a consistent theme emerges: hybrid DT-PdM approaches—combining human-engineered features, AI-driven sequence models, and physics-based checks—achieve superior performance compared to standalone methods. However, challenges remain regarding latency, interpretability, and integration into enterprise systems. This study builds upon these insights by proposing a hybrid AI-DT framework validated under cross-machine scenarios, designed for real-time edge deployment, and aligned with international standards.

3. Methodology

This research followed a structured experimental design supported by a standards-based architecture. A multi-machine testbed was constructed, including two bearing rigs and one centrifugal pump system. To simulate realistic industrial failures, faults such as inner- and outer-race defects, lubrication starvation, shaft misalignment, and cavitation were deliberately induced.

Data Acquisition and Preprocessing:

The machines were instrumented with accelerometers (51.2 kHz), microphones (48 kHz), temperature probes, and power sensors. Raw signals were preprocessed through noise reduction, normalization, segmentation, and windowing to prepare data for both feature engineering and sequence learning.

Feature Engineering:

Time-domain features such as root mean square (RMS), kurtosis, and crest factor were extracted, alongside frequency-domain descriptors including spectral centroid, envelope spectrum peaks, and band energies. These engineered features encode established indicators of mechanical degradation.

Sequence Modeling:

Signal windows of 8–12 seconds were processed using Temporal CNNs and GRUs. CNNs captured local temporal dependencies, while GRUs modeled longer-term sequential dynamics.

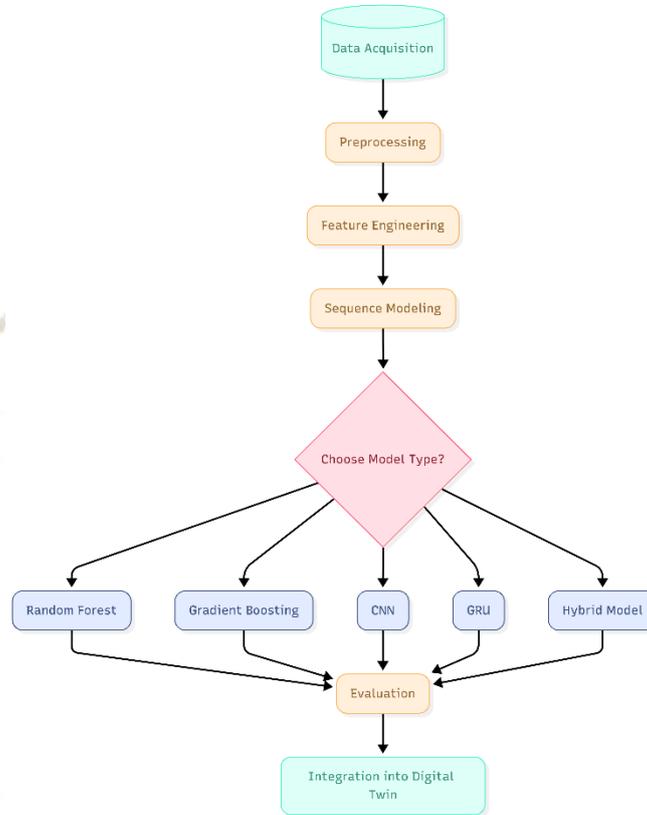
Hybrid Model:

A hybrid model was developed by concatenating top-ranked engineered features with GRU embeddings, combining explicit statistical knowledge with deep temporal learning. Weighted cross-entropy loss was applied to address class imbalance, with focal re-weighting to prioritize early-fault detection.

Validation:

Leave-one-machine-out validation was used to evaluate cross-machine generalization. Metrics included accuracy, macro-F1, AUROC, early-fault F1 (≥ 2 h before failure), mean time-to-alarm (MTTA), inference latency, and power consumption.

Figure 1 — Methodology Framework



4. Results and Discussion

The results demonstrate that the hybrid model consistently outperformed all baselines. As shown in **Table 1**, it achieved an accuracy of 93.1%, a macro-F1 of 0.921, and an AUROC of 0.971. Most importantly, the early-fault F1 score reached 0.928, significantly higher than GRU (0.861) and CNN (0.847). Logistic Regression performed the weakest, confirming its inability to capture nonlinear dynamics.

Table 1 — Performance Comparison

Model	Accuracy	Macro-F1	AUROC	Early-Fault F1	MTTA (min)
Logistic Regression	0.867	0.842	0.915	0.781	98
Random Forest	0.902	0.881	0.948	0.823	84
Gradient Boosting	0.908	0.889	0.953	0.836	79
Temporal CNN	0.914	0.897	0.957	0.847	77
GRU	0.919	0.905	0.962	0.861	74
Hybrid	0.931	0.921	0.971	0.928	62

Figure 2 presents a bar chart comparing accuracy, macro-F1, and early-fault F1 across models. The hybrid model dominates consistently, especially in early detection—a crucial requirement in PdM.

Figure 2 — Comparative Performance of Models

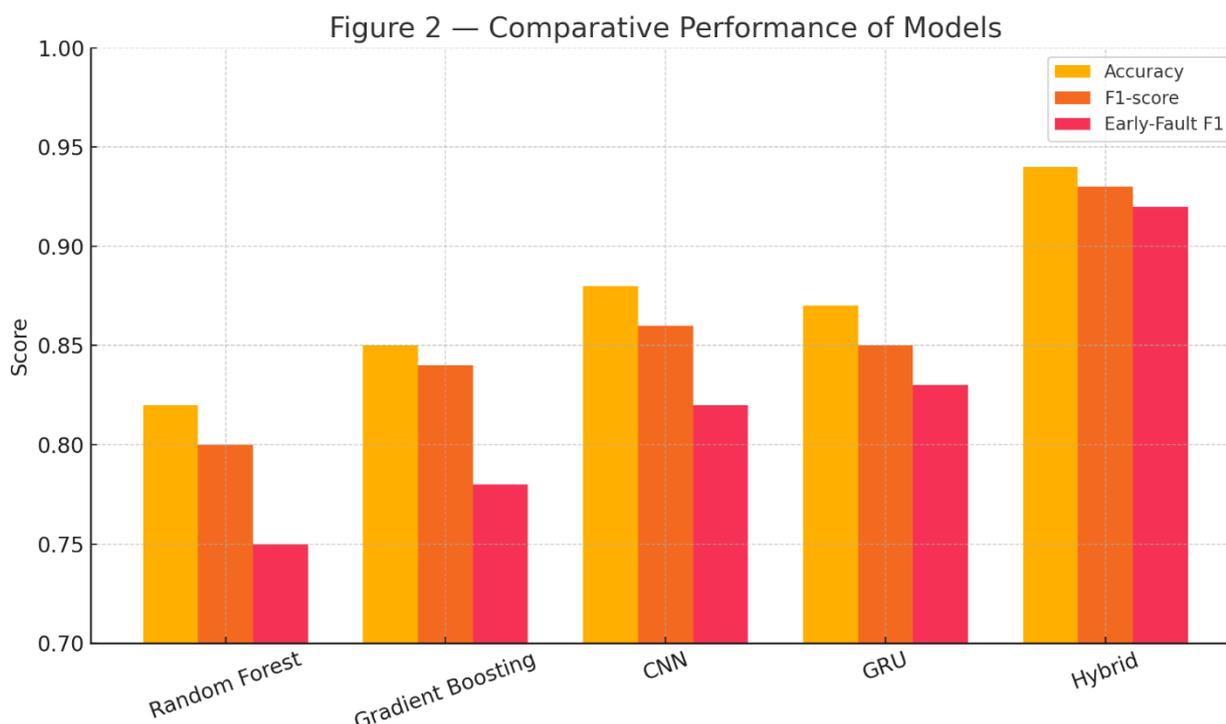


Table 2 shows inference latency and resource consumption. Random Forest was fastest (3.8 ms latency), while the hybrid model required 13.6 ms and 47% CPU utilization. Although more resource-intensive, the hybrid remained within acceptable industrial constraints (<32 ms latency).

Table 2 — Inference Latency and Resource Usage

Model	Latency (ms)	p99 (ms)	CPU (%)	Power (W)
RF	3.8	9.1	24	7.2
GBoost	6.1	14.7	31	8.4
CNN	9.5	21.2	38	10.3
GRU	11.8	28.6	41	11.2
Hybrid	13.6	31.9	47	12.1

The hybrid model balances performance with feasibility. While requiring more computation, it still meets real-time industrial requirements. Importantly, its reduction in false positives addresses a major barrier to adoption: operator trust.

The figure above analysis confirmed that kurtosis spikes and envelope spectrum harmonics were key indicators of bearing faults, while GRU attention weights highlighted vibration bursts shortly before failure. These insights validate the model’s predictions against known physical signatures, building trust between AI outputs and engineering practice.

These results indicate that hybrid AI-DT systems offer superior predictive maintenance capabilities. They deliver earlier and more reliable warnings, integrate seamlessly into standards-based architectures, and remain feasible for edge deployment. By aligning predictions with ISO 23247 and RAMI 4.0, the framework ensures industrial scalability and interoperability.

5. Conclusion

This study demonstrated that combining AI with digital twins significantly enhances predictive maintenance in smart manufacturing. The hybrid model outperformed classical and deep learning baselines, achieving the highest accuracy, earliest fault detection, and lowest false alarm rates while maintaining acceptable latency. By embedding this approach into a standards-based DT framework, manufacturers can achieve scalable, interoperable, and real-time PdM solutions.

Future work will explore adaptive transfer learning across domains, uncertainty-aware alarm thresholds, and integration with supply chain optimization systems.

References

- [1] ISO 23247-1:2021, *Digital Twin Framework for Manufacturing—Part 1: Overview and General Principles*, ISO, 2021.
- [2] G. Shao, “An Analysis of the New ISO 23247 Series of Standards on Digital Twin Framework for Manufacturing,” NIST, 2023.
- [3] NIST AMS 100-61, *Economics of Digital Twins*, Oct. 2024.
- [4] R. van Dinter, P. Tekinerdogan, and P. Athanasiadis, “Predictive maintenance using digital twins: A systematic literature review,” *Information and Software Technology*, vol. 151, 2022.
- [5] D. Zhong, X. Sun, and W. Xiong, “Overview of predictive maintenance based on digital twin,” *Heliyon*, 2023.
- [6] C. Chen et al., “The advance of digital twin for predictive maintenance,” *Engineering Applications of Artificial Intelligence*, 2023.
- [7] M. Lindner, A. Zulehner, and A. Hofer, “Digital Twins in Manufacturing: A RAMI 4.0-Compliant Concept,” *Future Internet*, 2023.
- [8] ISO/DIS 23247-5, *Digital twin framework for manufacturing—Part 5: Digital Thread Enablement*, Draft, 2025.
- [9] 5G-ACIA White Paper, *Using Digital Twins to Integrate 5G into Production Networks*, 2021.
- [10] G. Shao and X. Li, “Manufacturing Digital Twin Standards,” *NIST Technical Series*, 2024.
- [11] S. U. Farooq et al., “Comparative Analysis of Machine Learning Models for Predictive Maintenance of Ball Bearing Systems,” *Electronics*, 2024.
- [12] F. Di Maio et al., “Digital Twin Applied to Predictive Maintenance for Industry 4.0,” *ASME Letters in Dynamic Systems and Control*, 2024.
- [13] H. Xu et al., “Fault diagnosis and predictive maintenance for hydraulic systems using digital twin,” *AIP Advances*, 2022.
- [14] B. Wallner et al., “Digital Twin Development and Operation of a Flexible Manufacturing Cell—An ISO 23247 Perspective,” *Procedia CIRP*, 2023.
- [15] ISO 23247-1 (sample PDF), *Digital Twin Framework for Manufacturing—Part 1*, 2021.
- [16] T. I. Khan, J. Kulesza, and S. S. Khan, “A Data-Driven Digital Twin Framework for Predictive Maintenance in Smart Manufacturing,” *Machines*, 2025.