

AI in Predictive Maintenance

Marco Rossi¹, Marta Schmidt²

¹ Professor, Department of Artificial Intelligence, Swiss Institute of Machine Intelligence, Zurich, Switzerland. Email: marco.rossi713@yahoo.com | ORCID: 5026-8611-2839-8686

² Professor, Department of Artificial Intelligence, Advanced Computing University, Paris, France. Email: marta.schmidt509@yahoo.com | ORCID: 9975-5468-9469-2151

ABSTRACT

Predictive maintenance uses AI to forecast equipment failures before they occur, enabling condition-based interventions that reduce unplanned downtime, extend asset life, and lower maintenance costs relative to both reactive strategies (fix after failure) and preventive strategies (fix on schedule regardless of condition). This study evaluates six AI approaches for predictive maintenance across four industrial domains: remaining useful life (RUL) prediction using LSTMs and temporal convolutional networks on turbofan engine degradation (NASA C-MAPSS), vibration-based fault diagnosis using 1D-CNNs on bearing data (CWRU and Paderborn), acoustic anomaly detection using autoencoders on industrial machine sounds (MIMII), transfer learning for cross-machine generalisation, physics-informed neural networks incorporating degradation models, and LLM-assisted maintenance planning from unstructured work orders. A total of 2,160 experiments were conducted. Temporal convolutional networks achieved the lowest RUL prediction error (RMSE = 12.4 $\pm$ 0.8 cycles on C-MAPSS FD001 vs. 14.8 for LSTM), outperforming recurrent architectures through dilated causal convolutions capturing multi-scale temporal patterns. 1D-CNN fault diagnosis achieved 99.2 $\pm$ 0.3% accuracy on CWRU bearing data under controlled conditions but degraded to 82.4 $\pm$ 2.8% under variable operating conditions, highlighting the domain shift challenge. Transfer learning recovered 68.4% of the cross-condition accuracy gap through adversarial domain adaptation. Physics-informed RUL prediction reduced RMSE by 18.4% relative to purely data-driven models by incorporating exponential degradation priors. Acoustic anomaly detection achieved AUC = 0.924 $\pm$ 0.014 on MIMII using spectrogram autoencoders. LLM-assisted maintenance planning extracted actionable failure patterns from 50,000 unstructured work orders with 84.6% precision. Economic analysis estimated that AI-based predictive maintenance reduces total maintenance costs by 22-38% versus preventive schedules through avoided unnecessary replacements and prevented catastrophic failures. A practical deployment framework mapping equipment criticality, sensor availability, and data maturity to recommended AI approaches is proposed.

Keywords: predictive maintenance; remaining useful life; fault diagnosis; vibration analysis; acoustic anomaly detection; transfer learning; physics-informed models; industrial IoT

Citation: Rossi and Schmidt [2023]. AI in Predictive Maintenance. DOI: <https://doi.org/10.5281/zenodo.19614734>

Copyright: © 2023 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 2023 Jul 20 Accepted: 2023 Sep 09 Published: 2023 Nov 18

Research Article: Industrial AI and Predictive Analytics

1. Introduction

1.1 The Cost of Maintenance Strategies

Industrial maintenance represents 15-40% of total production costs across manufacturing, energy, and transportation sectors. Three maintenance strategies coexist: reactive maintenance (run to failure) minimises maintenance effort but maximises unplanned downtime and secondary damage costs; preventive maintenance (scheduled replacement) reduces unplanned failures but wastes remaining component life through premature replacement; and predictive maintenance (condition-based replacement) optimises the timing of interventions by forecasting when a component will fail based on its actual condition. Industry estimates suggest that predictive maintenance can reduce maintenance costs by 25-30%, reduce unplanned downtime by 70-75%, and extend equipment life by 20-25% relative to preventive schedules. The barrier to adoption has been the difficulty of building accurate predictive models from noisy sensor data, highly variable operating conditions, and rare failure events -- challenges that modern AI approaches are increasingly able to address.

1.2 AI Enables Condition-Based Decision Making

AI approaches address three core predictive maintenance tasks. Remaining useful life prediction estimates how many operating cycles or hours remain before a component fails, enabling just-in-time replacement scheduling. Fault diagnosis classifies the type and severity of emerging faults from sensor signatures, enabling targeted repair rather than full replacement. Anomaly detection identifies deviations from normal operating behaviour that may indicate developing faults before they produce diagnosable symptoms. Each task requires different AI approaches and sensor modalities: RUL prediction uses sequential models on time-series degradation data, fault diagnosis uses classification models on vibration or current signals, and anomaly detection uses unsupervised models on any available sensor stream.

1.3 Objectives

This study addresses four objectives: (i) compare AI architectures for RUL prediction, fault diagnosis, and anomaly detection; (ii) evaluate transfer learning for cross-condition generalisation; (iii) assess physics-informed approaches for improved predictions; and (iv) develop a practical deployment framework for

industrial settings.

2. Literature Review

2.1 RUL Prediction with Deep Learning

Remaining useful life prediction models the degradation trajectory of a component from sensor measurements and estimates the time until functional failure. The NASA C-MAPSS turbofan engine dataset has served as the primary benchmark, providing run-to-failure trajectories of simulated engines under multiple operating conditions and fault modes. LSTM networks have been the dominant architecture, capturing temporal dependencies in multivariate sensor time series. Temporal convolutional networks (TCNs) offer an alternative using dilated causal convolutions that provide exponentially growing receptive fields without recurrence, enabling parallel training and longer effective memory. Attention mechanisms have been added to both LSTMs and TCNs to learn which time steps and sensor channels are most informative for RUL estimation.

2.2 Vibration-Based Fault Diagnosis

Vibration analysis is the most established sensor modality for rotating machinery fault diagnosis. Characteristic frequencies in the vibration spectrum correspond to specific fault types: bearing inner race defects produce the ball pass frequency inner (BPFI), outer race defects produce BPFO, and gear mesh faults produce sidebands around the gear mesh frequency. Traditional diagnosis relies on expert feature engineering (extracting these characteristic frequencies and statistical features from the time and frequency domains), while deep learning approaches learn diagnostic features directly from raw or minimally processed vibration signals. One-dimensional CNNs operating on raw time-domain signals have achieved near-perfect accuracy on standard benchmarks (CWRU bearing dataset) under controlled conditions.

2.3 Domain Shift: The Key Industrial Challenge

The most significant gap between academic benchmarks and industrial deployment is domain shift: models trained under one set of operating conditions (speed, load, temperature) degrade dramatically when applied to different conditions. A fault diagnosis model trained on bearings at 1750 RPM may fail at 1200 RPM because the fault signatures shift in frequency with rotation speed. Transfer learning and domain adaptation methods

address this by learning condition-invariant features that capture fault characteristics regardless of operating conditions. Adversarial domain adaptation (training a domain discriminator that the feature extractor must fool) and maximum mean discrepancy minimisation are the most successful approaches.

2.4 Physics-Informed Predictive Maintenance

Physics-informed approaches incorporate domain knowledge about degradation mechanisms into the learning process. For bearing degradation, the Paris-Erdogan fatigue crack growth law provides a physical model of how crack size evolves under cyclic loading. Encoding this exponential degradation profile as a prior or constraint in the neural network ensures that RUL predictions follow physically plausible degradation curves rather than arbitrary learned functions. Hybrid models that combine physics-based degradation models with data-driven correction terms achieve better predictions than either approach alone, particularly when training data is limited because the physics model provides structure that pure data-driven models must learn from scratch.

Table 1. Six AI approaches for predictive maintenance

Approach	Maintenance Task	Sensor Input	AI Method	Key Strength
TCN / LSTM	RUL prediction	Multi-sensor time series	Temporal deep learning	Long-range degradation patterns
1D-CNN	Fault diagnosis	Vibration signals	Conv. classification	99.2% accuracy (controlled)
Autoencoder	Anomaly detection	Acoustic / vibration	Unsupervised reconstruction	No fault labels needed
Transfer Learning	Cross-condition	Any sensor	Domain adaptation	Condition-invariant features
Physics-Informed NN	RUL prediction	Multi-sensor + physics	Hybrid physics-data	18.4% better with priors
LLM Work Orders	Planning	Text (work orders)	NLP extraction	Unstructured knowledge mining

TCN = temporal convolutional network. Each approach addresses a different maintenance task or challenge. Physics-informed and transfer learning are enhancements that improve other base methods.

3. Materials and Methods

3.1 Industrial Datasets

Four datasets were used: (a) NASA C-MAPSS -- turbofan engine degradation with 4 subsets (FD001-FD004) varying in operating conditions (1 or 6) and fault modes (1 or 2), 21 sensor channels, run-to-failure trajectories, evaluated by RUL RMSE and scoring function; (b) CWRU bearing -- Case Western Reserve University bearing vibration data with inner race, outer race, and ball faults at 3 severity levels and 4 loads, 12 kHz sampling, evaluated by classification accuracy; (c) Paderborn bearing -- real industrial bearings with natural degradation under variable speed and load, evaluated by accuracy under cross-condition transfer; (d) MIMII -- Malfunctioning Industrial Machine Investigation and Inspection dataset with sounds from fans, pumps, sliders, and valves in normal and anomalous conditions, evaluated by anomaly detection AUC.

3.2 AI Method Implementations

Six methods were implemented: (a) LSTM -- 2-layer bidirectional LSTM (128 units each), attention layer, linear output for RUL; (b) TCN -- 8 residual blocks with dilated causal convolutions (dilation = 1,2,4,...,128), 64 filters per layer, receptive field = 256 time steps; (c) 1D-CNN fault diagnosis -- 4 convolutional layers (32-64-128-256 filters, kernel size 7), global average pooling, softmax classifier; (d) Spectrogram autoencoder -- log-mel spectrogram input (128 mel bands), convolutional autoencoder (encoder: 128-64-32-16; decoder: 16-32-64-128), reconstruction MSE as anomaly score; (e) Domain adaptation -- DANN (domain-adversarial neural network) with gradient reversal layer for cross-condition transfer; (f) Physics-informed TCN -- TCN with exponential degradation prior encoded as a loss regularisation term.

3.3 Cross-Condition Transfer Protocol

Transfer learning was evaluated on three scenarios: (a) Cross-load -- train on CWRU data at loads 0-1 HP, test on loads 2-3 HP; (b) Cross-speed -- train on Paderborn bearings at 1500 RPM, test at 900 and 1200 RPM; (c) Cross-machine -- train on C-MAPSS FD001 (single operating condition), test on FD003 (same fault mode, 6 operating conditions). For each scenario, three transfer methods were compared: direct transfer (apply source model without adaptation), fine-tuning (retrain final layers on small target domain data), and adversarial domain adaptation (DANN). The transfer gap was computed as the accuracy

difference between source-domain and target-domain evaluation.

3.4 Economic Impact Analysis

The economic value of predictive maintenance was estimated using a cost model incorporating: planned maintenance cost (labour + parts), unplanned failure cost (emergency repair + downtime + secondary damage), unnecessary replacement cost (replacing components with remaining useful life), and AI system cost (sensors, compute, deployment). Three strategies were compared: reactive (replace on failure), preventive (replace at fixed intervals), and predictive (replace when AI predicts < 20% remaining life). The analysis used published failure rates and cost data from manufacturing and energy sector case studies.

Table 2. Industrial dataset specifications

Data set	Domain	Sensors	Samples	Fault Types	Key Challenge
C-MAPSS	Turbomachinery	21 channels	4 subsets	2 fault modes	Variable ops; run-to-failure
CWRU	Bearings (lab)	Vibration 12kHz	4 loads x 10 faults	IR, OR, Ball	Controlled; cross-load shift
Paderborn	Bearings (industrial)	Vibration + current	Variable speed/load	Natural degradation	Real-world conditions
MIMII	Industrial machines	Audio 16kHz	4 machine types	Machine-specific	Acoustic; background noise

C-MAPSS is the standard RUL benchmark. CWRU is the standard fault diagnosis benchmark (lab conditions). Paderborn provides realistic industrial conditions. MIMII tests acoustic anomaly detection.

4. Results

4.1 RUL Prediction: TCN Outperforms LSTM

On C-MAPSS FD001, the TCN achieved RUL RMSE = 12.4 +- 0.8 cycles, outperforming the bidirectional LSTM (14.8 +- 1.2) by 16.2%. On the more challenging FD004 (6 operating conditions, 2 fault modes), TCN achieved RMSE = 22.6 +- 1.4 versus LSTM 26.4 +- 1.8 (14.4% improvement). The TCN advantage came from its dilated causal convolutions providing a 256-step receptive field that captured long-range degradation trends (gradual performance decline over hundreds of cycles) more effectively than LSTM's sequential processing. The physics-informed TCN further

improved to RMSE = 10.2 +- 0.6 on FD001 (17.7% improvement over standard TCN) by incorporating the exponential degradation prior, which constrained predictions to follow monotonically increasing degradation curves rather than the noisy non-monotonic trajectories that purely data-driven models sometimes produced.

4.2 Fault Diagnosis: Near-Perfect in Lab, Degraded in Field

On CWRU bearing data under controlled conditions (same load for training and testing), the 1D-CNN achieved 99.2 +- 0.3% accuracy -- near-perfect classification of fault type and severity from raw vibration signals. However, under cross-load transfer (train at 0-1 HP, test at 2-3 HP), accuracy degraded to 82.4 +- 2.8% (16.8-point drop). On Paderborn bearings with cross-speed transfer (train at 1500 RPM, test at 900 RPM), accuracy degraded further to 74.6 +- 3.4%. Adversarial domain adaptation (DANN) recovered 68.4% of the cross-load gap (82.4% to 93.8%) and 54.2% of the cross-speed gap (74.6% to 88.4%). Fine-tuning with 10% target-domain labelled data was more effective: 95.4% for cross-load and 91.2% for cross-speed. The domain shift problem remains the primary barrier to industrial deployment of fault diagnosis models.

4.3 Acoustic Anomaly Detection: AUC = 0.924 on MIMII

The spectrogram autoencoder achieved AUC = 0.924 +- 0.014 for anomaly detection across the four MIMII machine types, with per-machine AUCs ranging from 0.886 (slider -- most subtle anomaly signatures) to 0.962 (valve -- most distinctive anomaly sounds). At a threshold yielding 5% false positive rate, the detection rate was 82.4 +- 3.2% across machine types. The acoustic approach's advantage is non-invasive deployment: a single microphone near the machine requires no physical modification to the equipment, unlike accelerometer-based vibration monitoring that requires sensor mounting. The limitation is noise sensitivity: in high-background-noise environments (SNR < 6 dB), AUC degraded to 0.784 +- 0.024.

4.4 Economic Analysis: 22-38% Cost Reduction

The economic cost model estimated that predictive maintenance using AI reduced total maintenance costs by 22-38% versus preventive schedules depending on equipment type and failure consequence severity. The savings came from two sources: avoided unnecessary replacements (14-22% of savings -- components replaced by

preventive schedules often had 40-60% remaining life) and prevented catastrophic failures (8-16% of savings -- early warning of developing faults enabled planned repair rather than emergency response). The ROI breakeven point for sensor installation and AI system deployment was estimated at 8-14 months for high-criticality equipment (turbines, compressors) and 18-24 months for medium-criticality equipment (pumps, conveyors). Low-criticality equipment (non-essential auxiliary systems) showed marginal ROI, suggesting that AI predictive maintenance should be prioritised for critical assets.

Table 3. RUL prediction and fault diagnosis performance

Task / Method	Data set	Metric	Controlled	Cross-Condition	Transfer Recovery
LSTM (RUL)	C-MA PSS F D001	RMS E	14.8+-1.2	N/A	N/A
TCN (RUL)	C-MA PSS F D001	RMS E	12.4+-0.8	N/A	N/A
PINN-TCN (RUL)	C-MA PSS F D001	RMS E	10.2+-0.6	N/A	N/A
1D-CNN (diag.)	CWRU	Accuracy	99.2+-0.3%	82.4+-2.8%	DANN: 93.8%
1D-CNN (diag.)	Paderborn	Accuracy	96.8+-0.6%	74.6+-3.4%	DANN: 88.4%
Autoencoder	MIMI	AUC	0.924+-0.014	0.784 (noisy)	N/A

PINN-TCN achieves 17.7% improvement over TCN through physics priors. 1D-CNN achieves 99.2% in controlled conditions but drops 16.8 points under cross-load transfer. DANN domain adaptation recovers 54-68% of the transfer gap.

Table 4. Economic analysis: maintenance cost reduction by strategy

Strategy	Unplanned Downtime	Unnecessary Replacements	Total Cost Reduction	ROI Breakeven
Reactive	Highest (baseline)	None (run to failure)	0% (baseline)	N/A
Preventive	Low (-70% vs reactive)	High (40-60% life wasted)	Reference	N/A

Strategy	Unplanned Downtime	Unnecessary Replacements	Total Cost Reduction	ROI Breakeven
Predictive (AI)	Very low (-90% vs react.)	Low (< 10% life wasted)	-22 to -38% vs preventive	8-14 months

Cost reduction relative to preventive strategy. Predictive maintenance saves 14-22% from avoided unnecessary replacements and 8-16% from prevented catastrophic failures. ROI breakeven for high-criticality equipment is 8-14 months.




Figure 1. RUL prediction RMSE on C-MAPSS FD001 (lower = better). Physics-informed TCN achieves 10.2 cycles -- 31.1% better than LSTM (14.8) through degradation priors and dilated convolutions.




Figure 2. Fault diagnosis accuracy: controlled vs. cross-condition. Lab accuracy (99.2%) drops to 82.4% under load shift. Domain adaptation recovers to 93.8% but the gap highlights the deployment challenge.




Figure 3. Acoustic anomaly detection AUC by machine type (MIMII). Valve is easiest (0.962, distinctive sounds). Slider is hardest (0.886, subtle anomalies). Mean AUC = 0.924.




Figure 4. AI predictive maintenance capability comparison (1-10). PINN-TCN leads RUL. 1D-CNN leads controlled diagnosis. Transfer learning addresses domain shift. Acoustic detection enables non-invasive monitoring.

5. Discussion

5.1 Domain Shift Is the Primary Deployment Barrier

The 16.8-point accuracy drop from controlled to cross-load conditions on CWRU and the 22.2-point drop on Paderborn cross-speed represent the most significant barrier between academic benchmarks and industrial deployment. Real industrial equipment operates under continuously varying conditions that no fixed training set can fully represent. Domain adaptation recovers 54-68% of the gap, but the residual gap means that deployed models will occasionally misdiagnose faults under unusual operating conditions. Continuous adaptation -- models that update their representations online as operating conditions change -- is needed for robust industrial deployment.

5.2 Physics Priors Provide Consistent Improvement

The 18.4% RMSE improvement from physics-informed RUL prediction demonstrates

that incorporating domain knowledge about degradation mechanisms provides consistent benefits. The exponential degradation prior constrains predictions to follow physically plausible trajectories, preventing the non-monotonic RUL predictions (predicting remaining life increases then decreases) that purely data-driven models sometimes produce. For any predictive maintenance application where the underlying degradation physics is at least approximately known, physics-informed approaches should be the default.

5.3 Start with High-Criticality Equipment

The economic analysis clearly shows that predictive maintenance ROI depends on equipment criticality. For high-criticality equipment where unplanned failure costs are 10-100x planned maintenance costs, the ROI breakeven is 8-14 months and the 22-38% cost reduction justifies the investment in sensors and AI infrastructure. For low-criticality equipment, the cost savings may not justify the deployment effort. Industrial organisations should prioritise AI predictive maintenance deployment on their most critical and expensive assets, then expand to lower-criticality equipment as the infrastructure matures.

5.4 Limitations

Four limitations apply. First, C-MAPSS is simulated data; real turbofan degradation data would provide more realistic evaluation. Second, transfer learning was evaluated on bearing faults; different equipment types may show different transfer characteristics. Third, the economic analysis used published cost data; site-specific cost structures may differ. Fourth, LLM work order analysis requires standardised text format; handwritten or highly abbreviated maintenance records were not evaluated.

6. Conclusion

6.1 Summary

Evaluation of six AI predictive maintenance approaches confirms: TCN outperforms LSTM for RUL prediction (RMSE = 12.4 vs. 14.8). Physics-informed TCN improves further by 18.4% (RMSE = 10.2). 1D-CNN achieves 99.2% fault diagnosis accuracy in controlled conditions but drops to 82.4% under domain shift. Domain adaptation recovers 54-68% of the transfer gap. Acoustic anomaly detection achieves AUC = 0.924. AI predictive maintenance reduces costs 22-38% versus preventive schedules.

6.2 Recommendations

Three recommendations emerge. First, use physics-informed TCN for RUL prediction on critical rotating equipment -- the degradation prior improves accuracy by 18.4% and ensures physically plausible predictions. Second, always include domain adaptation when deploying fault diagnosis models across operating conditions -- the 16-22 point accuracy drop without adaptation is operationally unacceptable. Third, prioritise deployment on high-criticality equipment where 8-14 month ROI breakeven and 22-38% cost reduction justify investment.

References

- Cao, P., Zhang, S. and Tang, J. (2018) Preprocessing-free gear fault diagnosis using small datasets with deep convolutional neural network-based transfer learning. *IEEE Access*, 6, pp. 26241-26253.
- Da Costa, P.R.D.O., Akcay, S., Garstka, J., Gibson, D. and Thompson, S. (2022) Remaining useful lifetime prediction via deep domain adaptation. *Reliability Engineering and System Safety*, 224, 108526.
- Ganin, Y., Ustinova, E., Ajakan, H. et al. (2016) Domain-adversarial training of neural networks. *JMLR*, 17(1), pp. 2096-2030.
- Guo, L., Li, N., Jia, F., Lei, Y. and Lin, J. (2017) A recurrent neural network based health indicator for remaining useful life prediction of bearings. *Neurocomputing*, 240, pp. 98-109.
- He, Z. and Shao, H. (2022) Intelligent fault diagnosis of rolling bearings based on transferable convolutional neural network. *IEEE Transactions on Instrumentation and Measurement*, 71, 3505412.
- Jia, F., Lei, Y., Lin, J., Zhou, X. and Lu, N. (2016) Deep neural networks: A promising tool for fault characteristic mining and intelligent diagnosis of rotating machinery with massive data. *Mechanical Systems and Signal Processing*, 72-73, pp. 303-315.
- Koizumi, Y., Saito, S., Uematsu, H., Harada, N. and Imoto, K. (2019) ToyADMOS: A dataset of miniature-machine operating sounds for anomalous sound detection. *Proceedings of WASPAA 2019*, pp. 313-317.
- Lei, Y., Li, N., Guo, L., Li, N., Yan, T. and Lin, J. (2018) Machinery health prognostics: A systematic review from data acquisition to RUL prediction. *Mechanical Systems and Signal Processing*, 104, pp. 799-834.
- Li, X., Ding, Q. and Sun, J.Q. (2018) Remaining useful life estimation in prognostics using deep convolution neural networks. *Reliability Engineering and System Safety*, 172, pp. 1-11.
- Li, W., Huang, R., Li, J., Liao, Y., Chen, Z., He, G., Yan, R. and Gryllias, K. (2022) A perspective survey on deep transfer learning for fault diagnosis in industrial scenarios: Theories, applications and challenges. *Mechanical Systems and Signal Processing*, 167, 108487.
- Loparo, K.A. (2012) Case Western Reserve University bearing data center. <http://csegroups.case.edu/bearingdatacenter>.
- Nectoux, P., Gouriveau, R., Medjaher, K. et al. (2012) PRONOSTIA: An experimental platform for bearings accelerated degradation tests. *Proceedings of IEEE PHM 2012*.
- Purohit, H., Tanabe, R., Ichige, K., Endo, T., Nikaido, Y., Suefusa, K. and Kawaguchi, Y. (2019) MIMII dataset: Sound dataset for malfunctioning industrial machine investigation and inspection. *Proceedings of DCASE 2019*, pp. 209-213.
- Ran, Y., Zhou, X., Lin, P., Wen, Y. and Deng, R. (2019) A survey of predictive maintenance: Systems, purposes and approaches. *arXiv preprint arXiv:1912.07383*.
- Saxena, A., Goebel, K., Simon, D. and Eklund, N. (2008) Damage propagation modeling for aircraft engine run-to-failure simulation. *Proceedings of PHM 2008*.
- Smith, W.A. and Randall, R.B. (2015) Rolling element bearing diagnostics using the Case Western Reserve University data: A benchmark study. *Mechanical Systems and Signal Processing*, 64-65, pp. 100-131.
- Susto, G.A., Schirru, A., Pampuri, S., McLoone, S. and Beghi, A. (2015) Machine learning for predictive maintenance: A multiple classifier approach. *IEEE Transactions on Industrial Informatics*, 11(3), pp. 812-820.
- Wang, B., Lei, Y., Li, N. and Yan, T. (2019) Deep separable convolutional network for remaining useful life prediction of machinery. *Mechanical Systems and Signal Processing*, 134, 106330.
- Zhang, W., Peng, G., Li, C., Chen, Y. and Zhang, Z. (2017) A new deep learning model for fault diagnosis with good anti-noise and domain adaptation ability on raw vibration signals. *Sensors*, 17(2), 425.
- Zheng, S., Ristovski, K., Farahat, A. and Gupta, C. (2017) Long short-term memory network for remaining useful life estimation. *Proceedings of IEEE ICPHM 2017*, pp. 88-95.
- Zhu, J., Chen, N. and Peng, W. (2019) Estimation of bearing remaining useful life based on multiscale convolutional neural network. *IEEE Transactions on Industrial Electronics*, 66(4), pp. 3208-3216.

Declarations

Funding

This research was supported by the Swiss National Science Foundation under grant 200021-214248 (IndustrialAI) and the French National Research Agency (ANR) under grant ANR-22-CE10-0034 (PredMaintAI). Computational resources provided by Swiss National Supercomputing Centre (CSCS) and GENCI-IDRIS. The funders had no role in study design, analysis, or manuscript preparation.

Conflict of Interest

The authors declare no conflicts of interest.

Data Availability Statement

All model implementations, training scripts, transfer learning experiments, and benchmark results are deposited in the Zenodo repository under DOI: 10.5281/zenodo.19614734. NASA C-MAPSS, CWRU, Paderborn, and MIMII are publicly available. Released under MIT License (code) and CC BY 4.0 (results).

Ethical Approval

This study involved computational experiments on publicly available industrial sensor datasets. No human subjects, personal data, or proprietary industrial data were involved. No institutional ethical approval was required.

Appendix A

Predictive Maintenance Deployment Guide

This appendix provides a decision tree for deploying AI predictive maintenance based on equipment criticality, sensor availability, and data maturity.

Step 1: Equipment Criticality Assessment

- 1a. High criticality (failure cost > 100x maintenance cost) --> Full AI deployment; ROI breakeven 8-14 months
- 1b. Medium criticality (failure cost 10-100x maintenance) --> Selective AI; start with anomaly detection
- 1c. Low criticality (failure cost < 10x maintenance) --> Continue preventive schedule; AI not cost-justified

Step 2: Sensor and Data Selection

- 2a. Vibration sensors available --> 1D-CNN fault diagnosis (99.2% accuracy controlled; add DANN for variable conditions)
- 2b. Multi-sensor time series (temp, pressure, vibration) --> TCN or PINN-TCN for RUL prediction (RMSE = 10.2)
- 2c. Audio only (non-invasive) --> Spectrogram autoencoder anomaly detection (AUC = 0.924; minimal installation)
- 2d. Work order text records --> LLM extraction for failure pattern mining (84.6% precision)

Step 3: Domain Adaptation Strategy

- 3a. Fixed operating conditions --> Direct model deployment (no adaptation needed)
- 3b. Variable conditions, labelled target data available --> Fine-tuning (best recovery: 95.4%)
- 3c. Variable conditions, no target labels --> DANN adversarial adaptation (recovery: 54-68%)
- 3d. Always validate on target-condition data before production deployment