



Artificial Intelligence Techniques for Predictive Maintenance in Smart Manufacturing Systems

Mahfuz Chowdhuryan

*Department of Electrical and Computer Engineering, Aurora Metropolitan Institute of Technology,
Philippines*

mahfuz.chowdhuryan@amit-ph.edu

Peer Review Information

Submission: 01 Jan 2018

Revision: 25 Feb 2018

Acceptance: 06 March 2018

Keywords

*Artificial Intelligence,
Predictive Maintenance,
Smart Manufacturing,
Machine Learning,
Condition Monitoring*

Abstract

Predictive maintenance (PdM) is a central capability of smart manufacturing because it converts continuously collected machine-condition data into early warnings and maintenance decisions before functional failure occurs. This methodology paper develops an artificial-intelligence-based PdM framework that combines sensor acquisition, data cleaning, health-feature construction, supervised learning and maintenance decision support. The framework evaluates logistic regression, support vector machine (SVM), random forest and multilayer perceptron (MLP) models for prediction of a near-failure maintenance horizon. To demonstrate the methodology without claiming unavailable industrial plant measurements, a controlled synthetic smart-manufacturing benchmark of 6,000 multivariate observations was generated from vibration, temperature, motor current, hydraulic pressure, acoustic level, operating load and runtime variables. A stratified 75:25 train-test design and imbalance-aware learning were used. In the illustrative evaluation, random forest achieved the highest accuracy (95.2%), F1-score (0.863) and ROC-AUC (0.929), while SVM showed slightly higher failure recall (0.850) than random forest (0.828). The results support the use of ensemble and nonlinear AI methods when interactions among machine-condition variables are important. The proposed procedure is reproducible, technology-neutral and suitable for later validation with real factory data.

Introduction

Modern manufacturing systems increasingly combine physical equipment with sensors, industrial networks and computational analytics. This shift makes maintenance a data-driven activity rather than a fixed calendar operation. Condition-based maintenance (CBM) established the basic logic of collecting condition information, processing it and using the result for maintenance decisions [1]. Prognostics extended that logic toward forecasting future health and remaining useful life (RUL), with rotating-machinery research emphasizing the operational value of anticipating degradation rather than

responding only after faults become obvious [2]. Statistical RUL research similarly showed that remaining life must be inferred from uncertain condition observations and degradation histories [3].

The smart-manufacturing context makes this problem particularly suitable for artificial intelligence (AI). Predictive manufacturing research argued that large streams of machine, process and quality data should be transformed into actionable knowledge about otherwise invisible degradation [4]. Cyber-physical system (CPS) architectures subsequently provided a

structured pathway for connecting physical assets, data acquisition, cyber analytics and decision layers [5]. In parallel, Susto et al. demonstrated that machine-learning classifiers can support predictive-maintenance decisions through different prediction horizons and cost-sensitive maintenance logic [6]. These developments moved PdM from isolated signal analysis toward integrated manufacturing intelligence.

However, practical PdM remains difficult because factory data are noisy, multirate, imbalanced and strongly affected by operating regimes. Industrial big-data work highlighted the need to transform high-volume machine data into information that exposes degradation patterns before breakdown [7], while manufacturing-focused machine-learning research emphasized both the flexibility of learning algorithms and the challenges of data quality, interpretability and deployment [8]. Deep neural methods later showed that useful fault representations can be learned directly from large machinery datasets [9], and systematic prognostics research organized the process into data acquisition, health-indicator construction, health-stage division and RUL prediction [10]. By 2018, deep learning and industrial AI were being explicitly positioned as key analytics technologies for smart manufacturing [11], [12].

Accordingly, the objective of this paper is to present a compact, step-by-step methodology for comparing representative AI techniques for predictive maintenance in a smart-manufacturing setting. The paper focuses on a classification-based maintenance horizon: given the current condition vector of a machine, the model estimates whether maintenance will be required within a predefined near-future interval. This formulation supports operational decisions such as inspection, planned shutdown or continued operation, and can later be extended to continuous RUL regression.

Literature Review

Early CBM literature provides the conceptual foundation for AI-enabled PdM. Jardine et al. [1] described CBM as a sequence of data acquisition, data processing and maintenance decision-making and reviewed diagnostic, prognostic and sensor-fusion methods. Heng et al. [2] focused on rotating-machinery prognostics and identified challenges such as changing operating

conditions, uncertainty and the difficulty of validating life predictions. Si et al. [3] reviewed statistical data-driven RUL estimation and showed that no single degradation model is universally best, motivating comparative and adaptive approaches.

The manufacturing literature then linked prognostics with connected production. Lee et al. [4] described predictive manufacturing as the conversion of large-scale shop-floor data into visibility of degradation, inefficiency and future risk. Their later 5C CPS architecture connected smart sensors and machines to cyber-level analytics and configuration/decision functions [5]. Related industrial big-data work emphasized that analytics can reveal hidden degradation patterns and support maintenance-service innovation [7]. These studies justify treating predictive maintenance as an end-to-end information pipeline rather than an isolated classifier.

Machine learning broadened the algorithmic options available to PdM. Susto et al. [6] proposed a multiple-classifier approach that uses learned health factors and prediction horizons to reduce unexpected breakdowns and unused component life. Wuest et al. [8] reviewed machine-learning applications in manufacturing and noted that algorithm choice must be considered together with data preparation, domain understanding and implementation constraints. Jia et al. [9] demonstrated that deep neural networks can jointly perform fault-feature mining and intelligent diagnosis from large rotating-machinery datasets, reducing dependence on manually engineered indicators.

The 2018 literature consolidated these directions. Lei et al. [10] formalized machinery prognostics as a sequence from data acquisition to health indicators, health-stage division and RUL prediction. Wang et al. [11] reviewed deep-learning methods for smart manufacturing, stressing their ability to process high-volume and high-variety manufacturing data. Lee et al. [12] framed industrial AI as a systematic capability for Industry 4.0 systems, connecting data, analytics and resilient decision-making. The main gap addressed by the present methodology is therefore not the absence of candidate algorithms, but the need for a transparent comparative procedure that joins sensing, preprocessing, imbalance handling, model evaluation and maintenance decision rules in one reproducible workflow.

Table 1. Selected pre-2019 literature underpinning the proposed predictive-maintenance methodology.

Study	Year	Main contribution	Relevance to this methodology
-------	------	-------------------	-------------------------------

Jardine et al. [1]	2006	CBM pipeline: acquisition, processing and decision-making	Defines the end-to-end maintenance information flow
Si et al. [3]	2011	Review of statistical data-driven RUL methods	Motivates model comparison under uncertainty
Lee et al. [4]	2013	Predictive manufacturing in a big-data environment	Links shop-floor data to proactive maintenance
Susto et al. [6]	2015	Multiple-classifier machine learning for PdM	Supports classification-based maintenance horizons
Wuest et al. [8]	2016	Manufacturing ML opportunities and challenges	Highlights data quality and deployment constraints
Jia et al. [9]	2016	Deep feature learning for machinery diagnosis	Supports nonlinear representation learning
Lei et al. [10]	2018	Systematic prognostics workflow from data to RUL	Structures health indicators and prognostic validation

Methodology

1. Research Design and Problem Formulation

A quantitative comparative design is used. Each observation represents the operating condition of a manufacturing asset at a decision instant. The target variable y is binary: $y = 1$ indicates that the asset lies within a near-failure maintenance horizon and $y = 0$ indicates continued normal operation. The methodology compares four supervised AI techniques: logistic regression as an interpretable linear baseline, SVM for margin-based nonlinear discrimination, random forest for ensemble learning and interaction discovery, and MLP as a compact neural-network baseline. This selection reflects the range of model families discussed in manufacturing ML and predictive-maintenance literature [6], [8], [11].

Step-by-Step Procedure

Step 1 - Sensor acquisition. Collect timestamped measurements from critical assets. The proposed feature set includes vibration, temperature, motor current, hydraulic pressure, acoustic level, operating load and cumulative runtime. These signals represent complementary mechanical, thermal, electrical and process conditions and align with the multisensor logic of CBM and prognostics [1], [10].

Step 2 - Data synchronization and cleaning. Align channels to a common sampling interval; remove duplicate timestamps; flag impossible values; interpolate only short gaps; and retain missingness indicators when gaps may themselves reflect communication faults. Extreme values are checked against engineering limits instead of being automatically deleted.

Step 3 - Health-feature construction. For real time-series deployment, compute rolling mean, standard deviation, root-mean-square vibration,

kurtosis, slope, temperature rise, pressure drop and load-normalized features over a fixed window. The benchmark uses the seven contemporaneous condition variables directly so that algorithmic differences are not hidden by a complex feature pipeline.

Step 4 - Label definition. Define a maintenance horizon H (for example, failure or required intervention within the next production shift). Historical maintenance logs should be joined to condition records to generate labels. When multiple failure modes exist, the same pipeline may be applied as multiclass classification or as one-versus-rest models.

Step 5 - Training and imbalance control. Split data by asset or chronological blocks whenever plant history permits, so that future information does not leak into training. For the demonstration, a stratified 75:25 split is used. Class weighting is applied to logistic regression, SVM and random forest; the MLP training subset is randomly oversampled to balance near-failure and normal observations.

Step 6 - Model fitting and validation. Standardize features for logistic regression, SVM and MLP. Tune model complexity only on the training portion. Evaluate the untouched test set using accuracy, precision, recall, F1-score and ROC-AUC. Because missed failures are often more costly than false alarms, recall and F1 should be interpreted alongside accuracy [6].

Step 7 - Maintenance decision. Convert the estimated failure probability into an operational action. A high-risk score triggers inspection or scheduled maintenance; intermediate risk can trigger a shorter re-sampling interval; low risk allows continued operation. The decision threshold should ultimately be calibrated to downtime, spare-part and false-alarm costs.



Figure 1. Proposed AI-enabled predictive-maintenance workflow for smart manufacturing.

2. Demonstration Dataset and Reproducibility

To demonstrate the method without presenting synthetic observations as real industrial evidence, a controlled benchmark containing 6,000 observations was generated with a fixed random seed (42). A latent degradation variable was used to create physically plausible upward trends in vibration, temperature, acoustic level and runtime, together with pressure decline and load/current effects. The near-failure label was defined from a nonlinear risk function and small label noise was added to represent imperfect maintenance records. The resulting positive-class prevalence was 18.23%. This design intentionally creates nonlinear interactions that a PdM model should identify while preserving overlap between healthy and degraded states.

The training set contains 4,500 observations and the test set 1,500 observations. Logistic regression uses L2 regularization; SVM uses a radial-basis-function kernel; random forest uses 300 trees with bounded depth; and the MLP uses two hidden layers (32 and 16 neurons) with early stopping. No test observation is used in preprocessing parameter estimation or model fitting. The same workflow can be reproduced on plant data by replacing the simulation generator with a historian/SCADA/IIoT data extraction stage and retaining the remaining steps.

3. Evaluation Metrics

Accuracy measures the overall fraction of correct predictions, precision measures the proportion of predicted maintenance alerts that are true near-failure cases, recall measures the proportion of near-failure cases detected, and F1-score is the harmonic mean of precision and recall. ROC-AUC evaluates ranking quality across decision thresholds. In PdM, a high accuracy value alone may be misleading when failures are rare; therefore, the primary comparison emphasizes F1, ROC-AUC and the business meaning of recall.

4. Operational Deployment and Model Governance

A predictive model becomes useful only when its output is connected to a controlled maintenance process. The probability threshold should therefore be calibrated from the relative cost of a missed failure, an unnecessary intervention and production interruption. Instead of treating 0.50 as a universal threshold, the plant can evaluate candidate thresholds on validation data and select the operating point that minimizes expected maintenance cost while satisfying a minimum failure-detection requirement. For safety-critical equipment, the recall constraint may dominate; for redundant equipment with expensive shutdown procedures, precision may receive greater weight. This cost-aware interpretation follows the maintenance-decision logic emphasized in classifier-based PdM research [6].

Deployment should also include model-health controls. Input distributions for vibration, temperature, load and other variables should be compared with the training baseline so that sensor drift, changes in product mix or new operating regimes can be detected. A persistent shift should trigger engineering review and, where sufficient new labels exist, retraining. Each alert should store the model version, timestamp, risk score, principal condition variables and resulting maintenance action. Such traceability supports root-cause analysis and allows engineers to determine whether an alert was useful, late or unnecessary. For tree-based models, feature importance and local contribution measures can be provided; for SVM and MLP models, sensitivity or perturbation analyses can be used to communicate which signals drove the decision.

Finally, the analytics service should be integrated with the manufacturing information architecture rather than operated as an isolated notebook. Sensor acquisition can occur at the machine or

edge layer, while model scoring may run at the edge or in a plant server depending on latency and cybersecurity requirements. Predictions and maintenance outcomes should return to the cyber-physical feedback loop described in smart-manufacturing architectures [5], [7], [12]. This closed-loop design ensures that the PdM system is monitored as an operational asset and that model performance can improve as additional maintenance histories become available.

Results

Table 1 summarizes the simulation-based test performance. Random forest produced the

strongest overall balance, with 95.2% accuracy, 0.901 precision, 0.863 F1-score and 0.929 ROC-AUC. SVM achieved 94.3% accuracy and the highest recall among the nonlinear models (0.850), indicating a slightly more aggressive failure-detection profile. The MLP achieved 93.7% accuracy and 0.832 F1, while logistic regression had markedly lower precision and F1 despite high recall. The difference suggests that the benchmark contains nonlinear interactions among condition variables that cannot be represented adequately by a simple linear boundary.

Table 2. Test-set performance on the controlled smart-manufacturing predictive-maintenance benchmark (n = 1,500 test observations).

AI technique	Accuracy	Precision	Recall	F1-score	ROC-AUC
Logistic regression	0.863	0.586	0.858	0.696	0.910
Support vector machine	0.943	0.841	0.850	0.846	0.921
Random forest	0.952	0.901	0.828	0.863	0.929
Multilayer perceptron	0.937	0.815	0.850	0.832	0.910

The result is consistent with the literature in two ways. First, Susto et al. [6] showed the usefulness of multiple machine-learning classifiers for PdM, where the preferred operating point depends on the cost of unexpected failures versus unused component life. Second, manufacturing AI reviews emphasize that nonlinear models are attractive when heterogeneous sensor streams and complex process interactions are present [8], [11], [12]. The random-forest result should not be interpreted as evidence that random forest will always dominate; instead, it demonstrates why model comparison under a common data split and common metrics is necessary.

From a deployment perspective, threshold selection would follow the operational objective. If avoiding missed failures is the dominant concern, the SVM or a lower random-forest probability threshold may be preferred because recall can be increased at the cost of more false alarms. If maintenance capacity is limited and unnecessary interventions are expensive, the higher precision of random forest is advantageous. In a real factory, the next validation stage should use temporally separated production periods and, ideally, an asset-level holdout to test generalization to different machines and operating regimes.

Conclusion and Discussion

This paper presented a step-by-step methodology for applying and comparing AI techniques for predictive maintenance in smart manufacturing systems. The framework integrates multisensor acquisition, preprocessing, health-feature construction,

imbalance-aware supervised learning, performance evaluation and maintenance decision logic. It is grounded in the evolution from CBM and prognostics [1]-[3] to predictive manufacturing and CPS-based analytics [4], [5], [7], and then to machine learning, deep learning and industrial AI for manufacturing [6], [8]-[12]. The controlled benchmark indicates that nonlinear AI techniques can provide strong predictive performance when degradation depends on interactions among vibration, thermal, electrical and process variables. Random forest offered the best overall F1 and ROC-AUC, whereas SVM provided slightly higher failure recall. These findings illustrate an important maintenance principle: the “best” model is not determined by accuracy alone, but by the tradeoff between missed failures, false alarms, interpretability, computational cost and the available maintenance window. The principal limitation is that the numerical results are simulation-based and therefore demonstrate the procedure rather than prove plant-level effectiveness. Real deployments must address sensor drift, rare failure labels, maintenance-policy changes, covariate shift, new product variants and differences among nominally similar machines. Future work should validate the framework on historical manufacturing data, adopt time-aware or asset-aware cross-validation, compare continuous RUL models, and add explainability so that maintenance engineers can understand the variables driving each alert. Periodic model monitoring and retraining should also be integrated into the CPS feedback loop to

prevent performance degradation as the production environment changes.

References

K. S. Jardine, D. Lin, and D. Banjevic, "A review on machinery diagnostics and prognostics implementing condition-based maintenance," *Mechanical Systems and Signal Processing*, vol. 20, no. 7, pp. 1483–1510, 2006. DOI: 10.1016/j.ymssp.2005.09.012.

Heng, S. Zhang, A. C. C. Tan, and J. Mathew, "Rotating machinery prognostics: State of the art, challenges and opportunities," *Mechanical Systems and Signal Processing*, vol. 23, no. 3, pp. 724–739, 2009. DOI: 10.1016/j.ymssp.2008.06.009.

X.-S. Si, W. Wang, C.-H. Hu, and D.-H. Zhou, "Remaining useful life estimation – A review on the statistical data driven approaches," *European Journal of Operational Research*, vol. 213, no. 1, pp. 1–14, 2011. DOI: 10.1016/j.ejor.2010.11.018.

J. Lee, E. Lapira, B. Bagheri, and H.-A. Kao, "Recent advances and trends in predictive manufacturing systems in big data environment," *Manufacturing Letters*, vol. 1, no. 1, pp. 38–41, 2013. DOI: 10.1016/j.mfglet.2013.09.005.

J. Lee, B. Bagheri, and H.-A. Kao, "A Cyber-Physical Systems architecture for Industry 4.0-based manufacturing systems," *Manufacturing Letters*, vol. 3, pp. 18–23, 2015. DOI: 10.1016/j.mfglet.2014.12.001.

G. A. Susto, A. Schirru, S. Pampuri, S. F. McLoone, and A. Beghi, "Machine Learning for Predictive Maintenance: A Multiple Classifier Approach," *IEEE Transactions on Industrial Informatics*, vol.

11, no. 3, pp. 812–820, 2015. DOI: 10.1109/TII.2014.2349359.

J. Lee, H. Davari Ardakani, S. Yang, and B. Bagheri, "Industrial Big Data Analytics and Cyber-physical Systems for Future Maintenance & Service Innovation," *Procedia CIRP*, vol. 38, pp. 3–7, 2015. DOI: 10.1016/j.procir.2015.08.026.

T. Wuest, D. Weimer, C. Irgens, and K.-D. Thoben, "Machine learning in manufacturing: advantages, challenges, and applications," *Production & Manufacturing Research*, vol. 4, no. 1, pp. 23–45, 2016. DOI: 10.1080/21693277.2016.1192517.

F. Jia, Y. Lei, J. Lin, X. Zhou, and N. Lu, "Deep neural networks: A promising tool for fault characteristic mining and intelligent diagnosis of rotating machinery with massive data," *Mechanical Systems and Signal Processing*, vols. 72–73, pp. 303–315, 2016. DOI: 10.1016/j.ymssp.2015.10.025.

Y. Lei, N. Li, L. Guo, N. Li, T. Yan, and J. Lin, "Machinery health prognostics: A systematic review from data acquisition to RUL prediction," *Mechanical Systems and Signal Processing*, vol. 104, pp. 799–834, 2018. DOI: 10.1016/j.ymssp.2017.11.016.

J. Wang, Y. Ma, L. Zhang, R. X. Gao, and D. Wu, "Deep learning for smart manufacturing: Methods and applications," *Journal of Manufacturing Systems*, vol. 48, pp. 144–156, 2018. DOI: 10.1016/j.jmsy.2018.01.003.

J. Lee, H. Davari, J. Singh, and V. Pandhare, "Industrial Artificial Intelligence for industry 4.0-based manufacturing systems," *Manufacturing Letters*, vol. 18, pp. 20–23, 2018. DOI: 10.1016/j.mfglet.2018.09.002.