Contemporary Education Frontiers, 2026, Volume 4, Issue 7
https://bbwpublisher.com/journal/CEF
Online ISSN: 3029-1860
Print ISSN: 3029-1879

Artificial Intelligence for Fault Early Warning in Equipment Manufacturing

Zhiqiang Li*

Nanjing University, Nanjing 210008, Jiangsu, China

Author to whom correspondence should be addressed.

Copyright: © 2026 Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY 4.0), permitting distribution and reproduction in any medium, provided the original work is cited.

Abstract: Artificial intelligence can identify equipment deterioration, but accurate fault classification does not by itself establish effective early warning. This narrative review examines how condition monitoring can support maintenance decisions in machinery-manufacturing plants. It connects diagnostic and prognostic research with data quality, operating context, uncertainty, and alert evaluation. A decision-oriented framework is developed through synthesis, linking observations to a prediction target, an event-level warning, a maintenance response, and subsequent verification. Two illustrative applications, machining-centre spindle monitoring and hydraulic press monitoring, explain possible implementation choices without implying completed field trials. The review argues that practical evaluation should consider useful warning time, missed events, false-alert burden, and maintenance consequences alongside model performance. Controlled faults and simulated degradation provide valuable methodological evidence, but their relevance to factory operation requires separate assessment. No new experimental results are reported.

Keywords: Artificial intelligence; Equipment manufacturing; Fault early warning; Predictive maintenance; Condition monitoring

Online publication: July 26, 2026

1. Introduction

Production equipment in machinery-manufacturing plants supports connected machining, assembly, and testing operations. A deteriorating component can interrupt several processes, yet replacing it at the first unusual observation may also waste usable life. Maintenance teams must decide when to inspect, intervene, or continue operating.

Condition-based maintenance connects data acquisition and analysis to maintenance decisions [1]. Machine learning expands the available methods, although predictive-maintenance research addresses different targets and operating conditions [2, 3]. A multi-sector mapping also identifies variation in inputs, predicted variables, and applications [4]. Consequently, a model described as an early-warning system requires more evidence than accurate recognition of an existing fault.

This review examines that evidence and proposes a structure for implementation. Its scope is production assets within equipment-manufacturing factories. Evidence from other asset classes is used only with stated limitations. The contribution is a synthesis linking observations, prediction targets, alerts, and maintenance actions, rather than a new algorithm or experimentally validated system.

2. Review approach

Sources were selected purposively for relevance to machinery diagnostics, prognosis, manufacturing maintenance, representative datasets, and supporting methods. Bibliographic details were checked against publisher records, official proceedings, or institutional repositories. This narrative selection is not exhaustive; no systematic search count or pooled effect is claimed.

Diagnostic and health-monitoring reviews provide the methodological background [5, 6]. Work on prognostics and a recent diagnostic roadmap help frame transfer and deployment limitations [7, 8]. Original papers clarify specific methods and datasets. The synthesis distinguishes condition recognition, future-event prediction, and operational maintenance evidence. The scenarios below are proposed applications, not projects performed for this review.

3. Defining an actionable warning

Detection asks whether equipment currently departs from an accepted condition; diagnosis concerns the fault involved. Prognosis estimates future condition or failure. An actionable early warning additionally requires an event definition and enough time for a specified response.

The target should identify the component, prediction window, and intended maintenance action. Susto et al. connect classifiers with different prediction horizons to maintenance decisions [9]. Remaining useful life is uncertain rather than an exact countdown [10]; hidden-state models provide another formulation of degradation for prognosis [11]. These formulations still require a warning interval appropriate to the equipment.

A technically correct alert may arrive too late to arrange inspection. A much earlier alert may be too uncertain to justify replacement. Access, spare parts, production commitments, and failure consequences determine the acceptable interval. **Table 1** distinguishes the evidence needed for progressively stronger claims.

Table 1. Evidence requirements for early warning claims

Claim	Evaluation focus	Interpretation limit
Condition recognition	Errors across machines and conditions	Does not establish warning time
Future-event prediction	Event detection and warning-time distribution	Needs a defined event and horizon
Operational warning	Alert burden and response outcomes	Does not establish financial benefit
Maintenance benefit	Comparison including interventions and costs	Requires attention to confounding

4. Data and evidence quality

Vibration, current, temperature, pressure, and maintenance records can inform equipment assessment. Their meaning depends on operating conditions. Vibration changes during a new machining task may reflect load

rather than damage. Speed, task, operating mode, sensor configuration, and recent maintenance should therefore remain traceable. Missing observations should not silently imply healthy operation.

Benchmarks support reproducibility but have limited scope. The Paderborn bearing dataset includes artificial damage and damage from accelerated lifetime tests [12]. This is valuable diversity, not evidence of continuous factory deployment. Engine degradation data described by Saxena et al. were generated through simulation [13]; they support prognostic-method development rather than manufacturing-field validation.

Evaluation should separate machines, acquisition runs, or periods according to the intended use. Randomly distributing overlapping windows can expose a model to near-duplicate training and test observations. Many windows do not create many independent failures. Reporting should identify the machines and events represented. Predictors must also exclude information available only after inspection or repair, which could disclose the outcome indirectly.

5. Selecting AI methods

Selection should follow the target and available evidence. Engineered features with statistical models, support-vector machines, or tree ensembles provide useful baselines. Deep models can learn representations from raw or transformed signals [5,6], but their flexibility does not establish universal superiority.

Anomaly detection can be appropriate when normal operation is well represented and fault examples are scarce. Its score indicates departure from a reference, not necessarily damage or failure probability. Sensor changes and unfamiliar production tasks can also produce deviations. A system should allow uncertain observations to prompt investigation without forcing a known fault label.

For prognosis, future loading and the available degradation history affect inference [7]. Domain adaptation may help address differences between training and deployment. Domain-adversarial learning, for example, seeks task-relevant features that reduce domain discrimination [14]. This general method is not itself an industrial warning trial. Target-domain evaluation remains necessary because alignment can also remove differences relevant to equipment health.

6. From scores to maintenance actions

A score should be interpreted as a failure probability only after that interpretation has been evaluated. Neural-network calibration research shows that confidence and observed correctness can differ [15]. Success on general classification tasks does not establish calibrated machinery risks, especially after operating conditions change.

Alert logic may combine persistence, context, and escalation. Repeated deviation during comparable cycles could warrant inspection, while a transient during a mode change could warrant further observation. This is a proposed design principle, not a universal threshold rule. Persistence periods must reflect equipment behaviour and response requirements.

Explanations should support inspection of the evidence. LIME uses a local surrogate to explain predictions [16]; SHAP assigns feature contributions within an additive framework [17]. Attribution can reveal reliance on an irrelevant variable, but it does not prove mechanical causation. An alert record should retain the signal, context, model version, explanation, responsible reviewer, and subsequent disposition.

7. Illustrative applications

For a machining-centre spindle, vibration, current, and temperature could be associated with speed, load, task, and maintenance history. Persistent departure from comparable operating periods could trigger remeasurement and inspection of lubrication, tooling, mounting, or the spindle assembly. Similar signals can have different causes so that component-specific diagnosis would require additional evidence. Evaluation would compare useful warning time, alert burden, and inspection findings with the existing maintenance procedure.

For an assembly hydraulic press, pressure, flow, oil temperature, and cycle timing could be compared within the same task and operating phase. Persistent changes in pressure build-up or holding behaviour could support an inspection request. Leakage, valve behaviour, sensor drift, and task changes would remain competing explanations. Product-quality deviation and equipment deterioration must also remain separate outcomes. Records connecting alerts with inspection findings and later operation would help refine the target.

Both scenarios are illustrative proposals. Neither establishes achieved accuracy, avoided downtime, or financial returns. Initial use would support review; changes to machine settings would require a separately authorized control procedure.

8. Validation and deployment

Offline evaluation should compare an existing engineering rule, a simple statistical baseline, and the learning method using the same information available at prediction time. Model selection and threshold tuning belong in training or validation data, not the final test set. Results should be separated by machine and condition where sample size permits.

Event-level reporting is essential: hundreds of positive windows before one fault remain one event. Useful measures include event detection, warning-time distribution, missed events, and false alerts per unit of operating time. Precision-recall measures can supplement these when faults are rare. Threshold selection should expose the trade-off between review workload and missed events.

Intervention complicates assessment. Replacing a component after an alert leaves its eventual failure time unobserved. The absence of a later breakdown establishes neither that the alert was false nor that a failure was prevented. Inspection findings and intervention histories must be retained. Economic claims need a credible comparison rather than counting every avoided stoppage as a model benefit.

A prospective shadow phase can examine data continuity and alert behaviour before recommendations influence maintenance. Later evaluation should record response delays, inspection effort, and production interruptions. Edge or centralized processing should be chosen against measured latency and reliability needs. Recoverable model and configuration versions are necessary when performance changes.

9. Discussion and conclusions

The synthesis places the maintenance decision at the centre of fault early warning. An implementation should define the event, horizon, operating context, and intended response before selecting its model. Benchmark accuracy supports only part of this process; operational consequences require separate evaluation.

Limitations include purposive source selection, heterogeneous evidence, and the absence of new experiments. No pooled performance estimate is justified. The scenarios require local validation, and findings

for bearings or simulated engines cannot establish performance for a particular factory. Existing application maps should guide investigation rather than replace site evidence [4, 8].

Future work should prioritize prospective event-level studies with transparent operating conditions and intervention records. Useful progress would connect warnings to what engineers subsequently observe and do. Where failure events remain scarce, the uncertainty should be reported rather than hidden behind precise aggregate scores.

Disclosure statement

The author declares no conflict of interest.

References

[1] Jardine AKS, Lin D, Banjevic D, 2006, A Review on Machinery Diagnostics and Prognostics Implementing Condition-Based Maintenance. Mechanical Systems and Signal Processing, 20(7): 1483–1510.

[2] Carvalho TP, Soares FAAMN, Vita R, et al., 2019, A Systematic Literature Review of Machine Learning Methods Applied to Predictive Maintenance. Computers & Industrial Engineering, 137: 106024.

[3] Zonta T, da Costa CA, da Rosa Righi R, et al., 2020, Predictive Maintenance in Industry 4.0: A Systematic Literature Review. Computers & Industrial Engineering, 150: 106889.

[4] Mallioris P, Aivazidou E, Bechtsis D, 2024, Predictive Maintenance in Industry 4.0: A Systematic Multi-Sector Mapping. CIRP Journal of Manufacturing Science and Technology, 50: 80–103.

[5] Lei Y, Yang B, Jiang X, et al., 2020, Applications of Machine Learning to Machine Fault Diagnosis: A Review and Roadmap. Mechanical Systems and Signal Processing, 138: 106587.

[6] Zhao R, Yan R, Chen Z, et al., 2019, Deep Learning and Its Applications to Machine Health Monitoring. Mechanical Systems and Signal Processing, 115: 213–237.

[7] Fink O, Wang Q, Svensen M, et al., 2020, Potential, Challenges and Future Directions for Deep Learning in Prognostics and Health Management Applications. Engineering Applications of Artificial Intelligence, 92: 103678.

[8] Vashishtha G, Chauhan S, Sehri M, et al., 2025, A Roadmap to Fault Diagnosis of Industrial Machines via Machine Learning: A Brief Review. Measurement, 242: 116216.

[9] Susto GA, Schirru A, Pampuri S, et al., 2015, Machine Learning for Predictive Maintenance: A Multiple Classifier Approach. IEEE Transactions on Industrial Informatics, 11(3): 812–820.

[10] Si XS, Wang W, Hu CH, et al., 2011, Remaining Useful Life Estimation: A Review on the Statistical Data-Driven Approaches. European Journal of Operational Research, 213(1): 1–14.

[11] Tobon-Mejia DA, Medjaher K, Zerhouni N, et al., 2012, A Data-Driven Failure Prognostics Method Based on Mixture of Gaussians Hidden Markov Models. IEEE Transactions on Reliability, 61(2): 491–503.

[12] Lessmeier C, Kimotho JK, Zimmer D, et al., 2016, Condition Monitoring of Bearing Damage in Electromechanical Drive Systems by Using Motor Current Signals of Electric Motors: A Benchmark Data Set for Data-Driven Classification. PHM Society European Conference, 3(1): 1577.

[13] Saxena A, Goebel K, Simon D, et al., 2008, Damage Propagation Modeling for Aircraft Engine Prognostics. NASA Technical Reports Server, Record 20090029214.

[14] Ganin Y, Ustinova E, Ajakan H, et al., 2016, Domain-Adversarial Training of Neural Networks. Journal of Machine Learning Research, 17(59): 1–35.

[15] Guo C, Pleiss G, Sun Y, et al., 2017, On Calibration of Modern Neural Networks. Proceedings of the 34th International Conference on Machine Learning, PMLR 70: 1321–1330.

[16] Ribeiro MT, Singh S, Guestrin C, 2016, "Why Should I Trust You?": Explaining the Predictions of Any Classifier. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.

[17] Lundberg SM, Lee SI, 2017, A Unified Approach to Interpreting Model Predictions. Advances in Neural Information Processing Systems, 30.

Publisher's note

Bio-Byword Scientific Publishing remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.