

Predictive Maintenance of Electrical Equipment Using Sensor Data and Machine Learning

Jieyi Chen

*Department of Electrical and Electronic Engineering, University of Southampton, Southampton, UK
chenjieyi0616@gmail.com*

Abstract. Motors, transformers, and generators run for long periods under electrical, thermal, and mechanical stress, and when they fail without warning the result is often costly downtime and sometimes a safety risk. Predictive maintenance (PdM) tries to avoid this by estimating the health of the equipment from condition monitoring data, rather than relying on fixed inspection schedules or waiting for a breakdown. This paper reviews recent work on sensor-based PdM for electrical equipment, with the emphasis on machine learning (ML) methods. The review starts with the basic ideas of condition monitoring, such as typical sensors, the physical quantities they measure, and the usual data pipeline. It then summarises the main ML methods, from classical algorithms like support vector machines and random forests to deep learning models such as convolutional neural networks, recurrent networks, and autoencoders. After that, representative studies from recent years are compared in terms of equipment type, sensor setup, feature extraction, model choice, and evaluation metrics, and some public datasets from the prognostics and health management community are mentioned. The paper also looks at current limitations, including data imbalance, scarce labelled fault data, weak cross-domain generalisation, interpretability, and deployment constraints, and outlines possible directions such as domain generalisation, multi-sensor fusion, explainable artificial intelligence (AI), and human-centred Industry 5.0 maintenance. Overall, ML-based PdM is technically feasible and increasingly mature, although further research is needed to ensure dependable performance in real industrial environments.

Keywords: predictive maintenance, electrical equipment, condition monitoring, machine learning, sensor data

1. Introduction

Electrical equipment sits at the core of modern power systems and industrial production. Rotating machines, power transformers, switchgear, and generators are expected to run more or less continuously for years, often in tough conditions with high temperature, vibration, humidity, and electrical stress. When such equipment fails unexpectedly, the consequences can range from unplanned production stops and damage to connected systems to safety hazards and large economic losses. Maintenance strategy is therefore a real concern for both utilities and manufacturers [1].

Traditional maintenance can be roughly split into two kinds: corrective and preventive. Corrective maintenance simply repairs the equipment after it fails. It is easy to manage, but the downtime is unpredictable, and secondary damage can make the repair much more expensive. Preventive maintenance replaces parts and carries out inspections at fixed time intervals. This lowers the chance of sudden failure, but it may also throw away components that are still healthy, which wastes labour and materials. It also cannot catch faults that develop between two scheduled inspections [2]. Predictive maintenance (PdM), which is closely related to condition-based maintenance, is a third option. Predictive maintenance monitors equipment condition through sensors, detects early signs of degradation, and schedules maintenance only when evidence indicates that intervention is necessary [3]. Done well, this cuts both unnecessary intervention and catastrophic failure.

A few technical trends have made this practical. Low-cost sensors, the Industrial Internet of Things (IIoT), and edge computing now allow vibration, current, voltage, temperature, acoustic, and partial discharge signals to be collected continuously and sent to local servers or cloud platforms [4]. Meanwhile, machine learning (ML) methods can learn complex nonlinear links between sensor measurements and equipment health that may be difficult to capture with purely model-based approaches. Data-driven PdM for industrial equipment has already been surveyed from several angles, such as data acquisition, feature engineering, fault diagnosis, and remaining useful life (RUL) prediction [2]. More recent reviews look at deep learning in the IIoT context [4], fault diagnosis for electrical machines [5], and maintenance approaches for the power industry in particular [1].

Even so, using ML-based PdM on real electrical equipment is not straightforward. Faults are rare compared with normal operation, so the training data are heavily imbalanced. Labels are expensive, because they usually need expert inspection or even destructive testing. A model trained on one machine or one operating condition often does badly on another machine, which limits generalisation [6]. On top of that, maintenance engineers want to know why a model raised an alarm, so interpretability has become an active topic [7]. Because of these practical problems, the gap between academic results and industrial deployment is still significant [8, 9].

This paper reviews recent progress in sensor-based predictive maintenance for electrical equipment using machine learning and identifies the principal limitations and promising research directions. It focuses on literature that deals directly with electrical or electromechanical assets, while also drawing on broader PdM surveys where they offer useful methodology. The rest of the paper is organised as follows. Section 2 introduces the basic concepts of condition monitoring for electrical equipment. Section 3 summarises the ML methods commonly used in this field. Section 4 compares representative recent studies in terms of equipment, sensors, features, models, and evaluation metrics. Section 5 discusses limitations and future prospects. Section 6 concludes the paper.

2. Basic concepts of electrical equipment condition monitoring

Condition monitoring (CM) is the foundation of predictive maintenance. Condition monitoring involves measuring physical quantities that reflect equipment health and analysing these measurements to detect abnormal behaviour [3]. For electrical equipment, the monitored quantities fall into a few broad groups.

The first group is mechanical signals, mainly vibration and acoustic emission. Although electrical equipment is primarily defined by its electrical function, faults such as bearing damage, rotor imbalance, misalignment, and looseness often first appear as changes in vibration. Accelerometers

mounted on the housing are the most common sensors here [5]. The second group is electrical signals, such as stator current, voltage, and power. A common example is motor current signature analysis (MCSA), which identifies characteristic fault frequencies in the current spectrum. Current-based monitoring is attractive because the sensors can sit in the motor control cabinet and never touch the machine itself. The third group is thermal signals. Infrared thermography and embedded temperature sensors are widely used on transformers, cables, and switchgear, since overheating often points to insulation degradation or a poor electrical contact [1]. The fourth group is specific to high-voltage equipment: partial discharge (PD) measurement, dissolved gas analysis (DGA) in transformer oil, and insulation resistance testing. These give fairly direct evidence of insulation condition, which matters most for transformers and rotating machine windings.

From the data side, a typical PdM pipeline has four stages [2]. Data acquisition comes first: sensor signals are sampled at a suitable rate and stored. The choice of sampling rate is a compromise, because high-frequency vibration analysis needs fast sampling while temperature trends can be recorded quite slowly. Next is preprocessing — filtering, denoising, segmentation, and dealing with missing values. Real industrial data are often noisy and incomplete, so preprocessing quality strongly influences all subsequent analyses [8]. The third stage is feature extraction. Common time-domain features include root mean square (RMS), kurtosis, and crest factor; frequency-domain features come from the fast Fourier transform (FFT); and time-frequency methods like the short-time Fourier transform and wavelet transform handle non-stationary signals. Deep learning models can also learn features on their own from raw or lightly processed signals [4]. The final stage is decision making, which includes fault detection, fault diagnosis, and prognosis, with prognosis commonly expressed as remaining useful life.

It also helps to place PdM inside the wider maintenance picture. Corrective maintenance is reactive, preventive maintenance is time-based, and predictive maintenance is condition-based and may become prescriptive by recommending appropriate actions [1, 10]. These basics are needed before the specific machine learning methods can be discussed.

3. Machine learning methods for sensor-based predictive maintenance

The machine learning methods used in PdM can be divided into three broad groups: classical supervised methods, unsupervised and anomaly-detection methods, and deep learning methods [4].

Classical supervised algorithms need labelled data, where each sample is tied to a known health state. Support vector machines (SVMs) are widely used for fault classification because they cope reasonably well with small datasets. Decision trees and random forests are popular too, since they are easy to train and give some idea of feature importance. k-nearest neighbours and naive Bayes classifiers show up in earlier studies as well. For prognosis, regression versions of these models, together with artificial neural networks, are used to map extracted features to RUL [2]. The catch is that classical methods depend heavily on handcrafted features, so feature engineering often ends up being the most time-consuming part of the whole workflow.

Unsupervised methods matter because labelled fault data are rare in practice. Clustering methods such as k-means can group operating states without any labels. Autoencoders, discussed below, can be trained only on healthy data and then used to flag anomalies through their reconstruction error. One-class SVM follows a similar logic. These approaches simply accept the reality that normal data are plentiful while fault data are not [8].

Convolutional neural networks (CNNs) can identify local patterns directly from raw vibration signals or time-frequency images and are widely used for fault diagnosis of rotating machinery [4, 5]. Recurrent neural networks (RNNs), especially long short-term memory (LSTM) networks,

model how signals evolve over time and are often used for RUL prediction. Autoencoders are used both for dimensionality reduction and for anomaly detection. More recently, transfer learning and domain adaptation methods have tried to reuse knowledge learned on one machine or operating condition on another, which helps with the data scarcity problem [6]. Hybrid architectures, such as CNN layers stacked with LSTM layers, are also common in recent papers.

Method selection should consider labelled-data availability and quality, signal sampling rate and duration, edge-computing resources, and the interpretability of model outputs, rather than relying only on small differences in benchmark accuracy [7, 9]. These points keep coming up in the comparison of recent studies in the next section.

4. Recent studies of sensor-based predictive maintenance

Recent literature shows a shift from small, single-signal experiments towards more integrated studies that consider the entire maintenance workflow. This section compares representative work in terms of equipment type, sensor configuration, features, models, and evaluation metrics, drawing mainly on publications from 2023 onwards.

Rotating electrical machines are still the most studied assets. Akbar et al. reviewed fault diagnosis techniques for electrical machines, covering induction motors, permanent magnet machines, and their typical faults in bearings, stator windings, and rotors [5]. Their comparison shows that vibration and stator current remain the dominant signals, but thermal imaging and stray flux measurement are getting more attention because they can reveal faults that vibration misses. Molęda et al. looked at maintenance for the power industry as a whole, including transformers, circuit breakers, and transmission assets [1]. They point out that the power sector is moving from corrective and time-based schemes towards condition-based and predictive ones, although the speed of adoption differs a lot between utilities. This is a useful reminder that PdM maturity depends not only on algorithms but also on organisational factors.

On the sensor and data side, there is a clear trend towards multi-sensor monitoring. Single-sensor systems are cheaper and simpler to install, but they only give one view of the equipment state. Fusing several sensors can combine complementary information and make the system more robust. Yang et al. proposed a decision-level fusion scheme combining ultrasonic guided-wave and vibration measurements for damage identification, and the combined sensing channels outperformed either channel used alone [11]. Their study is not limited to electrical machines, but the fusion idea transfers naturally to motor and generator monitoring, where electrical and mechanical signals exist side by side.

As for features and models, deep learning is prominent in recent publications. Wang et al. surveyed deep-learning-enabled PdM in the IIoT and grouped the methods by network type and application scenario [4]. They note that CNNs and their variants are widely used, while autoencoders are commonly used when labelled fault data are limited or unavailable. Hoffmann and Lasch reviewed artificial intelligence (AI) approaches for reducing industrial downtime and observed that many studies stop at the diagnosis stage, while full prognostic workflows — with RUL estimation and maintenance decision optimisation — are less common [9]. Murtaza et al. compared PdM under the Industry 4.0 and Industry 5.0 visions, and argued that human-centred, resilient maintenance needs models that work together with engineers instead of making fully automatic decisions [10]. Their case study also shows that deployment architecture, including edge computing and connectivity, matters just as much as model choice in practice.

Evaluation practices vary substantially across studies. Most diagnosis studies report accuracy, precision, recall, and F1-score, while prognostic studies use RUL-oriented metrics such as root

mean square error and scoring functions that punish late predictions more than early ones [2]. However, Nunes et al. point out that inconsistent datasets and metrics make it hard to compare studies fairly, and they call for more standardised evaluation [8]. Interpretability has also entered the evaluation discussion. Gawde et al. applied explainable AI techniques, including local interpretable model-agnostic explanations (LIME), SHapley additive explanations (SHAP), partial dependence plots (PDP), and individual conditional expectation (ICE), to predictive maintenance of rotating machines, and showed how feature-level explanations can make classifier decisions more transparent to engineers [7]. Work like this responds to the common complaint that deep models are black boxes.

Two other issues come up again and again in recent studies. One is cross-domain generalisation. Zhao et al. studied domain generalisation for fault diagnosis from an application-oriented angle and built a benchmark [6]. Their results show that models trained under one operating condition often degrade sharply under another, and that carefully designed domain generalisation methods can narrow this gap but not close it. The other issue is the availability of public datasets. Well-known sets like the Intelligent Maintenance Systems (IMS) bearing dataset and the National Aeronautics and Space Administration (NASA) Commercial Modular Aero-Propulsion System Simulation (C-MAPSS) turbofan dataset are used widely for benchmarking PdM pipelines. It is worth noting, though, that these datasets come from mechanical bearings and aircraft engines rather than electrical equipment. They are useful mainly as stand-ins for developing and testing prognostic workflows, and conclusions drawn from them should not be transferred to motors or transformers without care.

Taken together, the recent literature suggests that sensor-based PdM for electrical equipment is heading towards deeper models, richer sensor combinations, and more attention to deployment and interpretability. At the same time, standardised evaluation and generalisation across real operating conditions are still open problems.

5. Limitations and prospects

The reviewed studies show real progress, but several limitations still hold back the industrial use of ML-based predictive maintenance for electrical equipment. This section goes through them, together with possible future directions.

The most obvious limitation is the scarcity and imbalance of fault data. Electrical equipment is built to be reliable, so genuine failures are rare. Most of the collected data show healthy operation, and the few fault samples that do exist are often poorly documented. As a result, supervised models can reach high test accuracy on carefully balanced laboratory datasets and then fail to detect rare faults in the field [8]. Possible remedies include unsupervised and semi-supervised learning, anomaly detection trained only on healthy data, data augmentation, and physics-informed simulation that generates synthetic fault signals. Each of these has its own problems — the realism of simulated faults, for example — and they still need to be validated on real assets.

A second limitation is weak generalisation across machines, operating conditions, and sites. A model trained on one motor at constant speed may do poorly on a motor of different size, or on the same motor under a variable load. Domain adaptation and domain generalisation methods try to fix this, and benchmark studies show encouraging but still incomplete results [6]. In practice, some site-specific fine-tuning may still be needed, which adds to the deployment cost. Future work should aim for methods that need little or no labelled data from the target domain, and for reporting standards that make cross-domain performance visible instead of hiding it behind single-condition accuracy.

Interpretability and trust form a third limitation. Maintenance decisions can carry serious safety and economic consequences, so engineers are unlikely to trust a model that only outputs a fault

probability. Explainable AI methods such as SHAP and LIME have been shown to give useful feature-level explanations for rotating machine diagnosis [7]. Still, the explanations themselves can be unstable or even misleading, and there is no agreement on how to judge their quality. One promising direction is to combine data-driven models with physical knowledge, so that the model's behaviour can be checked against known fault mechanisms, such as the characteristic frequencies in vibration and current spectra [5].

A fourth limitation is the gap between prediction and decision. Many studies stop at fault classification or RUL estimation and never consider how the prediction should feed into maintenance scheduling, spare part management, or workforce planning [9]. A predicted failure with a very wide confidence interval may be less useful than a rougher prediction whose uncertainty is well quantified. Bringing uncertainty estimation and cost-sensitive decision making into PdM systems is therefore an important direction. Related to this, the Industry 5.0 discussion pushes for human-centred systems where ML supports maintenance engineers rather than replacing them [10]. Designing the interfaces and workflows for that kind of collaboration is still at an early stage.

The fifth limitation is practical deployment. Industrial environments put hard constraints on sensor installation, communication bandwidth, edge computing resources, cybersecurity, and integration with existing supervisory control and data acquisition (SCADA) or enterprise systems [3, 4]. A deep model that runs well on a workstation may be far too heavy for an embedded device, and streaming high-frequency raw signals all the time is expensive. Edge computing, model compression, and smarter data reduction — transmitting features instead of raw waveforms, for instance — are active areas that can narrow this gap. Multi-sensor fusion raises its own practical questions about synchronisation and sensor reliability, since a failed sensor can corrupt the fused result if the system has no redundancy built in [11].

There is also a methodological limitation in the literature itself. Datasets, preprocessing choices, and evaluation metrics differ widely from study to study, which makes fair comparison hard [8]. Public benchmark datasets help, but as noted in Section 4, the available ones mostly come from mechanical or aerospace test rigs rather than from electrical equipment in real plants. The community would benefit from open datasets collected from motors, transformers, and switchgear under realistic operating and failure conditions, together with agreed evaluation protocols. That kind of standardisation would make it much easier to tell genuine progress apart from small gains on a single test rig.

To sum up, the more promising prospects are: hybrid physics-informed and data-driven models; robust domain generalisation and adaptation; explainable and uncertainty-aware prediction; decision-level integration with maintenance planning; and lightweight edge deployment with multi-sensor fusion. Progress on these fronts would move PdM of electrical equipment from promising prototypes towards dependable industrial practice.

6. Conclusion

This paper has reviewed sensor-based predictive maintenance of electrical equipment with machine learning. The basic concepts of condition monitoring were introduced first, covering mechanical, electrical, thermal, and insulation-related measurements, along with the standard pipeline from data acquisition to prognosis. Classical and deep learning methods were then summarised, and recent literature increasingly uses deep models that learn features on their own. The comparison of recent studies shows growing interest in multi-sensor fusion, cross-domain generalisation, explainability, and deployment-oriented design. At the same time, major limitations remain, including scarce and imbalanced fault data, weak transfer across machines and operating conditions, limited

interpretability, and the gap between prediction and maintenance decisions. Dealing with these issues will be essential if research results are to become reliable industrial applications. Future work should also tie predictive models more closely to maintenance scheduling and cost analysis, so that the economic benefit of condition-based maintenance is shown in practice rather than simply assumed. Reliable deployment also requires evaluation under changing loads, sensor noise, missing data, and machine-to-machine variation rather than only balanced laboratory conditions. For high-consequence electrical assets, model accuracy should therefore be considered together with uncertainty, interpretability, computational cost, and the feasibility of sensor installation. Multi-sensor fusion and edge deployment can improve practical coverage, but they also introduce synchronisation, communication, and sensor-reliability constraints. Progress will depend on linking diagnostic and prognostic outputs to maintenance decisions while validating models on representative electrical equipment across sites and operating conditions. Such validation should also include rare faults and changing loads when these conditions are relevant to deployment.

References

- [1] Mołęda, M., Małysiak-Mrozek, B., Ding, W., Sunderam, V., & Mrozek, D. (2023). From corrective to predictive maintenance-A review of maintenance approaches for the power industry. *Sensors*, 23(13), Article 5970. <https://doi.org/10.3390/s23135970>
- [2] Lei, Y., Li, N., Guo, L., Li, N., Yan, T., & Lin, J. (2018). Machinery health prognostics: A systematic review from data acquisition to RUL prediction. *Mechanical Systems and Signal Processing*, 104, 799-834. <https://doi.org/10.1016/j.ymssp.2017.11.016>
- [3] Pech, M., Vrchota, J., & Bednář, J. (2021). Predictive maintenance and intelligent sensors in smart factory: Review. *Sensors*, 21(4), Article 1470. <https://doi.org/10.3390/s21041470>
- [4] Wang, H., Zhang, W., Yang, D., & Xiang, Y. (2023). Deep-learning-enabled predictive maintenance in Industrial Internet of Things: Methods, applications, and challenges. *IEEE Systems Journal*, 17(2), 2602-2615. <https://doi.org/10.1109/JSYST.2022.3193200>
- [5] Akbar, S., Vaimann, T., Asad, B., Kallaste, A., Sardar, M. U., & Kudelina, K. (2023). State-of-the-art techniques for fault diagnosis in electrical machines: Advancements and future directions. *Energies*, 16(17), Article 6345. <https://doi.org/10.3390/en16176345>
- [6] Zhao, C., Zio, E., & Shen, W. (2024). Domain generalization for cross-domain fault diagnosis: An application-oriented perspective and a benchmark study. *Reliability Engineering & System Safety*, 245, Article 109964. <https://doi.org/10.1016/j.res.2024.109964>
- [7] Gawde, S., Patil, S., Kumar, S., Kamat, P., Kotecha, K., & Alfarhood, S. (2024). Explainable predictive maintenance of rotating machines using LIME, SHAP, PDP, ICE. *IEEE Access*, 12, 29345–29361. <https://doi.org/10.1109/ACCESS.2024.3367110>
- [8] Nunes, P., Santos, J., & Rocha, E. (2023). Challenges in predictive maintenance-A review. *CIRP Journal of Manufacturing Science and Technology*, 40, 53–67. <https://doi.org/10.1016/j.cirpj.2022.11.004>
- [9] Hoffmann, M. A., & Lasch, R. (2024). Tackling industrial downtimes with artificial intelligence in data-driven maintenance. *ACM Computing Surveys*, 56(4), Article 82, 1–33. <https://doi.org/10.1145/3623378>
- [10] Murtaza, A. A., Saher, A., Zafar, M. H., Moosavi, S. K. R., Aftab, M. F., & Sanfilippo, F. (2024). Paradigm shift for predictive maintenance and condition monitoring from Industry 4.0 to Industry 5.0: A systematic review, challenges and case study. *Results in Engineering*, 24, Article 102935. <https://doi.org/10.1016/j.rineng.2024.102935>
- [11] Yang, X., Fang, C., Kundu, P., Yang, J., & Chronopoulos, D. (2024). A decision-level sensor fusion scheme integrating ultrasonic guided wave and vibration measurements for damage identification. *Mechanical Systems and Signal Processing*, 219, Article 111597. <https://doi.org/10.1016/j.ymssp.2024.111597>