



Predictive Maintenance Framework for Smart Electrical Distribution Infrastructure Using Artificial Intelligence and Data Analytics

Fahad Khan^{1*} 

Mechanical Engineering, Open University of Sri Lanka: Colombo, Western Province, LK

*Corresponding Author

Fahad Khan

Mechanical Engineering,
Open University of Sri
Lanka: Colombo, Western
Province, LK

Article History

Received: 14.08.2026

Accepted: 06.10.2026

Published: 13.10.2026

Abstract: Smart electrical distribution networks are becoming more observable, automated, and data intensive, yet maintenance practice often remains dominated by periodic inspection or intervention after alarms and failures. This review develops a predictive maintenance framework that connects artificial intelligence and data analytics with the operational realities of distribution transformers, on-load tap changers, switchgear, circuit breakers, feeders, cables, substations, smart meters, and distributed-energy interfaces. The evidence base is restricted to peer-reviewed work published from 2020 to 2025 and is synthesized through a structured integrative review focused on data provenance, condition indicators, learning tasks, validation practice, interpretability, uncertainty, deployment constraints, and maintenance decision linkage. The review shows that useful predictive maintenance cannot be reduced to selecting a high-performing classifier. Reliable deployment requires a chain of asset-aware sensing, data-quality control, representation learning, anomaly detection, fault diagnosis, health forecasting, uncertainty estimation, consequence assessment, and human-governed maintenance orchestration. Deep learning and ensemble methods improve representation of nonlinear and temporal degradation, while digital twins can provide synchronized context, scenario testing, and traceability; however, scarce fault labels, changing operating regimes, class imbalance, fragmented semantics, cyber risk, and weak external validation remain major obstacles. A layered framework is therefore proposed in which analytical outputs are converted into risk-calibrated maintenance actions rather than isolated predictions. The framework emphasizes evidence quality, cross-asset generalization, explainable reasoning, model monitoring, and feedback from completed work orders. It provides a research and implementation architecture for moving smart distribution maintenance from condition visibility toward dependable, auditable, and economically meaningful predictive decision support.

Keywords: Predictive Maintenance; Smart Distribution Networks; Artificial Intelligence; Data Analytics; Distribution Transformers; Condition Monitoring; Fault Diagnosis.

Copyright © 2026 The Author(s): This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY-NC 4.0) which permits unrestricted use, distribution, and reproduction in any medium for non-commercial use provided the original author and source are credited.

1. INTRODUCTION

Electrical distribution infrastructure is undergoing a structural transition from predominantly passive networks toward digitally

instrumented, bidirectional, and increasingly autonomous systems. Advanced metering, intelligent electronic devices, feeder automation, distributed energy resources, and dense sensing have expanded

Citation: Khan, F. (2026). Predictive Maintenance Framework for Smart Electrical Distribution Infrastructure Using Artificial Intelligence and Data Analytics. *Glob Acad J Econ Buss*, 8(5), 1742-1753. <https://doi.org/10.36348/gajeb.2026.v08i05.060>

the volume and resolution of operational data available to utilities. At the same time, distribution assets are exposed to thermal loading, insulation deterioration, mechanical wear, environmental stress, switching duty, harmonics, and rapidly changing power flows. These conditions make purely corrective maintenance costly and time-based preventive maintenance increasingly inefficient because calendar intervals do not represent the actual degradation state of heterogeneous assets. Recent reviews consequently position predictive maintenance as a key mechanism for anticipating faults, localizing abnormal behavior, and planning intervention before service continuity is compromised.

Predictive maintenance is often described as the application of condition data and predictive models to estimate the probability, mode, or timing of failure. For electrical distribution infrastructure, that definition is necessary but incomplete. A useful system must also determine whether data are trustworthy, distinguish transient operational variation from degradation, express uncertainty, consider asset criticality, and connect analytical evidence to a maintenance action that can be justified to engineers. Studies of digital-twin-enabled maintenance further show that decision quality depends on the synchronization of physical states, digital representations, analytics, and operational feedback rather than on algorithm performance alone. This wider system view is especially important in safety-critical distribution environments, where a false negative can contribute to an outage or equipment damage and a false positive can trigger unnecessary switching, inspection, or replacement.

The literature has advanced rapidly across several fronts, but the evidence remains fragmented. General predictive-maintenance surveys discuss deep learning architectures and implementation barriers, smart-grid reviews classify distribution faults and prediction methods, transformer studies emphasize dissolved-gas analysis, thermal indicators, partial discharge, and health indices, and digital-twin research focuses on architectural synchronization and interoperability. These streams are complementary, yet they are rarely integrated into a distribution-specific maintenance logic spanning data acquisition, prognostics, risk, and work execution. Model comparisons also frequently use different assets, sampling rates, labeling rules, failure definitions, and evaluation splits, making headline accuracy values difficult to compare directly.

This review addresses that gap by treating predictive maintenance as an end-to-end socio-technical capability for smart electrical distribution infrastructure. Rather than ranking algorithms, it

examines the conditions under which analytical methods can support dependable maintenance decisions. Particular attention is given to data quality, class imbalance, temporal leakage, cross-asset transferability, interpretability, uncertainty, cybersecurity, and the feedback required to keep models valid after deployment. The resulting framework is intended to support researchers designing evaluation studies and utilities planning scalable predictive-maintenance programs.

2. AIM, OBJECTIVES AND REVIEW QUESTIONS

The aim of this review is to develop an evidence-grounded framework for predictive maintenance of smart electrical distribution infrastructure by integrating artificial intelligence, multimodal condition data, risk analytics, and maintenance decision processes.

Five objectives guide the synthesis. First, the review identifies the asset classes, failure mechanisms, sensing modalities, and data-quality conditions that determine whether predictive models are technically meaningful. Second, it evaluates how machine learning, deep learning, anomaly detection, prognostics, and digital-twin methods contribute to fault recognition, health assessment, and remaining-useful-life estimation. Third, it examines the methodological weaknesses that limit transfer from experimental studies to operational distribution networks. Fourth, it defines how predictive outputs should be combined with uncertainty, consequence, and asset criticality to create risk-calibrated maintenance priorities. Fifth, it proposes a closed-loop framework connecting sensing, analytics, engineering interpretation, work execution, and model learning.

The review addresses four questions: What data and condition indicators provide defensible evidence of degradation across major distribution assets? Which analytical approaches are appropriate for anomaly detection, fault diagnosis, health forecasting, and maintenance timing? What validation, governance, and deployment controls are necessary for trustworthy use in live networks? How can model outputs be translated into maintenance decisions that improve reliability without creating excessive false alarms or avoidable intervention?

3. REVIEW METHODOLOGY

A structured integrative review was adopted because the topic combines heterogeneous engineering evidence rather than a single clinical-style intervention or a uniform outcome suitable for meta-analysis. The review design was informed by recent maintenance and smart-energy syntheses that use explicit databases, search strings, eligibility rules, and quality criteria to make evidence selection

auditable. The review period was restricted to 2020-2025 to capture contemporary artificial intelligence, digital-twin, edge-computing, and smart-grid practices while maintaining a coherent technological window.

The principal search domains were Scopus, Web of Science Core Collection, IEEE Xplore, and ScienceDirect. Search logic combined maintenance, distribution-infrastructure, and analytics concepts. A representative expression was: ("predictive maintenance" OR prognostic* OR "condition monitoring") AND ("distribution network" OR "smart grid" OR transformer* OR switchgear OR feeder* OR substation*) AND ("artificial intelligence" OR "machine learning" OR "deep learning" OR "data analytics" OR "digital twin"). Variants were adapted to database syntax and supplemented with targeted combinations for dissolved-gas analysis, partial discharge, anomaly detection, remaining useful life, health index, and fault classification.

Eligible studies were peer-reviewed journal or full conference papers written in English, published during the defined period, assigned a DOI, and concerned with maintenance prediction, condition assessment, fault diagnosis, degradation forecasting, or a directly transferable analytical architecture for electrical or comparable cyber-physical assets. Priority was given to distribution transformers and tap changers, switchgear and breakers, feeders and cables, substation equipment, smart-grid monitoring, and enabling architectures. Studies were excluded when they lacked a maintenance interpretation, reported only steady-state control or protection performance, provided insufficient methodological detail, duplicated another publication, fell outside the date range, or had no verifiable DOI. General predictive-maintenance and digital-twin papers were retained only when their methods addressed a deployment issue directly relevant to distribution infrastructure.

Quality appraisal emphasized engineering validity rather than a single numerical score. Each source was examined for clarity of asset and failure definition, provenance of measurements, labeling procedure, treatment of missing or noisy data, class balance, separation of training and test information, validation across time or equipment, appropriateness of metrics, uncertainty handling, explainability, computational requirements, and linkage between predictions and maintenance actions. Particular caution was applied to studies reporting random sample splits on strongly autocorrelated time series, since neighboring observations from the same degradation episode can inflate apparent generalization. Accuracy alone was also treated as insufficient for rare-fault settings;

recall, precision, F1 measures, false-negative behavior, calibration, and lead time are more informative when operational consequences are asymmetric.

Data were synthesized thematically rather than statistically because datasets, failure taxonomies, sampling frequencies, and evaluation endpoints were not sufficiently comparable for a defensible pooled effect size. The synthesis therefore maps evidence across six interacting themes: asset observability, data engineering, diagnostic analytics, prognostic analytics, digital context, and maintenance decision integration. Citation chaining was used to strengthen coverage of relevant 2020-2025 studies, but no unverified article counts or artificial screening totals are reported. This approach preserves methodological transparency without implying a precision that the heterogeneous literature cannot support.

4. ASSET DEGRADATION AND DATA FOUNDATIONS

Predictive maintenance begins with an observable degradation process, summarized across major assets in Table 1. In distribution transformers, useful evidence can include dissolved gases, top-oil and winding temperatures, loading, moisture, dielectric characteristics, vibration, acoustic emissions, partial discharge, and tap-changer operating behavior. Reviews of transformer condition monitoring show that no single measurement captures all failure modes; diagnostic confidence increases when electrical, thermal, chemical, and mechanical indicators are interpreted together. Machine-learning studies similarly demonstrate that fault labels derived from dissolved-gas patterns can support automated diagnosis, but performance depends heavily on the representativeness of historical cases and on how ambiguous or mixed faults are handled.

Switchgear and circuit breakers create a different observability problem. Their degradation may be episodic rather than continuously expressed, with information distributed across operation counts, coil current, travel signatures, contact resistance, timing, temperature, and protection-event logs. Feeders and cables add spatial complexity: incipient insulation problems, high-impedance faults, vegetation contact, weather stress, and joint degradation can produce weak or nonstationary signatures. Smart-grid fault studies therefore emphasize the value of voltage and current measurements, synchronized events, smart meters, and communication-enabled sensing for fault detection and localization [1]. Deep neural approaches can classify distribution faults effectively when training data represent the relevant operating

configurations, but a laboratory or simulation result should not be assumed to transfer unchanged to a feeder with different grounding, topology, distributed generation, sensor placement, or noise.

Data quality is consequently a first-class maintenance variable. Missing samples, clock drift, sensor recalibration, communication loss, outliers, changes in firmware, and inconsistent asset identifiers can create patterns that resemble degradation. A model may learn maintenance practice instead of physical condition when labels are generated from work orders that themselves reflect inspection policy. Likewise, using measurements collected after an alarm or intervention can create information leakage. Effective pipelines need time alignment, quality flags, range and plausibility checks, event reconciliation, provenance records, and

explicit distinction between raw measurements, derived features, labels, and engineer annotations.

Multimodal fusion should be purposeful rather than maximal. High-frequency waveforms can reveal transient signatures that minute-level supervisory data cannot, whereas slower thermal or chemical trends may be sufficient for long-horizon deterioration. Edge and industrial Internet-of-Things platforms provide a practical mechanism for collecting, preprocessing, and streaming distributed measurements, particularly where bandwidth or latency constraints make continuous cloud transfer undesirable. The design question is therefore not simply how much data can be collected, but which measurements provide incremental information about a specific failure mechanism at a time horizon useful for maintenance.

Table 1: Asset-specific evidence for predictive maintenance in smart electrical distribution infrastructure

Asset / subsystem	Representative data	AI / analytics tasks	Maintenance output	Key deployment concern
Distribution transformer / OLTC	DGA, oil and winding temperature, load, moisture, partial discharge, vibration, acoustic and tap-operation records	Fault classification; anomaly detection; health index; RUL / risk forecasting	Targeted test, inspection, oil treatment, tap maintenance, load management or replacement planning	Mixed faults, sparse failures, fleet differences, laboratory-to-field transfer
Switchgear / circuit breaker	Operation count, coil current, contact resistance, timing, travel signatures, thermal data and event logs	Signature classification; trend deviation; mechanism diagnosis	Inspection, lubrication, contact service, timing test or planned replacement	Intermittent evidence, limited labels, maintenance-induced distribution shift
Feeders / cables / joints	Voltage-current waveforms, smart-meter events, fault records, partial discharge, thermal and weather context	Fault detection / location; anomaly detection; degradation screening	Patrol, cable testing, joint inspection, vegetation or component intervention	Topology change, weather effects, spatial sparsity, high-impedance faults
Substation / smart-grid edge	SCADA / IED values, alarms, harmonics, power quality, communications and environment	Multivariate anomaly detection; event correlation; asset risk screening	Prioritized engineering review and coordinated maintenance window	Semantic heterogeneity, time synchronization, cyber integrity and alert volume
DER interface equipment	Inverter telemetry, temperature, harmonics, alarms, loading and firmware / configuration records	Drift detection; anomaly classification; fleet comparison	Inspection, firmware/configuration review, thermal or component service	Rapid technology change, firmware effects and nonstationary operating regimes

5. ARTIFICIAL INTELLIGENCE AND DATA-ANALYTICS METHODS

The analytical methods used in predictive maintenance can be grouped according to the kind of decision they aim to support. Anomaly detection seeks to determine if the current behaviour differs

from the normal state that has been learned. Fault diagnosis attempts to identify the most plausible failure mode. Health assessment provides an estimate of the condition state, either as a continuous or ordinal value. Prognostics gives an estimate of the degradation trajectory, the probability of failure, or

the remaining useful life. Since these tasks call for different types of labels and evaluation criteria, it is inappropriate to combine them without consideration and this can result in models that appear to be statistically impressive but are operationally unclear.

Even when the data are structured and the features that have been engineered are meaningful, classical supervised methods remain useful. Machines such as support-vector machines, decision trees, random forests, gradient boosting algorithms, and kernel methods are able to function with a moderate number of samples and usually allow for clearer attribution of features than deep networks do. Research into transformer fault diagnosis shows how optimised kernel machines and neural methods can be used to map dissolved-gas features to fault types. The main limitation of these methods is not their age but rather their dependence on stable feature distributions; if patterns of loading, the design of the equipment, or the sampling systems change, a classifier that has been trained on historical conditions can become miscalibrated even though the fundamental failure physics are still similar.

Deep learning has eliminated the need for manual feature engineering and is especially useful when dealing with images, waveforms, and multivariate sequences. With convolutional architectures it is possible to detect local patterns, recurrent models can account for temporal dependence, and attention-based transformers are able to capture long-range as well as inter-variable relationships. Various surveys have found that deep learning performs well in all the areas of predictive maintenance, yet they also point out that data hunger, interpretability, and deployment complexity remain ongoing obstacles. Hybrid prognostic models have been created for the purpose of predicting remaining useful life, and transformer-based anomaly detectors have improved the representation of the interactions between correlated variables. The practical conclusion is that greater model complexity should be based on an information advantage rather than simply on novelty.

Semi-supervised and self-supervised learning are particularly important since there are few confirmed electrical failures. Although utility companies have a large amount of data from normal operation, they have only a small number of well-labeled degradation trajectories. With transfer learning it is possible to reuse representations among similar equipment, and semi-supervised methods can make use of the unlabeled records while still relying on the smaller number of expert-labeled examples. Studies into transformer diagnosis have shown the benefits of transfer mechanisms in cases where fault

data are limited. Yet the domain shift has to be measured rather than simply ignored. Factors such as equipment rating, manufacturer, age, loading pattern, climate, and protection settings may all change the distribution of features.

The situation presents an extra difficulty in that the model has to make predictions early enough to provide maintenance value. A remaining-usable-life estimate which is only accurate just before failure may be less effective in practice than one that is less accurate but gives an earlier indication of risk. It is important to use Bayesian and probabilistic methods since they take uncertainty into account rather than providing a single definite estimate. An example of a prognostic framework using Bayesian deep learning shows how uncertainty can be combined with the prediction of degradation. In practical applications, prediction intervals and calibrated probabilities can lead to different actions: further observation when the level of uncertainty is high and the consequences are minor, a targeted inspection when risk is appearing, and immediate action when both the probability and the consequences are significant.

The evaluation of models should align with these maintenance decisions. Since random train-test splits are a poor indication of temporal degradation they might end up assigning similar operating states to both the training and test sets. More reliable approaches involve time-forward validation, testing with one asset left out, cross-site evaluation, and assessing the model across seasons or different operating conditions. In the case of classification, along with the area-under-curve figures class-specific recall, precision, F1, false-alarm rate and confusion structure should be provided. With regard to prognostics, absolute error should be supplemented by lead-time usefulness, interval coverage, calibration, and the economic cost of intervening too early or too late. A model which lowers the average error but fails to detect the few high-consequence failures could be unacceptable.

Explainability belongs to the process of engineering validation and should not just be included as a decorative element on a dashboard. In order to carry out their duties, maintenance engineers need to know which measurements, trends, or events led to an alert and also have to verify that the reasoning is physically plausible. Techniques such as feature attribution, counterfactual analysis, the visualisation of trends, and the identification of the nearest historical analogues can reveal spurious dependencies. Explainability must therefore be checked against existing domain knowledge and an understanding of possible failure mechanisms. It is especially important that human reviewers examine the situation when an algorithm suggests a switch, an

outage, or the replacement of a component since the cost of taking such action can be comparable to the cost of making a prediction error.

6. DIGITAL TWINS AND CONTEXTUAL INTELLIGENCE

Digital twins can be used to improve predictive maintenance by linking analytical models to a synchronized version of the asset and its network context. Recent reviews have identified a functional twin as being different from a static model by highlighting the importance of continuous data exchange, regular updates of the state, and the capacity to support diagnosis, simulation, or decision-making processes. In distribution systems, a twin is able to associate condition indicators with factors such as topology, asset configuration, loading, protection status, environmental context, and maintenance history. The relevance of this context lies in the fact that the same temperature rise or harmonic level could indicate different levels of risk depending on the loading and ambient conditions.

The fact that a twin replaces machine learning is not the case; instead, it offers a structured setting in which data, physics, and analytics can be brought together. By incorporating physics into the learning process, predictions can be restricted to plausible behaviour, and by means of scenario simulation it is possible to examine the effects of delaying maintenance or altering the loading. Moreover, hierarchical power-system twins show how the representations of individual assets can be coordinated with higher-level network models. In the context of predictive maintenance, this kind of hierarchy enables a move from individual component scores to a form of prioritisation that takes account of the network.

For a system to be scalable, interoperability is essential. The data concerning assets can come from supervisory systems, intelligent electronic devices, historian databases, laboratory systems, enterprise asset management tools, and mobile inspection records. A study into the interoperability of IEC 61850 stresses the need for consistent information models and for integration that is based on certification. If there is no semantic alignment, the same variables may have different names, different units, different timestamps, or different quality codes at various sites, which means that each predictive project has to develop its own fragile custom mappings. Thus, before advanced analytics are incorporated, a scalable framework must include a controlled asset identifier, a metadata model, agreed-upon unit conventions, defined time semantics, and traceable transformations.

The literature on digital twins also warns against going beyond what is necessary in terms of architecture. Since high-fidelity simulation is computationally costly, it is not essential that every maintenance decision be based on a full physical replica. The design should ensure that the level of fidelity matches the decision horizon; for risk screening, simple state synchronization will do, whereas in the case of important assets and high-consequence situations detailed electrothermal or mechanical simulation may be appropriate. The correct criterion is maintenance value per unit of complexity, not the degree of digital sophistication.

7. PROPOSED PREDICTIVE MAINTENANCE FRAMEWORK

The framework that has been proposed divides predictive maintenance into seven interconnected layers, as shown in Table 2 and in Figure 1. The physical layer includes transformers, tap changers, switchgear, breakers, feeders, cables, substations, smart meters, and interfaces with distributed energy sources. The acquisition layer collects data relating to electrical, thermal, chemical, mechanical, environmental, event, and maintenance aspects. The data-governance layer ensures that the timestamps are aligned, that the data quality is verified, that the asset identities are resolved, that the provenance is recorded, and that the data lineage is protected. It is only after these measures have been carried out that the analytical layer carries out feature learning, anomaly detection, fault diagnosis, health estimation, and prognostics.

The fifth layer involves the combination of uncertainty and risk. Analytical results are merged with the model's confidence level, the criticality of the asset, the consequences of a failure, the presence of redundancy, the loading conditions, the safety implications, and the possible times for intervention. This avoids a typical mistake made by predictive projects—that is, treating the model's probability as the basis for maintenance priority. For example, a fault probability of 0.7 on a redundant asset with a low impact might call for monitoring, whereas a lower probability on a single critical feeder transformer during peak demand could be sufficient to warrant an inspection. Thus, risk fusion converts statistical evidence into something that has engineering significance.

The sixth level is known as maintenance orchestration. Alerts are classified into various action categories, for example observe, inspect, schedule, derate, test, or intervene urgently. The recommendations contain the evidence in support of the action, the degree of uncertainty, the expected lead time, and the measurements that should be checked. By integrating with computerized

maintenance or enterprise asset management systems, approved recommendations are able to generate or enhance work orders rather than staying on a separate analytics dashboard. Industrial predictive-maintenance platforms show how important it is to link data collection with the operational workflow rather than treating analytics as a standalone application.

The seventh layer involves feedback and model operations. Once inspections have been carried out, test results obtained, components replaced, false alarms recorded, faults confirmed, and the performance after maintenance has been assessed, all this information is sent back to the evidence store. The labels thus produced are more valuable than those automatically inferred from alarms since they are based on engineering confirmation. The models are kept under observation for any drift in the distribution of inputs, in calibration, in the false-alarm rate, in missing data, and in performance when compared across different asset groups. Retraining is initiated on the basis of evidence of a change, not simply because a fixed period of time has elapsed. Version control ensures that the model, the feature pipeline, the training data window, the thresholds, and the approval status used for each past decision are preserved.

The decision logic employed in this architecture is shown in Figure 2. The measurements first go through a quality check; in the case of degraded data, confidence may be reduced or automated recommendations suppressed. An anomaly stage is used to detect deviations from normal behaviour. The diagnostic models then assess possible failure modes, after which health conditions or remaining lives are forecasted. Uncertainty is measured and combined with consequence and criticality. The final result is not a raw score but an

action band together with escalation rules, and the loop is closed by engineer confirmation, thus preventing autonomous escalation on the basis of an unverified model while at the same time permitting quick prioritisation.

Instead of insisting on a single analytical approach, this framework allows for the use of multiple analytical methods. For example, a tree ensemble would be suitable for tabular transformer data, a convolutional network for thermal images, an attention model for multivariate time series, and a physics-informed twin for scenario evaluation. The research into digital twins backs up this modular approach by demonstrating that synchronized representations can act as integration layers for various analytics and maintenance functions. Moreover, modularity also lowers the risk of having to replace components since a improved prognostic model can be incorporated without the need to rebuild the asset identity, the data-quality, the governance, or the work-order interfaces.

A key aspect of the design is to keep the generation of evidence separate from the power to make decisions. The analytics layer provides probabilistic evidence; the risk layer assesses the consequences; the orchestration layer suggests actions; and duly authorized engineers then approve the necessary interventions in accordance with operational procedures. It is important to have this separation for reasons of safety, accountability, and auditability, and it also makes any disagreements evident. An engineer might decide not to act on an alert because of known sensor maintenance, due to a planned switch, or because of field observations which are not available to the model. Recording the reason for this decision results in valuable feedback and highlights missing contextual factors.

Table 2: Functional layers of the proposed predictive maintenance framework

Layer	Core function	Primary inputs	Decision output	Assurance focus
1. Physical assets	Define failure mechanisms and operational context	Asset design, duty, topology, environment	Observable maintenance problem	Failure-mode relevance
2. Acquisition	Capture multimodal condition evidence	Sensors, IEDs, SCADA, inspections, work history	Time-stamped evidence streams	Coverage, latency, sensor health
3. Data governance	Create trustworthy analytical data	Raw measurements and metadata	Aligned, quality-tagged, traceable dataset	Identity, provenance, leakage prevention
4. AI analytics	Detect, diagnose and forecast condition	Features / sequences / images / events	Anomaly, fault, health and RUL estimates	Temporal / cross-asset validation
5. Risk fusion	Translate predictions into engineering significance	Prediction, uncertainty, criticality, consequence	Risk-calibrated priority	Calibration and asymmetric costs

Layer	Core function	Primary inputs	Decision output	Assurance focus
6. Maintenance orchestration	Convert priority into controlled action	Risk band, constraints, work capacity	Observe / inspect / schedule / urgent action	Human authorization and audit trail
7. Feedback & model operations	Learn from completed interventions and changing assets	Inspection findings, work results, drift indicators	Updated labels, thresholds and approved model versions	Drift, versioning, rollback and lifecycle governance




Figure 1: Layered AI-enabled predictive maintenance architecture for smart electrical distribution infrastructure




Figure 2: Risk-calibrated decision logic converting analytical evidence into governed maintenance action

8. DEPLOYMENT, CYBERSECURITY AND GOVERNANCE

The shift to operational deployment changes the focus from model development to the reliability of the managed service. The data pipelines have to run continuously, a sensor failure must not be taken to mean that the asset has failed, and the latency of inference has to correspond to the decision horizon. Edge processing can reduce the amount of bandwidth used and offer local resilience, whereas centralised platforms are able to combine evidence at the fleet level and assist with model governance. Architectural decisions should therefore separate out fast protection functions from maintenance analytics. Predictive maintenance should be used to guide either planned or condition-based interventions; it must not be confused with deterministic protection which has to trigger within strict real-time limits.

Predictive maintenance is closely linked with cybersecurity since compromised measurements can corrupt both the diagnosis and the prioritization process. The use of digital energy systems increases the attack surface by means of connected sensors, gateways, remote access, and cloud services. The controls that should be put in place include the use of authenticated data sources, implementation of least-privilege access, secure methods for updating systems, network segmentation, integrity checks, logging, and the ability to detect unexpected changes in measurements. It is also necessary for the inputs and outputs of the model to have provenance so that an engineer can work out which data led to a particular recommendation.

Governance should deal with issues of model ownership, approval, monitoring, and retirement. A model which has performed well during the commissioning phase may deteriorate if a feeder is reconfigured, if a new inverter fleet causes changes in harmonics, or if the sensor range is modified. Performance should therefore be divided according to asset class and operating context. Any changes to thresholds, retraining, or feature modifications must be accompanied by documented validation. This applies likewise to fairness in an engineering context: assets with sparse historical data should not be systematically given lower-quality predictions unless that limitation is made visible to the planners.

Explainability and uncertainty offer governance interfaces, enabling engineers to tell the difference between a believable degradation pattern and a correlation that has no physical basis. Standardized evidence packages can illustrate recent trends, the significant variables, similar historical events, confidence intervals, and data-quality alerts. Because of these records, maintenance recommendations become auditable and post-event

learning is improved. When digital-twin and predictive-maintenance reviews are carried out repeatedly, they consistently identify interoperability, implementation complexity, and validation as barriers to reliable deployment; governance is the means by which these obstacles are turned into specific engineering controls.

9. DISCUSSION: FROM PREDICTION ACCURACY TO MAINTENANCE VALUE

The literature shows that artificial intelligence is able to identify fault patterns and model degradation, but the most significant unresolved issue is not whether some other algorithm could increase benchmark accuracy. Rather, it is whether the predictions stay useful when they are applied to different assets, seasons, locations, and with changing network configurations. Smart distribution infrastructure is a constantly changing system: distributed generation causes the loading direction to change, automation modifies the switching patterns, sensors are replaced, and maintenance alters the group of surviving equipment. Models should therefore be seen as adaptive elements of an asset-management system rather than as fixed classifiers.

Another point is the difference between statistical performance and maintenance value. False positives use up inspection capacity and may cause operators to lose trust in the alarms. False negatives can allow damage to develop. Although early predictions offer more opportunities for planning they also involve a higher level of uncertainty. The appropriate threshold thus varies according to the consequences, the cost of intervention, the amount of redundancy available, and the lead time. Research into combining predictive maintenance with operational planning shows the importance of linking degradation forecasts with resource decisions rather than trying to optimize the predictions on their own. The framework suggested here goes beyond that principle when applying it to distribution networks by carrying out risk fusion and maintenance orchestration after the prediction has been made.

Another point concerns evidence portability. Although transformer studies offer a variety of techniques specific to a domain, more general research in deep learning and digital twins provides transferable architectures. To combine these approaches it is necessary to apply disciplined abstraction. It is reasonable to be able to transfer methods for uncertainty quantification, data lineage, temporal validation, or model monitoring. It is not reasonable, however, to suppose that a feature learned from one transformer fleet or from a simulated feeder will keep the same meaning in other contexts. Studies in the future should therefore

include sufficient contextual information to allow replication: such information should cover the asset type and rating, the operating environment, the sampling frequency, the source of the labels, the split strategy, the class distribution, the preprocessing steps, the hyperparameters, and the intervention horizon.

In the end, the review states that the level of maturity achieved in predictive maintenance should be assessed according to the closure of the decision loop. A project is considered immature if it simply concludes with an accuracy table or a dashboard. It only becomes operationally significant when alerts are connected to risk, engineers are able to examine the evidence, maintenance actions are logged, and the results lead to better future predictions. This criterion based on a closed loop offers a more rigorous and more realistic standard for the deployment by utilities.

10. RESEARCH GAPS AND FUTURE DIRECTIONS

The initial research gap lies in external validation. A great deal of research is carried out using only one laboratory, one simulated network, or one group of pieces of equipment. Cross-utility datasets are hard to share due to security concerns, commercial sensitivity, and inconsistent schemas, and without testing across multiple sites it is not clear whether the performance figures reported actually show general degradation patterns or merely local artifacts. Federated or privacy-preserving evaluation might offer a way of obtaining more extensive evidence while still keeping the raw operational data within the boundaries of the utility.

The second area involves rare-event and uncertainty-aware learning. Since faults are deliberately uncommon in well-kept networks, the datasets are mainly composed of normal operation. Research should go beyond simple synthetic balancing and instead assess calibrated probabilities, conformal or Bayesian uncertainty, positive-unlabeled learning, self-supervised representations, and expert-in-the-loop labeling. When making maintenance decisions, there should be a clear distinction between the absence of evidence and evidence of health.

The fourth area is network-aware prognostics. While most models assess the condition of individual assets, distribution risk is influenced by topology, redundancy, switching options, weather conditions, loading, and the presence of distributed resources. Approaches based on graphs and those based on twin models could combine information about component health with details of power flow and contingency situations, so that maintenance priorities take into account the impact on the system

rather than considering the condition of the asset alone.

The fifth area involves lifecycle evaluation; most studies do not follow up to see if the predictive recommendations led to a reduction in outages, maintenance costs, or risk over a number of years. In future research it is necessary to report the yield of the intervention, the number of failures avoided, the number of false inspections, the distribution of lead times, the amount of calibration drift, and the net economic value. Stronger evidence could be obtained through controlled field trials coupled with gradually increasing human supervision.

The fifth area is that of secure and sustainable AI. Since predictive models need to be monitored, versioned, have access control implemented, allow for rollbacks, and include procedures for dealing with incidents, research should look into how cyber anomalies, sensor spoofing, or failures in the data pipeline affect maintenance decisions and how resilient analytics can fail safely. It is essential that utilities be able to trust the increasingly automated recommendations.

11. PRACTICAL AND MANAGERIAL IMPLICATIONS

Utilities should start by considering failure modes and making maintenance decisions, not by carrying out an exercise to procure an algorithm. A practical approach involves identifying the most important assets, determining which decisions need to be improved, checking the data that is available, and then choosing a small number of measurable objectives such as earlier detection, fewer emergency interventions, or a better inspection yield. During the early stages, investment in data engineering and in asset identification should be given greater priority than in model complexity.

The implementation must take place in stages. This should begin with offline retrospective testing, then move on to a shadow-mode deployment in which the model produces alerts without carrying out any actions. The engineers then assess the quality of the alerts, the lead time, and the explanatory evidence. Recommendations should not be incorporated into the formal maintenance planning until stable performance has been achieved. For important decisions, human approval must be retained and clear rollback procedures must be in place.

Management should also regard predictive maintenance as a portfolio capability, since common data standards, model governance and feedback procedures enable the various methods to be applied across transformers, switchgear and feeders. Although separate pilot projects may show that a

technique is technically feasible they tend to yield only minor gains in terms of organisational learning. A reusable architecture together with a common evidence model help to reduce duplication and make it possible to compare performance between different types of assets.

12. CONCLUSION

Predictive maintenance in the case of smart electrical distribution infrastructure ought to be viewed as a decision system that is governed rather than as a fault-classification task. Current evidence demonstrates a considerable level of capability in the areas of machine learning, deep learning, prognostics, transformer diagnostics, smart-grid fault recognition, industrial data platforms, and digital twins; the value of these capabilities is seen only when they are linked together through reliable sensing, data provenance, asset context, uncertainty, risk, and the maintenance workflow.

The framework proposed in this review separates data acquisition, governance, analytics, risk fusion, maintenance orchestration, and learning feedback while preserving human decision authority. This structure makes model performance only one element of operational trust. For research, it shifts evaluation toward temporal and cross-asset validation, calibrated uncertainty, lead time, and intervention outcomes. For utilities, it provides a path from fragmented condition monitoring to auditable, risk-based maintenance intelligence. The most consequential advances will therefore come not from increasingly complex models alone, but from stronger evidence, interoperable data, secure deployment, and closed-loop learning that converts predictions into demonstrable reliability and maintenance value.

REFERENCES

1. Mahmoud, M.A., Md Nasir, N.R., Gurunathan, M., Raj, P., Mostafa, S.A. The current state of the art in research on predictive maintenance in smart grid distribution network: fault's types, causes, and prediction methods—a systematic review. *Energies* 14, 5078 (2021). <https://doi.org/10.3390/en14165078>
2. Serradilla, O., Zugasti, E., Rodriguez, J., Zurutuza, U. Deep learning models for predictive maintenance: a survey, comparison, challenges and prospects. *Applied Intelligence* 52, 10934–10964 (2022). <https://doi.org/10.1007/s10489-021-03004-y>
3. Li, Z., He, Q., Li, J. A survey of deep learning-driven architecture for predictive maintenance. *Engineering Applications of Artificial Intelligence* 133, 108285 (2024). <https://doi.org/10.1016/j.engappai.2024.108285>
4. Zhong, D., Xia, Z., Zhu, Y., Duan, J. Overview of predictive maintenance based on digital twin technology. *Heliyon* 9, e14534 (2023). <https://doi.org/10.1016/j.heliyon.2023.e14534>
5. Zhuang, L., Xu, A., Wang, X.-L. A prognostic driven predictive maintenance framework based on Bayesian deep learning. *Reliability Engineering & System Safety* 234, 109181 (2023). <https://doi.org/10.1016/j.res.2023.109181>
6. Chen, C., Fu, H., Zheng, Y., Tao, F., Liu, Y. The advance of digital twin for predictive maintenance: the role and function of machine learning. *Journal of Manufacturing Systems* 71, 581–594 (2023). <https://doi.org/10.1016/j.jmsy.2023.10.010>
7. Ma, S., Flanigan, K.A., Bergés, M. State-of-the-art review and synthesis: a requirement-based roadmap for standardized predictive maintenance automation using digital twin technologies. *Advanced Engineering Informatics* 62, 102800 (2024). <https://doi.org/10.1016/j.aei.2024.102800>
8. Abd Wahab, N.H., Hasikin, K., Lai, K.W., Xia, K., Bei, L., Huang, K., Wu, X. Systematic review of predictive maintenance and digital twin technologies challenges, opportunities, and best practices. *PeerJ Computer Science* 10, e1943 (2024). <https://doi.org/10.7717/peerj-cs.1943>
9. Yassin, M.A.M., Shrestha, A., Rabie, S. Digital twin in power system research and development: principle, scope, and challenges. *Energy Reviews* 2, 100039 (2023). <https://doi.org/10.1016/j.enrev.2023.100039>
10. Heluany, J.B., Gkioulos, V. A review on digital twins for power generation and distribution. *International Journal of Information Security* 23, 1171–1195 (2024). <https://doi.org/10.1007/s10207-023-00784-x>
11. Mchirgui, N., Quadar, N., Kraiem, H., Lakhssassi, A. The applications and challenges of digital twin technology in smart grids: a comprehensive review. *Applied Sciences* 14, 10933 (2024). <https://doi.org/10.3390/app142310933>
12. Aghazadeh Ardebili, A., Zappatore, M., Ramadan, A.I.H.A., Longo, A., Ficarella, A. Digital twins of smart energy systems: a systematic literature review on enablers, design, management and computational challenges. *Energy Informatics* 7, 94 (2024). <https://doi.org/10.1186/s42162-024-00385-5>
13. Ayello, M., Lopes, Y. Interoperability based on IEC 61850 standard: systematic literature review, certification method proposal, and case study. *Electric Power Systems Research* 220, 109355 (2023). <https://doi.org/10.1016/j.epr.2023.109355>
14. Ruhe, S., Schaefer, K., Branz, S., Nicolai, S., Bretschneider, P., Westermann, D. Design and implementation of a hierarchical digital twin for

- power systems using real-time simulation. *Electronics* 12, 2747 (2023). <https://doi.org/10.3390/electronics12122747>
15. Zhang, J., Zhao, X. Digital twin of wind farms via physics-informed deep learning. *Energy Conversion and Management* 293, 117507 (2023). <https://doi.org/10.1016/j.enconman.2023.117507>
16. Lu, Y., Liu, C., Wang, K.I.-K., Huang, H., Xu, X. Digital Twin-driven smart manufacturing: connotation, reference model, applications and research issues. *Robotics and Computer-Integrated Manufacturing* 61, 101837 (2020). <https://doi.org/10.1016/j.rcim.2019.101837>
17. Dehghan Shoorkand, H., Nourelfath, M., Hajji, A. A hybrid deep learning approach to integrate predictive maintenance and production planning for multi-state systems. *Journal of Manufacturing Systems* 74, 397–410 (2024). <https://doi.org/10.1016/j.jmsy.2024.04.005>
18. Kang, H., Kang, P. Transformer-based multivariate time series anomaly detection using inter-variable attention mechanism. *Knowledge-Based Systems* 290, 111507 (2024). <https://doi.org/10.1016/j.knosys.2024.111507>
19. Xie, Z., Du, S., Lv, J., Deng, Y., Jia, S. A hybrid prognostics deep learning model for remaining useful life prediction. *Electronics* 10, 39 (2021). <https://doi.org/10.3390/electronics10010039>
20. Yang, J., Sun, Y., Cao, Y., Hu, X. Predictive maintenance for switch machine based on digital twins. *Information* 12, 485 (2021). <https://doi.org/10.3390/info12110485>
21. Resende, C., Folgado, D., Oliveira, J., Franco, B., Moreira, W., Oliveira-Jr, A., Cavaleiro, A., Carvalho, R. TIP4.0: Industrial Internet of Things platform for predictive maintenance. *Sensors* 21, 4676 (2021). <https://doi.org/10.3390/s21144676>
22. Esmaeili Nezhad, A., Samimi, M.H. A review of the applications of machine learning in the condition monitoring of transformers. *Energy Systems* 15, 463–493 (2024). <https://doi.org/10.1007/s12667-022-00532-5>
23. Jin, L., Kim, D., Abu-Siada, A., Kumar, S. Oil-immersed power transformer condition monitoring methodologies: a review. *Energies* 15, 3379 (2022). <https://doi.org/10.3390/en15093379>
24. Ismail, F.B., Mazwan, M., Al-Faiz, H., Marsadek, M., Hasini, H., Al-Bazi, A., Ghazali, Y.Z.Y. An offline and online approach to the OLTC condition monitoring: a review. *Energies* 15, 6435 (2022). <https://doi.org/10.3390/en15176435>
25. Wang, J., Zhang, X., Zhang, F., Wan, J., Kou, L., Ke, W. Review on evolution of intelligent algorithms for transformer condition assessment. *Frontiers in Energy Research* 10, 904109 (2022). <https://doi.org/10.3389/fenrg.2022.904109>
26. Mao, W., Wei, B., Xu, X., Chen, L., Wu, T., Peng, Z., Ren, C. Fault diagnosis for power transformers through semi-supervised transfer learning. *Sensors* 22, 4470 (2022). <https://doi.org/10.3390/s22124470>
27. Kim, S., Park, J., Kim, W., Jo, S.-H., Youn, B.D. Learning from even a weak teacher: bridging rule-based Duval method and a deep neural network for power transformer fault diagnosis. *International Journal of Electrical Power & Energy Systems* 136, 107619 (2022). <https://doi.org/10.1016/j.ijepes.2021.107619>
28. Han, X., Ma, S., Shi, Z., An, G., Du, Z., Zhao, C. A novel power transformer fault diagnosis model based on Harris-Hawks-Optimization algorithm optimized kernel extreme learning machine. *Journal of Electrical Engineering & Technology* 17, 1993–2001 (2022). <https://doi.org/10.1007/s42835-022-01000-x>
29. Alhanaf, A.S., Balik, H.H., Farsadi, M. Intelligent fault detection and classification schemes for smart grids based on deep neural networks. *Energies* 16, 7680 (2023). <https://doi.org/10.3390/en16227680>
30. Mbey, C.F., Foba Kakeu, V.J., Boum, A.T., Souhe, F.G.Y. Fault detection and classification using deep learning method and neuro-fuzzy algorithm in a smart distribution grid. *The Journal of Engineering* 2023, e12324 (2023). <https://doi.org/10.1049/tje2.12324>