# Predictive Maintenance of Power Transformers Using Quantum Machine Learning Algorithms: A Framework and Literature-Based Review

Damini Choukse[1], Prathmesh Kolte[2]
[1]*Department of Electronics and Power Mauli Group of Institutions' College of Engineering & Technology Shegaon, India*
[2]*Department of Information Technology Mauli Group of Institutions' College of Engineering & Technology Shegaon, India*

***Abstract*—Power transformers are among the most critical and capital-intensive assets in electrical transmission and distribution networks, and their unplanned failure can trigger cascading outages and substantial economic loss. Traditional maintenance philosophies, whether time-based or reactive, are increasingly being displaced by predictive maintenance (PdM) strategies that use condition-monitoring data, most notably dissolved gas analysis (DGA), to anticipate incipient faults before they escalate. Classical machine learning (ML) methods such as support vector machines, decision trees, and artificial neural networks have demonstrated strong results in this domain, yet they face scalability and feature-space limitations as sensor networks generate increasingly high-dimensional, noisy data streams. This paper surveys classical ML-based transformer PdM and examines how quantum machine learning (QML) algorithms, including quantum support vector machines, variational quantum classifiers, quantum neural networks, and quantum annealing-based optimization, could extend these capabilities by exploiting high-dimensional Hilbert spaces and quantum parallelism. Building on this literature, a hybrid quantum-classical framework for transformer health monitoring is proposed and illustrated with diagrams, comparative tables, and literature-reported performance figures. The paper concludes by outlining the hardware, noise, and scalability barriers that currently constrain quantum-enhanced PdM and identifies promising directions for near-term research.**

***Index Terms*—Predictive maintenance, power transformer, dissolved gas analysis, quantum machine learning, quantum support vector machine, variational quantum classifier, NISQ**

## I. INTRODUCTION

Power transformers form the backbone of every electrical grid, stepping voltage up for efficient transmission and down again for safe distribution. A single large power transformer can represent an investment of several million dollars, and its unplanned failure can cause extended outages, safety hazards from oil fires, and cascading disruption across an interconnected network. For decades, utilities relied on time-based maintenance, in which equipment is inspected or replaced on a fixed schedule regardless of its actual condition, or on reactive maintenance, in which repairs occur only after failure. Both strategies are economically inefficient: the former wastes resources on healthy assets, while the latter risks catastrophic and costly failures [1].

Predictive maintenance (PdM) addresses this gap by continuously monitoring an asset's actual condition and forecasting when intervention is required. In the case of oil-immersed power transformers, the most widely used diagnostic signal is dissolved gas analysis (DGA), which detects trace gases produced when insulating oil and paper degrade under thermal or electrical stress. Over the last decade, classical machine learning techniques, including support vector machines (SVM), k-nearest neighbors (KNN), decision trees, random forests, and increasingly deep neural networks, have been applied to DGA and sensor data to classify incipient faults and estimate remaining useful life with considerable success [1].

At the same time, quantum computing has progressed from a purely theoretical pursuit to small- and medium-scale hardware accessible through cloud platforms, giving rise to the field of quantum machine learning (QML). QML algorithms attempt to exploit superposition, entanglement, and high-dimensional Hilbert-space representations to process data in ways that may, for certain problem classes, be difficult for classical computers to replicate efficiently [2]. Direct applications of QML to power-transformer diagnostics remain scarce in the published literature; the closest and most relevant body of work applies hybrid quantum-classical deep learning to fault diagnosis in broader electrical power systems and industrial process systems [3], [4].

This paper has three goals. First, it reviews classical ML-based predictive maintenance of power transformers and the diagnostic standards that underpin it. Second, it introduces the core families of quantum machine learning algorithms and explains, at a conceptual level, how each could plausibly be adapted to transformer condition monitoring and maintenance scheduling. Third, it proposes a hybrid quantum-classical framework and illustrates it with architecture diagrams, a schematic quantum circuit, comparison tables, and a chart of performance figures drawn from published case studies in adjacent domains. Because no large-scale, transformer-specific QML experiments currently exist in the public literature, the quantitative figures presented here are explicitly literature-derived and comparative rather than the results of new experiments conducted for this paper; this distinction is maintained throughout.

## II. TRANSFORMER FAILURE MECHANISMS AND CONDITION MONITORING

### A. Common Failure Modes

Power transformers fail for a range of interrelated reasons. Typical malfunctions include dielectric breakdown of the insulating oil and paper, thermal losses arising from copper (I2R) heating, mechanical distortion of windings caused by short-circuit forces, bushing failures, tap-changer malfunctions, core faults, tank integrity loss, and failure of the associated cooling or protection systems [5]. Because many of these degradation processes are gradual, they leave measurable chemical and physical traces long before a catastrophic failure occurs, which is precisely what makes condition-based and predictive maintenance strategies viable.

### B. Dissolved Gas Analysis and Diagnostic Standards

When the mineral oil and cellulose insulation inside a transformer are subjected to thermal or electrical stress, they decompose and release characteristic gases into the oil, including hydrogen (H2), methane (CH4), ethane (C2H6), ethylene (C2H4), acetylene (C2H2), carbon monoxide (CO), and carbon dioxide (CO2). Dissolved gas analysis interprets the relative concentrations of these gases to infer the type and severity of an incipient fault, and it remains the most trusted non-invasive diagnostic tool in the industry [6]. Several standardized interpretation schemes have been developed, including the IEC 60599 gas-ratio method, the IEEE C57.104 guide, the Rogers ratio method, and graphical techniques such as the Duval Triangle and Duval Pentagon, which plot the relative proportions of key gases to classify faults into categories such as partial discharge (PD), low-energy discharge (D1), high-energy discharge (D2), and thermal faults of increasing severity (T1, T2, T3) [7], [8], [9]. Table I summarizes representative fault categories and their associated key gases, and Fig. 1 illustrates, in simplified schematic form, the general layout of fault zones used in triangular DGA interpretation methods.

TABLE I Representative Transformer Fault Categories and Associated Key Dissolved Gases (adapted from [6], [7])

| Fault Code | Fault Type | Dominant Key Gases |
|---|---|---|
| PD | Partial discharge | Hydrogen (H2) |
| D1 | Low-energy discharge | Acetylene (C2H2), Hydrogen (H2) |
| D2 | High-energy discharge (arcing) | Acetylene (C2H2), Ethylene (C2H4) |
| T1 | Thermal fault, < 300°C | Methane (CH4), Ethane (C2H6) |
| T2 | Thermal fault, 300–700°C | Ethylene (C2H4), Methane (CH4) |
| T3 | Thermal fault, > 700°C | Ethylene (C2H4), trace Acetylene (C2H2) |

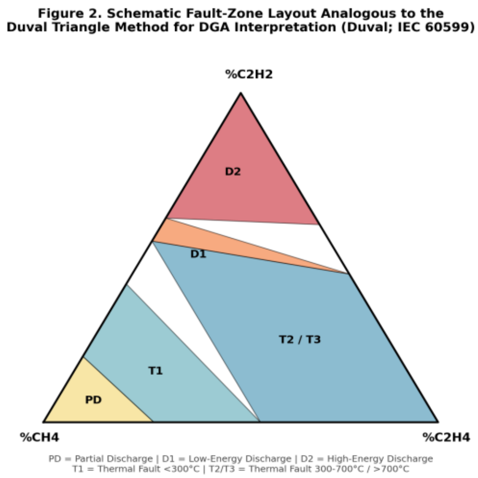


Fig. 1. Schematic fault-zone layout analogous to the Duval Triangle interpretation method (conceptual illustration; not to precise IEC scale).

## III. CLASSICAL MACHINE LEARNING FOR TRANSFORMER PREDICTIVE MAINTENANCE

### A. Supervised Learning Approaches

Supervised classifiers trained on labeled DGA and sensor data are the most common classical ML tool in transformer PdM. A representative case study applying SVM, KNN, and decision-tree classifiers to transformer fault data reported testing accuracies of 95.65% for SVM, 95.65% for KNN, and 89.13% for the decision tree, with SVM and KNN outperforming the decision tree model [5]. A broad systematic review of 126 peer-reviewed studies published between 2015 and 2024 found that neural networks, support vector machines, decision trees, and ensemble methods dominate the published AI approaches to transformer PdM [1].

### B. Deep Learning and Hybrid AI-IoT Systems

More recent work integrates deep learning with Internet of Things (IoT) sensor networks, cloud-edge computing, and digital twin frameworks to fuse dissolved gas analysis with vibration data, thermal imaging, and partial-discharge measurements into a unified transformer health index [1]. These hybrid systems aim to move beyond single-signal diagnostics toward continuous, multi-modal condition assessment that can also estimate remaining useful life rather than simply classifying a fault after it has already begun.

### C. Limitations of Classical Approaches

Despite these advances, classical ML-based PdM faces recurring limitations. Fault classes are often severely imbalanced because healthy operation vastly outnumbers fault events in real fleets, and adding more diagnostic modalities increases the dimensionality of the feature space, which can degrade the performance of distance-based and kernel-based classifiers. Deep models, in turn, require large labeled datasets that are difficult to obtain for rare, severe fault types and often behave as opaque “black boxes” that are hard to trust in safety-critical grid operations. These constraints motivate exploratory interest in whether quantum-enhanced feature spaces and quantum-assisted training could offer a complementary, rather than a proven superior, path forward, a distinction the quantum machine learning community itself emphasizes remains an open empirical question [10].

## IV. FOUNDATIONS OF QUANTUM COMPUTING AND QUANTUM MACHINE LEARNING

### A. Qubits, Superposition, and Entanglement

Classical bits take a definite value of 0 or 1, whereas a quantum bit, or qubit, can exist in a superposition of both states simultaneously, and multiple qubits can become entangled such that their joint state cannot be described independently. Quantum machine learning is the study of how these properties, together with parameterized quantum circuits, can be combined with learning algorithms to represent and process data [2]. Encoding classical data into an exponentially large Hilbert space is the central mechanism that most QML classification algorithms rely on to seek patterns that may be difficult to separate in the original feature space.

### B. The NISQ Era

Present-day quantum processors fall into what is commonly called the Noisy Intermediate-Scale Quantum (NISQ) era: devices with on the order of tens to a few hundred qubits, without full error correction, and subject to gate noise and decoherence [11]. This constrains circuit depth and qubit count and is the primary reason that most practical QML proposals, including those discussed below, are hybrid algorithms in which a shallow quantum circuit is combined with a classical optimizer rather than relying on a purely quantum pipeline.

C. Core Quantum Machine Learning Algorithms

1) Quantum Support Vector Machines (QSVM): The original quantum support vector machine proposal showed that, under certain data-loading assumptions, a least-squares SVM could in principle be solved with exponential speed-up relative to classical methods [12]. A more hardware-feasible variant maps classical feature vectors into a quantum-enhanced feature space using a shallow feature-map circuit and then estimates a quantum kernel that is passed to a classical SVM solver, an approach demonstrated on real near-term quantum hardware [13]. For transformer PdM, the classical features would be DGA gas concentrations or ratios, and the quantum kernel would attempt to expose nonlinear fault-class separability that is difficult to achieve with a classical kernel.

2) Variational Quantum Classifiers and Quantum Neural Networks: A variational quantum classifier (VQC), sometimes described as a form of quantum neural network, replaces the kernel-plus-classical-solver structure of QSVM with an end-to-end trainable quantum circuit. Classical data are encoded through a feature-map circuit, passed through a parameterized ansatz whose rotation angles are the trainable weights, and measured; a classical optimizer then updates the parameters based on the observed loss, forming a hybrid quantum-classical training loop [13]. Fig. 3 illustrates a schematic four-qubit circuit of this type as it might be applied to a DGA feature vector.

3) Quantum Annealing for Maintenance Scheduling: Beyond classification, predictive maintenance also involves combinatorial optimization: deciding which transformers to inspect first, how to route maintenance crews, and how to allocate spare parts under budget and time constraints. Quantum annealers, most notably D-Wave systems, are specialized for solving such optimization problems when formulated as quadratic unconstrained binary optimization (QUBO) problems. In the energy sector specifically, D-Wave has partnered with the utility E. ON to apply quantum annealing to electric grid partitioning and configuration optimization, demonstrating a real-world precedent for quantum-annealing use in grid operations [14]. Maintenance scheduling and crew-routing problems share the same combinatorial structure and could plausibly be formulated the same way, although this remains a conceptual extension rather than a demonstrated application specific to transformers.

4) Hybrid Quantum-Classical Deep Learning: The most directly relevant precedent for this paper's proposal comes from hybrid quantum-classical deep learning applied to fault diagnosis in electrical power systems. Ajagekar and You combined a conditional restricted Boltzmann machine with quantum-assisted generative training on an adiabatic quantum computer to diagnose faults on IEEE 30-bus and 14-bus test systems, reporting improved diagnostic performance and faster response time relative to classical artificial neural network and decision-tree baselines [3]. In a related industrial process study, the same authors reported average fault detection rates of 79.2% for a continuous stirred tank reactor process and 99.39% for the Tennessee Eastman process using a comparable quantum-assisted deep belief network [4]. While neither study addresses power transformers directly, both illustrate that hybrid QC-classical deep learning is a demonstrated, hardware-tested approach to fault diagnosis in structurally similar industrial and power-system settings.

## V. PROPOSED HYBRID QUANTUM-CLASSICAL FRAMEWORK

Drawing on the classical PdM literature and the QML precedents described above, this paper proposes a four-stage hybrid framework, shown in Fig. 2. The framework is a conceptual synthesis intended to guide future empirical work rather than a system that has been built and validated for this paper.

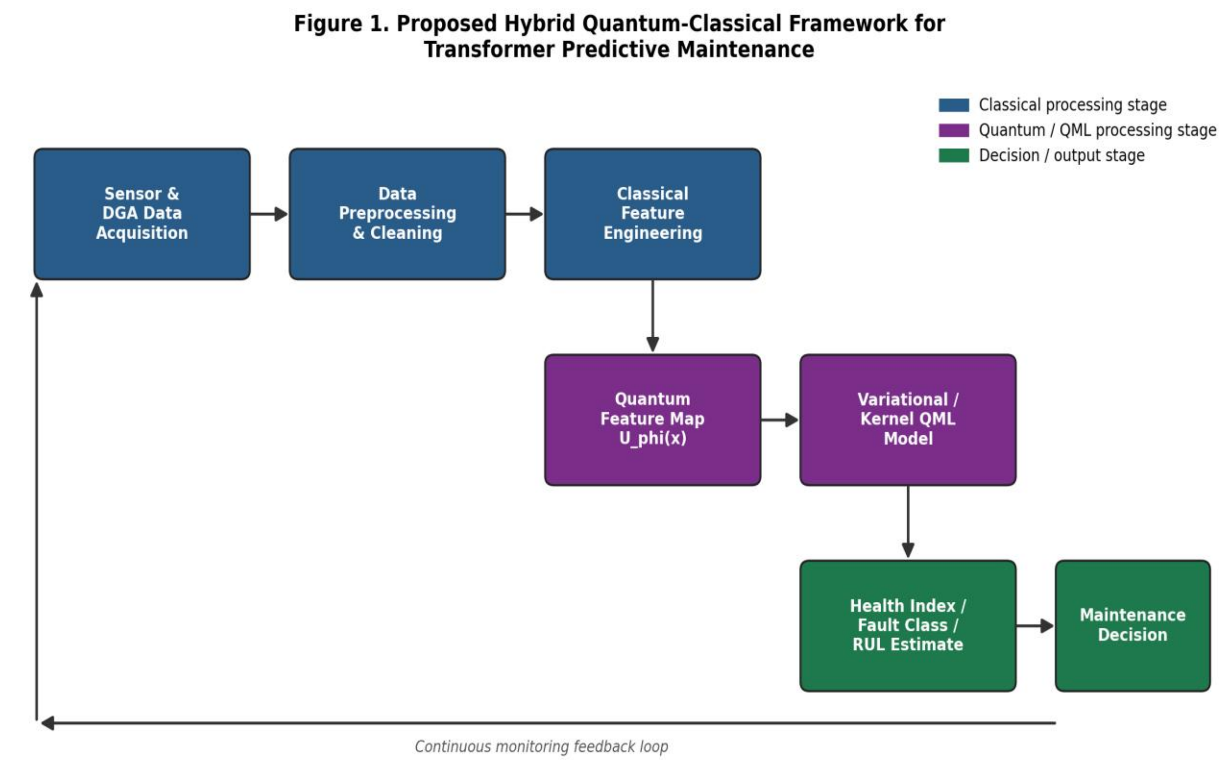


Fig. 2. Proposed hybrid quantum-classical framework for transformer predictive maintenance.

A. Data Layer

The data layer aggregates dissolved gas analysis results, online DGA sensor streams, thermal and load data, partial-discharge measurements, and SCADA time series, followed by classical preprocessing: missing-value handling, denoising, normalization, and computation of standard diagnostic ratios such as the IEC 60599 gas ratios used in Duval-style interpretation.

B. Quantum Feature Mapping and Learning Layer

Preprocessed features are encoded into a shallow quantum circuit using an encoding scheme such as angle encoding, in which each feature is mapped to a qubit rotation angle [13]. Depending on the deployment scenario, the framework can use either a QSVM-style quantum kernel evaluated on a quantum device and passed to a classical support vector solver, or an end-to-end variational quantum classifier trained through the hybrid loop illustrated in Fig. 3. The output is a fault-class probability distribution or a continuous transformer health index.

C. Decision and Maintenance-Scheduling Layer

The health index and fault classification feed a decision layer that prioritizes inspection and repair actions. For fleet-wide scheduling and resource allocation across many transformers, this layer could optionally formulate the scheduling problem as a QUBO and solve it with a quantum annealer, following the precedent set by quantum-annealing grid-optimization deployments [14]. A continuous monitoring feedback loop returns new sensor data to the data layer as conditions evolve.

## VI. ILLUSTRATIVE COMPARATIVE ANALYSIS

This section compares classical and quantum approaches at a conceptual level and situates published performance figures from the literature side by side. Because no controlled, transformer-specific classical-versus-quantum benchmark currently exists in the public literature, the comparisons below should be read as a synthesis of separate studies rather than a single controlled experiment.

TABLE II IEC 60599 / Duval-Style DGA Fault Codes Referenced in Fig. 1

| Code | Meaning | Typical Diagnostic Indication |
|---|---|---|
| PD | Partial discharge | Low-energy electrical discharges, e.g., in voids or bubbles |
| D1 | Discharge of low energy | Sparking or low-current arcing |
| D2 | Discharge of high energy | Continuous arcing, power follow-through |
| T1 | Thermal fault < 300°C | General overheating, e.g., poor connections |
| T2 | Thermal fault 300–700°C | Localized core or winding overheating |
| T3 | Thermal fault > 700°C | Severe localized overheating / carbonization |

**Figure 3. Schematic Variational Quantum Classifier (VQC) Circuit for DGA Feature Classification**

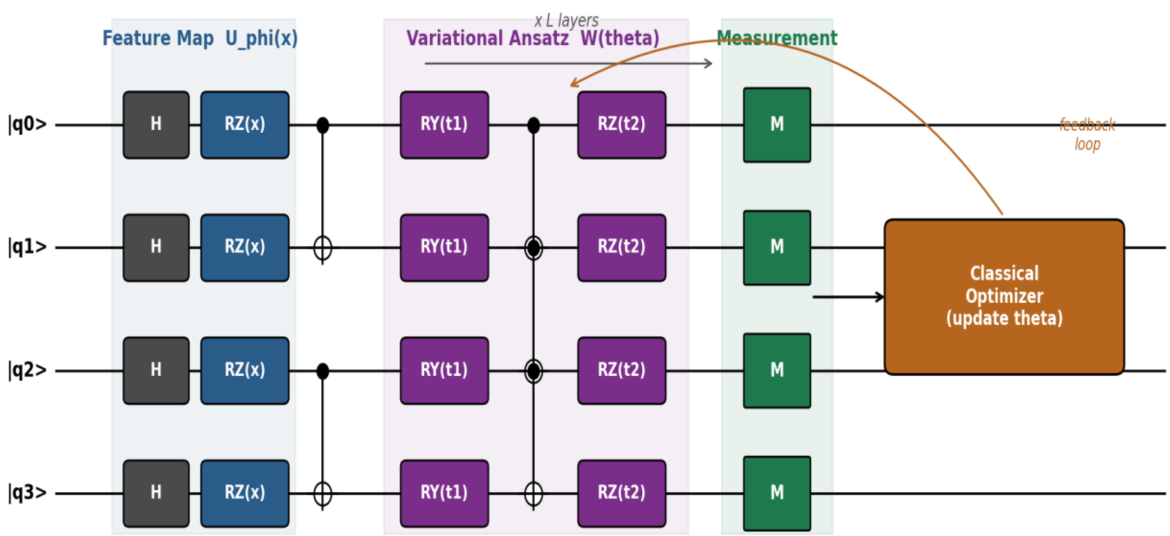


Fig. 3. Schematic variational quantum classifier (VQC) circuit for DGA feature classification, showing a quantum feature map, a repeated variational ansatz, measurement, and the classical optimization feedback loop.

TABLE III Classical Machine Learning Algorithms Commonly Used in Transformer PdM

| Algorithm | Typical Role | Strengths | Limitations |
|---|---|---|---|
| SVM | DGA/fault classification | Strong margin-based separation; reported 95.65% accuracy in a transformer case study | Kernel choice is manual; scales poorly with very high-dimensional features |
| KNN | DGA/fault classification | Simple, interpretable; matched SVM at 95.65% in | Sensitive to feature scaling and class imbalance; |

| Algorithm | Typical Role | Strengths | Limitations |
|---|---|---|---|
| | | the same case study | slow at large scale |
| Decision Tree / Random Forest | Fault classification, ranking | Interpretable rules; ensembles reduce variance | Single trees reported lower accuracy (89.13%) than SVM/KNN |
| Deep Neural Networks / DBN | Multi-modal health index, RUL | Learns complex, multi-sensor patterns | Needs large labeled data; limited interpretability |

TABLE IV Quantum Machine Learning Algorithms Relevant to Transformer PdM

| Algorithm | Mechanism | Potential PdM Role | NISQ Feasibility |
|---|---|---|---|
| QSVM (kernel form) | Quantum feature map + classical SVM solver | Nonlinear DGA fault-class separation | Demonstrated on near-term hardware [13] |
| Variational Quantum Classifier / QNN | Trainable parameterized circuit, hybrid loop | End-to-end fault classification, health index | Feasible but limited by barren plateaus and noise [10] |
| Quantum Annealing / QUBO | Adiabatic optimization on specialized hardware | Maintenance scheduling, crew routing, spare allocation | Commercially available (D-Wave); demonstrated for grid optimization [14] |
| Hybrid QC-Assisted Deep Learning | Quantum sampling within deep generative training | Feature extraction and classification for power-system faults | Demonstrated on IEEE 30-bus and industrial process case studies [3], [4] |

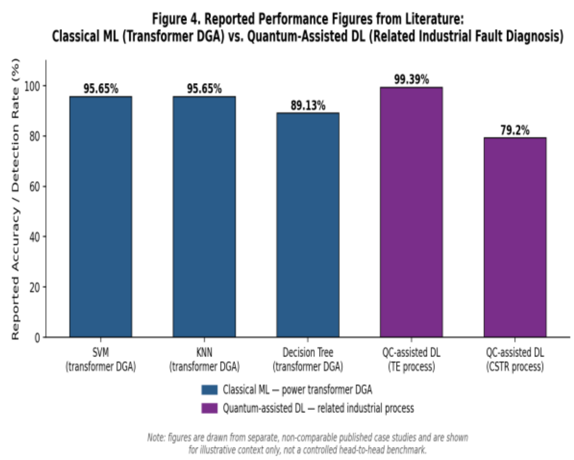


Fig. 4. Reported performance figures from separate published case studies: classical ML on power-transformer DGA data versus quantum-assisted deep learning on related industrial fault-diagnosis case studies.

TABLE V Source Data for Fig. 4

| Ref. | Domain | Method | Reported Metric |
|---|---|---|---|
| [5] | Power transformer DGA | SVM | 95.65% testing accuracy |
| [5] | Power transformer DGA | KNN | 95.65% testing accuracy |
| [5] | Power transformer DGA | Decision Tree | 89.13% testing accuracy |
| [4] | Tennessee Eastman process | QC-assisted deep belief network | 99.39% avg. fault detection rate |
| [4] | CSTR process | QC-assisted deep belief network | 79.2% avg. fault detection rate |

## VII. CHALLENGES AND LIMITATIONS

Several barriers currently limit the practical deployment of quantum machine learning for transformer predictive maintenance. First, NISQ-era devices offer limited qubit counts, shallow achievable circuit depth, and non-negligible gate error rates, constraining the complexity of feature maps and ansatze that can be reliably executed [11]. Second, variational algorithms such as VQCs are prone to "barren plateaus," regions of the training landscape where gradients vanish exponentially with system size, making optimization difficult as circuits scale [10]. Third, loading classical data into a quantum state efficiently, often called the data-loading or input problem, remains a bottleneck that can offset theoretical speed-ups if not carefully engineered. Fourth, no large-scale, transformer-specific, publicly available QML benchmark currently exists; the most relevant demonstrated work addresses broader power-system or industrial-process fault diagnosis rather than transformer DGA specifically [3], [4]. Finally, access to quantum hardware remains costly and queue-limited, meaning most near-term research and the framework proposed in this paper would necessarily rely on classical simulation of small quantum circuits for validation. Taken together, the quantum machine learning community itself acknowledges that a rigorous, general quantum advantage for practical, structured classification tasks of the kind found in transformer DGA data has not yet been established [10].

## VIII. FUTURE RESEARCH DIRECTIONS

Future work should prioritize the creation of labeled, transformer-specific DGA and multi-sensor datasets suitable for QML benchmarking, potentially building on existing repositories such as the IEC TC10 DGA database. Hardware-efficient ansatz designs tailored to the low qubit counts needed for compact gas-ratio feature vectors could reduce circuit depth and mitigate barren-plateau effects. Quantum-enhanced digital twins that combine classical physical simulation with quantum-accelerated combinatorial optimization for fleet-wide maintenance scheduling represent a promising integration point with existing AI-IoT PdM systems [1]. Federated or distributed quantum learning approaches could also allow multiple utilities to collaboratively train fault-classification models without sharing raw, commercially sensitive asset data. Finally, developing explainability tools for quantum and hybrid models will be essential for building the operational trust that grid operators require before adopting any AI-driven diagnostic recommendation in a safety-critical setting.

## IX. CONCLUSION

Predictive maintenance has matured considerably for power transformers, with classical machine learning and AI-IoT systems now well established for interpreting dissolved gas analysis and multi-sensor condition data. Quantum machine learning remains an early-stage but conceptually promising complement to these methods, particularly through quantum-enhanced feature spaces for fault classification and quantum annealing for the combinatorial optimization problems inherent in fleet-wide maintenance scheduling. The hybrid framework proposed in this paper synthesizes these threads into a four-stage pipeline grounded in published precedents from power-system and industrial fault diagnosis. Substantial barriers remain, including limited near-term hardware, unresolved training difficulties such as barren plateaus, and the absence of transformer-specific QML benchmarks, and a rigorous quantum advantage for this class of problem has not yet been demonstrated. The framework and comparative analysis presented here are therefore best understood as a literature-grounded roadmap intended to orient future empirical research rather than as a validated deployment-ready system.

## REFERENCES

[1] M. Nuruzzaman, G. Q. Limon, A. R. Chowdhury, and M. A. M. Khan, "Predictive maintenance in power transformers: A systematic review of AI and IoT applications," SSRN, Jan. 2025. [Online]. Available: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5346270

[2] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, "Quantum machine learning," Nature, vol. 549, no. 7671, pp. 195–202, 2017.

[3] A. Ajagekar and F. You, "Quantum computing

based hybrid deep learning for fault diagnosis in electrical power systems," Appl. Energy, vol. 303, art. no. 117628, 2021.

[4] A. Ajagekar and F. You, "Quantum computing assisted deep learning for fault detection and diagnosis in industrial process systems," Comput. Chem. Eng., vol. 143, art. no. 107119, 2020.

[5] "Predictive maintenance of power transformer using machine learning: A case study," Zenodo, 2026. [Online]. Available: https://doi.org/10.5281/zenodo.19356496

[6] J. Faiz and M. Soleimani, "Dissolved gas analysis evaluation in electric power transformers using conventional methods: A review," IEEE Trans. Dielectr. Electr. Insul., vol. 24, no. 2, pp. 1239–1248, 2017.

[7] M. Duval, "New techniques for dissolved gas-in-oil analysis," IEEE Electr. Insul. Mag., vol. 19, no. 2, pp. 6–15, 2003.

[8] M. Duval and L. Lamarre, "The Duval pentagon: A new complementary tool for the interpretation of dissolved gas analysis in transformers," IEEE Electr. Insul. Mag., vol. 30, no. 6, pp. 9–12, 2014.

[9] International Electrotechnical Commission, Mineral Oil-Impregnated Electrical Equipment in Service: Guide to the Interpretation of Dissolved and Free Gases Analysis, IEC 60599, 2007.

[10] M. Cerezo, G. Verdon, H.-Y. Huang, L. Cincio, and P. J. Coles, "Challenges and opportunities in quantum machine learning," Nat. Comput. Sci., vol. 2, pp. 567–576, 2022.

[11] J. Preskill, "Quantum computing in the NISQ era and beyond," Quantum, vol. 2, p. 79, 2018.

[12] P. Rebentrost, M. Mohseni, and S. Lloyd, "Quantum support vector machine for big data classification," Phys. Rev. Lett., vol. 113, no. 13, art. no. 130503, 2014.

[13] V. Havlíček et al., "Supervised learning with quantum-enhanced feature spaces," Nature, vol. 567, no. 7747, pp. 209–212, 2019.

[14] Quantum Algorithms Institute, "Quantum computing and renewable energy: Optimizing the grid for a sustainable future," Quantum Algorithms Institute. [Online]. Available: https://www.qai.ca/resource-library/quantum-computing-and-renewable-energy-optimizing-the-grid-for-a-sustainable-future