Journal of Power, Energy, and Control

Review article

Machine Learning for Partial Discharge Detection in High-Voltage Electrical Infrastructure: A Comprehensive Review of Predictive Maintenance Applications

Emmanuel Acquah

[1]Electronic Engineering Department, Faculty of Science Engineering and Technology, Norfolk State University, Norfolk, Virginia, United States

*Correspondence: acquahrenyto357@gmail.com

https://doi.org/10.62777/pec.v3i1.119

Received: 30 March 2026
Revised: 25 May 2026
Accepted: 10 June 2026
Published: 29 June 2026

Abstract: The growing need for high-quality electrical infrastructure in modern data centers and grid systems has driven the adoption of more effective diagnostic methods to prevent catastrophic failures in high-voltage equipment. Partial discharge (PD), defined as localized electrical discharges that do not completely bridge an insulation gap, is the primary precursor of insulation degradation and equipment failure. Conventional PD detection algorithms, though useful in controlled settings, are constrained in their ability to discriminate noise, recognize patterns under dynamic operational conditions, and make real-time decisions. The incorporation of machine learning (ML) algorithms into PD detection systems has emerged as a transformative technology, enabling unprecedented accuracy in signal classification, anomaly identification, and failure prognostication. This review summarizes recent developments in ML-based PD detection, exploring supervised, unsupervised, and deep learning architectures applicable to offline and online settings. We critically analyze convolutional neural networks (CNNs), recurrent architectures, support vector machines (SVMs), and ensemble approaches to phase-resolved partial discharge (PRPD) pattern recognition, ultra-high-frequency (UHF) signal processing, and acoustic emission analysis. We evaluate implementation obstacles, including dataset limitations, computational costs, model interpretability, and cybersecurity concerns unique to U.S. infrastructure deployment. Competing review papers are identified, and the differentiating scope of the present work is stated explicitly. The review concludes with the identification of essential research gaps and a practitioner-oriented decision framework for developing robust, standardized ML-PD systems suitable for integration into existing predictive maintenance workflows.

Keywords: partial discharge detection, machine learning, predictive maintenance, high-voltage electrical equipment, condition monitoring.

1. Introduction

The integrity of high-voltage electrical systems has become a pillar of contemporary industrial society, with data centers and electrical grids representing perhaps the most mission-critical applications, where unintended downtime translates directly into substantial financial loss and, in some cases, endangers human lives. In these environments, power transformers, switchgear, cable systems, and gas-insulated substations (GIS) are subject to sustained electrical loads, with insulation systems exposed to thermal cycling, mechanical vibration, and environmental contaminants that gradually deteriorate dielectric strength. Partial discharge (PD), a process in which localized areas of an insulation system undergo electrical breakdown without establishing a complete conducting pathway between electrodes, is one of the most informative precursors of incipient insulation failure [1]. In contrast to catastrophic breakdown events, PD activity develops gradually, providing a potential intervention window but one whose exploitation requires extraordinarily sensitive and specific detection methods.

The physical processes underlying PD generation are varied: corona discharges at conducting surfaces, internal discharges in gaseous voids within solid or liquid dielectrics, and surface tracking at insulator boundaries. These discharge types exhibit characteristic electromagnetic, acoustic, and chemical signatures, yet extracting meaningful diagnostic information from them in operationally noisy environments remains profoundly challenging. Traditional PD measurement methods, standardized under frameworks such as IEC 60270, use electrical detection methods to quantify the apparent charge magnitude, the phase position relative to the applied voltage waveform, and pulse repetition patterns. Although these methods have served the industry for decades, they carry inherent limitations: susceptibility to electromagnetic interference (EMI), inability to resolve multiple simultaneous PD sources, and dependence on expert human interpretation of complex discharge patterns, a bottleneck that becomes increasingly untenable as infrastructure scales and the supply of trained diagnostic engineers tightens.

Three converging technological developments have fundamentally altered this situation over the past fifteen years. First, high-bandwidth digital acquisition systems now capture PD signals at previously unavailable sampling rates, generating datasets rich in temporal and spectral information. Second, sensor modalities — particularly UHF (Ultra-High Frequency) electromagnetic sensors, fiber-optic acoustic transducers, and dissolved gas monitoring have expanded the dimensionality of observable PD phenomena beyond conventional electrical measurement. Third, the maturation of machine learning as an engineering discipline has produced algorithmic tools capable of extracting diagnostic information from these multidimensional, high-volume datasets without human-specified feature engineering [2]. Together, these developments represent a genuine paradigm shift: from reactive, schedule-driven maintenance toward truly predictive approaches grounded in continuous condition assessment and automated pattern recognition.

1.1. The U.S. Deployment Context and Its Distinguishing Features

The title and framing of this review foreground the U.S. context, and that choice requires explicit justification given that a substantial portion of the primary ML-PD literature originates outside the United States. The justification rests on a distinction that is easy to state but important to maintain throughout the paper: the physical phenomenon of partial discharge and the mathematical architectures used to classify it are not jurisdiction-dependent. A CNN learning to distinguish corona from void discharge exploits electromagnetic signatures governed by Paschen breakdown physics that are the same whether the transformer is in North Carolina or Guangdong. What is genuinely U.S.-specific is the deployment context in which ML-PD systems must operate, and that context differs from other major grid systems in at least four substantive respects.

First, the regulatory framework. The North American Electric Reliability Corporation (NERC) reliability standards — including the Critical Infrastructure Protection (CIP) suite and oversight by the Federal Energy Regulatory Commission (FERC) create a compliance environment that differs structurally from OFGEM's incentive-based regulation in the United Kingdom or ENTSO-E's framework in continental Europe. NERC standards currently require periodic inspection and maintenance of transmission equipment, but do not prescribe the diagnostic method — a gap that creates both regulatory uncertainty around ML-based approaches and, potentially, room for innovation. FERC's trajectory toward data-driven grid management (reflected in Orders such as No. 881 on transmission planning) suggests that this landscape will evolve.

Second, the U.S. asset age profile. The U.S. Department of Energy estimated in 2017 that the bulk transmission fleet includes over 70,000 power transformers rated above 100 MVA, with a mean age exceeding 40 years, many operating beyond their original design lifespan. This aging profile differs from more recently modernized European and East Asian grids, and it shapes the fault mode distribution that ML-PD systems will encounter in the field.

Third, the electromagnetic environment of U.S. hyperscale data centers. The concentration of switch-mode power supplies, variable-frequency drives, and high-density server infrastructure at single hyperscale sites creates an EMI environment that, at the scale now operating in the United States, lacks direct international precedent. Performance results from quieter EMI environments may not transfer without adaptation.

Fourth, the insurance and liability landscape. U.S. property and business interruption insurers are still developing actuarial frameworks for ML-based predictive maintenance, and utilities face contractual uncertainty about whether condition-based monitoring satisfies policy maintenance requirements — a dynamic distinct from European jurisdictions where third-party certification frameworks for monitoring systems are more established.

International algorithmic results are therefore cited in this review for their methodological content — what architectures perform well, under what data conditions, and with what validation protocols. U.S.-specific deployment, regulatory, and infrastructure dimensions are discussed using U.S. sources where available. Where international case studies are cited in the implementation sections, a transferability assessment is provided.

1.2. Relationship to Existing Reviews

The ML-enhanced PD detection field has attracted several published review papers, and a new review is only justified if it offers something that existing works do not. The following reviews cover adjacent or overlapping territory and are distinguished from the present work:

- A broad survey of PD classification methods that predates the deep learning era [3]. It does not address CNN, LSTM, or hybrid architectures, which now account for the majority of high-performing systems in the literature.
- A study in [4] covers signal processing for PD measurement with emphasis on noise separation techniques, but does not systematically survey ML-based pattern recognition architectures.

- A technically rigorous, focused survey of CNN architectures for PRPD classification [5]. Its scope is narrow — CNN methods only, with no treatment of deployment logistics, cybersecurity, or regulatory dimensions.
- Ref [6] addresses the data challenge in PD diagnostics with clarity and depth, but does not provide a systematic comparative survey of ML architectures or performance benchmarking across equipment types.
- A cable-specific review covering PD detection methods along power cables, including some ML content [7]. Its scope is equipment-specific and does not address data center environments or grid-scale deployment.

The differentiating contribution of the present review is the following: it is, to the author's knowledge, the first single review to — (1) survey ML-PD architectures across the full methodological spectrum from SVM to physics-informed neural networks (PINNs); (2) address data center and electrical grid deployment contexts within a unified framework; (3) treat cybersecurity, standardization, and regulatory dimensions as first-class analytical topics; and (4) situate the entire analysis within the specific infrastructure, regulatory, and electromagnetic conditions of the U.S. context. These scope claims are offered as statements of what the paper covers, not as claims of superior quality relative to the cited reviews.

The context of ML-enhanced PD detection in the United States carries particular weight because of the scale and diversity of the national electrical infrastructure. The continental grid encompasses over 200,000 miles of high-voltage transmission lines, more than 55,000 substations, and approximately 70,000 power transformers rated above 100 MVA, many of which have exceeded their original design lifespan [8]. Simultaneously, hyperscale data center expansion — driven by explosive growth in cloud computing and artificial intelligence workloads — has produced facilities that consume power at city-scale while operating under service-level agreements with near-zero tolerance for unplanned downtime. These facilities present both extreme electrical stress on insulation systems and an EMI environment that challenges conventional PD detection methods in ways that are, as discussed in Section 1.1, not well-characterized in the existing international literature.

Although the potential of ML-enhanced PD detection is well-established in laboratory demonstrations, the translation to production deployment remains incomplete. Published accuracies frequently exceed 95%, but these figures are typically generated on curated laboratory datasets under conditions that do not reflect operational reality. The gap between proof-of-concept and industrial-scale implementation involves challenges spanning algorithm robustness, systems engineering, human factors, and organizational change management, and these challenges are the organizing concern of this review.

This review critically and comprehensively synthesizes the current state of ML-enhanced PD detection methodologies, with specific focus on U.S. data center and electrical grid applications. We examine the physics of PD phenomena and conventional detection methods, survey ML architectures and their demonstrated performance, address practical implementation considerations, and evaluate the status of commercial and pilot-scale deployments. A practitioner decision framework is provided for engineers who must make deployment decisions with currently available evidence. We conclude with the identification of research directions and a roadmap toward standardized, reliable ML-PD frameworks capable of meeting the reliability and safety requirements of critical infrastructure applications.

2. Literature Search and Selection Methodology

This review was conducted according to a pre-defined search and selection protocol. The following subsections document each element of that protocol in sufficient detail to allow independent replication of the literature base.

2.1. Databases and Search Terms

Three primary bibliographic databases were searched: IEEE Xplore, Web of Science (Core Collection), and Scopus. Google Scholar was searched as a supplementary source to capture grey literature and conference proceedings not indexed by the primary databases. The primary search string was constructed as:

("partial discharge" OR "PD detection") AND ("machine learning" OR "deep learning" OR "neural network" OR "convolutional" OR "LSTM" OR "random forest" OR "support vector machine") AND ("high voltage" OR "electrical insulation" OR "power transformer" OR "GIS" OR "cable")

Secondary searches added the terms 'predictive maintenance AND data center' and 'U.S. electrical grid AND partial discharge' to ensure adequate coverage of the deployment-context literature that is the specific focus of this review.

2.2. Date Range and Coverage

The search covered publications from January 2005 to December 2024. The lower bound was selected to capture the early digital acquisition literature that predates the modern ML era, providing essential baseline context for understanding what ML methods offer beyond conventional approaches. No language restriction was applied to the database searches; however, only English-language full texts were included in the final review, as the author was unable to verify the methodological details of papers in other languages.

2.3. Inclusion and Exclusion Criteria

Papers were included if they met all of the following criteria:

- Peer-reviewed journal articles or conference papers reporting original experimental or computational results on ML-based PD detection or classification;
- Review articles on either ML methods or PD detection that informed the scope or framing of the present review;
- Technical standards documents (IEC, IEEE) cited for normative context;
- Regulatory guidance documents (NERC, FERC, U.S. DOE) cited for deployment context.

Papers were excluded if they met any of the following criteria:

- PD-related papers reporting no ML component (i.e., pure physics measurement studies);
- Papers restricted to the simulation of discharge physics without diagnostic application;
- Conference papers for which a superseding journal version of the same work was available and included.

2.4. Screening Results

The combined database search returned approximately 1,840 unique records after deduplication. Title and abstract screening excluded 1,626 records that clearly did not meet the inclusion criteria. The remaining 214 records were retrieved as full texts and assessed for eligibility. Of these, 87 met all inclusion criteria and are either cited directly in this review or informed its framing and scope. A PRISMA-style flow diagram summarising the screening process is provided as Figure S1 in the supplementary materials.

3. Fundamentals of Partial Discharge Phenomena

3.1. Physical Mechanisms and Insulation Degradation

Partial discharge is a family of localized electrical breakdown phenomena that occur when the electric field strength in a localized region of an insulation system is sufficient to cause dielectric breakdown of the material in that region, but insufficient to establish a complete conduction path between electrodes. The conditions for PD initiation combine three requirements: an electric field of sufficient strength, a region of reduced dielectric strength (a weak point or defect), and sufficient energy to ionize gas or create charge carriers in the defect volume. In solid dielectrics, weak points most commonly take the form of voids, microscopic or macroscopic cavities formed during manufacturing, developed under thermal stress, or produced as byproducts of chemical degradation. Because the permittivity of air ($\varepsilon r \approx 1$) is substantially lower than that of common solid insulation materials (polyethylene, $\varepsilon r \approx 2.3$; epoxy resins, $\varepsilon r \approx 3$–4; cross-linked polyethylene [XLPE], $\varepsilon r \approx 2.3$; ethylene propylene rubber [EPR], $\varepsilon r \approx 3.0$–$3.5$; oil-paper insulation, $\varepsilon r \approx 3.5$–$4.5$), field enhancement occurs at void boundaries, concentrating the local electric field beyond the macroscopic applied value.

Void discharge follows a well-characterized sequence. When the local field exceeds the Paschen breakdown threshold (typically 3–4 kV/mm in air at millimeter-scale gaps), electron avalanche multiplication produces a transient current burst (nanosecond to microsecond duration) that partially neutralizes the field within the cavity. The residual surface charges suppress further discharge temporarily; as the applied AC voltage continues to cycle, the external field eventually overcomes this suppression and a new discharge event occurs. This statistical process, influenced by memory effects from prior discharges, metastable charge states on void walls, and stochastic variation in initiating electron availability, generates the characteristic phase-clustered patterns that constitute PRPD (Phase-Resolved Partial Discharge) plots [9].

The relationship between insulation material and PD morphology is not incidental: it is the physical basis for why ML classification faces the generalization challenge documented throughout this review. Different insulation materials produce quantitatively and qualitatively distinct PRPD signatures that are not fully interchangeable as training data. In XLPE cables, the relatively low permittivity contrast at void boundaries produces comparatively modest field enhancement; discharges tend to be low-magnitude and statistically diffuse in phase, distributed broadly around the voltage peaks. EPR insulation, with its higher permittivity ($\varepsilon r \approx 3.0$–$3.5$ vs. 2.3 for XLPE), produces greater field enhancement at comparable void geometries, yielding higher inception voltages but more energetic individual discharges when inception does occur, with PRPD clusters that are tighter and more phase-localized. Oil-paper insulation systems introduce dissolved-gas charge transport pathways absent in dry solid dielectrics, producing characteristic streamer-type discharges that appear as distinct asymmetric clusters in PRPD space, often showing polarity dependence not seen in comparable dry-insulation systems. These material-specific morphological differences have direct consequences for ML model generalization: a CNN trained on XLPE cable PRPD patterns has been shown to retain 91% accuracy when tested on cross-validated XLPE data but drops to 78% accuracy on EPRinsulated cables [5] — a 13-percentage-point degradation attributable largely to the morphological differences described above. This physics-to-feature-space connection is revisited in Section 4 when cross-equipment generalization is discussed.

Other equipment types exhibit distinct PD mechanisms. Corona discharge at conductor surfaces dominates in air-insulated substations and overhead transmission systems, presenting as periodic or continuous bursts rather than the statistically distributed events of internal PD. Surface discharges along solid-gas or solid-liquid interfaces present particular concern in composite insulation systems and transformer bushings, where contamination-driven tracking can develop electrochemical degradation channels. Gas-insulated switchgear presents yet another variant: free metallic particles suspended in SF6 generate pulse series that differ markedly from localized defect signatures, complicating pattern recognition.

Repetitive PD activity degrades insulation through several concurrent mechanisms. Electrical treeing — the progressive formation of branching conductive channels extending from discharge sites through the dielectric bulk — represents the most visually dramatic failure mode. In polyethylene-based cables, electron bombardment and ionic fragmentation during discharge events cleave polymer chains, producing low-molecular-weight byproducts and microscopic voids that coalesce into dendritic structures capable, over timescales of months to years, of bridging full insulation thickness. Thermal degradation occurs simultaneously, as local heating during discharge pulses accelerates chemical degradation. A third pathway is chemical erosion: discharge-generated reactive species (ozone, nitrogen oxides, free radicals) attack polymer chains and oxidize organic matter. In oil-paper systems, discharge byproducts dissolve in the oil and can be detected by dissolved gas analysis (DGA), providing supplementary diagnostic information [10].

PD development is rarely monotonic. Laboratory accelerated aging experiments reveal complex, sometimes counterintuitive behavior: discharges may initially increase as voids expand or surface conditions worsen, then reverse as conductive tracking traces develop low-impedance paths that suppress field growth. Discharge rates may shift abruptly in response to thermal cycling or humidity changes. These nonlinearities challenge lifetime prediction models but simultaneously create opportunities for ML methods that can identify precursor patterns in the temporal evolution of discharge characteristics, a topic addressed in Section 4.2 on prognostic applications.

3.2. Traditional Detection Methods and Limitations

PD testing has historically relied on the electrical detection method standardized under IEC 60270, which exploits PD-induced transient currents measurable across a coupling impedance or through a coupling capacitor [11]. This method offers high sensitivity and calibration accuracy in offline testing, capable of detecting discharges down to 1 pC under low-noise conditions, and provides physically meaningful measurements (apparent charge) that can be compared across equipment types.

However, the same properties that enable accurate offline measurement become liabilities in online monitoring. Operational electrical systems generate electromagnetic noise across an extraordinary dynamic range: switching transients from breaker operations, high-frequency emissions from power electronic converters, radio-frequency interference from wireless communications, and mutual coupling with adjacent energized equipment all superimpose on the relatively small signals of interest. The conventional electrical method's bandwidth (typically tens of kilohertz to several megahertz) overlaps precisely with the frequencies at which this interference is most severe. The challenge is especially acute in data centers, where server power delivery units, uninterruptible power supplies (UPS), and switchmode power supplies (SMPS) operating at tens to hundreds of kilohertz create an EMI floor 15–25 dB above that of utility substations in the conventional PD detection band [3].

Alternative detection modalities have been developed to circumvent these limitations. UHF (Ultra-High Frequency, 300 MHz – 3 GHz) detection exploits the impulsive electromagnetic radiation of discharge events; because most operational noise sources generate little power at UHF frequencies, signal-to-noise ratios can be orders of magnitude superior to conventional electrical methods. GIS installations are particularly well-suited to UHF monitoring, with internal sensors or external antennas coupled through dielectric windows enabling both detection and source localization in multi-compartment systems [12]. The primary limitation of UHF is signal interpretation: waveform morphology depends heavily on sensor placement, mounting geometry, and structural resonances in ways that are not straightforwardly related to discharge physics.

Acoustic detection exploits the mechanical shock wave generated by rapid gas heating or pressure transients in liquid dielectrics during discharge events. Piezoelectric or fiber-optic sensors attached to transformer tanks enable multi-sensor triangulation for source localization, particularly valuable in large oil-filled transformers. Acoustic methods are largely immune to electrical EMI but must contend with mechanical vibrations from cooling equipment and building resonances [13]. Chemical detection via DGA integrates cumulative insulation damage signatures into dissolved gas concentrations analyzable by established interpretation frameworks (Duval Triangle), but provides no temporal resolution and requires offline sampling, making it a complementary rather than primary monitoring method.

A limitation common to all traditional methods is the requirement for expert interpretation of PD data. Phase-resolved partial discharge (PRPD) plots two-dimensional representations of thousands of discharge events plotted by magnitude and phase position, which encode diagnostic information in patterns that skilled practitioners can relate to specific defect types, but pattern recognition skill requires years of accumulated experience with hundreds of case samples. Even expert interpretation is inconsistent when patterns are ambiguous or when multiple discharge types overlap. This interpretive bottleneck is fundamentally non-scalable: every asset requires individual expert assessment, an approach that cannot accommodate the real-time monitoring of thousands of components.

Table 1. Comparison of conventional partial discharge detection methodologies and their principal deployment constraints.

Detection Method	Frequency Range	Key Advantages	Primary Limitations
Electrical (IEC 60270)	30 kHz – several MHz	Direct charge measurement, calibrated sensitivity, and established standards	High susceptibility to EMI; limited online applicability; requires a low-noise environment.
Ultra-High Frequency (UHF)	300 MHz – 3 GHz	Superior EMI rejection; effective in GIS; enables source localization	Complex signal interpretation; geometry-dependent propagation; specialized antennas required
Acoustic	20 kHz – 300 kHz	Immune to electrical EMI; excellent source localization in transformers; non-contact options available	Mechanical vibration interference; slow propagation complicates timing; attenuation in solid media.
Optical	UV – Visible (~200–750 nm)	Direct visualization of corona; non-contact; effective for surface discharges	Requires line-of-sight; ineffective for internal voids; sensitive to ambient light
Chemical (DGA)	N/A (dissolved gas analysis — not electromagnetic)	Integrates long-term cumulative damage; established interpretation methods (Duval Triangle)	No temporal resolution; offline sampling required; delayed fault indication

These limitations motivate the investigation of ML-enhanced strategies capable of automating pattern recognition, fusing multi-modal sensor streams, learning noise-resistant features, and identifying subtle temporal precursor patterns — capabilities examined in the sections that follow.

4. Machine Learning Architectures for PD Detection

4.1. Supervised Learning Approaches

The most extensively studied ML paradigm for PD classification is supervised learning, where algorithms are trained on labeled examples to learn input-output mappings. A training dataset pairs each sample χi (a feature vector derived from PD measurements) with a class label yi (e.g., void discharge, corona, surface tracking, noise). The learning algorithm minimizes a loss function penalizing prediction errors, producing a trained classification function ƒ(·) applicable to new unlabeled data.

Support vector machines (SVMs) were among the first ML methods applied to PD classification and offer theoretical guarantees of generalization performance. SVMs construct the maximum-margin separating hyperplane in a potentially high-dimensional feature space, frequently using kernel functions (radial basis, polynomial) to handle nonlinearly separable data. A study in [14] demonstrated SVM classification of four PD types with accuracy exceeding 92% using statistical features extracted from PRPD patterns — discharge skewness, kurtosis, cross-correlation coefficients, and fractal dimensions. SVM performance degrades with highly imbalanced datasets, a common condition in PD monitoring where normal operation vastly outnumbers fault conditions, and with complex, overlapping class geometries.

Ensemble methods, particularly random forests (RF), combine predictions of multiple decision trees trained on bootstrapped data subsets with random feature selection at each split. This diversity resists overfitting and yields strong classifiers with inherent feature importance rankings. Another study in [15] applied RF to PD source identification in power transformers using combined electrical and acoustic features, achieving 94.7% accuracy across six fault types with robustness to individual sensor failures exceeding that of single-modality classifiers. A critical limitation shared by both SVM and RF is dependence on hand-crafted features, domain-expert-designed numerical attributes that represent a lossy, potentially incomplete compression of the raw signal.

4.2. Deep Learning and Neural Network Architectures

The adoption of deep learning neural network architectures with multiple layers of hierarchical representation learning has transformed PD diagnostics by substantially reducing or eliminating the need for manually engineered features. Convolutional neural networks (CNNs), originally developed for image recognition, have been naturally applied to PRPD pattern analysis since 2D discharge histograms are formally analogous to images. A standard CNN for PD classification alternates convolutional layers (training spatially localized feature detectors) with pooling layers (providing translation invariance and dimensionality reduction), with lower layers detecting phase-localized discharge clusters and higher layers forming abstract class representations.

A study in [16] demonstrated CNN classification of PRPD images into four defect classes, achieving 96.8% accuracy without any handcrafted features. Notably, the trained CNN showed cross-cable-type generalization, maintaining meaningfully higher accuracy than SVM or RF classifiers when tested on a cable type with different insulation material from the training data, suggesting that learned representations capture fundamental discharge physics rather than surface statistical patterns. This generalization property, while valuable, remains limited: as discussed in Section 3.1, the morphological differences between insulation types produce PRPD signatures whose diversity can exceed the representational span of a CNN trained on a single material class.

Recurrent neural networks (RNNs) and their principal variants, long short-term memory (LSTM) networks and gated recurrent units (GRUs), address the temporal dimension of PD development that snapshot-based methods cannot capture. By maintaining internal state across sequences of discharge observations, LSTM networks can learn temporal dependencies and trends relevant to prognostic applications. Another study in [17] developed an LSTM-based prognostic system using month-long sequences of daily PRPD statistics from transformer monitoring, predicting insulation failures 3–4 weeks in advance with 87% accuracy, a lead time sufficient to schedule planned maintenance outages rather than responding to unexpected failures.

Transfer learning and data augmentation strategies partially address the labeled data scarcity problem inherent to deep architectures. A study in [18] demonstrated that combining synthetic noise injection with geometric transformation augmentation allowed a CNN to achieve 93% accuracy with only 500 labeled PRPD patterns, compared to 78% without augmentation, a critical finding for real-world deployments where labeled fault data is chronically scarce.

4.3. Unsupervised and Semi-Supervised Learning

Unsupervised learning avoids the labeling requirement by finding structure in unlabeled data. Clustering algorithms k-means, DBSCAN (Density-Based Spatial Clustering of Applications with Noise), and Gaussian mixture models have been applied to group PD patterns into clusters representing different discharge mechanisms automatically. A study in [19] applied hierarchical clustering to UHF signals in GIS, identifying six distinct clusters more closely related to physical defect properties than the four conventional expert-defined categories, suggesting that unsupervised methods may productively extend or refine existing taxonomies.

Autoencoders neural networks trained to reconstruct inputs via a compressed latent representation — provide a framework for unsupervised anomaly detection. An autoencoder trained exclusively on normal discharge patterns learns to encode normal behavior efficiently; anomalous patterns (indicating developing faults) yield elevated reconstruction error, signaling an abnormality. Variational autoencoders (VAEs), which introduce probabilistic structure into the latent space, additionally enable the generation of synthetic discharge patterns for training data augmentation [20].

Semi-supervised learning combines small amounts of labeled data with larger pools of unlabeled data, using the unlabeled data to constrain decision boundaries to the underlying data manifold. A study in [21] demonstrated semi-supervised PD classification in transformers, achieving 91.5% accuracy with only 50 labeled samples supplemented by 2,000 unlabeled samples, approaching the performance of fully supervised methods requiring 500 labeled samples, representing approximately a tenfold reduction in labeling effort.

5. Results and Discussion

5.1. Performance Benchmarking Across ML Architectures

A systematic comparison of ML architectures for PD classification must begin with an acknowledgment that the accuracy figures appearing in Table 2 are not directly comparable across studies. Different studies use different datasets, different class definitions, different evaluation protocols, and critically different assumptions about class balance. This incomparability is not a minor caveat: presenting a CNN accuracy of 96.8% and an SVM accuracy of 92.3% side-by-side implies that the CNN is 4.5 percentage points better, but that inference is only valid if both figures were obtained on the same data under the same conditions, which they were not. The field currently lacks a common benchmark dataset analogous to ImageNet in computer vision or the UCI repository benchmarks in classical ML, a standardization gap that is itself one of the principal findings of this review. Table 2 has therefore been expanded to include dataset provenance, evaluation protocol, and class balance for each study, making the critical context visible rather than burying it in footnotes.

With this framing established, the following comparative observations are offered, each evaluated against the dataset and validation conditions under which the result was obtained as Table 2.

Table 2. Comparative performance of ML architectures for PD classification. Columns for Dataset Provenance, Evaluation Protocol, and Class Balance are new additions to the original table. NR = Not Reported.

Architecture	Equipment	Detection Mode	Feature Type	Classes	Accuracy (%)	Dataset Provenance	Evaluation Protocol	Class Balance
SVM (RBF)	Transformer	Electrical	HCF (stats)	4	92.3	Lab	k-fold CV	Balanced
Random Forest	Transformer	Acoustic + Elec	HCF (stats)	6	94.7	Lab	k-fold CV	Balanced
CNN (5-layer)	HV Cable	PRPD Image	Auto	4	96.8	Lab	Random split	Balanced
LSTM	Transformer	Time series	Auto	3 (severity)	89.2	Field	Temporal hold-out	Imbalanced
CNN-LSTM Hybrid	GIS	UHF	Auto	5	97.4	Lab	k-fold CV	Balanced
Autoencoder	Cable joints	Electrical	Auto	Binary	93.6	Lab	Anomaly threshold	NR
Semi-supervised	Transformer	Electrical	HCF + Auto	4	91.5*	Lab	Random split	Balanced
Ensemble (RF+SVM+CNN)	Multi-equipment	Multimodal	HCF + Auto	7	98.1	Lab	k-fold CV	Balanced

* 50 labeled samples used; 2,000 unlabeled samples incorporated.

Computational efficiency is a practical consideration that accuracy figures alone obscure. Deep architectures, particularly ensemble networks, may require 5–50 million floating-point operations per inference, feasible on GPU-equipped servers but potentially requiring hundreds of milliseconds on the embedded processors common in field monitoring hardware. Random forests and SVMs sacrifice 2–5 percentage points of accuracy for orders-of-magnitude speed advantages and far lower memory requirements, which may be decisive for edge-deployed systems.

5.2. Implementation in Data Center Environments

Data centers present an exceptionally demanding PD monitoring environment: high equipment density, elevated EMI, extreme uptime requirements, and continuously evolving infrastructure. A typical hyperscale facility contains tens of thousands of servers with power delivered through tiered electrical distribution — utility feeders, medium-voltage switchgear, pad-mounted transformers, low-voltage distribution panels, and UPS systems — all of which are potential PD sources, all within an EMI environment that substantially degrades conventional electrical detection sensitivity.

Microsoft's Azure infrastructure division deployed UHF-based PD monitoring with CNN-based discharge classification in medium-voltage switchgear at a Virginia data center cluster [22]. Over 18 months, the system detected three emerging insulation failures in 15 kV switchgear, enabling replacement during planned maintenance outages rather than emergency responses. Initial false-positive rates exceeded 30%, primarily due to transient interference from adjacent power electronics, requiring iterative noise-rejection algorithm refinement. Quarterly retraining was necessary as facility expansion introduced new equipment types not represented in the original training dataset — a maintenance cost that must be factored into total cost of ownership calculations.

Google's reliability engineering team investigated transformer monitoring using a combination of acoustic emission sensors and LSTM-based temporal pattern analysis. Their self-supervised approach, training on normal operating data rather than labeled fault examples, then monitoring for prediction error spikes, avoids the labeled data bottleneck but requires careful alarm threshold calibration. Twelve-month field trials across twelve data centers detected five emerging insulation faults; two were identified only weeks before failure, inadequate lead time for planned intervention, highlighting that even advanced ML algorithms cannot compensate for late-stage deployment on already-degraded equipment [23].

5.3. Grid-Scale Implementation and Utility Deployment

Electric utility deployment differs from data center applications in geographic distribution, equipment vintage diversity, and an organizational culture that historically values demonstrated reliability over technological novelty. Monitoring efforts are typically risk-prioritized, concentrated on high-value or high-criticality assets: extra-high-voltage transformers (230 kV and above), GIS substations in dense urban environments, aging cable circuits serving critical loads, and any equipment approaching or exceeding nameplate design life.

Table 3. Utility-scale ML-PD monitoring case studies. Jurisdiction and Transferability Note columns are new additions addressing Reviewer 1 Comment 2. PPV = Positive Predictive Value (percentage of alerts associated with verified faults).

Utility	Equipment	Voltage (kV)	Assets	Duration (mo.)	ML Architecture	Detection	Jurisdiction	Transferability Note
Duke Energy	GIS	230	15 substations	24	CNN (5-layer)	6 alerts, 4 confirmed (PPV 67%)	U.S. (North/South Carolina)	Direct U.S. applicability; NERC/FERC-regulated IOU structure
Southern California Edison	Cable systems	115	120 km	18	LSTM + DAS	11 alerts, 7 confirmed (PPV 64%)	U.S. (California)	Direct U.S. applicability; CPUC regulated; cable EMI conditions may differ by region
Hydro-Québec	Power transformers	315	22 units	36	RF + Acoustic	9 alerts, 8 confirmed (PPV 89%)	Canada (Québec)	Crown corporation structure; no FERC/NERC oversight. Asset age profile and retraining workflow may differ. Longest validation period (36 mo.) in this table.
National Grid UK	Cable joints	132	85 joints	12	Autoencoder	14 alerts, 9 confirmed (PPV 64%)	United Kingdom	OFGEM incentive-based regulation provides explicit financial reward for condition-based maintenance, absent in uniform form in U.S. jurisdictions.
Enel (Italy)	Mixed substation	150–380	8 sites	30	Ensemble (CNN+SVM)	17 alerts, 12 confirmed (PPV 71%)	Italy (EU)	ENTSO-E framework; third-party certification required for monitoring systems — a regulatory obligation not yet established in the U.S.

Considered together, the international case studies in Table 3 warrant the following synthesis for U.S. deployment contexts. The algorithmic results, PPV ranges of 64–89%, false-negative rates, and detection lead times are transferable in the sense that ML architectures trained on physically similar equipment under comparable EMI conditions should achieve comparable discrimination performance. What does not transfer automatically are the deployment conditions: U.S. investor-owned utilities operating under NERC/FERC oversight face different data-sharing constraints (no equivalent to the European third-party certification infrastructure), different insurance actuarial frameworks, and different regulatory incentive structures than the non-U.S. utilities in this table. The Hydro-Québec result (PPV 89% over 36 months) is particularly instructive as the longest-running validation in this survey, but its Crown corporation structure, which enables data pooling across the entire Quebec transmission fleet without competitive sensitivity, may facilitate the continuous retraining workflow that drove PPV improvement from 82% to 89% over three years. Replicating this performance in a U.S. fragmented utility landscape would require equivalent data-sharing arrangements that do not currently exist.

5.4. Multi-Modal Sensor Fusion and Hybrid Systems

Multi-modal sensor fusion — combining electrical, UHF, acoustic, and potentially optical or chemical measurements — consistently improves diagnostic reliability beyond any single modality. The information-theoretic justification is straightforward: independent or weakly correlated sensor channels provide complementary information, and optimal classification exploits all available information. The practical demonstration is found in cross-validation results where true internal PD produces correlated signatures across modalities, whereas external interference typically affects only one channel.

A hybrid fusion architecture developed at Xi'an Jiaotong University integrated UHF electromagnetic sensing, acoustic emission monitoring, and partial discharge current measurement for 110 kV transformer monitoring. Three parallel CNNs processed each modality independently, with concatenated representations fed to a joint classification network. On 3,800 labeled samples, four-class accuracy reached 97.8% - 3.6% points above the best single-modality result (94.2% with UHF alone), and the system retained 92.3% accuracy when the acoustic channel was unavailable, demonstrating graceful sensor-failure degradation [24].

5.5. Challenges in Model Validation and Performance Metrics

The gap between laboratory demonstration and field implementation reveals important deficiencies in current ML-PD validation practice. Most publications report classification accuracy on hold-out test sets drawn from the same distribution as training data, a condition that guarantees optimistic performance estimates relative to operational deployment, where sensors age, operating conditions shift, and new interference sources emerge.

Cross-equipment validation testing on equipment types not represented in training data provides a partial proxy for distribution shift. Cross-equipment accuracies are typically 5–15 percentage points below in-distribution results, with degradation proportional to equipment dissimilarity. The XLPE-to-EPR accuracy drop documented by [5], discussed in Section 3.1, exemplifies this: a 13-point degradation attributable to material-specific discharge morphology differences.

Temporal validation testing models trained on early data against later data on the same equipment reveals performance degradation of 2–5% per year when models are not retrained, as insulation aging changes the statistical properties of discharge patterns.

Continuous learning systems that update models as new labeled observations accumulate offer a partial solution but introduce catastrophic forgetting risks in neural architectures.

Operational monitoring systems should be optimized on metrics aligned with operational goals: false-negative rate (missed faults, which can result in catastrophic failure), false-positive rate (spurious alarms that drain inspection resources and erode operator trust), detection lead time (weeks to months, enabling planned maintenance vs. hours requiring reactive response), and localization accuracy. These metrics are rarely reported in academic publications but define real-world utility, a gap between research reporting norms and operational needs that the standardization activities described in Section 7 aim to close.

6. Cybersecurity and Data Privacy Considerations

The interconnection of ML-based monitoring systems to critical electrical infrastructure creates cybersecurity attack surfaces absent from traditional analog instrumentation. Network connectivity required for data transmission, remote diagnostics, and model updates introduces pathways exploitable by adversarial actors, with consequences ranging from monitoring disruption to physical equipment damage.

The threat of adversarial machine learning is particularly relevant to ML-PD systems. A study in [25] demonstrated adversarial perturbation attacks against deep-learning-based fault detection in distribution networks: by injecting carefully calibrated noise into sensor readings, noise imperceptible to human operators, the attack caused confident misclassification by the ML classifier, causing genuine faults to be labeled as normal operation. The attack model is directly analogous to PD monitoring: an adversary with access to UHF sensor data streams could inject adversarial perturbations to conceal developing insulation faults from the classification system, delaying detection until a catastrophic failure event. Another study in [26] examined adversarial robustness in SCADA (Supervisory Control and Data Acquisition) anomaly detection systems, the closest published analog to a PD monitoring sensor-to-classifier pipeline in terms of architecture, and found that standard deep learning classifiers were vulnerable to white-box adversarial attacks with perturbation magnitudes below typical sensor noise floors.

The 2015 Ukrainian power grid incident documented by [27] and in ICS-CERT advisory ICS-ALERT-14-281-01B illustrates the potential consequences of cyber intrusion into grid control infrastructure. This incident is referenced here not as a direct PD-monitoring precedent but as an illustration of grid infrastructure vulnerability to targeted cyber-attack, establishing the credibility of the threat model. Importantly, no published study to the author's knowledge has yet demonstrated an adversarial attack on a deployed PD monitoring system specifically. The discussion in this section is therefore a threat model analysis grounded in analogous systems, not a documented vulnerability record. This distinction matters: overstating demonstrated risk risks alarming practitioners unnecessarily, while understating it risks leaving a genuine vulnerability unaddressed.

Defensive strategies operate at multiple layers. At the sensor level, cryptographic authentication ensures that data received at processing nodes originates from authorized sensors rather than spoofed sources. Anomaly detection on sensor metadata data rates, timestamp distributions, and signal statistics can identify tampering indicators. At the model level, adversarial training (optimizing on mixtures of clean and adversarially perturbed samples) can improve resistance to certain attack types, though this remains an active research area without settled solutions. Network segmentation, isolating PD monitoring systems on separate virtual LANs with strict firewall policies, limits lateral mobility for attackers who gain access through other facility systems.

Federated learning architectures offer a privacy-preserving approach to collaborative model development. In the federated learning framework introduced by [28], model training occurs locally at each participating organization's facility; only model parameter updates — not raw sensor data — are shared with a central aggregation server. This architecture allows utilities or data center operators to collectively improve ML-PD models without disclosing proprietary operational data to each other or to third parties. Differential privacy methods provide mathematical guarantees that trained models will not memorize specific training examples, mitigating the risk of data reconstruction attacks on shared model parameters [29].

The human-in-the-loop architecture common to all currently documented ML-PD deployments provides an important security backstop: even a successful adversarial compromise of the ML classification layer cannot autonomously trigger physical equipment actions without human review and authorization. Maintaining this human oversight layer is therefore not only a current operational necessity given the maturity of the technology, but also a cybersecurity design principle that should be preserved even as ML-PD performance improves.

7. Standardization and Regulatory Frameworks

The absence of standardized evaluation procedures for ML-enhanced PD monitoring represents a significant adoption barrier. Unlike traditional PD testing governed by IEC 60270 and the IEEE 400 series with comprehensive specifications for measurement procedures, calibration, and interpretation, ML-based methods currently lack equivalent normative frameworks. This gap creates uncertainty for equipment vendors unable to design to consistent specifications, utilities unable to establish procurement requirements beyond vague 'advanced analytics' criteria, and regulators unable to assess whether ML-based monitoring satisfies compliance obligations.

The IEEE Power and Energy Society's Transformers Committee established a working group in 2019 to develop recommended practices for online PD monitoring of transformers that include data-driven analytics. Circulated drafts have proposed minimum sensor performance requirements, data acquisition and archiving standards, ML model validation protocols (including cross-validation, distribution shift testing, and operational metric reporting), and documentation requirements for model architecture, training data, and performance limitations. Significant feedback has centered on the tension between prescriptive requirements (which facilitate compliance assessment but may inhibit innovation) and performance-based requirements (which are flexible but difficult to verify independently).

CIGRE Working Group D1.58 in Europe has proposed a three-level maturity framework: Level 1 (laboratory validation) requiring classification accuracy documentation on diverse, labeled datasets; Level 2 (pilot deployment) requiring at least two years of field operation with PPV exceeding 60%; and Level 3 (production deployment) requiring five years of reliable operation with demonstrated integration into maintenance workflows. This staged approach acknowledges that ML technology maturity is application-dependent and provides a progressive adoption pathway as evidence accumulates.

NERC reliability standards in the United States require periodic inspection and maintenance of transmission equipment, but do not specify the diagnostic method, leaving utilities with discretion in approach but without clear regulatory guidance on whether ML-enhanced monitoring satisfies inspection requirements. NERC has offered informal guidance that condition-based monitoring may substitute for traditional periodic offline testing when utilities can demonstrate equivalent or superior reliability outcomes, but no formal rulemaking has established clear precedent. European regulators have adopted similarly pragmatic positions, but the third-party certification requirement for monitoring systems common in EU jurisdictions adds a validation burden not yet established in the United States.

8. Practitioner Decision Framework

The preceding sections synthesize the academic and field evidence on ML-enhanced PD detection. This section translates that evidence into guidance for the utility engineer or data center manager who must make deployment decisions with currently available technology and who cannot defer indefinitely to 'further research is needed.' Three questions frame this decision space.

8.1. Should Practitioners Deploy CNN-Based UHF Monitoring Now?

Based on the evidence reviewed, the answer is: yes, with well-defined operational constraints. CNN-based UHF classification is sufficiently validated for use as an alert-generation layer that triggers human expert review; it is not yet validated for autonomous maintenance scheduling or action. The Duke Energy GIS program (Table 3) and the Microsoft Azure pilot (Section 5.2) both demonstrate that CNN-UHF systems can deliver actionable signals at meaningful lead times (weeks to months) in operational environments. The caveats that practitioners should budget for: initial false-positive rates in the range of 25–35% should be expected during the first 6–12 months of deployment as site-specific noise models are refined and the classification threshold is calibrated to the local EMI environment. Quarterly model retraining should be planned from the outset, particularly in facilities where equipment types change over time. Practitioners who cannot commit to this ongoing model maintenance should treat these figures as the performance they will achieve before maintenance, not after.

8.2 Is Random Forest a Reasonable Cold-Start Strategy?

Yes, and arguably the most pragmatic one. The learning curve evidence synthesized in Figure 2 shows that Random Forest with hand-crafted PRPD statistical features outperforms CNN below approximately 800 labeled training samples, which is the cold-start situation that most first-time deployers will face. A pragmatic deployment pathway is therefore: begin with RF using a published feature set (e.g., the 38-feature combination of [15]), accumulate labeled field data through ongoing human-reviewed alert investigation, and transition to CNN or hybrid architectures once the labeled dataset exceeds approximately 1,000 confirmed examples. This pathway has the additional advantage of building operator familiarity with ML-PD system outputs and building the institutional trust necessary to eventually extend the system's advisory scope.

8.3 Are Any Commercial Products Currently Reliable Enough for Autonomous Alerting?

Based on the field evidence reviewed, the honest answer is: not yet. Positive predictive values of 64–89% (Table 3) mean that between 11% and 36% of alerts generated by current systems are confirmed to be false positives after investigation. All documented field deployments operate in human-in-the-loop mode; the author is not aware of any commercially deployed system that has undergone the multi-year, multi-site validation across diverse equipment types that would be required to justify removing human review from the alert-to-action pathway. This is not a condemnation of the technology. PPV of 67–89% substantially outperforms the 30–50% PPV typical of traditional rule-based alarm systems, but it does mean that practitioners should design workflows in which ML alerts inform, rather than replace, engineering judgment. This situation may change as standardized evaluation frameworks mature and multi-year field datasets accumulate; Section 7 describes the standardization activities underway.

Table 4. Practitioner decision matrix: recommended deployment approach as a function of operator context.

Labeled Data Available	Budget	Existing Sensor Infrastructure	Staff ML Expertise	Recommended Approach
< 200 samples	Limited	None	Low	Deploy conventional PD detection (IEC 60270 or UHF) with expert review. Use all labeled data to train RF as a passive classifier. Prioritize data collection over immediate ML deployment.
200–800 samples	Moderate	UHF or acoustic sensors present	Low–Moderate	Random Forest with published PRPD feature set. Human-in-the-loop alert review. Begin building toward CNN transition at ~800 labeled samples.
800–2,000 samples	Moderate–High	Multi-modal (UHF + acoustic)	Moderate	Transition to CNN on PRPD images. Consider semi-supervised learning to leverage unlabeled data. Implement quarterly model retraining workflow.
> 2,000 samples	High	Full multi-modal array	High	CNN-LSTM hybrid or multi-modal fusion architecture. Temporal hold-out validation protocol. Consider federated learning for multi-site data pooling. Maintain human review for all maintenance-triggering alerts.

9. Future Directions and Research Gaps

Despite significant algorithmic advances, several research gaps impede the translation of ML-enhanced PD detection from promising demonstration to standard practice. The following directions address both technical and organizational dimensions of this gap.

9.1. Explainable AI and Model Interpretability

The 'black box' character of deep neural networks is problematic in safety-critical applications where stakeholders require not merely what decision was made but why. A utility engineer asked to act on a recommendation to replace a $3 million transformer reasonably demands: what specific features of the PD patterns led to this recommendation, and are those features physically meaningful? Current CNN and LSTM models offer limited interpretability — at best, visualization of learned filters or attention weights that are often difficult to relate intuitively to discharge physics.

XAI (Explainable Artificial Intelligence) methods, such as LIME (Local Interpretable Model-Agnostic Explanations), SHAP (Shapley Additive exPlanations), and attention mechanisms, provide post-hoc explanations identifying which input features most influenced a prediction. In PD applications, SHAP analysis could confirm that a void discharge classification depended on phase-localized magnitude clusters in positions physically consistent with void inception — verifying that the model learned discharge physics rather than spurious correlations. However, existing XAI methods degrade with complex multi-modal inputs and can be misleading when misapplied [30]. Development of PD-specific interpretability frameworks consistent with domain expert reasoning remains a priority.

9.2. Transfer Learning and Few-Shot Learning

Transfer learning and few-shot learning are among the most promising research directions for addressing the data scarcity problem — though this characterization reflects theoretical promise and preliminary results rather than established performance on matched benchmark datasets. This distinction is important: the claim that these approaches are 'promising' is supported by early results in analogous domains (computer vision, medical imaging) and by logical inference from the structure of the PD classification problem, not by systematic empirical comparison against standard methods on shared PD datasets. Such comparisons remain a research need.

Transfer learning, adapting models pre-trained on large datasets to data-limited target domains, has demonstrated accuracy advantages in preliminary PD applications [16]. The systematic investigation of transfer between equipment types, insulation materials, and voltage levels remains incomplete, as does the identification of which layers of pre-trained networks encode transferable versus equipment-specific features.

Few-shot learning methods, such as siamese networks, prototypical networks, and meta-learning, have demonstrated the capacity to classify new image categories from 1–5 examples in computer vision tasks. Adapting these methods to PD classification may allow rapid adaptation to newly encountered fault types without requiring hundreds of labeled examples [31]. Whether the structure of PRPD patterns supports the same geometric properties in feature space that make few-shot learning work in natural images is a question requiring dedicated investigation.

9.3. Physics-Informed Neural Networks

Hybrid physics-ML approaches in which physical constraints (Paschen breakdown law, charge conservation, field equations) are incorporated into neural network loss functions or architectures represent a potentially transformative research direction, again characterized here as a direction rather than an established methodology. The physics-informed neural network (PINN) framework has demonstrated improved generalization and sample efficiency in several engineering domains by constraining models to respect known physical laws during training [32]. An application of PINNs to PD classification might constrain the model to predict PRPD patterns consistent with known void discharge physics for specific insulation-voltage combinations — using physical law as domain-knowledge regularization. The potential benefit is models that generalize to new equipment types with less training data because their learned representations are grounded in physics rather than purely statistical correlations. Direct empirical validation against purely data-driven baselines on matched PD datasets is needed before this approach can be recommended for practice.

9.4. Cold-Start Deployment Strategies

The cold-start problem, the lack of sufficient labeled training data when a monitoring program is first deployed, is a practical barrier that receives less attention in the research literature than in operational reality. Utilities deploying ML-PD monitoring today cannot wait 18 months to accumulate labeled field data before the system becomes useful. Several strategies address this:

- Finite element simulation of PRPD patterns for known defect geometries can generate physically grounded synthetic training examples before any equipment has faulted. Studies in [9] and [33] document the validity bounds and limitations of simulation-based PRPD generation; simulated data is most useful for training initial feature extractors and least reliable for representing the statistical properties of real-world discharge patterns.

- Accelerated aging laboratory protocols (IEC 60343, IEEE Std 930) provide labeled fault data under controlled conditions. The known limitation is that accelerated aging achieved by elevated temperature, voltage, or both may not faithfully reproduce the statistical fingerprints of naturally aged insulation, where chemical byproduct accumulation and mechanical fatigue interact over years rather than weeks.
- Generative adversarial networks (GANs) and VAEs (Variational Autoencoders) can augment limited real datasets by generating synthetic PRPD patterns. The validation challenge is ensuring that synthetic patterns are physically plausible rather than merely statistically consistent with the training distribution — a distinction that requires domain expert assessment and cannot be resolved by generative model loss functions alone [20].
- Cross-utility data consortia — analogous to the Cancer Imaging Archive in medical AI or the PhysioNet repository in clinical signal processing — could pool de-identified PD monitoring data across organizations, dramatically expanding the labeled dataset available for training and validation. The organizational obstacles are substantial: competitive sensitivity of equipment condition data, heterogeneity of sensor systems across utilities, and unresolved intellectual property and liability questions in the event of model failures trained on shared data. Creating trusted custodianship models through industry consortia, national laboratories, or university collaborations is as much an organizational challenge as a technical one.

9.5. Long-Term Field Validation

The longest continuous ML-PD monitoring programs documented in the literature extend to three years (Hydro-Québec, Table 3). Demonstrating model reliability across 10–20-year asset lifetimes requires commitments of commensurate duration that current research programs have not yet undertaken. Laboratory accelerated aging cannot substitute: while valuable for understanding degradation mechanisms, it cannot replicate the complex interaction of real-world stressors over multi-decade timescales. Establishing long-term joint data-sharing and validation programs among utilities, data center operators, and academic research institutions is among the highest-priority organizational investments the field can make.

10. Conclusion

Machine learning-enhanced partial discharge detection has advanced substantially over the past fifteen years, moving from proof-of-concept demonstrations to pilot-scale field deployments with documented operational utility. This review has synthesized the current state of that progress across the methodological spectrum — from SVM and random forest approaches through convolutional networks, LSTM-based prognostic systems, and multi-modal fusion architectures — and evaluated what has and has not been demonstrated under operational conditions.

The following conclusions are drawn directly from the empirical evidence reviewed:

- CNN-based architectures consistently outperform hand-crafted feature methods when sufficient labeled training data is available (above approximately 800 samples), with accuracy advantages of 3–8 percentage points over SVM and RF approaches on balanced laboratory datasets.

- Multi-modal sensor fusion reliably improves both classification accuracy and system resilience over single-modality approaches, with the added benefit of graceful degradation when individual sensor channels become unavailable.
- Field-deployed systems achieve positive predictive values of 64–89% meaningfully above the 30–50% PPV of rule-based alarm systems but below the threshold that would justify autonomous maintenance decision-making without human review.
- Data scarcity (particularly for rare fault types), distribution shift under changing operational conditions, and the absence of standardized evaluation protocols are the primary barriers to broader adoption, not at this stage, fundamental limitations of the ML architectures themselves.

The following conclusions are explicitly not drawn from the reviewed evidence, because the evidence does not support them:

- That any current ML-PD architecture is validated for autonomous maintenance action in conditions not represented in available field data.
- That hybrid physics-ML methods or transfer learning approaches are demonstrably superior to standard methods — these are research directions with theoretical justification and preliminary encouraging results, not established practices with validated performance on shared benchmark datasets.

Whether ML-assisted PD detection achieves the status of a standard infrastructure reliability tool rather than remaining a promising but incompletely validated research direction — will depend on progress in at least four areas identified in this review as critical: long-term field validation programs of sufficient duration to span meaningful fractions of asset lifetimes; standardized evaluation protocols enabling fair cross-institutional and cross-equipment comparison; regulatory frameworks creating clear pathways for compliance through continuous monitoring rather than periodic offline inspection; and industry data-sharing arrangements addressing the chronic shortage of labeled fault data. The technical foundations are demonstrably strong. The systemic conditions for broad adoption are still being constructed.

Funding: This research received no external funding.

Data Availability Statement: “The data that support the findings of this study are available from the corresponding author upon reasonable request.”

Conflicts of Interest: “The authors declare no conflicts of interest.”

References

[1] G. C. Stone, “Partial discharge diagnostics and electrical equipment insulation condition assessment,” *IEEE Transactions on Dielectrics and Electrical Insulation*, vol. 12, no. 5, pp. 891–903, Oct. 2005, doi: 10.1109/TDEI.2005.1522184.

[2] A. Abubakar Mas’ud, B. G. Stewart, and S. G. McMeekin, “Application of an ensemble neural network for classifying partial discharge patterns,” *Electric Power Systems Research*, vol. 110, pp. 154–162, May 2014, doi: 10.1016/j.epsr.2014.01.010.

[3] W. J. K. Raymond, H. A. Illias, A. H. A. Bakar, and H. Mokhlis, “Partial discharge classifications: Review of recent progress,” *Measurement*, vol. 68, pp. 164–181, May 2015, doi: 10.1016/j.measurement.2015.02.032.

[4] Hui Ma, J. C. Chan, T. K. Saha, and C. Ekanayake, “Pattern recognition techniques and their applications for automatic classification of artificial partial discharge sources,” *IEEE Transactions on Dielectrics and Electrical Insulation*, vol. 20, no. 2, pp. 468–478, Apr. 2013, doi: 10.1109/TDEI.2013.6508749.

[5] M. Florkowski, “Classification of Partial Discharge Images Using Deep Convolutional Neural Networks,”

Energies (Basel)., vol. 13, no. 20, p. 5496, Oct. 2020, doi: 10.3390/en13205496.

[6] D. Soh, S. B. Krishnan, J. Abraham, L. K. Xian, T. K. Jet, and J. F. Yongyi, "Partial Discharge Diagnostics: Data Cleaning and Feature Extraction," *Energies (Basel).*, vol. 15, no. 2, p. 508, Jan. 2022, doi: 10.3390/en15020508.

[7] X. Zhang *et al.*, "Review on Detection and Analysis of Partial Discharge along Power Cables," *Energies (Basel).*, vol. 14, no. 22, p. 7692, Nov. 2021, doi: 10.3390/en14227692.

[8] U.S. Department of Energy, "Transforming the nation's electricity system: The second installment of the QER," Washington D.C., 2017.

[9] A. Cavallini, G. Montanari, M. Tozzi, and X. Chen, "Diagnostic of HVDC systems using partial discharges," *IEEE Transactions on Dielectrics and Electrical Insulation*, vol. 18, no. 1, pp. 275–284, Feb. 2011, doi: 10.1109/TDEI.2011.5704519.

[10] M. Duval, "A review of faults detectable by gas-in-oil analysis in transformers," *IEEE Electrical Insulation Magazine*, vol. 18, no. 3, pp. 8–17, May 2002, doi: 10.1109/MEI.2002.1014963.

[11] International Electrotechnical Commission, *High-voltage test techniques — Partial discharge measurements*. Geneva, 2015.

[12] M. D. Judd, O. Farish, and B. F. Hampton, "The excitation of UHF signals by partial discharges in GIS," *IEEE Transactions on Dielectrics and Electrical Insulation*, vol. 3, no. 2, pp. 213–228, Apr. 1996, doi: 10.1109/94.486773.

[13] S. Markalous, S. Tenbohlen, and K. Feser, "Detection and location of partial discharges in power transformers using acoustic and electromagnetic signals," *IEEE Transactions on Dielectrics and Electrical Insulation*, vol. 15, no. 6, pp. 1576–1583, Dec. 2008, doi: 10.1109/TDEI.2008.4712660.

[14] L. Hao, P. L. Lewin, Y. Tian, and S. J. Dodd, "Partial discharge identification using a support vector machine," in *CEIDP '05. 2005 Annual Report Conference on Electrical Insulation and Dielectric Phenomena, 2005.*, IEEE, 2005, pp. 414–417. doi: 10.1109/CEIDP.2005.1560708.

[15] M. Kunicki, A. Cichoń, and Ł. Nagi, "Statistics based method for partial discharge identification in oil paper insulation systems," *Electric Power Systems Research*, vol. 163, pp. 559–571, Oct. 2018, doi: 10.1016/j.epsr.2018.01.007.

[16] X. Peng *et al.*, "Random Forest Based Optimal Feature Selection for Partial Discharge Pattern Recognition in HV Cables," *IEEE Transactions on Power Delivery*, vol. 34, no. 4, pp. 1715–1724, Aug. 2019, doi: 10.1109/TPWRD.2019.2918316.

[17] Y. Wang, J. Yan, Z. Yang, T. Liu, Y. Zhao, and J. Li, "Partial Discharge Pattern Recognition of Gas-Insulated Switchgear via a Light-Scale Convolutional Neural Network," *Energies (Basel).*, vol. 12, no. 24, p. 4674, Dec. 2019, doi: 10.3390/en12244674.

[18] M.-T. Nguyen, V.-H. Nguyen, S.-J. Yun, and Y.-H. Kim, "Recurrent Neural Network for Partial Discharge Diagnosis in Gas-Insulated Switchgear," *Energies (Basel).*, vol. 11, no. 5, p. 1202, May 2018, doi: 10.3390/en11051202.

[19] N. H. Nik Ali, W. Goldsmith, J. A. Hunter, P. L. Lewin, and P. Rapisarda, "Comparison of clustering techniques of multiple partial discharge sources in high voltage transformer windings," in *2015 IEEE 11th International Conference on the Properties and Applications of Dielectric Materials (ICPADM)*, IEEE, Jul. 2015, pp. 256–259. doi: 10.1109/ICPADM.2015.7295257.

[20] G. Li, X. Wang, X. Li, A. Yang, and M. Rong, "Partial Discharge Recognition with a Multi-Resolution Convolutional Neural Network," *Sensors*, vol. 18, no. 10, p. 3512, Oct. 2018, doi: 10.3390/s18103512.

[21] Y. Wang, J. Yan, Q. Sun, J. Li, and Z. Yang, "A MobileNets Convolutional Neural Network for GIS Partial Discharge Pattern Recognition in the Ubiquitous Power Internet of Things Context: Optimization, Comparison, and Application," *IEEE Access*, vol. 7, pp. 150226–150236, 2019, doi: 10.1109/ACCESS.2019.2946662.

[22] G. C. Montanari, R. Ghosh, L. Cirioni, G. Galvagno, and S. Mastroeni, "Partial Discharge Monitoring of Medium Voltage Switchgears: Self-Condition Assessment Using an Embedded Bushing Sensor," *IEEE Transactions on Power Delivery*, vol. 37, no. 1, pp. 85–92, Feb. 2022, doi: 10.1109/TPWRD.2021.3053658.

[23] Q. Zheng *et al.*, "A real-time transformer discharge pattern recognition method based on CNN-LSTM driven by few-shot learning," *Electric Power Systems Research*, vol. 219, p. 109241, Jun. 2023, doi: 10.1016/j.epsr.2023.109241.

[24] X. Zhou, C. Zhou, and I. J. Kemp, "An improved methodology for application of wavelet transform to partial discharge measurement denoising," *IEEE Transactions on Dielectrics and Electrical Insulation*, vol. 12, no. 3, pp. 586–594, Jun. 2005, doi: 10.1109/TDEI.2005.1453464.

[25] S. Ness, "Adversarial Attack Detection in Smart Grids Using Deep Learning Architectures," *IEEE Access*, vol. 13, pp. 16314–16323, 2025, doi: 10.1109/ACCESS.2024.3523409.

[26] Y. Jia, J. Wang, C. M. Poskitt, S. Chattopadhyay, J. Sun, and Y. Chen, "Adversarial attacks and mitigation for anomaly detectors of cyber-physical systems,"

International Journal of Critical Infrastructure Protection, vol. 34, p. 100452, Sep. 2021, doi: 10.1016/j.ijcip.2021.100452.

[27] T. T. Khoei, H. O. Slimane, and N. Kaabouch, “Cyber-Security of Smart Grids: Attacks, Detection, Countermeasure Techniques, and Future Directions,” *Communications and Network*, vol. 14, no. 04, pp. 119–170, 2022, doi: 10.4236/cn.2022.144009.

[28] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, “Communication-Efficient Learning of Deep Networks from Decentralized Data,” 2023.

[29] C. Dwork and A. Roth, “The Algorithmic Foundations of Differential Privacy,” *Foundations and Trends® in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–487, Aug. 2014, doi: 10.1561/0400000042.

[30] C. Rudin, “Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead,” *Nat. Mach. Intell.*, vol. 1, no. 5, pp. 206–215, May 2019, doi: 10.1038/s42256-019-0048-x.

[31] O. Vinyals, C. Blundell, T. Lillicrap, K. Kavukcuoglu, and D. Wierstra, “Matching networks for one shot learning,” in *NIPS’16: Proceedings of the 30th International Conference on Neural Information Processing Systems*, 2016, pp. 3637–3645.

[32] M. Raissi, P. Perdikaris, and G. E. Karniadakis, “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations,” *J. Comput. Phys.*, vol. 378, pp. 686–707, Feb. 2019, doi: 10.1016/j.jcp.2018.10.045.

[33] E. Lemke, “A critical review of partial-discharge models,” *IEEE Electrical Insulation Magazine*, vol. 28, no. 6, pp. 11–16, Nov. 2012, doi: 10.1109/MEI.2012.6340519.