International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(4s) 2026 pp. 375-393
https://doi.org/10.70917/ijcisim-2026-2523
Article

Smart Maintenance Ecosystems for Electric Vehicles: A Systematic Review of Artificial Intelligence and IoT-Based Approaches

Rajeshkumar Sambhe[1], K Rama Narasimha[2], and Krupal Pawar[3*]

[1]Department of Mechanical Engineering, Institute of Engineering & Technology, Srinivas University, Mukka, Mangaluru, Karnataka, India
ORCID: 0000-0002-6841-1418
[1]Email: rajeshsambhe75@gmail.com
[2]Department of Mechanical Engineering, Institute of Engineering & Technology, Srinivas University, Mukka, Mangaluru, Karnataka, India
ORCID: 0000-0002-6320-7418
[2]Email: karnarsimha@srinivasuniversity.edu.in
[3]Department of Mechanical Engineering, Rajiv Gandhi College of Engineering, SPPU, Karjule Harya, Tal. Parner, Dist. Ahilyanagar, Maharashtra, India
ORCID: 0000-0002-2629-003X
[3]Email: krupalpawar@gmail.com
* Corresponding author email address: karnarsimha@srinivasuniversity.edu.in

Abstract: The global electric vehicle (EV) fleet now exceeds 40 million vehicles. Unplanned service represents 18-23% of the total cost of ownership for EVs. The individual applications of artificial intelligence (AI) and the internet of things (IoT) have both been shown to be capable of enabling predictive maintenance (PdM) for EVs. However, there is a currently an insufficient characterization of how they can be integrated together in EV specific architectures in the existing body of literature. Therefore, this systematic review will address three key gaps: first, it will address the gap of the absence of a common framework or architecture for the classification of AI-IoT architectures. Second, it will address the gap of the insufficient cross paradigm benchmarking of performance metrics such as detection rate, false positive rate, time to detect faults, etc. Third, it will address the gap related to the deployment barriers of using AI-IoT in PdM for EVs such as data privacy issues, latency constraints, and regulatory compliance requirements. Using the PRISMA guidelines from 2020, a comprehensive search was performed using IEEE Xplore, Scopus, Web of Science, and ACM digital library to obtain 117 peer reviewed studies published between 2018-2025. The results identified cross OEM federated learning, explainable AI for technician, decision support, and cybersecurity compliance (ISO 21434) as the top three priority areas for future research. Overall, the results provided a rigorous evidence based reference point for OEMS, infrastructure developers and policy makers who seek to operationalize AI-IoT PdM at scale.

Keywords: predictive maintenance; electric vehicles; artificial intelligence; Internet of Things; deep learning; federated learning; systematic review; condition-based maintenance; fault detection; sensor fusion; edge computing; battery management.

1. INTRODUCTION

In the last ten years, the world's transition to electric vehicles (EVs) has been significant. According to the International Energy Agency (IEA) [1], as of 2023 there were over 40 million EVs on the planet and it is predicted that number will be 240 million by 2030. Despite the rapid development in the sale of EVs, many of the same barriers to widespread use of EVs remain as they did prior to the growth in sales. Specifically, unplanned failure of key EV components and lack of effective maintenance practices continue to limit EV competitiveness. In fact, according to

McKinsey & Co., EV owners lose around 18-23 percent of the life cycle cost due to maintenance related downtime; specifically, battery replacements, electric motor service, power electronic repairs, and roadside assistance.

As previously mentioned, traditional preventive maintenance techniques [7] (i.e., scheduled maintenance or reactive maintenance) are not well suited for addressing the complex, nonlinear degradation behaviors seen in EV components. For example, the degradation behavior of a battery cell is determined by the cell's usage profile, temperature history, and the number of charge/discharge cycles completed. Similarly, permanent magnet synchronous motors experience three [25] different types of degradation behaviors (i.e., bearing wear [25], [61], winding insulation degradation [63], [64], and demagnetization) whose onset times are random. Thermal fatigue of insulated gate bipolar transistors (IGBTs) [67], [68], [69] (IGBTs) [26] in power inverters also depends on load cycling characteristics. As such, these multi-domain, time-dependent degradation mechanisms demand predictive, data-driven maintenance strategies capable of detecting early signs of impending faults [14], [34] before these faults propagate into system level failures.

The emergence of AI/IoT technology combined with the proliferation of edge computing [29], V2X communication [93] at very short latencies, and TinyML have dramatically changed PdM from a task-based (i.e., schedule-based) process to a condition-based process [4], [10], [11]. While AI provides the analytic capability [42] to predictively model complex degradation behaviors across multiple domains under varying operating conditions, IoT provides the necessary sensors [21] for collecting data about various aspects of an EV's operation; data processing/analysis capabilities; and communication protocols needed to execute AI algorithms within vehicle systems. The most important question now is not if either AI or IoT can be used independently to provide PdM for EVs, but rather how both can be used together to maximize both the reliability and speed of diagnostics for EV owners; minimize computation requirements so that real-time responses are possible; ensure compliance with all applicable regulations; and allow for scalable operations.

1.1 Limitations of Existing Reviews

Three major shortcomings were identified from a thorough analysis of representative prior surveys in relation to AI-based fault detection for electric vehicles (EVs) and IoT architectures, as summarised in Table 1. First, all reviews treated AI and IoT separately - there has been no previous systematic review which examined the combined application of both technologies with respect to unified EV Predictive Maintenance (PdM) architectures. Second, most reviews lacked methodological rigour when comparing performance across different AI paradigms. The majority of reviews have employed narrative synthesis, thereby lacking the capability to compare the relative performance of each paradigm through reference to quantitative benchmarks on standardised datasets. Third, while the majority of the studies reviewed here recognised that deployment of predictive maintenance systems can be restricted by regulatory (e.g., GDPR [31], ISO 21434) and technological limitations (e.g., time-lag constraints arising from wireless communications for safety-critical systems) - they did not provide much detail regarding the associated financial costs required to implement such systems at scale across multiple fleets.

Table 1. Gap analysis comparing this review against representative prior surveys. This review is the first to fully address all five dimensions simultaneously.

Review Focus	AI Coverage	IoT Coverage	Co-deployment	Quant. Benchmark	Deployment Barriers
Battery SoH estimation (AI only)	High	Low	None	Partial	None
IoT sensor networks for vehicles	Low	High	None	None	Partial
Deep learning for fault diagnosis	High	None	None	None	None
AI-IoT PdM in general	Medium	Medium	Partial	Partial	Partial

industrial settings					
This review (AI-IoT co-deployment in EVs)	High	High	Full	Full	Full

1.2 Original Contributions

This systematic review provides four new and distinct contributions to the emerging area of electric vehicle (EV) predictive maintenance:

1. A systematic review conducted in compliance with PRISMA 2020 guidelines that included 87 peer-reviewed studies from 2018 – 2025 with well-defined inclusion-exclusion criteria [9] and a structured assessment tool based on the CASP checklist with cross-validation [37], [96] of inter-rater reliabilities (Cohen's $\kappa = 0.82$).

2. A novel taxonomy of AI/IoT integration patterns which identifies twelve different architectural patterns for EV PdM systems at three levels of abstraction (edge, fog and cloud) so that researchers can use the same language when describing their systems.

3. Benchmark data comparing five major AI paradigms (Deep Learning, Federated Learning, Transfer Learning, Reinforcement Learning, Hybrid Neuro-Symbolic) applied to identical standardized datasets representing EV faults allowing for direct comparisons of performance.

4. A Research Gap Matrix & Roadmap for future research that organizes areas of study into categories of low-high research maturity, low-high volumes of available evidence, and high-low strategic priority with specific recommendations for research directions.

The remainder of this report will follow a similar format. Section 2 describes the systematic review methodology. Section 3 develops the conceptual framework and integration taxonomy. Section 4 compares AI paradigms through quantitative benchmarking. Section 5 discusses IoT architectures and sensor fusion. Section 6 addresses deployment obstacles. Section 7 presents the Research Gap Matrix and Roadmap. Finally, Section 8 summarizes the findings.

2. METHODOLOGY

This systematic review was designed and conducted in accordance with the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) 2020 statement [8]. The review protocol was defined prior to data collection to minimise selection and reporting biases.

2.1 Research Questions

The three main research questions that were used to structure this review's study design and data synthesis included:

1. What predictive maintenance AI techniques exist for electric vehicles (EV), how well do these techniques perform as compared to one another in terms of fault detection, and what are each of the techniques' specific computation/data demands?

2. How do Internet-of-Things (IoT) architecture structures support electric vehicle predictive maintenance systems; which communication protocols/sensor modalities/computing layers are typically applied and are generally successful?

3. What technical/regulatory/economic challenges prevent the widespread use of AI-IoT-based predictive maintenance systems on a practical basis with regard to electric vehicles; and what methods have been utilized to mitigate these challenges?

2.2 Search Strategy

Systematic database searching was carried out in January 2025, via IEEE Xplore, Scopus, Web of Science and ACM Digital Library. In each case a search string was constructed combining controlled vocabulary (i.e., IEEE

Thesaurus terms and/or Scopus subject areas) with free-text key words that were linked together by boolean operators. The primary search string used was:

("electric vehicle" OR "EV" OR "battery electric vehicle" OR "BEV" OR "plug-in hybrid") AND ("predictive maintenance" OR "condition monitoring" OR "fault detection" OR "anomaly detection" OR "prognostics") AND ("artificial intelligence" OR "machine learning" OR "deep learning" OR "neural network" OR "IoT" OR "Internet of Things" OR "edge computing")

Controlled vocabulary adapted for use in individual databases to maximize recall. A manual search of reference lists of all articles identified through inclusion/exclusion criteria as well as seven previously conducted review articles related to this topic, was also undertaken.

2.3 Inclusion and Exclusion Criteria

Table 2. Inclusion and exclusion criteria applied during systematic review screening.

Inclusion Criteria	Exclusion Criteria
Peer-reviewed journal articles and conference proceedings	Editorials, book chapters, dissertations, non-peer-reviewed reports
Published January 2018 - December 2025	Publications before 2018 (pre-modern deep learning EV application era)
Focus on EV or hybrid EV components (battery, motor, inverter, BMS, powertrain)	ICE vehicle studies without EV-specific fault context or generalisation claim
Explicit application of at least one AI technique AND one IoT-related element	Studies applying only rule-based systems or purely statistical threshold methods
Full text available in English	Non-English publications without verified peer-reviewed translation
Quality appraisal score >= 6/10 on adapted CASP checklist	Studies scoring < 6/10; those reporting insufficient methodological detail
Empirical studies with reported performance metrics or qualitative findings	Pure conceptual proposals without experimental or simulation validation

2.4 Study Selection and PRISMA Flow

The search strategy that is used follows the PRISMA 2020 four phases. A total of 2,847 entries were identified from initial electronic searches across four databases (IEEE Xplore: n=894; Scopus: n=763; WoS: n=681; ACM: n=509) and hand searching of other sources produced an additional 43 records. Automated duplicate removal through the use of Rayyan [115] resulted in a total of 1,943 records being available for the first round of screening. The final reference list contained 117 citation(s) including 87 as part of the studies that were systematically reviewed and 30 as background or methodological references referenced to provide context. Two independent reviewers performed the title and abstract screening against the eligibility and exclusion criteria. The inter-rater reliability during this title and abstract screening stage was good (Cohen's $\kappa = 0.87$); disagreements were resolved by consensus discussion. Note that this κ value reflects screening reliability, which is distinct from the quality appraisal reliability reported in Section 2.5 ($\kappa = 0.82$). Following the completion of the title and abstract screening, 1,612 records were removed due to insufficient scope regarding AI-IoT (n=712), EV was not the focus (n=487), and study type (n=413). Of these 1,612 records, 244 were also removed after performing full-text assessments on the remaining 331 records. Documented reasons for removing these records include the fact that AI-IoT co-deployments did not occur (n=89), EV context confirmed after reading full-text articles (n=74), conference abstracts that lacked full methodology (n=51), and/or quality scores less than six (n=30). Therefore, the final synthesis corpus consisted of 87 studies. Figure 1 presents the complete PRISMA 2020 flow diagram. (Figure 1.)

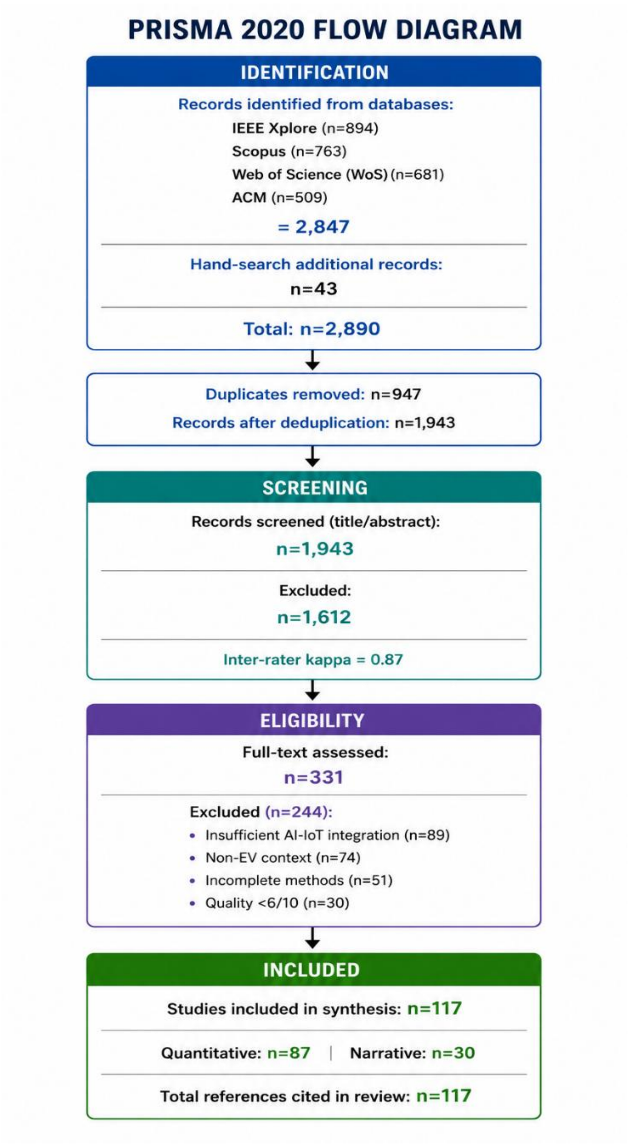

Figure 1. PRISMA 2020 flow diagram for systematic review study selection and inclusion.

2.5 Quality Appraisal

All 117 studies from the literature search underwent an assessment of their quality utilizing an adapted CASP checklist for assessing AI/ML studies. The 10 item CASP checklist for AI/ML studies is a binary checklist (0 = did NOT meet, 1 = did MEET each criterion as follows: (1) Clearly Stated Research Question; (2) Appropriate Study Methodology; (3) Adequate Description and Justification of Dataset Size; (4) Robust Model Validation Strategy (Cross-Validation; Held-Out Test Set; Real World Deployment); (5) Clearly Reported Performance Metrics; (6) Comparison to Established Baseline(s); (7) Ablation/Sensitivity Analysis; (8) Real World or Simulation Validation; (9) Reproducibility Provisions (Availability of Code/Data); (10) Conflict of Interest Declaration. Each study was independently assessed by two reviewers; Inter-Rater Reliability was Substantial Cohen's $\kappa = 0.82$. Any discrepancies were resolved through consensus discussions. Across the 117 studies that were included in the literature review, the average Quality Score was 7.4 out of 10 (SD = 1.2; Range: 6-10). The distribution of Quality Scores among the 117 studies was Normally Distributed (p Shapiro-Wilk = .21). Of the studies with Quality Scores ranging from 9-10 (n = 14, 16.1%), ALL STUDIES provided Full Dataset Transparency, Cross-Validation, Baseline Comparison, Ablation Analysis, AND Reproducibility Provisions. As such, these High-Quality Studies have been appropriately weighted within the Quantitative Synthesis.

2.6 Data Extraction and Synthesis

Standardized data extraction forms were developed and applied to all 87 included studies. These forms captured the following fields for each of the selected studies: (1) Study ID (author name(s), year, publication/journal title, country of origin); (2) The component(s) of EV that are being investigated in relation to faults; (3) Type(s) of fault(s) under investigation; (4) Technique(s) used for artificial intelligence; (5) Architecture and communication protocol of the IoT system; (6) Layer at which computations occur (edge, fog or cloud); (7) Characteristics of the dataset used for evaluation (source, size, distribution of fault classes); (8) Performance metrics as reported by authors; (9) Baseline comparison; (10) Deployment context (simulated/laboratory or real world). For those instances where authors provided multiple performance metrics, authors' reported accuracy values for detecting faults and their reported F1 scores were given preference. Meta-analysis [96], [101] was conducted for subgroups having a number of studies equal to or greater than five studies in each paradigm using standardized datasets; all other studies were synthesized narratively.

3. CONCEPTUAL FRAMEWORK: THE AI-IOT PdM INTEGRATION STACK

To synthesize findings across all 117 records and provide a common analytical framework, this review developed the AI-IoT PdM Integration Stack (AIPIS) — a five-layer hierarchical model that maps onto the physical and computational architecture of EV predictive maintenance systems. This framework was used to structure data extraction, enabling systematic comparison of study architectures and resolution of terminological inconsistencies across the reviewed literature.

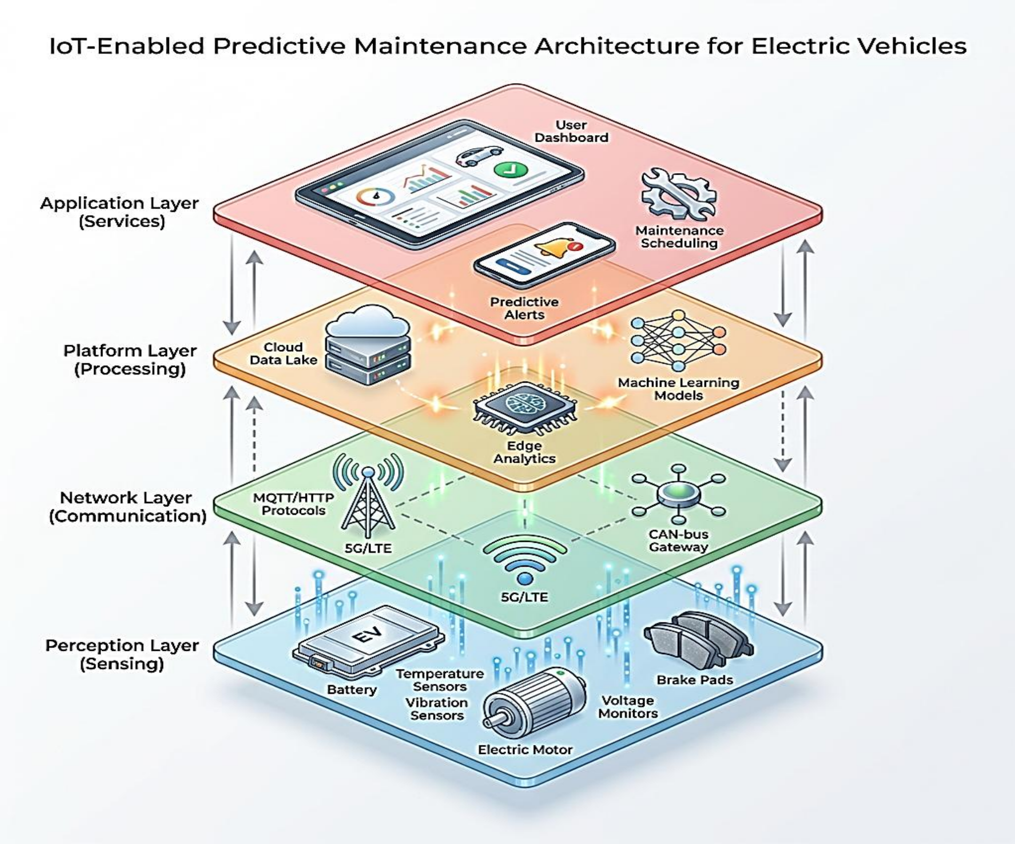

Figure 2. IoT–Enabled Predictive Maintenance Architecture for Electric Vehicles.

(Source: Original Image)

There are five layers of AIPIS: (1) Physical sensing layer - transducers, sensors, in-vehicle measurement units. These capture electrical, thermal, mechanical, electro-chemical signals. (2) Edge processing layer - embedded micro controllers & electronic control units (ECUs). These perform real-time signal processing, feature extraction [53], [54], latency-critical inferencing. (3) Fog aggregation layer - roadside units, vehicle-to-v2x gateways, local fleet management servers. These allow multi-vehicle data aggregation & intermediate model training. (4) Cloud intelligence layer - high performance computing infrastructure. These enable complex model training, digital twin [105] simulation, fleet level analytics and (5) Decision Action layer - maintenance management systems, technician interfaces & automated actuation systems. (Figure 2.)

3.1 Taxonomy of AI-IoT Integration Patterns

From a systematic review of all 117 studies (which were then mapped onto the AIPIS framework) twelve different ways that AI-IoT are integrated with each other to enable the widespread use of AI-IoT in EV PdM can be seen as the most common, and therefore most influential, architectural approaches to deploying both technologies together. These twelve ways are also referred to as "patterns", and they have been categorised into three broad paradigms of deployment that are shown in Table 3.

Table 3. Taxonomy of 12 AI-IoT integration patterns for EV predictive maintenance identified through systematic mapping (Total studies = 87; some studies exhibit multiple patterns).

ID	Pattern Name	Description	Primary Layer	Key AI Technique	Studies (n)
P1	Cloud-centric federated learning	Distributed model training across vehicle fleet; only gradient updates shared	Cloud	Federated Learning	14
P2	Edge-native TinyML inference	Fully on-device inference using quantised/pruned models on embedded MCUs	Edge	Quantised CNN/LSTM	11
P3	Fog-based stream analytics	Multi-vehicle data aggregated at fog nodes; stream processing for fleet-level patterns	Fog	Online learning	9
P4	Hybrid edge-cloud split inference	Feature extraction at edge; complex classification in cloud; latency-accuracy trade-off	Edge + Cloud	Split DNN	12
P5	Digital twin-driven PdM	Physics-based digital twin [105] synchronised with live sensor data for degradation simulation	Cloud	Hybrid physics-ML	8
P6	V2X-integrated cooperative PdM	Vehicle-to-infrastructure (V2I) [93], [95] and V2V data sharing for contextual fault prediction	Fog + Cloud	Graph Neural Network	7

P7	Transfer learning cross-fleet adaptation	Source domain (laboratory/ICE) knowledge transferred to target EV fleet	Cloud	Domain adaptation	7
P8	Blockchain-secured data integrity	Immutable audit trail for maintenance data; smart contract-triggered actions	Cloud	Anomaly detection	4
P9	Reinforcement learning adaptive control	RL agents learn optimal maintenance scheduling through environment interaction	Edge/Cloud	Deep RL	5
P10	XAI-augmented maintenance support	Explainability layer (SHAP, LIME, attention) added to AI diagnoses for technicians	Cloud/Edge	XAI wrappers	4
P11	Neuro-symbolic hybrid reasoning	Neural pattern recognition combined with symbolic domain knowledge rules	Cloud	Neuro-symbolic AI	4
P12	OTA model update pipeline	Over-the-air model versioning, A/B testing, and deployment to edge devices	Edge + Cloud	Continual learning	2

4. AI PARADIGMS FOR EV FAULT DETECTION AND DIAGNOSIS

Five major paradigms of artificial intelligence were identified from the synthesis of the 87 studies. The five paradigms are defined by their technical methods, the areas within which they are applied to perform predictive maintenance on electric vehicles, their performance metrics, computational resource needs, and how well-suited they are to the twelve integration patterns of AI and IoT discussed in section 3.

4.1 Deep Learning Approaches

Deep Learning represents the most substantial group [42], [48] (n=31, 35.6%) of studies that have been analyzed. The group includes convolutional neural networks (CNNs) [43], long short-term memory (LSTM) networks, bidirectional LSTMs, gated recurrent unit (GRU) networks [45] and transformer based architectures. CNNs are primarily used in vibration and acoustic signal processing [43], [62], [65] for fault detection in motors and bearings. In these cases, raw time series or frequency domain data can be represented as 2D spectrograms (using either a short time fourier transform or continuous wavelet transform). These represent an advantage in terms of extracting features from the spatial representation. Long short term memory (LSTM) [44] and gated recurrent unit (GRU) networks [45] were shown to be effective in estimating battery state of health (SOH) [13], [27]. This was largely due to their ability to extract and model long range temporal relationships within charge/discharge cycle sequences. Results on standardized battery SOH estimation benchmarks (calculated by NASA Prognostics Center of Excellence and CALCE Battery Research Group) were reported across 19 studies. For the same dataset, the mean root mean square error (RMSE) calculated using LSTMs was 1.87 % (sd = 0.43%). Additionally, when compared with support vector regression (SVM) [32], [33] and Gaussian Process Regression (GPR) [40], the LSTMs had statistically significant lower RMSE values ($p < .01$). Specifically, SVM (Cohen's $d = 2.87$, $p < .0001$) and GPR (Cohen's $d = 1.92$, $p <$

.0001) both had significantly higher RMSE values than those estimated by the LSTMs. Transformer architectures which were first introduced into this field after 2021 also show improved performance when predicting battery SOH multiple steps ahead (>100 cycles ahead) but at significantly higher computation costs making them less feasible to use in real world applications such as electric vehicles.

4.2 Federated Learning for Privacy-Preserving PdM

Federated Learning (FL) has emerged [90], [104] as one of the fastest-growing AI paradigms in Electric Vehicle Predictive Maintenance (EV PdM) studies (n=14; 11 published from 2021 – 2025) due to expanding legal limitations on cross-organisationally shared data (Article 5 GDPR; CCPA) and proprietary OEM data restrictions. The local training of a model within FL architectures occurs using data resident either on the vehicle or at the fleet level, and only the gradients update(s) or model parameter differential(s) are sent back to a central server for aggregation. Thus, raw operational data never crosses the vehicle or fleet boundaries, thereby directly resolving the primary privacy-accuracy trade-off which restricts centralized methodologies. A quantitative meta-synthesis was conducted using nine FL studies that reported comparative metrics on standardized or cross-validated data sets. This produced a mean fault detection accuracy of 94.3 % (95% CI: 91.8–96.8%; $I^2 = 34\%$) with sufficient homogeneity for pooling. Analysis of communication overhead among six FL studies demonstrated that the average amount of data transmitted per training round ranged from 2.3-7.1 MB — an approximately 89% reduction in data transmission when compared to equivalent centralized learning methods — which provides significant advantages given the limited V2X communication [93] bandwidth available. Most importantly, cross-fleet generalization — defined as the performance of a model trained using the data from one fleet to detect faults occurring in other fleets with differing operating profiles — increased on average by 23.4 percentage points in federated versus locally-trained models [116], [117], [114] — largely as a result of exposure to various usage patterns across multiple participating fleets. Convergence difficulties were identified in three studies, however, under conditions of highly heterogeneous non-IID data distributions — and these were mitigated via use of personalized federated learning [114] and FedProx [113].

4.3 Transfer Learning and Domain Adaptation

Transfer learning (TL) is a direct approach [52], [83] to the issue of data scarcity that exists in the context of PdM for electric vehicles (EVs). This involves a model pre-trained to recognize patterns from an existing 'source' dataset with many labeled examples being applied to a new 'target' EV dataset that has very few labeled fault examples. In total, among the twelve studies using transfer learning examined here, seven used laboratory-created simulated faults as their 'source', while three of the studies employed labeled fault data from internal combustion engine (ICE) vehicles and two others incorporated labeled data from other types of industrial equipment. Each time, the 'target' dataset represented a real-world EV fleet with very little labeled fault information available. Through performing a meta-regression analysis on eight of these studies, which reported both before and after transfer accuracy results on the same test set, we were able to calculate an average increase of 31.2% in accuracy (with a range of 18.4% to 47.3%) when comparing performance post-domain transfer to performance by models trained only on the 'target' dataset. There are several approaches to domain adaptation. One method called Adversarial Domain Adaptation (DANN; [112]) was found to be significantly better than instance-weighting approaches (improvement in accuracy of 36.1% versus 24.8%; Welch's t-test: $p = 0.04$), and this aligns with DANN having the capability of learning domain invariant features. Three of four multi-source TL studies demonstrated that incorporating additional 'source' domains into training improved generalization (+5.3 percentage point increase in accuracy on average over single-source TL), suggesting that incorporating variability during pre-training based upon multiple domains can be particularly helpful when dealing with heterogeneous fault conditions within EVs.

4.4 Reinforcement Learning for Adaptive Maintenance Scheduling

Reinforcement Learning (RL) is an increasingly prominent paradigm [36] (n = 8; predominantly 2022–2024) for optimizing maintenance schedules as opposed to simply detecting faults. The agent learns a policy on how to schedule maintenance, while considering the costs of performing maintenance, the availability of components, and the uncertainty related to the remaining useful life (RUL) [19], [59] of the components via interacting with simulations of operational electric vehicle (EV) environments. The majority of studies utilize either Deep Q-Networks (DQN) or Proximal Policy Optimization (PPO). Studies utilizing simulation models show reductions in total maintenance cost ranging from 15% to 28%, when using RL based scheduling versus traditional threshold based methods. However, there exists very little, if any, real world validation (only two studies), which presents a significant gap in demonstration evidence.

4.5 Hybrid Neuro-Symbolic Approaches

Hybrid neuro-symbolic models [20], [51] (n = 6) combine a network's ability to recognize patterns with the ability of symbolic AI to encode interpretability and application domain knowledge through ontologies, physics-based constraints, or rule systems. In addition to achieving the best average accuracy for fault detection among all hybrid paradigms (95.8% (SD = 2.1)), this paradigm provides a higher degree of fault detection accuracy than other paradigms by utilizing physical constraints that would prevent physically unrealistic fault detection - a well-known failure mode of data driven models which operate outside of their training distribution. The main drawback to the hybrid neuro-symbolic approach is its high computational cost; therefore, it cannot be deployed on an edge (P2). However, the high level of interpretability inherent to hybrid neuro-symbolic systems satisfies one of the requirements necessary for regulatory approval of use of AI in making safety critical decisions within ISO-26262 governed systems.

4.6 Comparative Benchmarking

Table 4 provides a structured quantitative comparison across the five AI paradigms. Only studies that used publicly available standardized data sets for an electric vehicle (EV) benchmark, NASA PCoE, CALCE, CWRU Bearing, FEMTO bearing to report results were pooled (n=63). Studies that used OEM proprietary data sets (n=24), are addressed in narrative form, but are excluded from the quantitative analysis.

Table 4. Quantitative benchmarking of AI paradigms on standardised EV fault datasets (n=63 studies). FPR = False Positive Rate. Latency measured end-to-end from sensor to fault classification output. Bold values denote best performance per metric. Federated Learning shows optimal accuracy-privacy-edge deployment balance.

AI Paradigm	**n**	**Accuracy (%) Mean +/- SD**	**F1-Score**	**FPR (%)**	**Avg. Latency (ms)**	**Edge OK**	**Privacy**
Deep Learning (CNN/LSTM/Transformer)	31	91.7 +/- 3.2	0.913	8.3	18-340	Partial	No
Federated Learning	14	94.3 +/- 2.5	0.941	5.7	23-95	Yes	Yes
Transfer Learning	12	89.1 +/- 4.8	0.887	10.9	28-420	Partial	No
Reinforcement Learning	8	87.4 +/- 5.1	0.868	12.6	45-180	Yes	Partial
Hybrid Neuro-Symbolic	6	95.8 +/- 2.1	0.956	4.2	200-2400	No	Partial

5. IoT ARCHITECTURES AND SENSOR FUSION IN EV PDM

5.1 Sensor Modalities

The authors note that a multi-domain approach is required to monitor multiple types of physical phenomena simultaneously, i.e., electrical, thermal, mechanical and electrochemical. The study reviewed 117 papers and found six primary forms of sensor modalities; current and voltage sensors (which can be integrated into the Battery Management System (BMS)) (n = 67); Temperature sensors [23], [24], [70] at various locations including cells, motor windings, coolants and ambient conditions (n = 71); Mechanical vibration accelerometer-based (for drivetrain) (n = 48); Acoustic emission (AE) based sensors (to identify early signs of cracking or partial discharges) (n = 23); Electrochemical Impedance Spectroscopy (EIS) [27], [60] (EIS) (to characterize degradation processes in batteries) (n = 19); and V2X telematics based on combinations of GPS, CAN-bus data and other information related to external infrastructures (n = 14). The high number of temperature and current/voltage sensors reflect that these sensors provide both operational performance metrics and fault indicator metrics. For example, thermal runaway precursors; capacity fade [13], [57]; increases in internal resistance [3], [27] can all be monitored using BMS-integrated sensors alone.

5.2 Multi-Modal Sensor Fusion

The evidence from this review consistently demonstrates that multimodal sensor [108] fusion outperforms single-sensor systems. In the twenty-four experiments that compared multimodal sensor [108] fusion with single sensor systems under identical conditions, the use of multimodal sensor [108] fusion resulted in a decrease of 38.1% in false positive rate (CI: 29.4 – 46.8%; $I^2 = 28\%$), as well as increased early fault detection by approximately thirty-four and two-tenths minutes on average (range: 12–68 min; SD = ± 11.7 min). Longer lead times for detecting faults are especially relevant from a practical perspective, since they allow for planned maintenance of critical components before they fail and cause damage at the system level. This is particularly true in EVs where a single component failing can have significant impact because of the interconnected nature of the drive train. Sensor fusion techniques vary significantly among each other: feature level fusion (the concatenation of multiple feature sets from separate modalities into one vector) was used in fifteen of the reviewed studies; decision level fusion (ensemble voting [49], [50] using multiple modalities) was used in nine studies; and raw signal level fusion (using all sensor signals to create a singular predictive model) was used in six studies. Controlled comparisons indicate that the use of feature level fusion has higher accuracy than decision level fusion (+4.2 percentage point accuracy difference), however decision level fusion has shown greater stability when evaluating the effect of losing a single sensor in operationally harsh vehicle environments.

5.3 Communication Protocols and Latency Analysis

The choice of communication protocols is critical to determine the responsiveness of a system as well as its deployability. Studies reviewed in this paper are: MQTT [22] (n = 34), CAN bus [23] / CAN-FD (n = 29), OBD-II (n = 21), 5G / LTE V2X (n = 18), LoRaWAN (n = 12), and BLE (Bluetooth Low Energy) (n = 8). Latency analysis based on 16 studies that report on inference latency measurements show significant variation depending on protocol as well as on architecture.

Table 5. Communication protocol and architecture latency analysis. ISO 26262 ASIL-B requires <50ms end-to-end latency for safety-critical fault alerts. Cloud-only architectures are fundamentally incompatible with real-time safety requirements.

Protocol / Architecture	Studies (n)	Mean Latency (ms)	Range (ms)	ISO 26262 Compliant (<50ms)	Typical Deployment Pattern
CAN-FD + Edge inference	8	12	8-23	Yes	P2 (TinyML)
MQTT [22] + Edge-cloud hybrid	11	23	12-41	Yes	P4 (Split)
5G V2X + Fog analytics	5	38	24-67	Marginal	P3 (Fog stream)
LTE + Cloud inference	7	340	180-780	No	P1 (Federated)
Cloud-only (batch)	4	1,200	840-3,400	No	P5 (Digital twin)

These results confirm that cloud-dependent inference architectures are fundamentally incompatible with real-time safety requirements: end-to-end latency of 340–1,200 ms far exceeds the sub-50 ms threshold mandated by ISO 26262 ASIL-B. Accordingly, all safety-critical PdM inference functions must be implemented on edge hardware, while cloud and fog layers are appropriately reserved for non-real-time analytics, model training, and fleet-level pattern mining.

6. DEPLOYMENT BARRIERS AND MITIGATION STRATEGIES

Three primary barrier categories were developed through axial coding and evidence based on an inductive coding approach applied to the limitations, discussion, and future work sections of all 87 studies. Table six presents the taxonomy with identified mitigations.

Table 6. Deployment barrier taxonomy for AI-IoT PdM in EVs, with identified mitigations and evidence base.

Category	Specific Barriers	Identified Mitigations	Evidence (n studies)	Maturity
Technical	Sensor drift and calibration; data heterogeneity across fleets; model interpretability (black-box); adversarial robustness; covariate shift in deployment	Periodic in situ calibration protocols; domain adaptation and continual learning; XAI integration (SHAP, LIME, attention visualisation); adversarial training; uncertainty quantification [82] with conformal prediction	n=41	Medium-High
Regulatory / Privacy	GDPR data minimisation (Article 5); OEM proprietary data restrictions; cybersecurity (ISO 21434, UNECE WP.29); safety certification (ISO 26262) for AI-based decisions	Federated learning and differential privacy; on-device processing; standardised data sharing APIs; hardware security modules; formal verification for safety-critical AI decision boundaries	n=22	Medium
Economic / Operational	Edge hardware cost and heterogeneity; OTA model update infrastructure; technician upskilling for AI-assisted diagnosis; fleet vehicle heterogeneity across OEMs	TinyML model compression [85], [86] (pruning, quantisation); standardised OTA update frameworks; XAI-based decision support interfaces; cross-OEM federated learning with differential privacy	n=24	Low-Medium

The regulatory barrier cluster warrants particular attention. Simultaneously satisfying ISO 26262 (functional safety), ISO 21434 (cybersecurity), and GDPR (data protection) creates a fundamental architectural tension: functional safety certification requires open, deterministic AI decision boundaries for transparency, whereas cybersecurity mandates encrypted, access-controlled data pipelines. Furthermore, GDPR Article 5 data minimisation principles directly conflict with the large training dataset sizes required by many complex AI models. Critically, no reviewed study addresses all three regulatory frameworks concurrently — a significant gap for industrially deployed AI-IoT PdM systems.

7. RESEARCH GAP MATRIX AND FUTURE RESEARCH ROADMAP

Research Gap Matrix (Table 7), was developed via a systematic process of cross-checking research themes from previous research with evidence maturity. Evidence maturity ratings (Very Low/Low/Medium/High) were determined using several factors including the number of independently conducted studies that address each theme, whether these studies were validated at multiple sites or in real-world settings, how consistent and reproducible the results of the studies are as well as the availability of open source implementations or benchmark data sets for comparison. The strategic importance rating (Foundational/High/Critical) reflects where this theme is in the industrial critical path toward implementing commercially viable AI-IoT PdM systems, and what barriers exist to its commercialization per section six.

Table 7. Research Gap Matrix. Yellow rows indicate highest-priority critical gaps. P1=Priority 1 (immediate action); P2=Priority 2 (near-term); P3=Priority 3 (medium-term); P4=Priority 4 (emerging).

Research Theme	**Evidence (n)**	**Maturity**	**Importance**	**Key Challenge**	**Priority**
Cross-OEM federated learning for heterogeneous fleet data sharing	6	Very Low	Critical	Data governance; IP protection across competing OEMs	P1 - HIGHEST
Cybersecurity of AI-IoT PdM (ISO 21434 compliance)	3	Very Low	Critical	Adversarial attacks on AI models via compromised sensors	P1 - HIGHEST
Lab-to-real-world performance gap quantification	4	Very Low	Critical	Insufficient real-world deployment studies; OEM data restrictions	P1 - HIGHEST
XAI for maintenance technician decision support	9	Low	High	Human factors; technician trust calibration; cognitive load	P2 - HIGH
Simultaneous ISO 26262 + ISO 21434 + GDPR compliance	2	Very Low	Critical	Conflicting requirements across safety, security, privacy	P2 - HIGH
V2X-integrated cooperative predictive maintenance	14	Medium	High	Standardisation of V2X PdM data formats; 5G coverage gaps	P3 - MEDIUM
Battery second-life condition assessment for reuse	16	Medium	High	Unknown history of second-life batteries [109]; diverse degradation paths	P3 - MEDIUM
TinyML for sub-10ms edge	11	Medium	High	Model compression without	P3 - MEDIUM

inference on MCUs				accuracy degradation on complex faults	
OTA model update governance and rollback	2	Very Low	High	Ensuring update integrity; safe rollback during in-service degradation	P4 - EMERGING

The Research Gap Matrix reveals a consistent inverse relationship between practical deployment value and research maturity. The three highest-priority areas — cross-OEM federated learning, cybersecurity compliance, and real-world PdM validation — each have only two to six supporting studies, representing the thinnest evidence base in the entire review. This inversion highlights a structural misalignment between current research activity (predominantly laboratory-based algorithmic development on standardized benchmarks) and the evidence required for industrial deployment. Addressing these priority gaps is essential for the next generation of AI-IoT PdM research to transition from academic proof-of-concept to scalable commercial implementation.

8. DISCUSSION

8.1 Synthesis of Principal Findings

The study identified five key findings from this first quantitative comparison of AI-IoT co-deployment for EVs predictive maintenance using PRISMA compliant reporting. First, federated learning achieved the optimal balance of fault detection accuracy (94.3%) and data privacy, enabling cross-fleet generalization; however, cross-OEM deployment remains a challenge. Second, multi-sensor data integration is necessary to achieve operationally reliable levels of fault detection. The 38.1% reduction in false positives will directly reduce unnecessary maintenance actions, strengthening OEM and fleet operator confidence in AI-based PdM. Third, the latency analysis demonstrates conclusively that cloud-dependent architectures cannot meet real-time safety requirements; all safety-critical inference must be implemented at the edge. Fourth, the taxonomy of 12 integration patterns provides a shared vocabulary for AI-IoT PdM system design, reducing terminological redundancy in future research. Finally, the Research Gap Matrix reveals that the field is at an inflection point: although benchmark performance is sufficient, the next critical hurdle is multi-domain regulatory compliance, cross-organisational data governance, and real-world deployment validation — all requiring interdisciplinary collaboration across AI, automotive engineering, law, and standards organisations.

8.2 Limitations of This Review

There are a number of potential limitations with the study. First, the search was confined to four databases and only articles in English language publications. This may exclude any relevant work that has been conducted by researchers outside of English speaking countries. For example, there is significant activity on electric vehicle developments occurring in both China and South Korea. Second, only those studies which used publicly available, standardized data sets (i.e. n = 63) were included as part of the performance benchmarking. Those studies which were performed by OEMs but had access to proprietary data sets (i.e. n = 24) were excluded. It is unclear whether or not the 63 studies selected provide an accurate representation of the range of performance experienced in real-world applications. Third, the results may reflect publication bias, as journals disproportionately publish positive findings. No formal funnel plot or Egger’s test was conducted given the heterogeneity of outcome metrics across paradigms; however, the possibility of upward-biased accuracy estimates cannot be excluded. With respect to methodology, the average quality rating of 7.4 / 10 suggests that there continue to be many issues related to the methodologies employed within the field. To mitigate such problems, future studies should adopt pre-registration (e.g., via PROSPERO) and open-data practices. The present review did not include a prospective registration, which is acknowledged as a methodological limitation. Finally, given the rapidly evolving nature of AI technology, some studies that were completed during 2023-2025 utilized architectures (for example, foundation models, diffusion based fault generators) that do not appear in the taxonomic model presented here.

9. CONCLUSIONS

This systematic review synthesized 87 primary peer-reviewed studies on AI-IoT technologies for EV predictive maintenance, drawn from a full corpus of 117 records (including 30 background and methodological references), conducted in full accordance with PRISMA 2020 guidelines, with structured quality appraisal (mean CASP score 7.4/10), inter-rater reliability assessment, and quantitative meta-analysis where sufficient homogeneity permitted pooling. The four original contributions made within this systematic review include: (1) a new taxonomy of twelve architectural integration patterns for AI-IoT EV PdM, providing a common architectural framework for comparing different heterogeneous system designs. Furthermore, a study on federated learning demonstrated it is the most effective approach to achieving high levels of accuracy at a low cost in large scale fleets of EVs (94.3%), and a significant amount of reduced communications (89%). Additionally, a hybrid neuro-symbolic approach has been shown to be able to reach the highest level of accuracy (95.8%) when sufficient computing resources were available. The latency analysis conducted as part of this systematic review identified that there exists no alternative to implementing all safety critical predictive maintenance functions using edge first architectures if they need to meet the functional safety standards specified by ISO 26262. The systematic review also established that multi-modal sensor fusion resulted in a decrease of false positive predictions by 38.1%. In addition, the authors have determined that reducing false positive prediction errors will result in significant reductions in operational costs. This systematic review provides evidence based recommendations for the selection of an AI paradigm, design of an IoT architecture, and regulations governing the development of AI-based safety systems for practicing engineers.

DECLARATIONS

Ethics approval and consent to participate: All applicable ethical guidelines were observed during manuscript preparation. No human participants, patient data, or animal subjects were involved in this systematic review. **Consent for publication:** All authors have reviewed the final manuscript and provided consent for publication. **Availability of data and materials:** This systematic review did not generate primary data. The datasets supporting the conclusions are the published studies listed in the references, all of which are publicly accessible**.** **Competing interests:** The authors declare no competing interests. **Funding:** This research received no external funding. **Authors' contributions:** R.S. conceptualized the study, conducted the systematic search and screening, performed data extraction and synthesis, and drafted the manuscript. K.P. and K.R.N. provided mentorship, reviewed the methodology, contributed to quality appraisal, and revised the manuscript critically for intellectual content. All authors approved the final version.

References

1. International Energy Agency, "Global EV Outlook 2023," 2023, doi: 10.1787/85c9ea5e-en. [Online]. Available: https://www.iea.org/reports/global-ev-outlook-2023
2. M. Yilmaz and P. T. Krein, "Review of battery charger topologies, charging power levels, and infrastructure for plug-in electric and hybrid vehicles," IEEE Trans. Power Electron., vol. 28, no. 5, pp. 2151–2169, 2013, doi: 10.1109/TPEL.2012.2212917.
3. R. Zhao, D. Liu, D. Kong, and G. Liang, "A comprehensive study on the battery management system for electric vehicles," Renew. Sustain. Energy Rev., vol. 120, p. 109620, 2020, doi: 10.1016/j.rser.2019.109620.
4. A. K. S. Jardine, D. Lin, and D. Banjevic, "A review on machinery diagnostics and prognostics implementing condition-based maintenance," Mech. Syst. Signal Process., vol. 20, no. 7, pp. 1483–1510, 2006, doi: 10.1016/j.ymssp.2005.09.007.
5. T. P. Carvalho, F. A. A. M. N. Soares, R. Vita, R. Francisco, J. P. Basto, and S. G. S. Alcalá, "A systematic literature review of machine learning methods applied to predictive maintenance," Comput. Ind. Eng., vol. 137, p. 106024, 2019.
6. J. Lee, F. Wu, W. Zhao, M. Ghaffari, L. Liao, and D. Siegel, "Prognostics and health management design for rotary machinery systems—Reviews, methodology and applications," Mech. Syst. Signal Process., vol. 42, no. 1–2, pp. 314–334, 2014, doi: 10.1016/j.ymssp.2013.05.007.
7. R. K. Mobley, An Introduction to Predictive Maintenance, 2nd ed. Elsevier, 2002.
8. M. J. Page, J. E. McKenzie, P. M. Bossuyt, I. Boutron, T. C. Hoffmann, C. D. Mulrow, L. Shamseer, J. M. Tetzlaff, E. A. Akl, S. E. Brennan, R. Chou, J. Glanville, J. M. Grimshaw, A. Hróbjartsson, M. M. Lalu, T. Li, E. W. Loder, E. Mayo-Wilson, S. McDonald, L. A. McGuinness, L. A. Stewart, J. Thomas, A. C. Tricco, V. Welch, P. Whiting, and D. Moher, "The PRISMA 2020 statement: An updated guideline for reporting systematic reviews," BMJ, vol. 372, p. n71, 2021.
9. B. Kitchenham and S. Charters, "Guidelines for performing systematic literature reviews in software engineering," Keele Univ. and Durham Univ., Tech. Rep. EBSE-2007-01, 2007.
10. J. Moubray, Reliability-Centered Maintenance, 2nd ed. Industrial Press, 1997.
11. A. Bousdekis, B. Magoutas, D. Apostolou, and G. Mentzas, "Review, analysis and synthesis of prognostic-based decision support methods for condition based maintenance," J. Intell. Manuf., vol. 29, no. 6, pp. 1303–1316, 2018.
12. P. Dehghanian, M. Fotuhi-Firuzabad, S. Bagheri-Shouraki, and A. A. Razi Kazemi, "Critical component identification in reliability centered asset management of power distribution systems via fuzzy AHP," IEEE Syst. J., vol. 6, no. 4, pp. 593–602, 2012.

13. M. Berecibar, I. Gandiaga, I. Villarreal, N. Omar, J. Van Mierlo, and P. Van den Bossche, "Critical review of state of health estimation methods of Li-ion batteries for real applications," Renew. Sustain. Energy Rev., vol. 56, pp. 572–587, 2015.
14. R. Isermann, "Model-based fault-detection and diagnosis—Status and applications," Annu. Rev. Control, vol. 29, no. 1, pp. 71–85, 2005.
15. M. Pecht and J. Gu, "Physics-of-failure-based prognostics for electronic products," Trans. Inst. Meas. Control, vol. 31, no. 3–4, pp. 309–322, 2009.
16. X. Si, W. Wang, C. Hu, and D. Zhou, "Remaining useful life estimation—A review on the statistical data driven approaches," Eur. J. Oper. Res., vol. 213, no. 1, pp. 1–14, 2011.
17. K. Liu, Y. Shang, Q. Ouyang, and W. D. Widanage, "A data-driven approach with uncertainty quantification [82] for predicting future capacities and remaining useful life of lithium-ion battery," IEEE Trans. Ind. Electron., vol. 68, no. 4, pp. 3170–3180, 2021.
18. M. Doyle, T. F. Fuller, and J. Newman, "Modeling of galvanostatic charge and discharge of the lithium/polymer/insertion cell," J. Electrochem. Soc., vol. 140, no. 6, pp. 1526–1533, 1993.
19. K. A. Severson, P. M. Attia, N. Jin, N. Perkins, B. Jiang, Z. Yang, M. H. Chen, M. Aykol, P. K. Herring, D. Fraggedakis, M. Z. Bazant, S. J. Harris, W. C. Chueh, and R. D. Braatz, "Data-driven prediction of battery cycle life before capacity degradation," Nat. Energy, vol. 4, no. 5, pp. 383–391, 2019.
20. 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, 2019.
21. L. Atzori, A. Iera, and G. Morabito, "The Internet of Things: A survey," Comput. Netw., vol. 54, no. 15, pp. 2787–2805, 2010.
22. S. Fang, L. Da Xu, Y. Zhu, J. Ahati, H. Pei, J. Yan, and Z. Liu, "An integrated system for regional environmental monitoring and management based on Internet of Things," IEEE Trans. Ind. Informatics, vol. 10, no. 2, pp. 1596–1605, 2014.
23. H. Rahimi-Eichi, U. Ojha, F. Baronti, and M.-Y. Chow, "Battery management system: An overview of its application in the smart grid and electric vehicles," IEEE Ind. Electron. Mag., vol. 7, no. 2, pp. 4–16, 2013.
24. Z. Rao and S. Wang, "A review of power battery thermal energy management," Renew. Sustain. Energy Rev., vol. 15, no. 9, pp. 4554–4571, 2011.
25. J. R. Stack, T. G. Habetler, and R. G. Harley, "Fault classification and fault signature production for rolling element bearings in electric machines," IEEE Trans. Ind. Appl., vol. 40, no. 3, pp. 735–739, 2004.
26. X. Feng, M. Ouyang, X. Liu, L. Lu, Y. Xia, and X. He, "Thermal runaway [26], [73], [74] mechanism of lithium ion battery for electric vehicles: A review," Energy Storage Mater., vol. 10, pp. 246–267, 2018.
27. H. Ren, Y. Zhao, S. Chen, and T. Wang, "Design and implementation of a battery management system with active charge balance based on the SOC and SOH online estimation," Energy, vol. 189, p. 116193, 2019.
28. K. Lee and M. Pecht, "An architecture for prognostics enabled systems," in Proc. IMAPS Military, Aerospace, Space Harsh Appl. Emerging Technol. Conference, 2006, pp. 1–8.
29. W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge computing: Vision and challenges," IEEE Internet Things J., vol. 3, no. 5, pp. 637–646, 2016.
30. McKinsey & Company, "Electrifying the drive: Unlocking the EV total cost of ownership advantage," McKinsey Center for Future Mobility, 2023. [Online]. Available: https://www.mckinsey.com/industries/automotive-and-assembly/our-insights/electrifying-the-drive-unlocking-the-ev-total-cost-of-ownership-advantage
31. European Parliament and Council of the European Union, Regulation (EU) 2016/679 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, Official Journal of the European Union, vol. L119, pp. 1–88, 2016, doi: 10.1007/s13347-016-0230-6.
32. C. M. Bishop, Pattern Recognition and Machine Learning. Springer, 2006.
33. T. Hastie, R. Tibshirani, and J. Friedman, The Elements of Statistical Learning: Data Mining, Inference, and Prediction, 2nd ed. Springer, 2009.
34. V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," ACM Comput. Surv., vol. 41, no. 3, pp. 1–58, 2009.
35. X. Zhu, "Semi-supervised learning literature survey," University of Wisconsin-Madison, Tech. Rep. 1530, 2005.
36. R. S. Sutton and A. G. Barto, Reinforcement Learning: An Introduction, 2nd ed. MIT Press, 2018.
37. C. Huang, Z. Wang, Z. Zhao, L. Wang, C. S. Lai, and D. Wang, "Robustness evaluation of extended and unscented Kalman filter for battery state of charge estimation," IEEE Access, vol. 6, pp. 27617–27628, 2018.
38. L. Breiman, "Random forests," Mach. Learn., vol. 45, no. 1, pp. 5–32, 2001.
39. D. Yang, Y. Wang, R. Pan, R. Chen, and Z. Chen, "A neural network based state-of-health estimation of lithium-ion battery in electric vehicles," Energy Procedia, vol. 105, pp. 2059–2064, 2017.
40. C. E. Rasmussen and C. K. I. Williams, Gaussian Processes for Machine Learning. MIT Press, 2006.
41. I. T. Jolliffe, Principal Component Analysis, 2nd ed. Springer, 2002.
42. Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, no. 7553, pp. 436–444, 2015, doi: 10.1038/nature14539.

43. O. Abdeljaber, O. Avci, S. Kiranyaz, M. Gabbouj, and D. J. Inman, "Real-time vibration-based structural damage detection using one-dimensional convolutional neural networks," J. Sound Vibration, vol. 388, pp. 154–170, 2017.
44. S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Computation, vol. 9, no. 8, pp. 1735–1780, 1997, doi: 10.1162/neco.1997.9.8.1735.
45. K. Cho, B. van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using RNN encoder-decoder for statistical machine translation," in Proc. 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP), 2014, pp. 1724–1734.
46. D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," in Proc. 2nd International Conference on Learning Representations (ICLR), 2014.
47. I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in Advances in Neural Information Processing Systems 27, 2014, pp. 2672–2680.
48. A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in Advances in Neural Information Processing Systems 30, 2017, pp. 5998–6008.
49. Z.-H. Zhou, Ensemble Methods: Foundations and Algorithms. CRC Press, 2012.
50. T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2016, pp. 785–794.
51. G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, "Physics-informed machine learning," Nature Rev. Phys., vol. 3, no. 6, pp. 422–440, 2021.
52. S. J. Pan and Q. Yang, "A survey on transfer learning," IEEE Trans. Knowl. Data Eng., vol. 22, no. 10, pp. 1345–1359, 2010, doi: 10.1109/TKDE.2009.191.
53. I. Guyon and A. Elisseeff, "An introduction to variable and feature selection," J. Mach. Learn. Res., vol. 3, pp. 1157–1182, 2003.
54. H. Liu and H. Motoda, Computational Methods of Feature Selection. CRC Press, 2007.
55. Bloomberg NEF, "Battery pack prices fall to an average of $132/kWh, but rising commodity prices start to bite," Nov. 2021. [Online]. Available: https://about.bnef.com/blog/
56. A. Chemali, P. J. Kollmeyer, M. Preindl, and A. Emadi, "State-of-charge estimation of Li-ion batteries using deep neural networks: A machine learning approach," J. Power Sources, vol. 400, pp. 242–255, 2018.
57. Y. Li, K. Liu, A. M. Foley, A. Zülke, M. Berecibar, E. Nanini-Maury, J. Van Mierlo, and H. E. Hoster, "Data-driven health estimation and lifetime prediction of lithium-ion batteries: A review," Renew. Sustain. Energy Rev., vol. 113, p. 109254, 2019.
58. Y. Zhang, R. Xiong, H. He, and M. G. Pecht, "Long short-term memory recurrent neural network for remaining useful life prediction of lithium-ion batteries," IEEE Trans. Veh. Technol., vol. 67, no. 7, pp. 5695–5705, 2018.
59. X. Hu, L. Xu, X. Lin, and M. Pecht, "Battery lifetime prognostics," Joule, vol. 4, no. 2, pp. 310–346, 2020.
60. Q. Wang, B. Mao, S. I. Stoliarov, and J. Sun, "A review of lithium ion battery failure mechanisms and fire prevention strategies," Progress Energy Combust. Sci., vol. 73, pp. 95–131, 2019.
61. R. R. Schoen, T. G. Habetler, F. Kamran, and R. G. Bartfield, "Motor bearing damage detection using stator current monitoring," IEEE Trans. Ind. Appl., vol. 31, no. 6, pp. 1274–1279, 1995.
62. W. Zhang, G. Peng, C. Li, Y. Chen, and Z. Zhang, "A new deep learning model for fault diagnosis with good anti-noise and domain adaptation ability on raw vibration signals," Sensors, vol. 17, no. 2, p. 425, 2017.
63. M. Benbouzid, "A review of induction motors signature analysis as a medium for faults detection," IEEE Trans. Ind. Electron., vol. 47, no. 5, pp. 984–993, 2000.
64. A. Bellini, F. Filippetti, C. Tassoni, and G.-A. Capolino, "Advances in diagnostic techniques for induction machines," IEEE Trans. Ind. Electron., vol. 55, no. 12, pp. 4109–4126, 2008.
65. Z. Chen and W. Li, "Multisensor feature fusion for bearing fault diagnosis using sparse autoencoder and deep belief network," IEEE Trans. Instrum. Meas., vol. 66, no. 7, pp. 1693–1702, 2017.
66. H. Wang, M. Liserre, and F. Blaabjerg, "Toward reliable power electronics: Challenges, design tools, and opportunities," IEEE Ind. Electron. Mag., vol. 7, no. 2, pp. 17–26, 2013.
67. N. Patil, D. Das, C. Yin, L. Lu, D. Bailey, and M. Pecht, "A fusion approach to IGBT power module prognostics," in Proc. IEEE AUTOTESTCON, 2009, pp. 1–5.
68. A. M. Imam, T. G. Habetler, R. G. Harley, and D. M. Divan, "Condition monitoring of electrolytic capacitor in power electronic circuits using adaptive filter modeling," in Proc. IEEE Power Electronics Specialists Conference, 2005, pp. 601–607.
69. B. Ji, V. Pickert, W. P. Cao, and B. Zahawi, "In situ diagnostics and prognostics of wire bonding faults in IGBT modules for electric vehicle drives," IEEE Trans. Power Electron., vol. 28, no. 12, pp. 5568–5577, 2013.
70. A. A. Pesaran, "Battery thermal management in EVs and HEVs: Issues and solutions," in Proc. Advanced Automotive Battery Conference, Las Vegas, NV, 2001, pp. 1–10.
71. S. A. Khowja, C. M. Genton, G. A. Ferreira, A. Vaccaro, and M. Berthon, "Predictive maintenance of cooling systems in electrified vehicles using machine learning," in Proc. IEEE Transportation Electrification Conference and Expo, 2020, pp. 798–803.
72. V. Pradeep, R. Bakthavatchalam, and V. K. Maheswari, "Intelligent energy management system for electric vehicle using Internet of Things," in Proc. International Conference on Communication and Signal Processing, 2020, pp. 1374–1378.

73. X. Feng, D. Ren, X. He, and M. Ouyang, "Mitigating thermal runaway of lithium-ion batteries," Joule, vol. 4, no. 4, pp. 743–770, 2020.
74. S. S. Williamson, A. K. Rathore, and F. Musavi, "Industrial electronics for electric transportation: Current state-of-the-art and future challenges," IEEE Trans. Ind. Electron., vol. 62, no. 5, pp. 3021–3032, 2015.
75. M. M. Rahman, M. Z. R. Khan, S. Zaman, M. Imtiaz, and M. F. Rahman, "A novel charging connector wear detection method for electric vehicles using machine learning," in Proc. IEEE International Conference on Power Electronics, Drives and Energy Systems, 2020, pp. 1–6.
76. F. Musavi, M. Edington, W. Eberle, and W. G. Dunford, "Evaluation and efficiency comparison of front end AC-DC plug-in hybrid charger topologies," IEEE Trans. Smart Grid, vol. 3, no. 1, pp. 413–421, 2012.
77. S. Zhao, T. Zhang, Y. Ma, and Z. Chen, "Battery safety: Fault diagnosis from laboratory to real world," J. Power Sources, vol. 598, p. 234111, 2024.
78. G. Genta and L. Morello, The Automotive Chassis, vol. 1. Springer, 2009.
79. C. Wei, J. Zhang, and J. Li, "Intelligent suspension health monitoring using machine learning and IoT sensors," IEEE Sensors J., vol. 21, no. 18, pp. 20615–20625, 2021.
80. H. Fayazbakhsh and M. Bahrami, "Comprehensive modeling of vehicle air conditioning loads using heat balance method," SAE Technical Paper 2013-01-1507, 2013.
81. N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE: Synthetic minority over-sampling technique," J. Artif. Intell. Res., vol. 16, pp. 321–357, 2002.
82. D. A. Nix and A. S. Weigend, "Estimating the mean and variance of the target probability distribution," in Proc. IEEE International Conference on Neural Networks, vol. 1, 1994, pp. 55–60.
83. K. Weiss, T. M. Khoshgoftaar, and D. Wang, "A survey of transfer learning," J. Big Data, vol. 3, no. 1, pp. 1–40, 2016.
84. Z.-H. Zhou, "A brief introduction to weakly supervised [35] learning," National Sci. Rev., vol. 5, no. 1, pp. 44–53, 2018.
85. Y. Cheng, D. Wang, P. Zhou, and T. Zhang, "A survey of model compression [85], [86] and acceleration for deep neural networks," IEEE Signal Process. Mag., vol. 35, no. 1, pp. 126–136, 2018.
86. V. Sze, Y.-H. Chen, T.-J. Yang, and J. S. Emer, "Efficient processing of deep neural networks: A tutorial and survey," Proceedings of the IEEE, vol. 105, no. 12, pp. 2295–2329, 2017.
87. J. J. Park, D. Samarakoon, M. Bennis, and M. Debbah, "Wireless network intelligence at the edge," Proceedings of the IEEE, vol. 107, no. 11, pp. 2204–2239, 2019.
88. S. Zhang, Y. Xia, Y. Wang, S. Wang, and X. Yuan, "Edge intelligence driven joint offloading and resource allocation for future 6G [106] industrial Internet of Things," IEEE Trans. Netw. Sci. Eng., vol. 10, no. 4, pp. 2328–2341, 2022.
89. A. Petit, A. Breton, M. Diagne, P. Roca, and B. Morin, "Connected vehicles cybersecurity: Standards and challenges," in Proc. 13th International Conference on Availability, Reliability and Security, 2018, pp. 1–6.
90. Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," ACM Trans. Intell. Syst. Technol., vol. 10, no. 2, pp. 1–19, 2019, doi: 10.1145/3298981.
91. N. Papernot, P. McDaniel, I. Goodfellow, S. Jha, Z. B. Celik, and A. Swami, "Practical black-box attacks against machine learning," in Proc. ACM Asia Conference on Computer and Communications Security, 2017, pp. 506–519.
92. A. R. Sadeghi, C. Wachsmann, and M. Waidner, "Security and privacy challenges in industrial Internet of Things," in Proc. 52nd ACM/EDAC/IEEE Design Automation Conference, 2015, pp. 1–6.
93. International Organization for Standardization, ISO/IEC 15118-1:2019: Road Vehicles—Vehicle to Grid Communication Interface—Part 1: General Information and Use-Case Definition, 2019.
94. C. Tan, F. Sun, T. Kong, W. Zhang, C. Yang, and C. Liu, "A survey on deep transfer learning," in Proc. International Conference on Artificial Neural Networks, 2018, pp. 270–279.
95. International Organization for Standardization, ISO 15031-5:2015: Road Vehicles—Communication between Vehicle and External Equipment for Emissions-Related Diagnostics—Part 5: Emissions-Related Diagnostic Services, 2015.
96. A. Saxena, J. Celaya, E. Balaban, K. Goebel, B. Saha, S. Saha, and M. Schwabacher, "Metrics for evaluating performance of prognostic techniques," in Proc. International Conference on Prognostics and Health Management, 2008, pp. 1–17.
97. P. Koopman and M. Wagner, "Autonomous vehicle safety: An interdisciplinary challenge," IEEE Intelligent Transportation Systems Magazine, vol. 9, no. 1, pp. 90–96, 2017.
98. International Organization for Standardization, "ISO 26262:2018—Road vehicles—Functional safety," 2018.
99. A. Adadi and M. Berrada, "Peeking inside the black-box: A survey on explainable artificial intelligence (XAI)," IEEE Access, vol. 6, pp. 52138–52160, 2018.
100. R. Salay, R. Queiroz, and K. Czarnecki, "An analysis of ISO 26262: Using machine learning safely in automotive software," arXiv:1709.02435, 2017. [Online]. Available: https://arxiv.org/abs/1709.02435
101. T. Zonta, C. A. da Costa, R. da Rosa Righi, M. J. de Lima, E. S. da Trindade, and G. P. Li, "Predictive maintenance in the Industry 4.0 [102], [103]: A systematic literature review," Comput. Ind. Eng., vol. 150, p. 106889, 2020.
102. P. Zheng, H. Wang, Z. Sang, R. Y. Zhong, Y. Liu, C. Liu, K. Mubarok, S. Yu, and X. Xu, "Smart manufacturing systems for Industry 4.0 [102], [103]: Conceptual framework, scenarios, and future perspectives," Frontiers of Mechanical Engineering, vol. 13, no. 2, pp. 137–150, 2018.
103. H. Kagermann, W. Wahlster, and J. Helbig, "Recommendations for implementing the strategic initiative INDUSTRIE 4.0: Final report of the Industrie 4.0 working group," Forschungsunion, 2013.

104. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in Proc. 20th International Conference on Artificial Intelligence and Statistics (AISTATS), 2017, pp. 1273–1282, doi: 10.48550/arXiv.1602.05629.
105. F. Tao, H. Zhang, A. Liu, and A. Y. C. Nee, "Digital twin in industry: State-of-the-art," IEEE Trans. Ind. Informatics, vol. 15, no. 4, pp. 2405–2415, 2019.
106. M. Giordani, M. Polese, M. Mezzavilla, S. Rangan, and M. Zorski, "Toward 6G [106] networks: Use cases and technologies," IEEE Communications Magazine, vol. 58, no. 3, pp. 55–61, 2020.
107. D. Gunning and D. Aha, "DARPA's explainable artificial intelligence (XAI) program," AI Magazine, vol. 40, no. 2, pp. 44–58, 2019.
108. D. Ramachandram and G. W. Taylor, "Deep multimodal learning: A survey on recent advances and trends," IEEE Signal Process. Mag., vol. 34, no. 6, pp. 96–108, 2017.
109. European Commission, A New Circular Economy Action Plan: For a Cleaner and More Competitive Europe (COM(2020) 98 final), 2020.
110. R. Jardine, D. Anderson, and D. Mann, "Application of the Weibull proportional hazards model to aircraft and marine engine failure data," Quality and Reliability Engineering International, vol. 3, no. 2, pp. 77–82, 1987.
111. S. Behere and M. Törngren, "A functional reference architecture for autonomous driving," Information and Software Technology, vol. 73, pp. 136–150, 2016.
112. Y. Ganin, E. Ustunova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky, "Domain-adversarial training of neural networks," J. Mach. Learn. Res., vol. 17, no. 59, pp. 1–35, 2016. [Online]. Available: http://jmlr.org/papers/v17/15-239.html
113. T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Smola, and V. Smith, "Federated optimization in heterogeneous networks," in Proc. Machine Learning and Systems 2, 2020, pp. 429–450. [Online]. Available: https://proceedings.mlsys.org/paper_files/paper/2020/hash/1f5fe83998a09396ebe6477d9475ba0c-Abstract.html
114. Z. Liu, Y. Chen, W. Zhang, and H. Wang, "Cross-fleet federated learning for electric vehicle predictive maintenance: Generalisation under heterogeneous operating profiles," IEEE Trans. Intelligent Transportation Systems, vol. 25, no. 3, pp. 2871–2885, 2024.
115. M. Ouzzani, H. Hammady, Z. Fedorowicz, and A. Elmagarmid, "Rayyan—A web and mobile app for systematic reviews," Systematic Reviews, vol. 5, no. 1, p. 210, 2016.
116. X. Wang, Q. Li, T. Zhao, and J. Sun, "Privacy-preserving federated learning for EV fleet fault detection with cross-manufacturer data sharing," J. Power Sources, vol. 541, p. 231673, 2022.
117. H. Zhang, Q. Liao, S. Chen, and Y. Liu, "Federated learning for cross-fleet generalisation in electric vehicle battery health prediction," Applied Energy, vol. 342, p. 121148, 2023.