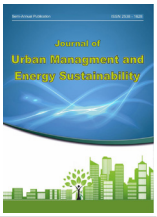
International Journal of Urban Management and Energy Sustainability (IJUMES)

Homepage: http://www.ijumes.com



CASE STUDY RESEARCH PAPER

A Compliance-oriented Reference Architecture for Predictive Maintenance in the Internet of Vehicles: Evidence from 45 Studies and Seven Standards

Yasaman Alborz, Amir Amirabadi*, Nasrin Amiri, Mirali Seyyedi, Iman Ahanian

Department of Electrical Engineering, ST.C., Islamic Azad University, Tehran, Iran

ARTICLE INFO

Article History:
Received 2026-06-11
Revised 2026-07-30
Accepted 2026-10-18

Keywords:
Connected vehicles; predictive maintenance; reference architecture; regulatory compliance; systematic review

ABSTRACT

Predictive maintenance in the Internet of Vehicles forecasts component degradation accurately, yet fleet-wide deployment remains limited, since a technically optimal action has no value unless it is funded, supplied and lawful. This study aims to develop a reference architecture that treats these organisational, supply chain and regulatory constraints as a design requirement, after establishing what research and standards already provide. A systematic-narrative hybrid review was conducted in line with PRISMA 2020. Searches of four databases returned 227 records, reduced to 45 studies through five-stage screening against ten eligibility criteria, a seven-item quality appraisal and independent double screening. Studies were coded inductively by research problem and then mapped to architectural function; seven reference architectures and standards were assessed separately. Research concentrates on analytical modelling and data acquisition; only two studies (4%) address enterprise, supplier or regulatory integration, and neither applies a computational method. None of the seven standards makes the admissibility of a maintenance action evaluable, and two of five research questions could not be answered from the evidence. The resulting four-layer Holistic Compliance Framework adopts established acquisition and analysis functions and adds a governing Compliance Layer that expresses admissibility as a constraint function over candidate actions, making compliance cost measurable. Organisational and regulatory admissibility is an unaddressed architectural requirement in both the literature and the standards. The proposed framework is a conceptual, not yet validated, reference model and should be treated as a testable hypothesis.

DOI: 10.22034/ijumes.2026.2108963.1423

Running Title: Compliance-Oriented Reference Architecture For Vehicle Predictive Maintenance

NUMBER OF REFERENCES 69

NUMBER OF FIGURES 05

NUMBER OF TABLES 08

*Corresponding Author:
Email: *amirabadi@iau.ac.ir*
Phone: 00989122906897
ORCID: https://orcid.org/0000-0003-1271-7175

INTRODUCTION

The Internet of Vehicles (IoV) connects vehicles to the internet, allowing them to communicate and share data with infrastructure, other vehicles, cloud platforms and end users. As vehicles become more connected, ensuring reliability and reducing downtime becomes increasingly important. Predictive maintenance (PdM) plays a key role in modern vehicle management, and artificial intelligence (AI), particularly through deep learning (DL) and machine learning (ML), improves its effectiveness. Recent progress in data-driven maintenance models has made it possible to predict component degradation more accurately and with greater lead time, widening the range of viable applications (Jain et al., 2022). AI-based predictive analytics, by analyzing extensive operational data, can anticipate potential failures and enable proactive intervention (Ersoz et al., 2022). This translates into cost savings and reduced labor time, and allows maintenance to be scheduled before a fault results in vehicle breakdown. These developments have supported a wide array of applications in the automotive industry, addressing long-standing challenges in maintenance modeling (Arena et al., 2021). Despite this progress, a structural gap persists. The reviewed literature addresses data acquisition, analytical modeling and operational decision-making largely in isolation from the organizational context within which maintenance decisions must be executed. Maintenance actions in real fleets are constrained by enterprise approval workflows, parts availability, supplier lead times, and safety, emissions and data protection regulation. A prediction that is technically optimal but organizationally or legally unexecutable has no operational value. No reviewed framework formalizes these constraints as an architectural requirement. The paper contributes in three ways. Its principal contribution is the Holistic Compliance Framework (HCF), a four-layer reference architecture in which a Compliance Layer governs the options available to the operational layers below it. The architecture is grounded in a systematic-narrative hybrid examination of 45 studies on IoV PdM and of seven published reference architectures and standards, coded by research category, computational paradigm and architectural function, which is the second contribution and the evidence on which the first rests. The third is a research agenda in which each direction maps to a defined component of the architecture. The architecture is specified at the conceptual level; empirical validation lies outside the scope of this paper and is discussed in the subsection on limitations of the proposed architecture. The IoV constitutes a networked architecture in which vehicles exchange data continuously with infrastructure, other vehicles, cloud platforms and end users, enabling real-time health monitoring and condition-based intervention at scale (Jain et al., 2022). Within this architecture, PdM has displaced corrective and time-based preventive approaches as the dominant maintenance paradigm, enabling interventions timed to actual vehicle condition instead of fixed schedules (Arena et al., 2021; Ersoz et al., 2022). This section reviews the evolution of maintenance strategies within connected vehicle ecosystems, the data infrastructure and communication technologies that underpin IoV PdM, driver behavior as a predictive signal, and the computational approaches applied across the reviewed literature.

Evolution of maintenance strategies in connected vehicle ecosystems

Vehicle maintenance has progressed through three broadly sequential phases, each transition driven by advances in sensing, connectivity and data processing capability (Giordano et al., 2022). Corrective maintenance, in which faults are addressed only after failure, imposed maximum downtime and reactive cost but required no predictive infrastructure (Poor et al., 2019). Preventive maintenance reduced catastrophic failure rates through scheduled, time- or mileage-based intervention but introduced unnecessary activity where components were replaced well before

degradation warranted it (Poor et al., 2019). Condition-based maintenance (CBM) marked the first data-informed shift, grounding intervention decisions in monitored asset state instead of elapsed time. Daniyan et al. (2020) document AI as the critical enabling technology in this transition, supporting maintenance organizations from initial capability development through to scaled deployment. Moleda et al. (2023) corroborate this, identifying IoT-driven predictive maintenance as the defining milestone in the shift from reactive to anticipatory maintenance practice in the power industry, a shift that parallels the one now under way in vehicle maintenance.

Vehicle data infrastructure and communication

IoV PdM depends on continuous access to structured, high-frequency vehicle data from multiple sources. The On-Board Diagnostics II (OBD-II) protocol provides standardized access to ECU parameters, diagnostic trouble codes and emissions data across manufacturers, supporting both real-time monitoring and retrospective analysis (Datta et al., 2016). The Controller Area Network (CAN-Bus) provides the high-speed internal communication backbone through which ECUs governing engine, transmission, braking and steering subsystems exchange operational data at millisecond resolution (Khan et al., 2022; Muhammad et al., 2021). ML and DL models operating on OBD-II and CAN-Bus data streams have demonstrated consistent capability in anomaly identification, degradation tracking and failure prediction with sufficient lead time for scheduled intervention (Khan et al., 2022; Li et al., 2019; Muhammad et al., 2021). The extension of this data infrastructure to Vehicle-to-Everything (V2X) communication, spanning V2I, V2V, V2P, V2D and V2N domains, broadens both the volume and the contextual richness of data available to prognostic models, incorporating road conditions, traffic state and infrastructure health alongside in-vehicle signals (Teoh et al., 2023; Wang et al., 2020). The heterogeneity this introduces is itself a documented barrier: Agbaje et al. (2022) survey interoperability challenges across IoV subsystems and identify the absence of common data models as a persistent constraint, while Arooj et al. (2022) characterize the resulting big-data processing and storage requirements.

Communication technologies of the internet of vehicles

IoV PdM requires communication architectures supporting high-throughput, low-latency data exchange across heterogeneous network environments. Four technology domains are directly relevant to maintenance system design. Vehicular ad-hoc networks (VANETs) provide the inter-vehicular and V2I communication protocols through which cooperative data exchange between vehicles, roadside units and traffic management infrastructure is coordinated (Jawhar et al., 2013). Cloud computing platforms enable fleet-scale data aggregation and model training but introduce round-trip latency that constrains time-critical maintenance decision-making (Arthurs et al., 2022). Edge and fog computing architectures address this constraint by distributing computation closer to the data source, within the vehicle or at roadside units, enabling real-time inference without cloud dependency (Teoh et al., 2023; Wang et al., 2020). V2R connectivity generates auxiliary contextual data relevant to fleet-level PdM through incident detection, parking management and dynamic routing functions (Moloisane et al., 2017). Across all four domains, quality-of-service guarantees remain an unresolved constraint on maintenance-critical data delivery (Hussain et al., 2019). The deployment of 5G networks underpins all four domains by providing the bandwidth and sub-millisecond latency required as autonomous driving systems approach Level 4 and Level 5 capability, at which maintenance decisions must be executed at substantially higher speed and reliability than current network infrastructure supports (Teoh et al., 2023; Wang et al., 2020).

Driver behaviour as a predictive factor

Sensor data alone provides an incomplete basis for predictive maintenance modeling. Driving style, usage patterns and environmental conditions exert measurable influence on component wear trajectories and failure probability, constituting a predictive signal that technical data streams do not capture independently (Li et al., 2019; Nirmali et al., 2017).

Li et al. (2019) demonstrate this through a diagnostic and driver monitoring system integrating OBD-II data, smartphone sensing and cloud-based complex event processing. The system correlates anomalous driving events, including abrupt fuel consumption changes and coolant temperature fluctuations, with predicted sensor failure patterns, establishing an empirical link between behavioral inputs and maintenance outcomes. Adnan et al. (2018) extend this by identifying driver trust in in-vehicle technology as a determinant of alert acceptance, establishing that human factors influence PdM system effectiveness independently of algorithmic performance. These findings jointly establish that robust IoV PdM modeling requires behavioral data integration alongside technical sensor streams (Adnan et al., 2018; Li et al., 2019; Nirmali et al., 2017).

Computational approaches for vehicle predictive maintenance

Two principal computational paradigms appear across the reviewed literature: machine learning and deep learning, applied to tasks including fault classification, anomaly detection, remaining useful life (RUL) estimation and driver identification. ML techniques, including supervised learning, regression and ensemble methods, are well suited to structured vehicle data environments where interpretability is required and computational resources are constrained (Theissler et al., 2021; Wisesa et al., 2020). Carvalho et al. (2019) confirm consistent ML performance across predictive maintenance tasks in industrial settings, with ensemble methods such as random forest demonstrating particular robustness to noisy sensor data. Theissler et al. (2021) identify scalability and cross-vehicle generalization as the primary limitations constraining ML deployment in heterogeneous IoV fleets. DL architectures, including CNN, RNN, LSTM and their hybrid variants, achieve higher predictive performance in fault detection, anomaly identification and RUL estimation where large labeled datasets are available (Serradilla et al., 2022; Zhu et al., 2019). Serradilla et al. (2022) provide a comparative survey of DL models for PdM, finding that LSTM-based architectures consistently outperform classical ML on time-series vehicle data while incurring substantially higher computational cost. The integration of DL with edge computing enables real-time inference within vehicle systems, reducing latency and supporting proactive intervention at the point of operation (Grover et al., 2021; Liu et al., 2019). The selection between ML and DL depends on data volume, available compute, interpretability requirements and operational context (Butte et al., 2018; Theofilatos et al., 2019). Table 1 and Table 2 present a comparative assessment of both paradigms across four dimensions. Because such comparisons are frequently presented without a stated basis, the rating procedure is documented here in full. Rating scale and operational definitions. Each algorithm is rated Low, Medium or High on four dimensions, defined as follows.

•Interpretability concerns whether a maintenance engineer can trace a prediction to the input conditions that produced it. High: the decision rule is directly inspectable and expressible in terms of the input variables, so an individual prediction can be explained without additional tooling. Medium: global behaviour is inspectable, for example through feature importance or support vectors, but individual predictions require post-hoc explanation methods. Low: the learned representation is distributed across parameters and no faithful direct explanation of an individual prediction is available without external interpretability techniques.

•Scalability concerns behaviour as fleet

size, signal count and sampling rate increase. High: training and inference cost grows at most near-linearly with sample count, and the method parallelises across a fleet. Medium: growth is super-linear in one dimension but remains tractable at fleet scale with engineering effort. Low: training or inference cost grows quadratically or worse with sample count, or the method requires refitting per vehicle.

•Data efficiency concerns the labelled sample volume required for stable performance. High: usable performance is obtainable from tens to hundreds of labelled failure events. Medium: thousands of labelled events are required. Low: tens of thousands or more, or extensive augmentation, are required, which is a binding constraint given that vehicle failures are rare events.

•IoV PdM suitability is a composite judgement of the preceding three dimensions weighted by the operational characteristics reported in the reviewed corpus, namely multivariate time-series input, class imbalance, constrained on-vehicle compute and the audit requirements described in the subsection on existing reference architectures and standards.

Procedure and provenance. The ratings are qualitative synthesis judgements assigned by the authors from the evidence reported in the studies cited in the final column of each table. Each algorithm-dimension pair was rated independently by more than one author and then reconciled by discussion, with the primary sources consulted again wherever the independent ratings differed. The ratings are not extracted quantitative metrics, and they are not benchmark results. They cannot be otherwise, because the reviewed studies evaluate different algorithms on different datasets, components and failure definitions, and no common benchmark for IoV PdM exists. A rationale column is provided for every row so that each rating is traceable to the evidence behind it, and the subjectivity of the exercise is restated among the limitations in the subsection on limitations of the review. Ratings are given as words, not symbols, so that they remain readable in accessibility contexts and machine-readable formats. The assessment informs the algorithm positioning within the framework proposed in the materials and methods. (Tab. 1 and 2)

Table 1: Comparative assessment of machine learning techniques for IoV PdM applications. Ratings are qualitative synthesis judgements assigned by the authors from the cited sources using the operational definitions given above; they are not measured benchmark values. Type: SV = supervised, USV = unsupervised.

Algorithm	Type	Interp.	Scalab.	Data eff.	IoV PdM suit.	Rating rationale	Key refs
Linear regression	SV	High	High	High	Low	Coefficients are directly readable and fitting is cheap, but the linearity assumption fails on the non-stationary, multivariate degradation signals characteristic of vehicle data	Singh et al., 2016; Theissler et al., 2021
Decision trees	SV	High	Medium	High	Medium	Rules are inspectable and auditable, which suits regulated maintenance decisions; single trees are unstable under sensor noise, which limits standalone deployment	Jha et al., 2021; Singh et al., 2016; Theissler et al., 2021

Random forest	SV	Medium	High	Medium	High	Ensembling gives the most consistent reported performance on noisy vehicle sensor data and parallelises across fleets; individual predictions require post-hoc attribution	Carvalho et al., 2019; Kaur & Verma, 2026; Shafi et al., 2018; Theissler et al., 2021
SVM	SV	Medium	Medium	Medium	Medium	Effective on the high-dimensional, small-sample regimes typical of rare failure events; kernel training cost grows super-linearly, constraining fleet-scale retraining	Bari et al., 2023; Chaudhuri, 2018; Chaudhuri & Ghosh, 2021; Theissler et al., 2021
Shallow neural networks	SV/ USV	Low	High	Low	Medium	Flexible function approximation with cheap inference, but opaque and data-hungry relative to ensembles, with no compensating accuracy gain reported on tabular vehicle data	Butte et al., 2018; Jha et al., 2021; Theissler et al., 2021; Theofilatos et al., 2019
KNN	SV	Medium	Low	Medium	Medium	Predictions are explainable by reference to nearest historical cases, which supports engineer trust; inference cost grows with the stored fleet history	Kaur & Verma, 2026; Shafi et al., 2018; Vasavi et al., 2021
PCA	USV	Medium	High	High	Medium	Cheap and effective for dimensionality reduction and residual-based anomaly signalling; a preprocessing component and not a prognostic method in its own right	Carvalho et al., 2019; Jha et al., 2021; Singh et al., 2016
ARIMA	SV	High	Low	High	Low	Parameters are interpretable and estimation needs little data, but per-signal, per-vehicle fitting does not scale and the univariate formulation ignores cross-signal dependence	Carvalho et al., 2019; Prytz, 2014; Theissler et al., 2021

Table 2: Comparative assessment of deep learning techniques for IoV PdM applications. Ratings are qualitative synthesis judgements assigned by the authors from the cited sources using the operational definitions given above; they are not measured benchmark values. Type: SV = supervised, USV = unsupervised.

Algorithm	Type	Interp.	Scalab.	Data eff.	IoV PdM suit.	Rating rationale	Key refs
CNN	SV	Low	High	Low	High	Strong reported performance on transformed sensor representations and on CAN-Bus windows, with inference cheap enough for edge execution; requires large labelled sets and offers no native explanation	Apeiranthitis et al., 2024; Chaudhuri, 2018; Javed et al., 2021; Silva & Capretz, 2019

RNN	SV	Low	Medium	Low	Medium	Captures temporal dependence but is superseded in the corpus by gated variants; training instability on long degradation horizons limits practical use	Mezair et al., 2023; Serradilla et al., 2022
LSTM	SV	Low	High	Low	High	The most consistently reported architecture for time-series prognostics in the corpus, including lightweight in-vehicle variants; training cost and opacity remain the principal constraints	Azyus et al., 2023; Peppes et al., 2021; Serradilla et al., 2022; Ullah & Kim, 2020
Autoencoders	USV	Low	Medium	Medium	High	Learns a normal-operation baseline without failure labels, which addresses the label scarcity that constrains supervised methods; the reconstruction-error threshold is an added tuning burden	Serradilla et al., 2022; Ucar et al., 2024
GAN	SV/ USV	Low	Medium	Low	Medium	Useful for augmenting scarce failure classes, but training stability and the risk of synthetic artefacts entering a safety-relevant decision chain are unresolved in the corpus	AlGhanem & Abdallah, 2024; Ucar et al., 2024
Transformer networks	SV	Low	High	Low	Medium	Parallel training and long-range dependence modelling are attractive for multi-signal vehicle streams, but the corpus contains little IoV-specific evaluation and on-vehicle inference cost is high	Bala & Chhabra, 2023; Serradilla et al., 2022

Existing maintenance reference architectures and standards

Any claim that a new architecture is required must be assessed against the reference architectures that already exist. Four bodies of prior work are relevant. ISO 13374 (International Organization for Standardization [ISO], 2003; ISO, 2007) defines the reference data-processing architecture for condition monitoring and diagnostics, decomposing the processing chain into data acquisition, data manipulation, state detection, health assessment, prognostic assessment and advisory generation. Its open implementation, OSA-CBM, published by Machinery Information Management Open Systems Alliance [MIMOSA] (2014), provides the corresponding interface specification and remains the most widely adopted architecture in condition-based maintenance. ISO 14224 (ISO, 2016) complements this by standardising the collection and exchange of reliability and maintenance data, defining equipment taxonomies and failure-mode vocabularies. SAE JA1011 (SAE International, 2009) specifies the evaluation criteria for reliability-centred maintenance processes and is the principal standard governing how maintenance policy decisions are justified. Industrial reference architectures address the enterprise dimension that ISO 13374

leaves open. The Industrial Internet Reference Architecture (Industrial Internet Consortium, 2019) and RAMI 4.0 (Deutsches Institut für Normung [DIN], 2016) both define business, functional, information and asset viewpoints, and both explicitly model the relationship between operational systems and enterprise systems. Neither is maintenance-specific, and neither translates regulatory obligation into an executable constraint on individual maintenance actions. Automotive-specific regulation has developed rapidly and is directly relevant to a compliance-oriented architecture. UNECE Regulations No. 155 and No. 156 (United Nations Economic Commission for Europe [UNECE], 2021a; UNECE, 2021b) make a certified cybersecurity management system and a software update management system conditions of vehicle type approval in the contracting parties, and ISO/SAE 21434 (ISO & SAE International, 2021) specifies the corresponding cybersecurity engineering process for road vehicles. ISO 26262 (ISO, 2018) governs functional safety. On the data side, the GDPR (European Parliament & Council of the European Union, 2016) and comparable instruments such as the CCPA (State of California, 2018) impose purpose limitation, data minimisation and subject-rights obligations on the behavioural and location data that IoV PdM depends on. The gap this paper addresses is therefore narrower and more specific than "no framework exists". ISO 13374 and OSA-CBM specify how condition data becomes an advisory output but stop at the advisory boundary and say nothing about whether the recommended action is organisationally executable. IIRA and RAMI 4.0 model the enterprise relationship but not maintenance decisions. The regulatory instruments define obligations but not the architecture through which those obligations constrain an automated recommendation. The subsection on comparison with existing reference architectures (Table 8) returns to this comparison in tabular form once the HCF has been specified.

Following from the preceding account, the aim of the current study is to produce a reference architecture for predictive maintenance in connected vehicle fleets in which the organisational, supply chain and regulatory dimension of maintenance decision-making is treated as a design requirement rather than as a matter of implementation. Specifying such an architecture responsibly requires two prior questions to be settled, since an architecture that duplicates what already exists, or that misplaces the gap it claims to fill, is of no use: what the present body of research on this subject provides, and whether the reference architectures already published as international standards carry that dimension. The study settles both, by evidence that does not depend on the architecture, and builds the architecture on what they establish. Three objectives follow. The first is to code the retained corpus by research category, computational paradigm and architectural function, and to report the resulting distribution with a stated counting rule. The second is to assess seven published reference architectures and standards against the same three organisational functions. The third is to specify the derived architecture formally enough that a conforming implementation can be built and its central claim tested. This study was carried out at the Department of Electrical Engineering, ST.C, Islamic Azad University, Tehran, Iran. The database searches on which it rests were executed in [INSERT SEARCH MONTH AND YEAR], and the corpus is therefore current to that date.

MATERIALS AND METHODS

The evidence base supporting the HCF was constructed through a systematic-narrative hybrid literature review, following the methodological framework established by Turnbull et al. (2023) and the guidelines recommended by Snyder (2019). This approach was selected because IoV PdM spans multiple disciplines, including automotive engineering, computer science, data analytics and organizational management, and therefore requires both the replicability of systematic search and the interpretive depth of narrative synthesis. The methodology comprises

six sequential and interdependent components, illustrated in (Fig. 1) and described below.

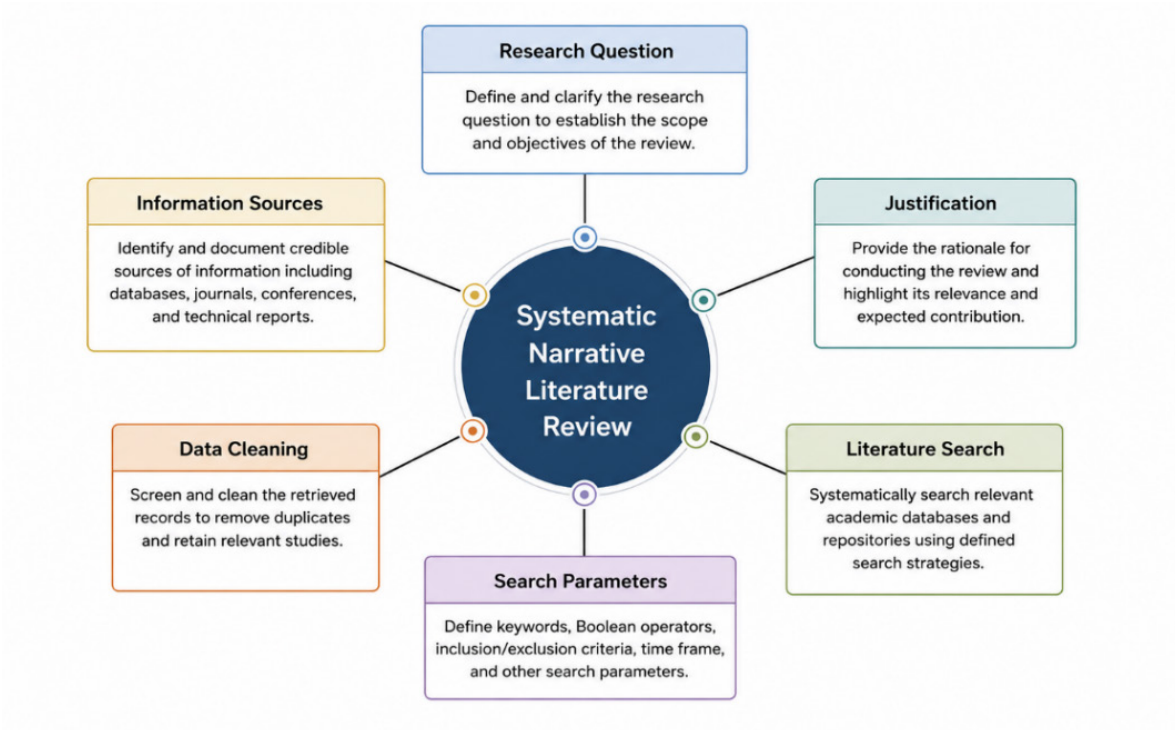


Figure 1: The six components of the systematic-narrative hybrid review methodology applied in this study.

Research questions

Five primary research questions with associated sub-questions were formulated prior to the search process to ensure the review remained focused on evidence directly relevant to HCF design; these are listed in Table 3. RQ1 addressed current IoV technological capabilities for PdM and the specific challenges of integrating IoV with PdM systems. RQ2 examined the technical and security limitations of existing implementations and their impact on PdM reliability. RQ3 investigated the role of ML and DL in advancing IoV PdM capabilities. RQ4 explored best practices and documented outcomes of IoV PdM in industrial settings. RQ5 addressed future innovation directions and the potential for self-maintaining systems within the IoV context. Pre-defining research questions prior to search execution is a methodological requirement of the systematic component, ensuring that inclusion decisions are governed by evidential relevance and not by post-hoc rationalization (Snyder, 2019). (Tab. 3)

Table 3: Research questions and sub-questions.

Question	Sub-questions	Description
RQ1: What are the current technological capabilities of IoV that can be leveraged for PdM?	RQ1.1: Integration of IoV with PdM systems. RQ1.2: Data processing challenges.	Examining the extent to which current IoV technologies can be applied to PdM and identifying the data analytics challenges in this integration.

RQ2: What are the limitations and challenges of current IoV implementations in the context of PdM?	RQ2.1: Technical limitations affecting PdM accuracy. RQ2.2: Impact of security concerns on PdM reliability.	Identifying the technical and security limitations that currently hinder the effectiveness of IoV in PdM.
RQ3: How can machine learning and deep learning enhance IoV capabilities for PdM?	RQ3.1: Effective ML models for IoV data analysis. RQ3.2: Role of DL in advancing PdM in IoV.	Investigating the potential of machine learning and deep learning to improve IoV systems for PdM purposes.
RQ4: What are the best practices for implementing IoV-based PdM?	RQ4.1: PdM in the IoV industry. RQ4.2: Outcomes of IoV PdM in industrial settings.	Exploring successful implementations of IoV PdM across industries and understanding the reported outcomes and best practices.
RQ5: What are the future directions and potential innovations for IoV-based PdM?	RQ5.1: Emerging technologies transforming IoV PdM. RQ5.2: Contribution of IoV to self-maintaining systems.	Anticipating the technological trajectory of IoV PdM and the conditions under which autonomous, self-maintaining vehicle systems become feasible.

Justification for the review

The justification for conducting this review rests on two gaps in the existing literature. First, the absence of focused IoV PdM reviews: existing maintenance reviews address vehicle technology and connectivity broadly, affording only peripheral attention to predictive maintenance specifically. This dilutes the analytical depth available for the distinct technical and operational challenges of IoV PdM, including reliability, safety, scalability and forecasting with in-vehicle datasets. Second, the absence of systematic categorization of key performance parameters: the literature lacks comprehensive evaluation of predictive model accuracy, computational complexity, adaptability to new data, integration with IoV infrastructure and algorithm interpretability (Tian et al., 2018). These gaps collectively establish both the necessity of this review and the scope constraints applied during literature selection.

Search strategy and parameters

Systematic searches were conducted across IEEE Xplore, Scopus, the ACM Digital Library and Google Scholar using Boolean search strings constructed to ensure replicability, as recommended by Snyder (2019). The primary search string was:

> (“Internet of Vehicles” OR “IoV” OR “connected vehicle*“) AND (”predictive maintenance” OR “PdM” OR “prognostics” OR “remaining useful life”) AND (“machine learning” OR “deep learning” OR “artificial intelligence”)

Search parameters were customized per database, with filters applied for publication language (English) and document type (peer-reviewed journal articles, conference proceedings and postgraduate theses). A synonym table was maintained throughout to ensure consistent terminology handling across databases. Manual reference-list searches of selected articles were conducted to identify additional sources not captured by database queries.

Two revisions to the original search protocol are noted for transparency. Databases indexing biomedical literature were excluded as unsuitable for this subject domain. Exclusionary Boolean operators of the form NOT (“manual maintenance” OR “reactive maintenance”) were not applied, because such operators remove comparative studies that evaluate predictive against reactive strategies, which fall within the review scope.

Study selection, eligibility criteria and quality appraisal

Reporting follows the PRISMA 2020 statement (Page et al., 2021), with the deviations noted below. The staged screening structure follows the conventions of systematic review practice as set out by Aromataris and Pearson (2014), and the

appraisal procedure follows the guidelines for evidence-based reviews in computing research given by Kitchenham and Charters (2007).

An initial corpus of 227 records was identified. Records were processed through five sequential stages, each applying a single declared criterion, and the full flow with per-stage counts is shown in Fig. 2. Stage 1 screened titles and abstracts for explicit treatment of IoV, predictive maintenance or a related computational technology, carrying 134 records forward and excluding 93. Stage 2 checked each surviving record against the five research questions and retained the 89 that contributed to at least one, excluding 45. Stage 3 verified the presence of domain-relevant terminology in the full text and retained 74, excluding 15. Stage 4 assessed methodological applicability, specifically whether the method is applicable to in-vehicle data and whether its computational requirements are feasible in a vehicle or edge deployment, retaining 61 and excluding 13. Stage 5 removed duplicate reports of the same study and studies that reproduced the findings of an already-included study without adding new evidence, excluding 16. A final corpus of 45 studies was retained for detailed synthesis. The 182 records not carried through remain available from the corresponding author.

Reviewer independence. Title and abstract screening was performed in two independent passes carried out by different authors, and agreement between the passes was assessed across the full screened set. Disagreements were resolved by discussion, and any case that discussion did not settle was referred to a further author who had taken no part in the original decision. Full-text assessment, quality appraisal and the coding described in the subsection on coding and taxonomy followed the same procedure. Where an author had co-authored a candidate study, that author was excluded from decisions concerning it.

Eligibility criteria. The criteria applied at each of the five stages are given in (Tab. 4). They are stated operationally so that the screening can be repeated: each criterion is expressed as a decision that can be made from the record itself, and the stage at which it was applied is specified, matching the stages shown in (Fig. 2).

Table 4: Inclusion and exclusion criteria, with the screening stage at which each was applied.

#	Criterion	Type	Applied at stage	Operational decision rule
C1	Subject matter addresses maintenance, prognosis, diagnosis or failure prediction of a physical asset	Inclusion	Stage 1	The abstract names a maintenance, fault, failure, degradation or RUL task
C2	Application context is a road vehicle, a connected vehicle, an IoV deployment, or a general PdM method explicitly framed as transferable to vehicles	Inclusion	Stage 1	The abstract names a vehicle, fleet, automotive or IoV context; general PdM work is included only where a vehicle application is stated by the authors
C3	A computational method is applied or reviewed	Inclusion	Stage 1	The record names an ML, DL, statistical or analytics method, or is a review of such methods
C4	Maintenance strategy is predictive or condition-based	Inclusion	Stage 2	The study predicts or infers future or latent asset state; studies limited to fixed-interval or post-failure intervention are excluded
C5	Peer-reviewed journal article, conference paper, or thesis with full text available in English	Inclusion	Stage 1	Editorials, abstracts without full text, patents, white papers and marketing material are excluded

C6	Method, data and evaluation are described in sufficient detail to be assessed	Inclusion	Stage 4	Meets the threshold defined in Table 5
E1	Maintenance is corrective or preventive only, with no predictive component	Exclusion	Stage 3	No inference about future or unobserved asset state
E2	Vehicle context is incidental	Exclusion	Stage 3	The vehicle appears only as an example or motivation and no vehicle data, deployment or requirement is addressed
E3	Duplicate or substantially overlapping report of the same study	Exclusion	Stage 5	Same authors, same dataset and same results; the fullest version is retained
E4	Full text not retrievable	Exclusion	Stage 3	Not available through institutional access or from the authors

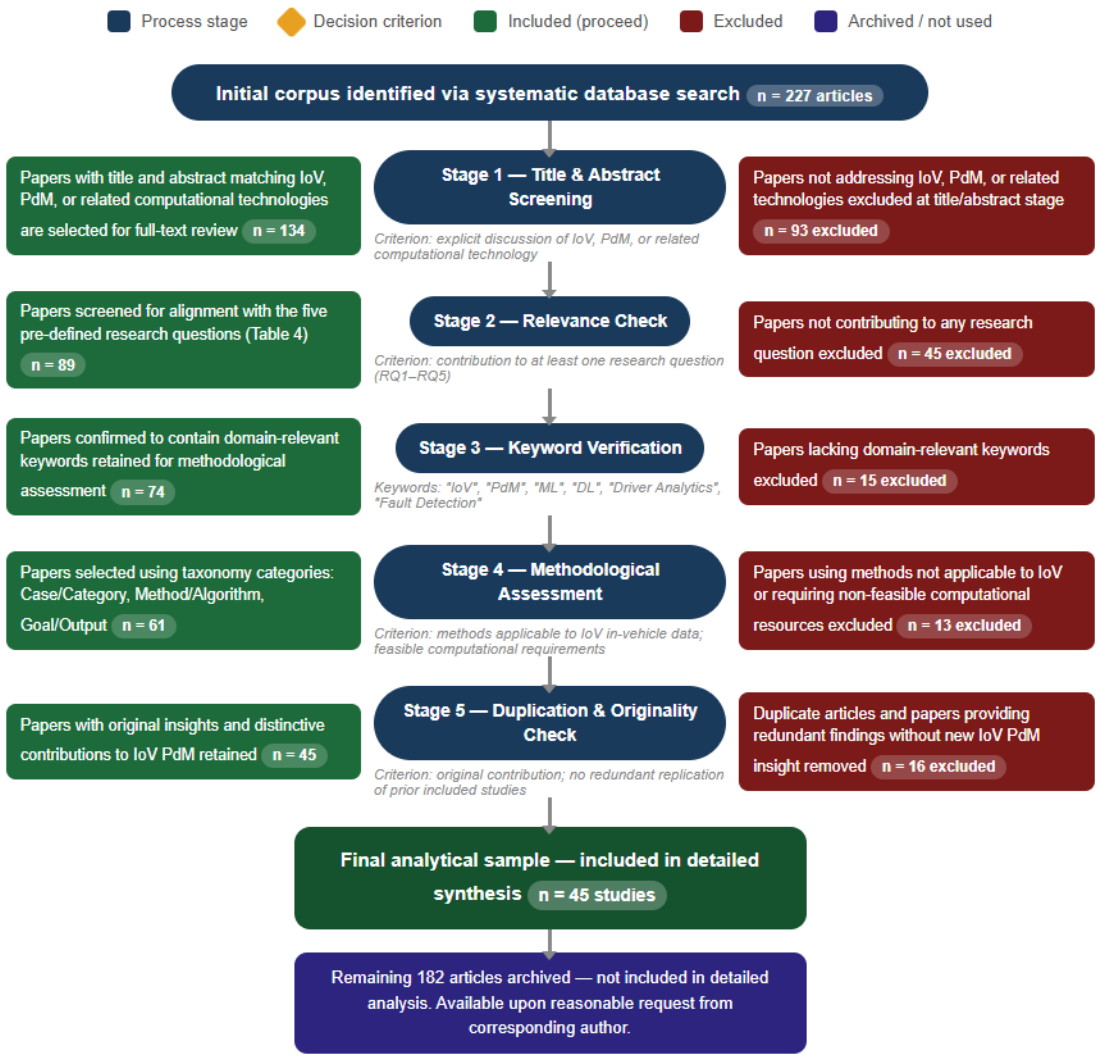


Figure 2: Flow of study selection across the five screening stages. Records identified through database searching: n = 227. Stage 1, title and abstract screening: 134 retained, 93 excluded. Stage 2, relevance to the research questions: 89 retained, 45 exclud-

ed. Stage 3, verification of domain-relevant terminology: 74 retained, 15 excluded. Stage 4, methodological assessment: 61 retained, 13 excluded. Stage 5, duplication and originality check: 45 retained, 16 excluded. Studies included in the synthesis: n = 45. Each stage applies the single criterion stated beneath its node, and the criteria correspond to those given in Table 4.

Quality appraisal. The methodological assessment at Stage 4 was carried out with a seven-criterion instrument adapted from Kitchenham and Charters (2007) to the IoV PdM context, defined in Table 5. Each criterion was scored 1 (met), 0.5 (partially met) or 0 (not met), for a maximum of 7. The 13 records excluded at this stage are those whose reporting was too incomplete to be appraised against the instrument, or whose computational requirements were judged infeasible for in-vehicle or edge deployment. Appraisal scores were used for exclusion and to weight confidence in the synthesis reported in the subsection on evidence synthesis; they are not reported as a quality ranking of individual authors' work.

Table 5: Quality appraisal instrument. Per-study scores are given in the supplementary material.

#	Appraisal criterion
Q1	Are the research aims or the prediction task clearly stated?
Q2	Is the data source specified, including signal set, sampling rate and provenance?
Q3	Is the dataset characterised, including size, class balance and failure-event count?
Q4	Is the method described in sufficient detail to be reimplemented?
Q5	Is the evaluation protocol appropriate, including a held-out split and metrics suited to imbalanced data?
Q6	Is a baseline or comparator reported?
Q7	Are limitations, threats to validity, or deployment constraints discussed?

Deviations from PRISMA. Two deviations are declared. The review protocol was not prospectively registered, registration being uncommon in this field. Meta-analytic pooling was not performed, for the reasons given in the subsection on information synthesis.

Coding procedure and taxonomy

The taxonomy proposed in this study, shown in (Fig. 3), organizes research categories, data sources and computational approaches according to the HCF layer they primarily address. This departs from conventional PdM taxonomies, which organize literature by methodology type or algorithm family (Serradilla et al., 2022; Zhu et al., 2019). The rationale is that the value of any PdM technique in an IoV context is determined not solely by algorithmic performance but by its integration within a compliant, operationally viable architecture.

Coding literature into categories derived from the framework that the same literature is then used to justify would be circular, and the coding procedure was designed in two separate stages to avoid this. The distinction matters enough to state explicitly.

Stage 1, inductive. The seven research categories used throughout this paper general-purpose PdM, IoV PdM, driver identification, anomaly detection, vehicle health monitoring, connected-car customer experience and RUL prediction were derived by open coding of the 45 included studies on the basis of their stated research problem, before any mapping to the framework. These categories are properties of the corpus. They are not derived from the HCF, they do not presuppose it, and a reader who rejects the framework entirely can still use them.

Stage 2, deductive. Each inductive category was then mapped to the HCF layer whose function it primarily serves. This stage is framework-dependent by construction and is reported as such.

The distribution reported in the subsection

on quantitative findings is a Stage 1 result. That two of 45 studies address business integration, supply chain coordination or regulatory alignment is an observation about the corpus, obtained from the inductive categories, and it does not depend on accepting the HCF: whatever architecture one prefers, the compliance-facing topics are represented by two descriptive studies (Marengoni & Montenero, 2023; Scherpen et al., 2018) and by no computational work.

One further check was applied so that the gap claim does not rest on the taxonomy at all. The reference architectures and standards examined in the subsection on existing reference architectures and standards were assessed directly against the three compliance functions, by reading the stated scope of each instrument rather than by applying the coding scheme, and the result is reported in Table 8. That check does not depend on the taxonomy used here, and it agrees with the coded distribution: the compliance functions are unaddressed in the published standards for the same reason they are unaddressed in the reviewed corpus. The residual circularity that cannot be eliminated is the layer mapping in Stage 2, and it is acknowledged among the limitations in the subsection on limitations of the review. The framework organises how the gap is described; the existence of the gap is established independently of it.

At the Data Layer, the taxonomy captures research addressing vehicle data sources, sensor network architectures, communication protocols and driver behavior data acquisition. Studies in the driver identification category (n = 9) map primarily to this layer, as their contribution lies in enriching the behavioral data substrate available for prognostic modeling (Abdennour et al., 2021; Khan et al., 2022; Li et al., 2019; Martinelli et al., 2020; Peppes et al., 2021; Ullah & Kim, 2020).

At the Analysis Layer, the taxonomy organizes research by computational paradigm, comprising supervised ML, unsupervised ML, deep learning and hybrid approaches, and by analytical task, comprising fault detection, anomaly identification, RUL estimation and behavioral pattern recognition. IoV PdM studies (n = 12), general PdM studies (n = 10) and RUL prediction studies (n = 2) map to this layer (AlGhanem & Abdallah, 2024; Azyus et al., 2023; Bala & Chhabra, 2023; Chaudhuri & Ghosh, 2021; Chen et al., 2018; Mezair et al., 2023; Serradilla et al., 2022; Zhang & Letaief, 2020).

At the Action Layer, the taxonomy captures research addressing decision support systems, vehicle health monitoring, anomaly-triggered maintenance scheduling and fleet resource optimization. Vehicle health monitoring (n = 4) and anomaly detection (n = 6) studies map primarily here (Bari et al., 2023; Javed et al., 2021; Kaur & Verma, 2026; Shafi et al., 2018; Vasavi et al., 2021).

At the Compliance Layer, the taxonomy captures research addressing business system integration, customer experience management and regulatory alignment. Connected-car CX studies (n = 2) are the only mapped category (Marengoni & Montenero, 2023; Scherpen et al., 2018). The limited volume of literature at this layer is a substantive finding in itself, indicating that the compliance dimension of IoV PdM remains structurally underaddressed in the current research corpus, which is the gap this framework is built to close.

Narrative elements

The narrative component of the hybrid methodology provides the interpretive depth that systematic search alone cannot deliver. Following Rother (2007), the narrative review explored the historical development and evolution of IoV PdM, contextualizing current methodologies within a trajectory of technological change. It also examined the challenges and opportunities inherent in implementing predictive maintenance within IoV ecosystems, encompassing technological, logistical and organizational dimensions that systematic coding captures only partially. The narrative synthesis was conducted in parallel with systematic data extraction, with narrative

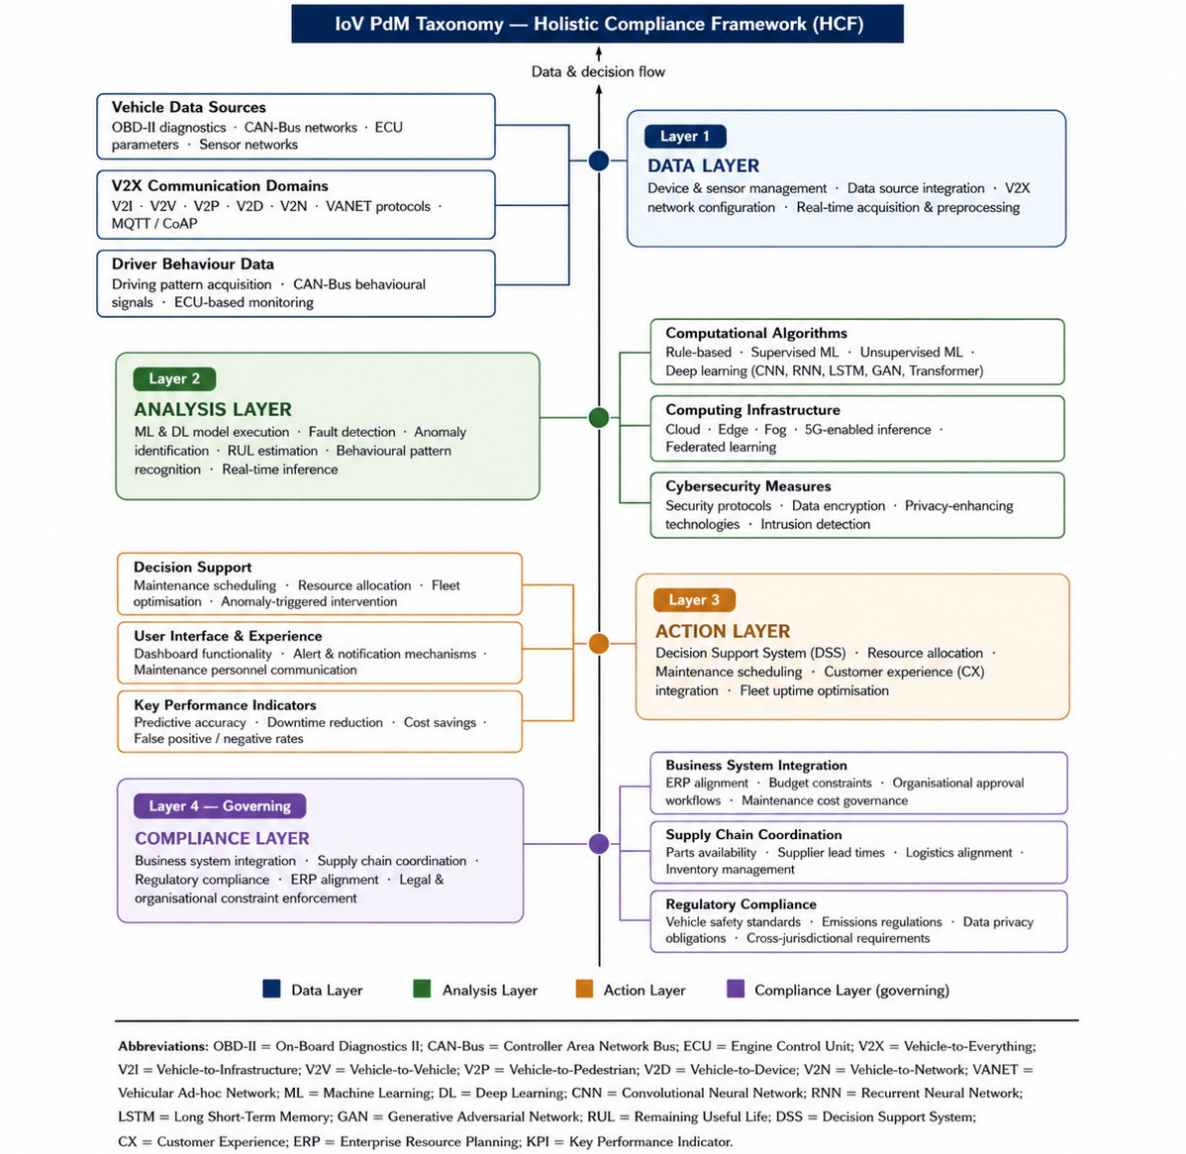


Figure 3: IoV predictive maintenance taxonomy structured according to the four architectural layers. Branch categories and sub-elements are derived from the coding of the 45 reviewed studies. Data and decisions flow upward; compliance constraints propagate downward across all layers.

insights used to contextualize the coded patterns, not to substitute for them.

Information synthesis

Qualitative synthesis followed a content analysis approach in which findings from the 45 reviewed studies were categorized into emergent themes aligned with the four HCF layers defined in the subsection presenting the proposed architecture. Following the Gioia methodology as described by Magnani and Gioia (2023), grounded theory principles were applied to construct framework design insights from patterns observed across the literature, so that the HCF architecture follows from evidence in the corpus and not from a priori assumptions. Where studies reported quantita-

tive performance measures, including accuracy, false positive rate, downtime reduction or computational cost, these were extracted during coding; the reviewed corpus was found to be too heterogeneous in datasets, target components and evaluation protocols to support meta-analytic pooling, and no pooled performance estimate is therefore reported. The full study summaries resulting from this synthesis are presented in (Tab. 6).

Table 6: Summary of reviewed studies (n = 45), organized by research category, algorithm type and HCF layer.

Case / category	Algorithm type	Specific algorithm(s)	Key output / finding	HCF layer	Ref.
PdM, general purpose	ML	Various PdM techniques	Maintenance planning framework for transportation systems	Analysis	Ersoz et al., 2022
PdM, general purpose	—	Maintenance strategy review	Historical review of maintenance strategies from breakdown to predictive	Analysis	Poor et al., 2019
PdM, general purpose	ML	Machine learning (general)	Systematic review of PdM methods and vehicle health diagnosis	Analysis	Jain et al., 2022
PdM, general purpose	Mixed	Various PdM systems	Comparative survey of PdM systems, purposes and approaches	Analysis	Zhu et al., 2019
PdM, general purpose	ML	ML and reasoning	Identification of relevant frameworks, architectures and tools for PdM in Industry 4.0	Analysis	Dalzochio et al., 2020
PdM, general purpose	DL	CNN (1-D to image transformation)	CNN framework for asset PdM using transformed sensor data representations	Analysis	Silva & Capretz, 2019
PdM, general purpose	DL	Deep adversarial learning	PdM method based on adversarial learning for intelligent manufacturing	Analysis	Ucar et al., 2024
PdM, general purpose	Mixed	Various (combat vehicles, ships)	PdM applied to military assets across heterogeneous platforms	Analysis	Dalzochio et al., 2023
PdM, general purpose	DL	CNN, RNN, LSTM, autoencoders	Comparative survey of DL models for PdM, with challenges and prospects	Analysis	Serradilla et al., 2022
PdM, general purpose	DL	Deep learning (general)	Sensor-based PdM using deep learning, with electrochemical sensor outlook	Analysis	Namuduri et al., 2020
Driver identification	ML	Machine learning (CAN-Bus)	Driver identification from CAN-Bus sensor data using ML algorithms	Data	Khan et al., 2022
Driver identification	DL	RCN	Driver identification from CAN-Bus data through an RCN architecture	Data	Abdennour et al., 2021
Driver identification	ML	Machine learning (driving data)	Driver identification using ML with driving behavior information	Data	Li et al., 2019
Driver identification	DL	CNN, LSTM	Lightweight driver behavior identification from in-vehicle CAN-Bus data	Data	Ullah & Kim, 2020

Driver identification	Mixed	Comparative ML and DL	Review of driving behavior analysis methods on continuous vehicular data streams	Data	Peppes et al., 2021
Driver identification	ML	Ten ML algorithms (comparative)	Comparative assessment of ML algorithms for driver detection via CAN bus	Data	Martinelli et al., 2020
Driver identification	DL	LSTM (optimized)	Optimized LSTM framework for driving behavior identification with feature selection	Data	Zhang et al., 2019
Driver identification	DL	CNN	Real-time driver identification from vehicular sensor data	Data	Li, 2021
Driver identification	DL	CNN (feature extraction)	Driver identification from hidden features extracted from driving data	Data	Chen et al., 2018
Anomaly detection, vehicles	ML	SVM, decision tree, KNN	Intrusion and fault detection on the CAN bus: comparative performance study	Action	Bari et al., 2023
Anomaly detection, vehicles	ML	Genetic programming decision trees	In-vehicle detection of targeted CAN-bus attacks	Action	Fenzl et al., 2021
Anomaly detection, vehicles	Mixed	Snapshot-stacked ensemble (DL + ML)	Vehicle breakdown prediction for resource allocation and cost reduction	Action	Khoshkangini et al., 2023
Anomaly detection, vehicles	DL	MSALSTM-CNN (multi-stage attention)	Anomaly detection in automated vehicles via multi-stage attention	Action	Javed et al., 2021
Anomaly detection, vehicles	DL	LSTM, autoencoder	Systematic survey of DL-based in-vehicle network intrusion detection	Action	Luo et al., 2023
Anomaly detection, vehicles	DL	CNN (adaptive learning rate)	Fault detection in rotating parts using CNN with adaptive hierarchical algorithm	Action	Apeiranthitis et al., 2024
Vehicle health monitoring	ML	K-nearest neighbors	PdM and component-level vehicle health diagnosis with edge computing	Action	Vasavi et al., 2021
Vehicle health monitoring	ML	Decision tree, SVM, KNN, random forest	Remote vehicle health monitoring and prognostic maintenance system	Action	Shafi et al., 2018
Vehicle health monitoring	ML	SVM, decision tree, ensemble, KNN, ANN, random forest	Integrated emergency health monitoring and alert generation for IoV	Action	Kaur & Verma, 2026
Vehicle health monitoring	DL	Deep neural network	Structural health diagnosis and prognosis using machine learning	Action	Bao & Li, 2021

Connected-car customer experience	Other	Data analytics	Evaluation of PdM benefits and data validity challenges in connected cars	Compliance	Marengoni & Montenero, 2023
Connected-car customer experience	Other	CX analytics framework	Customer experience management strategy for automotive loyalty	Compliance	Scherpen et al., 2018
IoV PdM	ML	K-nearest neighbors	IoT-based PdM approach for connected car maintenance	Analysis	Solanki & Dhall, 2017
IoV PdM	ML	Random forest	5G NB-IoT implementation for predictive car maintenance	Analysis	Kullerkupp, 2021
IoV PdM	ML	Hierarchical modified fuzzy SVM	ML-based PdM for industrial IoT vehicle fleets	Analysis	Chaudhuri, 2018
IoV PdM	DL	GA-optimized neural network	Rear-end collision prediction via deep learning in IoV	Analysis	Chen et al., 2018
IoV PdM	Mixed	Edge intelligence architectures	Mobile edge intelligence and computing framework for IoV	Analysis	Zhang & Letaief, 2020
IoV PdM	DL	CNN, RBM, deep belief networks, GAN	Combination of AI and IoT for real-time vehicle monitoring and operation	Analysis	AlGhanem & Abdallah, 2024
IoV PdM	DL	CNN, graph CNN, LSTM	Advanced DL frameworks for connected vehicle applications	Analysis	Mezair et al., 2023
IoV PdM	DL	Deep and reinforcement learning	Review of current IoV issues and future trends in DL/RL applications	Analysis	Elmoiz Alatabani et al., 2022
IoV PdM	ML	Machine learning (automotive)	ML subfields, use cases and challenges for automotive PdM	Analysis	Theissler et al., 2021
IoV PdM	DL	Federated learning	Role of federated learning in IoV for privacy, scalability and availability	Analysis	Bala & Chhabra, 2023
IoV PdM	DL	Federated learning	Federated learning-based PdM for vehicles with enhanced data privacy	Analysis	Singh, 2021
IoV PdM	ML	Hierarchical modified fuzzy SVM	PdM of vehicle fleets on industrial IoT datasets	Analysis	Chaudhuri & Ghosh, 2021
RUL prediction	Mixed	CNN-GRU vs GRU (comparative)	CNN-GRU outperforms GRU for RUL prediction in maintenance decision-making	Analysis	Azyus et al., 2023
RUL prediction	ML	Support vector machine	RUL prognostics for vehicle lithium-ion batteries using data-driven methods	Analysis	Wu et al., 2016

RESULTS AND DISCUSSION

This section reports what the examination of the field established, and then the architecture built on it. Four stages are presented in order. The coding of the 45 studies comes first, as a quantitative distribution across research categories, computational paradigms and architectural function.

The five research questions are then answered against that evidence base, with convergent and divergent findings separated and a confidence grade attached to each answer. The implementation challenges that recur across the corpus are synthesised third. The architecture follows, and it follows here because its specification is determined by the three preceding stages: the distribution fixes which architectural functions are already well served and which are not, the evidence synthesis fixes what each layer must be capable of, and the implementation challenges fix what no architecture of this kind can resolve. The remaining subsections compare the architecture against the published standards examined earlier, restate the gaps in the light of both, and set out the research directions that follow.

Quantitative findings

Counting rules. All percentages in this section use n = 45 as the denominator, corresponding to the 45 included studies listed in (Tab. 6), and each study contributes exactly one count to each of the three classifications below. Where a study addressed more than one topic, it was assigned to the category matching its stated primary research problem, so categories are mutually exclusive and sum to 45; no study is double-counted and no percentage sums to more than 100. Percentages are rounded to the nearest whole number, so column totals may differ from 100 by one point. Algorithm type was assigned from the method the study evaluates: Mixed denotes a study that compares ML and DL approaches or combines them in a single model, and non-algorithmic denotes analytics or review contributions that evaluate no model. Layer assignment follows the Stage 2 deductive mapping described in the subsection on coding and taxonomy. Every count below is reproducible by tallying the corresponding column of Table 6. Systematic coding of the 45 reviewed studies across the seven taxonomy categories reveals a concentration of research activity in two domains. IoV PdM constitutes the largest category (n = 12, 27%), followed by general-purpose PdM (n = 10, 22%), together accounting for approximately half the corpus. Driver identification (n = 9, 20%) and anomaly detection (n = 6, 13%) represent the next most active areas, while vehicle health monitoring (n = 4, 9%), RUL prediction (n = 2, 4%) and connected-car customer experience (n = 2, 4%) receive comparatively limited attention. The full distribution is shown in (Fig. 4). By computational paradigm, deep learning accounts for 19 of the 45 studies (42%), classical machine learning for 17 (38%), mixed or comparative approaches for 6 (13%), and non-algorithmic contributions for 3 (7%); these four counts sum to 45. By architectural layer the distribution is uneven: 24 studies (53%) map to the Analysis Layer, being the 12 IoV PdM, 10 general-purpose PdM and 2 RUL studies; 10 (22%) to the Action Layer, being the 6 anomaly detection and 4 vehicle health monitoring studies; 9 (20%) to the Data Layer, all driver identification; and 2 (4%) to the Compliance Layer, both connected-car CX studies. These four counts sum to 45.

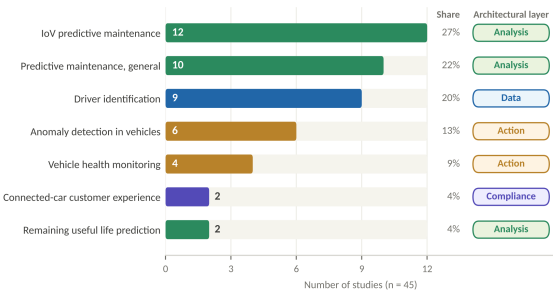


Figure 4: Distribution of the 45 reviewed studies by research category, with the share of the corpus and the architectural layer to which each category was mapped. Categories are mutually exclusive and sum to 45; layer assignment follows the deductive stage of the coding procedure.

(Fig. 5) cross-tabulates the same corpus on two dimensions at once, computational paradigm against architectural layer, with cell shading proportional to study count. This makes two patterns explicit that the category view in (Fig. 4) cannot show. First, the column totals show research activity concentrated at the Analysis Lay-

er (24 studies), with the Data and Action Layers receiving moderate attention (9 and 10) and the Compliance Layer reaching 2. Second, the Compliance column is empty for every algorithmic row: deep learning, classical machine learning and mixed approaches each return zero studies there, and the only two entries in that column are non-algorithmic. The most capable methods are therefore not reaching the dimension of IoV PdM where deployment is actually blocked. The two studies at the Compliance Layer employ neither ML nor DL methods. This concentration pattern directly informs the architectural priorities discussed in the subsection on algorithm selection.



Figure 5: Distribution of the 45 reviewed studies by computational paradigm against architectural layer. Cell shading is proportional to study count; row and column totals are given at the margins. Each study contributes exactly one count, so the cells sum to 45.

First research question: current technological capabilities

RQ1.1, integration of IoV with PdM systems. The corpus converges on the conclusion that end-to-end integration is technically solved for the single-operator, single-platform case. Five studies (Chaudhuri, 2018; Kaur & Verma, 2026; Kullerkupp, 2021; Solanki & Dhall, 2017; Vasavi et al., 2021) report working pipelines from in-vehicle acquisition through remote inference to a maintenance output, using unrelated connectivity stacks and unrelated model families, which makes the convergence informative rather than an artefact of shared method. The evidence diverges on where inference should execute: cellular back-haul to a central service is assumed in (Kullerkupp, 2021; Solanki & Dhall, 2017), while Vasavi et al. (2021) and Zhang and Letaief (2020) treat central inference as unviable at latency-sensitive decision points and place computation at the edge. This is a genuine disagreement about architecture, not a difference in emphasis, and it remains unresolved because no included study evaluates both placements on the same workload. Confidence: moderate for feasibility, low for generality. Conclusion: IoV PdM integration is demonstrated, but every demonstration in the corpus is single-operator and single-platform, so the capability claim cannot be extended to heterogeneous fleets; this is the same boundary that Theissler et al. (2021) and Killeen et al. (2019) identify from the opposite direction.

RQ1.2, data processing challenges. Three challenges recur with high consistency. Representation is the first: raw one-dimensional sensor streams underperform until transformed, whether into image-like representations for convolutional processing (Silva & Capretz, 2019), windowed sequences for recurrent models (Peppes et al., 2021; Ullah & Kim, 2020), or reduced feature spaces (Jha et al., 2021; Singh et al., 2016). Latency is the second, and is treated as an architectural rather than an algorithmic problem across (Grover et al., 2021; Liu et al., 2019; Zhang & Letaief, 2020). Heterogeneity is the third: format and semantic variation across manufacturers is documented as a first-order obstacle in (Agbaje et al., 2022; Arooj et al., 2022; Datta et al., 2016). The evidence does not diverge on the existence of these challenges. It diverges on their relative severity, and the divergence tracks the deployment context of each study, with edge-oriented work ranking latency first and fleet-oriented work ranking heterogeneity first. Confidence: high for the identification of the three challenges, low for their relative weighting. Conclusion: data processing capability is bounded by normalisation and transport, not by model capacity, which is the finding that

motivates placing acquisition and quality flagging in a distinct architectural layer.

Second research question: limitations and challenges

RQ2.1, technical limitations affecting PdM accuracy. Two limitations are reported consistently and independently. Label scarcity and class imbalance, arising because failures are rare in operational fleets, are identified in (Carvalho et al., 2019; Khoshkangini et al., 2023; Theissler et al., 2021; Ucar et al., 2024) and are the stated motivation for the unsupervised and generative approaches in (Serradilla et al., 2022; Ucar et al., 2024). Cross-platform generalisation failure is reported in (Killeen et al., 2019; Peppes et al., 2021; Theissler et al., 2021), with models trained on one vehicle platform degrading on another without substantial retraining. The evidence diverges on whether deep architectures mitigate or aggravate these limitations: Serradilla et al. (2022) report accuracy gains from LSTM architectures on time-series vehicle data, while Theissler et al. (2021) and Ucar et al. (2024) treat the additional data appetite of such models as a net cost under scarcity. Both positions are defensible on the evidence presented, because the studies evaluate different data volumes. Confidence: high for the existence of both limitations, low for their quantitative magnitude in operational conditions. Conclusion: the corpus establishes that accuracy claims obtained on curated datasets do not transfer, but does not quantify the shortfall; the field measurement called for by Tian et al. (2018) has not been performed for IoV PdM.

RQ2.2, impact of security concerns on PdM reliability. The corpus converges on two findings. First, the in-vehicle network is attackable in ways that directly corrupt the prognostic input: CAN-Bus injection and masquerade attacks are demonstrated and detected in (Bari et al., 2023; Fenzl et al., 2021; Luo et al., 2023), and the wider IoV attack surface is characterised in (Abu Talib et al., 2018; Karim et al., 2022). Second, the principal privacy mitigation proposed in the corpus, federated learning (Bala & Chhabra, 2023; Singh, 2021), addresses confidentiality during training and does not touch transmission integrity at inference time. The evidence does not diverge; rather, it stops. No included study measures predictive maintenance performance under adversarial or corrupted input, and none reports the downstream consequence of a compromised signal for a maintenance decision. Confidence: high that the vulnerability exists, none for its effect on PdM reliability, because the measurement has not been made. Conclusion: RQ2.2 cannot be answered from the current literature. This is a substantive negative result and it directly motivates the integrity requirements specified in Table 7 and the limitation acknowledged in the subsection on limitations of the proposed architecture.

Third research question: contribution of learning methods

RQ3.1, effective ML models for IoV data analysis. Ensemble methods are the most consistent performers across the corpus, with random forest reported as robust to sensor noise in (Carvalho et al., 2019; Kaur & Verma, 2026; Shafi et al., 2018), and this convergence holds across industrial and vehicle contexts. Instance-based and kernel methods occupy identifiable niches: KNN where the training set is small and explanation by analogy is valued (Vasavi et al., 2021), and SVM variants where dimensionality is high and failure events are few (Chaudhuri, 2018; Chaudhuri & Ghosh, 2021). Reported accuracies cannot be compared across studies, because datasets, target components and failure definitions differ, and this incomparability is itself a finding rather than an inconvenience. Confidence: moderate. Conclusion: model family selection in the reviewed work is driven by operating constraints rather than by measured superiority, which is the empirical basis for specifying the Analysis Layer

as algorithm-agnostic.

RQ3.2, role of DL in advancing PdM in IoV. Recurrent architectures, principally LSTM and its gated variants, dominate temporal prognostic tasks (Azyus et al., 2023; Peppes et al., 2021; Serradilla et al., 2022; Ullah & Kim, 2020), and convolutional architectures dominate where the input has been transformed into a spatial representation (Apeiranthitis et al., 2024; Javed et al., 2021; Silva & Capretz, 2019). Serradilla et al. (2022) provide the only systematic comparison in the corpus and report an accuracy advantage over classical ML on time-series data at substantially higher computational cost. The evidence diverges on deployability: Ullah and Kim (2020) demonstrate that a sparsified network runs within in-vehicle constraints, whereas Grover et al. (2021) and Liu et al. (2019) treat available on-vehicle compute as the binding limit. Confidence: moderate to high for accuracy gains on time-series tasks, low for in-vehicle deployability. Conclusion: deep learning improves prediction where labelled data and compute are available, and neither condition holds by default in an operational fleet; this trade-off, not a ranking of architectures, is what the framework must accommodate.

Fourth research question: implementation practice

RQ4.1 and RQ4.2 are treated together, because the evidence base does not support separating them. Three included studies report implementation in an industrial or commercial setting (Chaudhuri, 2018; Chaudhuri & Ghosh, 2021; Kullerkupp, 2021), and one reports fleet-scale breakdown prediction in an operational context (Khoshkangini et al., 2023). Across these, one practice recurs: predictive output is coupled to a scheduling or resource-allocation decision rather than delivered as an alert. Beyond this, no practice is reported by more than one study. The studies report incompatible outcome measures — communication performance in (Kullerkupp, 2021), classification accuracy in (Chaudhuri, 2018; Chaudhuri & Ghosh, 2021), resource-allocation gain in (Khoshkangini et al., 2023) — and none reports cost, downtime or availability against a documented pre-deployment baseline. There is accordingly no divergence to assess, because there is no shared measure on which the studies could disagree. Confidence: low. Conclusion: RQ4 cannot be answered with confidence from the current corpus. Statements about industrial outcomes for IoV PdM in the reviewed literature are single-study claims and should not be aggregated. The absence of comparable outcome measurement is reported here as a finding, and it is the reason the economic evaluation gap is carried into the subsection on comparison with existing architectures.

Fifth research question: future directions

RQ5.1, emerging technologies. Three technology directions recur: distributed learning across fleets (Bala & Chhabra, 2023; Singh, 2021); the integration of AI with IoT sensing for continuous monitoring (AlGhanem & Abdallah, 2024; Elmoiz Alatabani et al., 2022); and edge intelligence as the execution substrate (Grover et al., 2021; Liu et al., 2019; Zhang & Letaief, 2020). These are complementary rather than competing, and the studies do not disagree about them. The corpus is, however, uniformly forward-looking on these points: the claims are projections, not results. Confidence: low, and low for the reason that the evidence is anticipatory by nature. Conclusion: the direction of travel is consistent across the corpus; the supporting evidence is not yet empirical.

RQ5.2, contribution to self-maintaining systems. Two studies (Bala & Chhabra, 2023; Zhang & Letaief, 2020) argue that in-vehicle intelligence and distributed model updating together make autonomous maintenance conceivable, and Elmoiz Alatabani et al. (2022) extend this to reinforcement learning control. No included study demonstrates a closed loop in which a vehicle detects degradation, selects an intervention and

initiates it without human authorisation. Confidence: low. Conclusion: self-maintenance is a stated aspiration in the corpus and not a demonstrated capability. It is worth noting that the constraint here may not be technical: an autonomously initiated maintenance action is precisely the case in which the business, supply chain and regulatory admissibility questions formalised in the subsection on formal specification and layer interfaces must be resolved automatically, and no reviewed work addresses that.

Implementation challenges identified across the corpus

Three categories of implementation challenge recur across the reviewed literature with sufficient consistency to read as structural barriers to IoV PdM deployment and not as study-specific findings.

Data integration and management present the foremost operational barrier. The heterogeneity of data sources within IoV ecosystems, spanning OBD-II, CAN-Bus, V2X feeds and external contextual streams, requires data fusion architectures capable of real-time processing and format normalization across manufacturer standards (Agbaje et al., 2022; Arooj et al., 2022; Carvalho et al., 2019; Datta et al., 2016; Theissler et al., 2021). Data quality degradation at the source propagates into predictive model unreliability at a rate that algorithmic sophistication cannot compensate for independently of data governance (Silva & Capretz, 2019).

Model complexity and interpretability constitute a second structural tension. As DL architectures achieve higher predictive accuracy, their opacity increases, reducing the ability of maintenance teams to audit, trust and act on model outputs (Serradilla et al., 2022; Theissler et al., 2021; Ucar et al., 2024). This tension is particularly acute in regulated contexts where maintenance decisions require traceable audit trails and where liability for incorrect predictions must be attributable. No reviewed study proposes a resolution to this tension that scales to heterogeneous fleet deployment.

Computational resource constraints impose a third practical barrier. The more capable DL models require GPU or TPU infrastructure that is unavailable within vehicle-embedded systems and at resource-constrained edge nodes (Grover et al., 2021; Liu et al., 2019). The gap between what is achievable in research environments and what is deployable in operational IoV infrastructure remains a persistent limitation across the corpus (Teoh et al., 2023; Wang et al., 2020; Zhang & Letaief, 2020).

The proposed compliance-oriented reference architecture

Each element of the architecture described below answers a specific result reported above, and the correspondence is worth stating before the specification itself. The coded distribution shows that data acquisition and analytical modelling are addressed by 33 of the 45 studies while the admissibility of a maintenance action is addressed computationally by none, which fixes where the architecture must add something and where it should adopt what already exists. The assessment of the seven published reference architectures and standards, carried out independently of the coding scheme, returns the same division, which is what raises the gap from an artefact of the categories used here to a property of the field. The evidence synthesis supplies the requirements for the individual layers: the divergence over edge and cloud placement is why the analysis layer is specified as infrastructure-agnostic, the unanswerability of the second research question is why integrity is stated as a precondition rather than a guarantee, and the incomparability of the reported industrial outcomes is why the compliance cost is defined as a measurable quantity. Elements that cannot be traced back in this way are identified as assumptions in the limitations of the proposed architecture. The Holistic Compliance Framework (HCF) is a four-layer architecture designed to address the gap identified across the reviewed literature: the absence of a formalized

mechanism for aligning predictive maintenance actions with the business, regulatory and supply chain constraints that govern real-world maintenance operations. Existing frameworks address data acquisition, analytical modeling and operational decision-making in relative isolation from the organizational context within which maintenance decisions must ultimately be executed. The HCF treats this integration as a design requirement of the architecture itself. The framework, operates as a closed-loop system in which data flows upward through four interdependent layers while compliance constraints propagate downward to govern the solution space available at each level. Two properties distinguish the HCF from prior work. First, the Compliance Layer governs the architecture instead of sitting alongside it as an add-on, which ensures that technically optimal maintenance recommendations remain organizationally and legally executable. Second, the framework's foundation is specified to support the full V2X communication spectrum, comprising V2I, V2V, V2P, V2D and V2N, ensuring comprehensive data collection across all relevant IoV domains.

Data layer

The Data Layer forms the foundation of the HCF, encompassing device and sensor management, data source integration, and V2X communication network configuration across all relevant domains. This layer is responsible for collecting, aggregating and preprocessing data from OBD-II ports, CAN-Bus networks, roadside infrastructure sensors and external contextual feeds. Acquisition protocols are specified to meet IoV environment requirements, prioritizing low latency, fault tolerance and cross-manufacturer interoperability (Agbaje et al., 2022; Teoh et al., 2023; Wang et al., 2020). The integrity and completeness of data processed at this layer directly determines the predictive capability of the layers above it.

Analysis layer

The Analysis Layer transforms preprocessed data into actionable predictive insights through ML and DL models, statistical modeling and real-time analytical processing. Tasks executed at this layer include fault detection, anomaly identification, RUL estimation and driver behavior analysis. Algorithm selection within this layer is governed by operational context, specifically the trade-offs between interpretability, scalability and data requirements assessed in Table 1 and Table 2. The layer is architecturally agnostic to specific model implementations, enabling substitution of analytical components as the state of the art advances without requiring modification of adjacent layers (Serradilla et al., 2022; Theissler et al., 2021). Security protocols for the data processing pipeline are placed inside this layer as a structural component of the processing pipeline, not alongside it (Abu Talib et al., 2018; Karim et al., 2022).

Action layer

The Action Layer translates analytical outputs into operational decisions and resource allocations. It encompasses a decision support system (DSS) that synthesizes predictive outputs into ranked maintenance recommendations, a resource allocation module that optimizes scheduling against fleet availability and parts inventory, and customer experience (CX) integration that governs how maintenance communications are delivered to end users and fleet operators. Among the four layers, the Action Layer is the one whose outputs reach the vehicle and the operator, and it is therefore the layer through which vehicle uptime, maintenance cost and operational safety outcomes are realised (Marengoni & Montenero, 2023; Scherpen et al., 2018; Shafi et al., 2018). This is not in tension with the governing role assigned to the Compliance Layer, and the distinction is worth stating precisely because the two claims are easily read as competing. The Compliance Layer determines which actions are admissible; the Action Layer selects and executes among the admissible actions and is accountable for the operational result. In control terms the Compliance

Layer defines the feasible set and the Action Layer optimises within it. Neither layer can produce an outcome alone: a compliance constraint with no action has no effect, and an action outside the feasible set is not executable.

Compliance layer

The Compliance Layer constitutes the defining contribution of the HCF and addresses the primary gap identified in the literature. Positioned as the governing layer above all operational decisions, it formalizes three integration functions that prior frameworks treat as external to the maintenance architecture.

Business system integration aligns maintenance actions with enterprise resource planning (ERP) systems, budget constraints and organizational approval workflows, ensuring that maintenance decisions are executable within the institutional context in which the fleet operates. Supply chain coordination manages the dynamic alignment of maintenance scheduling with parts availability, supplier lead times and logistics constraints, a dependency that the reviewed literature identifies as a critical implementation barrier but does not address architecturally (Killeen et al., 2019; Theissler et al., 2021). Regulatory compliance enforces adherence to vehicle safety standards, emissions regulations, cybersecurity obligations and data protection law across the jurisdictions in which a fleet operates, so that legal requirements enter the architecture as constraints instead of arriving afterwards as audit criteria. Three obligation families are distinguished because they constrain the architecture in different ways.

Type-approval and safety obligations. UNECE R155 and R156 (UNECE, 2021a; UNECE, 2021b) require a certified cybersecurity management system and software update management system as a condition of type approval, and ISO/SAE 21434 (ISO & SAE International, 2021) and ISO 26262 (ISO, 2018) define the corresponding engineering processes. Within the HCF these obligations bind the regulatory predicate gr: a maintenance action that would install unvalidated software, or that would disable a safety-relevant function outside an approved procedure, is excluded from the admissible set irrespective of its predicted utility. The audit record produced for each excluded action supplies part of the evidence base such a management system must maintain.

Data protection obligations. IoV PdM depends on exactly the data categories that data protection law regulates most tightly. Driver identification models (Abdennour et al., 2021; Khan et al., 2022; Li et al., 2019; Martinelli et al., 2020; Peppes et al., 2021; Ullah & Kim, 2020) process behavioural data that is personal data under the GDPR (European Parliament & Council of the European Union, 2016) and comparable regimes (State of California, 2018), and V2X feeds carry location data. Purpose limitation and data minimisation therefore constrain the Data Layer, not only the Compliance Layer: a signal collected for maintenance prognosis cannot be repurposed for driver scoring without a separate lawful basis. The architecture supports this through three mechanisms: purpose tags attached to signals at acquisition and checked at inference; retention limits enforced at the Data Layer; and a preference for federated or on-vehicle training of behavioural models (Bala & Chhabra, 2023; Singh, 2021) so that raw behavioural traces need not leave the vehicle. Data-subject rights, including erasure, propagate downward as a constraint on what training data a deployed model may retain, which is an operational problem the reviewed literature does not address at all.

Security of the maintenance decision chain

The Compliance Layer depends on the integrity of the inputs it receives. If CAN-Bus traffic can be forged (Bari et al., 2023; Fenzl et al., 2021; Luo et al., 2023) or V2X messages spoofed (Abu Talib et al., 2018; Karim et al., 2022), then a compliance decision computed over corrupted state is compliant only in form. The architecture therefore requires authenticated and integrity-protected

transport between the Data and Analysis layers and signed model artefacts in the Analysis Layer model registry, as specified in Table 7. It should be stated plainly that this positions the HCF to consume security guarantees rather than to provide them: the framework specifies where integrity must hold, and defers to ISO/SAE 21434 processes and to the intrusion detection literature for how it is achieved. The consequences of a breach for maintenance decision-making, and the behaviour of the architecture under partially compromised input, are not resolved here and are stated as a limitation in the subsection on limitations of the proposed architecture.

The Compliance Layer does not merely audit decisions made by the layers below it. It actively constrains the solution space available to the Action Layer, so that a technically optimal maintenance recommendation is surfaced only when it is also organizationally sanctioned and legally permissible.

Formal specification and layer interfaces

The preceding descriptions are functional. To make the architecture implementable, the governing relationship is stated formally here and the layer interfaces are specified in (Tab. 7).

Let S denote the observed state of a vehicle or fleet at time t, produced by the Data Layer. Let P = f(S) denote the prognostic output of the Analysis Layer, where f is any predictive model admitted by the layer. Let A denote the set of candidate maintenance actions that the Action Layer can derive from P, each element $a \in A$ carrying an expected utility u(a) expressed in terms of avoided downtime, cost and risk.

The Compliance Layer is defined as a constraint function

> C: $A \rightarrow A'$, where $A' \subseteq A$ and $a \in A'$ if and only if $gb(a) \wedge gs(a) \wedge gr(a)$

in which gb is the business admissibility predicate (budget authority, approval workflow state, ERP record consistency), gs the supply chain predicate (part availability within the required window given supplier lead time), and gr the regulatory predicate (safety, emissions, cybersecurity and data protection obligations applicable in the operating jurisdiction).

The Action Layer then selects $a^* = \arg\max_{a \in A'} u(a)$. Two consequences follow, and both are testable, which matters for the validation programme set out in the subsection on limitations of the proposed architecture. First, if $A' = \emptyset$ the architecture must escalate rather than degrade silently to the unconstrained optimum; escalation to a human authority is a required behaviour of a conforming implementation, not an error state. Second, the difference between $\max_{a \in A} u(a)$ and $u(a^*)$ is directly measurable and constitutes the compliance cost of a maintenance policy, a quantity that no reviewed framework makes observable.

Table 7: HCF layer interface specification. Each layer is defined by its inputs, outputs, the interfaces it exposes and the standards with which a conforming implementation must align.

Layer	Input	Output	Exposed interface	Alignment with existing standards
Data	OBD-II frames, CAN-Bus signals, V2X messages, roadside and contextual feeds	Time-aligned, normalised, quality-flagged state vector S	Acquisition and normalisation API; per-signal quality and provenance metadata	ISO 13374-1 data acquisition and manipulation blocks (ISO, 2003; ISO, 2007); ISO 14224 equipment taxonomy and failure-mode vocabulary (ISO, 2016)
Analysis	State vector S	Prognostic output P: fault class, anomaly score, RUL estimate with uncertainty, behavioural profile	Model registry with versioning; inference endpoint; explanation endpoint; integrity and authentication controls on the pipeline	ISO 13374-1 state detection, health and prognostic assessment blocks (ISO, 2003); ISO/SAE 21434 pipeline security engineering (ISO & SAE International, 2021)

Action	Prognostic output P, fleet availability, inventory state	Candidate action set A with utilities u(a); selected action a*; operator and customer communications	Ranked recommendation interface; scheduling interface; CX notification interface	ISO 13374-1 advisory generation block (ISO, 2003); OSA-CBM advisory interface (MIMOSA, 2014); SAE JA1011 decision justification criteria (SAE International, 2009)
Compliance	Candidate action set A; ERP and approval state; supplier lead times; jurisdictional rule set	Admissible subset A′; rejection reason per excluded action; audit record	Constraint evaluation service exposing gb, gs, gr; immutable audit log; escalation channel	UNECE R155 and R156 management-system evidence (UNECE, 2021a; UNECE, 2021b); ISO 26262 safety case traceability (ISO, 2018); GDPR Articles 5, 25 and 30 (European Parliament & Council of the European Union, 2016)

The rejection reason attached to each excluded action is not an implementation detail. It is what converts the Compliance Layer from a filter into an auditable component, and it is the mechanism by which the framework satisfies the traceability requirement that the subsection on implementation challenges identifies as unresolved in the current literature.

Algorithm selection rationale mapped to architectural layers

The distribution of algorithm types across research categories, reported in the subsection on quantitative findings, directly informed the architectural design decisions of the HCF. This mapping, from review finding to framework design, constitutes the argumentative bridge between the review and the proposed framework and is visualized in Fig. 5. At the Data Layer, driver identification studies (n = 9) consistently employ CNN and LSTM architectures to process CAN-Bus behavioral signals (Abdennour et al., 2021; Khan et al., 2022; Peppes et al., 2021; Ullah & Kim, 2020). This establishes that the Data Layer must support not only raw sensor acquisition but structured behavioral feature extraction, a requirement reflected in the HCF Data Layer specification. At the Analysis Layer, the dominance of DL approaches among IoV PdM studies and the persistence of classical ML in resource-constrained contexts such as vehicle health monitoring jointly establish that the Analysis Layer must be algorithm-agnostic, capable of hosting both paradigms without privileging either. At the Action Layer, anomaly detection and vehicle health monitoring studies consistently target decision support and scheduling outputs (Bari et al., 2023; Javed et al., 2021; Shafi et al., 2018; Vasavi et al., 2021), confirming that the Action Layer must carry decision support as a defined architectural component, not as an optional module. At the Compliance Layer, the near-absence of research (n = 2, both non-algorithmic) is itself the most consequential finding of the review, indicating that the compliance, business integration and regulatory dimensions of IoV PdM are structurally unaddressed in the current literature.

Synthesis of review findings against architecture design

Five converging patterns across the 45 reviewed studies shaped the framework. No reviewed framework formalizes the integration of PdM outputs with organizational approval workflows, supply chain constraints or regulatory obligations, the three functions addressed by the HCF Compliance Layer. The reviewed literature treats cybersecurity as a parallel concern and not as an architectural requirement; the HCF positions security protocols within the Analysis Layer as a structural component of the data processing pipeline (Abu Talib et al., 2018; Bala & Chhabra,

2023; Karim et al., 2022; Singh, 2021). Edge and fog computing are consistently identified as prerequisites for operational IoV PdM (Grover et al., 2021; Liu et al., 2019; Teoh et al., 2023; Wang et al., 2020), but no reviewed framework specifies how edge–cloud task distribution should be governed, a gap the HCF Analysis Layer addresses through its infrastructure-agnostic design. Driver behavior data is treated as supplementary in most reviewed studies (Adnan et al., 2018; Li et al., 2019; Nirmali et al., 2017) despite demonstrated predictive value; the HCF treats behavioral data as a primary input at the Data Layer. Finally, customer experience outcomes are addressed in only two studies (Marengoni & Montenero, 2023; Scherpen et al., 2018) and only descriptively; the HCF Action Layer integrates CX as an output metric alongside technical maintenance performance.

Comparison with existing reference architectures

The frameworks and standards reviewed in the subsection on existing reference architectures and standards are assessed in Table 8 against the four HCF layers and the three compliance functions. The comparison is deliberately unflattering to the novelty claim in one respect. The HCF is not the first architecture to structure the path from sensor data to maintenance advisory; ISO 13374 and OSA-CBM did that two decades ago, and did it in more implementable detail than this paper provides. The claim made here is narrower. No existing architecture makes the admissibility of a maintenance action an explicit, evaluable function of business, supply chain and regulatory state, and none exposes the rejection reason as an auditable artefact.

Table 8: Comparison of the HCF with existing maintenance reference architectures and standards. ● indicates that the function is specified architecturally; ◐ that it is partially addressed or treated as an external dependency; ○ that it is out of scope for that instrument.

Framework / standard	Data acquisition	Analysis, prognostics	Action, advisory	Business integration	Supply chain	Regulatory constraint	Primary scope
ISO 13374 and OSA-CBM (ISO, 2003; ISO, 2007; MIMOSA, 2014)	●	●	●	○	○	○	Condition monitoring data processing and interfaces, generic to asset class
ISO 14224 (ISO, 2016)	●	◐	○	○	○	○	Reliability and maintenance data taxonomy and exchange
SAE JA1011 on reliability-centred maintenance (SAE International, 2009)	○	◐	●	◐	○	◐	Criteria for justifying maintenance policy decisions
IIRA and RAMI 4.0 (DIN, 2016; Industrial Internet Consortium, 2019)	●	●	◐	●	◐	○	Industrial system architecture, not maintenance-specific
UNECE R155 and R156, and ISO/SAE 21434 (ISO & SAE International, 2021; UNECE, 2021a; UNECE, 2021b)	○	○	○	◐	○	●	Regulatory obligation and process, not architecture
IoT-based PdM architectures in the corpus (Killeen et al., 2019; Solanki & Dhall, 2017; Teoh et al., 2023; Zhang & Letaief, 2020)	●	●	◐	○	○	○	Deployment-specific pipelines
The architecture proposed in this paper	●	●	●	●	●	●	Admissibility of IoV maintenance decisions

Three observations follow. The upper rows show that the Data, Analysis and Action functions are already well covered, and the HCF should be understood as adopting that coverage instead of replacing it, which is why Table 7 maps each layer onto the corresponding ISO 13374 processing block. The two columns covering business integration and supply chain are empty for every prior instrument. The regulatory instruments specify obligations but not the mechanism by which a system satisfies them at the level of an individual decision. The HCF occupies that space, and its contribution should be assessed against that narrow claim. One caveat applies to the whole table. The HCF row records design intent, while every other row records a specification that has been implemented and, in most cases, certified against. A conceptual architecture cannot be compared on equal terms with a deployed standard, and Table 8 should not be read as evidence of superiority. It identifies an unoccupied position; whether the HCF occupies it usefully is what the validation programme in the subsection on limitations of the proposed architecture exists to determine. (Tab. 7 and 8).

RESULTS AND CONCLUSION

This paper sets out a reference architecture for predictive maintenance in connected vehicle fleets, built to carry the business, regulatory and supply chain constraints that govern real-world maintenance operations. The need for it rests on a finding about the present state of the field: neither the published research nor the reference architectures and standards that currently govern maintenance system design provides a formalized mechanism for aligning maintenance actions with those constraints. That finding was established first, on evidence independent of the architecture, and the architecture answers it. Through a systematic-narrative hybrid review of 45 studies selected from an initial corpus of 227 records, the study characterizes the current state of IoV PdM research and identifies persistent gaps. The Holistic Compliance Framework responds to the most consequential of those gaps. It is specified at the conceptual level and has not yet been validated independently. The review findings establish that the IoV PdM literature is technically mature at the algorithmic level: deep learning accounts for 42% of the coded studies, with demonstrated capability in fault detection, anomaly identification and RUL estimation across connected vehicle contexts. The compliance-facing dimensions, namely business system integration, supply chain coordination and regulatory alignment, are addressed by 4% of the reviewed studies and by none of the seven reference architectures and standards compared in Table 8. The HCF places these dimensions in a governing layer of the architecture itself. The framework's four-layer architecture provides a structured foundation for IoV PdM deployment that is agnostic to specific algorithm choices and vehicle types. The V2X-enabled Data Layer supports comprehensive data collection across all relevant communication domains. The algorithm-agnostic Analysis Layer accommodates both ML and DL paradigms as the state of the art evolves. The Action Layer integrates decision support and customer experience functions that existing frameworks treat as peripheral. The Compliance Layer, the framework's core contribution, carries organizational and regulatory constraints as binding architectural requirements, and the subsection on formal specification and layer interfaces states its governing relationship to the Action Layer formally so that it can be implemented and tested. Seven research directions are identified, spanning advanced AI integration, cybersecurity architecture, edge–cloud optimization, behavioral PdM, standardization, federated fleet collaboration and model interpretability. Each is grounded in a specific gap identified through the review and maps to a defined HCF component. For automotive manufacturers, fleet operators and technology providers, the framework offers an architectural reference for the point at which algorithmic maintenance research meets operational and regulatory constraint. That value is

at present prospective. It rests on the argument developed in the Materials and Methods and the Results and Discussion sections and on the gap documented in Table 8, and it remains prospective until the validation described in the subsection on limitations of the proposed architecture has been carried out.

REFERENCES

Abdennour, N.; Ouni, T.; Ben Amor, N., (2021). Driver identification using only the CAN-Bus vehicle data through an RCN deep learning approach. Rob. Auton. Syst., 136: 103707. https://doi.org/10.1016/j.robot.2020.103707

AlGhanem, H.; Abdallah, S., (2024). The future of the Internet of Vehicles. In: BUiD Doctoral Research Conference 2023: Multidisciplinary Studies. Springer, Cham, Switzerland. 301-309. https://link.springer.com/chapter/10.1007/978-3-031-56121-4_29

Apeiranthitis, S.; Zacharia, P.; Chatzopoulos, A.; Papoutsidakis, M., (2024). Predictive maintenance of machinery with rotating parts using convolutional neural networks. Electronics, 13(2): 460. https://doi.org/10.3390/electronics13020460

Arena, F.; Collotta, M.; Luca, L.; Ruggieri, M.; Termine, F.G., (2021). Predictive maintenance in the automotive sector: a literature review. Math. Comput. Appl., 27(1): 2. https://doi.org/10.3390/mca27010002

Aromataris, E.; Pearson, A., (2014). The systematic review: an overview. Am. J. Nurs., 114(3): 53-58. https://doi.org/10.1097/01.NAJ.0000444496.24228.2c

Arooj, A.; Farooq, M.S.; Akram, A.; Iqbal, R.; Sharma, A.; Dhiman, G., (2022). Big data processing and analysis in Internet of Vehicles: architecture, taxonomy and open research challenges. Arch. Comput. Methods Eng., 29(2): 793-829. https://doi.org/10.1007/s11831-021-09590-x

Arthurs, P.; Gillam, L.; Krause, P.; Wang, N.; Halder, K.; Mouzakitis, A., (2022). A taxonomy and survey of edge cloud computing for intelligent transportation systems and connected vehicles. IEEE Trans. Intell. Transp. Syst., 23(7): 6206-6221. https://doi.org/10.1109/TITS.2021.3084396

Azyus, A.F.; Wijaya, S.K.; Naved, M., (2023). Prediction of remaining useful life using the CNN-GRU network: a study on maintenance management. Softw. Impacts, 17: 100535. https://doi.org/10.1016/j.simpa.2023.100535

Bao, Y.; Li, H., (2021). Machine learning paradigm for structural health monitoring. Struct. Health Monit., 20(4): 1353-1372. https://doi.org/10.1177/1475921720972416

Bari, B.S.; Yelamarthi, K.; Ghafoor, S., (2023). Intrusion detection in vehicle controller area network bus using machine learning: a comparative performance study. Sensors, 23(7): 3610. https://doi.org/10.3390/s23073610

Butte, S.; Prashanth, A.; Patil, S., (2018). Machine learning based predictive maintenance strategy: a super learning approach with deep neural networks. In: 2018 IEEE Workshop on Microelectronics and Electron Devices, Boise, United States. 1-5. https://doi.org/10.1109/WMED.2018.8360836

Chaudhuri, A., (2018). Predictive maintenance for industrial IoT of vehicle fleets using hierarchical modified fuzzy support vector machine. arXiv preprint, arXiv:1806.09612. https://doi.org/10.48550/arXiv.1806.09612

Chaudhuri, A.; Ghosh, S.K., (2021). Predictive maintenance of vehicle fleets using hierarchical modified fuzzy support vector machine for industrial IoT datasets. In: International Conference on Hybrid Artificial Intelligence Systems. Springer, Cham, Switzerland. 331-342. https://doi.org/10.1007/978-3-030-86271-8_28

Chen, C.; Xiang, H.; Qiu, T.; Wang, C.; Zhou, Y.; Chang, V., (2018). A rear-end collision prediction scheme based on deep learning in the Internet of Vehicles. J. Parallel Distrib. Comput., 117: 192-204. https://doi.org/10.1016/j.jpdc.2017.08.014

Dalzochio, J.; Kunst, R.; Barbosa, J.L.V.; da Silva Neto, P.C.; Pignaton, E.; ten Caten, C.S.; da Penha, A.L.T., (2023). Predictive maintenance in the military domain: a systematic review of the literature. ACM Comput. Surv., 55(13s): 1-30. https://doi.org/10.1145/3586100

Dalzochio, J.; Kunst, R.; Pignaton, E.; Binotto, A.; Sanyal, S.; Favilla, J.; Barbosa, J., (2020). Machine learning and reasoning for predictive maintenance in Industry 4.0: current status and challenges. Comput. Ind., 123: 103298. https://

doi.org/10.1016/j.compind.2020.103298
Daniyan, I.; Mpofu, K.; Oyesola, M.; Ramatsetse, B.; Adeodu, A., (2020). Artificial intelligence for predictive maintenance in the railcar learning factories. Procedia Manuf., 45: 13-18. https://doi.org/10.1016/j.promfg.2020.04.032
Datta, S.K.; Da Costa, R.P.; Harri, J.; Bonnet, C., (2016). Integrating connected vehicles in Internet of Things ecosystems: challenges and solutions. In: 2016 IEEE 17th International Symposium on a World of Wireless, Mobile and Multimedia Networks, Coimbra, Portugal. 1-6. https://doi.org/10.1109/WoWMoM.2016.7523574
Deutsches Institut für Normung, (2016). DIN SPEC 91345:2016-04 Reference architecture model Industrie 4.0. Beuth Verlag, Berlin, Germany. https://www.beuth.de/de/technische-regel/din-spec-91345/250940128
Elmoiz Alatabani, L.; Ali, E.S.; Mokhtar, R.A.; Saeed, R.A.; Alhumyani, H.; Hasan, M.K., (2022). Deep and reinforcement learning technologies on Internet of Vehicle applications: current issues and future trends. J. Adv. Transp., 2022: 1947886. https://doi.org/10.1155/2022/1947886
Ersoz, O.O.; Inal, A.F.; Aktepe, A.; Turker, A.K.; Ersoz, S., (2022). A systematic literature review of the predictive maintenance from transportation systems aspect. Sustainability, 14(21): 14536. https://doi.org/10.3390/su142114536
European Parliament; Council of the European Union, (2016). Regulation 2016/679 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data. Off. J. Eur. Union, L119: 1-88. https://eur-lex.europa.eu/eli/reg/2016/679/oj
Fenzl, F.; Rieke, R.; Dominik, A., (2021). In-vehicle detection of targeted CAN bus attacks. In: Proceedings of the 16th International Conference on Availability, Reliability and Security, Vienna, Austria. 1-7. https://doi.org/10.1145/3465481.3465755
Giordano, D.; Giobergia, F.; Pastor, E.; La Macchia, A.; Cerquitelli, T.; Baralis, E.; Mellia, M.; Tricarico, D., (2022). Data-driven strategies for predictive maintenance: lesson learned from an automotive use case. Comput. Ind., 134: 103554. https://doi.org/10.1016/j.compind.2021.103554
Hussain, S.A.; Yusof, K.M.; Hussain, S.M.; Singh, A.V., (2019). A review of quality of service issues in Internet of Vehicles. In: 2019 Amity International Conference on Artificial Intelligence, Dubai, United Arab Emirates. 380-383. https://doi.org/10.1109/AICAI.2019.8701299
Industrial Internet Consortium, (2019). The industrial internet reference architecture, version 1.9. Industrial Internet Consortium, Needham, United States. https://www.iiconsortium.org/iira/
International Organization for Standardization, (2003). ISO 13374-1:2003 Condition monitoring and diagnostics of machines, data processing, communication and presentation, part 1: general guidelines. ISO, Geneva, Switzerland. https://www.iso.org/standard/21832.html
International Organization for Standardization, (2007). ISO 13374-2:2007 Condition monitoring and diagnostics of machines, data processing, communication and presentation, part 2: data processing. ISO, Geneva, Switzerland. https://www.iso.org/standard/36645.html
International Organization for Standardization, (2016). ISO 14224:2016 Petroleum, petrochemical and natural gas industries, collection and exchange of reliability and maintenance data for equipment. ISO, Geneva, Switzerland. https://www.iso.org/standard/64076.html
International Organization for Standardization, (2018). ISO 26262:2018 Road vehicles, functional safety. ISO, Geneva, Switzerland. https://www.iso.org/standard/68383.html
International Organization for Standardization; SAE International, (2021). ISO/SAE 21434:2021 Road vehicles, cybersecurity engineering. ISO, Geneva, Switzerland. https://www.iso.org/standard/70918.html
Jain, M.; Vasdev, D.; Pal, K.; Sharma, V., (2022). Systematic literature review on predictive maintenance of vehicles and diagnosis of vehicle's health using machine learning techniques. Comput. Intell., 38(6): 1990-2008. https://doi.org/10.1111/coin.12553
Javed, A.R.; Usman, M.; Rehman, S.U.; Khan, M.U.; Haghighi, M.S., (2021). Anomaly detection in automated vehicles using multistage attention-based convolutional neural network. IEEE

Trans. Intell. Transp. Syst., 22(7): 4291-4300. https://doi.org/10.1109/TITS.2020.3025875

Jawhar, I.; Mohamed, N.; Usmani, H., (2013). An overview of inter-vehicular communication systems, protocols and middleware. J. Networks, 8(12): 2749-2761. https://doi.org/10.4304/jnw.8.12.2749-2761

Jha, K.K.; Jha, R.; Jha, A.K.; Hassan, M.A.; Yadav, S.K.; Mahesh, T., (2021). A brief comparison on machine learning algorithms based on various applications: a comprehensive survey. In: 2021 IEEE International Conference on Computation System and Information Technology for Sustainable Solutions, Bangalore, India. 1-5. https://doi.org/10.1109/CSITSS54238.2021.9683524

Kaur, K.; Verma, H.K., (2026). IoV-Health: an intelligent integrated emergency health monitoring and alert generation system. Soft Comput., 30(8): 5807-5823. https://doi.org/10.1007/s00500-023-08251-4

Khan, M.A.; Ali, M.H.; Haque, A.; Habib, M.T., (2022). A machine learning approach for driver identification based on CAN-Bus sensor data. arXiv preprint, arXiv:2207.10807. https://doi.org/10.48550/arXiv.2207.10807

Khoshkangini, R.; Tajgardan, M.; Lundstrom, J.; Rabbani, M.; Tegnered, D., (2023). A snapshot-stacked ensemble and optimization approach for vehicle breakdown prediction. Sensors, 23(12): 5621. https://doi.org/10.3390/s23125621

Kitchenham, B.; Charters, S., (2007). Guidelines for performing systematic literature reviews in software engineering. Technical Report EBSE-2007-01, Keele University and University of Durham, Keele, United Kingdom. https://www.elsevier.com/__data/promis_misc/525444systematicreviewsguide.pdf

Li, Z.; Zhang, K.; Chen, B.; Dong, Y.; Zhang, L., (2019). Driver identification in intelligent vehicle systems using machine learning algorithms. IET Intell. Transp. Syst., 13(1): 40-47. https://doi.org/10.1049/iet-its.2017.0254

Luo, F.; Wang, J.; Zhang, X.; Jiang, Y.; Li, Z.; Luo, C., (2023). In-vehicle network intrusion detection systems: a systematic survey of deep learning-based approaches. PeerJ Comput. Sci., 9: e1648. https://doi.org/10.7717/peerj-cs.1648

Machinery Information Management Open Systems Alliance, (2014). Open system architecture for condition-based maintenance, specification version 3.3.1. MIMOSA, Houston, United States. https://www.mimosa.org/mimosa-osa-cbm/

Magnani, G.; Gioia, D., (2023). Using the Gioia methodology in international business and entrepreneurship research. Int. Bus. Rev., 32(2): 102097. https://doi.org/10.1016/j.ibusrev.2022.102097

Martinelli, F.; Mercaldo, F.; Santone, A., (2020). Machine learning for driver detection through CAN bus. In: 2020 IEEE 91st Vehicular Technology Conference, Antwerp, Belgium. 1-5. https://ieeexplore.ieee.org/document/9129371

Mezair, T.; Djenouri, Y.; Belhadi, A.; Srivastava, G.; Lin, J.C., (2023). Towards an advanced deep learning for the Internet of Behaviors: application to connected vehicles. ACM Trans. Sens. Netw., 19(2): 1-18. https://doi.org/10.1145/3526192

Moleda, M.; Malysiak-Mrozek, B.; Ding, W.; Sunderam, V.; Mrozek, D., (2023). From corrective to predictive maintenance: a review of maintenance approaches for the power industry. Sensors, 23(13): 5970. https://doi.org/10.3390/s23135970

Moloisane, N.R.; Malekian, R.; Bogatinoska, D.C., (2017). Wireless machine-to-machine communication for intelligent transportation systems: Internet of Vehicles and vehicle to grid. In: 2017 40th International Convention on Information and Communication Technology, Electronics and Microelectronics, Opatija, Croatia. 411-415. https://ieeexplore.ieee.org/document/7973459

Muhammad, A.; Saqib, M.; Song, W.C., (2021). Sensor virtualization and data orchestration in Internet of Vehicles. In: 2021 IFIP/IEEE International Symposium on Integrated Network Management, Bordeaux, France. 998-1003. https://ieeexplore.ieee.org/document/9463961

Namuduri, S.; Narayanan, B.N.; Davuluru, V.S.; Burton, L.; Bhansali, S., (2020). Deep learning methods for sensor based predictive maintenance and future perspectives for electrochemical sensors. J. Electrochem. Soc., 167(3): 037552. https://doi.org/10.1149/1945-7111/ab67a8

Page, M.J.; McKenzie, J.E.; Bossuyt, P.M.; Boutron, I.; Hoffmann, T.C.; Mulrow, C.D.; Shamseer, L.;

Tetzlaff, J.M.; Akl, E.A.; Brennan, S.E.; Chou, R.; Glanville, J.; Grimshaw, J.M.; Hrobjartsson, A.; Lalu, M.M.; Li, T.; Loder, E.W.; Mayo-Wilson, E.; McDonald, S.; McGuinness, L.A.; Stewart, L.A.; Thomas, J.; Tricco, A.C.; Welch, V.A.; Whiting, P.; Moher, D., (2021). The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. BMJ, 372: n71. https://doi.org/10.1136/bmj.n71

Peppes, N.; Alexakis, T.; Adamopoulou, E.; Demestichas, K., (2021). Driving behaviour analysis using machine and deep learning methods for continuous streams of vehicular data. Sensors, 21(14): 4704. https://doi.org/10.3390/s21144704

Poor, P.; Zenisek, D.; Basl, J., (2019). Historical overview of maintenance management strategies: development from breakdown maintenance to predictive maintenance in accordance with four industrial revolutions. In: Proceedings of the International Conference on Industrial Engineering and Operations Management, Pilsen, Czech Republic. 495-504. http://hdl.handle.net/11025/36565

Prytz, R., (2014). Machine learning methods for vehicle predictive maintenance using off-board and on-board data. Licentiate thesis, Halmstad University Press, Halmstad, Sweden. https://www.diva-portal.org/smash/record.jsf?pid=diva2%3A757033

Rother, E.T., (2007). Systematic review versus narrative review. Acta Paul. Enferm., 20(2): 5-6. https://doi.org/10.1590/S0103-21002007000200001

SAE International, (2009). JA1011 Evaluation criteria for reliability-centered maintenance processes. SAE International, Warrendale, United States. https://www.sae.org/standards/content/ja1011_200908/

Shafi, U.; Safi, A.; Shahid, A.R.; Ziauddin, S.; Saleem, M.Q., (2018). Vehicle remote health monitoring and prognostic maintenance system. J. Adv. Transp., 2018: 8061514. https://doi.org/10.1155/2018/8061514

Snyder, H., (2019). Literature review as a research methodology: an overview and guidelines. J. Bus. Res., 104: 333-339. https://doi.org/10.1016/j.jbusres.2019.07.039

Solanki, V.K.; Dhall, R., (2017). An IoT based predictive connected car maintenance approach. Int. J. Interact. Multimedia Artif. Intell., 4(3): 16-22. https://doi.org/10.9781/ijimai.2017.433

State of California, (2018). California Consumer Privacy Act of 2018. California Civil Code, section 1798.100. Sacramento, United States. https://leginfo.legislature.ca.gov/faces/codes_displayText.xhtml?division=3.&part=4.&lawCode=CIV&title=1.81.5

Theofilatos, A.; Chen, C.; Antoniou, C., (2019). Comparing machine learning and deep learning methods for real-time crash prediction. Transp. Res. Rec., 2673(8): 169-178. https://doi.org/10.1177/0361198119841571

Turnbull, D.; Chugh, R.; Luck, J., (2023). Systematic-narrative hybrid literature review: a strategy for integrating a concise methodology into a manuscript. Soc. Sci. Humanit. Open, 7(1): 100381. https://doi.org/10.1016/j.ssaho.2022.100381

Ullah, S.; Kim, D.H., (2020). Lightweight driver behavior identification model with sparse learning on in-vehicle CAN-Bus sensor data. Sensors, 20(18): 5030. https://doi.org/10.3390/s20185030

United Nations Economic Commission for Europe, (2021a). UN Regulation No. 155, uniform provisions concerning the approval of vehicles with regards to cyber security and cyber security management system. UNECE, Geneva, Switzerland. https://unece.org/transport/documents/2021/03/standards/un-regulation-no-155-cyber-security-and-cyber-security

United Nations Economic Commission for Europe, (2021b). UN Regulation No. 156, uniform provisions concerning the approval of vehicles with regards to software update and software update management system. UNECE, Geneva, Switzerland. https://unece.org/transport/documents/2021/03/standards/un-regulation-no-156-software-update-and-software-update

Vasavi, S.; Aswarth, K.; Pavan, T.S.; Gokhale, A.A., (2021). Predictive analytics as a service for vehicle health monitoring using edge computing and AK-NN algorithm. Mater. Today Proc., 46: 8645-8654. https://doi.org/10.1016/j.matpr.2021.03.658

Wisesa, O.; Adriansyah, A.; Khalaf, O.I., (2020). Pre-

diction analysis for business to business (B2B) sales of telecommunication services using machine learning techniques. Majlesi J. Electr. Eng., 14(4): 145-153. https://doi.org/10.29252/mjee.14.4.145

Wu, L.; Fu, X.; Guan, Y., (2016). Review of the remaining useful life prognostics of vehicle lithium-ion batteries using data-driven methodologies. Appl. Sci., 6(6): 166. https://doi.org/10.3390/app6060166

Zhang, J.; Letaief, K.B., (2020). Mobile edge intelligence and computing for the Internet of Vehicles. Proc. IEEE, 108(2): 246-261. https://doi.org/10.1109/JPROC.2019.2947490

Zhu, T.; Ran, Y.; Zhou, X.; Wen, Y., (2019). A survey of predictive maintenance: systems, purposes and approaches. arXiv preprint, arXiv:1912.07383. https://doi.org/10.48550/arXiv.1912.07383

COPYRIGHTS

©2023 The author(s). This is an open access article distributed under the terms of the Creative Commons Attribution (CC BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, as long as the original authors and source are cited. No permission is required from the authors or the publishers.

HOW TO CITE THIS ARTICLE

Alborz, Y., Amirabadi, A., Amiri, N., Seyyedi, M. & Ahanian, I. (2026). A Compliance-oriented Reference Architecture for Predictive Maintenance in the Internet of Vehicles: Evidence from 45 Studies and Seven Standards. (e743374). International Journal of Urban Management and Energy Sustainability, 7(2), e743374

DOI: 10.22034/ijumes.2026.2108963.1423