The Future of Vehicle Maintenance: Revolutionizing Fleet Management through Auto Diagnosis, Digital Twins, IoT, and AI

Ekerete Usak

Department of Mechanical Engineering, Federal University of Technology Owerri Imo State

DOI: https://doi.org/10.51244/IJRSI.2026.1308000038

Received: 27 July 2026; Accepted: 01 August 2026; Published: 29 August 2026

ABSTRACT

Vehicle maintenance is transitioning from time-based and reactive repair toward continuous, evidence-based fleet health management. This transition is enabled by the convergence of onboard diagnostics, Internet of Things (IoT) telemetry, artificial intelligence (AI), and digital twins. This paper presents a practical reference architecture for intelligent fleet maintenance that integrates diagnostic trouble codes (DTCs), electronic control unit (ECU) signals, battery and drivetrain measurements, environmental context, workshop records, and supply-chain data. The architecture combines edge analytics for safety-critical anomaly detection, cloud-scale digital twins for state estimation and simulation, and AI models for fault classification, remaining useful life (RUL) estimation, maintenance prioritization, and parts planning. Digital twins should not be treated merely as visual replicas: their value lies in continuously reconciling physical-vehicle observations with physics-based and data-driven models to support counterfactual maintenance decisions. Recent literature identifies computational burden, data heterogeneity, and model complexity as persistent barriers to predictive-maintenance digital twins [6]. This paper addresses these barriers through a layered, standards-aware design; hybrid physics-informed AI; context-aware peer comparison; human-in-the-loop decision controls; and cybersecurity-by-design. It also proposes a phased implementation roadmap and a fleet maintenance scorecard. The central conclusion is that successful AI-enabled maintenance is not a standalone prediction problem: it is an operational decision system connecting vehicle condition, uncertainty, safety, maintenance capacity, parts availability, and economic consequence.

Keywords— Fleet management, predictive maintenance, vehicle diagnostics, digital twin, Internet of Things, artificial intelligence, prognostics and health management, remaining useful life, connected vehicles, automotive cybersecurity.

INTRODUCTION

Fleet operators face a structural maintenance challenge: vehicles operate under heterogeneous loads, climates, routes, driver behaviors, road conditions, payloads, and duty cycles. Consequently, fixed service intervals can result in both unnecessary maintenance and late intervention. A component that appears healthy according to mileage may be approaching failure because of thermal stress, harsh braking, repeated high-load operation, or a latent electrical anomaly. Conversely, a lightly used vehicle may not require service simply because a calendar threshold has been reached.

The modern alternative is **condition-based and predictive maintenance**, in which maintenance decisions are based on observed asset state, predicted degradation, operational risk, and service constraints. Digital twins provide the computational representation required to convert distributed vehicle data into a continuously updated condition model. The concept of a digital twin links a physical asset, a virtual representation, and information flows across the asset lifecycle [1]–[5]. In predictive-maintenance applications, the twin integrates sensor observations with rules, physics, historical maintenance records, and AI models to estimate current health and forecast future behavior [5], [6].

Automotive predictive maintenance is particularly demanding because a vehicle is not a stationary industrial asset. Its behavior is context dependent, and changes in measured signals do not necessarily imply faults. A 2023

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XIII Issue VIII August 2026

fleet study emphasized that environmental and operational context must be accounted for in vehicle anomaly detection; it also demonstrated the value of unsupervised, streaming approaches when labeled fault data are sparse [8]. Similarly, a review of machine learning in automotive maintenance found that public datasets remain limited, many methods rely on supervised labels, and interpretable methods are essential for operational adoption [7].

This paper makes four contributions:

1. It synthesizes auto diagnosis, IoT, AI, and digital twins into a unified fleet-maintenance architecture.
2. It distinguishes a true operational digital twin from a telemetry dashboard or a fault-code repository.
3. It proposes an implementable analytics-to-action workflow, including risk-based work-order orchestration.
4. It identifies the technical, organizational, safety, privacy, and cybersecurity controls necessary for real-world deployment.

From Reactive Repair to Twin-Enabled Fleet Maintenance

A. Maintenance paradigms

Fleet maintenance has evolved from corrective repair to time-based preventive maintenance, condition-based maintenance (CBM), and predictive maintenance (PdM). The latter does not replace preventive maintenance entirely; instead, it allocates maintenance resources according to actual degradation, failure consequence, and predicted lead time.

Table I. Comparison of Fleet Maintenance Strategies

Strategy	**Decision trigger**	**Typical data requirement**	**Strength**	**Principal limitation**
Reactive/corrective	Failure has occurred	Repair record, inspection	Low initial technology cost	Downtime, recovery cost, safety exposure
Time- or mileage-based preventive maintenance	Calendar date, odometer, engine hours	Usage history	Simple, auditable, widely established	Under- and over-maintenance
Condition-based maintenance	Threshold crossing or inspection finding	Sensor signals, diagnostics, inspection data	Responds to observed condition	Thresholds may not capture complex degradation
Predictive maintenance	Predicted fault probability/RUL	Multimodal longitudinal data, failure labels or peer context	Enables earlier planning and targeted intervention	Requires data quality, model governance, and workflow integration
Digital twin-enabled maintenance	Estimated state plus simulated future scenarios	Real-time telemetry, maintenance history, physics and AI models	Supports “what-if” analysis and fleet-level optimization	Higher integration, cybersecurity, and validation burden

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XIII Issue VIII August 2026

Condition-based maintenance is especially valuable when it is linked to reliability-centered maintenance and computerized maintenance management systems (CMMS/EAM). The output of analytics must become an actionable work order, not merely an alert on a dashboard [30]–[32].

B. Why a fleet is more than a collection of vehicles

A fleet creates a powerful comparative learning environment. Vehicles of the same platform, powertrain, component revision, and duty class can act as peer references. A temperature rise in a traction inverter may be normal during a mountain route but abnormal when compared with similar vehicles operating on equivalent routes and loads.

This enables a context-aware anomaly score:

$$A_{i,t} = D\left(x_{i,t},\ P_{i,t}\right)$$

Ai,t = D (xi,t, Pi,t)

where:

- $A_{i,t}$Ai,t is the anomaly score for vehicle ii at time tt;
- $x_{i,t}$xi,t is the vehicle state vector;
- $P_{i,t}$Pi,t is a peer group defined by platform, age, route, environment, payload, and operating mode;
- $D(\cdot)$D(·) is a distance-, density-, reconstruction-, or probability-based function.

This approach is preferable to applying one static threshold across an entire fleet. Fleet research has shown that contextual peer comparison can reduce maintenance cost relative to approaches that ignore operating context [8].

Reference Architecture for Intelligent Fleet Maintenance

A. Layered architecture

The proposed architecture combines diagnostic standards, edge intelligence, cloud data engineering, digital-twin services, AI models, and maintenance operations.

A conceptual illustration is developed with the intent to provide technology-neutral reference architecture rather than to represent a proprietary vehicle platform, OEM implementation, or vendor-specific software product. The diagram synthesizes the functional relationships among vehicle diagnostics, IoT connectivity, edge computing, fleet data management, digital twins, AI analytics, and maintenance execution systems.

Figure 1 presents the proposed end-to-end architecture for AI-enabled fleet maintenance. The framework is designed as a closed-loop cyber-physical maintenance system that connects physical vehicles, edge-based diagnostic intelligence, cloud-scale data infrastructure, digital-twin modeling, and human maintenance operations. At the vehicle level, heterogeneous information streams—including ECU fault codes, controller-area-network signals, battery and powertrain parameters, thermal measurements, tire data, and route-specific operating conditions—are collected through secure telematics gateways. Edge processing performs preliminary data validation, local fault-rule execution, event prioritization, and safety-critical alerting, thereby reducing latency and unnecessary data transmission.

Validated telemetry and contextual information are subsequently integrated with maintenance records, warranty data, vehicle configuration histories, parts availability, and workshop capacity in the fleet data platform. The digital twin transforms these data into a continuously updated representation of vehicle condition through state estimation, physics-based models, AI prognostic algorithms, and maintenance-scenario simulation. The resulting decision-intelligence layer estimates fault probability, remaining useful life, maintenance urgency, and resource requirements. Importantly, the architecture incorporates a feedback loop in which inspection outcomes, replaced

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XIII Issue VIII August 2026

components, technician observations, and confirmed repair results are returned to the data platform to improve diagnostic accuracy, prognostic calibration, and future maintenance decisions.

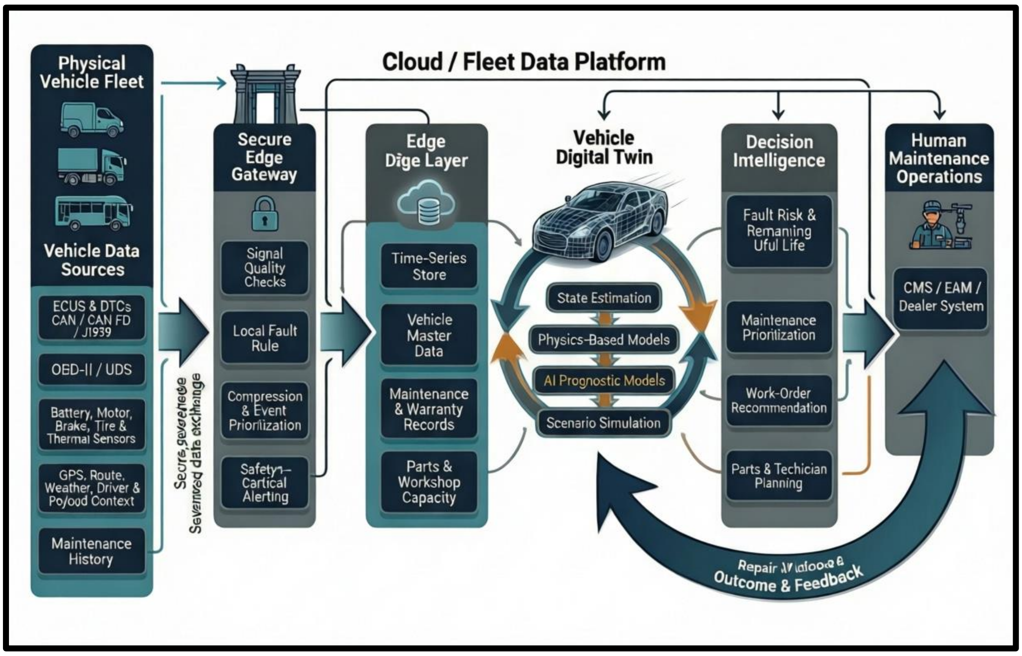

Figure 1. Conceptual Architecture for AI-Enabled Fleet Maintenance

The architecture should support a closed operational loop: In summary, AI-enabled fleet maintenance will be defined not by isolated diagnostic alerts or standalone predictive models, but by the ability to establish a reliable closed-loop connection between vehicle condition, digital-twin state estimation, maintenance decision-making, and verified repair outcomes. The architecture in Figure 1 demonstrates how auto diagnosis, IoT telemetry, cloud data infrastructure, AI prognostics, and human maintenance expertise can operate as an integrated system. When supported by secure data exchange, explainable analytics, standardized maintenance workflows, and continuous learning from real repair outcomes, this approach can reduce unplanned downtime, improve fleet availability, enhance safety, optimize parts and labor utilization, and support more sustainable vehicle operations.

The final "learn" stage is essential. Confirmed failures, mechanic notes, replaced parts, repair outcomes, and false-alarm adjudications must be returned to the fleet data platform to improve labels, calibration, and model validity.

B. Diagnostics and connectivity foundation

A reliable predictive-maintenance program begins with diagnostics interoperability. OBD communication between a vehicle and external test equipment is addressed by SAE J1979 and ISO 15031-5 [21], [22]. Unified Diagnostic Services (UDS), specified in ISO 14229-1, establishes data-link-independent diagnostic services for test equipment and ECUs [20]. ISO 15765 provides Diagnostic Communication over CAN (DoCAN), while SAE J1939 is important in heavy-duty and commercial-vehicle applications [23], [24].

Table II. Vehicle Data Domains and Their Maintenance Value

Data domain	Representative variables	Example maintenance use	Main data-quality issue
ECU diagnostics	DTCs, freeze-frame data, readiness monitors	Fault isolation and rule-based triage	DTC ambiguity and manufacturer-specific semantics

Powertrain telemetry	RPM, torque, boost pressure, fuel rate, exhaust temperature	Engine, transmission, turbocharger prognostics	Operating-condition dependence
Electric-vehicle battery data	Cell voltage, current, temperature, impedance, state of charge	State-of-health and RUL estimation	Sensor drift, balancing effects, partial charging
Braking and chassis	Pad wear, wheel speed, brake pressure, ABS events	Brake and suspension condition monitoring	Route and driver behavior confounding
Tire and wheel	Pressure, temperature, tread depth, vibration	Tire failure prevention and alignment detection	Sensor availability and tire replacement history
Thermal management	Coolant temperature, pump speed, fan current, HVAC load	Cooling system and inverter protection	Seasonal and ambient-temperature effects
Contextual data	Route, grade, weather, payload, traffic, driver events	Context-aware anomaly detection	Missing or inconsistent external data
Maintenance records	Work order, labor, parts, technician notes, warranty claim	Labeling, root-cause analysis, cost optimization	Unstructured text and inconsistent coding

Digital Twins as Maintenance Decision Systems

A. Defining the operational vehicle digital twin

A fleet digital twin is not simply a 3D vehicle model, dashboard, or data lake. It is a dynamic computational representation that maintains a traceable relationship with a specific physical vehicle and its operational environment. The twin should contain at least five components:

1. **Asset identity:** vehicle identification number, configuration, build date, software version, component serial numbers, and modification history.
2. **State model:** current health indicators, DTC states, battery state of health, thermal condition, mileage, operating mode, and estimated degradation.
3. **Behavior model:** physics-based, statistical, or hybrid models describing component behavior.
4. **Prediction model:** fault probability, anomaly score, confidence interval, and RUL distribution.
5. **Decision model:** maintenance priority, recommended action, timing, required parts, expected downtime, and economic consequence.

A digital twin may use physics-based models, data-driven models, or hybrid models. Physics-based models are useful when underlying degradation mechanisms are well understood, such as battery electrochemistry, thermal management, brake wear, and mechanical vibration. Data-driven AI is valuable when complex multi-sensor patterns are present but complete physical models are impractical. Hybrid models are often most appropriate because they constrain AI predictions with physical plausibility [5], [6].

Figure 2 illustrates the proposed digital-twin feedback cycle for fleet vehicle maintenance. The framework establishes a continuous bidirectional connection between the physical vehicle and its virtual representation. Vehicle-level operational conditions—including route characteristics, payload, ambient weather, and driver behavior—affect the performance and degradation behavior of key vehicle systems. These contextual factors are interpreted alongside telemetry streams and diagnostic information obtained from electronic control units,

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XIII Issue VIII August 2026

controller-area-network communications, battery-management systems, thermal-management systems, tire sensors, and diagnostic trouble codes.

Following data acquisition, telemetry is subjected to signal-quality verification and contextualization to distinguish genuine degradation signatures from normal operational variability. The validated data update the digital-twin state estimator, which combines physics-based component models with AI-driven diagnostic and prognostic algorithms. This process generates dynamic estimates of vehicle health, failure probability, and remaining useful life. The resulting outputs inform scenario-based maintenance decisions, including immediate intervention, enhanced monitoring, or planned service deferral.

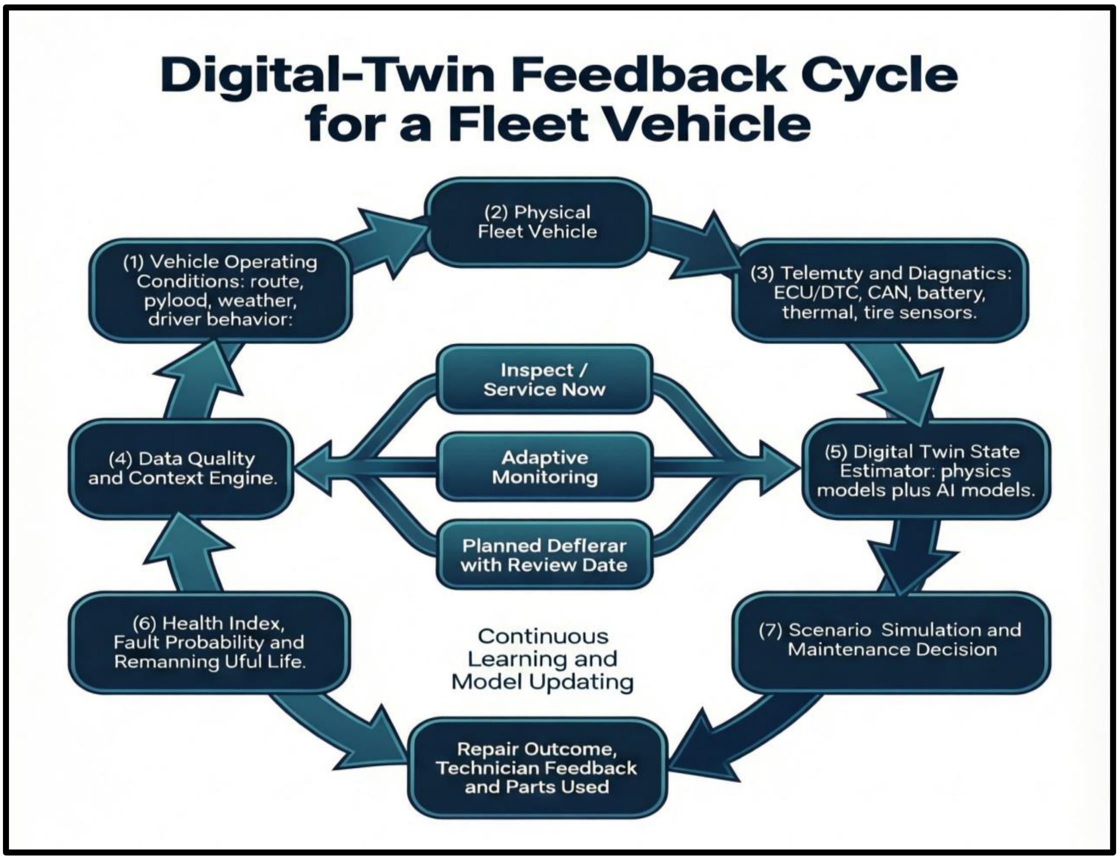

Figure 2. Digital-twin feedback cycle for a fleet vehicle

The feedback loop is completed through the integration of repair outcomes, technician observations, replaced-part records, and post-maintenance verification data. These real-world outcomes refine both the digital-twin state and the underlying AI models, thereby enabling continuous learning and progressive improvement in maintenance precision, fleet availability, and operational reliability

Figure 2 emphasizes that the value of a vehicle digital twin lies not only in its ability to represent current asset condition, but also in its capacity to learn continuously from maintenance outcomes. By connecting operating context, diagnostic telemetry, AI-based prognostics, scenario simulation, maintenance execution, and technician feedback, the digital twin becomes a living decision system rather than a static digital replica. This closed-loop approach enables fleet operators to move beyond reactive repair and fixed service schedules toward safer, more reliable, cost-effective, and evidence-based maintenance management.

B. Health indices and remaining useful life

A health index (HI) summarizes component condition on a defined scale:

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XIII Issue VIII August 2026

$$HI_t = \sum_{k=1}^{m} w_k z_{k,t}$$

HIt = k=1∑m wkzk,t

where $z_{k,t}$zk,t represents a normalized condition feature and $w_k$wk is a learned or expert-defined weight. For example, an electric-vehicle battery HI may incorporate capacity fade, internal resistance, temperature exposure, charging profile, and cell-voltage imbalance.

RUL should be treated as a probability distribution, not a single deterministic number:

$$P(RUL_{i,t} \leq h \mid x_{i,1:t})$$

P(RULi,t ≤ h ∣ xi,1:t)

This representation allows maintenance teams to compare uncertainty with operational consequence. A component with an estimated RUL of 30 days but very high uncertainty may require earlier inspection than one with a 20-day RUL and narrow uncertainty bounds.

Digital-twin-based battery-life prediction has demonstrated the potential of adaptive models that update as degradation data accumulate [17]. However, fleet deployment requires careful distinction between laboratory accuracy and real-world performance under heterogeneous driving and charging conditions.

AI for Auto Diagnosis and Prognostics

A. Core analytics tasks

AI should be selected according to the maintenance decision, data availability, and acceptable false-alarm burden—not because a specific algorithm is fashionable.

Table III. AI Methods Matched to Fleet-Maintenance Tasks

Maintenance task	Suitable methods	Primary output	Operational requirement
DTC interpretation and diagnosis	Rule engines, knowledge graphs, NLP, gradient boosting	Probable fault mode and recommended inspection	Diagnostic ontology and OEM knowledge
Sensor anomaly detection	Isolation forest, autoencoders, one-class models, transformers	Deviation score	Context-aware baselines
Fault classification	Random forest, XGBoost, CNN, LSTM	Failure category	High-quality labels
RUL prediction	Survival models, Kalman filters, LSTM, temporal transformers, hybrid models	RUL distribution	Longitudinal failure histories
Root-cause support	Causal models, Bayesian networks, SHAP, knowledge graphs	Contributing variables	Explainability and technician review
Maintenance scheduling	Mixed-integer optimization, reinforcement learning, simulation	Timing, resource, and parts plan	Cost and capacity constraints
Fleet learning	Federated learning, transfer learning, domain adaptation	Shared models without raw-data pooling	Privacy, security, and drift monitoring

Machine-learning methods for mechanical fault diagnosis span supervised, semi-supervised, and unsupervised approaches [9], [10]. Deep learning may improve feature extraction from complex time-series signals but can be difficult to validate, calibrate, and explain [11]. Transformer-based anomaly-detection approaches are increasingly useful for multivariate temporal data [13], though their deployment should be justified by measurable value over simpler models.

B. Context-aware anomaly detection

In fleet settings, an anomaly is not merely an unusual signal; it is an observation unusual **relative to comparable vehicles in comparable conditions**. A practical peer group may be stratified by:

- Vehicle platform and component revision;
- Powertrain type;
- Vehicle age and cumulative energy throughput;
- Route class and elevation profile;
- Ambient temperature;
- Payload class;
- Driver behavior indicators;
- Software/calibration version.

The fleet maintenance literature indicates that context-aware, unsupervised anomaly detection can be valuable where failures are rare and labels are incomplete [8]. This is particularly important because real maintenance datasets are often highly imbalanced: thousands of normal operating hours may precede a small number of confirmed failure events.

C. Explainability and human oversight

Maintenance personnel require more than a red-risk alert. They require a defensible reason for action, such as:

- "Battery module 6 shows persistent voltage dispersion relative to matched peer vehicles."
- "Turbocharger risk increased because boost-pressure deviation, exhaust-gas temperature, and repeated underboost DTCs co-occurred."
- "Brake-pad wear is projected to cross the service threshold before the next planned depot visit."

Local Interpretable Model-Agnostic Explanations (LIME) and Shapley Additive Explanations (SHAP) can help attribute model outputs to input features [14], [15]. However, explainability does not prove causality. It should support technician reasoning, engineering review, and root-cause investigation rather than replace them.

Maintenance Decision Optimization

A. From prediction to action

A highly accurate model can still fail operationally if it does not consider maintenance windows, safety risk, spare parts, workshop capacity, and service disruption. The decision layer should rank opportunities by expected risk-adjusted value:

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XIII Issue VIII August 2026

$$Priority_{i,t} = P(F_{i,t:t+h} \mid x_{i,t}) \times C_{failure,i} - \left(C_{planned,i} + C_{downtime,i} + C_{false\ alarm,i}\right)$$

Priorityi,t = P(Fi,t:t+h ∣ xi,t) × Cfailure,i − (Cplanned,i + Cdowntime,i + Cfalse alarm,i)

where:

- $P(F_{i,t:t+h})$P(Fi,t:t+h) is the predicted probability of failure within horizon hh;
- $C_{failure,i}$Cfailure,i is the consequence of in-service failure;
- $C_{planned,i}$Cplanned,i is the cost of planned service;
- $C_{downtime,i}$Cdowntime,i is operational downtime cost;
- $C_{false\ alarm,i}$Cfalse alarm,i captures unnecessary inspections or replacements.

A safety-critical component should not be managed exclusively by economic optimization. For braking, steering, high-voltage isolation, thermal runaway precursors, and other hazards, risk thresholds should be governed by safety engineering and approved maintenance policy.

B. Maintenance workflow

Risk-to-work-order workflow for AI-enabled fleet maintenance.

Vehicle telemetry, diagnostic trouble codes (DTCs), operating context, and historical maintenance records are first subjected to data-validation and sensor-quality checks. Validated inputs update the vehicle digital-twin state and associated uncertainty estimate, which are used to generate fault-probability, anomaly, and remaining-useful-life (RUL) predictions. A safety-rule gate separates safety-critical conditions requiring immediate inspection or vehicle hold from non-critical conditions that proceed to economic and operational prioritization. The resulting maintenance recommendation is evaluated against parts availability, workshop capacity, and route or depot scheduling constraints before a work order is issued. Technician confirmation, repair records, and feedback labels are returned to the digital twin, enabling continuous model learning and refinement of future maintenance decisions.

Figure 3 presents the proposed risk-to-work-order workflow that converts vehicle health information into feasible and auditable maintenance action. The workflow begins by integrating telemetry streams, diagnostic trouble codes, contextual operating variables, and maintenance-history data. Before analytical interpretation, these data undergo validation and sensor-quality assessment to identify missing values, communication errors, implausible readings, and context-dependent variations that could otherwise produce unreliable maintenance alerts.

The validated information updates the digital-twin representation of each vehicle, including its estimated health state and associated uncertainty. AI-enabled diagnostic and prognostic models then generate fault-probability scores, anomaly indicators, and remaining-useful-life estimates. Crucially, the workflow introduces a safety-rule gate before economic optimization. Safety-critical conditions, such as high-voltage isolation concerns, brake-system warnings, steering abnormalities, or severe thermal events, trigger immediate inspection or vehicle hold in accordance with fleet safety policy. Non-critical cases proceed to risk-based prioritization, where predicted failure consequences, planned-service costs, expected downtime, false-alarm burden, parts availability, workshop capacity, and route constraints are considered jointly.

The final output is a recommended work order supported by diagnostic evidence and confidence information. Technician confirmation, repair findings, replaced-component records, and post-repair verification are subsequently incorporated into the fleet data platform and digital twin. This feedback mechanism transforms the maintenance workflow from a one-way alerting process into a continuously learning operational system.

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XIII Issue VIII August 2026

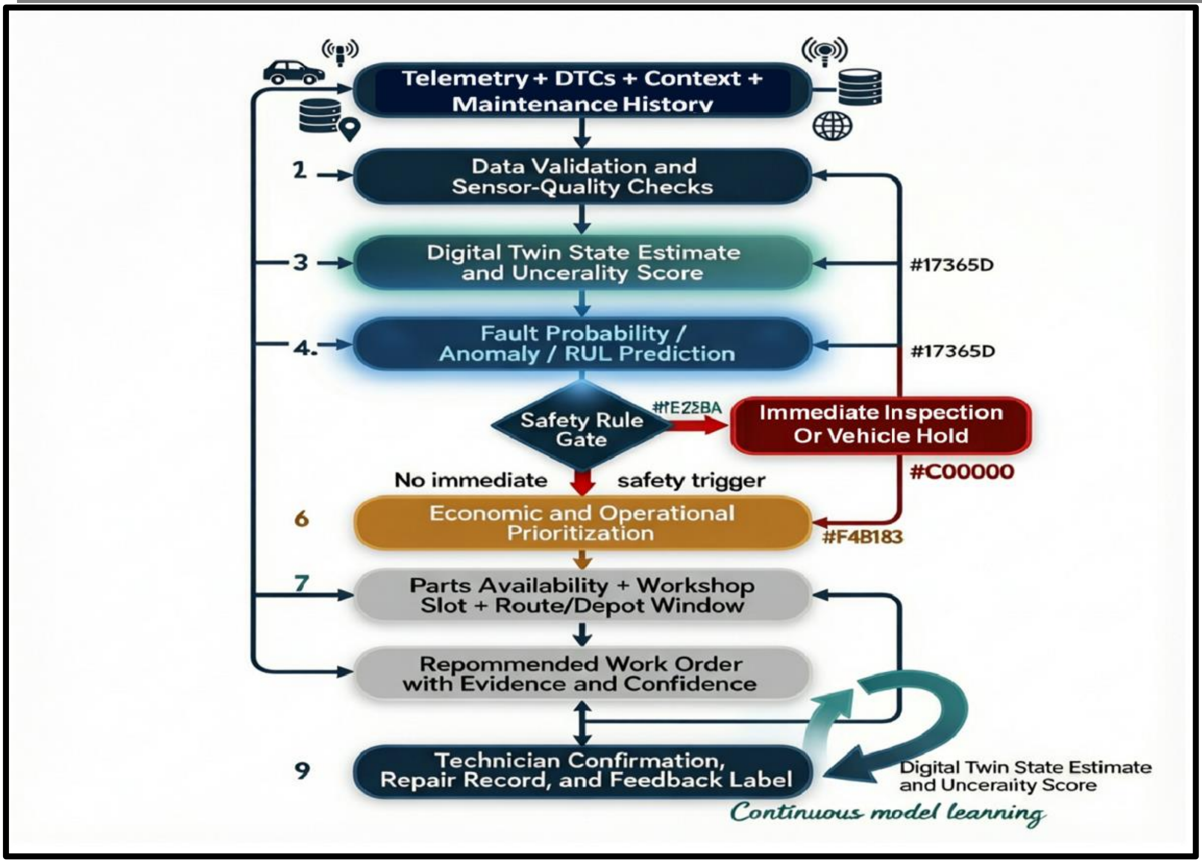

Figure 3. Risk-to-Work-Order Workflow

Figure 3 demonstrates that fleet maintenance is not improved merely by identifying anomalous signals or predicting component failure. Meaningful value arises when risk estimates are translated into safe, resource-aware, and evidence-supported work orders. By integrating safety-rule screening, digital-twin state estimation, AI-based prognosis, operational prioritization, workshop scheduling, and technician feedback, the proposed workflow closes the gap between predictive analytics and real maintenance execution. This closed-loop approach enables fleet operators to reduce unplanned downtime while preserving safety, improving maintenance traceability, optimizing labor and parts utilization, and continually enhancing the reliability of AI-enabled fleet-management systems.

Table IV. Fleet Maintenance Performance Scorecard

KPI category	Example metric	Why it matters
Reliability	Mean distance/time between failures; roadside failures per 100,000 km	Measures service continuity
Availability	Fleet availability; unplanned downtime hours	Reflects operational value
Prediction quality	Precision, recall, F1 score, false-alert rate, calibration error	Avoids misleading accuracy claims
Prognostics	RUL MAE, prediction-interval coverage, warning lead time	Measures actionable early warning
Economics	Avoided roadside cost, planned-to-unplanned maintenance ratio, parts expediting cost	Connects analytics to value

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XIII Issue VIII August 2026

Safety	Safety-critical alerts reviewed within SLA; post-repair recurrence rate	Ensures risk control
Maintenance execution	Work-order completion time; first-time fix rate; parts fill rate	Measures organizational readiness
AI governance	Drift alerts, percentage of explanations reviewed, model rollback events	Maintains reliable model operations
Sustainability	Fuel/energy impacts, tire-life extension, part reuse/remanufacturing rate	Captures lifecycle benefit

Interoperability, Cybersecurity, Safety, and Governance

A. Standards-aware interoperability

Interoperability should be designed from the outset. Relevant standards include ISO 14229 for UDS, ISO 15031 and SAE J1979 for OBD communication, ISO 15765 for diagnostics over CAN, and SAE J1939 for heavy-duty vehicle networks [20]–[24]. IoT gateways commonly use lightweight publish-subscribe protocols such as MQTT, while interoperability requires durable semantic mapping between OEM signals, component identities, maintenance codes, and work orders [33], [34].

B. Cybersecurity and software updates

Connected maintenance expands the attack surface. An attacker who manipulates telemetry, diagnostic access, over-the-air updates, or maintenance records could cause erroneous maintenance actions or conceal genuine degradation. Therefore, the architecture should enforce:

- Mutual authentication for vehicle gateways and cloud services;
- Cryptographic integrity protection for telemetry and software packages;
- Secure key management;
- Network segmentation between infotainment, telematics, and safety-critical domains;
- Role-based access control for diagnostics and maintenance data;
- Secure logging, auditability, and incident response;
- Continuous vulnerability management for gateways and cloud dependencies.

ISO/SAE 21434 specifies cybersecurity engineering requirements for road-vehicle electrical and electronic systems across the lifecycle, from concept through decommissioning [25]. It should be deployed alongside functional-safety practices under ISO 26262, software-update governance under UNECE R155/R156, and industrial-security controls such as IEC 62443 [26]–[29].

C. AI governance and privacy

AI governance should include dataset lineage, model versioning, performance monitoring, drift detection, independent validation, and rollback procedures. Technical debt in machine-learning systems often arises from hidden dependencies, undeclared data assumptions, and fragile pipelines [37]. The NIST AI Risk Management Framework provides a useful structure for govern-map-measure-manage practices [38].

Fleet data may contain personal or commercially sensitive information, including driver behavior, location, depot activity, customer routes, and proprietary vehicle signals. Federated learning can reduce raw-data

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XIII Issue VIII August 2026

centralization by training local models and sharing updates rather than source data [35], [36]. However, federated learning is not automatically private: model updates can still leak information unless combined with secure aggregation, access controls, and privacy-preserving methods.

Phased Implementation Roadmap

Phase 1: Establish the data and maintenance baseline

The first objective is not advanced AI. It is reliable asset identity, consistent work-order coding, DTC normalization, secure connectivity, and measurable baseline KPIs. Select one high-value component family, such as batteries, compressors, braking systems, tires, or cooling loops.

Phase 2: Deploy descriptive and condition-based analytics

Implement dashboards for DTC frequency, fleet peer comparison, threshold alerts, service compliance, and diagnostic completeness. Technicians should validate whether alerts correspond to observable conditions.

Phase 3: Introduce predictive models

Use a controlled pilot with clearly defined labels and maintenance outcomes. Start with interpretable models, simple baselines, and cost-sensitive evaluation. A model should be compared against current maintenance policy, not only against other AI models.

Phase 4: Build the digital twin and decision layer

Introduce component health states, hybrid models, RUL uncertainty, maintenance-scenario simulation, and work-order prioritization. Link the twin to parts inventory and workshop planning.

Phase 5: Scale through governance and continuous learning

Scale only after demonstrating safety, economic value, workflow adoption, and cybersecurity maturity. Monitor population drift caused by vehicle aging, firmware updates, new component suppliers, route changes, and electrification.

Table V. Practical Adoption Roadmap

Phase	**Main deliverable**	**Success criterion**	**Common failure mode**
1. Foundation	Unified vehicle and maintenance data model	Trusted asset and work-order records	Inconsistent identifiers and poor historical data
2. Visibility	Condition-monitoring dashboard	Technicians accept signal relevance	Dashboard without workflow ownership
3. Prediction	Validated fault-risk/RUL model	Better lead time than existing policy	Optimizing accuracy rather than maintenance cost
4. Twin-enabled optimization	Digital twin and maintenance decision engine	Risk-ranked work orders and fewer disruptions	Treating the twin as a visualization project
5. Enterprise scale	Governance, MLOps, cybersecurity controls	Reliable multi-site deployment	Model drift and uncontrolled vendor dependencies

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XIII Issue VIII August 2026

DISCUSSION

Real-World Applications and Research Priorities

A. Electric fleets

Electric fleets present major opportunities because battery, inverter, thermal, charging, and energy-use signals are digitally accessible. Battery digital twins can estimate state of health, detect cell imbalance, predict thermal stress, and schedule service before range loss or thermal derating disrupts operations [17]. Nevertheless, models must account for charging habits, climate, depth of discharge, battery chemistry, and software updates.

B. Commercial trucks, buses, and logistics fleets

Heavy-duty fleets can benefit from telematics, J1939 signals, air-system data, brake wear, tire monitoring, fuel-system diagnostics, and route context. The value proposition is often strongest where downtime causes missed service commitments, recovery costs, spare-vehicle requirements, or safety exposure. Context-aware fleet anomaly detection is particularly suitable for buses and logistics fleets because multiple comparable vehicles generate parallel operational data [8], [19], [20].

C. Autonomous and software-defined vehicles

As vehicles become more software-defined, maintenance increasingly includes software configuration, cybersecurity patches, sensor calibration, and confidence monitoring for automated-driving systems. The digital twin must therefore represent software versions, calibration states, and cybersecurity posture in addition to mechanical condition. AI-enabled diagnosis must be governed rigorously because incorrect recommendations can affect vehicle safety and regulatory compliance [25], [26], [39], [40].

D. Research gaps

Important research needs include:

1. **Cross-OEM data semantics:** standardized mappings from proprietary signals to maintenance-relevant concepts.
2. **Hybrid prognostics:** tighter integration of physical degradation models and AI.
3. **Uncertainty-aware RUL:** maintenance decisions should account for confidence intervals, not only point predictions.
4. **Causal maintenance intelligence:** separating correlation from true failure mechanisms.
5. **Human factors:** understanding how technicians interpret, trust, override, and improve AI recommendations.
6. **Fleet-scale validation:** prospective studies measuring downtime, safety, maintenance cost, and false-alarm burden.
7. **Sustainable maintenance:** linking predictive interventions to remanufacturing, parts circularity, energy efficiency, and lifecycle emissions.

CONCLUSION

The future of vehicle maintenance is a transition from isolated repairs to continuously learning fleet-health ecosystems. Auto diagnosis supplies standardized fault evidence; IoT provides persistent sensing and connectivity; AI detects patterns, forecasts degradation, and prioritizes intervention; and digital twins integrate these capabilities into an operational model of each vehicle and fleet.

The decisive innovation is not simply predicting that a component may fail. It is converting uncertain predictions into safe, explainable, economically rational, and operationally feasible maintenance actions. A mature system must integrate vehicle diagnostics, contextual fleet data, digital-twin state estimation, hybrid AI models, workshop capacity, parts availability, cybersecurity controls, and technician feedback.

Fleet operators should begin with data governance, diagnostic interoperability, and one high-value use case. They should evaluate success through reduced unplanned downtime, improved fleet availability, safer operations, credible maintenance lead time, and technician adoption—not model accuracy alone. With these conditions in place, digital twins and AI can transform maintenance from a cost center reacting to failures into a strategic capability for reliability, safety, service continuity, and sustainability.

REFERENCES

1. M. Grieves and J. Vickers, "Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems," in Transdisciplinary Perspectives on Complex Systems, F.-J. Kahlen, S. Flumerfelt, and A. Alves, Eds. Cham, Switzerland: Springer, 2017, pp. 85–113.
2. W. Kritzinger, M. Karner, G. Traar, J. Henjes, and W. Sihn, "Digital twin in manufacturing: A categorical literature review and classification," IFAC-PapersOnLine, vol. 51, no. 11, pp. 1016–1022, 2018.
3. F. Tao, H. Zhang, A. Liu, and A. Y. C. Nee, "Digital twin in industry: State-of-the-art," IEEE Transactions on Industrial Informatics, vol. 15, no. 4, pp. 2405–2415, 2019.
4. A. Fuller, Z. Fan, C. Day, and C. Barlow, "Digital twin: Enabling technologies, challenges and open research," IEEE Access, vol. 8, pp. 108952–108971, 2020.
5. I. Errandonea, S. Beltrán, and S. Arrizabalaga, "Digital twin for maintenance: A literature review," Computers in Industry, vol. 123, Art. no. 103316, 2020.
6. R. van Dinter, B. Tekinerdogan, and C. Catal, "Predictive maintenance using digital twins: A systematic literature review," Information and Software Technology, vol. 151, Art. no. 107008, 2022.
7. A. Theissler, J. Pérez-Velázquez, M. Kettelgerdes, and G. Elger, "Predictive maintenance enabled by machine learning: Use cases and challenges in the automotive industry," Reliability Engineering & System Safety, vol. 215, Art. no. 107864, 2021.
8. A. Giannoulidis and A. Gounaris, "A context-aware unsupervised predictive maintenance solution for fleet management," Journal of Intelligent Information Systems, vol. 60, pp. 521–547, 2023.
9. T. P. Carvalho, F. A. A. M. N. Soares, R. Vita, R. P. Francisco, J. P. Basto, and S. G. S. Alcalá, "A systematic literature review of machine learning methods applied to predictive maintenance," Computers & Industrial Engineering, vol. 137, Art. no. 106024, 2019.
10. Y. Lei, B. Yang, X. Jiang, F. Jia, N. Li, and A. K. Nandi, "Applications of machine learning to machine fault diagnosis: A review and roadmap," Mechanical Systems and Signal Processing, vol. 138, Art. no. 106587, 2020.
11. O. Fink, Q. Wang, M. Svensén, P. Dersin, W.-J. Lee, and M. Ducoffe, "Potential, challenges and future directions for deep learning in prognostics and health management applications," Engineering Applications of Artificial Intelligence, vol. 92, Art. no. 103678, 2020.
12. G. Pang, C. Shen, L. Cao, and A. van den Hengel, "Deep learning for anomaly detection: A review," ACM Computing Surveys, vol. 54, no. 2, pp. 1–38, 2021.
13. S. Tuli, G. Casale, and N. R. Jennings, "TranAD: Deep transformer networks for anomaly detection in multivariate time series data," Proceedings of the VLDB Endowment, vol. 15, no. 6, pp. 1201–1214, 2022.
14. S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in Advances in Neural Information Processing Systems, vol. 30, 2017.
15. M. T. Ribeiro, S. Singh, and C. Guestrin, "'Why should I trust you?' Explaining the predictions of any classifier," in Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2016, pp. 1135–1144.
16. P. K. Rajesh, N. Manikandan, C. S. Ramshankar, T. Vishwanathan, and C. Sathishkumar, "Digital twin of an automotive brake pad for predictive maintenance," Procedia Computer Science, vol. 165, pp. 18–24, 2019.

17. A. P. Das, K. Thomas, A. Mohan, A. S. Akhil, and K. P. Ganesh, "A battery digital twin framework for predictive maintenance and state of health estimation of electric vehicles," VDI-Berichte, vol. 2021, pp. 579–592, 2021.
18. Q. Feng, Y. Zhang, B. Sun, X. Guo, D. Fan, Y. Ren, Y. Song, and Z. Wang, "Multi-level predictive maintenance of smart manufacturing systems driven by digital twin: A matheuristics approach," Journal of Manufacturing Systems, vol. 68, pp. 443–454, 2023.
19. T. Rögnvaldsson, S. Nowaczyk, S. Byttner, R. Prytz, and M. Svensson, "Self-monitoring for maintenance of vehicle fleets," Data Mining and Knowledge Discovery, vol. 32, no. 2, pp. 344–384, 2018.
20. Y. Fan, S. Nowaczyk, and T. Rögnvaldsson, "Evaluation of self-organized approach for predicting compressor faults in a city bus fleet," Procedia Computer Science, vol. 53, pp. 447–456, 2015.
21. International Organization for Standardization, ISO 14229-1:2020: Road Vehicles—Unified Diagnostic Services (UDS)—Part 1: Application Layer. Geneva, Switzerland: ISO, 2020.
22. International Organization for Standardization, ISO 15031-5:2015: Road Vehicles—Communication Between Vehicle and External Equipment for Emissions-Related Diagnostics—Part 5: Emissions-Related Diagnostic Services. Geneva, Switzerland: ISO, 2015.
23. SAE International, SAE J1979: E/E Diagnostic Test Modes. Warrendale, PA, USA: SAE International, 2017.
24. International Organization for Standardization, ISO 15765-4:2016: Road Vehicles—Diagnostic Communication Over Controller Area Network (DoCAN). Geneva, Switzerland: ISO, 2016.
25. International Organization for Standardization and SAE International, ISO/SAE 21434:2021: Road Vehicles—Cybersecurity Engineering. Geneva, Switzerland: ISO, 2021.
26. International Organization for Standardization, ISO 26262:2018: Road Vehicles—Functional Safety. Geneva, Switzerland: ISO, 2018.
27. United Nations Economic Commission for Europe, UN Regulation No. 155: Cyber Security and Cyber Security Management System. Geneva, Switzerland: UNECE, 2021.
28. International Electrotechnical Commission, IEC 62443-3-3:2013: Industrial Communication Networks—Network and System Security—System Security Requirements and Security Levels. Geneva, Switzerland: IEC, 2013.
29. National Institute of Standards and Technology, The NIST Cybersecurity Framework (CSF) 2.0, NIST CSWP 29, 2024.
30. International Organization for Standardization, ISO 55000:2024: Asset Management—Vocabulary, Overview and Principles. Geneva, Switzerland: ISO, 2024.
31. International Organization for Standardization, ISO 13374-1:2019: Condition Monitoring and Diagnostics of Machines—Data Processing, Communication and Presentation—Part 1: General Guidelines. Geneva, Switzerland: ISO, 2019.
32. International Organization for Standardization, ISO 17359:2018: Condition Monitoring and Diagnostics of Machines—General Guidelines. Geneva, Switzerland: ISO, 2018.
33. International Organization for Standardization, ISO 23247-1:2021: Automation Systems and Integration—Digital Twin Framework for Manufacturing—Part 1: Overview and General Principles. Geneva, Switzerland: ISO, 2021.
34. OASIS, MQTT Version 5.0, OASIS Standard, 2019.
35. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, 2017, pp. 1273–1282.
36. P. Kairouz et al., "Advances and open problems in federated learning," Foundations and Trends in Machine Learning, vol. 14, no. 1–2, pp. 1–210, 2021.
37. D. Sculley et al., "Hidden technical debt in machine learning systems," in Advances in Neural Information Processing Systems, vol. 28, 2015.
38. National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1, 2023.
39. European Parliament and Council of the European Union, "Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence," Official Journal of the European Union, 2024.
40. International Organization for Standardization, ISO 21448:2022: Road Vehicles—Safety of the Intended Functionality. Geneva, Switzerland: ISO, 2022.