American Journal of Economics and Business Innovation (AJEBI)

ISSN: 2831-5588 (ONLINE), 2832-4862 (PRINT)

VOLUME 5 ISSUE 3 (2026)



PUBLISHED BY
E-PALLI PUBLISHERS, DELAWARE, USA

American Journal of Economics and Business Innovation (AJEBI)

Volume 5 Issue 3, Year 2026
ISSN: 2831-5588 (Online), 2832-4862 (Print)
DOI: https://doi.org/10.54536/ajebi.v5i3.8301
https://journals.e-palli.com/home/index.php/ajebi

Prioritization of Rental Vehicle Fleet Maintenance: From Predictive Condition Indexes to Autonomous Intervention Solutions

Alexey Viktorovich Gromov[1*]

Article Information

Received: December 25, 2025

Accepted: August 08, 2026

Published: November 16, 2026

Keywords

Autonomous Maintenance Intervention, Condition-Based Maintenance, Predictive Maintenance, Remaining Useful Life Estimation, Rental Fleet Management

ABSTRACT

Predictive maintenance research for vehicle fleets has advanced substantially in health index construction, remaining useful life (RUL) estimation, and uncertainty quantification. The operational translation of these outputs into scheduling decisions has received comparatively little structured attention in rental and car-sharing contexts. A rental vehicle changes drivers with every transaction, and its technical condition must be reconciled with reservation calendars, geographic coverage, and revenue considerations that have no analogue in owner-operated or logistics fleets. This paper synthesizes prognostics, health management, and fleet-optimization literature to characterize the gap between condition-index generation and autonomous intervention scheduling. Existing approaches are organized by the degree of operational context integration, the decision mechanism employed, and the scale at which validation has been reported. A composite scoring framework for rental fleets, validated across 847 vehicles and 2.31 million vehicle-hours of operation, is examined as the only located case combining all four rental-specific context variables into a single computable score. Condition-index generation has matured faster than the scheduling logic needed to act on it. Interoperability and data-governance constraints form the primary barrier to deployment at scale.

INTRODUCTION

Onboard diagnostics, known as OBD-II, has served as the default layer for detecting vehicle component degradation since the standard became mandatory in passenger cars sold from 1996 onward. The design is structurally reactive. A diagnostic trouble code appears only once a monitored parameter crosses a fixed threshold. By that point the underlying wear process is often close to failure. Brake pad wear illustrates the pattern well: a stored code for pad depletion typically appears only in the final days of the component's service life, long after the wear process itself began (Theissler *et al.*, 2021). For a privately owned vehicle this delay is an inconvenience. For a rental vehicle it is a distinct operational problem, because the asset sits in active customer custody during the highest-risk segment of its degradation curve, and the customer has no visibility into the underlying condition data.

The operational stakes of this delay are specific to the rental business model. A component failure that would simply generate a service ticket for an owner-operated vehicle produces, for a rental operator, an unplanned vehicle-off-road event, a customer relocation or replacement-vehicle assignment, and potential liability exposure if a safety-critical component degrades further before the customer returns the vehicle. Cost pressures compound this exposure. Unplanned corrective maintenance is consistently more expensive per event than scheduled preventive maintenance across the automotive predictive-maintenance literature (Mahale *et al.*, 2025), and a vehicle removed from a rental fleet without advance notice reduces available inventory precisely when reservation demand tends to be highest. These consequences give rental operators a direct commercial incentive to close the gap between condition-index generation and scheduling action, independent of the safety case alone.

Closing this gap requires more than incremental gains in estimation accuracy. The sections that follow review how condition-based maintenance research has evolved toward vehicle-specific and uncertainty-calibrated methods, describe the search and classification methodology used to organize that literature, and then examine how far the resulting decision mechanisms integrate the operational context that a rental fleet imposes.

LITERATURE REVIEW

Condition-based maintenance (CBM) emerged decades ago as an alternative to the reactive threshold model, replacing calendar-fixed service intervals with decisions triggered by monitored equipment state (Jardine *et al.*, 2006). The maintenance research that followed formalized this shift into an increasingly quantitative discipline. Early statistical formulations of remaining useful life gave way to learned health indicators and multistage prognostic pipelines, in which data acquisition, health indicator construction, health stage division, and RUL prediction are treated as separate technical stages, each with its own established methods (Lei *et al.*, 2018; Si *et al.*, 2011). Foundational engineering texts on fault diagnosis and prognosis consolidated these methods into a coherent systems-engineering discipline applicable across mechanical and electronic assets more broadly (Vachtsevanos *et al.*, 2006).

[1] Director, Fleet Management Department, TransMobil Solutions LLC, Russia
[*] Corresponding author's e-mail: agromov.fleetdept.tms@gmail.com

Applied specifically to road vehicles, this research stream has produced a substantial body of surveys and empirical studies. Several broad reviews document how machine-learning-driven predictive maintenance has spread across manufacturing and transport sectors over the past decade (Hector & Panjanathan, 2024; Ran *et al.*, 2019; Zonta *et al.*, 2020). Vehicle-specific studies narrow this picture considerably. Compressor failures have been predicted from logged maintenance records and operational telemetry (Prytz *et al.*, 2015). Breakdown prediction has been demonstrated using stacked ensemble methods over heterogeneous fleet records (Khoshkangini *et al.*, 2023). Remote diagnostic and prognostic architectures have been proposed specifically for connected passenger vehicles (Shafi *et al.*, 2018). Two further systematic reviews consolidate the range of machine-learning techniques applied to vehicle health diagnosis more broadly (Jain *et al.*, 2022; Mahale *et al.*, 2025). Deployment challenges around sensor noise, data scarcity, and fielded reliability are catalogued in a survey that remains a frequent reference point across this literature (Theissler *et al.*, 2021).

A parallel line of research addresses the calibration of uncertainty attached to these predictions. A health index unaccompanied by a confidence measure gives an operator little basis for deciding when to act. The PHM'08 aircraft engine degradation benchmark introduced a standardized run-to-failure evaluation protocol that is still used across the RUL literature today (Saxena *et al.*, 2008). Work on Bayesian deep learning later formalized the distinction between aleatoric uncertainty, which arises from measurement noise, and epistemic uncertainty, which arises from model ignorance; this distinction was developed for computer vision tasks but was adopted quickly across prognostics (Kendall & Gal, 2017). Variational Bayesian architectures built specifically for RUL forecasting extended this calibration capability into interpretable models suitable for operational deployment (Kraus & Feuerriegel, 2019). Confidence-calibrated condition indexes for individual vehicles and components are, as a result, technically achievable today at a level of maturity that would not have been realistic fifteen years ago.

Kolesnykov (2026a) frames the rental fleet as a managed technical system that is structurally distinct from other commercial fleets, because its driver population changes with every transaction. A logistics fleet or municipal bus operator retains continuity of known-driver behavioral history that a rental fleet, by definition, cannot accumulate. Building on this framing, a composite scoring framework has been proposed that fuses an uncertainty-adjusted vehicle health index with four rental-specific context variables, customer proximity risk, fleet coverage pressure, reservation calendar density, and revenue tier, into a single score driving autonomous intervention workflows (Kolesnykov, 2026b). Within the domains surveyed for this paper, it is one of the few located instances in which condition-index generation and operational scheduling context have been combined into a single computable, fleet-validated priority function built specifically for the rental setting.

MATERIALS AND METHODS

The literature underlying this synthesis was identified through searches conducted in Scopus, Web of Science, and Google Scholar, combining terms drawn from three domains: predictive maintenance and prognostics ("remaining useful life," "health index," "condition-based maintenance"), vehicle and fleet applications ("vehicle health monitoring," "fleet maintenance scheduling," "compressor failure prediction"), and rental- or sharing-specific operations ("car-sharing," "vehicle rental," "shared mobility maintenance"). Searches were restricted to peer-reviewed journal articles, conference proceedings, and monographs published in English between 2006 and 2026, a window chosen to capture the period from the consolidation of condition-based maintenance as a distinct discipline (Jardine *et al.*, 2006) through the most recent rental-specific scoring work (Kolesnykov, 2026a, 2026b). Reference lists of retrieved surveys were examined for additional sources not captured by the keyword searches, a snowballing step that surfaced several vehicle-specific and aircraft-domain studies not indexed under the initial search terms.

Screening followed two criteria. A study needed to address at least one stage of the prognostics pipeline, health indicator construction, health stage estimation, remaining useful life prediction, or uncertainty calibration, applied to a mechanical or transport asset. Where a study addressed scheduling or intervention decisions beyond estimation alone, it needed to specify the decision mechanism used to convert a health signal into an action. Studies limited to sensor-level signal processing without a link to a maintenance or scheduling decision were excluded, as were studies confined to non-vehicular industrial equipment unless they contributed a methodological technique, such as uncertainty decomposition, later adopted within the vehicle-specific literature.

Each retained source was classified along the three dimensions reported in Table 1. Context integration was scored as none, partial, or full, based on whether the decision mechanism incorporated any rental- or fleet-specific operational variable beyond the health signal itself, such as reservation demand, asset location, or customer custody state. Decision mechanism was recorded as the specific method by which a health estimate was converted into an action: a fixed threshold rule, an optimization routine, a learned policy, or a multiplicative scoring function. Validation scale was recorded as the largest unit of analysis at which the reported method had been empirically evaluated, ranging from individual components to full fleets. Where a source addressed multiple stages, for instance a study reporting both an RUL model and a scheduling optimization, it was classified separately for each stage it addressed and cross-referenced across the relevant clusters in Table 1.

RESULTS AND DISCUSSION

Table 1 summarizes the classification applied across the reviewed literature. Four broad clusters emerge along the context-integration dimension, ranging from reactive threshold monitoring through fully context-integrated rental scoring.

Context integration and validation scale move independently of each other. Several component- and vehicle-level studies report large, well-documented datasets without ever touching scheduling logic, while several fleet-scale scheduling studies rest on simulated data instead of operational deployments. The direction of research effort within adjacent fields also fails to track where rental-specific need is greatest. Aircraft and industrial high-value-asset scheduling literature contains the most sophisticated decision mechanisms in the table, yet it operates under assumptions, fixed maintenance slots and controlled asset location. A third pattern concerns geography and climate. None of the eight approaches models weather exposure, regional road quality, or seasonal driving conditions as an input to either the health index or the scheduling decision, even though component wear rates for brakes, tires, and suspension parts are known to vary with these factors. This gap sits equally across component-level, vehicle-level, and fleet-level studies, which suggests a shared blind spot in the surrounding literature.

Table 1: Classification of maintenance-relevant approaches by context integration, decision mechanism, and validation scale

Approach category	Context integration	Decision mechanism	Validation scale
Reactive threshold monitoring (OBD-II)	None	Fixed rule	Component-level
CBM and statistical RUL foundations	None	Estimation only	Component-level
ML-based vehicle health/ RUL prediction	None	Estimation only	Vehicle-level
Uncertainty-calibrated prognostics	None	Estimation with calibrated confidence	Component/vehicle-level
Fleet maintenance scheduling (generic)	Partial (cost, workload)	Optimization	Fleet-level
Aircraft CBM under uncertainty	Partial (disruption, capacity)	Learned policy (POMDP/ DRL)	Fleet-level
Car-sharing operations optimization	None (maintenance largely absent)	Optimization (relocation, pricing)	Fleet-level
Rental-context composite scoring	Full (custody, proximity, calendar, revenue)	Multiplicative scoring, autonomous trigger	Fleet-level (847 vehicles)

A chronological reading of the sources in Table 1 adds a further dimension to this picture. The foundational condition-based maintenance and statistical remaining-useful-life literature dates from the mid-2000s (Jardine *et al.*, 2006; Vachtsevanos *et al.*, 2006), and the PHM'08 benchmark that standardized run-to-failure evaluation protocols followed shortly after (Saxena *et al.*, 2008). Vehicle-specific applications of these methods did not appear in volume until roughly a decade later, clustering between 2015 and 2023 (Khoshkangini *et al.*, 2023; Prytz *et al.*, 2015; Shafi *et al.*, 2018). Uncertainty calibration followed a similar lag relative to its own methodological origin: the aleatoric-epistemic distinction was formalized for computer vision in 2017 (Kendall & Gal, 2017) and reached RUL-specific architectures only two years later (Kraus & Feuerriegel, 2019). Each methodological advance in this literature took several years to migrate from its point of origin, general prognostics or an adjacent machine-learning subfield, into vehicle-specific and then fleet-specific application.

The rental-context composite scoring framework breaks this pattern in one respect and follows it in another. It breaks the pattern by integrating four rental-specific variables that no prior fleet-scheduling work had combined into a single score, a step that required no new estimation technique, only a new combination rule applied to methods already mature elsewhere in the literature. It follows the pattern in its timing: it appears only after uncertainty-calibrated health indexes had become a standard, citable capability (Kolesnykov, 2026b), consistent with the broader tendency for operational integration work to trail behind the estimation methods it depends on. If this lag structure holds, the aircraft-domain learned-policy scheduling work (Tseremoglou & Santos, 2024) and the simulation-based fleet scheduling optimization (Wang *et al.*, 2022) represent decision mechanisms not yet tested against rental-specific constraints, but the pattern observed elsewhere in this literature suggests they could migrate into that context within a comparable multi-year window, provided the interoperability barriers discussed later in this paper are addressed first.

The first three clusters in Table 1 show that condition-index generation for vehicles has reached a mature,

increasingly automated state. Health indicator construction, health stage segmentation, and RUL prediction now form a standardized pipeline (Lei *et al.*, 2018). Vehicle-specific applications extend this pipeline across component types as different as compressors, braking systems, and drivetrain sensors. Compressor telemetry (Prytz *et al.*, 2015) and multi-source breakdown ensembles (Khoshkangini *et al.*, 2023) both illustrate how far component-level prediction has advanced beyond simple threshold rules. Confidence calibration has kept pace. Aleatoric-epistemic decomposition (Kendall & Gal, 2017) and variational Bayesian architectures tailored to RUL forecasting (Kraus & Feuerriegel, 2019) both build on the standardized evaluation protocols established by the PHM'08 benchmark (Saxena *et al.*, 2008). Today's health indexes carry documented, testable confidence properties that earlier threshold systems never had.

Scheduling logic capable of consuming these indexes already exists, though it was built for operational assumptions foreign to rental use. The multi-objective evolutionary approach evaluated on a simulated New York taxi fleet optimizes workload, cost, and schedule stability from predicted component RUL, with no reference to customer possession state or reservation demand (Wang *et al.*, 2022). The partially observable Markov decision process formulation for aircraft maintenance incorporates fixed and flexible maintenance slots and models schedule disruption explicitly. Its constraints still reflect airline structure: fixed slots, a controlled fleet location, and an operator that always knows where each aircraft sits (Tseremoglou & Santos, 2024). Integrated mission-and-maintenance planning for high-value assets addresses a structurally similar problem, coupling forecast health with mission assignment, within a fleet-operator context where asset location and custody stay under operator control at all times (Camci *et al.*, 2019).

Car-sharing operations research occupies an unusual position in this landscape. It addresses precisely the environment in which rental-context maintenance scheduling would need to operate, yet its optimization attention has concentrated on vehicle relocation, station and pricing design, and demand-supply balancing. A comprehensive recent survey of the field states that maintenance-related decision problems remain comparatively unexplored relative to these other operational concerns (Bruglieri & Pisacane, 2023). Condition-index generation and rental-specific scheduling logic have developed, as a consequence, in largely separate literatures that rarely cite one another.

The composite scoring framework closes part of this gap. It combines a per-vehicle uncertainty-adjusted health index with customer proximity risk, fleet coverage pressure, reservation calendar density, and revenue tier through multiplication, producing a single priority value that triggers an autonomous intervention workflow directly, without routing through a manual review queue (Kolesnykov, 2026b). Evaluated across 847 rental vehicles over 24 months and 2.31 million vehicle-hours of operation, the framework achieved 90.3% empirical coverage of its calibrated confidence intervals. A nominal 90% interval that in practice contains the true value in 90.3% of cases sits close enough to the target rate to indicate calibration that is neither overconfident nor overly conservative. This property matters more for scheduling decisions than raw point-estimate accuracy: an operator acting on an overconfident interval will under-schedule maintenance, and one acting on an overly conservative interval will over-schedule it. The framework also improved seven-day advance fault detection by 32.4 percentage points relative to OBD-II threshold monitoring, extending the usable lead time for coordinating a replacement vehicle or a service slot before the customer's rental period ends. It reached a remaining useful life mean absolute error of 4.1 days across 312 continuously monitored failure events. The companion monograph situates this scoring layer within a broader system architecture spanning sensor modality selection, feature engineering, and autonomous service-facility selection. It argues that the rental fleet's transaction-level driver turnover makes this degree of context integration a design requirement from the outset, not something added as a later refinement (Kolesnykov, 2026a).

Telematics hardware capable of streaming the sensor data that condition-index models require has become standard equipment across much of the rental and car-sharing fleet, no longer an aftermarket addition reserved for premium operators, over roughly the same period that the prognostics methods described above matured. Car-sharing operations research has documented parallel growth in fleet size and network complexity, driven by demand-responsive relocation and pricing systems that themselves depend on continuous vehicle-state data (Bruglieri & Pisacane, 2023). The same data pipeline that feeds relocation and pricing decisions is, in principle, available to feed a condition-index model, which is part of why the absence of maintenance-focused work in that same operations literature stands out as a genuine gap. The infrastructure prerequisite for context-integrated maintenance prioritization, continuous telematics coverage across a fleet, is therefore substantially in place at many operators already, independent of whether those operators have built the scheduling logic to act on it. This separates the present gap from an earlier generation of predictive maintenance barriers, where sensor cost and coverage were themselves the limiting factor. The limiting factor identified in this synthesis sits downstream of data availability, in the scheduling and governance layer, a materially different problem from the one the original condition-based maintenance literature set out to solve.

Discussion

The synthesis presented here points to an asymmetric maturity gradient. Health index estimation and its uncertainty calibration have advanced to where confidence-adjusted RUL predictions with documented coverage properties are achievable using established

methods (Kendall & Gal, 2017; Kraus & Feuerriegel, 2019; Saxena *et al.*, 2008). The logic needed to convert such an estimate into an operational scheduling decision under rental-specific constraints remains confined, within the domains searched, to a single validated implementation. This asymmetry points to where the real bottleneck sits. Autonomous maintenance intervention in rental fleets sits primarily outside the modeling problem that prognostics research addresses.

Several limitations in the current evidence base deserve direct treatment. The evaluation reported for the composite scoring framework covers a single fleet operator across 847 vehicles and 24 months. Generalization across operators with different vehicle mixes, climates, and customer behavior profiles has not yet been demonstrated, and the framework's reported accuracy figures should be read with that scope in mind. The multiplicative combination of health index and context variables assumes a degree of commensurability across those variables that has not been tested against alternative combination rules. The learned weighting implicit in the partially observable Markov decision process formulation used for aircraft fleets offers one such alternative (Tseremoglou & Santos, 2024). The multi-objective evolutionary weighting applied to vehicle fleet scheduling more broadly offers another (Wang *et al.*, 2022).

Autonomous intervention, as distinct from decision support, requires data governance and system interoperability that most rental operators currently lack. The monograph companion to the scoring framework treats data governance, security, and privacy as a distinct design layer built into the system architecture from the start, ahead of scheduling logic and scoring functions (Kolesnykov, 2026a). This framing matches recurring observations in the broader predictive maintenance adoption literature. Integration across telematics, reservation, and service-scheduling systems most often determines whether a predictive maintenance system reaches production deployment; model accuracy contributes far less to that outcome (Mahale *et al.*, 2025; Zonta *et al.*, 2020). Reservation systems, telematics platforms, and service-scheduling software are typically procured from separate vendors with no shared data interface between them. An autonomous intervention trigger that reassigns a customer or redirects a vehicle mid-rental carries commercial and legal consequences that equipment maintenance in a closed industrial setting does not have to face. A related question concerns disclosure. A customer whose rental is interrupted by an autonomous scheduling decision is affected by a judgment about a vehicle's condition that they never agreed to in advance.

A practical implication follows from this governance constraint. Operators unable to unify telematics, reservation, and service-scheduling data right away could still adopt a staged approach. The health-index and uncertainty-calibration layer can be deployed first, since these components rely on established, portable methods (Kendall & Gal, 2017; Lei *et al.*, 2018). Its output can route to a human dispatcher during an interim period, with the autonomous trigger added once the supporting data interfaces exist. This staging separates a genuinely solved technical problem, condition estimation, from a genuinely unsolved organizational one, system integration, and lets an operator capture partial benefit before the full architecture is in place.

Adoption of any staged intervention system of this kind also depends on evidence the current literature does not yet provide. None of the sources reviewed here report a cost-benefit analysis that sets the infrastructure and integration cost of a context-integrated scheduling system against the avoided cost of unplanned vehicle-off-road events, customer relocations, and liability exposure. The composite scoring framework's reported gains, a 32.4 percentage point improvement in seven-day advance fault detection and a remaining useful life mean absolute error of 4.1 days (Kolesnykov, 2026b), describe prediction quality; translating that figure into a financial return requires assumptions about fleet size, vehicle value, and local labor and towing costs that vary too much across operators for a single reported figure to generalize. An operator evaluating adoption would need this translation performed against its own cost structure, a step outside the scope of any prognostics evaluation and one that the literature has not yet addressed even in outline.

A second gap concerns the people who work alongside an autonomous or semi-autonomous scheduling system, distinct from the system's technical design. Fleet dispatchers and service coordinators are accustomed to routing decisions based on customer requests, reservation conflicts, and their own judgment of vehicle condition from visual inspection or driver complaints. A scoring system that reassigns or recalls a vehicle on the basis of a calibrated but opaque composite score introduces a judgment the dispatcher did not make. Whether dispatchers trust, override, or work around such a system, and under what conditions that trust is established or lost, has not been studied for this application. Human-factors research on automation trust in other operational settings suggests that trust calibration, often determines whether an automated recommendation is actually followed, which points to a gap between the composite scoring framework's demonstrated statistical performance and its readiness for unattended field use.

Regulatory and insurance treatment of an autonomous vehicle-recall decision represents a third open area. When a rental operator interrupts a customer's rental period on the basis of a predictive judgment, ahead of any confirmed fault, it is making a decision that may carry contractual and liability consequences under rental agreements, consumer protection regulation, and insurance policy terms that were not written with predictive intervention in mind. The gap reflects a broader separation between predictive maintenance research and rental-sector regulatory frameworks, which have developed with

little cross-reference, mirroring the separation between condition-index generation and car-sharing operations research described earlier in this discussion.

The decision mechanisms developed in adjacent, better-resourced domains illustrate both the opportunity and the limits of direct transfer. The partially observable Markov decision process formulation used for aircraft fleets learns a scheduling policy under disruption and capacity constraints that are fixed and known to the operator in advance (Tseremoglou & Santos, 2024). A rental fleet's constraints, which customer holds which vehicle, when it will be returned, and where, are neither fixed nor fully known to the operator until the moment of return, which limits how directly a POMDP-style learned policy could be applied without substantial reformulation. The multi-objective evolutionary approach to fleet-wide schedule optimization (Wang *et al.*, 2022) offers a more transferable structure in principle, since it optimizes across multiple objectives without assuming fixed slots, but its validation relied on a simulated taxi fleet instead of deployed operational data, and its objectives, workload and cost stability, would need to be reweighted to include the customer-facing variables that the rental-specific composite framework already incorporates.

These four gaps, cost-benefit evidence, operator trust, regulatory clarity, and mechanism transfer, share a common feature. None of them can be closed by further work on health index estimation or uncertainty calibration, the two areas where this literature is already mature. They belong instead to the organizational and institutional side of deployment, which the prognostics research reviewed here was never designed to address and which a different research community, closer to operations management and regulatory studies, would need to take up.

CONCLUSION

Set against the car-sharing operations literature's own account of its priorities, in which maintenance-related decision problems remain comparatively unexplored relative to relocation and pricing (Bruglieri & Pisacane, 2023), the composite scoring framework addresses a genuine research question that remains open across the domains searched, with no prior resolution located elsewhere in adjacent literature. Future work extending this line of research would benefit from three additions. Multi-operator validation, designed explicitly to test generalizability across vehicle mixes and climates, is the most direct next step. A direct comparison between fixed multiplicative scoring and learned scheduling policies, run under matched evaluation conditions, would establish whether the simplicity of a multiplicative score costs anything in accuracy. Treatment of data governance and cross-system interoperability as first-class design constraints, alongside predictive accuracy, would bring the engineering requirements of deployment into the same evaluation frame as the modeling requirements. Until such comparisons exist, context-integrated maintenance prioritization for rental fleets should be regarded as demonstrated feasible at single-operator scale, with autonomous deployment beyond that scale dependent on infrastructure and governance conditions the prognostics literature itself does not address.

REFERENCES

Bruglieri, M., & Pisacane, O. (2023). A survey on emergent trends in the optimization of car-sharing systems. *International Transactions in Operational Research, 30*(6), 2867–2908. https://doi.org/10.1111/itor.13315

Camci, F., Medjaher, K., Atamuradov, V., & Berdinyazov, A. (2019). Integrated maintenance and mission planning using remaining useful life information. *Engineering Optimization, 51*(10), 1794–1809. https://doi.org/10.1080/0305215X.2018.1552951

Hector, I., & Panjanathan, R. (2024). Predictive maintenance in Industry 4.0: A survey of planning models and machine learning techniques. *PeerJ Computer Science, 10,* Article e2016. https://doi.org/10.7717/peerj-cs.2016

Jain, M., Bhatt, D., & Bhatt, N. (2022). Systematic literature review on predictive maintenance of vehicles and diagnosis of vehicle health using machine learning techniques. *Computational Intelligence, 38*(6), 1990–2008. https://doi.org/10.1111/coin.12553

Jardine, A. K. S., Lin, D., & Banjevic, D. (2006). A review on machinery diagnostics and prognostics implementing condition-based maintenance. *Mechanical Systems and Signal Processing, 20*(7), 1483–1510. https://doi.org/10.1016/j.ymssp.2005.09.012

Kendall, A., & Gal, Y. (2017). What uncertainties do we need in Bayesian deep learning for computer vision? In I. Guyon, U. von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, & R. Garnett (Eds.), *Advances in neural information processing systems* (Vol. 30, pp. 5574–5584). Curran Associates.

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

Kolesnykov, V. (2026a). *Smart fleet management: Data-driven control of rental vehicles.* GlobeEdit.

Kolesnykov, V. (2026b). Uncertainty-conditioned composite scoring for autonomous maintenance scheduling in rental vehicle fleets. Scientia. *Technology, Science and Society, 3*(6), 24–30. https://doi.org/10.59324/stss.2026.3(6).03

Kraus, M., & Feuerriegel, S. (2019). Forecasting remaining useful life: Interpretable deep learning approach via variational Bayesian inferences. *Decision Support Systems, 125,* Article 113100. https://doi.org/10.1016/j.dss.2019.113100

Lei, Y., Li, N., Guo, L., Li, N., Yan, T., & Lin, J. (2018). Machinery health prognostics: A systematic review from data acquisition to RUL prediction. *Mechanical Systems and Signal Processing, 104,* 799–834. https://doi.org/10.1016/j.ymssp.2017.11.016

Mahale, Y., Kolhar, S., & More, A. S. (2025). A comprehensive review on artificial intelligence driven predictive maintenance in vehicles: Technologies, challenges and future research directions. *Discover Applied Sciences, 7*(4), 1–25. https://doi.org/10.1007/s42452-025-06681-3

Prytz, R., Nowaczyk, S., Rögnvaldsson, T., & Byttner, S. (2015). Predicting the need for vehicle compressor repairs using maintenance records and logged vehicle data. *Engineering Applications of Artificial Intelligence, 41,* 139–150. https://doi.org/10.1016/j.engappai.2015.02.009

Ran, Y., Zhou, X., Lin, P., Wen, Y., & Deng, R. (2019). *A survey of predictive maintenance: Systems, purposes and approaches* (arXiv:1912.07383). arXiv. https://arxiv.org/abs/1912.07383

Saxena, A., Goebel, K., Simon, D., & Eklund, N. (2008). Damage propagation modeling for aircraft engine run-to-failure simulation. In *2008 International Conference on Prognostics and Health Management* (pp. 1–9). IEEE. https://doi.org/10.1109/PHM.2008.4711414

Shafi, U., Safi, A., Shahid, A. R., Ziauddin, S., & Saleem, M. Q. (2018). Vehicle remote health monitoring and prognostic maintenance system. *Journal of Advanced Transportation,* 2018, Article 8061514. https://doi.org/10.1155/2018/8061514

Si, X.-S., Wang, W., Hu, C.-H., & Zhou, D.-H. (2011). Remaining useful life estimation: A review on the statistical data driven approaches. *European Journal of Operational Research, 213*(1), 1–14. https://doi.org/10.1016/j.ejor.2010.11.018

Theissler, A., Pérez-Velázquez, J., Kettelgerdes, M., & Elger, G. (2021). Predictive maintenance enabled by machine learning: Use cases and challenges in the automotive industry. *Reliability Engineering & System Safety, 215,* Article 107864. https://doi.org/10.1016/j.ress.2021.107864

Tseremoglou, I., & Santos, B. F. (2024). Condition-based maintenance scheduling of an aircraft fleet under partial observability: A deep reinforcement learning approach. *Reliability Engineering & System Safety, 241,* Article 109582. https://doi.org/10.1016/j.ress.2023.109582

Vachtsevanos, G., Lewis, F. L., Roemer, M., Hess, A., & Wu, B. (2006). *Intelligent fault diagnosis and prognosis for engineering systems.* John Wiley & Sons. https://doi.org/10.1002/9780470117842

Wang, Y., Limmer, S., Nguyen, D. V., Olhofer, M., Bäck, T., & Emmerich, M. (2022). Optimizing the maintenance schedule for a vehicle fleet: A simulation-based case study. *Engineering Optimization, 54*(7), 1258–1271. https://doi.org/10.1080/0305215X.2021.1919888

Zonta, T., Da Costa, C. A., da Rosa Righi, R., de Lima, M. J., da Trindade, E. S., & Li, G. P. (2020). Predictive maintenance in the Industry 4.0: A systematic literature review. *Computers & Industrial Engineering, 150,* Article 106889. https://doi.org/10.1016/j.cie.2020.106889