



Optimization of Vehicle Maintenance and Repair Processes Based on Operational Data Analytics

Roman Mokrynskyi  ¹ *

¹ Kryviy Rih Technical University, Kryviy Rih (Ukraine). Founder and Chief Engineer at Portland AutoTech Service, Specialist in Cars and Automotive.

* Corresponding Author, e-mail: mokrynskyi.r@gmail.com

ARTICLE INFO

Research Article

Received:

7 September 2025

Revised:

17 October 2025

Accepted:

27 November 2025

Published online:

25 December 2025

Copyright © 2025 by authors



This is an open access journal and all published articles are licensed under a Creative

Commons Attribution—NonCommercial 4.0

International

(CC BY-NC 4.0)

DOI: [10.5281/](https://doi.org/10.5281/zenodo.19154698)

[zenodo.19154698](https://doi.org/10.5281/zenodo.19154698)

ABSTRACT

The article looks at optimization of maintenance and repair services on the base of analysis of operational data as an implementation of real movement away from mileage/calendar maintenance towards the risk based condition based system, when servicing is prescribed not on the schedule, but on basis of actual operational behavior in actual service environment, proving, that data streams (telematic streams, OBD-II/CAN parameters, DTC codes, service history, repeat service records) build an “digital trace” of operation, which value is revealed when this trace is transformed into repetitive service operations (scheduling, further diagnosis, prevention maintenance or acceptable deferral with risk control) within a service loop (the data started to “work”). From a methodological viewpoint, the paper confirms an end-to-end proposal that mixes disparate data integration, time-series-based degradation-features engineering, anomaly detection, estimating the probability of a repair event over a certain horizon and an operations-and-economics model that determines a threshold for intervention, taking into account both downtime costs and costs of missed failure, capacity and logistics of service resources. In our opinion, it’s precisely the ‘price of error’ that gives the forecasting a practical meaning, defining when it’s useful to carry out early works and when the ‘cost-effective’ choice is to wait with an additional investigation during a short validation period or collect more features, thereby avoiding unnecessary jobs and not compromising confidence in suggestions. Scientific actuality would be the formalization of the movement from fragmented operational data toward a governed service process. Meanwhile we adopt standard vocabulary such as “symptom – diagnosis – work – result” and a feedback loop after repairing, because we need to update the system and account for changes in the data, like in the case of different software versions. Practical usefulness looks like being on determining rules on fleet segments, trigger levels, etc.

KEYWORDS

maintenance, repair, predictive maintenance, operational data, telematics, OBD-II, machine learning, time series, optimization, fleet management.

Introduction

It is obvious that it is these data, and not the set intervals of scheduled maintenance, which offer the useful insights which are currently so absent in technical maintenance and that only they allow for assessment of the actual state of a component or unit and its tendency to degrade over time, detect precursor symptoms of potential failures and uncover non-obvious patterns of why two visually identical, and technically equally aged, vehicles can lead significantly different “lives”, and accordingly, require differing maintenance strategies. In this context, data becomes the raw material for the rethinking of the fundamental logic of maintenance and repair, because the optimization lies not in reducing costs for repair itself, but in managing the risks of failure, time of non-operational period and of resources, but not on an intuitive level, and in the scope of reproducible, verifiable, and organized procedures.

It is essential to acknowledge the disciplinary effect of the classical regulatory framework based on calendar and mileage, but, as it seems to us, it is inherently “deaf” to the variability of real operation, in which load histories, thermal cycling, route structure, proportion of short-distance trips, even the form of acceleration or deceleration (which seem trivial, until you look at the wear data) will define disparate risk categories. This dictates, where the process of service has to be both economically justifiable and technically efficient, that a transition is required from a generic interval to an individualized one. In one case this solution might be very simple - “service ahead because operating conditions are rough” - in others it could take on a whole new degree of complexity where one has to consider available resource, anticipate anomalies, evaluate the root causes, compare the fault codes with the repair history, and, finally, determine whether a current signal represents noise or measurement error (a surprisingly common phenomenon). And so, optimizing M&R no longer presents merely a technical problem and begins to address the domain of operations management, where each service has a cost, each non-service carries a risk, and somewhere in between, decision space can be mined for information.

One must not naively suppose that all you have to do is to “take the data and train the model,” and the system will practically operate itself. In reality, the main problems are not always in the obvious first place. Data can be missing, conflicting, in a different format and have an additional quality of evolving with time because firmware, operating mode, fleet composition, or even drivers’ behavior change after putting the control system in service. Clearly, optimization of maintenance and repair by operation data implies not only a mathematical core, but also methods for maintaining data quality, post-repair feedback, results verification and periodical re-training of models, or analysis might be a ‘showcase’ and does not help in management. Thus, the emphasis of the article is on the practical integrity of a process from data to decision, from decision to a stable process that can withstand real-time service load.

Literature Review

Nowadays, among scholars’ debate on predictive maintenance, an awareness on machine learning for service are emerging which is not considered a tool of “automation”, but rather a means of transformation of the technical uncertainty into the “language” of controlled probabilities - suitable for making decisions, that has economic and safety implications. Many studies focused on predictive maintenance within automotive sector emphasizes that obstacles to implementation are more complicated than selection of algorithm and they included data quality, intelligibility of predictions and organizational capability to assimilate an analytical outcome into the service process (Theissler et al. 2021). At the beginning of looking, this fact looks clear but paradoxically here implementation very often “fails” although the model proves its ability on test sample by high metrics, then in the service process, conclusions of the model are not considered as a cause for action since “what exactly it sees” is not comprehensible and as the reason to “recommend” replacement or postponement.

Also the direction that the forecast of the state or of the remaining resource is viewed not as an outcome, but as an input to the planning problem. It is clear that this is an essential change in the

perspective, because a forecast that does not turn into a decision is rather an intellectual pursuit. This class of work unites time series models with optimization formulations, such as Markov decision models, in which the present state of the system and forecast of its evolution are input to determining the time and magnitude of service (Zheng et al., 2023). At the same time, it is asserted that forecasts will always have error and maintenance planning needs to “build in” a penalty for false alarms and a penalty for missed failures. This concept of thinking is embodied well in articles that address how to apportion risk under uncertainty in forecasting, as well as alarms that set service on the condition of achieving a probability threshold rather than that a failure actually has occurred (de Pater et al., 2022). We feel this is an area where much needs to be considered when dealing with the automobile industry, because of how readily you can slide into two extremes: either over-react and perform unnecessary maintenance or delay service and risk a greater repair cost, both damaging confidence in the analytic solution.

Second, we should point to a set of publications where data that has traditionally been associated with on-board diagnostics, i.e., OBD-II and associated telematic channels, is analyzed. From survey papers, we can see that not only are tasks for ‘error code recognition’ achievable with OBD data but also for more general tasks concerning the economy of driving, the safety and the stability of driving, and the application of ML to reveal hidden correlations that are precursors to a defect occurring (Michailidis et al., 2025). Along the lines of this is an applied study where an engine health monitoring task uses hybrid deep learning structures which incorporate convolutional and recurrent parts in an effort to deal with noisy time-series data with the need to capture both local and long-term dependencies (Rahim et al., 2024). This is technically demanding but in human terms, its purpose is clear - the algorithm is trying to “hear” precursors to a problem that cannot be detected by humans yet, where regulation does not yet demand anything.

A different research field is concerned with service demand forecasting; in this field we are not interested only in the node reliability probability but also on the volume of future work, the capacity load for service, the spare parts requirement. In these cases, the models are utilized with consideration on mileages, seasonality, usage patterns and other factors which determine request dynamics (Chen et al., 2024). The most complex models are based on multi-view approach or multi-source time models in which multiple information channels are exploited, increasing the forecast stability in conditions of non-complete information (Chen et al., 2025). At the same time it seems that these studies implicitly highlight a thing: if there is no rigor in data acquisition and a correct correspondence between work and defect dictionaries on the station and the manufacturer side, any model would soon “break” in the reality because it can not interpret events in a predictable way.

Over the last few years, interest has also arisen in the context of digital twins and Internet of Vehicles where predictive maintenance is treated as a part of the “overall digital ecosystem,” and the question of data and event trust is solved by mechanisms that make events auditable (such as block-chain solutions (Iqbal et al., 2025)). Besides this, a new kind of combined solutions appears where forecast, risk assessment and decision support for services are to be integrated into a single structure (Errezgouny et al., 2025). As seen, those publications expand the outlook, but at the same time pose the challenging practical problem of who within the company makes the decisions based on the model, how the critical limits are determined, how the repair result is documented and how the analyst receives the “feed-back” from outside, without which he is, as they put it, “blind.”

At the end, a trend that is progressively emphasized consists in the economic part, where the viability of an investment on condition monitoring, sensitization, or on analytics applied to fleets and service networks is evaluated, not at the level of a priori considerations but based on the models of cost, downtime, and risk (Crespo del Castillo & Parlikad, 2024). This perspective naturally requires a statistical “backing”, and at this stage, information on the trends in transport, on the structure of fleets, on the indices of motorization, on recall data (Eurostat, 2024; European Automobile Manufacturers’ Association, 2025; National Highway Traffic Safety Administration, 2026; World Bank, n.d.) is useful to permit calibration of the premises or to correlate an observation “local” with “global” trends. If we remove this comparison, there appears to be a risk of obtaining a locally optimal solution that does not hold to changes in conditions.

Problem Statement

The task of this article is to theoretically and practically substantiate a plan for a vehicle maintenance and repair optimization program in which operating data cease to be a passive store and a fully-fledged tool for making decisions on the timing, volume, and priority of service interventions and on ensuring a stock of spare parts and a service capacity use plan. We believe it is necessary not only to demonstrate the ability to train a prediction model using data (which has long since not cause surprise) but also to justify how analytics can organize the repair process in a technically sound, economically viable, and organizationally feasible manner (i.e., that the recommendations of the model will not float in the air and a handsome report, the nobody does not consider).

Methods and Materials

The suggested research is both built on time-series analysis and machine learning models for prediction of service status and demand, on the one hand, and economic and operational modeling of service solutions (within realistic bounds of the service stations), on the other. i.e., let's take that a technical status of a car is characterized by the dynamics of parameters and events, and a service solution always has a price, delay, logistic implications, and the possibility of mistakes. It follows directly from this that optimization isn't "finding a best algorithm", but coordinating various layers - data layer, forecasting layer, action layer.

First, formats of approaches are coordinated: synchronization in time, mileage, etc. Next, a sign transformation is done, and it is done not only on the level of "average values," but also on the level of dynamics' characteristics, variability, trend, transitions between modes, as the node failure is most of times not a single value but a change of system behavior structure during time (and, this is what seems to be slightly underestimated by the practitioner). Models of predicting technical state and probability of failures are applied according to formed features, and models of forecasting demand for service are built, with the purposes of planning of resources. Omalý detection methods for prognoses of pre-failure state and the models reproducing probabilistic course of degradation are used for forecasting resources, and it is necessary to take the forecast uncertainty into account because it's uncertainty that sets up a price of mistake.

To transition from forecast to service solution, an economic-operational model is proposed which contrasts the cost of a programmed intervention versus the cost and risk associated with an unplanned failure, the downtime cost, logistical supply constraints and the throughput service. We feel it is relevant that the optimality condition cannot be one dimensional, because to minimize costs in the tightest sense might entail increasing risk or downtime.

Results and Discussion

Optimization of maintenance and repair based on operation information does not start with a model; it starts with an inventory of what is actually meant by the word "information" for a given organization; because in real life, this word means completely different things: there is a telematics and a streams of parameters, there are diagnostic codes, service logs, warranty claims, data on replaced components, history of repeat visits, finally there are some information about recalls and systemic defects. And it is here, according to our opinion, that the first condition of success lies: if these sources exist in different "worlds" and they can't be summarized into a single "coordinated description of event", then analysis will be based on separate facts, and separate facts leads, as we all know, to separated solutions.

It is here where the appropriate level of risk and background costs need to be set, at which the prediction will be interpreted and threshold judgments made. That is the composition of the vehicle fleet by age and power source already defines initial expectations of the probability of failure and laboriousness of repair (in the case of the EU for example, the average age of passenger vehicles is 12.7 years, and the proportion of battery-powered vehicles in total fleet size is insignificant, which naturally influences the service load profile) (European Automobile Manufacturers' Association,

2026). Roadside assistance data, on the other hand, will give “hard” performance indications for reliability in real-time, broken down by age and propulsion system; there we see not general assumptions, but comparable metrics such as the number of breakdowns per 1000 vehicles (ADAC, 2025). Lastly, the arrays of mandatory technical inspection results help to determine the proportion of initial failures, as well as the structure of faults (by category - brakes, chassis, etc.) thereby translating “symptoms” into statistically-based diagnosis priorities (Driver and Vehicle Standards Agency, 2025; Driver and Vehicle Standards Agency, n.d.). These exterior series, combined with information about technical maintenance costs as a function of mileage, will make it possible to formalize the economic cost of forecast errors, rather than relying on general notions to make an optimization - the intervention threshold must not be subjective, but comparable to background risk and operational costs (American Automobile Association, 2025) (see Table 1).

Table 1. External statistical indicators for calibrating M&R thresholds and validating operational analytics

Indicator (what exactly it “calibrates”)	Example value (latest available snapshot)	Source	How to use in the model/process
Average age of vehicle fleet (baseline wear and tear)	EU: 12.7 years (passenger cars, 2024)	ACEA Vehicles on European Roads (Jan 2026)	Setting a priori risks of failure and labor intensity by age group; model segmentation
Fleet size (service load scale)	256 million passenger cars on EU roads (2024)	ACEA Vehicles on European Roads (Jan 2026)	Standardization of service demand indicators, recalculation of events per 1,000 cars/year
Share of BEVs in the fleet (change in the structure of typical work)	2.3% BEV in the total EU passenger car fleet	ACEA Vehicles on European Roads (Jan 2026)	Adjustment of defect profile/inventory composition; separate thresholds for BEV/ICE
Breakdowns per 1,000 cars (reliability in real traffic)	For 2-4-year-old cars: 9.4 (ICE) vs. 3.8 (BEV) breakdowns/1000	ADAC Pannestatistik 2025 (data for 2024)	Check whether your “repair event” forecast is consistent with independent statistics; false positive/false negative calibration
Initial technical inspection failures (prevalence of dangerous defects)	MOT Class 3/4: Total overall initial failure rate 28.58%; dangerous items 7.87% (2023-2024)	DVSA MOT open data (CSV)	Enhancement of the defect dictionary and diagnostic prioritization rules; selection of “control” systems (brakes, suspension, tires)
Cost benchmark for maintenance and repair per mile/km (threshold economics)	Maintenance/Repair/Tires: 11.04 cents/mile (5 years, 75,000 miles)	AAA Your Driving Costs 2025 (fact sheet)	Transition from probability to decision through expected value; setting thresholds where “it is better to do it now”

Source: Compiled based on analysis of sources (European Automobile Manufacturers’ Association, 2026; ADAC, 2025; Driver and Vehicle Standards Agency, 2025; Eurostat, n.d.).

A related aspect is that of data quality discipline: synchronizing time and mile, searching for data losses, removing anomalies, converting unit, and making sure that the same event isn’t represented in 3 ways according to the mechanic and the service location (since then the model is trained to the local routines, not the engineering patterns). The specifications for diagnostics modes and the semantics of some parameters, like those between OBD and UDS, then aren’t exactly “reference book” but “minimum grammar” which enables comparing measurements from different machines and different time span (SAE International, 2021). This is what we think is the boring, engineering effort that won’t fit in a slide-deck but the conclusions drawn on such data will have no value.

After having been harmonized the data, arises the question of what, precisely we are going to predict, because not only one “forecast” exists in the service, but rather many of them and for each one its

usefulness. One level is a forecast of the state of technical health, of the probability of failure of a particular node within an horizon of prediction. A further level, the evaluation of the remaining resource, when we are only interested not in the very occurrence of the risk but in its evolution. And finally, the third level which is sometimes down-valued: the prediction of demand for maintenance services, as such, is to say, the amount of work that the posts are going to generate in the future and what should be stock of spare parts (and now analytics are in relation not only to the engineer but to the dispatcher and logistician). The works in science show clearly this multi-level nature, in so far as the concept of predictive maintenance does not mean anything else than a model coupled with a procedure, not “push a button” (Theissler et al. 2021), and decision must also take into account the fact that the forecast will always be uncertain and that the errors are asymmetric in cost (de Pater et al., 2022). We are going therefore, not to simply train the model, but to implement a loop where the model gives the warning, the warning becomes the decision, the decision results in a consequence and the consequence back to the data by way of a feedback, in other words the system must be able to “learn” from the repairs itself.

The problem of feature formation is also important, because often there is not very much to tell from the naked value of the parameter. Degradation of nodes usually presents as the change in dynamic, the change in variability, the change in rate of transition between modes, accumulation of small discrepancies that are not conclusive on their own, but when taken together they produce a powerful signal. This is where hybrid deep learning approaches are particularly promising when dealing with sequences (time series) instead of a ‘table of averages,’ and there are applications to real world problems like condition monitoring for engines and similar systems (Rahim et al., 2024). On the other hand we should not substitute reality with a model, because however complex the model can be, we cannot deny the obvious truth; that some signals are noisy, some are dependent on the environment and some are not signs of defect, but just mode of operation. For this reason we must carry out an analysis checking the context; without it obviously every cold spell, or every instance of being stuck in a traffic jam will be an ‘anomaly.’

After gathering the forecast level, it comes down to a decision, and this is where, in our view, the fundamental leap occurs from “data science” to service management. If the model says that the probability of failure is growing, then the organization has to respond not with a generic sentence, but with a specific scenario. Will we make a reservation in a service center in the closest slot? Will we replace the part preventively? Will we carry out a more detailed diagnostics to reduce uncertainty? Or will we rather postpone the intervention, because the forecast is weak, and the cost of unwanted work too high? In the literature, this leap is defined well in models combining the forecast of a condition and decision support models, such as using Markov models, which help formally describe the trade-off between the risk and the costs of intervention (Zheng et al., 2023). In the practice of service, this may look simple enough (even mundane), but its simplicity doesn’t stem from any simplification, but from the fact that all complexity is embedded in the preparation phase: system defining the trigger levels, priority, acceptable delays, composing a task list, and technician receives not a ‘probability judgment’ but a logic that he can explain to the client or fleet manager.

An interesting but highly applicable part is the prediction of the need for servicing as part of resource planning. In this domain, research recommends taking into account mileage and context of use, as the demand for servicing does not follow a consistent and random distribution (Chen et al., 2024), and multichannel time series models can enable combining different information sources and smoothing out the forecast when data is patchy (Chen et al., 2025). It would appear that this block of processes is the one that has the highest potential for relatively rapid conversion into economic benefit: the improved planning of postal traffic and warehouse loads has a direct impact on reducing unproductive down times and costs that are not necessarily accounted for under ‘repair cost’ and thus appear directly as a company’s overall anxiety level, queue length, or missed deadlines.

Scientific articles can provide methods for analyzing whether investing in the maintenance system monitoring the fleets is justifiable; the problem becomes, in essence, the question of altering the cost structure and risks (Crespo del Castillo & Parlikad, 2024). We consider it useful to be “beyond” the model accuracy, as, on its own, accuracy will not yield benefits if there are any constraints (e.g., STO constraints) to implement recommendations, spare parts won’t reach the required location in time, or workers do not trust the system. This is the reason why we find it right to use external statistics,

whose values will permit better calibration of assumptions and present results within a broader context. This is what has been done using transport statistics publications (Eurostat, 2024), car industry market reports (European Automobile Manufacturers' Association, 2025), registered recalls (National Highway Traffic Safety Administration, 2026), and aggregate motorization levels (World Bank, n.d.). They are not a substitute for a technical study, but they definitely help define the order of magnitude, isolate the more relevant fields, and do not keep us stuck in a single data set context.

Lastly, let's briefly touch on the subject of organizational integration and trust. The problem of integrating the organization and maintaining the necessary trust is sometimes stated by digital twins/loV people in terms of event tracing and data trust, often with some elements of blockchain (Iqbal et al., 2025). In human terms, it's the same problem. Each analytic statement must point back to some transparent chain of evidence, each service decision to some transparent chain of responsibility. If not, the system might be mathematically elegant, but organizationally hollow, and then it does not influence service behavior, and it is behavior that we want to optimize.

Conclusion

The principal advantage of our proposed approach, described above, is the combination of a reduction in unplanned failures with a decrease in work that is done "unnecessarily" - merely because the technology demands it, rather than because the task requires it.

Simultaneously, analytics do not work in isolation and, arguably, that is the most grounded conclusion from this whole theme. A predictive model, no matter how well it works under lab conditions, does not by itself bring an effect if data has not been standardized to a uniform semantics, if events of repair are not logged to be returned to training, if the enterprise has not been structured by thresholds and response scenarios, if the spare part logistic is a thing separated from work orders, if the mechanics do not get explicable reason for a recommendation. Hence optimization comes from an integrated schema which alongside with time series models and anomaly detection techniques include data merging logic, post-repair feedback loop, data drift controls and economic/operational logic to transform likelihood into decision.

References

- ADAC. (2025). *ADAC Pannenstatistik 2025: Sind Elektroautos zuverlässiger?* <https://www.adac.de/rund-ums-fahrzeug/unfall-schaden-panne/adac-pannenstatistik>
- American Automobile Association. (2025). *Your driving costs 2025: Fact sheet.* <https://newsroom.aaa.com/wp-content/uploads/2025/09/UPDATE-AAA-Fact-Sheet-Your-Driving-Cost-9.2025-1.pdf>
- Chen, F., Shang, D., Zhou, G., Ye, K., & Ren, F. (2024). Mileage-aware for vehicle maintenance demand prediction. *Applied Sciences*, 14(16), Article 7341. <https://doi.org/10.3390/app14167341>
- Chen, F., Shang, D., Zhou, G., Ye, K., Ren, F., & Wu, G. (2025). Collaborative multiview time series modeling for vehicle maintenance demand prediction. *Scientific Reports*, 15(1), Article 13058. <https://doi.org/10.1038/s41598-025-96720-1>
- Crespo del Castillo, A., & Parlikad, A. K. (2024). Evaluating investment in condition monitoring for fleet maintenance. *IFAC-PapersOnLine*, 58(8), 377-382. <https://doi.org/10.1016/j.ifacol.2024.08.150>
- de Pater, I., Reijns, A., & Mitici, M. (2022). Alarm-based predictive maintenance scheduling for aircraft engines with imperfect Remaining Useful Life prognostics. *Reliability Engineering & System Safety*, (221), Article 108341. <https://doi.org/10.1016/j.ress.2022.108341>
- Driver and Vehicle Standards Agency. (2025, October 30). *MOT testing data for Great Britain* [Statistical data set]. <https://www.gov.uk/government/statistical-data-sets/mot-testing-data-for-great-britain>
- Driver and Vehicle Standards Agency. (n.d.). *MOT class 3 and 4 vehicles: Initial failures by defect category (data extract).* <https://www.gov.uk/csv-preview/68e682f5750fcf90fa6fff65/dvsa-mot-03-mot-class-3-and-4-vehicles-initial-failures-by-defect-category.csv>
- Errezgouny, A., Chater, Y., Barranco González, C. D., & Cherkaoui, A. (2025). An integrated deep learning approach for predictive vehicle maintenance. *Decision Analytics Journal*, (16), Article 100597. <https://doi.org/10.1016/j.dajour.2025.100597>

- European Automobile Manufacturers' Association. (2025). *Economic and market report: Global and EU auto industry – Full year 2024*. [https://www.acea.auto/files/Economic and Market Report-Full year-2024.pdf](https://www.acea.auto/files/Economic%20and%20Market%20Report-Full%20year%202024.pdf)
- European Automobile Manufacturers' Association. (2026, January). *Vehicles on European roads*. [https://www.acea.auto/files/ACEA Report-%E2%80%93-Vehicles on European roads 2026.pdf](https://www.acea.auto/files/ACEA%20Report-%E2%80%93-Vehicles%20on%20European%20roads%202026.pdf)
- Eurostat. (2024). *Key figures on European transport – 2024 edition*. Publications Office of the European Union. <https://doi.org/10.2785/9777356>
- Eurostat. (n.d.). *Passenger cars by age (road_eqs_carage)* [Data set]. https://ec.europa.eu/eurostat/databrowser/product/page/road_eqs_carage
- Iqbal, M., Suhail, S., Matulevičius, R., Shah, F. A., Malik, S. U. R., & McLaughlin, K. (2025). IoV-TwinChain: Predictive maintenance of vehicles in internet of vehicles through digital twin and blockchain. *Internet of Things*, (30), Article 101514. <https://doi.org/10.1016/j.iot.2025.101514>
- Michailidis, E. T., Panagiotopoulou, A., & Papadakis, A. (2025). A review of OBD-II-based machine learning applications for sustainable, efficient, secure, and safe vehicle driving. *Sensors*, 25(13), Article 4057. <https://doi.org/10.3390/s25134057>
- National Highway Traffic Safety Administration. (2026, January 24). *Recalls data (NHTSA campaigns)* [Data set]. <https://catalog.data.gov/dataset/recalls-data>
- Rahim, M. A., Rahman, M. M., Islam, M. S., Muzahid, A. J. M., Rahman, M. A., & Ramasamy, D. (2024). Deep learning-based vehicular engine health monitoring system utilising a hybrid convolutional neural network/bidirectional gated recurrent unit. *Expert Systems with Applications*, (257), Article 125080. <https://doi.org/10.1016/j.eswa.2024.125080>
- SAE International. (2021). *J1979-2_202104 – E/E Diagnostic Test Modes: OBDOnUDS*. https://www.sae.org/standards/j1979-2_202104-e-e-diagnostic-test-modes-obdonuds
- Theissler, A., Perez-Velazquez, 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.res.2021.107864>
- World Bank. (n.d.). *Motor vehicles (per 1,000 people) (indicator IS.VEH.NVEH.P3): Metadata and definition*. World Development Indicators. <https://databank.worldbank.org/metadataglossary/world-development-indicators/series/IS.VEH.NVEH.P3>
- Zheng, P., Zhao, W., Lv, Y., Qian, L., & Li, Y. (2023). Health status-based predictive maintenance decision-making via LSTM and Markov decision process. *Mathematics*, 11(1), Article 109. <https://doi.org/10.3390/math11010109>