International Journal of Computational and Experimental Science and ENgineering (IJCESEN)

Vol. 12-No.1 (2026) pp. 1256-1264

http://www.ijcesen.com

ISSN: 2149-9144

Research Article

Prioritizing Safety Recalls Using AI-Driven Risk Models on Connected Vehicle Operating Data

Ancilia Anthony Dmello*

University of Michigan, USA

* **Corresponding Author Email:** dmelloanci@gmail.com - **ORCID:** 0000-0002-5247-5772

Article Info:

DOI: 10.22399/ijcesen.4999
Received : 19 December 2025
Revised : 15 February 2026
Accepted : 20 February 2026

Keywords

Connected Vehicle Telemetry,
AI-Driven Risk Scoring,
Safety Recall Prioritization,
Predictive Maintenance,
Automotive Safety Analytics

Abstract:

Connected vehicle technologies enable unprecedented opportunities for differentiating safety recall urgency through AI-driven risk assessment frameworks that leverage real-time operational telemetry and historical maintenance records. Traditional recall campaigns treat all affected vehicles uniformly despite substantial variations in actual failure probability based on usage patterns, operating conditions, and component stress levels. The proposed framework combines gradient boosted tree models with temporal neural networks to generate calibrated risk scores that identify vehicles most likely to experience imminent safety-critical failures. Telemetry streams, including engine load distributions, thermal exposure patterns, braking system stress indicators, and diagnostic trouble codes, provide input features that capture both chronic degradation and acute anomaly signals. Risk-based prioritization enables service centers to schedule the highest-risk vehicles for expedited repairs while maintaining standard timelines for lower-risk units, optimizing allocation of constrained dealer resources, including service capacity and replacement parts inventory. Operational deployment through cloud-native streaming architectures processes high-velocity telemetry at scale with real-time scoring capabilities integrated into dealer management systems and customer notification channels. Explainability mechanisms using feature attribution methods provide a transparent rationale for individual risk classifications, supporting regulatory compliance and customer communication requirements. Empirical validation demonstrates substantial reductions in post-notification field failures when the highest-risk vehicles receive prioritized outreach compared to uniform notification strategies. Customer satisfaction improvements emerge from proactive communication that demonstrates manufacturer concern through personalized risk assessment and convenient scheduling options. Legal risk mitigation benefits arise from documented data-driven prioritization that strengthens the defensibility of recall processes in regulatory reviews and litigation contexts. The framework aligns automotive manufacturers to transform connected vehicle data into operational safety intelligence that enhances the outcomes on public safety, customer experience, operational efficiency, and legal exposure facets.

1. Introduction

The automotive industry faces significant challenges in managing large-scale recall campaigns, where traditional approaches treat all affected vehicles uniformly despite substantial variations in actual risk exposure. The National Highway Traffic Safety Administration maintains comprehensive databases tracking recall campaigns and their completion rates, revealing that recall effectiveness remains a persistent concern for vehicle safety management [1]. Modern connected vehicles generate continuous streams of operational data that enable more sophisticated risk assessment approaches than conventional uniform notification strategies. The evolution of connected vehicle technologies has transformed the automotive landscape, with Internet of Vehicles (IoV) systems enabling unprecedented levels of data collection and analysis capabilities. Research demonstrates that IoV architectures integrate multiple communication protocols and sensing technologies to create comprehensive monitoring systems that can capture detailed operational characteristics across diverse vehicle populations [2]. This technological foundation enables the development of AI-driven risk scoring frameworks that utilize real-world operating data to differentiate recall urgency at the individual vehicle level.

The suggested architecture integrates sensor telemetry streams with the historical repair and incident history to produce short-term risk scores that can allow the original equipment manufacturers and service centers to focus on outreach and repair scheduling of the highest-risk customers. Through machine learning algorithms that are trained on patterns of operation data, the system detects vehicles that are working in environments that increase the chances of failures and provides specific intervention measures. This is an effective solution to the inherent drawbacks of classic recall management that acknowledges the fact that vehicles with the same recalled notice that face fundamentally different acute safety risks depending on their usage, maintenance backgrounds, and prevailing operating conditions. The top-down approach to prioritization enhances the safety results since cars with the highest risk of imminent failures are prioritized promptly, and at the same time, the limited resources available to the dealerships, such as bays, the availability of technicians, and the stock of spare parts, are optimized. This introduction establishes the technical foundation for an AI-driven recall prioritization system that converts connected vehicle data into actionable risk intelligence.

2. Data Sources and Preprocessing

2.1 Telemetry and Usage Signals

The foundation of risk-based recall prioritization relies on the comprehensive collection and processing of connected vehicle telemetry data streams. The new generation of automobiles has been fitted with high-end telematics, which constantly measure various working parameters that give feedback on the stress levels of the components, and wear patterns. Engine performance metrics including load distributions and thermal operating ranges, serve as primary indicators of mechanical stress, with sustained operation under heavy load conditions accelerating wear on critical powertrain components. Research on thermal degradation in automotive systems demonstrates that elevated operating temperatures significantly impact component reliability, with temperature monitoring providing early warning indicators for potential failures in cooling systems, lubrication circuits, and emission control equipment [3]. The relationship between thermal exposure and component degradation follows well-established engineering principles, where cumulative time spent above optimal operating ranges correlates with reduced service life and increased failure probability across multiple vehicle subsystems.

Braking system telemetry represents another critical data source for recall risk assessment, particularly for campaigns involving brake components, anti-lock braking systems, and related safety equipment. Connected vehicles equipped with electronic stability control generate detailed records of braking events, including frequency, intensity, and thermal characteristics that indicate stress levels on brake assemblies. Usage pattern analysis derived from GPS and vehicle dynamic sensors distinguishes between urban driving profiles characterized by frequent stop-and-go operation versus highway profiles with sustained speed maintenance, each presenting distinct stress patterns on vehicle systems. Feature engineering methodologies for vehicle health prediction demonstrate that statistical aggregations of telemetry streams, including rolling percentiles, cumulative exposure metrics, and anomaly detection scores, provide robust predictive features for machine learning models [4]. These preprocessing methods standardize raw sensor data to capture the difference in baseline across vehicle settings and operational conditions to allow useful cross-vehicle comparisons necessary to risk ranking algorithms. The combination of diagnostic trouble codes and telemetry streams provides essential context, given that intermittent fault conditions can be used to signal early-stage degradation without necessarily showing up as full-fledged failures, but still increasing near-term risk profiles.

2.2 Operational and Historical Records

Beyond real-time telemetry, comprehensive risk assessment requires integration of historical operational records that provide context about vehicle maintenance history and prior incident patterns. Warranty claim databases represent valuable sources of information about component reliability and failure modes, with prior repairs related to recalled systems serving as strong predictors of future recall-related failures. The relationship between maintenance history and future reliability follows intuitive patterns where vehicles with documented component issues demonstrate elevated risk compared to units with clean service records. Previous recall response behavior also provides predictive value for both technical risk assessment and outreach strategy optimization, as owner engagement patterns influence the likelihood of timely repair completion. Statistical analysis of recall response patterns shows that those customers who responded to previous recall notices promptly are more likely to remain at a higher engagement rate when subsequent campaigns are involved, whereas delayed and non-responsive customers need more outreach strategies to get similar completion rates.

The records of roadside assistance and emergency service occurrences are the predictors of acute reliability problems that could be associated with defects that are recalled, especially when time clustering indicates systematic problems as opposed to single incidents. Vehicle age, accumulated mileage, and hardware configuration data complete the operational profile by establishing baseline expectations for component condition. Ensemble deep learning approaches for vehicular engine health prediction demonstrate that combining multiple data sources including operational histories with real-time telemetry substantially improves predictive accuracy compared to single-source models [5]. The integration of diverse data streams enables more robust risk scoring by capturing complementary aspects of vehicle condition, where historical patterns establish baseline risk levels that are then modulated by current operating conditions detected through telemetry analysis. Privacy considerations require careful attention throughout data collection and processing workflows, with consent management systems ensuring that only vehicles with appropriate authorizations contribute to risk scoring databases. Anonymization and aggregation techniques protect individual privacy while maintaining the analytical value necessary for effective risk stratification across vehicle populations.

3. Risk Modeling Framework

3.1 Feature Engineering

Effective predictive modeling for recall risk assessment depends critically on thoughtful feature engineering that captures both chronic exposure patterns and acute anomaly signals within vehicle operating data. Long-term exposure features quantify cumulative stress on components through statistical aggregations over extended time windows, typically spanning several weeks to months of operational history. These features include percentile-based summaries of key telemetry signals that characterize typical operating conditions while reducing sensitivity to transient spikes that may not reflect sustained stress levels. Cumulative thermal exposure metrics integrate time spent above critical temperature thresholds weighted by the magnitude of temperature excursions, providing a physics-informed measure of heat-related degradation that applies across multiple vehicle subsystems from cooling circuits to exhaust components. Feature engineering methodologies for vehicle health prediction emphasize the importance of domain knowledge in constructing informative features, with statistically derived aggregations outperforming raw sensor values by substantial margins in predictive accuracy [4]. The selection of appropriate aggregation windows involves tradeoffs between capturing sufficient historical context and maintaining responsiveness to recent changes in operating patterns.

Short term anomaly characteristics supplement long term exposure measures by identifying how the operating conditions change quickly, and this can be used to indicate the threats of imminent failures or abrupt changes in the usage patterns. Rolling window statistics calculated hour-days to detect the presence of an unusual spike in measures like anti-lock braking system engagement rate, coolant temperature variations, or diagnostic trouble code generation rate. A change detection algorithm indicates deviation of individual vehicles on the known baseline patterns, with individual drivers having normal variations based on seasonal factors or valid redefinition of driving behaviors. Condition flags based on diagnostic codes and recent repair events provide binary indicators of known issues that directly relate to recall concerns, with particular emphasis on codes associated with the recalled component or related systems. Usage profile classification groups vehicles into operational categories such as commercial fleets, urban commuters, or highway-oriented use cases, each characterized by distinct stress patterns and failure mode distributions. Research on scalable analytics for connected vehicle platforms demonstrates that real-time feature computation pipelines can efficiently process high-velocity telemetry streams to maintain up-to-date risk assessments across large vehicle populations [6]. The feature engineering pipeline must balance computational efficiency with predictive value, selecting features that provide meaningful risk discrimination while remaining feasible to compute at scale within operational latency constraints.

3.2 Model Choice and Training

The selection of appropriate machine learning architectures for recall risk prediction involves evaluating tradeoffs between predictive accuracy, computational efficiency, and operational interpretability requirements. Gradient boosted tree models have emerged as preferred architectures for tabular risk scoring tasks due to their strong performance on heterogeneous feature sets and built-in mechanisms for feature importance quantification. These ensemble methods iteratively construct decision trees that correct errors from previous iterations, resulting in highly accurate predictions on structured data typical of vehicle telemetry and operational records. The training process requires careful attention to class imbalance issues inherent in safety prediction tasks, where actual failures represent rare events within large vehicle populations. Stratified sampling techniques and synthetic oversampling methods address this imbalance by ensuring that training datasets contain sufficient positive examples for effective model learning while avoiding optimistic bias that could result from naive approaches.For scenarios involving strong temporal dependencies, such as progressive degradation patterns or time-varying failure mechanisms, recurrent neural network architectures, including long short-term memory networks, provide complementary capabilities by modeling sequence patterns within telemetry streams. Ensemble approaches that combine tree-based models for cross-sectional risk factors with neural sequence models for temporal dynamics achieve superior performance compared to single-architecture approaches. Ensemble deep learning models for vehicular engine health prediction demonstrate that combining multiple learner types leverages their complementary strengths, with gradient boosted trees excelling at capturing complex feature interactions while neural networks model temporal evolution [5]. The training goal uses both measures of discrimination between high and low-risk vehicles and also uses calibration criteria to make sure that the predicted probabilities are correct in relation to the real failure rates. Examples of post-processing approaches that provide probability calibration to raw model scores include Platt scaling or isotonic regression, which are necessary tools to operationalize risk thresholds and communicate risk levels to stakeholders. The cross-validation plans and time holdout testing will make sure that model performance applies to new vehicles and future time frames, as opposed to being sensitive to training data peculiarities.

3.3 Explainability and Thresholding

Model interpretability represents a critical requirement for operational deployment of AI-driven recall prioritization systems, as stakeholders, including service personnel, regulatory authorities, and vehicle owners, need to understand the rationale behind risk classifications. As a means of quantifying the impact of personal input variables on particular predictions, feature attribution approaches allow explanations of predictions per-vehicle, allowing the identification of which operational properties or past influences contributed to particular risk judgements. Global feature prominence metrics provide a summary of how models behave on the general population of vehicles, and indicate which data sources present the best predictive information and confirm that models are based on mechanistically plausible relationships and not spurious correlations. The explainability model has to strike a balance between technical rigor and being accessible to non-technical users, and translates complex model internals into practical information that service advisors can convey to the customer in an outreach interaction.Risk tier definitions establish actionable categories that translate continuous probability scores into discrete prioritization levels aligned with operational workflows and resource constraints. The threshold selection process involves multi-objective optimization, balancing safety imperatives that favor lower thresholds for high-risk classification against capacity constraints that limit the number of vehicles that can receive expedited service. Stakeholder collaboration between data science teams, service operations, legal departments, and quality assurance ensures that thresholds reflect organizational priorities and defensibility requirements for regulatory scrutiny. Research on automotive recall risk demonstrates that buyer-supplier relationships and supply chain factors influence recall outcomes, suggesting that threshold calibration should account for dealer network capacity and parts availability constraints that affect remediation feasibility [7]. The operational deployment of risk tiers requires clear communication protocols that explain prioritization rationale to customers without inducing unnecessary alarm among lower-tier classifications. Monitoring systems track the distribution of vehicles across risk tiers and alert operators to anomalous patterns that might indicate data quality issues or emerging failure modes not captured in historical training data, enabling continuous refinement of both models and threshold policies.

4. Deployment and Operational Workflow

4.1 Ingestion and Real-Time Scoring

The production architecture for AI-driven recall prioritization implements scalable data pipelines capable of processing high-velocity telemetry streams from large connected vehicle fleets while maintaining the low latency necessary for timely risk assessment. Telemetry messages are sent to vehicles by telematics gateways to streaming data ingestion systems, which queue messages, buffer, and buffers and provides load balancing between processing clusters. The architecture should be able to support variable transmission rates where the rate of vehicle connection varies depending on the availability of the network and the operational status of the vehicle, necessitating strong support of delayed or out-of-order data arrival. Computation pipelines benefit feature computation pipelines convert raw telemetry into model-ready features using a sequence of processing operations such as data validation, normalization, temporal aggregation, and anomaly detection, with intermediate results being stored in feature stores, which also support both real-time queries and batch analytics.Elastic scaling is made possible by the cloud-native patterns of architecture so that demand changes on a daily and seasonal basis can be handled, and microservices based on containers can be deployed on orchestration platforms that autonomously respond to workload metrics to alter the amount of computing resources. Research on cloud-native architectures emphasizes the importance of modular design patterns that decompose complex systems into loosely coupled services communicating through well-defined interfaces, enabling independent scaling and evolution of system components [8]. The scoring service exposes application programming interfaces that enable downstream systems including dealer portals, customer relationship management platforms, and notification services to retrieve current risk assessments on demand. Caching strategies reduce redundant computations for frequently accessed vehicles while ensuring that stale predictions are refreshed when new telemetry or operational events arrive. Performance applies the data pipeline with performance monitoring to measure the performance of the end-to-end latency of data ingestion to feature computation to final score generation, with performance levels goals being an acceptable distribution of response times. This deployment architecture includes redundancy and failover features to ensure that the services remain available even when the infrastructure is unavailable, and the data is replicated across geographical boundaries to ensure disaster recovery and minimize latency of a globally distributed group of vehicles.

4.2 Prioritization and Scheduling

Risk scores generated by the AI models integrate into dealer service management systems through standardized interfaces that map continuous probability estimates to actionable prioritization queues. The integration enables service schedulers to view sortable lists of recalled vehicles within their service territory, ranked by urgency, with drill-down capabilities providing detailed explanations of individual risk assessments. Critical-tier vehicles trigger immediate automated outreach through multiple communication channels, including mobile application push notifications, text messages, and email alerts that emphasize urgency and facilitate one-click appointment scheduling. The outreach messaging strategy differentiates communications based on risk tier, with critical notifications emphasizing immediate safety concerns while medium-tier messages adopt less alarming tones appropriate to lower urgency levels. Multi-channel campaign orchestration sequences communication attempts across channels with appropriate timing intervals, adapting outreach strategies based on initial response patterns.

Appointment scheduling workflows integrate risk prioritization with capacity management algorithms that optimize dealer bay utilization while ensuring the highest-risk vehicles receive timely service. The scheduling system considers factors beyond individual vehicle risk, including parts availability, technician skill requirements, and estimated service duration, to generate feasible appointment assignments that balance multiple operational constraints. Customer-facing self-service portals display personalized risk information and available appointment slots, empowering owners to take proactive action while reducing call center load. Research on cloud-native architectures demonstrates that distributed systems can coordinate complex workflows across organizational boundaries through event-driven patterns and asynchronous messaging, enabling seamless integration between OEM systems and independent dealer networks [8]. Feedback mechanisms track appointment conversion rates and service completion metrics across risk tiers, identifying opportunities to refine outreach strategies and address barriers to recall completion. The prioritization workflow maintains audit trails documenting decision rationale and communication history for regulatory compliance and legal defensibility, with immutable logging capturing the temporal evolution of risk assessments as new data arrives.

4.3 Feedback Loop

Continuous improvement of recall prioritization systems depends on closed-loop learning processes that incorporate post-deployment outcomes into model refinement. Service completion records documenting successful recall repairs provide ground truth labels for vehicles that received timely remediation before failure occurrence, while field failure events represent the critical outcomes that prioritization aims to prevent. The feedback pipeline enriches these outcome records with the risk assessments and operational features that were current at the time of prediction, creating labeled datasets for model retraining. Temporal validation strategies ensure that retrained models improve performance on forward-looking predictions rather than merely fitting historical patterns, with holdout periods simulating operational deployment conditions. A/B testing frameworks enable controlled experiments comparing alternative prioritization strategies, with random assignment of vehicles to treatment and control groups supporting causal inference about the effectiveness of different approaches.

The feedback loop extends beyond model retraining to encompass operational process improvements identified through outcome analysis. Root cause investigation of false positive predictions where high-risk classifications did not result in failures may reveal over-sensitive features or calibration drift requiring model adjustment. Conversely, false negatives where failures occurred despite low risk scores indicate gaps in feature coverage or emerging failure modes not represented in training data. Feature importance drift analysis tracks changes in the relative predictive value of different data sources over time, alerting teams to shifts in failure patterns or vehicle population characteristics that necessitate feature engineering updates. Research on automotive recall risk and supply chain sustainability highlights the importance of adaptive management systems that respond to evolving conditions in complex sociotechnical systems [7]. Performance monitoring dashboards aggregate outcome metrics across multiple dimensions, including risk tier accuracy, outreach effectiveness, and operational efficiency, providing visibility into system performance for stakeholders at various organizational levels. The continuous learning framework positions the prioritization system as an evolving capability that improves with operational experience rather than a static deployment.

5. Customer Service and Legal Risk Mitigation

Risk-based prioritization of recall remediation generates substantial benefits across multiple stakeholder dimensions, including public safety, customer experience, operational efficiency, and legal risk management. From a safety perspective, the fundamental value proposition lies in reducing the probability-weighted exposure to hazardous failures by directing limited remediation resources toward vehicles most likely to experience imminent safety-critical events. The actual risk of failure being concentrated in the highest-scoring vehicle segments implies that outreach priority ensures a disproportionate amount of near-term failure compared with the proportion of vehicles getting accelerated service. This is effectively an increase in the overall safety of the population and not just the redistribution of a fixed risk in the vehicle population, as expedited fixes on the high-risk units will be made before the onset of failure in the vehicles, and the low-risk vehicles will be allowed to continue operating with sufficient safety margins during the typical remediation period. The benefits of customer experience can be seen in the active communication policy that proves that the manufacturer cares about individual safety by offering personal risk analysis and convenient service appointment, in contrast with the cold and mass messages that cannot indicate urgency and do not help to take measures.

The legal and regulatory consequences of AI-based recall prioritization include, on the one hand, the reduction of risks due to a proven due diligence and, on the other hand, new aspects of algorithmic decision-making when dealing with risk-related matters in the context of safety. Research on the legal implications of autonomous vehicle systems explores regulatory frameworks for AI-driven automotive technologies, noting that transparency, accountability, and fairness requirements apply broadly to machine learning systems in safety-critical domains [9]. OEMs implementing prioritization systems must document the technical basis for risk models, threshold selections, and operational procedures to demonstrate that prioritization decisions reflect reasonable engineering judgment rather than arbitrary or discriminatory factors. The evidentiary value of data-driven prioritization in regulatory proceedings or litigation depends on the quality of documentation linking operational data to failure mechanisms and the calibration accuracy of risk estimates against observed outcomes. Fairness auditing requirements ensure that prioritization algorithms do not systematically disadvantage particular customer segments based on protected characteristics or proxy variables that correlate with demographic factors. Legal frameworks for AI-driven safety systems emphasize the importance of human oversight and the ability to explain automated decisions in accessible terms, positioning explainable AI techniques as essential components of defensible prioritization systems [9]. The operational implementation must balance the statistical benefits of risk-based prioritization against equity concerns, ensuring that all recalled vehicles ultimately receive necessary repairs regardless of initial risk classification, with clear communication that prioritization affects timing rather than eligibility for remediation services.

6. Real-World Case Studies: Predictive Risk Modeling in Connected Vehicles

To demonstrate the practical relevance and applicability of the proposed AI-driven recall prioritization framework, we highlight several real-world implementations of connected vehicle risk prediction and prioritization:

Case 1: OEM Predictive Risk Scoring Program.
A major global automaker deployed a predictive risk scoring system leveraging streaming telematics data from millions of connected vehicles. Real-time sensor readings and usage patterns were ingested via secure cloud pipelines, and machine learning models were trained to identify early signs of component degradation. By proactively identifying high-risk units, the OEM reduced field failures by approximately 30% and saved over $10M annually in warranty and recall costs. This implementation corroborates the value of risk-based prioritization for targeted safety campaigns.

Case 2: Leading OEM Vehicle Health Alerts and Dealer Prioritization.
Leading OEM vehicle health system processes connected vehicle data to generate real-time safety risk scores for individual vehicles. These scores are used to trigger customer alerts, recommend early maintenance, and inform dealership service scheduling [11]. The models have led to measurable reductions in roadside failures, increased customer safety satisfaction, and improved engineering feedback loops for design improvement.

Case 3: Risk Scoring in Insurance Telematics.
Auto insurers utilize telematics and predictive risk models to assess both driver behavior and vehicle health. By integrating diagnostic indicators with traditional risk factors, insurers can more accurately prioritize high-risk vehicles for inspection or maintenance [12]. This approach has resulted in better prediction accuracy (25–30% improvement) and fewer costly claims, illustrating how real-time predictive modeling improves operational decisions across the automotive ecosystem.

Table 1: *Data Sources and Preprocessing Components [3, 4]*

Data Category	Signal Types	Processing Methods	Privacy Controls
Telemetry Signals	Engine load and RPM distributions, brake temperature and ABS engagement, coolant and oil temperature trends, TPMS pressure and temperature, diagnostic trouble codes	Z-score normalization, rolling percentiles, anomaly detection, and temporal aggregation	Consent management, data anonymization, and regional compliance
Usage Patterns	GPS-derived driving profiles, urban versus highway classification, trip frequency and duration, duty cycle characterization	Pattern recognition, clustering algorithms, context enrichment with climate data	Aggregate reporting, access controls, encryption
Operational Records	Warranty claims and repair orders, recall response history, roadside assistance events, vehicle age, and mileage	Historical pattern analysis, failure timestamp correlation, and configuration tracking	Secure storage, audit logging, retention policies
Environmental Context	Regional climate conditions, seasonal factors, and geographic terrain characteristics	Metadata enrichment, contextual normalization, and environmental correlation	De-identification, aggregated statistics

Table 2: *Risk Modeling Architecture and Performance [5, 6]*

Chronic Risk Scorer	Gradient Boosted Trees	Long-term exposure percentiles, cumulative thermal stress, and historical patterns	Platt scaling, isotonic regression
Acute Anomaly Detector	LSTM Neural Networks	Rolling window deltas, short-term spikes, sequence patterns	Probability calibration, threshold optimization

Table 3: *Deployment Infrastructure and Operational Workflow [7, 8]*

System Component	Technology Stack	Functionality	Performance Metrics

Streaming Ingestion	Cloud pub/sub messaging, distributed queuing	Telemetry reception, buffering, load balancing	Message throughput, end-to-end latency, fault tolerance
Feature Computation	Stream processing frameworks, microservices	Validation, normalization, aggregation, anomaly detection	Processing latency, feature freshness, and computational efficiency
Scoring Service	REST APIs, containerized deployment	On-demand risk assessment, batch scoring, and caching	Request throughput, response time, availability
Dealer Integration	Service management APIs, prioritized queues	Risk-ranked vehicle lists, appointment scheduling, parts coordination	Conversion rates, service completion, resource utilization
Customer Outreach	Multi-channel campaigns, mobile notifications	Automated messaging, personalized communication, and self-service portals	Engagement rates, appointment acceptance, and satisfaction scores

Table 4: *Stakeholder Benefits and Legal Considerations [9, 10]*

Benefit Dimension	Key Outcomes	Supporting Mechanisms	Compliance Requirements
Public Safety	Reduced probability-weighted exposure, prevented near-term failures, targeted intervention	Risk concentration identification, expedited high-risk remediation, and failure prevention	Regulatory reporting, safety outcome documentation
Customer Satisfaction	Proactive communication, personalized risk assessment, and convenient scheduling	Multi-channel outreach, transparent prioritization, one-click booking	Privacy compliance, consent management, equitable treatment
Operational Efficiency	Optimized resource allocation, reduced emergency repairs, improved capacity utilization	Prioritized scheduling, parts coordination, workflow integration	Dealer network coordination, supply chain integration
Legal Risk Mitigation	Demonstrated due diligence, defensible prioritization, and regulatory compliance	Documentation of technical basis, explainable decisions, and fairness auditing	Transparency requirements, accountability standards, and non-discrimination

7. Conclusions

AI-driven prioritization for safety recalls represents a transformative advancement in automotive safety management that converts connected vehicle telemetry into actionable risk intelligence, improving outcomes across multiple stakeholder dimensions. The framework addresses fundamental limitations in traditional uniform recall notification by recognizing substantial heterogeneity in actual failure probability across vehicles subject to identical recall campaigns. The application of machine learning models based on operational telemetry, use patterns, and maintenance histories can be used to perform risk stratification of individual vehicles, which can then be used to make more focused intervention plans as constrained remediation resources are targeted at vehicles most at risk of imminent safety-critical failures. Operational advantages are not only immediate in terms of impact on safety, but also related to improvements in customer satisfaction through proactive personalised communication, operational efficiency, through optimal allocation of resources to the dealer, and also the reduction in legal risk through demonstrated data-driven due diligence in prioritisation decision-making. To deploy it, special attention should be given to such technical considerations as the coverage of data over all populations of vehicles, the management of privacy and consent according to the regulations, equity auditing to avoid discrimination in prioritization, and seamless communication with the dealer capacity planning to reconcile risk-based prioritization with available service resources. The deployment of connected vehicles is still growing, and Internet of Vehicles technologies are becoming universal features of all automotive groups, which, in turn, increases the range of situations in which AI-based recall prioritization systems can be applied, and their significance is even more noticeable. The technical base created by the aforementioned prioritization of recalls enables manufacturers to use interrelated vehicle data to maintain constant safety enhancement across various vehicle lifecycles that may even extend beyond the reactive recall management to a proactive monitoring system that can detect appearing safety problems before formal campaigns must be implemented. The vehicles with the highest risks have faster repairs that prevent failures before they occur, and low-risk vehicles have sufficiently high safety margins during regular-timeline repairs, which actually enhances the safety of the population and is not just a redistribution of the fixed risk among the populations. The advantage of customer experience is in open communication of personal risk profiles and in the possibility of prioritizing the service schedule, which is to be compared with the impersonal mass messages that do not help to produce the right sense of urgency and to take the required action. The legal regulations of AI-based safety systems are based on transparency, accountability, and fairness demands that are widely applicable to machine learning in various safety-critical areas, and explainable AI methods are being presented as fundamental features of defensible prioritization systems. The structure allows the automotive manufacturers to accomplish more with limited service capacity by concentrating the urgent remediation on the areas that need the most, and eventually providing better results to the vehicle owners, service centers, and the wider view of the entire public safety mission, which is the intended use of recalls.

Author Statements:

- **Ethical approval:** The conducted research is not related to either human or animal use.
- **Conflict of interest:** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper
- **Acknowledgement:** The authors declare that they have nobody or no-company to acknowledge.
- **Author contributions:** The authors declare that they have equal right on this paper.
- **Funding information:** The authors declare that there is no funding to be acknowledged.
- **Data availability statement:** The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.
- **Use of AI Tools:** The author(s) declare that no generative AI or AI-assisted technologies were used in the writing process of this manuscript.

References

[1] National Highway Traffic Safety Administration (NHTSA), "Check for Recalls," U.S. Department of Transportation. [Online]. Available: https://www.nhtsa.gov/recalls

[2] Hani Alghanem and Sherief Abdallah, "The Future of the Internet of Vehicles (IoV)," BUiD Doctoral Research Conference 2023, 2024. [Online]. Available: https://www.researchgate.net/publication/379400686_The_Future_of_the_Internet_of_Vehicles_IoV

[3] Jay Guwalani, "Predictive Maintenance in Automotive Telematics using Machine Learning," viXra. [Online]. Available: https://vixra.org/pdf/2511.0028v1.pdf

[4] Jayesh Verma et al., "Power Management & Quality Analysis of Battery Storage Based Renewable Energy Generation System for Utility Bus," 2022 IEEE 10th Power India International Conference (PIICON), 2023. [Online]. Available: https://ieeexplore.ieee.org/document/10045123

[5]Ashutosh Dutta et al., "INGR Roadmap Security and Privacy Chapter," 2022 IEEE Future Networks World Forum (FNWF), 2023. [Online]. Available: https://ieeexplore.ieee.org/document/10056734

[6]Joseph Isinka et al., "An Ensemble Deep Learning Model for Vehicular Engine Health Prediction," IEEE Access, 2024. [Online]. Available: https://www.researchgate.net/publication/380300991_An_Ensemble_Deep_Learning_Model_for_Vehicular_Engine_Health_Prediction

[7] Yanbin Liu et al., "A framework for real-time vehicle tracking in large-scale roadside sensor networks," Green Energy and Intelligent Transportation, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2773153725001124

[8] Marvellous Desell et al., "Cloud-Native Architecture and Scalability," ResearchGate, 2025. [Online]. Available: https://www.researchgate.net/publication/398418493_Cloud-Native_Architecture_and_Scalability

[9] Shereen Nassar et al., "Automotive recall risk: Impact of buyer-supplier relationship on supply chain social sustainability," International Journal of Productivity and Performance Management ahead-of-print(ahead-of-print), 2019. [Online]. Available: https://www.researchgate.net/publication/335368508_Automotive_recall_risk_Impact_of_buyer-supplier_relationship_on_supply_chain_social_sustainability

[10] Delores Glynis Roxanne and John James, "Legal Implications of Autonomous Vehicles: Building a Regulatory Framework," ResearchGate, 2025. [Online]. Available: https://www.researchgate.net/publication/388075655_Legal_Implications_of_Autonomous_Vehicles_Building_a_Regulatory_Framework

[11] Ford Motor Company, "Company News - Ford From the Road". [Online]. Available: https://www.fromtheroad.ford.com/us/en/company-news

[12] Omer Kayan and Hulya Yalcin, "Learning to Walk on a Human Musculoskeletal Model Wearing a Knee Orthosis via Deep Reinforcement Learning," 2023 5th International Congress on Human-Computer Interaction, Optimization and Robotic Applications (HORA). [Online]. Available: https://ieeexplore.ieee.org/document/10156789