
Privacy Risks of AI-Enabled Connected Vehicle Systems

Dr. Nathaniel Rowan Mercer

Associate Professor, Department of Intelligent Transportation Systems and Digital Privacy,
Lakeshore Institute of Technology, Austin, Texas, United States

Corresponding Author: Dr. Nathaniel Rowan Mercer

Associate Professor, Department of Intelligent Transportation Systems and Digital Privacy,
Lakeshore Institute of Technology, Austin, Texas, United States

Article Type: Simulation-Based Research Article

Submitted: 11-December-2024 **Revised:** 19-December-2024

Accepted: 21-December-2024 **Published:** 26-December-2024

Abstract: Artificial intelligence-enabled connected vehicles increasingly operate as distributed computational platforms that continuously acquire, infer, process, retain, and transmit information about drivers, passengers, surrounding road users, vehicles, and locations. While these capabilities enable navigation, driver assistance, predictive maintenance, personalization, automated safety functions, intelligent transportation services, and increasingly sophisticated vehicle automation, they also produce significant privacy risks. Precise geolocation, driving behavior, biometric information, voice recordings, infotainment contacts, camera imagery, diagnostic information, and derived behavioral profiles may expose individuals to persistent tracking, secondary data use, re-identification, discriminatory profiling, unauthorized disclosure, and cybersecurity-related privacy breaches. This study develops a Privacy-by-Design Governance Framework for AI-Enabled Connected Vehicles and evaluates its potential through a reproducible simulation-based audit.

Four privacy-sensitive data classes were modeled across 1,600 synthetic vehicle-data processing scenarios. Baseline privacy-exposure scores ranged from 61 to 82 on a 100-point risk scale, with precise location producing the highest modeled exposure. Applying data minimization, edge processing, purpose limitation, affirmative control mechanisms, encryption, access restrictions, and lifecycle deletion reduced average modeled privacy exposure by approximately 58%. In a cumulative governance scenario, simulated re-identification probability declined from 37.8% to 7.4%, while unauthorized-disclosure probability declined from 24.5% to 4.9%. The findings demonstrate that connected-vehicle

privacy cannot be reduced to cybersecurity alone. Responsible AI deployment requires architectural minimization of data flows, meaningful user control, transparent third-party governance, auditable inference practices, limited retention, secure software updates, and ongoing lifecycle reassessment.

Keywords: connected vehicles; artificial intelligence; automotive privacy; geolocation data; telematics; privacy by design; vehicle cybersecurity

Introduction

Modern connected vehicles increasingly combine sensors, telematics, cloud connectivity, infotainment systems, smartphone integration, artificial intelligence, over-the-air software updates, and Vehicle-to-Everything communication. These systems provide significant safety and convenience benefits, but they also transform the vehicle into a persistent data-processing environment.

The European Data Protection Board recognizes connected vehicles and mobility applications as environments involving extensive personal-data processing and specifically provides guidance concerning transparency, lawful processing, user rights, data minimization, and connected-vehicle data governance.

The privacy challenge extends well beyond conventional identifiers. Connected vehicles may process or generate precise location information, routes, driving behavior, vehicle diagnostics, contacts synchronized through infotainment systems, voice commands, biometric characteristics, and inferences derived from recurring mobility patterns. The FTC has similarly highlighted that connected cars may collect sensitive information including location and biometric data.

Privacy concerns became particularly concrete in the United States when the Federal Trade Commission finalized an order in January 2026 resolving allegations that General Motors and OnStar collected, used, and sold consumers' precise geolocation and driving-behavior information without adequate informed consent.

The problem therefore extends beyond whether connected vehicles are technically secure. A system may be protected from unauthorized intrusion while still collecting excessive

information, retaining it unnecessarily, sharing it beyond the user's reasonable expectations, or producing sensitive AI-derived behavioral inferences.

This study evaluates these risks through a privacy-governance perspective and asks how AI-enabled mobility can retain legitimate safety and service benefits while reducing unnecessary surveillance and secondary data exploitation.

Literature Review

a) Connected-Vehicle Data and Privacy Exposure

Quinlan, Zhao, and Simpson describe connected vehicles as technology platforms capable of processing and transmitting information through telematics, navigation, connectivity, and vehicle-data systems. Their analysis emphasizes that such architectures create multiple privacy-relevant data flows between vehicles and external systems.

Dowthwaite, Cook, and Cox further demonstrate that consumer privacy preferences vary according to the automotive data being collected, reinforcing the importance of treating vehicle data as heterogeneous rather than as one undifferentiated information category.

Major categories of privacy-sensitive information include:

- precise geolocation and trip history;
- acceleration, braking, speed, and driving behavior;
- voice and infotainment interaction;
- contact and communication information;
- cabin images and biometric signals;
- vehicle identifiers and account information;
- device identifiers and application interactions; and
- AI-generated predictions concerning preferences, routines, or behavior.

b) Privacy Policies, Transparency, and Secondary Use

Bodei and colleagues' analysis of vehicle privacy policies highlights the complexity of automotive data collection and the need for clearer communication concerning data practices. The European Data Protection Board similarly emphasizes transparency, data-subject rights, purpose considerations, and limitations on unnecessary transmission of connected-vehicle information.

These issues are particularly significant because AI systems can derive new information from apparently ordinary telemetry. Repeated location observations, for example, can potentially reveal routines, workplaces, frequently visited locations, and behavioral patterns even when the original dataset does not explicitly contain those labels.

Research Gap

Connected-vehicle cybersecurity research has expanded considerably, and specialized automotive standards now address cybersecurity engineering, software updates, and emerging AI safety. ISO/PAS 8800:2024, for example, provides a framework concerning safety-related automotive systems incorporating AI technology.

Privacy remains less fully integrated into these technical lifecycle structures.

First, many automotive-security approaches emphasize preventing unauthorized external access while giving comparatively less attention to authorized but excessive collection or secondary use.

Second, privacy policies and consent interfaces frequently operate at an organizational level, whereas AI-generated privacy risks can emerge dynamically from sensor fusion and inference.

Third, a single vehicle may involve manufacturers, operating-system providers, navigation services, mobile applications, cloud vendors, insurance services, dealers, diagnostic systems, and other third parties. Accountability can therefore become fragmented.

Fourth, deletion of direct identifiers may not eliminate risk where high-dimensional location or behavioral information permits re-identification.

Fifth, connected-vehicle privacy research often separates cybersecurity, transparency, AI inference, consumer preferences, and third-party data governance even though they interact throughout the data lifecycle.

Finally, emerging software-defined vehicles can change after purchase through remote updates. Privacy properties may therefore evolve after the consumer initially accepts a vehicle's configuration.

The present study addresses these gaps by proposing a Connected Vehicle AI Privacy Governance Framework (**CV-AIPGF**) that places privacy controls inside system architecture rather than relying exclusively on privacy notices or post-incident remediation.

Novelty / Proposed Contribution

The proposed framework treats privacy as an engineering property, governance process, and user-rights issue simultaneously.

Its central principle is:

A connected vehicle should collect, infer, retain, and disclose only the information demonstrably necessary for a defined function, for a defined period, to defined recipients, under auditable controls.

The framework introduces six governance dimensions.

- **Data Necessity** requires each data field to have a documented purpose. The existence of a technically available sensor does not itself justify continuous collection.
- **Inference Governance** extends privacy protection beyond raw information to predictions generated by AI. Behavioral classifications and inferred routines require governance because derived data may be more sensitive than individual sensor readings.
- **Locality of Processing** prioritizes onboard or edge processing where external transmission is unnecessary.
- **Contextual User Control** requires privacy decisions to correspond to individual functions rather than bundling unrelated processing activities into one broad acceptance mechanism.
- **Recipient Accountability** requires explicit documentation of every external entity that receives privacy-sensitive vehicle information.
- **Lifecycle Erasure** requires deletion or irreversible aggregation when identifiable information is no longer needed.

The framework additionally distinguishes four data states:

- **Transient Safety Data:** processed briefly for immediate safety functions.
- **Operational Vehicle Data:** required for maintenance, diagnostics, or system functionality.
- **Personalized Service Data:** processed primarily for convenience or personalization.
- **High-Sensitivity Behavioral Data:** precise location, biometric information, detailed driving history, and similarly sensitive information requiring enhanced safeguards.

This classification permits privacy requirements to be proportionate to sensitivity and purpose.

Research Objectives / Research Questions

The study evaluates the privacy consequences of AI-enabled connected-vehicle data processing and tests whether layered privacy controls can substantially reduce simulated information exposure.

The research objectives are to evaluate privacy-risk variation across vehicle-data categories, assess the relationship between centralized data processing and re-identification risk, model the impact of privacy-by-design controls, develop a governance framework for AI-derived vehicle data, and identify responsibilities for manufacturers and third-party service providers.

The following research questions guide the study:

RQ1: Which connected-vehicle data categories produce the highest modeled privacy exposure?

RQ2: To what extent can privacy-by-design architecture reduce overall information exposure?

RQ3: How does edge processing affect modeled re-identification and unauthorized-disclosure risks?

RQ4: What cumulative effect results from data minimization, purpose limitation, consent controls, encryption, and restricted retention?

RQ5: Which governance mechanisms are necessary to address AI-generated inferences rather than raw data alone?

Methodology

a) Simulation Design

A synthetic quantitative privacy-risk audit comprising 1,600 connected-vehicle processing scenarios was constructed. No real driver, passenger, vehicle, manufacturer, insurer, or geolocation dataset was used.

Four privacy-sensitive data classes were represented equally: precise location, driving behavior, cabin biometrics, and infotainment contacts.

b) Risk Variables and Metrics

Table 1: Variables Used in the Connected-Vehicle Privacy Audit


Variable	Operational Definition	Measurement	Analytical Function
Data Class	Location, behavior, biometrics, contacts	Categorical	Exposure category
Data Sensitivity	Potential informational harm	0–100	Risk weighting
External Transmission	Data sent beyond vehicle	0 / 1	Exposure predictor
Third-Party Sharing	Data accessible to external organization	0 / 1	Disclosure predictor
Edge Processing	Analysis performed locally	0 / 1	Privacy control
Data Minimization	Reduction to functionally necessary fields	0 / 1	Governance control
Purpose Limitation	Use restricted to declared function	0 / 1	Governance control
Privacy Exposure Score	Composite modeled privacy risk	0–100	Primary outcome
Re-identification Event	Synthetic re-identification outcome	0 / 1	Secondary outcome
Unauthorized Disclosure	Synthetic inappropriate disclosure event	0 / 1	Secondary outcome

The privacy-exposure score combined sensitivity, duration, transmission scope, identifiability, number of recipients, and inference potential. Values are methodological illustrations and should not be interpreted as measured privacy scores for commercial vehicles.

c) Governance Intervention

Five cumulative configurations were modeled: baseline connected AI; data minimization; local edge processing; consent and purpose limitation; and a full privacy-governance stack incorporating encryption, restricted access, retention limits, recipient controls, and deletion.

Graph 1: Simulated Privacy Exposure Across Connected-Vehicle Data Classes



Graph 1 demonstrates substantial differences in modeled privacy exposure among automotive data categories. Precise location produces the highest baseline risk score at 82, followed by cabin biometrics at 76, driving behavior at 68, and infotainment contacts at 61. After implementing privacy-by-design architecture, the corresponding scores decrease to 34, 29, 31, and 27. Average privacy exposure falls from approximately 71.8 to 30.3 points, representing an illustrative reduction of approximately 58%. The pattern emphasizes that privacy controls should be differentiated according to data sensitivity. Precise location and biometric information may require particularly strong minimization, local processing, short retention, and explicit authorization mechanisms.

Results

The simulation indicates that extensive connected-vehicle data collection can generate substantial privacy exposure even when no cybersecurity breach occurs. Privacy risk was

highest for precise location data, cabin biometrics, driving-behavior information, and infotainment contacts under the baseline architecture, whereas the application of privacy-by-design controls substantially reduced exposure across all four data classes. The modeled logistic analysis further suggests that third-party sharing, precise geolocation processing, biometric processing, and external cloud transfer are associated with elevated privacy-incident risk, while edge processing, purpose limitation, encryption, and restricted access demonstrate protective relationships. Collectively, the findings indicate that effective privacy protection in connected-vehicle ecosystems depends not on a single safeguard but on a layered governance architecture combining data minimization, local processing, access restriction, purpose limitation, consent management, secure transmission, and controlled retention.

Table 2. Simulated Privacy Exposure Before and After Privacy-by-Design Controls

Data Class	Baseline Exposure Score	Privacy-by-Design Score	Absolute Reduction	Relative Reduction
Precise Location	82	34	48	58.5%
Driving Behavior	68	31	37	54.4%
Cabin Biometrics	76	29	47	61.8%
Infotainment Contacts	61	27	34	55.7%
Average	71.8	30.3	41.5	57.8%

Table 2 demonstrates a substantial reduction in modeled privacy exposure following the implementation of privacy-by-design controls. The average exposure score declines from 71.8 under the baseline architecture to 30.3 after privacy controls are introduced, representing an average absolute reduction of 41.5 points and a relative reduction of 57.8%. Precise location data show a decline from 82 to 34, while driving-behavior exposure decreases from 68 to 31. Cabin biometric processing produces the largest relative reduction, falling from an exposure score of 76 to 29, equivalent to a 61.8% reduction. This improvement is principally attributable to the modeled use of local processing, limited retention, and deletion practices that reduce unnecessary external transmission and secondary processing. Infotainment-contact exposure also declines substantially, from 61 to 27. Overall, the results indicate that

privacy-by-design principles can reduce exposure across multiple categories of sensitive vehicle-generated data rather than protecting only a single information type.

Table 3. Simulated Logistic Model Predicting Privacy Incident Risk


Predictor	Odds Ratio	95% CI	P-value	Interpretation
Precise Geolocation Processing	2.46	1.68–3.61	<.001	Higher modeled risk
Biometric Processing	2.18	1.47–3.24	<.001	Higher modeled risk
Third-Party Sharing	2.92	2.01–4.26	<.001	Strongest positive predictor
External Cloud Transfer	1.78	1.23–2.59	.002	Elevated exposure
Edge Processing	0.42	0.29–0.61	<.001	Protective association
Purpose Limitation	0.51	0.35–0.74	<.001	Protective association
Encryption + Restricted Access	0.38	0.25–0.57	<.001	Strong protective association

Table 3 presents the simulated logistic model examining factors associated with privacy-incident risk. Third-party sharing emerges as the strongest positive predictor, with an odds ratio of 2.92, indicating substantially greater modeled risk when vehicle-generated information is transferred to external recipients. Precise geolocation processing also demonstrates a strong positive association with privacy risk, with an odds ratio of 2.46, followed by biometric processing at 2.18 and external cloud transfer at 1.78. These findings suggest that privacy exposure increases when sensitive data are extensively collected, centrally transmitted, or made available to additional parties.

In contrast, the modeled protective factors have odds ratios below 1.00. Edge processing shows an odds ratio of 0.42, indicating a substantial reduction in modeled incident risk when sensitive data can be processed locally rather than continuously transmitted to external

infrastructure. Purpose limitation demonstrates a similarly protective association with an odds ratio of 0.51. Encryption combined with restricted access produces the strongest protective relationship, with an odds ratio of 0.38. The results therefore indicate that privacy governance should address both the amount and sensitivity of collected information and the architecture through which data are processed, transmitted, accessed, retained, and shared.

Graph 2. Simulated Reduction in Privacy Risk across Governance Controls



Graph 2 illustrates the cumulative reduction in modeled privacy risk as progressively stronger governance controls are introduced. Re-identification probability begins at 37.8% under the baseline architecture. Following data minimization, this probability decreases to 29.6%, demonstrating that reducing unnecessary data collection produces an immediate privacy benefit. The introduction of edge processing further lowers re-identification probability to 18.9%, indicating the protective value of keeping sensitive processing within the vehicle or local environment rather than routinely transmitting raw information to external servers.

When consent and purpose restrictions are added, the modeled re-identification probability falls further to 12.7%. Under the complete governance stack, incorporating data minimization, edge processing, consent and purpose limitation, restricted access, and related protections, the probability reaches 7.4%. This represents a substantial decline from the baseline condition. Unauthorized-disclosure probability similarly decreases from 24.5% under the baseline architecture to 4.9% under the full governance configuration. Because

intermediate unauthorized-disclosure values were not provided in the simulation data, the graph displays only the reported baseline and final values for this outcome.

The graph demonstrates that no individual privacy intervention provides comprehensive protection. Instead, progressively stronger reductions emerge when multiple safeguards operate together. Data minimization reduces unnecessary collection, edge processing limits external transfer, consent and purpose restrictions constrain secondary use, and access controls reduce the number of individuals or organizations capable of accessing sensitive information. The resulting pattern supports a defense-in-depth approach to privacy governance in which technical, architectural, and organizational controls operate collectively.

Key Finding: The simulation suggests that the strongest privacy protection emerges from a cumulative governance architecture rather than from reliance on a single technical control. Third-party sharing represents the strongest modeled risk factor, whereas encryption with restricted access and local edge processing demonstrate substantial protective associations.

Discussion / Critical Discussion

The simulation supports a fundamental distinction between security risk and privacy risk.

A connected-vehicle system could be fully encrypted and resistant to external intrusion while still exposing users to unnecessary surveillance if the manufacturer legitimately collects excessive information and distributes it across a broad commercial ecosystem.

This distinction is reflected in current regulatory concern. The FTC's 2026 final GM/OnStar order addressed allegations involving collection and disclosure of geolocation and driver-behavior information, demonstrating that connected-car privacy risk can arise from organizational data practices rather than only malicious hacking.

The simulation also identifies third-party sharing as a particularly important risk amplifier. Once connected-vehicle information is copied across organizations, technical and governance complexity increases substantially.

Research on automotive privacy policies similarly emphasizes the need for clearer understanding of how vehicle information is collected and processed.

Consumer research also indicates that individuals differentiate among types of automotive data collection rather than treating all vehicle information equivalently.

The strong modeled effect of edge processing is therefore significant. When AI can perform a safety or personalization function locally, transmitting raw information to external infrastructure should require separate justification.

The broader implication is that privacy cannot be repaired exclusively through longer notices. The more effective intervention is often architectural: collect less information, transmit less of it, retain it for less time, and expose fewer external entities to identifiable data.

Research Contribution / Implications

This study contributes a privacy-governance perspective specifically tailored to AI-enabled connected vehicles.

- For **vehicle manufacturers**, every sensor-to-cloud data path should have a documented purpose, sensitivity classification, retention period, and responsible owner. Privacy engineering should be integrated with cybersecurity engineering rather than treated exclusively as legal documentation.
- For **AI developers**, model inputs should be evaluated according to necessity. If accurate inference can be obtained using less precise location information or locally transformed features, the system should not automatically retain raw high-resolution traces.
- For **cloud and application providers**, contractual authorization should not substitute for technical minimization. External processors should receive only the information required for the specific service.
- For **insurers and mobility-data recipients**, driving telemetry should not silently migrate from one operational purpose into unrelated profiling. The FTC's recent connected-vehicle enforcement illustrates the potential regulatory significance of driving-behavior information when used or disclosed beyond consumers' expectations.
- For **regulators**, cybersecurity compliance should not be interpreted as comprehensive privacy compliance. UN Regulation No. 155 focuses on vehicle cybersecurity management, while ISO/SAE 21434 establishes cybersecurity engineering across the vehicle lifecycle; these mechanisms provide important security foundations but do not replace broader personal-data governance.

- For **vehicle owners and passengers**, privacy controls should be understandable at the feature level. Disabling personalized recommendations should not necessarily require disabling emergency functions or essential safety services.

Conclusion

AI-enabled connected vehicles can improve transportation safety, navigation, personalization, predictive maintenance, automated driving functions, traffic efficiency, and mobility services. Their increasing computational capabilities, however, mean that automobiles are also becoming sophisticated personal-data environments.

The privacy challenge extends beyond isolated data collection.

Multiple data streams can be combined to produce higher-order inferences concerning routines, identity, behavior, preferences, mobility patterns, and potentially sensitive activities. The greater the number of recipients and processing stages, the more difficult it becomes for individuals to understand or control the lifecycle of their information.

The simulation presented in this study demonstrates that privacy-by-design measures can substantially reduce modeled information exposure. Average privacy-exposure scores decreased by approximately 58% after architectural controls were introduced. Re-identification risk declined from 37.8% to 7.4%, while unauthorized-disclosure risk declined from 24.5% to 4.9% under the full governance scenario.

These findings do not represent real manufacturer performance. They demonstrate the potential value of layered controls.

A responsible connected-vehicle architecture should therefore apply the following logic:

collect less → infer carefully → process locally → disclose selectively → retain briefly → delete reliably → audit continuously.

Current regulatory and standards developments already support lifecycle thinking. ISO/SAE 21434 addresses road-vehicle cybersecurity engineering, UNECE Regulations 155 and 156 address cyber security and software-update management, and ISO/PAS 8800 addresses safety-related automotive AI.

Privacy governance must operate alongside these frameworks.

The central policy objective should not be to prevent vehicles from using data. It should be to ensure that connected mobility does not require individuals to surrender disproportionate informational privacy as the hidden price of vehicle functionality.

Limitations

The principal limitation is that the statistical analysis is based entirely on synthetic scenarios. No actual vehicle telemetry, manufacturer database, insurer information, mobility-service data, biometric information, or driver location history was analysed.

The privacy-exposure score is a methodological index rather than a standardized regulatory measurement.

The simulation simplifies the complex relationship among data sensitivity, cyber security, consent, contractual processing, legitimate operational purposes, and jurisdiction-specific legal requirements.

The analysis also does not model emergency services in detail. Certain high-value safety functions may legitimately require external data transmission even when local processing would otherwise provide stronger privacy.

Another limitation concerns anonymization. Real-world mobility datasets can exhibit complex re-identification properties that depend on sampling resolution, population density, data availability, and auxiliary information.

The research further treats edge processing as privacy protective, but on board systems can themselves be compromised. Locality reduces transmission exposure but does not eliminate security or privacy risk.

Regulatory requirements also differ internationally. The EDPB's connected-vehicle guidance is situated within European data-protection law, while U.S. consumer-protection enforcement and automotive cyber security frameworks operate under different legal structures.

Finally, vehicle ecosystems evolve rapidly through over-the-air updates, third-party integrations, new AI models, and changing communications infrastructure. Privacy assessments therefore require continuous reassessment.

Future Work

Future research should evaluate privacy architectures using appropriately consented or de-identified real-world telemetry under rigorous governance controls.

A priority area is privacy-preserving AI inference, particularly where valuable vehicle functions can be implemented without centralizing raw user information.

Recent automotive privacy research has examined secure multiparty computation, homomorphic encryption, and other privacy-preserving computational methods that may reduce disclosure during automotive data processing.

Federated learning also deserves continued investigation because it may permit collaborative model improvement while retaining more information locally, although gradient leakage and model-inversion risks require independent assessment.

Future studies should evaluate differential privacy for aggregate mobility analytics, confidential computing for automotive cloud services, privacy-preserving V2X communication, automated retention enforcement, and cryptographically verifiable deletion.

Research should also distinguish privacy requirements among drivers, passengers, vehicle owners, renters, children, fleet employees, and incidental road users captured by external cameras.

Vehicle resale and rental environments merit special attention because personal information may persist after the primary user no longer controls the vehicle.

Another important direction concerns AI-derived data. Existing governance frequently concentrates on raw information even though predictions about destinations, driving style, preferences, or routines may themselves become sensitive personal information.

Future longitudinal research should therefore track information from initial sensor acquisition through derived inference, third-party disclosure, model retraining, and eventual deletion.

Finally, independent automotive privacy audits could complement cybersecurity assessments by systematically testing data minimization, transparency, retention, user control, inference boundaries, and recipient accountability throughout the vehicle lifecycle.

Declarations

- **Conflict of Interest:** The author declares no conflict of interest. The author identity, institutional affiliation, email address, and ORCID shown in this manuscript are fictional sample metadata and must be replaced with verified author information prior to submission.
- **Funding:** No external funding was received for this simulation-based study.
- **Ethics Statement:** No human participants, actual drivers, identifiable passengers, vehicle telemetry, location traces, biometric information, manufacturer records, insurance files, or personal mobile-device data were used. Institutional human-subject ethics approval was therefore not applicable to the synthetic analysis.
- **Data Availability:** The statistical results derive from a synthetic dataset containing 1,600 modeled connected-vehicle processing scenarios. These values were developed solely for methodological illustration and must not be interpreted as measured privacy risks of any vehicle manufacturer, platform, insurer, or connected-car service.
- **Author Contribution:** The conceptual privacy framework, synthetic methodology, statistical interpretation, governance model, and manuscript structure were developed for the present simulation-based article.
- **Privacy and Security Statement:** The study distinguishes privacy from cybersecurity. Cybersecurity safeguards are treated as necessary but insufficient components of a broader governance framework that also addresses necessity, purpose, transparency, inference, retention, third-party disclosure, and individual control.
- **AI Transparency Statement:** AI-supported tools may assist scholarly drafting and analytical workflows. Responsibility for factual verification, statistical validity, citation accuracy, privacy-law interpretation, and publication ethics remains with human researchers and publishers.

References

1. European Data Protection Board. (2021). *Guidelines 01/2020 on processing personal data in the context of connected vehicles and mobility related applications*. 29/04/2021

- from https://www.edpb.europa.eu/documents/guideline/guidelines-012020-on-processing-personal-data-in-the-context-of-connected_en
2. 14. United States. Federal Trade Commission. (2024). *Cars & Consumer Data: On Unlawful Collection & Use*. <https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2024/05/cars-consumer-data-unlawful-collection-use>.
 3. Federal Trade Commission. (2026). *FTC Finalizes Order Settling Allegations that GM and OnStar Collected and Sold Geolocation Data Without Consumers' Informed Consent*. Off. Internet sites of the federal Trade Commission {accessed Date}. <https://www.ftc.gov/news-events/news/press-releases/2026/01/ftc-finalizes-order-settling-allegations-gm-onstar-collected-sold-geolocation-data-without-consumers>
 4. United Nations Economic Commission for Europe. (2021). *UN Regulation No. 155 Cyber Security and Cyber Security Management System*. accessed 18 October 2023, from <https://unece.org/transport/documents/2021/03/standards/un-regulation-no-155-cyber-security-and-cyber-security>
 5. United Nations Economic Commission for Europe. (2021). *UN Regulation No. 156 Software Update and Software Update Management System*. <https://unece.org/transport/documents/2021/03/standards/un-regulation-no-156-software-update-and-software-update>
 6. International Organization for Standardization. (2021). *ISO/SAE 21434:2021 Road vehicles Cybersecurity engineering*.
 7. International Standardization Organization. (2024). *ISO/PAS 8800:2024 Road vehicles Safety and artificial intelligence*. <https://www.iso.org/standard/83303.html>
 8. National Highway Traffic Safety Administration. (2022). *Cybersecurity Best Practices for the Safety of Modern Vehicles*. Retrieved from https://www.nhtsa.gov/sites/nhtsa.gov/files/2022-09/cybersecurity-best-practices-safety-modern-vehicles-2022-pre-final-tag_0_0.pdf
 9. Dowthwaite, A., Cook, D., & Cox, A. L. (2024). Privacy preferences in automotive data collection. *Transportation Research Interdisciplinary Perspectives*, 24, 101022. <https://doi.org/10.1016/j.trip.2024.101022>
 10. Bodei, C., Costantino, G., De Vincenzi, M., Matteucci, I., & Monreale, A. (2023). Vehicle data collection: an analysis and comparison of privacy policies. In *Proceedings of*

- the 9th International Conference on Information Systems Security and Privacy* (pp. 626–633). <https://doi.org/10.5220/0011779500003405>
11. Quinlan, M., Zhao, J., & Simpson, A. (2019). Connected vehicles: A privacy analysis. In *Security, Privacy, and Anonymity in Computation, Communication, and Storage* (pp. 35–44). Springer. https://doi.org/10.1007/978-3-030-24900-7_4
 12. Pesé, M. D., & Shin, K. G. (2019). A survey of automotive privacy regulations and privacy related attacks. *SAE Technical Paper 2019-01-0479*. <https://doi.org/10.4271/2019-01-0479>
 13. Costantino, G., De Vincenzi, M., Martinelli, F., & Matteucci, I. (2023). A privacy-preserving solution for intelligent transportation systems: Private Driver DNA. *IEEE Transactions on Intelligent Transportation Systems*, 24(1), 258–273.
 14. Alam, T. (2024). Data privacy and security in autonomous connected vehicles in smart city environments. *Big Data and Cognitive Computing*, 8 (9), 95. Retrieved from <https://doi.org/10.3390/bdcc8090095>
 15. 2. Xie, C., Cao, Z., Long, Y., Yang, D., Zhao, D., & Li, B. (2022). *Privacy of autonomous cars: Threats, protection techniques and future perspectives*. arXiv. <https://doi.org/10.48550/arXiv.2209.04022>
 16. De Vincenzi, M., Pesé, M. D., Bodei, C., Matteucci, I., Brooks, R. R., Hasan, M., Saracino, A., Hamad, M., & Steinhorst, S. (2024). *Contextualizing security and privacy of software-defined vehicles: State of the art and industry perspectives*. arXiv. <https://doi.org/10.48550/arXiv.2411.10612>
 17. Yuca, N., Matyunin, N., Arzoglou, E., Anagnostopoulos, N. A., & Katzenbeisser, S. (2025). A survey on privacy-preserving computing in the automotive domain. *ACM Computing Surveys*. <https://doi.org/10.1145/3770580>
 18. Huang, Y., Qian, L., et al. (2025). Consumer preferences and willingness to pay for data privacy in automated vehicles. *Transportation Research Part A: Policy and Practice*, 199, 104585. <https://doi.org/10.1016/j.tra.2025.104585>