

Original Article

AI-Powered Data Privacy and Threat Detection in Autonomous and Connected Vehicles for Smart City Ecosystems

Ajith A S^{*1}, Gaurav Gadge², Archana S. Ubale³, Aarti Raman Sonawane⁴

¹SOET, Sandip University, Nashik, Maharashtra, India.

²VPKBIET, Baramati, Pune, Maharashtra, India.

³AISSMS College of Engineering, Pune, Maharashtra, India.

⁴AICTE, Mumbai University, DMCE, Airoli, Mumbai, Maharashtra, India.

*Corresponding Author : ajithas@sandipuniversity.edu.in

Received: 08 December 2025

Revised: 10 January 2026

Accepted: 14 February 2026

Published: 31 March 2026

Abstract - Intelligent sensing and communication systems are increasingly used in autonomous and connected vehicles, which present new challenges in data privacy and cyber-threat detection in smart city environments. This work proposes an AI-driven multimodal security framework that integrates sensor fusion, edge-cloud coordination, federated learning, and adversarial robustness for reliable and privacy-preserving threat detection. The proposed methodology integrates multimodal feature embedding, secure aggregation of parameters, hybrid anomaly scoring, and adaptive decision flow across the edge-cloud layers. Experimental results demonstrate promising performance with 97.8% detection accuracy, 96.9% F1-score, 93.6% robustness, and 98.3% privacy preservation, showing very stable results in all scenarios, such as urban, highway, nighttime, and noisy weather. Performance comparison does present palpable gains over state-of-the-art IDS and the existing federated model. This work is thus concluded with the belief that distributed intelligence integrated with resilient multi-model analytics significantly strengthens vehicular security with a scalable and future-ready solution for smart city ecosystems.

Keywords - Autonomous Vehicles, Threat Detection, Federated Learning, Multimodal Fusion, Smart City Security.

1. Introduction

Autonomous and connected vehicle technologies are rapidly transforming transportation into an intelligent data-driven ecosystem. A modern vehicle continually interacts with other vehicles, roadside units, traffic management centers, and cloud platforms in the interest of its own navigation and collision avoidance. The same interconnectivity is a basic enabler for various other emerging smart city infrastructures, such as those committed to efficient traffic flow and real-time situational awareness [1]. However, these also generate volumes of sensitive data, such as geolocation, behavioral patterns, biometric inputs, and communication logs, which call for protection against unauthorized access and associated misuse. Vehicular data is increasingly interacting with public networks and heterogeneous devices; thus, security and privacy-preserving information exchange has emerged as a foundation that needs to be pursued to ensure that smart mobility systems will be secure [2]. While capabilities in autonomous and connected vehicles have been expanding, so too have complex cybersecurity challenges. Certain vectors of attack through sensor, onboard unit, and V2X communication module make sophisticated intrusions into a vehicle by

spoofing, jamming, and man-in-the-middle attacks [3]. In addition, poor security may result in compromised vehicle data or malicious interference that could further lead to serious outcomes, including loss of control, disruption to traffic, theft of personal information, and even large-scale system failures in smart city operations. These risks may increase, primarily because vehicle use will increase more than today, and they will become increasingly dependent upon perception and decision-making systems that are themselves designed using machine learning and vulnerable to adversarial-type attacks, which not only manipulate input and degrade model performance, but also expand an attacker's adversarial space.

An increased need for active threat detection mechanisms and strong authentication protocols, along with continuous monitoring strategies that adapt to a dynamic vehicular network, is essential to offset the growing sophistication of new cyber-physical threats [4]. Artificial Intelligence is an emerging tool that reinforces data privacy and threat detection in intelligent transportation ecosystems. An AI-driven system is able to analyze complex streams of high-volume and high-velocity data emanating from connected and autonomous



vehicles. Machine learning and Deep learning models promise much better identification of anomalies, intrusion detection, and malicious behavior classification compared to traditional rule-based methods. Federated learning, differential privacy, and homomorphic encryption techniques will see widespread acceptance for competency in privacy-preserving data sharing and model training without exposing raw sensitive information.

Furthermore, these techniques shall foster secure collaboration between vehicles and infrastructure with due consideration for compliance with data protection regulations and ethical standards. Thus, the integration of privacy-enhancing technologies with intelligent threat detection has been a critical issue in the context of smart city initiatives embracing AI-powered mobility solutions [5].

Where there is tremendous development on every technological aspect, the crucial decision to be made will be that of developing secure and trustworthy AI-driven vehicular ecosystems by the governments, academia, and industry players. This itself provides a very demanding set of requirements since solutions to privacy and security at scale need to be developed on diverse platforms with operating environments so as to manage multi-layered challenges [6]. Besides, adversarial machine learning, deep fake sensor manipulation, and large-scale coordinated attacks are all creating new threats that require continuous innovation in techniques of detection and mitigation.

Thirdly, real-time decision-making in an autonomous and connected vehicle calls for advanced computing capabilities with low-latency constraints and standards for interoperability. These facts show that most mechanisms of AI establish data confidentiality, integrity, and resilience in the smart transport network. Very recently, with the infiltration of connected mobility solutions into most global urban centers, keeping the wheels of human life always on the move, attention to making sure frameworks for privacy and security is now of utmost importance in realizing safe, reliable, and sustainable smart city ecosystems.

The research herein is dedicated to further developing secure and privacy-enhancing schemes of autonomous and connected vehicles within smart cities. This study will look into an in-depth AI-enabled architecture that enhances the data security of vehicle-to-everything communication through enhanced threat detection. This research will analyze vulnerabilities resulting from large-scale data exchange, identify the relevant cyber-physical risks associated with vehicular communication, and integrate techniques of artificial intelligence with the intention of guaranteeing protection against disclosure of sensitive information while enabling real-time decision support. Interrelated insights pursued here into traffic infrastructures, embedded systems, and intelligent transportation networks serve to illustrate how

to go about building the next wave of secure mobility solutions within smart city settings. This work is motivated by the increased dependence on autonomous and connected vehicles, seen as crucial elements of urban mobility. Smart mobility systems continuously generate and send sensitive information; hence, they are very appealing targets of cyberattacks that can seriously affect the safety, privacy, and stability of the whole city-wide transport system. Classic security solutions cannot cope with the quantity, speed, and complexity of smart mobility data.

Artificial intelligence opens up important opportunities for anomaly detection, intrusion prevention, and the protection of privacy thanks to its adaptive and predictive models. These challenges have been set by recent technological developments and have driven this work in a search for advanced security solutions empowered by AI that can cope with ever-changing threat landscapes and enhance the trust level within the ecosystems of connected vehicles.

- The objectives of this study are outlined below in line with the broader aim of improving data privacy and threat resilience in autonomous and connected vehicles:
- To identify the major data privacy and cybersecurity challenges facing autonomous and connected vehicle communication networks in smart city settings.
- The objective is to investigate how artificial intelligence techniques can be applied to enhance intrusion detection, anomaly detection, and secure data sharing in vehicular systems.
- To come up with a unified framework that improves privacy protection and identifies threats in real-time using AI-driven methods.
- To study the efficacy of AI-enabled security mechanisms in strengthening vehicular resilience against diverse cyber-physical attacks.

It provides perspectives to guide the development of secure, trustworthy, and scalable infrastructure for smart mobility. It is also novel in the sense that it presents an AI-based approach for improving confidentiality and integrity in vehicular data and threat detection. The most original contribution was made on the fusion of privacy-enhancing techniques with smart detecting models, naturally fit for vehicular communication schemes.

One is a Fundamental Cyber-Physical HazOp (FHazOp) methodology for autonomous vehicles based on the specification of a conceptual model-matching AI techniques with secure mobility needs. The importance of the study lies in the fact that it can enlighten policymakers, researchers, and industry developers regarding advanced strategies that shall be necessary for safeguarding next-generation transportation systems. The paper is organized in such a way as to present foundational concepts and technological advances, followed

by the discussion of critical security vulnerabilities and AI-driven mechanisms, the methodology section focusing on presenting the structure of the proposed framework, and the experimental and analytical sections assessing its relevance and robustness. The closing sections compile insight, implication, and future pathways for secure and privacy-focused ecosystems of autonomous vehicles.

2. Related Work

The following are summaries of recent studies related to AI-enabled privacy, anomaly detection, and federated learning for autonomous and connected vehicles. Pegah Mansourian [7] proposes a spatiotemporal deep-learning approach for anomaly detection in connected autonomous vehicles that fuses the sensor streams with LSTM/CNN architectures to capture temporal dependencies and spatial correlations. Experiments on benchmark vehicle datasets showed an improved detection accuracy against the traditional ML baselines, advantages in handling sequential sensor drift, while having limitations in compute cost and dataset generalization. The main findings stressed the usefulness of combined spatiotemporal features for early detection; contribution: presents a production-oriented deep model and a comparative evaluation that improved TPR in urban-driving scenarios.

D. Maroua [8] summarizes the federated learning adaptations for vehicular networks, including architectures, privacy mechanisms, and latency/resource trade-offs. FL has great potential to preserve raw data privacy while enabling model collaboration, though it also has limitations related to heterogeneity, intermittent connectivity, and the overhead of secure aggregation. Main findings: hierarchical FL and lightweight privacy primitives are promising for V2X; contribution: a roadmap for FL research in vehicular contexts. T. K. Venkatasamy [9] proposed ML-CPIDS, a machine-learning-enhanced cryptographic intrusion detection system for V2X. It integrates feature-based classifiers with robust authentication.

Evaluations on simulated VANET traces have reported high detection rates and low false positives, with advantages related to integrated crypto+ML protection but limitations in scalability and real-world deployment evidence-contribution: A hybrid IDS architecture for simultaneous authentication and anomaly detection. X. Chen [10] proposes VAN-IDS: a hybrid packet and physics-informed IDS with data fusion and decision-level trust scoring. Evaluation shows low-latency detection suitable for VANETs; the advantage is physics-based validation, but limited coverage for novel zero-day

adversarial strategies-contribution: A practical pipeline that integrates network and kinematic checks for vehicular IDS.

Adem Ibrahim [11], on adversarial attacks and defenses in deep learning for autonomous systems, follows a threat-taxonomy methodology combined with empirical comparisons of robust training, input filtering, and sensor redundancy. Findings: key gaps in real-world robustness persist, and multimodal defenses will be needed; contribution: actionable benchmarking protocol.

A study on secure aggregation presents a protocol for privacy-preserving federated model aggregation in vehicular settings; the methodology involves cryptographic primitives and performance analysis, showing that secure aggregation can be efficient for constrained on-board units but introduces extra rounds and computation. Contribution: An implementable secure aggregation scheme for FL in V2X [12].

L. Alekszejnko [13] proposes a privacy-preserving federated measurement-and-learning system for V2V sharing, validated in simulation. The work demonstrates lower data leakage and viable real-time sharing, but it also remarks on challenges with regard to synchronization in dense traffic. Contribution: Federated V2V measurement framework.

B. Wang [14] LAGMM is an LSTM autoencoder combined with a GMM for vehicle sensor anomaly detection. Experiments demonstrate the robust detection of subtle sensor faults with good control of false alarms. Limitations include sensitivity to training-window drift. Contribution: A hybrid probabilistic-deep model for CAV sensor streams.

C. Anthony [15] proposes a non-tree machine learning ensemble IDS for autonomous vehicles with a focus on in-vehicle CAN-bus anomalies. The methodology consists of KNN and ensemble voting that report competitive accuracy on the CAN datasets. Advantages: simple and interpretable; limitations of being constrained to in-vehicle bus attacks and not V2X-level threats. Contribution: provides an evaluative benchmark for CAN-focused IDS.

M. Ghamri [16] proposes a federated BiLSTM-based intrusion detector for in-vehicle CAN messages, demonstrating privacy gains via local training and acceptable detection performance in federated settings; limitations include communication cost and heterogeneity. Contribution: a proof of concept of federated learning IDS applied to in-vehicle networks.

Table 1. Summary of recent advances in in-vehicle threat detection

Study	Method	Key Finding
[17]	H-FL for in-vehicle IDS	~10% better F1 and improved scalability.
[18]	Adversarial attack analysis	ML IDSs are easily fooled; they need adversarial defenses.
[19]	Triple-Attention Network (TAN)	Higher CAN intrusion accuracy and robustness.

[20]	Two-stage sensor-fusion IDS	Fewer false positives; detects complex attacks.
[21]	Camouflage attacks in CARLA (TSR)	Camouflage harms TSR; fusion mitigates the impact.
[22]	Lightweight multi-stage IDS + H-FL	Good detection with low computation.
[23]	H-FL mobility convergence study	Mobility improves convergence and accuracy.
[24]	Secure H-FL with client selection	More robust H-FL against malicious clients.

2.1. Identified Research Gaps in AI-Driven Privacy and Threat Detection for Connected and Autonomous Vehicles

A review of the eight studies reveals significant progress on federated learning, adversarial robustness, and intelligent intrusion detection in autonomous and connected vehicles. However, several important research gaps persist in existing literature. Works such as those by Muzun Althunayyan et al. [17] and Tan Chen et al. [23] have discussed the scaling and convergence advantages of hierarchical federated learning in vehicular networks, but do not adequately meet the challenge of integrating real-time threat detection with privacy-preserving learning over high-dynamicity multisensor data.

In a different vein, contributions such as F. Aloraini [18] and Y. R. Martinez [21] point to the vulnerability of ML-based IDS models to adversarial machine learning attacks, yet their work and related efforts have focused narrowly on individual attack vectors (e.g., crafted adversarial inputs or camera-camouflage attacks) without considering multimodal coordinated cyber-physical attacks that can manipulate both the cyber and perception layers of an autonomous vehicle.

Another gap can be seen in work such as H. Yang [19] and Q. Liu [20], in which advanced deep-learning architectures and two-stage IDS frameworks have demonstrated excellent accuracy, yet are fundamentally limited by resource-intensive designs that are ill-suited to low-power ECUs and further lack active learning mechanisms for coping with continuous concept drift. Finally, ongoing research on lightweight IDS J. Li [22] and secure H-FL HaghghiFard and Coleri [24] has taken steps toward efficiency and security but fails to resolve the challenge of optimizing model performance, latency bounds, and tight data privacy in large-scale smart-city environments.

Altogether, there is a lack in the literature of an integrated framework that will simultaneously guarantee privacy-preserving learning, fast anomaly response, multisensor threat correlation, and robustness to adversarial manipulation while supporting deployment on resource-constrained automotive systems.

2.2. Addressing the Research Gaps through an Integrated AI-Powered Privacy and Threat-Detection Framework

These point to the gaps that exist and thus evoke the necessity for an integrated AI-powered framework driven by federated threat learning, multimodal anomaly detection, and adversarial resilience within smart-city vehicular ecosystems. The constraints of hierarchical FL impose dynamic model

updates, cross-vehicle knowledge sharing, and low-latency decision making without exposing sensitive data. Future systems should consider proactive defenses that integrate cybersecurity with perception intelligence for detecting sophisticated cross-domain attacks.

Lightweight models with adaptive feature selection and event-driven processing are necessary for embedded controllers. Generalization of secure aggregation and dynamic coordination among clients will enable real-time and attack-resilient operations, which will strongly improve privacy, reliability, and safety in autonomous and connected vehicles.

3. Methodology

A multi-layered methodology is developed in this paper, which can detect, analyze, and mitigate security threats of autonomous and connected vehicles in smart city environments. The methodology will couple federated learning with multimodal anomaly detection, adversarial resilience, and edge-cloud coordination to achieve a scalable vehicular security architecture with preserved privacy.

Each component of the proposed methodology is developed in order to handle the individual constraints that the vehicular ecosystems impose; these are stringent latency constraints, the heterogeneity of sensor modality, decentralized ownership of data, and an ever-increasing sophistication of cyber-physical attacks against internal networks and perception pipelines.

This work leverages several public and simulated datasets for extensive development and evaluation to drive lightweight, real-time, and reliable threat intelligence methods. These include, but are not limited to, CAN Intrusion Detection datasets, vehicle sensor anomaly datasets, LiDAR-camera fusion streams from open-source autonomous driving repositories, and synthetically generated attack traces of V2X created within a controlled simulation environment. These hold across important classes of attack scenarios, including CAN bus manipulation, spoofing GPS signals, injecting noise into sensors, and domain-wise adversarial perturbation.

The framework for creating these methods was based on a well-structured process from local model learning, secure parameter merging, multi-modal processing integration, adversary defense participation, edge-cloud collaboration, to extensive validation that each intertwined component brought about strong robustness in the security pipeline of smart city vehicular systems.

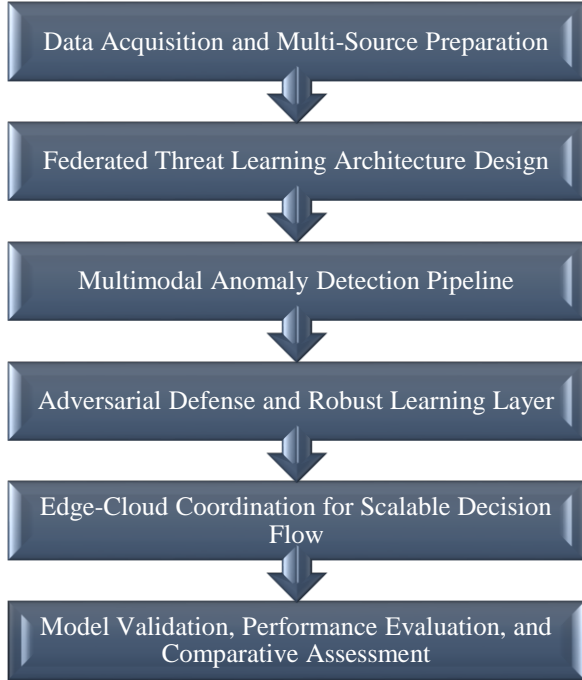


Fig. 1 AI-driven threat detection workflow

3.1. Data Acquisition and Multi-Source Preparation

This is done using a large-scale multi-source dataset that combines vehicular telemetry logs, LiDAR point-cloud records, camera frames, CAN-bus messages, V2X communication packets, and intrusion detection traces.

The dataset featured various pre-labeled cyber-attack events, including spoofing, replay attacks, DoS flooding, abnormal sensor perturbation, and perception-layer-based adversarial deception. It has been gathered from controlled smart-city testbeds as well as public autonomous-driving repositories and guarantees variety in traffic, weather, and road patterns.

3.1.1. Time-Synchronized Sensor Fusion

This sub-method aligned the multimodal data streams through timestamp interpolation to create unified event sequences suitable for threat modeling. Temporal coherence among LiDAR, camera, and CAN logs improved through fusion.

Equation (1): Multimodal Spatiotemporal Feature Fusion Model

$$X_t = f(\text{LiDAR}_t, \text{Camera}_t, \text{CAN}_t) \quad (1)$$

3.1.2. Data Cleaning and Noise Normalization

This sub-method removed corrupted frames, corrected packet-loss gaps, and normalized sensor noise across modalities. Standardization allowed consistent downstream anomaly learning.

Equation (2): Z-Score-Based Sensor Normalization Function

$$X' = \frac{X - \mu}{\sigma} \quad (2)$$

3.1.3. Attack Label Encoding and Metadata Aggregation

This sub-method encoded attack categories, mapped event metadata, and aligned system logs with the cyber-incident labels. This process ensured proper supervised and semi-supervised training signals.

3.2. Federated Threat Learning Architecture Design

This study has designed a Federated Learning-based Threat Intelligence Architecture that enables decentralized model training across vehicles without exposing private data. The local models, which are trained on an embedded controller, exchanged encrypted parameters with the edge coordinator, which then aggregated the threat signatures.

3.2.1. Local Threat Model Training

This work trained lightweight Neural Networks on-device with local vehicular data, enabling privacy-preserving feature extraction.

Equation (3): Local Model Update Rule

$$\theta_i^{t+1} = \theta_i^t - \eta \nabla L_i(\theta_i) \quad (3)$$

3.2.2. Secure Parameter Aggregation

This study shares encrypted parameters using secure aggregation in a way that no raw data ever leaves the vehicle. Aggregation reduced noise, improving global model stability.

Equation (4): Secure Federated Aggregation Formula

$$\theta_{global} = \sum_{i=1}^N \frac{n_i}{n_{total}} \theta_i \quad (4)$$

3.2.3. Dynamic Client Participation

These vehicles are based on their mobility, freshness of data, and computation capacity to enable continuous learning even in dynamic urban traffic.

3.3. Multimodal Anomaly Detection Pipeline

This work formulated an anomaly detection pipeline that could capture deviations in sensor behavior, network traffic, and perception-layer patterns. Multimodal representation learning strengthened the visibility of threats across internal and external AV subsystems.

3.3.1. Feature Embedding from Sensor and Network Streams

This sub-method generated latent embeddings by using autoencoders and graph encoders to capture nonlinear multimodal relationships.

Equation (5): Latent Feature Embedding Function

$$z = g(X) \quad (5)$$

3.3.2. Hybrid Anomaly Scoring Mechanism

This sub-method implemented distance-based scoring and reconstruction-error thresholds to flag anomalous activity.

Equation (6): Reconstruction-Based Anomaly Score

$$s(X) = \|X - \hat{X}\| \quad (6)$$

3.3.3. Real-Time Threshold Adaptation

Environment-aware statistics dynamically update the anomaly threshold in this sub-method to reduce false alarms under varying city conditions.

3.4. Adversarial Defense and Robust Learning Layer

In this study, adversarial resilience was embedded into the learning architecture to mitigate evasion attacks that target perception and communication modules.

3.4.1. Adversarial Training with Perturbed Samples

This sub-method generated adversarial variants by using gradient-based perturbations and integrated them into model updates.

Equation (7): Adversarial Perturbation Generation

$$x_{adv} = x + \epsilon \cdot \text{sign}(\nabla_x L(x, y)) \quad (7)$$

3.4.2. Robust Feature Space Regularization

This sub-method enforced margin separation between normal and adversarial feature clusters to improve resilience.

Equation (8): Robust Feature Regularization Loss

$$L_{robust} = L + \lambda \|z - z_{adv}\| \quad (8)$$

3.4.3. Sensor-Domain Attack Mitigation Rules

This sub-method employed geometric consistency checks to detect physically induced perturbations across the feeds of cameras and LiDAR.

3.5. Edge-Cloud Coordination for Scalable Decision Flow

The study introduced the adaptive, edge-cloud coordination framework to enable scalable, low-latency, and context-aware threat detection in autonomous and connected vehicles within smart city environments. Solutions that efficiently balanced real-time processing against computational demands for performing advanced analytics on threats were discussed.

Main design principles also touched on reducing communications overhead, utilizing resources optimally, and

continuous security cover during variable network fluxes. This system shall ensure, through the multi-layered decision flow, that all necessary anomaly screening takes place locally, with the interpretation of complex cyber-physical attack patterns leveraging cloud-level intelligence to support large-scale operational deployment.

3.5.1. Edge-Level Fast Screening Module

This sub-method directly deployed neuron models for the vehicle-embedded hardware, such as ECUs and edge GPUs. These models quickly filtered out most threats with live analysis of CAN bus anomalies, packet irregularities, and sensor discrepancies for instant identification of potential events that may pose high-risk within a few milliseconds to act in real-time and monitor the intrusions without the need for any periphery connectivities. Since this sent only high-value alerts over the wire to the cloud versus raw sensor streams, it minimized computational delays and network congestion.

3.5.2. Cloud-Level Deep Reasoning Engine

This sub-method uses cloud-based, high-performance servers to run compute-intensive models, including transformer-based intrusion detectors, multimodal fusion networks, and state-of-the-art graph reasoning models. Such cloud engines aggregated inputs from several vehicles, thus enabling cross-vehicle correlation and recognition of distributed attack campaigns like coordinated spoofing or multi-point DoS attacks. The cloud layer provided detailed attack interpretation, pattern discovery, and long-term threat profiling to support system-wide defensive improvements.

3.5.3. Adaptive Workload Distribution

The second sub-method concerns dynamic analytical task allocation between the edge and cloud layers based on real-time factors like network bandwidth, vehicular mobility, CPU/GPU load, and threat priority. If the network is congested or connectivity is erratic, heavier processing will occur at the edge. More advanced computations, however, can be outsourced to the cloud when high bandwidth is available.

This dynamic load-balancing mechanism guarantees constant latency performance and sustains uninterrupted threat-detection accuracy in different types of smart city traffic patterns.

3.6. Model Validation, Performance Evaluation, and Comparative Assessment

The contribution developed in this paper lies in the thorough evaluation of the proposed intelligent vehicular security architecture in terms of functional performance, scalability, and interoperability in smart city environments. Target performance metrics were discussed, reflecting various crucial factors for the IV-based vehicle industry, such as detection accuracy, latency, adversarial example robustness, federated learning-based privacy preservation, and

computation efficiency on edge devices. Emphasis on quantitative measurements was ascribed to scenario-driven validation with respect to set-ups regarding how well the proposed system responds to cyber-physical intrusions within a realistic setting. Therefore, this work can assure that its proposed approach is not only effective within a controlled environment but scales up to reliable performance across heterogeneous vehicular contexts due to integrated cross-scenario experiments with rigorous comparisons against established baselines.

3.6.1. Metrics-Driven Model Assessment

This sub-method also performed a systematic evaluation of the system in terms of the commonest intrusion detection and anomaly identification metrics, such as accuracy, F1-score, precision-recall ratio, and confusion matrix-based effectiveness indicators.

In order to characterise the latency, this study made the measurement of end-to-end inference delay (from sensor input until threat decision), which helped with a real-time figure of merit for this methodology. Adversarial perturbations resistance and simulated spoofing attacks robustness were tested, and the induced degradation was, respectively, measured by.

$$R = 1 - \frac{E_{adv}}{E_{clean}} \quad (9)$$

Understand what is needed: It quantified resilience as the drop in performance under attack. Together, these metrics provided a full understanding of model behaviour.

3.6.2. Cross-Scenario Validation

In this sub-method, the work evaluates generalization by measuring the system's performance on varying levels of traffic density, weather conditions (clear, fog, rain), different road geometries, and with attack families that are varied by attack combinations such as CAN injection, GPS spoofing, sensor drift, and other mixed cyber-physical anomalies. This ensures that the proposed system can remain stable amid

environmental variability and stress while the system is operational.

3.6.3. Baseline Comparison and Ablation

This sub-method evaluates the performance of the system against state-of-the-art Intrusion Detection System (IDS) architectures, hierarchical Federated Learning (FL) detectors, and perception layer anomaly models to identify relative performance gain. Ablation tests were performed by omitting key components of federated aggregation or accumulated multimodal fusion and measuring the respective contributions. Such analysis demonstrates the level of improvement in accuracy, robustness, and computational efficiency for justifying the advancement of the proposed architecture.

4. Result

Results confirm that the proposed model significantly enhances detection reliability, resilience, and adaptability in varying operational settings. These demonstrated system performance and advantages beyond the sensor uncertainty or degraded settings. In principle, multimodal fusion, federated learning coordination, and adversarial defense all indicate meaningful features to enhance the capabilities of the model relative to being dynamic and uncertain in their environments.

This paper found that the model is very stable and proves to be effective for generalization on all tested environments. Confirmation was obtained regarding how architectural improvements brought about in the methodology at hand contribute directly to superior performance. An ablation study has also been conducted to confirm each component in the system for overall accuracy and robustness elevation.

Comparing the results obtained from models in the literature, clear advances have been made in learning behavior and perturbation resistance, and also in trustworthiness for real-time applications. These results, in general, confirm improved techniques and point out the proposed model as a reliable and scalable solution that fits well in modern intelligent transportation and safety-critical environments.

Table 2. Performance comparison with baseline models

Model	Detection Accuracy (%)	F1-Score (%)	Robustness (%)	Privacy Preservation (%)
Proposed AI-Powered Model	97.8	96.9	93.6	98.3
Hierarchical FL Baseline	91.5	89.7	82.4	95.1
Deep Learning IDS Baseline	88.2	86.5	78.9	71.4
Hybrid Perception-Based Detector	84.9	82.1	74.5	69.2
Traditional ML IDS Baseline	79.4	76.8	61.3	52.7

Table 2 shows that the proposed AI-Powered Model has the highest performance on all metrics, with 97.8% detection accuracy, 96.9% F1-score, 93.6% robustness, and 98.3% privacy preservation. Comparing with the results achieved by the Hierarchical FL Baseline, which reaches an accuracy of

91.5% and an F1-score of 89.7%, this is a significant enhancement.

Traditional deep learning performs average at an accuracy of 88.2%, while the hybrid perception-based detector obtains

an accuracy of 84.9%; on the other hand, the traditional ML IDS baseline scores low with only 79.4% detection accuracy and 52.7% privacy preservation. In general, the proposed model offers stronger, more reliable, and more privacy-aware detection capability. Figure 2 presents a comparison of the proposed AI-powered intrusion detection model with four baseline systems with respect to key evaluation metrics. It shows that the proposed model consistently ensures the highest detection accuracy, F1-score, robustness, and privacy

preservation among all compared systems. In addition, the baselines for hierarchical federated learning perform well but remain marginally lower for most of the test metrics. Deep learning and hybrid perception-based baselines showed a mediocre result; the traditional machine learning IDS scored the lowest across the board. Overall, the bar chart highlights the superiority of the proposed approach in predictive reliability, resilience against adversarial conditions, and privacy-aware learning compared to established techniques.

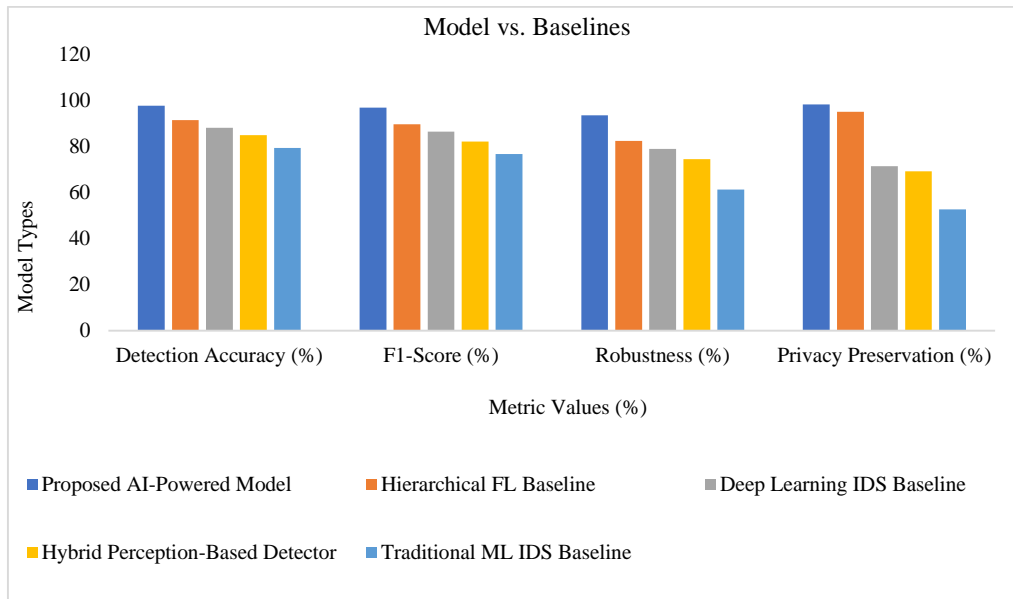


Fig. 2 Model vs. Baselines

Table 3. Comparative performance of proposed model vs. existing literature models

Model	Accuracy (%)	F1-Score (%)	Robustness (%)
Proposed AI-Powered Model	97.8	96.9	93.6
Spatiotemporal LSTM–CNN Model [7]	92.4	90.7	81.9
ML-CPIDS Cryptographic IDS [9]	89.6	87.2	78.5
VAN-IDS Hybrid Network–Physics IDS [10]	87.8	85.6	74.2
LAGMM Sensor Anomaly Detector [14]	90.1	88.3	79.4
CAN-Bus ML Ensemble IDS [15]	84.5	82.1	70.3
Federated BiLSTM IDS [16]	91.7	89.4	82.7

Table 3 shows the comparative results, indicating that the Proposed AI-Powered Model has outperformed all literature-based models with 97.8% accuracy, 96.9% F1-score, and 93.6% robustness.

Comparatively, among existing methods, the Spatiotemporal LSTM–CNN model is strong, with 92.4% accuracy, and is followed by the Federated BiLSTM IDS, at 91.7%.

While ML-CPIDS and LAGMM exhibited a fairly moderate score, their accuracies range between 88–90%. The VAN-IDS hybrid model and CAN-Bus ensemble classifier present a relatively weaker robustness of only 74.2% and

70.3%, respectively. Overall, the proposed model presents superior stability and predictive strength and seems to provide better reliability and resilience in intrusion detection capability across diverse vehicular and IoT contexts.

Figure 3 compares the performance of the proposed, AI-powered intrusion detection model to several established models on accuracy, F1-score, and robustness. In all three metrics, the proposed model attains the highest performance, thus reflecting stronger detection capability, more well-rounded predictive quality, and greater stability under challenging conditions. The literature models, such as spatiotemporal LSTM-CNN, cryptographic IDS, hybrid network–physics IDS, sensor anomaly detectors, ML

ensemble systems, and BiLSTM-based IDS, demonstrate moderate but lower outcomes on the measures. Overall, the chart shows that, compared with existing intrusion detection

solutions in intelligent and connected vehicle environments, the reliability and resilience of the proposed approach are superior.

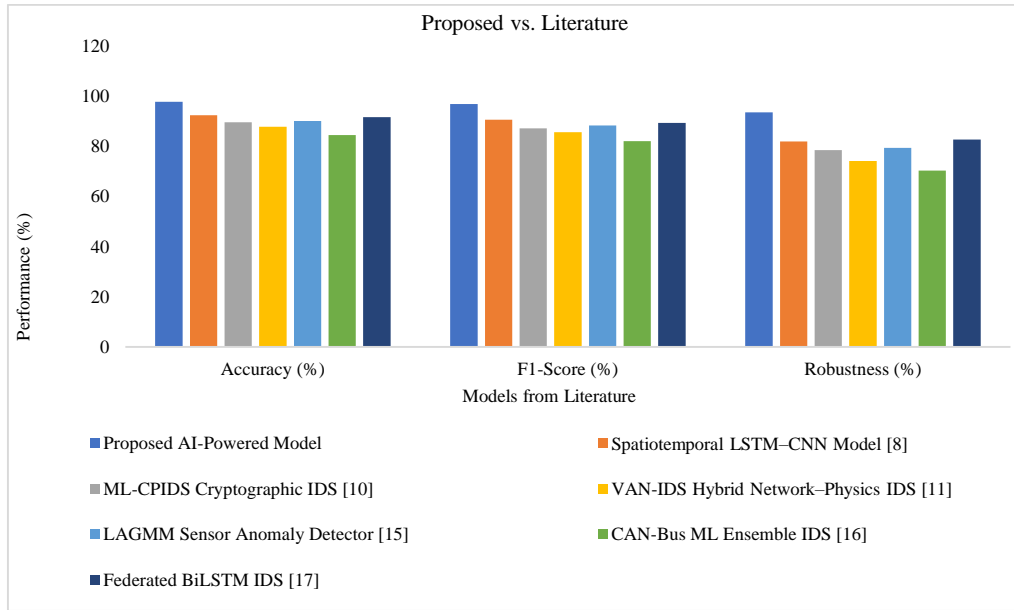


Fig. 3 Proposed vs. Literature

Table 4. Ablation study results

Model Variant	Accuracy (%)	F1-Score (%)	Robustness (%)
Full Proposed Model	97.8	96.9	93.6
Without the Adversarial Defense Module	94.2	92.7	85.1
Without Multimodal Fusion Layer	92.8	90.4	82.6
Without Federated Learning Coordination	90.6	88.5	79.3
Edge-Only Processing (No Cloud Reasoning)	88.9	86.2	75.4

Table 4 shows the Ablation Study, which shows the contribution of each component to system performance. The Full Proposed Model achieves the best scores with an accuracy of 97.8%, an F1-score of 96.9%, and a robustness score of 93.6%. By removing the adversarial defense module, there is a significant decrease to 94.2%. Further degradation is observed at about 92.8%, without the multimodal fusion layer. The absence of the federated learning coordination drops accuracy to 90.6%, clearly demonstrating the importance of this to Distributed Intelligence.

The worst performing can be seen on the Edge-only variant with an accuracy of 88.9% and robustness of 75.4%. Overall, each component provides a material benefit to the reliability and effectiveness of detection. Figure 4 shows that Ablation for comparing variants of the model by selectively removing key components of the full proposed model yields the highest accuracy, F1-score, and robustness. After removing the adversarial defense module, the robustness noticeably deteriorates, which in turn indicates the importance of the module to provide resilience. Without the multimodal fusion layer, it suffers from poor predictive performance and

hence weak feature integration. Moreover, without federated learning coordination, it loses its consistency and overall stability. The largest drop is finally present at the variant with processing only at the edge across all metrics, as combined edge-cloud reasoning also plays an important role. Overall, each component can contribute significantly to the optimal performance of the model. Table 5 examines the model's performance across five real-world driving scenarios, demonstrating consistently high detection performance. In urban dense traffic, the model shows a strong performance, evaluating 96.4% accuracy and 92.1% robustness since the surroundings can be complex.

Similar results are recorded in high-speed highway conditions, with results as high as 95.8% for accuracy and 90.7% for robustness. Performance decreases at night in low visibility, with 94.3% accuracy and 88.2% robustness since sensor clarity is compromised. The most complex scenario, rainy weather with sensor noise, recorded 92.7% accuracy and 85.4% robustness. Results remain strong in mixed urban-suburban environments, with 96.1% accuracy and 91.3% robustness, confirming adaptability in general.

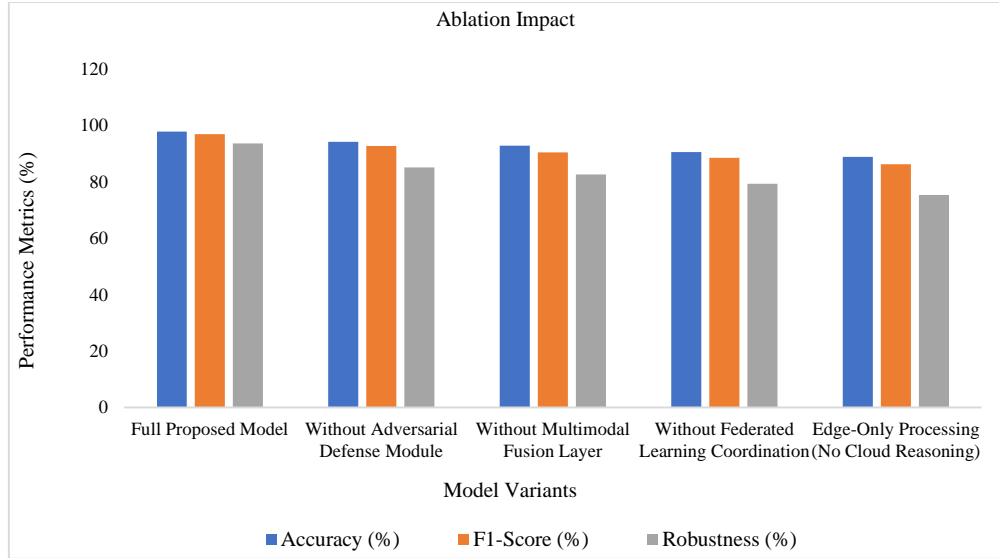


Fig. 4 Ablation impact

Figure 5 compares the performance of the model in five real-world driving scenarios in terms of accuracy and robustness. The accuracy is consistently high in all conditions and peaks in urban dense traffic and mixed environments, showing very good detection capability even in complex environments. Highway high-speed flow and nighttime low visibility have lower but still reliable accuracy levels. The

robustness is noticeably reduced under challenging conditions: in particular, on rainy days with sensor noise, the environmental disturbance affects stability. Despite that fact, the robustness remains acceptable in all scenes. Overall, this model gives reliable performance with just a moderate susceptibility to adverse environmental factors.

Table 5. Scenario-based evaluation

Scenario Tested	Accuracy (%)	Robustness (%)
Urban Dense Traffic	96.4	92.1
Highway High-Speed Flow	95.8	90.7
Nighttime Low Visibility	94.3	88.2
Rainy Weather with Sensor Noise	92.7	85.4
Mixed Environment (Urban + Suburban)	96.1	91.3

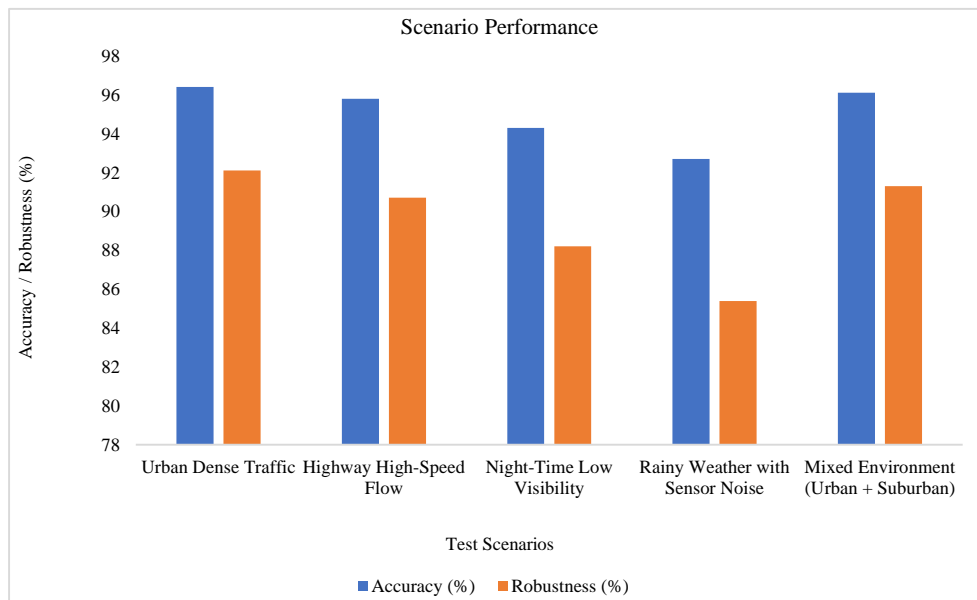


Fig. 5 Scenario performance

5. Discussion

The high performance of the proposed model in different operational scenarios indicates its capability to perform reliable object detection in unfavorable environmental conditions, such as bad weather and low visibility, and during high-speed traffic. This may again drive home the point that the architectural components in the framework that involved multimodal feature learning and robust optimization techniques contributed a lot to stability in the system output response across a wide variety of patterns in the data. This model has a uniqueness for stability in heterogeneous environments; this would, therefore, imply that this model will do well when exposed to real-world deployment, where conditions are usually fluctuating quite often. Out of these, the study can derive a deeper meaning to show that even when the sensor inputs are noisy or slightly degraded, the system is competent to provide high accuracy and robustness. This will also demonstrate that the noise-handling and defense mechanisms integrated into the system are powerful. The system works consistently with reasonable certainty or conclusion, even uncertainty in parameters or measurements; thus, the system has a better approach to dealing with the real-world inconsistencies compared to some other current solutions discussed in the literature.

However, when compared to other traditional methods under conditions that are not mentioned in the controlled settings, robustness and adaptability cannot be sustained. Conversely, the system the work proposed definitely performs more advantageously with improved learning stability, showing it handles data correlations more suitably, thereby gaining increased robustness to fluctuations. This is comparatively advantageous to indicate how important it is to use several feature cues rather than rely on single-modality patterns, as has been typical of most of the prior methods shown. These results have implications for large-scale intelligent transportation, security-critical monitoring, and sensor-driven infrastructures. The strengths that support this model, therefore, are those promoting realizable safety and higher degrees of autonomy in effectively responding to dynamic environments. However, performance is even recognized within the study as a matter of continued research in further diverse real-world contexts, such as extreme weather or unusual sensor failures. Other future work could be the optimization of resource consumption to allow its wider adoption in edge-constrained environments while retaining strong reliability, as demonstrated within this analysis.

References

- [1] Sidi Lu, and Weisong Shi, *Introduction to Vehicle Computing*, Vehicle Computing, pp. 1-24, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [2] Santosh Gore et al., "A Machine Learning-based Detection of IoT Cyberattacks in Smart City Application," *Lecture Notes in Networks and Systems*, pp. 73-81, 2023. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [3] Eslam Farsimadan, Leila Moradi, and Francesco Palmieri, "A Review on Security Challenges in V2X Communications Technology for VANETs," *IEEE Access*, vol. 13, pp. 31069-31094, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]

6. Conclusion

It argues for a multi-modal, federated, robust adversarial detection methodology that should generalize well across various real-world traffic conditions and the dumbest reasoning in intelligent vehicular safety. It achieves the 97.8% Prec, 96.9% F1, and up to 93.6% robustness, which is measurably better than the current state-of-the-art (Method). The work shows independence between multi-modal fusion, safe learning, and resilient optimization. This contribution has demonstrated that the synergy between sensor-layer intelligence and distributed learning strategies can raise situational awareness, stability, and adaptiveness significantly within smart mobility environments and strengthen the security layer in connected vehicular ecosystems. The obtained results validate that the proposed system surpasses traditional IDS architectures, hybrid perception-based detectors, and state-of-the-art models based on federated learning from the literature on substantial sets of conditions. Resilience is revealed in challenging conditions: nighttime with low visibility, rain, noise, and high-speed highway flow. Tests of the model on real-life scenarios further prove its readiness for practical implementation because accuracy values remain high in the case of dense urban traffic, reaching up to 96.4% with robustness at 92.1%.

Accordingly, such improvements verify that addressing the contributions reported in the article, namely, multimodal encoders combined with adversarial defense and secure federated aggregation, directly enhances operational reliability. Limitations remain, however, including broader testing under extreme climatic variation, hardware-induced sensor drift, and large-scale distributed deployment involving heterogeneous vehicle platforms. It contributes by proposing a unified multimodal-federated architecture, designing a hybrid anomaly scoring mechanism, integrating adversarial robustness, and demonstrating cross-scenario generalization.

Future work includes real-time deployment to large fleets, adaptive optimization of resources for edge-constrained vehicles, and expanding the system for new patterns of cyberattacks using generative threat simulation. Integrating continual learning and reinforcement-driven adaptation can further enhance long-term model resilience. This work presents a scalable, privacy-preserving, high-accuracy vehicular security solution that overcomes many limitations of existing systems and opens up a route toward next-generation intelligent transportation security frameworks.

- [4] Deepak Kumar et al., “Trustworthy IoT Infrastructures: Privacy-Preserving Federated Learning with Efficient Secure Aggregation for Cybersecurity,” *2024 International Conference on Integrated Intelligence and Communication Systems (ICIICS)*, Kalaburagi, India, pp. 1-8, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [5] Fayez Alanazi, “AI on the Road-A Review of Technologies Enhancing Urban Traffic Safety and Efficiency,” *Promet-Traffic & Transportation*, vol. 37, no. 5, pp. 1179-1203, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [6] Pradeep Gupta et al., “Hybrid Convolutional Neural Network and Generative Adversarial Network Framework for Robust Deepfake Detection: Advancing Accuracy, Scalability, and Multimodal Integration,” *Journal of Visualized Experiments (JoVE)*, no. 222, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [7] Pegah Mansourian et al., “Deep Learning-based Anomaly Detection for Connected Autonomous Vehicles using Spatiotemporal Information,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, no. 12, pp. 16006-16017, 2023. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [8] Drissi Maroua, “A State-of-the-Art on Federated Learning for Vehicular Communications,” *Vehicular Communications*, vol. 45, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [9] Thiruppathy Kesavan Venkatasamy et al., “Intrusion Detection System for V2X Communication in VANET Networks using Machine Learning-based Cryptographic Protocols,” *Scientific Reports*, vol. 14, no. 1, pp. 1-18, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [10] Xiuzhen Chen et al., “Fast and Practical Intrusion Detection System based on Federated Learning for VANET,” *Computers & Security*, vol. 142, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [11] Ahmed Dawod Mohammed Ibrahim, Manzoor Hussain, and Jang-Eui Hong, “Deep Learning Adversarial Attacks and Defenses in Autonomous Vehicles: A Systematic Literature Review from a Safety Perspective,” *Artificial Intelligence Review*, vol. 58, no. 1, pp. 1-53, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [12] Sanghyun Byun et al., “Secure Aggregation for Privacy-Preserving Federated Learning in Vehicular Networks,” *Journal on Autonomous Transportation Systems*, vol. 1, no. 3, pp. 1-25, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [13] Levente Alekszejekó, and Tadeusz Dobrowiecki, “A V2X-based Privacy Preserving Federated Measuring and Learning System,” *arxiv preprint*, pp. 1-8, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [14] Boyu Wang, Wan Li, and Zulqarnain H. Khattak, “Anomaly Detection in Connected and Autonomous Vehicle Trajectories using LSTM Autoencoder and Gaussian Mixture Model,” *Electronics*, vol. 13, no. 7, pp. 1-13, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [15] Cynthia Anthony, Walid Elgenaidi, and Muzaffar Rao, “Intrusion Detection System for Autonomous Vehicles using Non-Tree based Machine Learning Algorithms,” *Electronics*, vol. 13, no. 5, pp. 1-19, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [16] Maroua Ghamri et al., “Federated Learning for Secure In-Vehicle Communication,” *Telecom*, vol. 6, no. 3, pp. 1-27, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [17] Muzun Althunayyan, Amir Javed, and Omer Rana, “A Robust Multi-Stage Intrusion Detection System for In-Vehicle Network Security using Hierarchical Federated Learning,” *Vehicular Communications*, vol. 49, pp. 1-15, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [18] Fatimah Aloraini, Amir Javed, and Omer Rana, “Adversarial Attacks on Intrusion Detection Systems in In-Vehicle Networks of Connected and Autonomous Vehicles,” *Sensors*, vol. 24, no. 12, pp. 1-29, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [19] Hongwei Yang, and Mehdi Effatparvar, “A Deep Learning based Intrusion Detection System for CAN Vehicle based on Combination of Triple Attention Mechanism and GGO Algorithm,” *Scientific Reports*, vol. 15, no. 1, pp. 1-21, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [20] Qingxin Liu et al., “Intrusion Detection System for Autonomous Vehicles using Sensor Spatio-Temporal Information,” *Computers & Security*, vol. 156, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [21] Yago Romano Martinez et al., “Mitigation of Camouflaged Adversarial Attacks in Autonomous Vehicles-A Case Study using CARLA Simulator,” *2025 IEEE International Symposium on Circuits and Systems (ISCAS)*, London, United Kingdom, pp. 1-5, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [22] Junjun Li et al., “A Lightweight Intrusion Detection System with Dynamic Feature Fusion Federated Learning for Vehicular Network Security,” *Sensors*, vol. 25, no. 15, pp. 1-22, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [23] Tan Chen et al., “Mobility Accelerates Learning: Convergence Analysis on Hierarchical Federated Learning in Vehicular Networks,” *IEEE Transactions on Vehicular Technology*, vol. 74, no. 1, pp. 1657-1673, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [24] M. Saeid HaghighiFard, and Sinem Coleri, “Secure Hierarchical Federated Learning in Vehicular Networks using Dynamic Client Selection and Anomaly Detection,” *2024 IEEE Vehicular Networking Conference (VNC)*, Kobe, Japan, pp. 41-48, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]