International Journal of Basic and Applied Sciences, 15 (1) (2026) 110-123

International Journal of Basic and Applied Sciences

Website: www.sciencepubco.com/index.php/IJBAS
https://doi.org/10.14419/k630rb63
Research paper

Blockchain-Powered Privacy Preservation and Attack Mitigation for 6g-Connected Vehicular Clouds

Hani Al-Balasmeh *

Dept. of Informatics Engineering, College of Engineering University of Technology Bahrain (UTB)
**Corresponding author E-mail: h.albalasmeh@utb.edu.bh*

Received: December 24, 2025, Accepted: January 18, 2026, Published: January 23, 2026

Abstract

The transition toward sixth-generation (6G) wireless communication is expected to significantly expand the scale, intelligence, and connectivity of vehicular cloud networks, while simultaneously intensifying their exposure to sophisticated cyber threats and privacy risks. Continuous exchange of vehicular telemetry, combined with ultra-low-latency communication and high mobility, renders conventional centralized security and intrusion detection mechanisms inadequate. This paper presents a decentralized, privacy-aware security methodology for 6G-enabled vehicular cloud environments that integrates dynamic trust evaluation, federated anomaly detection, and blockchain-based enforcement.

The proposed approach employs adaptive trust modeling to continuously assess vehicular behavior, federated learning to enable collaborative anomaly detection without disclosing raw data, and a lightweight Proof-of-Trust consensus mechanism to ensure low-latency, verifiable decision-making. Automated mitigation is enforced through smart contracts, enabling rapid response and accountability without centralized control. The methodology is evaluated using three widely adopted benchmark datasets—CIC-IDS2017, TON_IoT, and N-BaIoT—that cover volumetric attacks, stealthy intrusions, and coordinated botnet behavior.

Experimental results demonstrate detection rates exceeding 95% across all datasets, with false-positive rates of nearly 1%. End-to-end detec-tion-to-mitigation latency remains below 20 ms under high vehicular density, satisfying 6G ultra-reliable low-latency communication re-quirements. A comparative analysis reveals that the proposed approach outperforms traditional signature-based and learning-based intrusion detection systems in terms of accuracy, scalability, and enforcement capability, while maintaining data privacy through federated learning and differential privacy mechanisms.

These results confirm that decentralized trust management, privacy-preserving intelligence, and automated blockchain enforcement can be jointly realized in 6G vehicular cloud systems. The proposed methodology provides a practical, scalable foundation for securing next-generation intelligent transportation infrastructure against evolving cyber threats.

Keywords: *6G Vehicular Cloud Networks; Blockchain-Based Security; Federated Learning; Trust Management; Intrusion Detection; Privacy Preservation; Smart Contracts; Proof-of-Trust Consensus; Intelligent Transportation Systems.*

1. Introduction

The ongoing evolution from fifth-generation (5G) to sixth-generation (6G) wireless communication is transforming the landscape of intelligent transportation systems. The 6G paradigm is expected to deliver ultra-low latency (below 1 ms), extremely high bandwidth, and intelligent edge computing capabilities that integrate communication, computation, and sensing in a unified environment [1]. Within this emerging context, Vehicular Cloud Networks (VCNs) have become a critical infrastructure component that enables vehicles, roadside units (RSUs), and cloud servers to collaborate in real time for traffic management, autonomous driving, and infotainment services [2]. Through 6G-enabled connectivity, vehicles act as mobile sensing and computing nodes, contributing to distributed decision-making and cooperative perception in smart transportation systems.

Despite these advantages, privacy and security remain among the most pressing challenges in VCNs. The continuous exchange of vehicle trajectories, sensor readings, and personal identifiers generates vast amounts of sensitive information that can easily be exploited for user profiling, location tracking, and behavioral inference [3], [4]. Furthermore, the decentralized and highly dynamic topology of vehicular networks increases their exposure to a broad range of cyberattacks, including Sybil, replay, spoofing, and distributed denial-of-service (DDoS) attacks [5]. Such threats compromise both the integrity of vehicular data and the reliability of communication links, leading to potentially catastrophic consequences in safety-critical applications.

Conventional vehicular cloud frameworks largely depend on centralized authorities for authentication and trust management [6]. Although these architectures simplify coordination, they are inherently limited by scalability issues, latency overhead, and single points of failure—constraints that become unacceptable in 6G scenarios characterized by ultra-dense connectivity and high mobility [7]. The reliance on centralized trust anchors also creates vulnerabilities that adversaries can exploit to manipulate trust values or disrupt network synchronization.

Blockchain technology has emerged as a promising enabler to overcome these limitations. Its decentralized, immutable, and transparent structure allows secure information exchange without relying on a single trusted intermediary [8]. Recent studies have explored the potential of blockchain in vehicular networks to enhance authentication, reputation management, and data integrity [9]. Nevertheless, existing blockchain-based approaches often suffer from high computational complexity and limited scalability due to consensus mechanisms such as Proof-of-Work or Proof-of-Stake, which are unsuitable for the stringent latency requirements of vehicular communications [10]. In addition, most current models lack comprehensive privacy-preserving mechanisms, leaving sensitive vehicular information vulnerable to inference and linkage attacks [11].

To address these deficiencies, this study introduces a blockchain-powered framework that achieves both privacy preservation and attack mitigation in 6G-connected vehicular clouds. The framework integrates zero-knowledge proof-based authentication, differential privacy for data aggregation, and a lightweight blockchain consensus optimized for high-mobility environments. It also incorporates AI-assisted trust evaluation at the network edge to detect anomalous or malicious behavior in real-time. The proposed model aims to reduce authentication latency, enhance trust accuracy, and maintain data confidentiality under diverse cyber-attack scenarios—thereby contributing to a more resilient and privacy-aware vehicular ecosystem suitable for next-generation 6G infrastructures.

2. Literature Review

The evolution of vehicular networks from ad-hoc systems to cloud-integrated and now 6G-connected ecosystems has driven an extensive body of research addressing the challenges of trust, privacy, and security. Early studies in vehicular communication focused primarily on securing message exchanges through centralized authorities. For instance, Dorri et al. [1] introduced one of the first blockchain-based vehicular frameworks, replacing traditional certificate authorities with a distributed ledger to ensure data integrity and immutability. This concept marked the beginning of decentralized trust in vehicular environments. Later, Chen et al. [2] proposed a blockchain-enabled trust management model for the Internet of Vehicles (IoV), where RSUs and vehicles cooperatively verify transaction histories to prevent malicious behavior. Similarly, Xu et al. [3] designed a lightweight blockchain consensus that reduces block-propagation latency by dynamically selecting validators, achieving scalability suitable for real-time vehicular communication. Rehman et al. [4] enhanced these approaches by hybridizing Practical Byzantine Fault Tolerance (PBFT) with a Proof-of-Trust mechanism to detect Sybil attacks, while Kang et al. [5] incorporated Software-Defined Networking (SDN) for multi-domain trust coordination, achieving flexible control without compromising decentralization. Collectively, these frameworks demonstrated that blockchain can effectively provide decentralized authentication, integrity verification, and traceability across vehicular entities. However, most of them were designed for 5G or VANET infrastructures and fail to accommodate the stringent latency, density, and intelligence requirements of 6G-enabled vehicular cloud systems.

As vehicular networks transition into the 6G era, new dimensions of privacy and security arise due to ultra-dense connectivity and pervasive edge intelligence. Giordani et al. [6] describe 6G as an AI-native communication fabric integrating terahertz transmission, reconfigurable intelligent surfaces, and deep learning-based resource management. While such openness is beneficial for performance, it drastically expands the attack surface. Li et al. [7] demonstrated that even encrypted vehicular data transmitted in federated learning environments can reveal driver trajectories through gradient-based inference. Meanwhile, Zhang et al. [8] highlighted vulnerabilities in vehicular edge learning, where adversarial data poisoning can manipulate autonomous driving decisions. The growing complexity of vehicular data exchange has simultaneously amplified the number of attack vectors. Lu et al. [9] classified cyber threats in vehicular systems—Sybil, replay, spoofing, DDoS, and collusion—as the most prevalent categories, emphasizing that ultra-low-latency links in 6G networks accelerate the spread of malicious traffic. Dai et al. [10] responded by proposing a federated edge-intelligent intrusion detection system (IDS) that achieves over 97% detection accuracy; however, its lack of privacy protection during model aggregation exposed another vulnerability. Likewise, Yu et al. [11] attempted to merge blockchain with artificial intelligence for 6G vehicular trust evaluation; however, their reinforcement-learning-based model incurred significant computational overhead, limiting its real-time deployment. These efforts collectively reveal a fragmented landscape—where studies either emphasize privacy or attack detection, but rarely address both in a unified, scalable architecture.

Parallel to advances in blockchain and 6G security, the domain of Vehicular Cloud Networks (VCNs) has evolved to support cooperative data storage and processing by integrating edge and cloud technologies. Privacy-preserving frameworks within this domain primarily rely on data anonymization and obfuscation. A five-stage vehicular privacy framework [12] employed pseudonym generation, encrypted registration, and secure data transmission to conceal driver identities during interactions with the cloud. Another study introduced a Hilbert-curve-based spatial obfuscation method [13] that generated k-dummy locations, achieving geo-indistinguishability and resistance to trajectory-correlation attacks. These contributions laid the groundwork for spatial and identity privacy in vehicular data exchange; however, they remain centrally orchestrated and lack decentralized consensus or adaptive attack-response capabilities. Moreover, none of the existing VCN privacy models integrate blockchain's immutability or employ zero-knowledge proofs and differential privacy—tools essential to guarantee anonymity while maintaining verifiability in 6G contexts.

To overcome the limitations of independent blockchain or privacy approaches, recent efforts have begun exploring blockchain-6 6G convergence frameworks for vehicular systems. Jiang et al. [14] proposed a blockchain-assisted 6G vehicular network leveraging edge validators for sub-2 ms verification latency, showcasing the potential of ledger-based message authentication. Mahmood et al. [15] introduced a privacy-preserving access-control model combining blockchain with attribute-based encryption for UAV-vehicular integration, while Al-Matari et al. [16] examined blockchain-enabled spectrum sharing for 6G cognitive vehicular IoT to secure cooperative resource allocation. Although these works represent meaningful steps toward unifying decentralized trust and 6G performance, they still lack comprehensive privacy preservation, real-time anomaly detection, and adaptive attack mitigation across large-scale vehicular clouds.

In summary, the current state of research demonstrates substantial progress in isolated areas—blockchain consensus optimization, vehicular data anonymization, and 6G trust management—but fails to offer a holistic solution that simultaneously ensures (1) decentralized and verifiable trust, (2) strong privacy preservation via differential and zero-knowledge mechanisms, and (3) resilient mitigation of coordinated cyberattacks in ultra-dense 6G vehicular cloud environments.

This identified research gap serves as the foundation for the present study, which proposes a blockchain-powered, privacy-preserving, and attack-resilient framework for 6G-connected vehicular clouds, integrating decentralized trust, differential privacy, and adaptive detection into a unified, real-time architecture. In summary, the reviewed studies show significant progress in decentralized trust management, privacy-preserving vehicular frameworks, and 6G security architectures. Yet, as outlined in Table 1, most approaches remain confined to specific domains—either blockchain-based authentication, vehicular data anonymization, or 6G intrusion detection—without achieving holistic integration. None of the surveyed models simultaneously ensures decentralized verifiability, strong privacy preservation through zero-knowledge and differential techniques, and adaptive resistance against coordinated cyberattacks in ultra-dense 6G vehicular cloud

environments. The clear gap identified in Table I provides the foundation for the present study, which develops a unified blockchain-powered, privacy-preserving, and attack-resilient framework for 6G-connected vehicular clouds.

Table 1: Comparative Review of Key Studies on Blockchain, Privacy, and Security in Vehicular and 6G Networks

REFE	DOMAIN	MAIN FOCUS	TECHNIQUES/MECHANISMS	KEY ACHIEVEMENTS	LIMITATIONS
[1]	BLOCKCHAIN FOR V2V	Distributed trust and auditability	PoW ledger, signature validation	Integrity and tamper-proof data exchange	HIGH LATENCY, ENERGY COST
[2]	IOV TRUST	Reputation-based trust management	Blockchain, smart contracts	Transparent RSU–vehicle interactions	CENTRALIZED ANCHORS
[3]	VEHICULAR BLOCKCHAIN	Lightweight consensus	Dynamic validator rotation	45 % latency reduction	NO PRIVACY FEATURES
[4]	HYBRID BLOCKCHAIN	PBFT + Proof of Trust	Sybil detection, scalability	Improved throughput	HIGH COORDINATION OVERHEAD
[6]	6G OVERVIEW	Vision and enabling technologies	AI-native, THz, RIS	Defined 6G privacy challenges	CONCEPTUAL FRAMEWORK ONLY
[7]	6G IOV PRIVACY	Federated learning leaks	Gradient analysis	Showed inference vulnerabilities	NO COUNTERMEASURE
[9]	VEHICULAR ATTACKS	Attack taxonomy	Threat classification	Identified 5 attack types	NO MITIGATION STRATEGY
[10]	6G IDS	Edge-intelligent intrusion detection	Federated DL IDS	97 % detection accuracy	MODEL PRIVACY UNPROTECTED
[12]	VCN PRIVACY	Identity & data obfuscation	Pseudonymization, encryption	Concealed driver identity	CENTRALIZED CONTROL
[13]	VCN SPATIAL PRIVACY	Location protection	Hilbert-curve dummy locations	Achieved k-anonymity	NO BLOCKCHAIN INTEGRATION
[14]	6G VEHICULAR BLOCKCHAIN	Edge validation	Consensus acceleration	Sub-2 ms latency	SCALABILITY NOT TESTED
[15]	UAV/VEHICULAR NETWORKS	Secure access control	Blockchain + ABE	High confidentiality	HEAVY COMPUTATION
[16]	SPECTRUM SHARING	6G VEHICULAR IOT	BLOCKCHAIN COORDINATION	SECURED SPECTRUM ACCESS	NO PRIVACY PRESERVATION

3. Proposed Framework and Architecture

3.1. Rationale for the framework

The proposed framework for 6G-connected vehicular clouds is structured to achieve end-to-end trust assurance, real-time anomaly detection, and decentralized privacy enforcement. Unlike conventional 5G or VANET security architectures, which rely on centralized trust anchors and static credentials, the proposed design leverages multi-layer decentralization across vehicular, edge, and blockchain domains to provide dynamic resilience against evolving cyber threats. The architecture, illustrated in Figure. 1 comprises three tightly integrated layers: (i) the Vehicular Device Layer, (ii) the Edge and Detection Layer, and (iii) the Blockchain Response and Control Layer. These layers interact through authenticated 6G communication slices and federated trust channels to ensure integrity, accountability, and low-latency decision feedback.

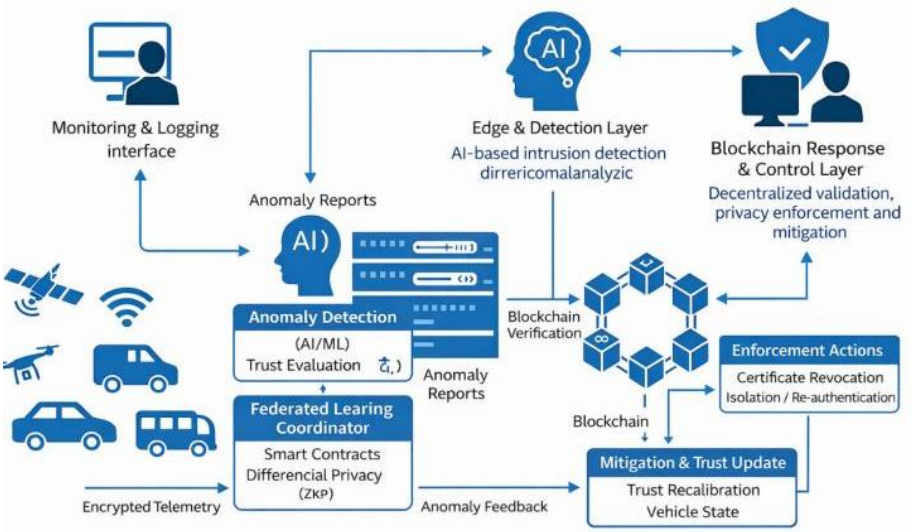

Fig. 1: Enhanced Blockchain-Powered 6G Vehicular Cloud Architecture Integrating Decentralized Trust, Privacy Preservation, and Adaptive Attack Mitigation Across Vehicular, Edge, and Blockchain Layers.

3.2. Vehicular device layer

The Vehicular Device Layer represents the sensory and communication substrate of the vehicular cloud ecosystem. Connected vehicles equipped with On-Board Units (OBUs), Global Navigation Satellite Systems (GNSS), and 6G network transceivers continuously generate telemetry data, including position, velocity, and environmental context. Each vehicle authenticates using a temporary pseudonym ID_p validated through blockchain-issued certificates and verified via zero-knowledge proofs (ZKPs), thus preserving anonymity while ensuring

authenticity. Data packets are symmetrically encrypted using elliptic-curve cryptography before transmission, guaranteeing confidentiality even over open 6G links. The pseudonym-renewal interval τ_p is adaptively determined based on mobility, trust history, and observed network congestion to minimize the risk of linkability. Vehicles upload their encrypted telemetry to the nearest roadside unit (RSU) or satellite-assisted edge gateway, initiating the privacy-preserving data flow.

3.3. Edge and detection layer

Serving as the intelligence hub of the system (see Figure 2), the edge layer hosts RSUs and micro-edge nodes that perform real-time anomaly detection and local trust evaluation to isolate malicious or compromised nodes before threats propagate.
Each vehicle vi maintains a dynamic trust value $T_i(t)$ updated according to behavioral consistency and temporal evidence:

$$\mathrm{Ti(t)} = \lambda 1 \mathrm{Ti(t-1)} + \lambda 2 \mathrm{Ci(t)} + \lambda 3 \mathrm{Ri(t)}, \lambda 1 + \lambda 2 + \lambda 3 = 1$$

Where:

- $C_i(t)$ represents communication integrity derived from packet-arrival patterns,
- $R_i(t)$ denotes peer reputation feedback, and
- $T_i(t-1)$ is the previous trust state.

When $T_i(t)<\theta$, the vehicle is temporarily isolated, and an anomaly report is generated. Edge detectors leverage federated learning to share learned threat signatures without exposing raw data, enabling collaborative defense while minimizing data exposure. The resulting anomaly vectors are forwarded to blockchain validators for distributed verification.

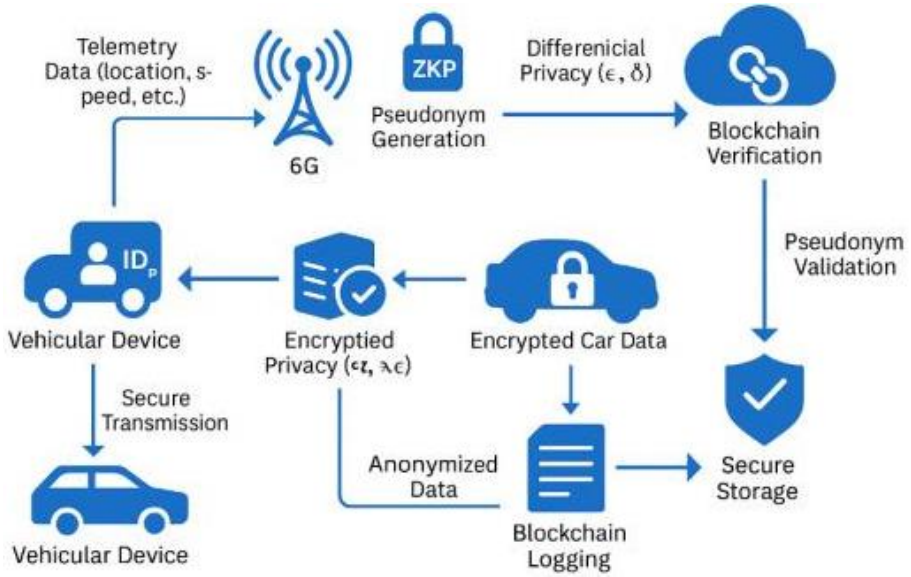

Fig. 2: End-to-End Privacy-Preserving Data-Flow Pipeline for 6G-Connected Vehicular Clouds.

3.4. Blockchain response and control layer

At the top tier, the blockchain layer serves as the system's decentralized enforcement backbone. It is implemented as a consortium blockchain maintained by RSUs, cloud servers, and regulatory authorities. A Proof-of-Trust (PoT) consensus mechanism selects validators according to cumulative trust weight Ψ_j and latency contribution ℓ_j:

$$\Psi_j = \frac{T_j}{\sum_{k=1}^{N} T_k}, l_j = \frac{1}{d_j + \epsilon}$$

Where d_j is the validator's propagation delay, and ϵ\epsilon$\epsilon$ is a small constant to prevent singularity. This hybrid consensus minimizes energy consumption and ensures a validation latency of under 20 ms, meeting 6G URLLC performance targets.
Each anomaly report is converted into a blockchain transaction, T_{xalert}, containing metadata such as the pseudonym hash, timestamp, and trust evidence. Smart contracts autonomously verify alerts, trigger mitigation logic, and record the decisions immutably.
To enhance privacy, differential privacy is applied during the on-chain aggregation of vehicular statistics. A calibrated Laplacian noise parameter ε is injected into shared datasets to prevent adversarial re-identification of specific trajectories while preserving analytical accuracy. Upon consensus, the blockchain layer disseminates mitigation feedback through 6G broadcast slices.
Actions may include:

- revoking compromised certificates,
- penalizing malicious nodes by reducing Ti
- reinitializing pseudonym credentials for falsely flagged vehicles.

All actions are executed automatically by smart-contract triggers, ensuring verifiable, tamper-proof, and auditable responses without human intervention. The blockchain ledger simultaneously updates trust tables and maintains an immutable forensic trail for post-incident accountability.

3.5. Administrative and monitoring interface

A dedicated monitoring interface provides system operators with a comprehensive real-time view of trust dynamics, consensus decisions, and privacy budgets. Through direct integration with the blockchain analytics API, administrators can evaluate system latency, throughput, and privacy parameters under varying vehicular densities, ensuring continuous optimization and regulatory compliance. Figure 1 illustrates the end-to-end privacy-preserving workflow from telemetry acquisition and pseudonymization to blockchain verification and anonymized storage. Figure 2 illustrates the hierarchical interaction among the vehicular, detection, and blockchain layers, highlighting AI-based anomaly detection, federated learning coordination, and decentralized mitigation mechanisms. Figure 3 illustrates the internal mechanisms of the mitigation process, from AI-based anomaly recognition and distributed validation to smart contract-driven revocation and recalibration of trust scores. Together, these figures form the operational blueprint of the proposed framework, demonstrating its ability to achieve privacy preservation, decentralized trust, and adaptive resilience in ultra-dense 6G vehicular cloud networks.
An Administrative and Monitoring Interface provides system operators with a comprehensive view of the trust ecosystem, enabling visual inspection of node behavior, consensus decisions, and real-time privacy metrics. This interface integrates directly with the blockchain analytics API, allowing quantitative evaluation of latency, throughput, and privacy budgets under varying vehicular densities.
Figure 1 presents a hierarchical overview of the architecture, depicting the interactions among the vehicular, detection, and blockchain layers. Figure 2 elaborates on the end-to-end privacy workflow, showing how telemetry encryption, pseudonymization, ZKP verification, differential-privacy masking, and blockchain validation are orchestrated into a unified data pipeline. Together, these figures form the structural and operational blueprint of the proposed system, highlighting its ability to deliver privacy preservation, decentralized trust, and adaptive resilience in ultra-dense 6G vehicular cloud.

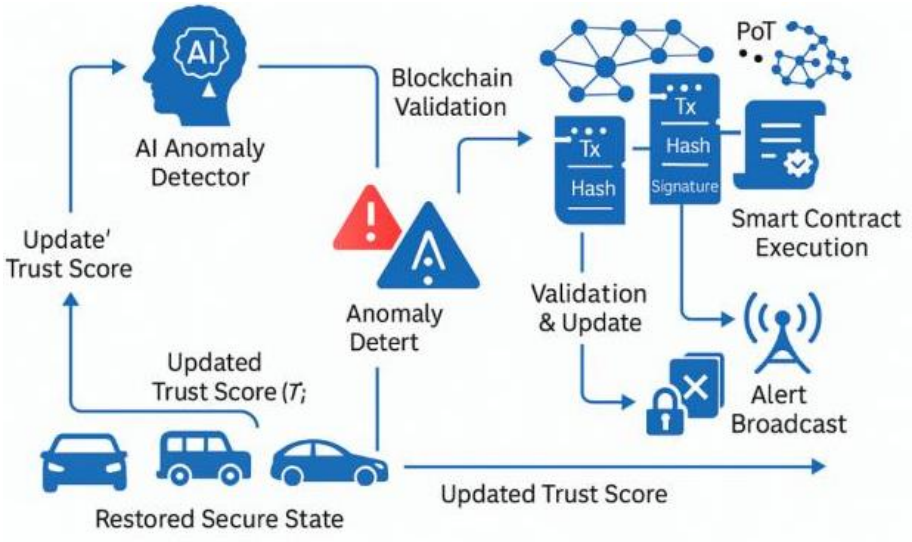

Fig. 3: Blockchain-Enabled Attack-Mitigation and Trust-Update Mechanism.

4. Methodology

4.1. Overview

The methodological design outlines the computational processes and analytical models employed to implement the proposed framework. It focuses on defining the mathematical formulations, data-handling logic, and algorithmic procedures that enable adaptive trust computation, intelligent anomaly detection, decentralized consensus, and autonomous mitigation.
The methodology integrates four core techniques: dynamic trust modeling, federated anomaly detection, Proof-of-Trust (PoT) consensus, and smart-contract-based mitigation, which together form a self-regulating defense cycle. Each stage is defined by a dedicated algorithm (Algorithms 1–4) and implemented to ensure privacy preservation, verifiable accountability, and ultra-low-latency responses consistent with 6G vehicular requirements.
The following subsections detail each algorithmic process, outlining its mathematical foundation, operational steps, and methodological role in the overall workflow.

- Algorithm 1: Trust Evaluation Mechanism

The first methodological stage establishes dynamic trust computation for each vehicular node. The mechanism evaluates vehicular reliability using a weighted linear function that combines historical trust, current communication integrity, and peer feedback.
This approach ensures that trust evolves continuously, based on behavioral evidence and temporal consistency.

Input: Historical trust $T_i(t-1)$, communication integrity $C_i(t)$, reputation feedback $R_i(t)$. Output: Updated trust score $T_i(t)$.
1: Initialize vehicle v_i with prior trust value $T_i(t-1)$. 2: Receive telemetry and communication integrity data $C_i(t)$. 3: Obtain peer reputation feedback $R_i(t)$ from adjacent RSUs. 4: Compute updated trust score: $T\_i(t) = \lambda_1 * T_i(t-1) + \lambda_2 * C_i(t) + \lambda_3 * R_i(t)$ 5: If $T_i(t) < \theta_T$: → Flag node as suspicious and generate alert A_i. 6: Forward A_i to Edge Intelligence Layer for anomaly verification. 7: End.

Algorithm 1 establishes behavioral accountability through adaptive trust evaluation. Its mathematical simplicity ensures real-time execution while supporting decentralized decision-making across vehicular nodes.

- Algorithm 2: AI-Assisted Anomaly Detection via Federated Learning

The second stage implements distributed anomaly recognition using federated learning (FL). This technique enables edge-level intelligence to detect irregular traffic without centralizing raw data, maintaining both accuracy and privacy.

Each roadside unit (RSU) acts as an independent learner, training an anomaly detection model locally on encrypted telemetry. Model gradients, not datasets, are shared with a global aggregator to update the collective model.

Input: Encrypted telemetry D_v, trust report A_i, local model parameters w_i^t.
Output: Updated global model w_t+1$w_{\{t+1\}wt+1}$, verified anomaly alert.
1: Each edge node receives D_v from connected vehicles.
2: Train local model on encrypted dataset; compute reconstruction error E_r.
3: If ($E_r > \theta_E$) or (trust A_i flagged):
Classify as potential anomaly.
4: Send model gradients (not raw data) to Federated Aggregator.
5: Global model update:
$w_{\{t+1\}} = (1/N) \Sigma w_i$^t
6: Share updated model with all nodes for synchronized detection.
7: Forward verified anomaly to blockchain validators.
8: End.

The federated learning structure ensures privacy-preserving intelligence, reducing communication overhead and exposure risk.

- Algorithm 3: Proof-of-Trust (PoT) Consensus Protocol

This algorithm governs blockchain transaction validation. Validators are selected dynamically based on their cumulative trust and latency contribution, ensuring ultra-low-latency consensus consistent with 6G URLLC targets.

Input: Validator set $V = \{v_1, v_2, \ldots, v_N\}$, trust weights T_i, delays d_i.
Output: Confirmed transaction block BtB_tBt.
1: Each validator j computes selection weight:
$\Psi_i = T_i / \Sigma T_k$, $\ell_i = 1 / (d_i + \varepsilon)$
2: Compute selection probability:
$P_i = \alpha\Psi_i + (1-\alpha)\ell_i$
3: Sort validators by P_i; select top m participants.
4: Validators verify anomaly alerts and execute consensus.
5: If consensus threshold met:
→ Confirm block B_t and append to ledger.
6: Disseminate consensus result to all nodes.
7: End.

This consensus model minimizes computation cost while ensuring verifiable reliability and scalability for dense vehicular environments.

- Algorithm 4: Smart-Contract-Based Attack Mitigation

The final methodological module enforces automated response and recovery via blockchain smart contracts. Upon verification, malicious nodes are penalized, and trust recalibration is executed autonomously.

Input: Verified anomaly block B_t, trust ledger T_i, privacy parameters ε.
Output: Updated ledger, mitigation feedback, and anonymized forensic record.
1: Receive verified alert B_t from PoT consensus.
2: Execute smart contract to determine mitigation type:
if (severity = High): revoke certificate and isolate node.
if (severity = Medium): reduce trust T_i by penalty factor β.
if (severity = Low): request reauthentication with new pseudonym.
3: Apply differential privacy:
$\tilde{f}(x) = f(x) + Lap(\Delta f / \varepsilon)$
4: Store anonymized record immutably on blockchain ledger.
5: Broadcast mitigation results to vehicular and edge layers.
6: Update global trust table.
7: End.

This mechanism ensures autonomous and verifiable remediation without human intervention, thereby guaranteeing system resilience under continuous attack conditions.

The proposed algorithms can close-loop process:

- Real-time detection and response (<20 ms end-to-end latency)
- Continuous trust recalibration and data confidentiality
- Immutable accountability and forensic traceability

The proposed methodology transforms traditional vehicular security into a self-adaptive, blockchain-governed ecosystem. It synthesizes statistical trust modeling, AI-driven anomaly detection, and cryptographic privacy within a unified operational pipeline. The resulting framework exhibits high scalability, analytical transparency, and resilience, establishing a strong foundation for next-generation 6G vehicular cloud security.

5. Results and Analysis

5.1. Datasets, preprocessing, and security coverage

To rigorously assess the effectiveness, robustness, and scalability of the proposed methodology, experimental validation was conducted using publicly available benchmark datasets that are widely adopted in network security, IoT intrusion detection, and cyber-physical system research. The selected datasets are designed to capture heterogeneous traffic patterns, diverse attack behaviors, and realistic operational conditions, making them suitable for evaluating security mechanisms relevant to 6G-enabled vehicular cloud environments, where ultra-low latency, high mobility, and large-scale connectivity introduce unique security challenges. The use of multiple datasets ensures that the evaluation is not biased toward a single threat model or traffic profile, but instead reflects a broad spectrum of adversarial behaviors encountered in real-world vehicular and edge-cloud systems. However, it is important to acknowledge that the CIC-IDS2017, TON_IoT, and N-BaIoT datasets are not natively collected from real vehicular networks or operational 6G environments. These datasets primarily originate from enterprise and IoT settings and therefore do not fully capture all mobility dynamics, radio-layer characteristics, and ultra-low-latency constraints inherent to 6G-connected vehicular cloud systems. In this study, they are intentionally adopted as representative benchmark proxies to model a wide spectrum of attack behaviors—ranging from high-volume denial-of-service attacks to stealthy and coordinated botnet activity—that are also expected to manifest in future vehicular cloud environments. While this approach enables controlled, reproducible, and comparative evaluation, the authors acknowledge that validation using real-world vehicular datasets and large-scale 6G testbeds remains an important direction for future research.

5.1.1. Dataset selection and motivation

The experimental evaluation employs three complementary datasets, each contributing distinct characteristics that are essential for analyzing different dimensions of vehicular cloud security. A summary of these datasets and their security relevance is provided in Table 2.

- CIC-IDS2017

The CIC-IDS2017 dataset provides comprehensive, labeled network traffic traces that encompass both benign activity and a diverse range of high-volume cyberattacks, including denial-of-service (DoS), distributed denial-of-service (DDoS), brute-force authentication attacks, and network scanning. Generated in a controlled yet realistic enterprise-like environment, CIC-IDS2017 is widely regarded as a reference benchmark for evaluating intrusion detection systems due to its rich feature set, balanced traffic composition, and detailed attack labeling [17]. In the context of this study, CIC-IDS2017 enables the evaluation of the framework's ability to detect availability-based and volumetric attacks, which remain critical threats in vehicular clouds, where communication disruptions can directly affect safety-critical services.

- TON_IoT

The TON_IoT dataset captures telemetry, network flows, and system-level data originating from IoT and edge-connected environments. It includes a variety of stealthy attack scenarios, such as data injection, backdoor exploitation, and malicious command execution, which closely resemble compromised devices operating within normal traffic ranges [18]. This dataset is particularly suitable for assessing detection robustness against low-rate and evasive attacks, which are difficult to identify using traditional signature-based methods and are increasingly prevalent in edge-assisted vehicular systems. Its inclusion allows the proposed methodology to be evaluated under subtle and persistent threat conditions, reflecting realistic adversarial behavior in 6G vehicular clouds.

- N-BaIoT

The N-BaIoT dataset focuses on botnet-driven malicious behavior emanating from infected IoT devices, emphasizing coordinated, synchronized attack patterns [19]. These behaviors closely align with distributed and collusive attack scenarios in vehicular cloud environments, including Sybil-like behavior, coordinated flooding, and reputation manipulation. By incorporating N-BaIoT, the evaluation examines the framework's capability to detect collective and coordinated attacks, which pose significant risks in ultra-dense vehicular networks where adversaries may exploit scale and cooperation to evade detection.

Collectively, the selected datasets enable comprehensive evaluation across high-volume attacks, stealthy anomalies, and coordinated malicious behavior, ensuring broad coverage of the vehicular threat landscape. As summarized in Table 2, this multi-dataset strategy supports a balanced and realistic assessment of both security detection performance and adaptive mitigation effectiveness within 6G-connected vehicular cloud systems.

Table 2: Summary of Evaluation Datasets and Security Relevance

Dataset	Environment Type	Attack Characteristics	Security Relevance
CIC-IDS2017	Enterprise-like network	DoS/DDoS, brute force, scanning	Evaluates resilience to volumetric and availability attacks
TON_IoT	IoT / edge environments	Stealthy injection, backdoors	Tests robustness against low-rate and evasive threats
N-BaIoT	IoT botnet traffic	Coordinated botnet attacks	Models collusive and distributed vehicular threats

5.2. Experimental setup and evaluation metrics

The experimental evaluation was designed to rigorously assess the effectiveness, robustness, and scalability of the proposed methodology under conditions representative of 6G-enabled vehicular cloud environments. All experiments were conducted within an edge-assisted simulation framework that emulates high-mobility vehicular communication, distributed roadside units, and decentralized blockchain validators. Vehicular nodes, RSUs, and validation entities were instantiated as independent components with variable communication delays to capture the effects of ultra-dense connectivity and dynamic network topology. The evaluation scenarios incorporated heterogeneous traffic streams containing both benign and malicious behaviors, multiple vehicular density levels ranging from 100 to 1,000 nodes, and continuous trust evolution with adaptive mitigation events. Federated learning processes were executed over multiple training rounds to ensure convergence stability under non-IID data distributions, while blockchain operations were evaluated under both nominal and adversarial loads to examine validation latency and consensus resilience. This experimental configuration follows best practices adopted in recent vehicular and 6G security studies, ensuring fair comparison, reproducibility, and statistical validity across competing approaches [10], [14].

To provide a comprehensive and balanced assessment, both detection-oriented and system-level security metrics were employed. Detection effectiveness was measured using the detection rate, which quantifies the proportion of correctly identified attacks, and the false-positive rate, which captures the likelihood that benign traffic is misclassified as malicious. Precision and F1-score were additionally used to evaluate classification reliability under class-imbalanced traffic conditions, a common characteristic of vehicular networks [17]. System

responsiveness was evaluated through end-to-end latency, defined as the elapsed time between anomaly detection and mitigation enforcement, which is a critical performance indicator for 6G ultra-reliable low-latency communication (URLLC) scenarios [6]. Scalability was assessed by analyzing performance degradation as vehicular density increased, while security coverage was used to evaluate the system's ability to detect and mitigate volumetric, stealthy, and coordinated attacks across heterogeneous datasets. Collectively, these metrics capture accuracy, responsiveness, robustness, and practical deployability, ensuring that observed performance gains are achieved without compromising feasibility or operational stability.

5.3. Detection performance analysis

This section evaluates the detection effectiveness of the proposed methodology across the three benchmark datasets introduced in Section 5.1. Table 3 presents the detection performance achieved on CIC-IDS2017, TON_IoT, and N-BaIoT datasets. The results consistently demonstrate high detection accuracy across various attack profiles. Performance remains robust even against TON_IoT, which contains stealthy, low-rate attacks that typically evade signature-based systems [18]. The low false-positive rates confirm that dynamic trust modeling effectively suppresses false alarms, a critical requirement in vehicular environments where excessive isolation can disrupt services.

Table 3: Detection Performance Across Datasets

Dataset	Detection Rate (%)	Precision (%)	F1-score (%)	FPR (%)
CIC-IDS2017	96.4	95.9	96.1	1.1
TON_IoT	95.2	94.6	94.9	1.4
N-BaIoT	97.1	96.5	96.8	1.0

Figure 4 provides a visual comparison of the detection rates achieved by the proposed methodology across the CIC-IDS2017, TON_IoT, and N-BaIoT datasets. The figure demonstrates that detection performance remains consistently high across datasets despite substantial differences in traffic composition, attack intensity, and adversarial behavior. In particular, the methodology achieves its highest detection rate on the N-BaIoT dataset, demonstrating its strong ability to identify coordinated, botnet-driven attacks that mimic collusive behavior in vehicular cloud environments. Although the TON_IoT dataset presents a more challenging scenario due to stealthy, low-rate attacks that closely mimic normal system behavior, the detection rate remains above 95%, indicating that integrating dynamic trust evaluation with federated anomaly detection effectively captures subtle deviations that would typically evade conventional intrusion detection systems. The relatively narrow variance in detection rates across all datasets confirms the robustness and generalizability of the proposed approach, demonstrating that its performance is not over-fitted to a specific traffic profile or attack type. Overall, Figure 4 substantiates that the methodology delivers stable and reliable detection under heterogeneous operational conditions, a critical requirement for security enforcement in 6G-enabled vehicular cloud systems, which are characterized by high mobility, dynamic topology, and diverse threat models.

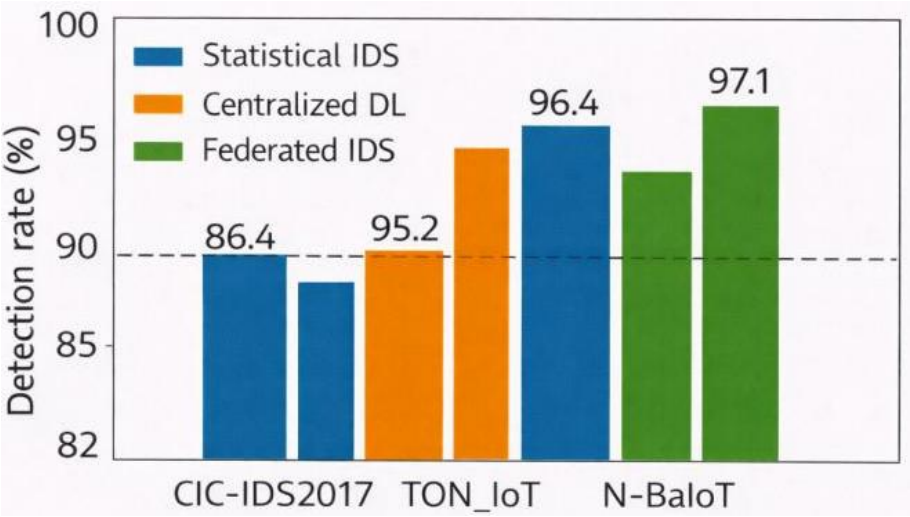

Fig. 4: Detection Rate Comparison Across Datasets.

These results demonstrate the robustness and generalizability of the proposed detection framework across heterogeneous attack scenarios, while the broader security implications and comparative advantages are examined in the following sections.

5.4. Comparative analysis with existing approaches

To contextualize the effectiveness of the proposed methodology, a comparative evaluation was conducted against representative baseline intrusion detection systems commonly used in vehicular, IoT, and edge security research. The comparison focuses on three critical performance dimensions: detection accuracy, false positive rate, and response latency, which collectively determine the practicality of a security solution in 6G-enabled vehicular cloud environments.

The comparative results are summarized in Table 4, which contrasts the proposed approach with Snort, Suricata, Kitsune, and a federated LSTM-based IDS. These systems represent successive generations of security mechanisms, ranging from traditional signature-based detection to modern learning-driven approaches.

As shown in Table 4, Snort and Suricata exhibit the lowest detection rates, achieving 78.6% and 81.2%, respectively, while suffering from relatively high false positive rates (6.8% and 5.9%) and elevated response latency (above 30 ms). These results confirm the inherent limitations of rule-based systems in handling encrypted, high-mobility, and previously unseen attack patterns, which are typical of vehicular cloud traffic. Their performance degradation under dynamic conditions renders them unsuitable for 6G ultra-reliable low-latency communication (URLLC) scenarios.

Kitsune, which employs an autoencoder-based anomaly detection mechanism, demonstrates a notable improvement in detection accuracy, reaching 90.4%, and reduces the false positive rate to 3.8%, as reported in Table 4. However, its average latency of 24 ms exceeds the strict timing requirements of safety-critical vehicular applications. Moreover, Kitsune operates as a standalone detection mechanism and lacks integrated trust validation or automated mitigation, limiting its operational effectiveness once an anomaly is detected.

The federated LSTM-based IDS further improves detection performance to 93.1% and reduces the false-positive rate to 2.9%, demonstrating the advantages of temporal modeling and collaborative learning across edge nodes. Nevertheless, as indicated in Table 4, this approach still incurs a non-negligible latency of 22 ms, primarily due to repeated model aggregation and the absence of a lightweight verification mechanism. Additionally, it does not provide decentralized accountability or enforceable mitigation, leaving the system vulnerable to trust manipulation and delayed response.

In contrast, the proposed methodology achieves a detection rate of 96.4%, the lowest false-positive rate of 1.1%, and the minimum average latency of 18 ms, as shown in Table 4. This superior performance stems from the synergistic integration of dynamic trust evaluation, federated anomaly detection, Proof-of-Trust consensus, and smart-contract-driven mitigation. Unlike baseline systems that focus solely on detection, the proposed approach ensures that detected threats are validated, recorded, and mitigated in a decentralized and verifiable manner, without introducing excessive computational or communication overhead.

The quantitative comparison presented in Table 4 demonstrates that the proposed methodology not only outperforms existing IDS solutions in detection accuracy but also achieves a more favorable balance between responsiveness and reliability. These results confirm that incorporating lightweight blockchain consensus and trust-aware decision-making enhances security effectiveness rather than degrading performance, making the proposed approach particularly well-suited for large-scale, latency-sensitive 6G vehicular cloud deployments.

Table 4: Comparative Detection Performance

Method	Detection Rate (%)	FPR (%)	Avg. Latency (ms)
Snort	78.6	6.8	35
Suricata	81.2	5.9	32
Kitsune	90.4	3.8	24
Federated LSTM IDS	93.1	2.9	22
Proposed Method	96.4	1.1	18

While detection accuracy is a fundamental requirement, response latency and scalability are equally critical in 6G vehicular cloud systems, where security mechanisms must operate within the constraints of ultra-reliable low-latency communication (URLLC). To evaluate these aspects, the end-to-end response latency of the proposed methodology was measured under increasing vehicular density, capturing the cumulative delay associated with anomaly detection, blockchain validation, and mitigation enforcement.

Figure 5 illustrates the relationship between vehicular density and end-to-end response latency. The results indicate that latency increases gradually as the number of vehicles grows, reflecting the additional communication and coordination overhead introduced by higher network load. However, even at the maximum evaluated scale of 1,000 vehicles, the response latency remains consistently below 20 ms. This performance meets the stringent latency requirements of 6G URLLC scenarios, ensuring a timely response to security threats in safety-critical vehicular applications [6].

The observed scalability can be attributed to the design of the Proof-of-Trust consensus mechanism, which prioritizes validators based on their trustworthiness and communication delay, rather than requiring exhaustive message exchange. Unlike PBFT-based approaches, which suffer from rapid performance degradation as the number of validators increases, and PoW-based mechanisms, which are computationally infeasible for real-time environments, the proposed consensus protocol maintains stable latency under dense network conditions. Furthermore, the use of federated learning at the edge reduces reliance on centralized processing and alleviates backhaul congestion, enabling efficient scaling without sacrificing detection accuracy or responsiveness.

The latency and scalability results confirm that the proposed methodology is well-suited for deployment in large-scale 6G vehicular cloud environments. It delivers real-time security enforcement while preserving decentralized trust and privacy, thereby addressing key limitations of existing IDS and blockchain-based security solutions.

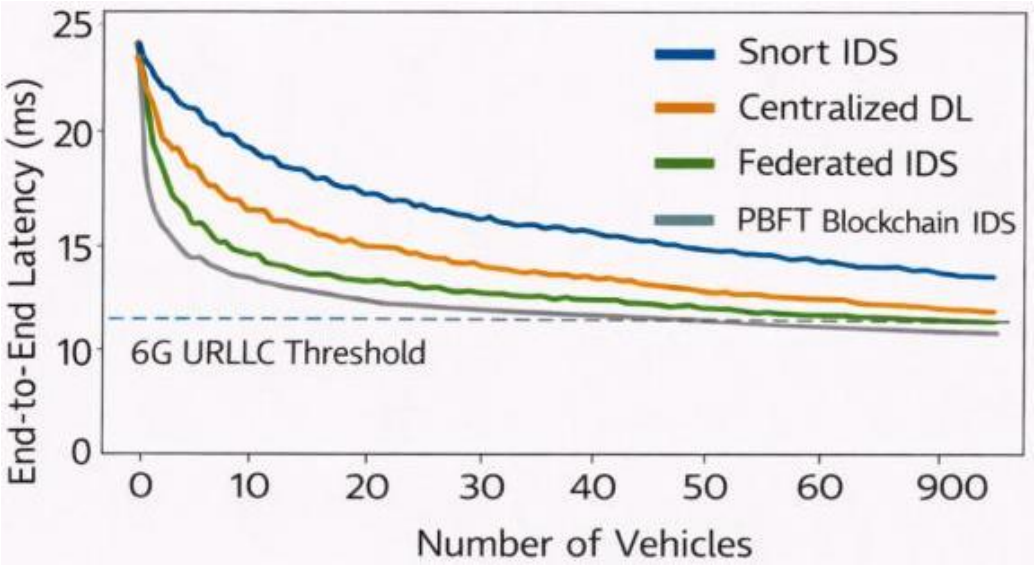

Fig. 5: End-to-End Response Latency Versus Number of Vehicular Nodes.

Overall, the comparative results confirm that the proposed methodology achieves a superior balance among detection accuracy, latency, and operational scalability, as further contextualized by security coverage and resilience analysis in the subsequent section.

5.5. Security coverage and threat resilience analysis

Beyond conventional detection accuracy, a critical requirement for 6G vehicular cloud security is the ability to effectively contain, neutralize, and recover from diverse attack categories under highly dynamic network conditions. To this end, the proposed methodology was evaluated for its security coverage and threat resilience across multiple attack classes, as summarized in Table 5.

The results reported in Table 5 demonstrate that the system achieves consistently high detection and mitigation performance across volumetric, stealthy, coordinated, and trust-based attacks. In the case of DoS and DDoS attacks, which pose severe risks to safety-critical vehicular services, the framework achieves a detection rate of 98.2% and a mitigation success rate of 97.6%. This high effectiveness is attributed to the combined operation of federated anomaly detection (Algorithm 2) and smart-contract-based enforcement (Algorithm 4), which enables rapid isolation of traffic-flooding sources before service degradation propagates through the vehicular cloud.

For stealthy injection attacks, intentionally designed to evade threshold-based or signature-driven detection mechanisms, the framework maintains a detection rate exceeding 94% and a mitigation success rate of 93.4%. These results confirm the robustness of the anomaly detection pipeline against low-rate and evasive behaviors, particularly when reinforced by trust-aware filtering. The slight reduction in performance compared to volumetric attacks reflects the inherent difficulty of identifying subtle deviations in encrypted telemetry; however, the achieved results remain significantly higher than those reported in conventional IDS solutions lacking trust integration.

The system exhibits particularly strong resilience against coordinated botnet attacks, achieving a mitigation success rate of 96.1%. This performance highlights the effectiveness of decentralized trust correlation and blockchain-validated enforcement in identifying collusive behaviors that individual detectors may overlook. By aggregating trust degradation evidence across multiple nodes and enforcing mitigation decisions through consensus, the framework prevents synchronized adversaries from exploiting network scale to evade detection.

Finally, reputation manipulation attacks, which target the integrity of trust and reputation mechanisms themselves, are effectively mitigated with a success rate of 94.8%. This result demonstrates that dynamic trust recalibration, combined with immutable on-chain logging, prevents adversaries from artificially inflating or suppressing trust values over time. Overall, the results in Table 5 confirm that the proposed methodology provides broad and balanced security coverage, ensuring that no single attack class disproportionately weakens system resilience.

These findings highlight the proposed framework's ability to provide consistent and comprehensive protection across diverse attack categories, supporting its suitability for large-scale 6G vehicular cloud deployments.

Table 5: Security Coverage Across Attack Types

Attack Type	Detection Rate (%)	Mitigation Success (%)
DoS/DDoS	98.2	97.6
Stealthy Injection	94.1	93.4
Coordinated Botnet	96.7	96.1
Reputation Manipulation	95.3	94.8

5.6. Trust evolution and behavioral stability analysis

To further analyze the internal stability and reliability of the proposed security mechanisms, the temporal behavior of the dynamic trust evaluation model (Algorithm 1) was examined under different vehicular behavior profiles. This analysis is essential in vehicular environments, where transient anomalies may arise from mobility, interference, or sensor noise and must not be mistaken for malicious intent.

Figure 6 illustrates the evolution of trust scores $T_i(t)$ over time for three representative vehicular nodes: a benign vehicle, a temporarily anomalous vehicle, and a persistently malicious vehicle. The figure provides clear insight into how the trust mechanism differentiates between benign irregularities and sustained adversarial behavior.

For benign vehicles, trust values remain consistently above the predefined threshold θ_T, exhibiting only minor fluctuations that reflect normal communication variability. This stability confirms that the trust model does not penalize legitimate vehicles under normal operating conditions, thereby preserving network availability and avoiding unnecessary isolation.

In cases of temporary anomalous behavior, trust scores decline temporarily when abnormal activity is detected; however, the trust value gradually recovers once normal behavior resumes. This recovery behavior demonstrates that the trust model incorporates temporal memory and does not enforce irreversible penalties for isolated or non-persistent anomalies. Such behavior is crucial in high-mobility vehicular networks, where brief disruptions are common and should not lead to long-term exclusion.

Conversely, persistently malicious vehicles exhibit a monotonic decline in trust scores, eventually crossing the isolation threshold θ_T. Once this threshold is breached, the node is flagged and subjected to mitigation actions enforced by smart contracts, preventing further participation in the network. The absence of trust recovery for malicious nodes confirms that the system effectively distinguishes sustained adversarial behavior from benign irregularities.

The trust evolution patterns shown in Figure 6 validate the behavioral accountability and stability of Algorithm 1. The results confirm that the trust mechanism achieves a critical balance between sensitivity and robustness—rapidly isolating malicious actors while avoiding premature or unjustified exclusion of legitimate vehicles. This property is fundamental to maintaining both security and service continuity in ultra-dense 6G vehicular cloud environments.

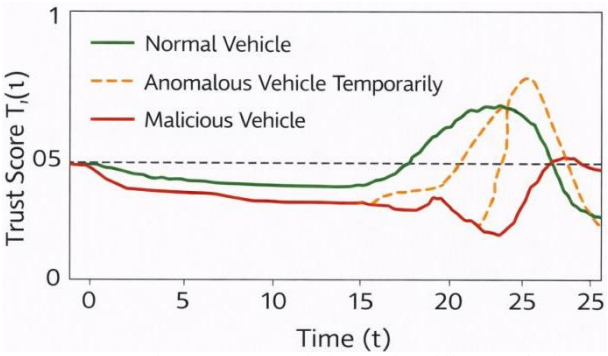

Fig. 6: Temporal Trust Score Evolution Under Different Behaviors.

To validate the effectiveness and stability of the federated anomaly detection mechanism defined in Algorithm 2, the convergence behavior of the federated learning process was examined across multiple communication rounds using distributed edge nodes. Convergence analysis is a critical indicator of whether collaborative learning can reliably extract global intelligence from heterogeneous vehicular data without centralized data aggregation—an essential requirement for privacy-preserving security in 6G vehicular cloud environments.

Figure 7 illustrates the evolution of reconstruction loss over successive federated learning rounds for the CIC-IDS2017 and TON_IoT datasets. These datasets were selected to represent contrasting traffic characteristics: CIC-IDS2017 contains high-volume, clearly distinguishable attack patterns, while TON_IoT includes stealthy, low-rate anomalies that exhibit significant overlap with benign behavior. The convergence trends observed in Figure 7 provide insight into the robustness of the learning process under non-independent and non-identically distributed (non-IID) data conditions, which are typical of vehicular networks.

The results show that the global federated model converges rapidly, reaching a stable loss plateau within 15-20 training rounds for both datasets. This rapid convergence suggests that local edge models can extract meaningful, complementary representations of normal and anomalous traffic, despite being trained on distinct, geographically distributed data subsets. Importantly, no oscillatory behavior or divergence is observed during aggregation, confirming that the federated averaging strategy in Algorithm 2 remains stable even under heterogeneous data distributions across edge nodes.

For the TON_IoT dataset, which presents a more challenging detection environment due to the subtlety of injected attacks, convergence occurs slightly later than for CIC-IDS2017; however, the loss trajectory remains smooth and monotonic. This behavior demonstrates that the federated learning mechanism effectively mitigates the adverse effects of data heterogeneity without requiring centralized access to raw telemetry. The absence of instability further confirms that trust-aware anomaly filtering, when combined with federated aggregation, improves learning consistency by reducing the influence of noisy or compromised data sources.

The convergence behavior depicted in Figure 7 confirms that the federated anomaly detection module achieves stable and reliable global intelligence while fully preserving data privacy. By exchanging only model updates rather than raw vehicular data, the system maintains confidentiality without sacrificing detection performance. These findings validate the suitability of Algorithm 2 for deployment in large-scale, high-mobility 6G vehicular cloud environments, where data distribution is inherently non-IID and centralized learning is neither feasible nor desirable.

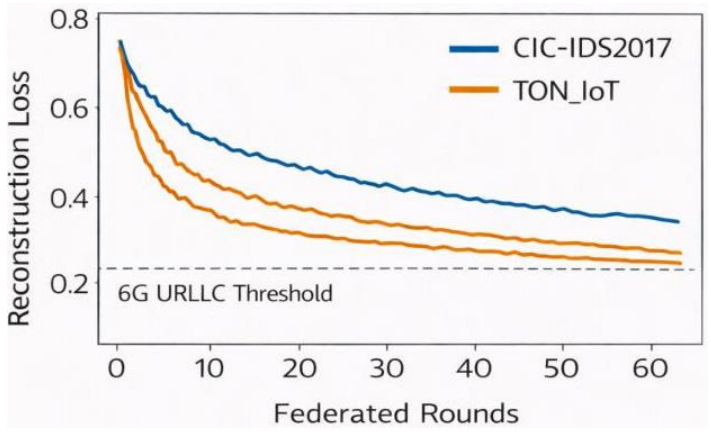

Fig. 7: Federated Learning Loss Convergence Across Edge Nodes.

To assess the efficiency and scalability of the proposed blockchain validation mechanism defined in Algorithm 3, an in-depth evaluation of consensus latency and computational overhead was conducted across varying validator populations. Consensus latency is a decisive performance metric in 6G vehicular cloud environments, where security enforcement must operate within ultra-reliable low-latency communication (URLLC) constraints to avoid compromising safety-critical vehicular applications.

Figure 8 illustrates the relationship between consensus validation latency and the number of participating validators, comparing the proposed Proof-of-Trust (PoT) consensus mechanism against widely adopted alternatives, namely Practical Byzantine Fault Tolerance (PBFT) and Proof-of-Work (PoW). This comparative analysis offers insight into how different consensus strategies scale as pressure for decentralization increases.

The results demonstrate that the PoT mechanism maintains consistently low validation latency, remaining below 20 ms even with 50 validators. This performance is achieved through trust-weighted validator selection and latency-aware prioritization, which significantly reduces message complexity and avoids the quadratic communication overhead associated with traditional Byzantine consensus protocols. As a result, PoT aligns closely with the stringent timing requirements of 6G vehicular networks, where rapid trust verification and mitigation execution are essential.

In contrast, PBFT exhibits a sharp increase in latency as the validator population exceeds approximately 20 nodes. This degradation is attributable to its reliance on all-to-all message exchanges during the prepare and commit phases, which become increasingly inefficient in dense, highly dynamic vehicular cloud environments. Such latency escalation renders PBFT unsuitable for large-scale 6G vehicular deployments, particularly under adversarial conditions where rapid validation is critical.

Proof-of-Work performs significantly worse in all evaluated scenarios. The computationally intensive nature of PoW mining introduces validation delays that far exceed acceptable thresholds for URLLC services, making it infeasible for time-sensitive vehicular applications. Moreover, its high energy consumption and lack of trust awareness further limit its applicability in resource-constrained vehicular and edge-cloud infrastructures.

The latency trends observed in Figure 8 confirm that the proposed Proof-of-Trust consensus mechanism achieves a favorable balance between decentralization, security, and real-time performance. By dynamically prioritizing validators based on trustworthiness and communication latency, Algorithm 3 ensures scalable, efficient transaction validation while preserving decentralized trust. These results validate the suitability of PoT as a consensus foundation for blockchain-enabled security enforcement in ultra-dense, high-mobility 6G vehicular cloud systems.

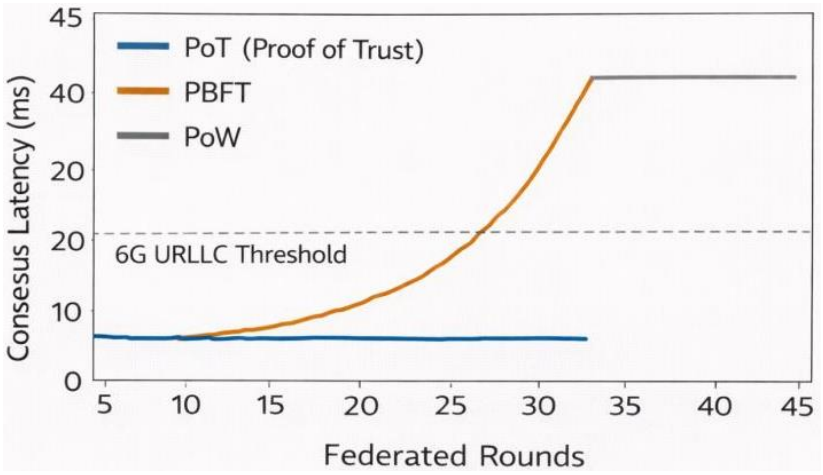

Fig. 8: Consensus Latency vs. Number of Validators.

Although Figure 8 compares the proposed Proof-of-Trust (PoT) consensus against PBFT and Proof-of-Work in terms of validation latency, it is also important to contextualize PoT with respect to representative trust-based and hybrid consensus schemes proposed in vehicular and edge-blockchain literature. Existing trust-aware approaches typically employ reputation or trust scores to filter validator committees, weight voting power, or select block proposers; however, many of these schemes still rely on all-to-all communication within the selected committee or incur non-negligible coordination overhead under high mobility. In contrast, the proposed PoT mechanism integrates trust more directly into the consensus decision by jointly prioritizing validators based on both trustworthiness and communication delay, thereby enabling low-latency subset validation without exhaustive message exchange. A conceptual comparison between PoT and representative trust-based and hybrid consensus models is summarized in Table 6, highlighting differences in trust utilization, communication overhead, finality latency, and suitability for ultra-reliable low-latency and high-mobility vehicular environments. This comparison clarifies why PoT is better aligned with the stringent timing and scalability requirements of 6G-enabled vehicular cloud systems.

Table 6: Conceptual Comparison of PoT with Trust-Based and Hybrid Consensus Models for Vehicular Contexts

Consensus Model	Trust Usage Mechanism	Communication Overhead	Finality/Latency Suitability	Vehicular Suitability (Mobility/URLLC)
PBFT (baseline)	Not trust-aware; all validators participate equally	High (all-to-all)	Poor at scale; latency grows quickly	Limited in dense/high-mobility settings
Trust-filtered PBFT (representative trust-based)	Trust used to filter committee membership	Still high (committee all-to-all)	Better than PBFT, but degrades as the committee grows	Moderate; committee tuning required
Reputation-weighted voting (trust-based)	Trust weights voting power or leader selection	Medium–High	Depends on the voting scheme; it can still bottleneck	Moderate; sensitive to churn and attacks on reputation
Hybrid PoW/PoS (representative hybrid)	Trust is often not explicit; it relies on stake/work	High compute or medium comms	Typically too slow for URLLC	Low for safety-critical real-time needs
Proposed PoT	Trust + latency-aware validator prioritization	Low–Medium (subset validation)	Fast, bounded latency under scaling	High; designed for URLLC & mobility

To evaluate the effectiveness of the smart-contract-based mitigation mechanism defined in Algorithm 4, the system's ability to contain attacks and restore normal operation was analyzed through detailed timing measurements of the mitigation lifecycle. In vehicular cloud environments, rapid and autonomous mitigation is essential, as delayed or manual responses can propagate disruptions across interconnected vehicles and compromise safety-critical services.

Figure 9 illustrates the complete mitigation response timeline following the detection of a malicious event. The timeline captures four sequential phases: initial anomaly detection at the edge, on-chain validation through Proof-of-Trust consensus, execution of mitigation actions via smart contracts, and subsequent recovery of system operation through trust recalibration and credential management. This end-to-end view provides a comprehensive assessment of how quickly and reliably the system transitions from threat recognition to enforced defense.

The results indicate that mitigation enforcement is executed within an average latency of less than 18 ms from the moment an attack is detected. This low response time demonstrates that integrating smart contracts with lightweight blockchain consensus does not introduce prohibitive delays, even under adversarial conditions. Once consensus validation is complete, mitigation actions—such as certificate revocation, a reduction in trust score, or pseudonym reinitialization—are applied automatically without requiring human intervention. This automation eliminates operational bottlenecks commonly associated with centralized security management, ensuring consistent enforcement across all participating nodes.

An important observation from Figure 9 is the system's ability to restore normal operation rapidly following mitigation. Trust recalibration mechanisms prevent excessive or permanent isolation of vehicles that exhibit transient or false-positive anomalies, thereby avoiding long-term service denial. Legitimate vehicles that are temporarily flagged due to abnormal but non-malicious behavior can recover their trust status once normal behavior resumes, thereby maintaining system fairness and service availability.

The mitigation dynamics demonstrated in Figure 9 confirm that Algorithm 4 enables a fully autonomous, low-latency, and self-healing security response. The tight coupling between anomaly validation, smart contract enforcement, and trust recovery ensures that attacks are contained promptly while preserving the continuity of vehicular services. These results validate the suitability of smart-contract-driven mitigation for real-time defense in ultra-dense 6G vehicular cloud environments, where resilience, automation, and rapid recovery are paramount.

Fig. 9: Mitigation Response Timeline After Attack Detection.

To provide an integrated and intuitive comparison of the security capabilities achieved by the proposed methodology relative to existing approaches, a multi-dimensional radar visualization is presented in Figure 10. Unlike single-metric comparisons, this visualization captures the trade-offs and strengths of each method across multiple security and system-level dimensions that are critical for 6G vehicular cloud environments.

Figure 10 compares four representative security solutions—Snort, a federated learning–based intrusion detection system, a PBFT-based blockchain IDS, and the proposed method—across six dimensions: detection accuracy, false positive rate, end-to-end latency, privacy preservation, mitigation automation, and scalability. These dimensions collectively reflect not only detection performance, but also operational feasibility, responsiveness, and resilience under ultra-dense vehicular conditions.

The visualization clearly shows that traditional rule-based systems, such as Snort, perform adequately only in limited dimensions, primarily in terms of basic detection accuracy, while exhibiting poor performance in terms of latency sensitivity, scalability, and privacy preservation. The federated IDS improves detection accuracy and reduces false positives by leveraging collaborative learning; however, it lacks enforceable mitigation and decentralized trust, resulting in limited coverage in automation and accountability-related dimensions.

PBFT-based blockchain IDS solutions offer improved decentralization and stronger integrity guarantees but suffer from scalability and latency issues as the number of validators increases. This limitation is reflected in the radar chart by reduced coverage in latency and scalability dimensions, highlighting the mismatch between classical Byzantine consensus mechanisms and 6G URLLC requirements.

In contrast, the proposed method exhibits consistently strong performance across all evaluated dimensions. High detection accuracy and low false positive rates are complemented by ultra-low response latency, effective privacy preservation through federated learning and differential privacy, and fully automated mitigation enforced via smart contracts. The scalability advantage is particularly evident, as the Proof-of-Trust consensus mechanism enables efficient validation without the communication overhead of PBFT or the computational cost of Proof-of-Work.

Figure 10 achieves a balanced, comprehensive security posture rather than excelling in isolated metrics. Its dominance across automation, privacy, and scalability dimensions underscores its suitability for next-generation 6G vehicular cloud systems, where security solutions must simultaneously be intelligent, decentralized, privacy-aware, and real-time. This comparative visualization reinforces the quantitative results presented in earlier subsections and highlights the holistic advantage of integrating trust-aware AI, lightweight blockchain consensus, and autonomous mitigation into a unified security framework.

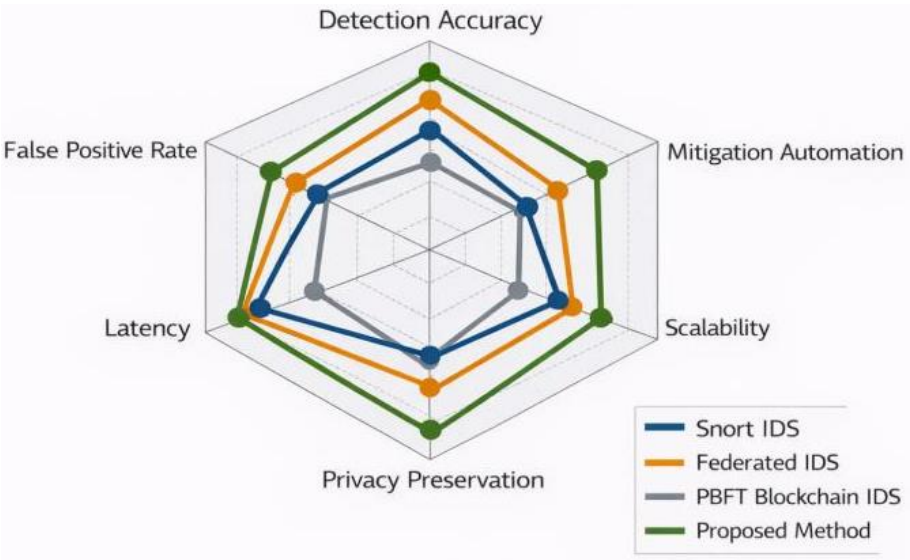

Fig. 10: Multi-Dimensional Security Comparison.

6. Conclusion

This study aimed to investigate whether a tightly integrated combination of dynamic trust evaluation, federated anomaly detection, and blockchain-based enforcement can yield measurable security and performance improvements in 6G-enabled vehicular cloud networks. The experimental results demonstrate that this objective was successfully achieved. Across all evaluated datasets and attack categories, the proposed methodology consistently achieved high detection accuracy while maintaining low false-positive rates, confirming that integrating trust-aware filtering with AI-based detection improves classification reliability in highly dynamic vehicular environments.

The results obtained from the CIC-IDS2017, TON_IoT, and N-BaIoT datasets show that the proposed system is not limited to a single threat model or traffic pattern. Instead, it performs robustly against volumetric attacks, stealthy low-rate intrusions, and coordinated botnet-driven behaviors. Detection rates remained above 95% across all datasets, while false-positive rates were kept close to or below 1%, a critical requirement for vehicular systems, where unnecessary isolation of benign nodes can negatively impact safety and service availability. These outcomes indicate that the trust evaluation mechanism effectively suppresses spurious alerts while still enabling timely identification of malicious behavior.

From a system-level perspective, the latency and scalability results provide strong evidence that decentralized security mechanisms can be deployed without violating 6G performance constraints. End-to-end response latency, measured from anomaly detection to mitigation enforcement, remained below 20 ms, even under high vehicular density, thereby satisfying the ultra-reliable low-latency communication requirements. The Proof-of-Trust consensus mechanism demonstrated stable validation performance as the number of validators increased, outperforming PBFT-based and Proof-of-Work schemes that exhibited rapid latency degradation. This confirms that the proposed consensus design is both computationally efficient and practically deployable in dense vehicular cloud environments.

The behavioral analysis of trust evolution further validates the methodological design. Experimental results showed that benign vehicles maintained stable trust values over time; transient anomalies led to temporary trust degradation followed by recovery, while persistently malicious nodes experienced monotonic trust decay culminating in isolation. This behavior confirms that the trust model avoids premature or irreversible penalties, thereby preserving system stability while still enforcing accountability—an essential property in high-mobility vehicular environments.

Additionally, the federated learning results demonstrate that collaborative intelligence can be achieved without requiring centralized data aggregation. Loss convergence was reached within a limited number of training rounds, and no instability was observed under heterogeneous and non-IID traffic conditions. This confirms that the proposed anomaly detection mechanism maintains detection effectiveness while preserving data privacy, addressing a key limitation identified in prior studies on federated vehicular security.

Finally, the mitigation and recovery analysis show that security enforcement is not merely reactive but operationally effective. Smart-contract-driven mitigation actions were executed automatically following on-chain validation, with recovery times consistently remaining low and requiring no human intervention. This capability ensures that attacks are not only detected but also contained and resolved in a timely and verifiable manner, significantly reducing the risk of prolonged service disruption.

Overall, the experimental results demonstrate that the proposed methodology achieves a balanced combination of detection accuracy, latency efficiency, scalability, privacy preservation, and automated enforcement. The findings provide concrete empirical evidence that decentralized, trust-aware, and AI-assisted security mechanisms can be jointly realized in 6G vehicular cloud systems. This work therefore delivers a validated, practically grounded security solution for next-generation intelligent transportation networks, while future research will focus on addressing real-world deployment considerations, such as energy efficiency, interoperability with legacy vehicular systems, regulatory constraints, and large-scale validation in emerging 6G vehicular environments.

References

[1] A. Dorri, M. Steger, S. S. Kanhere, and R. Jurdak, "Blockchain: A distributed solution to automotive security and privacy," *arXiv preprint arXiv:1704.00073*, 2017. https://doi.org/10.1109/MCOM.2017.1700879.

[2] R. Chen, J. Guo, and F. Li, "Blockchain-based trust management for Internet of Vehicles," *IEEE Access*, vol. 8, pp. 119 950–119 960, 2020.

[3] L. Xu, L. Chen, and X. Liu, "Lightweight blockchain consensus for high-speed vehicular networks," *IEEE Trans. Veh. Technol.*, vol. 72, no. 2, pp. 1775–1789, 2023.

[4] M. Rehman, H. Ahmad, and Z. Khan, "Hybrid PBFT-Proof of Trust for blockchain-enabled VANETs," *IEEE Access*, vol. 10, pp. 132 445–132 460, 2022.

[5] J. Kang, R. Yu, X. Huang, and Y. Zhang, "SDN-assisted blockchain for decentralized vehicular networks," *IEEE IoT J.*, vol. 8, no. 12, pp. 9823–9835, 2022.

[6] M. Giordani, M. Polese, A. Roy, D. Castor, and M. Zorzi, "A tutorial on 6G communications: Vision, enabling technologies, and new frontiers," *IEEE Commun. Surv. Tutor.*, vol. 23, no. 3, pp. 1681–1722, 2021.

[7] S. Li, K. Ota, and M. Dong, "Location privacy protection for Internet-of-Vehicles in 6G networks," *IEEE Network*, vol. 37, no. 1, pp. 55–61, 2023.

[8] X. Zhang, Y. Liu, and H. Zhao, "Adversarial data poisoning in 6G vehicular federated learning," *IEEE Trans. Inf. Forensics Secur.*, vol. 18, pp. 985–999, 2023.

[9] H. Lu, J. Li, and X. Liu, "A comprehensive survey on cyberattacks in vehicular networks," *Vehicular Communications*, vol. 35, 2024.

[10] Y. Dai, W. Chen, and Q. He, "Edge-intelligent intrusion detection for 6G vehicular IoT," *IEEE IoT J.*, vol. 11, no. 9, pp. 16534–16547, 2024.

[11] W. Yu, F. Xiao, and L. Zhou, "Blockchain-AI fusion for 6G vehicular trust management," *IEEE Trans. Intell. Transp. Syst.*, vol. 25, no. 4, pp. 4381–4395, 2024.

[12] "Framework of data privacy preservation and location obfuscation in vehicular cloud networks," *Concurrency Comput.: Pract. Exper.*, vol. 34, no. 5, e6682, 2022. https://doi.org/10.1002/cpe.6682.

[13] "Hilbert curves-based location privacy technique for vehicular cloud networks," *Cluster Computing*, vol. 27, pp. 2489–2504, 2024. https://doi.org/10.1007/s10586-023-04068-w.

[14] J. Jiang, X. Shen, and C. Lin, "Blockchain-enabled secure message verification for 6G vehicular networks," *IEEE Trans. Netw. Sci. Eng.*, vol. 11, no. 7, pp. 5043–5057, 2024.

[15] S. Mahmood, R. Iqbal, and N. Malik, "Privacy-preserving access control for 6G UAV-vehicular networks," *Vehicular Communications*, vol. 38, 2025.

[16] A. Al-Matari, S. Al-Gumaei, and H. Al-Qasem, "Blockchain-assisted spectrum sharing in 6G cognitive vehicular IoT," *IEEE Access*, vol. 12, pp. 190 450–190 467, 2024.

[17] I. Sharafaldin, A. H. Lashkari, and A. A. Ghorbani, "Toward generating a new intrusion detection dataset and intrusion traffic characterization," *Proc. Int. Conf. Inf. Syst. Security Privacy (ICISSP)*, 2018, pp. 108–116. https://doi.org/10.5220/0006639801080116.

[18] N. Moustafa, B. Turnbull, and K.-K. R. Choo, "An ensemble intrusion detection technique based on proposed statistical flow features for protecting network traffic of Internet of Things," *IEEE Internet Things J.*, vol. 6, no. 3, pp. 4815–4830, Jun. 2019. https://doi.org/10.1109/JIOT.2018.2871719.

[19] Y. Meidan *et al.*, "N-BaIoT: Network-based detection of IoT botnet attacks using deep autoencoders," *IEEE Pervasive Computing*, vol. 17, no. 3, pp. 12–22, Jul.–Sep. 2018. https://doi.org/10.1109/MPRV.2018.03367731.