A Multi-Tiered Hybrid Intrusion Detection Framework for Vehicular Networks: Integrating Supervised Ensembles and Unsupervised Clustering

Rahul A, Shivaraj M, Shriram S, Manohar M, Dr. Pushpa CN

Student, Student, Student, Student, Associate Professor
Department of Computer Science and Engineering,
University of Visvesvaraya College of Engineering, Bangalore, India

Abstract : The rapid proliferation of the Internet of Vehicles (IoV) has fundamentally revolutionized modern transportation. By integrating previously isolated intra-vehicle Controller Area Networks (CAN) with external Vehicle-to-Everything (V2X) communication infrastructures, modern vehicles are capable of advanced autonomous driving, dynamic routing, and cooperative traffic management. However, this unprecedented connectivity drastically expands the attack surface, exposing critical vehicle safety systems to highly sophisticated cyber threats. Traditional cryptographic countermeasures often violate the stringent realtime latency constraints inherent to vehicle networks, necessitating the deployment of intelligent, lightweight Intrusion Detection Systems (IDS). This paper presents an applied replication, extension, and structural evaluation of the Multi-Tiered Hybrid Intrusion Detection System (MTH-IDS) framework. The proposed architecture establishes a comprehensive hierarchical detection pipeline encompassing multi-layered supervised machine learning for binary anomaly detection and multi-class threat categorization, coupled with unsupervised clustering for zero-day threat discovery. Utilizing ensemble methods including Random Forests, Extra Trees, and Extreme Gradient Boosting alongside Mini-Batch K-Means clustering, the system is rigorously evaluated against the CICIDS2017 and CAN-Intrusion datasets. Experimental results demonstrate that the integrated stacking ensembles achieve near-perfect classification accuracies exceeding 99.9% for known threats, while maintaining execution latencies well below the mandatory 10-millisecond vehicular safety thresholds. This establishes the framework's viability for deployment within resource-constrained Electronic Control Units (ECUs).

IndexTerms **- Intrusion Detection System, Internet of Vehicles, Controller Area Network, Machine Learning, Ensemble Learning, Unsupervised Clustering, Random Forest, XGBoost, Cybersecurity**

I. INTRODUCTION

A. The Evolution of the Internet of Vehicles

The automotive industry is currently undergoing a historic paradigm shift, transitioning from conventional, mechanically isolated transportation units into highly intelligent, autonomous, and seamlessly interconnected systems via the Internet of Vehicles (IoV) [1]. The IoV architecture acts as the primary vehicular communication framework, facilitating highly reliable, high-speed data exchanges. This expansive communication ecosystem is broadly bifurcated into two primary network domains: the intra-vehicle network (IVN) and the external vehicular network.

Modern connected vehicles are equipped with an increasingly large array of Electronic Control Units (ECUs), often ranging from seventy to over one hundred discrete embedded systems per vehicle [2]. These ECUs govern a vast spectrum of both critical and non-critical functionalities, spanning engine timing control, anti-lock braking systems (ABS), autonomous steering, telemetry data aggregation, and passenger infotainment platforms. To coordinate these complex functions efficiently, ECUs communicate primarily via the Controller Area Network (CAN) bus. The CAN protocol is an international standard favored for its low manufacturing cost, minimal wiring complexity, high reliability in electromagnetically noisy environments, and robust fault-tolerance properties [3].

Simultaneously, the external vehicular network utilizes complex Vehicle-to-Everything (V2X) technologies. By leveraging cellular networks (such as 4G LTE and emerging 5G architectures), Dedicated Short-Range Communications (DSRC), and Wireless Local Area Networks (WLANs), V2X establishes a continuous bridge connecting the vehicle's internal embedded systems to external roadside infrastructure (V2I), smart city grids (V2G), pedestrian networks (V2P), and other vehicles in the immediate vicinity (V2V) [4]. While this profound integration unlocks unprecedented levels of traffic efficiency and passenger safety, it inadvertently bridges the gap between critical physical automotive actuators and the inherently hostile external cyber environment.

B. The Cybersecurity Conundrum in Vehicular Networks

The expanding connectivity of the IoV significantly amplifies the security risks associated with modern vehicles, transforming them from localized transit tools into high-value cyber targets. Cyber threats targeting vehicular networks can critically degrade the

stability and robustness of the IoV, potentially leading to vehicle unavailability, malicious remote hijacking, severe financial data breaches, and catastrophic physical traffic accidents.

The inherent vulnerabilities of vehicular networks stem from the foundational design of the legacy network protocols [5]. Within the IVN, the standard CAN bus relies on a broadcast-based communication strategy that entirely lacks built-in encryption, digital signatures, sender authentication, or secure priority arbitration schemes. This oversight renders the internal network highly susceptible to message injection, distributed denial-of-service, and hardware spoofing attacks once the perimeter is breached. In the external IoV network, the sheer volume of heterogeneous connected nodes dictates that every compromised smart traffic light, roadside unit, or unpatched mobile application serves as a potential vector for brute-force incursions, botnet propagation, and Man-in-theMiddle (MitM) interceptions [6].

Compounding these structural vulnerabilities is the reality that traditional information security mechanisms—such as heavy cryptographic encryption (e.g., RSA, AES-256) and rigorous multi-way authentication handshakes—are largely incompatible with intra-vehicle networks [1]. The computational overhead and transmission delays imposed by such cryptographic mechanisms violate the strict, deterministic timing constraints required by vehicular safety services. For autonomous and cooperative driving applications, V2X traffic safety requirements mandate a stringent processing latency bound of 10 to 20 milliseconds. Consequently, Intrusion Detection Systems (IDS) powered by sophisticated Machine Learning (ML) algorithms have emerged as an essential, lightweight defensive layer capable of identifying malicious traffic in real-time

C. Research Objectives and Paper Organization

This project aims to address the critical and rapidly evolving security requirements of the IoV by replicating, adapting, and structurally evaluating the Multi-Tiered Hybrid Intrusion Detection System (MTH-IDS) framework initially proposed by Yang et al. [1]. The primary objectives of this research are:

1) To systematically profile the structural vulnerabilities of both the CAN bus and external V2X networks.
2) To develop a robust ML-based intrusion detection pipeline capable of effectively identifying highly complex, overlapping cyber-attacks.
3) To mathematically model and implement diverse treebased ensemble algorithms (Random Forest, Extra Trees, XGBoost) and unsupervised clustering methods (Mini-Batch K-Means).
4) To evaluate the computational feasibility and real-time execution latencies of these algorithms within the strict constraints of vehicular network topologies.

The remainder of this paper is organized as follows: Section II reviews related work. Section III details the theoretical vulnerability analysis and threat modeling. Section IV presents the data preprocessing and feature engineering pipeline. Section V outlines the mathematical foundations and architecture of the multi-tiered machine learning framework. Section VI presents the experimental setup, validation strategies, and evaluation metrics. Section VII analyzes the experimental results, and Section VIII concludes the paper with directions for future research.

II. RELATED WORK

The application of Intrusion Detection Systems within vehicular networks has garnered significant academic attention, diverging primarily into signature-based and anomaly-based methodologies [8].

Signature-based IDSs operate by comparing incoming network packets against a static database of known threat signatures. While highly efficient and yielding virtually zero false positive rates for documented attacks, these systems are fundamentally incapable of identifying zero-day exploits or newly synthesized attack variants [9]. Conversely, anomaly based IDSs establish a statistical baseline of ”normal” network behavior and flag any significant deviations as potential intrusions.

Early research into vehicular anomaly detection heavily relied on statistical rule-based systems. For instance, Song et al. proposed an IDS analyzing the frequency and timing intervals of CAN messages, effectively detecting highfrequency Denial-of-Service (DoS) floods [10]. However, such simplistic timing analyses fail against sophisticated spoofing attacks where the adversary meticulously mimics legitimate message frequencies.

Consequently, recent literature has shifted toward supervised ML frameworks. Decision Trees (DT), Support Vector Machines (SVM), and Artificial Neural Networks (ANN) have been extensively deployed [11]. Kang et al. utilized a Deep Neural Network (DNN) structure to extract high-dimensional features from CAN packet payloads, achieving high detection accuracy but suffering from extensive computational training times and large memory footprints, making deployment on low-end ECUs impractical [12].

To balance the trade-off between detection accuracy and computational overhead, ensemble learning algorithms such as Random Forests (RF) and Gradient Boosting Machines (GBM) have emerged as optimal solutions for vehicular environments. However, a singular reliance on supervised ML algorithms leaves networks vulnerable to unknown, unlabelled attacks. To address this, the foundational MTH-IDS paper [1] proposed a hybrid architecture, merging supervised ensembles for known threat identification with unsupervised clustering to isolate zero-day anomalies—a methodology that forms the core of the applied replication in this project.

III. THREAT MODELING AND VULNERABILITY ANALYSIS

A comprehensive understanding of vehicular intrusion detection requires a rigorous analysis of the protocol-level vulnerabilities that plague underlying communication infrastructures.

A. Architectural Vulnerabilities of the CAN Bus

The Controller Area Network protocol utilizes differential voltage signaling across a tightly twisted pair of channels (CAN-High and CAN-Low). The standardized structure of a classical CAN data frame consists of several distinct fields, outlined in Table I.

table i: can data frame field structure and vulnerability analysis

Frame Field	Bit Size	Function & Exploitation Context
Start of Frame (SOF)	1 bit	Denotes the beginning of a transmission; vulnerable to synchronization disruption.
Arbitration ID	11 or 29 bits	Determines message priority. Attackers spoof low IDs (e.g., 0x000) to seize bus control continuously.
Control Field / Data Field	6 bits / 0 to 64 bits	Specifies the Data Length Code (DLC); can be manipulated to cause buffer overflows. Contains actual payload, the primary target for arbitrary malicious instruction injection.
CRC Field	16 bits	Cyclic Redundancy Check for error detection. Attackers easily recalculate valid CRCs for forged messages.
ACK Field	2 bits	Acknowledgment slot. Can be manipulated to simulate false network stability.
End of Frame (EOF)	7 bits	Delineates the frame conclusion.

The most critical vulnerability of this architecture lies within its Arbitration mechanism. This broadcast-oriented design gives rise to three primary intra-vehicle attack vectors:

- **Denial-of-Service (DoS) Attacks:** The attacker completely floods the CAN bus with massive volumes of highest-priority messages (ID 0x000).
- **Fuzzy Attacks:** Malicious actors rapidly inject entirely arbitrary messages containing randomly spoofed identifiers and chaotic payloads.
- **Impersonation and Spoofing Attacks:** A highly sophisticated attack wherein adversaries passively monitor the bus to identify the specific CAN IDs associated with targeted critical functions.

B. Threat Vectors in External Vehicular Networks

The external V2X architecture exposes the vehicle to the broader spectrum of conventional internet protocol (IP) based threats:

- **Distributed Denial-of-Service (DDoS):** Multiple compromised external nodes launch coordinated volumetric attacks against centralized IoV servers or specific vehicle IP addresses.
- **Brute-Force and Credential Stuffing:** Automated scripts attempt to crack weak, factory-default passwords protecting vehicle web interfaces.
- **Botnet Propagation:** Connected vehicles with unpatched infotainment systems are infected with self-propagating malware.
- **Web and Application Attacks:** Exploitations targeting the centralized IoV backend servers via Cross-Site Scripting (XSS) or SQL Injections.

IV. DATA PREPROCESSING AND FEATURE ENGINEERING

The efficacy of any complex machine learning model is strictly bounded by the mathematical quality, cleanliness, and representational accuracy of its training data matrix.

A. Dataset Profiling

1. **CICIDS2017 Dataset:** Explicitly engineered to provide a highly reliable benchmark for external network environments.
2. **CAN-Intrusion Dataset:** Represents physical layer communications captured directly via the OBD-II port of an actual test vehicle undergoing simulated cyber-attacks.

B. Data Cleansing and Missing Value Imputation

Raw network traffic matrices frequently contain mathematical anomalies resulting from capture errors or protocol malfunctions. The corrupted instances are systematically purged to prevent critical mathematical failures during the gradient descent optimization phases.

C. Label Encoding Paradigms

A 'LabelEncoder' is utilized to assign a unique, sequential integer to each distinct string category. For Tier 1 binary classification models, all distinct attack strings are aggregated into a singular positive integer representing an 'ATTACK' state, while benign traffic is mapped to zero.

D. Dimensionality Reduction and Feature Scaling

Vehicular network datasets are inherently high-dimensional. Standard-scaling (Z-score normalization) is utilized to ensure that large-magnitude features do not implicitly dominate the mathematical calculations of distance-based or gradient-based algorithms.

V. MULTI-TIERED DETECTION ARCHITECTURE AND MATHEMATICAL FOUNDATIONS

To effectively mitigate the complex, heterogeneous threat landscape, this project implements a hierarchical, multi-tiered architecture.

A. Tier 1: Anomaly Detection via Tree Ensembles

The first layer of the architecture operates as a rapid binary anomaly detection filter. This layer relies heavily on tree-based ensemble methods due to their parallel processing capabilities and high execution speeds.

1. **Decision Trees (DT):** The foundational building block of the architecture.
2. **Random Forest (RF):** To counteract the instability of singular trees, the Random Forest algorithm utilizes a rigorous 'bagging' (bootstrap aggregating) technique.
3. **Extra Trees (ET):** The Extremely Randomized Trees classifier pushes the Random Forest paradigm to its absolute limit, providing significantly faster computational training times.

B. Tier 2: Attack Family Classification via Gradient Boosting

Once a packet is mathematically flagged as malicious by the Tier 1 binary filters, it enters Tier 2, which functions as the primary multi-class attack classifier. Tier 2 employs the Extreme Gradient Boosting (XGBoost) algorithm.

C. Tier 3: Stacking Ensemble Methodology

To maximize overall system accuracy and eliminate isolated algorithmic blind spots, the project implements a rigorous Stacking architecture. The meta-classifier learns mathematically how to weight differing predictions dynamically.

D. Tier 4: Zero-Day Discovery via Unsupervised Clustering

To identify entirely novel, zero-day attacks, an unsupervised learning tier utilizing Mini-Batch K-Means clustering is incorporated. Mini-Batch K-Means utilizes small, random subsamples of the dataset at each iteration to update the centroids.

VI. EXPERIMENTAL SETUP AND EVALUATION METRICS

A. Development Environment and Hardware Simulation

The experimental validation of the proposed multi-tiered framework was executed within a Python 3.x based development environment utilizing standard, highly optimized data science libraries ('pandas', 'numpy', 'scikit-learn', 'xgboost').

B. Validation Strategy and Cross-Validation

To guarantee maximum algorithmic robustness and prevent data leakage, a comprehensive train-test split protocol was employed. The preprocessed datasets underwent a primary partition, allocating 70% of the data points for model training and setting aside an entirely mathematically isolated 30% hold-out test set for final evaluation.

C. Standardized Evaluation Metrics

The performance of the classification algorithms is rigorously quantified using an array of standardized metrics derived from the fundamental classification confusion matrix: Accuracy, Precision, Recall, and F1-Score.

VII. RESULTS AND DISCUSSION

A. Tier 1: Binary Anomaly Detection Efficacy

The implementation of the tree-based ensemble frameworks (Random Forest and Extra Trees) yielded near-perfect classification metrics on both the intra-vehicle and external network datasets. As illustrated in Table II, when evaluated on the structured, high-frequency physical layer CAN datasets, tree-based models routinely achieve accuracies exceeding 99.9%.

table ii: tier 1 performance evaluation (binary classification)

Algorithm	Dataset	Accuracy	Precision	F1-Score
Decision Tree	CAN-Intr.	0.9984	0.9981	0.9982
Random Forest	CAN-Intr.	0.9999	0.9998	0.9999
Extra Trees	CAN-Intr.	0.9999	0.9999	0.9999
XGBoost	CAN-Intr.	0.9996	0.9995	0.9996
Decision Tree	CICIDS17	0.9850	0.9810	0.9830
Random Forest	CICIDS17	0.9980	0.9970	0.9970
Extra Trees	CICIDS17	0.9990	0.9980	0.9980

B. Tier 2 & 3: Multi-Class Categorization and Stacking

In the complex, multi-class environment of Tier 2, the system must accurately categorize the specific family of attack to allow the vehicle's automated mitigation systems to execute the correct countermeasures. The XGBoost algorithm effectively handled the heterogeneous complexities of external V2X traffic.

table iii: tier 2 & 3 performance (multi-class categorization)

Algorithm	Precision	Recall	F1-Score	FPR
XGBoost (Standalone)	0.9950	0.9970	0.9960	0.0031
Random Forest	0.9920	0.9940	0.9930	0.0058
Extra Trees	0.9940	0.9950	0.9945	0.0042
MTH-IDS Stacking	0.9991	0.9993	0.9992	0.0008

C. Computational Complexity and Latency Analysis

The practical deployment of any machine learning IDS within an IoV environment depends entirely on strict adherence to low-latency processing constraints. Empirical benchmarks of the proposed Stacking architecture executing on simulated ECU-equivalent hardware consistently demonstrate average inference execution latencies of approximately 0.5 to 1.2 milliseconds per packet payload. This blazing processing speed comfortably clears the stringent 10-millisecond threshold required by autonomous vehicular safety applications.

VIII. CONCLUSION AND FUTURE SCOPE

The profound, expanding connectivity enabling the modern Internet of Vehicles necessitates a parallel, rapid evolution in embedded cybersecurity protocols. The exhaustive analysis of network traffic architectures and CAN protocol vulnerabilities confirms that traditional cryptographic defenses are insufficient. Advanced, low-latency, machine learning-driven Intrusion Detection Systems stand as a non-negotiable imperative to protect physical passenger safety and broader network integrity.

ACKNOWLEDGMENT

The authors express their sincere gratitude to Dr. Pushpa C N, Associate Professor, Department of CSE, UVCE, for her continuous guidance and technical expertise. We also thank Dr. Manjula S H, Chairperson, Department of CSE, and Dr. Subhasish Tripathy, Principal, UVCE, for their support.

REFERENCES

1. L. Yang, A. Moubayed, and A. Shami, "MTH-IDS: A Multi-Tiered Hybrid Intrusion Detection System for Internet of Vehicles," IEEE Internet of Things Journal, vol. 9, no. 1, pp. 616-632, Jan. 2022.
2. J. Petit and S. E. Shladover, "Potential cyberattacks on automated vehicles," IEEE Transactions on Intelligent Transportation Systems, vol. 16, no. 2, pp. 546-556, April 2015.
3. A. Taylor, N. Japkowicz, and S. Leblanc, "Frequency-based anomaly detection for the automotive CAN bus," in Proc. 2015 World Congress on Industrial Control Systems Security (WCICSS), 2015, pp. 45-49.
4. S. Zeadally, J. Guerrero, and J. Contreras, "A tutorial survey on vehicle-to-everything (V2X) communications," Telecommunication Systems, vol. 73, pp. 469-489, 2020.
5. K. Koscher et al., "Experimental security analysis of a modern automobile," in Proc. 2010 IEEE Symposium on Security and Privacy, Oakland, CA, USA, 2010, pp. 447-462.
6. M. Tomlinson, J. Bryans, S. A. Shaikh, and A. Kalutarage, "Detection of MAC layer spoofing and denial of service attacks in micro aerial vehicle networks," Computers & Security, vol. 77, pp. 175–194, 2018.
7. A. Loukas, G. Karopoulos, S. Kambourakis, "Cyber-physical attacks on vehicular networks," IEEE Communications Magazine, vol. 55, no. 7, pp. 104–110, 2017.
8. S. Hansman and R. Hunt, "A taxonomy of network and computer attacks," Computers & Security, vol. 24, no. 1, pp. 31–43, 2005.
9. M. Ring, S. Wunderlich, D. Scheuring, D. Landes, and A. Hotho, "A survey of network-based intrusion detection data sets," Computers & Security, vol. 86, pp. 147–167, 2019.
10. H. M. Song, H. R. Kim, and H. K. Kim, "Intrusion detection system based on the analysis of time intervals of CAN messages for in-vehicle network," in Proc. Int. Conf. Inf. Netw. (ICOIN), 2016, pp. 63–68.
11. M. J. Kang and J. W. Kang, "Intrusion detection system using deep neural network for in-vehicle network security," PLoS One, vol. 11, no. 6, 2016, Art. no. e0155781.
12. E. Seo, H. M. Song, and H. K. Kim, "GIDS: GAN based intrusion detection system for in-vehicle network," in Proc. 16th Annu. Conf. Privacy, Secur. Trust (PST), 2018, pp. 1–6.
13. P. S. Murvay and B. Groza, "Source identification using signal characteristics in controller area networks," IEEE Signal Processing Letters, vol. 21, no. 4, pp. 395–399, April 2014.
14. M. E. Hoque, S. R. Chowdhury, and M. R. Rahman, "Machine learning based intrusion detection system for vehicular ad hoc networks," in Proc. 15th Int. Conf. Comput. Inf. Technol. (ICCIT), 2012, pp. 1–6.
15. I. Sharafaldin, A. H. Lashkari, and A. A. Ghorbani, "Toward generating a new intrusion detection dataset and intrusion traffic characterization," in Proc. 4th Int. Conf. Inf. Syst. Secur. Privacy (ICISSP), Portugal, Jan. 2018, pp. 108–116.
16. L. Breiman, "Random forests," Machine Learning, vol. 45, no. 1, pp. 5–32, 2001.
17. P. Geurts, D. Ernst, and L. Wehenkel, "Extremely randomized trees," Machine Learning, vol. 63, no. 1, pp. 3–42, 2006.
18. T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining, 2016, pp. 785–794.