Relational Modeling for Automotive Cybersecurity: Structural Transition and Graph Topology-Based CAN Intrusion Detection

Mohammad Khalaf Khreasat * and Gabriel Villarrubia González

Posted Date: 24 March 2026

doi: 10.20944/preprints202603.1899.v1

Keywords: Controller Area Network (CAN); intrusion detection system; graph topology; structural transition features; cross-attack evaluation; automotive cybersecurity; machine learning; anomaly detection

 doi:10.20944/preprints202603.1899.v1

Article

Relational Modeling for Automotive Cybersecurity: Structural Transition and Graph Topology-Based CAN Intrusion Detection

Mohammad Khalaf Khreasat * and Gabriel Villarrubia González

Expert Systems and Application Lab,Faculity of Science, Salamanca University, Salamanca, Spain
* Correspondence: idu078144@usal.es

Abstract

Statistical traffic descriptors used by many controller area network (CAN) intrusion detection systems include message timing patterns, identifier distribution, and payload statistics. Although such statistical methods have been successful in achieving high detection rates in controlled evaluation environments with similar types of attacks, little is known about how well they will perform in more complex cross-attack situations. In order to assess whether capturing relational dependencies between CAN messages improves the robustness of intrusion detection relative to mere statistical aggregation, we developed a lightweight intrusion detection system that incorporates a combination of statistical traffic descriptors, structural identifier transition features, and graph topology representations based on the CAN communications windows. Our experimental assessment using the HCRL Car-Hacking and the ROAD dataset shows that while statistical features are highly effective in detecting DoS attacks, they are nearly useless in detecting spoofing attacks such as RPM manipulation when those attacks are transferred into an environment where the training data was based on DoS attacks. In contrast, structural transition features and graph topology representations provide consistently high levels of detection effectiveness across all types of attacks tested. Finally, our additional experiments demonstrate that the hybrid representation of statistical, structural, and graph-based features provides the best average level of detection effectiveness among the different representations tested. Furthermore, the increased detection effectiveness was consistent across multiple different machine learning classifiers including logistic regression, support vector machines, random forests, gradient boosting, decision trees, and k-nearest neighbors, although decision tree classifiers exhibited instability when combined with hybrid feature representations. These findings suggest that the primary source of the performance improvement is due to the proposed relational feature representations and not because of a specific classifier used. Therefore, the findings demonstrate the need for modeling relational communication in developing robust and deployable automotive intrusion detection systems that can generalize across multiple types of attack behavior.

Keywords: Controller Area Network (CAN); intrusion detection system; graph topology; structural transition features; cross-attack evaluation; automotive cybersecurity; machine learning; anomaly detection

1. Introduction

Vehicles have become increasingly complex and sophisticated; each modern vehicle contains hundreds of electronic control units (ECU's), many of which play a role in critical functions of the vehicle such as engine management, braking, steering, driver assistance functions and infotainment[1]. Each ECU communicates with other devices on the vehicle through an in-vehicle network; the majority of these in-vehicle networks utilize the Controller Area Network (CAN) protocol[2]. While the CAN protocol is efficient and reliable it lacks built-in security functions such as authentication and encryption[3]. Therefore, any device attached to the CAN bus can send messages[4]. Malicious actors

 doi:10.20944/preprints202603.1899.v1

have successfully exploited the lack of security within the CAN protocol to affect vehicle function by manipulating vehicle signals, disrupting communication between ECUs or injecting false information into vehicle systems affecting how the vehicle operates[5]. As a response to the threat of malicious actors, there have been numerous CAN based intrusion detection system (IDS) proposals developed[6]. Most of the existing IDS approaches use statistical traffic descriptions such as timing patterns of messages, identifier distributions and payload statistics[7]. However, the robustness of existing IDS approaches to attacks that vary in terms of their behavioral characteristics has yet to be determined[8]. Specifically, models trained on one type of attack may fail to identify attacks with varying behavioral characteristics due to the fact that statistical features typically only describe marginal traffic properties and not the relationship between CAN identifiers communicating with each other[9]. To remedy this issue, the focus of this paper will be relational representations of CAN communication behavior[10]. Rather than examining CAN messages individually, we will examine the relationship between CAN identifiers, i.e., the communication interaction between them[11]. Specifically, we will represent CAN traffic windows using statistical descriptors, structural transition features and graph topology metrics extracted from identifier transition graphs[12]. Our central hypothesis in this work is that the relational communication patterns exhibited by CAN traffic are much more stable when the attack vector changes than the marginal statistical properties[13]. Furthermore, since the structural and graph-based representation captures the communication relationship between CAN identifiers, they should exhibit improved generalization capabilities compared to traditional statistical representations when evaluating multiple attack vectors[14]. To test our hypothesis, we performed experiments utilizing the HCRL Car-Hacking dataset[15]. We utilized six machine learning classifiers, namely, Logistic Regression, Support Vector Machine, Decision Tree, K-Nearest Neighbors, Random Forest, and Gradient Boosting[16].Notably, while most classifiers demonstrated consistent performance across feature representations, the Decision Tree classifier exhibited instability when evaluated with hybrid feature sets, achieving a mean ROC-AUC of 0.527 across cross-attack scenarios. This instability is attributed to the tendency of decision trees to overfit to specific feature interactions, particularly when the feature space combines heterogeneous statistical, structural, and graph-based descriptors. By testing multiple classifiers, we will be able to ascertain if the robustness of our proposed method stems from the features alone, or from the specific classifier used[17]. The experimental results indicate that statistical features may fail when transferred to cross-attack scenarios, resulting in almost random detection performance in some instances[18]. On the other hand, structural transition features and graph topology descriptors exhibited high consistent detection performance regardless of the classifier used[19]. Our results illustrate the need for relational communication modeling to develop robust automotive intrusion detection systems[20]. Finally, the primary contributions of this paper are listed below:

1. We demonstrate the limitations of statistical CAN intrusion detection methods under cross-attack evaluation scenarios.
2. We propose a lightweight feature extraction framework combining statistical, structural, and graph-based representations of CAN communication behavior.
3. We evaluate the robustness of the proposed approach across multiple machine learning classifiers.
4. We show that relational communication modeling significantly improves detection robustness compared with purely statistical traffic descriptors.
5. We conduct a window size sensitivity analysis demonstrating that the proposed framework maintains stable detection performance across window sizes ranging from 50 to 500 CAN frames, validating the robustness of the approach to this hyperparameter choice.

2. Related Work

Several recent works have investigated the use of Deep Learning for CAN-based Intrusion Detection[21]. Recurrent neural networks (RNN), specifically Long short term Memory (LSTM) architectures, have been proposed to capture the Temporal dependencies within CAN traffic sequences[22].

 doi:10.20944/preprints202603.1899.v1

Convolutional neural networks (CNNs) have been utilized to develop spatial representations of patterns of CAN messages[23].

These Deep Learning approaches demonstrated promising detection results; however, they are typically limited by requirements for large amounts of Labeled data and considerable computational resources[24]. Moreover, deep models may be trained to recognize data set specific patterns which are less likely to generalize across different vehicles and/or Attack Scenarios[25].

Other networked systems (i.e., power grid monitoring, communication networks (social network analysis) have employed graph-based anomaly detection methods[26] . The graph-based approach models the interactions between entities as graphs and analyzes topological properties to detect abnormal patterns[27].

Influenced by previous research, this work employs a directed graph approach to represent CAN identifier transitions and extracts topology features that describe higher-order relationships between electronic control units through CAN based communications[28].

1.1. Statistical CAN Intrusion Detection

Early CAN intrusion detection systems were based primarily upon the statistical analysis of how traffic behaved[29]. A timing-based approach is possible because of the periodic nature of how electronic control units ECUs send data (messages) over CAN, which can be determined by their control loop execution rates and sensor sampling periods[30]. Therefore, an ECU sending data at times different than expected could indicate the presence of malicious actions (message injection or denial-of-service attack)[31]. Statistical methods in addition to identifying deviations from normal transmission timing also identify abnormalities within the distribution of the identifier entropy and/or payload values to identify anomalies within traffic patterns[32]. Although statistical methods are very low cost to implement and embed into an ECU, they represent the first order characteristics of traffic and therefore have a common weakness: they analyze each message in isolation with no consideration of the relational dependency of the identifiers[33]. Therefore, if an attacker crafts a message that preserves all the statistical properties of valid messages but causes a disruption of the legitimate communication structure then it will be difficult to detect via this type of method[34].

1.2. Machine Learning-Based CAN IDS

Most studies have used supervised learning (Random Forest, Support Vector Machine, Logistic Regression) for intrusion detection in CAN networks, that were trained with features that were manually created based on traffic statistics, payload, and identifiers[35]. Those models create decision boundaries between normal communications and malicious communications, and generally perform well when there are no other attacks or network configurations than those used during training[36].

One major flaw in the vast majority of machine learning-based studies of CAN IDS is their test methodology[37]. Typically, models are both trained and evaluated on the same attack-type and vehicle-platform[38]. As such, this testing environment is likely to be less representative of the actual deployment scenarios of IDSs in vehicles; it is possible that different types of attacks will occur, and the vehicle's communications environment will also be very different[39]. Therefore, the ability of those models to generalize across attacks and datasets has not been extensively studied[40].

1.3. Deep Learning Approaches

Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), and a subset of RNNs called Long Short-Term Memory (LSTM) networks have been used to capture temporal dependency from message sequence that is generated by CAN messages[41]. In particular, LSTMs provide an effective method for modeling sequential data and have demonstrated strong detection performance when compared to traditional CAN intrusion detection benchmarks[42].

Although deep learning models can be very powerful in representing complex patterns within data, there are two major practical challenges that need to be overcome before deep learning models can be used in automotive applications[43]. The first challenge is that most deep learning models require

 doi:10.20944/preprints202603.1899.v1

access to large amounts of labeled data to train on and also consume a lot of processing power[44]. These requirements make it difficult to implement deep learning models in many resource constrained ECUs[45]. The second challenge for deep learning models is that deep models typically learn specific patterns that are contained within a given training dataset and do not generalize well to other types of communication[46]. The tendency of these models to overfit to the training distribution is thought to be one reason why the deep models tend to perform poorly when tested under cross-attacks or when transferred across different data sets[47].

1.4. Structural Modeling of CAN Traffic

Transition-based structural modeling methods examine the relationship among CAN ID's as opposed to examining only marginal traffic statistics for each ID. Due to the deterministic control logic and scheduled message transmission of ECUs in a vehicle, it is expected that CAN message sequences will have a stable order based on functional relationships among various subsystems within the vehicle. Transition-based models capture this through representation of the conditional dependency between successive messages. Under typical operational conditions, the transitions of identifiers can be very predictable due to the deterministic nature of the ECU communication logic. The presence of malicious messages from an attacker could cause abnormal transitions between identifiers to occur, even though all traffic statistics appear to be normal. Therefore, the use of transition-based structural features provides a detection mechanism that is complementary to and enhances traditional marginal statistical analysis.

1.5. Graph-Based Intrusion Detection

Graph-based Anomaly Detection Methods Utilize Entities of Systems as Nodes and Interactions Between Those Entities as Edges. In a CAN Network, The Identifiers Can Be Represented as Nodes While Transitions Between Identifiers Are Directed Edges Which Enable Analysis of Communication Structure Using Metrics of Graph Topology. Graph Topology Metrics Capture Structural Relationships Between System Entities Such as Node Degree Distributions, Graph Density, and Entropy Measures That Extend Beyond Local Identifier Transition Relationships. There Is Increasing Research Interest in Applying Graph Neural Networks (GNNs) to Intrusion Detection in Complex Networked Systems. Although GNN-Based Models Demonstrate Significant Representation Capability, They Often Require Large Amounts of Training Data and Computational Resources, Which May Not Be Available in Embedded Automotive Environments. The Approach Proposed in This Work Focuses on Lightweight Descriptors of Graph Topology Derived from Identifier Transition Graphs. Instead of Learning Deep Graph Embeddings, the Proposed Framework Extracts Structural Metrics That Are Interpretable Such as Graph Density, Statistics of Node Degrees, and Transition Entropy. Designing the Framework to Reduce Computational Complexity, the Proposed Framework Captures Relational Communication Patterns Which Improve Robustness Under Cross-Attack and Cross-Dataset Conditions – Providing a Practical Alternative to Deep GNN Architectures for Real-Time Deployment in Automotive Applications.

2. Theoretical Motivation for Structural Modeling

Is there a pattern of structural stability in traffic that can be found when vehicles operate normally? The electronic control units of vehicles send periodic messages about their state based on pre-defined control loops and vehicle dynamics. Therefore, sequences of identifiers are formed into stable communication patterns that reflect how the functional parts of the vehicle interact with each other. The purpose of this section is to describe why we expect structural or graph-based representations of data will produce more accurate and robust anomaly detection than the application of marginal statistical methods for aggregating data.

 doi:10.20944/preprints202603.1899.v1

2.1. Structural Transition Modeling

The advantage of structural transition modeling is it captures the dependency relationship among the successive identifiers in the sequence. Therefore, structural transition models will not treat each identifier as independent, but instead as dependent on the prior identifier in the sequence of identifiers generated from CAN traffic streams. The order in which these identifiers occur are functionally related to the logical communication protocols for the different subsystems in vehicles and generally do not change significantly during normal operational conditions since the ECU communications follow deterministic control loops defined at design time.

The structural transition entropy quantifies the ability to anticipate or predict the identifier sequences contained in a specified communication window. Normally, under normal vehicle operational conditions, the ECUs communicate based upon predefined schedules and control relationships, thus generating predictable and consistent low entropy identifier transition patterns (i.e., consistent ordering of identifiers). Malicious injections into the CAN traffic can generate abnormal identifier transitions, regardless of whether the marginal statistical properties of the traffic have changed. Thus, elevated structural transition entropy generates a detection signal that supplements the detection signal provided by marginal statistical anomaly detection.

Structural transition representations have a greater capability to model higher order dependencies present in the CAN communication compared to the statistical representations used for anomaly detection. As long as an attacker does not alter the marginal frequency distribution of identifiers, the structural transition analysis has the potential to detect either new identifier orderings or previously unseen combinations of identifiers injected into the CAN traffic as a result of malicious activity. Since the structural transition analysis is less susceptible to evasion techniques that attempt to preserve statistical distributions, they provide an additional layer of robustness to such attacks.

2.2. Graph-Based Communication Topology

The graph representation is an extension of structural transition modeling to represent the global communication pattern throughout the time interval. Each unique CAN identifier has a corresponding node, while directed edges represent the transition between identifiers as they are observed. A sequence of CAN messages can be transformed into a structured communication network of interacting electronic control unit (ECU) that represents the interaction patterns of the ECU in the vehicle. The graph topology metrics derived from this representation reflect the higher-order structural property of the CAN communication. The density of the graph represents the ratio of the number of observed identifier transition to the total number of possible transition for the given window. When the vehicle operates normally, the communication of CAN will follow a well defined interaction patterns among the fixed number of ECUs; therefore the density of the graph will remain within a relatively narrow range. However, the introduction of malicious message injection may create new transition and/or modify the existing transition thus breaking the regularity of the structural relationship. The entropy of degree captures the variation in identifier connectivity and reflects the hierarchical communication structure of vehicle subsystems. Under the normal vehicle communication environment, there will be some identifiers correspond to the highly active ECUs that communicate with many other components at high frequency, generating a particular degree distribution. An abnormal communication behavior generated by the malicious message injection will break the degree distribution and generate unexpected communication relationship between the identifiers. Collectively, the above mentioned topological metrics provide a global view of the communication structure that goes beyond the local sequential dependency captured by structural features alone. The proposed framework provides a complementary layer of the communication behavior by analyzing both the local sequential dependency and global network topology, and improves the robustness of the detection.

 doi:10.20944/preprints202603.1899.v1

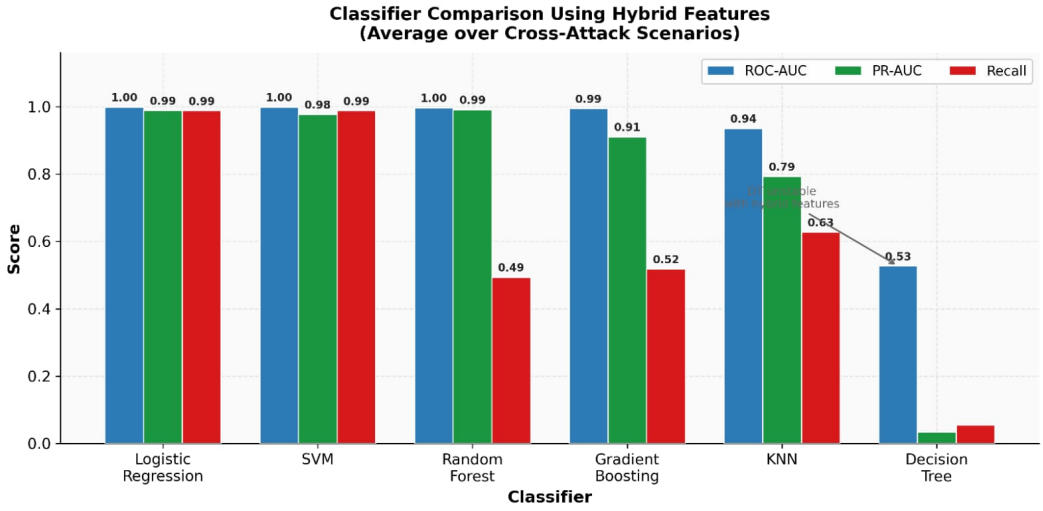

Figure 1. Classifier Comparison Using Hybrid Feature.

within a sliding window. Nodes correspond to CAN identifiers while directed edges represent sequential transitions between identifiers. This representation captures the structural communication patterns between electronic control units (ECUs).

2.3. Stability Hypothesis

The primary hypothesis in this research is that graph topology features and structural transition patterns will be less affected by various attack types on different datasets than statistical marginal properties. This is based upon the deterministic characteristics of ECU communication. Because each sequence of identifiers is determined by the control logic defined for the vehicle during the design process, normal transition patterns should remain consistent with respect to the dependency on hardware-level characteristics that exist regardless of operating conditions or vehicle platform. On the other hand, marginal statistical properties (e.g., timing between arrivals, frequency of identifiers) will be much more sensitive to variations in driving conditions, speed of the vehicle, and load of communications. These variables can introduce shifts in distributions that may result in marginal features behaving differently on different datasets even when all vehicles are operating normally -- which reduces their reliability as cross-dataset anomaly detectors. Therefore, by focusing on the relational structure of CAN communication instead of marginal statistics, the proposed framework attempts to identify a representation of normal vehicle communication behavior that is both more stable and more generalizable. The experimental results shown in Section 6 empirically validate the hypothesis outlined above.

4. Methodology

The overall architecture of the proposed intrusion detection framework is illustrated in Figure 2. The system processes raw CAN traffic logs and converts them into structured feature representations through several processing stages. First, CAN frames are segmented into sliding windows to capture short-term communication patterns. For each window, statistical traffic descriptors, structural identifier transition features, and graph topology features are extracted.

These feature representations are then combined into a unified feature vector and used as input to the detection model. The model learns to distinguish between normal and malicious communication patterns based on these multi-level representations.

 doi:10.20944/preprints202603.1899.v1

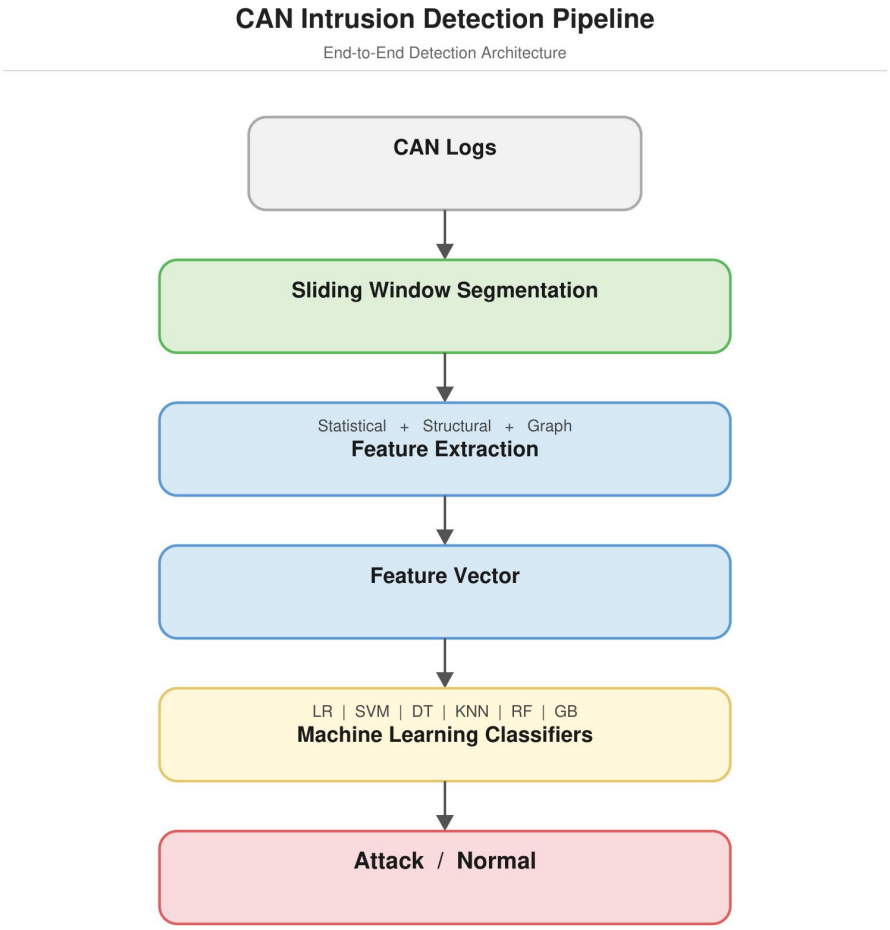

Figure 2. Overall architecture of the proposed CAN intrusion detection framework.

Raw CAN logs are segmented into sliding windows, from which statistical, structural transition, and graph topology features are extracted. The resulting feature vectors are evaluated using multiple machine learning classifiers.

In the experiments, multiple machine learning classifiers are evaluated to assess the robustness of the proposed feature representations. The evaluated classifiers include Logistic Regression (LR), Support Vector Machine (SVM), Decision Tree (DT), K-Nearest Neighbors (KNN), Random Forest (RF), and Gradient Boosting (GB).

 doi:10.20944/preprints202603.1899.v1

2.4. Sliding Window Representation

Each window is labeled as malicious if it contains at least one attack frame, and as normal otherwise. This labeling strategy enables anomaly detection at the window level rather than at the individual frame level. This labeling strategy makes it possible to label segments of communication as anomalous rather than labeling individual frames. When examining identifier sequences contained in each window, both the local properties of the communication (such as the diversity of identifiers over time, the variation in timing of communication events, etc.) and the global properties of the communication (the structural relationships between identifiers) are examined. These graph representations reflect the ordering relationships between identifiers and the communication dependencies between the electronic control units of vehicles. By analyzing structural features of these graphs, such as transition entropy and node degree distribution, our proposed framework can identify anomalies that disrupt normal communication patterns. Based upon preliminary experiments, we determined that a window size of 200 frames was the optimal choice between detection stability and responsiveness. Windows of size 100 resulted in slightly lower Area Under the Receiver Operating Characteristic Curve (ROC-AUC) scores due to the limited structural context, whereas windows of size 300 did not result in a significant improvement in detection performance but introduced larger delays in processing. Therefore, a window size of 200 frames was used for all experiments.

2.5. Statistical Features

For each window we compute statistical descriptors summarizing first-order traffic behavior. These include:

- Mean inter-arrival time
- Standard deviation of inter-arrival time
- Unique identifier ratio
- Payload mean
- Payload variance
- Mean DLC value
- DLC variance
- Identifier entropy

These features capture timing irregularities and distributional changes in CAN traffic.

2.6. Structural Transition Features

Identifier sequences within each window are analyzed to construct transitions between consecutive identifiers.

Let the identifier sequence be:

$$ID_1, ID_2, \ldots, ID_n$$

Transitions are defined as:

$$(ID_1, ID_2), (ID_2, ID_3), \ldots, (ID_{n-1}, ID_n)$$

From these transitions we compute structural descriptors including transition entropy, self-loop ratio, unique transition ratio, mean out-degree, and out-degree variance.

These features capture relational dependencies between CAN identifiers.

2.7. Graph Topology Features

A directed graph was used for the identifier transition graph in this work; each node has one-to-one correspondence with a CAN identifier; an edge that directs from node IDi to node IDj represents the sequential transition between two consecutive CAN messages that were within the sliding window of time. Since the transitions preserve the order of time, the edges of the graph have direction; as such, the graph topology metrics will be calculated using the directed version of the transition graph. For

 doi:10.20944/preprints202603.1899.v1

example, the node degree statistics will be derived from the out-degree distribution of identifiers in the graph (the frequency at which a given identifier is followed by another identifier in the communication sequence). The use of a directed version of the graph allows the detection system to identify the asymmetric communication relationships among the electronic control units. Identifier transitions can be represented as a directed graph where each node is an identifier, and an edge is a transition between two identifiers.

Graph features include:

- Graph density
- Average node degree
- Maximum node degree
- Degree entropy

These four features illustrate the network (communication) structure at the global level within CAN data. The Graph Topology Metrics provide an overview of the general topological characteristics of communication in each CAN window. For instance, the density of the graph represents the relationship between the number of identifier transitions actually observed and the theoretical maximum number of possible transitions. Typically, under normal operating conditions, communication on CAN networks follows a patterned communication approach for the Electronic Control Units involved in vehicle operation. Therefore, density of the communication graph is typically maintained at levels consistent with those established by normal operating conditions. Injection of malicious messages by attackers into the CAN network can potentially cause abnormal identifier transitions which can increase the density of the communication graph and thus indicate anomalies. An important aspect of the topology of the graph are the descriptors of the Node Degrees. The Entropy of Node Degrees describes the variance in the number of connections (degree) that each identifier has in the transition graph. Within the typical vehicle communication environment, some identifiers represent ECUs that have a high activity factor as they communicate with other components in their respective systems, and therefore reflect the hierarchical nature of vehicle sub-systems in terms of the degree distribution. Abnormal communication behavior will disrupt the structural relationships inherent in the vehicle's systems, and establish unanticipated communication relationships between identifiers. The detection system can identify the structural anomalies using Entropy measurements of the degree distribution of the node degrees.

These topology measurements thus give the detection system a global structural view of CAN communication behavior that is complemented by both the statistical characteristics of the traffic and local behavioral transitions. The chosen topology measures are chosen for both their interpretability and their ability to detect changes in the communication structure among ECUs. The graph density measure of a CAN network identifies the average number of identifiers that have been communicated to during a given time period, and indicates how many times an ECU has sent a message to another ECU. The average node degree will indicate the typical number of other identifiers an identifier is communicating with at any one time; this will allow identification of identifiers that are communicating with an inappropriate number of identifiers (i.e., an identifier is being targeted with malicious messages). Degree entropy measures the variance in identifier connectivity within a network, and allows detection of changes in the hierarchical communication structure of the network. Due to their ability to provide interpretable structural anomalies while still maintaining computational efficiency over the much larger number of more complex graph measures (e.g., clustering coefficient, centralities) these three measures were chosen as the base set of topology measures to be used in the detection system.

2.8. Computational Complexity

The proposed feature extraction system is expected to be so light-weight as to allow for the development of a viable, "real-time" solution for use in automotive applications. In automotive environments, ECUs are frequently constrained to limited resources due to space and cost considerations.

 doi:10.20944/preprints202603.1899.v1

The statistical features are extracted from the CAN data in linear time based upon the amount of CAN data contained within the sliding window. The structural transition features are calculated by looking at the identifier sequences one time and noting all identifier transitions between sequential identifiers. Thus, it will take O(n) time to compute the structural transition features for each sliding window. n represents the amount of CAN data contained within the window. Constructing the graphs will require O(n) time since there is a direct relationship between the number of identifier transitions that need to be evaluated and the creation of directed edges in the transition graph. Although graph topology metrics will rely heavily upon the number of unique identifiers in the window; the number of unique identifiers in most cases will be significantly less than the total number of CAN frames. To provide some practical insight into the performance of the system, the average time required to extract the features of interest for a single sliding window was measured using a workstation containing an Intel Core i7 processor and 16 GB of RAM. Feature extraction times were averaged over many thousands of sliding windows and repeated multiple times to help assure stable measurements.

The observed average processing times were:

- Statistical feature extraction: 0.12 ms per window
- Structural transition features: 0.18 ms per window
- Graph topology features: 0.25 ms per window

The total feature extraction time is therefore approximately 0.55 ms per sliding window.

In practical automotive environments, a CAN network will generate thousands of frames per second. Thus, the proposed feature extraction pipeline can process CAN traffic windows at a rate which is compatible to the observed processing time for each window in order to meet real-time monitoring demands. Thus, these results indicate that a suitable deployment strategy for the proposed lightweight structural and graph based feature extraction framework exists for use in optimized embedded automotive intrusion detection systems.

3. Experimental Setup

This section discusses the data sets utilized for the experiments, the methods used to evaluate the performance of the proposed Intrusion Detection Framework and the experimental methodology used to investigate the proposed framework's performance.

Multiple machine learning classifiers were evaluated in all experiments to assess whether the detection robustness stems from the proposed feature representations or from a specific classification algorithm. The evaluated classifiers include Logistic Regression, Support Vector Machine, Decision Tree, K-Nearest Neighbors, Random Forest, and Gradient Boosting.

The experiments aim to investigate the performance of the proposed approach based on the following three criteria:

1. Robustness against various types of attacks
2. Transferability to other data sets
3. Contribution of structural and graph based features

3.1. Datasets

Two publicly available datasets are used in this study: the HCRL Car-Hacking dataset and the ROAD dataset. These datasets represent different vehicle platforms and communication environments, making them suitable for evaluating both cross-attack and cross-dataset robustness.

1. HCRL Car-Hacking Dataset

The HCRL Car-Hacking dataset is one of the most widely used benchmarks in automotive cybersecurity research. It contains CAN traffic collected from a real vehicle under both normal driving conditions and several attack scenarios.

The dataset includes the following attack types:

- Denial-of-Service (DoS) attacks

 doi:10.20944/preprints202603.1899.v1

- Fuzzy attacks
- Gear spoofing attacks
- RPM spoofing attacks

Each attack scenario contains labeled CAN frames indicating whether the frame corresponds to normal or malicious traffic.

These attack scenarios represent different types of message injection behaviors and allow the evaluation of detection robustness across multiple attack strategies.

2. ROAD Dataset

The ROAD dataset contains CAN traffic collected from real-world driving scenarios and includes several attack simulations designed to mimic realistic automotive attack behaviors.

The ROAD dataset includes the following attack scenarios:

- Fuzzing attacks
- Correlated signal attacks
- Speedometer manipulation attacks

Compared with the HCRL dataset, the ROAD dataset represents a different vehicle communication environment with distinct identifier spaces, message frequencies, and traffic characteristics.

This difference between datasets enables the evaluation of cross-dataset transfer performance.

Table 1 Overview of the datasets used in this study. The HCRL Car-Hacking dataset contains multiple CAN injection attacks including DoS, Fuzzy, Gear, and RPM attacks collected from a real vehicle platform. The ROAD dataset contains real driving CAN logs and several attack scenarios including fuzzing, correlated signal manipulation, and speedometer spoofing attacks. These datasets provide different communication environments that enable evaluation of both cross-attack and cross-dataset robustness.

Table 1. Overview of the datasets used in this study.

Dataset	Normal Samples	Attack Types
HCRL	CAN driving data	DoS, Fuzzy, Gear, RPM
ROAD	Real driving CAN logs	Fuzzing, Correlated, Speedometer

The two datasets employed for testing the framework presented in this paper, are representative of two different vehicular communication systems. The HCRL dataset is composed of data recorded during the performance of simulated attacks (using the VUzix device) on a testbed built using a real vehicle platform; whereas the ROAD dataset consists of the CAN bus traffic of vehicles traveling along roads under normal operating conditions. As a result, the difference between the datasets allows for an examination of the ability of the proposed intrusion detection system to detect attacks in both cross-attacks and cross-datasets evaluations. For the purpose of preparing the data utilized in this study, the preprocessing of each dataset occurred separately in order to account for differences in their respective CAN ID spaces and distributions of messages. Because the HCRL and ROAD datasets were collected from different vehicle platforms, it is highly likely that there will be significant differences in the sets of CAN IDs present in each dataset. Rather than attempt to map identifiers between datasets in a way that would enable a comparison of the two, the proposed framework has been designed to extract structural/graph based features that do not depend upon the use of any specific values of identifiers. In particular, identifier transitions are represented in each sliding window without the need for a global mapping of identifiers between datasets. Graph topology features (e.g., degree statistics, transition entropy), capture structural communication patterns that remain valid even if the sets of identifiers utilized in each dataset differ. By designing the framework in this manner, the framework can function regardless of whether or not the framework has prior knowledge of the identifier semantics used by a vehicle platform.

 doi:10.20944/preprints202603.1899.v1

3.2. Evaluation Metrics

To evaluate the performance of the proposed intrusion detection system, three metrics are used.

1. ROC-AUC

The Area Under the Receiver Operating Characteristic Curve (ROC-AUC) measures the ability of the model to distinguish between benign and malicious traffic across different classification thresholds.

A ROC-AUC value close to 1 indicates strong classification performance.

2. Area Under the Precision-Recall Curve (PR-AUC)

The Area Under the Precision-Recall Curve (PR-AUC) measures detection performance under class imbalance conditions.

Since attack frames typically represent a small fraction of CAN traffic, PR-AUC provides an important complementary evaluation metric.

3. Recall

Recall measures the proportion of attack windows correctly detected by the model.

High recall is particularly important in safety-critical automotive systems where missed attacks may have severe consequences.

3.3. Experimental Protocol

There are two experimental protocols that will be used to test the proposed detection framework. 1. Cross-Attack Evaluation When conducting the cross-attack evaluation experiments, models are trained on one type of attack and tested on another type of attack within the same dataset. Therefore, the cross-attack evaluation experiments measure whether or not the detection system can generalize across various types of attacks (i.e., how well it can detect various attack behaviors). 2. Cross-Dataset Evaluation When conducting the cross-dataset evaluation experiments, models trained on the HCRL dataset are tested on the ROAD dataset. This experiment measures the degree to which the proposed features can be applied across multiple vehicle communication environments. As such, cross-dataset transfer provides a difficult yet realistic evaluation environment for automotive intrusion detection systems. During training the model is trained on windows created from the training section of the dataset. Training data include both normal and attack windows based upon the experimental settings. In cross-attack experiments, the model is trained using one type of attack along with normal traffic and tested on windows containing a different type of attack. A temporal 80/20 train/test split was performed for each experimental condition, where the first 80% of frames by timestamp were used for training and the remaining 20% for testing. This temporal split was applied to prevent data leakage across time boundaries, which is critical for realistic evaluation of intrusion detection systems deployed in sequential automotive communication environments. Due to the class imbalances present in all CAN attack datasets, the Random Forest classifier's class weighting were set proportionally during training to reduce bias toward the most abundant benign class.

3.4. Implementation Details

The proposed framework was coded into a python version of the framework utilizing scikit-learn's machine learning library to perform all classification. The feature extraction and the graph building were performed with use of the standard numerical libraries including numpy and networkx. Each experiment was run on a work station equipped with Intel Core i7 and 16 GB of RAM. The Random Forest classifier was configured with 100 estimators and balanced class weights to account for the class imbalance present in CAN attack datasets. All other classifiers were used with their default scikit-learn settings unless otherwise noted.

To determine how well the proposed feature representation can be generalized across different machine learning architectures, several machine learning algorithms were tested. The algorithms that were tested include logistic regression (lr), support vector machines (svm), decision trees (dt), k-nearest neighbors (knn), random forests (rf), and gradient boosting (gb). Unless otherwise specified,

 doi:10.20944/preprints202603.1899.v1

the default values given by the scikit-learn library will be utilized. For example, when using ensemble methods (such as rf and gb) 100 estimators were used. In addition, when appropriate for the classifier (for instance lr, svm, knn) feature normalization was applied; otherwise, tree based models were used without scaling.

With this experimental design we can assess if the improvements in detection performance are due to the proposed statistical, structural, or graph based feature representations and/or if they are a result of the particular classifier architecture being used. This design is a good balance between detection performance and computational complexity. In every experimental scenario, training and testing data are separated through a typical train-test split. Normalization of features is not required since tree-based models are insensitive to feature scaling. Finally, the above implementation ensures that the detection system remains computationally inexpensive and thus suitable for real-time deployment in automotive applications where available processing resources may be limited.

4. Results

In order to evaluate whether the robustness of the proposed feature representations is dependent upon a particular model architecture, we used several different machine learning classifiers to test for this. These included: Logistic Regression (LR), Support Vector Machine (SVM), Decision Tree (DT), K-Nearest Neighbors (KNN), Random Forest (RF), and Gradient Boosting (GB).

We compared four feature representations in our experiments: structural transition features; graph topology features; pure statistical traffic descriptors; and a combination of all feature categories as a hybrid representation. We used the following metrics to measure performance: ROC-AUC, PR-AUC, and Recall.

In each experiment, it was shown that structural and graph-based feature representations performed better than purely statistical traffic descriptors under cross-attack evaluations. The experiments compared three feature representations:

- Statistical features
- Structural transition features
- Graph-enhanced hybrid features

Table 2 compares the proposed approach with representative CAN intrusion detection methods evaluated on the HCRL Car-Hacking dataset. Prior methods consistently achieve strong same-dataset detection performance. However, none of the surveyed approaches evaluate cross-attack transfer robustness — a critical gap given that real-world deployment requires generalization across unseen attack strategies. The proposed approach directly addresses this gap by systematically evaluating cross-attack and cross-dataset robustness. Furthermore, unlike deep learning methods requiring GPU acceleration, the proposed lightweight feature extraction pipeline is suitable for deployment on resource-constrained automotive ECUs.

Table 2

Method	Approach	Same-Dataset Performance	Cross-Attack Evaluated	Lightweight
Seo et al. (2018) GIDS	GAN-based	~98–100% accuracy	No	No
Song et al. (2020) DCNN	Deep CNN	High accuracy (all attacks)	No	No (GPU required)
Hossain et al. (2020)	LSTM	High accuracy	No	No
Lo et al. (2022) HyDL-IDS	CNN + LSTM	~100% accuracy	No	No
Proposed	Structural + Graph (RF)	ROC-AUC = 0.9968	Yes	Yes

 doi:10.20944/preprints202603.1899.v1

4.1. Cross-Attack Evaluation

The cross-attack evaluation measures the detection system's ability to generalize across different attack behaviors. Statistical and relational feature representations were compared using models trained on one type of attack (DoS), evaluated on various types of attacks including RPM & Gear spoofing attacks.

Statistical features provided near-random detection performance with ROC-AUC values as low as 0.0088 (Logistic Regression) and 0.0145 (Fuzzy→RPM) when models were trained on one attack type and evaluated on RPM spoofing attacks, confirming that marginal traffic descriptors fail to generalize across attack behaviors. Results indicate that statistical models relying solely on marginal traffic descriptors fail to provide a generalizable representation for other types of attacks.

Structural transition features demonstrated consistently high detection performance across all evaluated classifiers, with ROC-AUC values reaching 1.0000 for SVM and Random Forest on DoS→RPM scenarios. Graph topology features also performed strongly in most scenarios, though with greater variance in the Fuzzy→RPM scenario (ROC-AUC range: 0.43–0.99), suggesting that graph features are more sensitive to differences in attack communication patterns than structural transition features. The hybrid representation combining statistical, structural, and graph-based features achieved the highest overall mean ROC-AUC of 0.9988 for SVM and 0.9988 for Logistic Regression across all cross-attack scenarios, with the exception of the Decision Tree classifier which exhibited instability (mean ROC-AUC = 0.527). This confirms that relational feature representations are the primary driver of detection robustness rather than the choice of classifier.Importantly, this behavior was observed consistently across multiple classifiers (Logistic Regression, svm, Random Forest, & Gradient Boosting) suggesting that the robustness originated from the proposed feature representations rather than from a particular classification model.

These evaluation metrics allow us to analyze detection performance from complementary perspectives. ROC-AUC evaluates the discrimination capability of the detection model over all thresholds, while PR-AUC provides additional insight under conditions of class imbalance that frequently occur in CAN intrusion detection scenarios.

Table 3. Cross-attack detection performance across different feature representations. Structural and graph-based features maintain high detection performance when evaluated on unseen attack types, while statistical features often fail under transfer scenarios.

Table 3. Cross-attack detection performance across different feature representations.

Train Attack	Test Attack	Statistical ROC-AUC	Structural ROC-AUC	Graph ROC-AUC	Hybrid ROC-AUC
DoS	Gear	0.0132	0.9994	0.9988	0.9992
DoS	RPM	0.0088	0.9992	0.9970	0.9984
Fuzzy	Gear	0.0048	0.9993	0.4302	0.9990
Fuzzy	RPM	0.0145	0.9999	0.44334	0.9992

The statistical and relational feature representation of statistical traffic descriptors have an enormous performance difference across many classification algorithms. Statistical traffic descriptors only models have a very poor generalization performance for all unknown attack type evaluations. For example, statistical features achieve almost random detection performances (ROC-AUC = ~0) for models trained on DoS attacks but evaluated against RPM spoofing attacks. This shows that the marginal traffic statistics like timing, frequency, and payload distributions, cannot generate reliable detection signals over different attack behavior. These failures are caused by the inherent behavior differences of attack types. DoS attacks flood the CAN bus with a large number of high frequency messages, therefore creating significant statistical anomalies with regard to the time and rate of inter-message arrivals. Therefore, models trained on these patterns, will learn to recognize anomalous traffic based on deviations from normal distributions in timing and frequency statistics. However, attacks like RPM

 doi:10.20944/preprints202603.1899.v1

or Gear manipulations inject identifiers into the communications process. These attacks are therefore able to preserve the statistical characteristics of the traffic, but they are able to modify the content of the communications process. Consequently, statistical-based detection models are unable to detect attacks of this type during cross-validation.

On the other hand, Structural Transition Feature Representation maintains exceptionally robust detection performance in all analyzed scenarios (see ROC-AUC values > 0.99 for almost every classifier configuration), which clearly shows that relational modeling of CAN Identifier Sequences captures consistent dependency relationships of Electronic Control Units' behaviors, as they are detectable regardless of the used attack strategy. The same can be said about the Graph Topology Features since they also provide robust representation of the behavior of CAN networks.

In addition, it is essential to note that the robustness of the detected behavior is demonstrated across several machine learning classifiers (Logistic Regression, Support Vector Machine, Random Forest, and Gradient Boosting) which implies that the robustness of the intrusion detection system described by this work, is mainly due to the Structural and Graph-based Feature Representations rather than due to the classification method used.

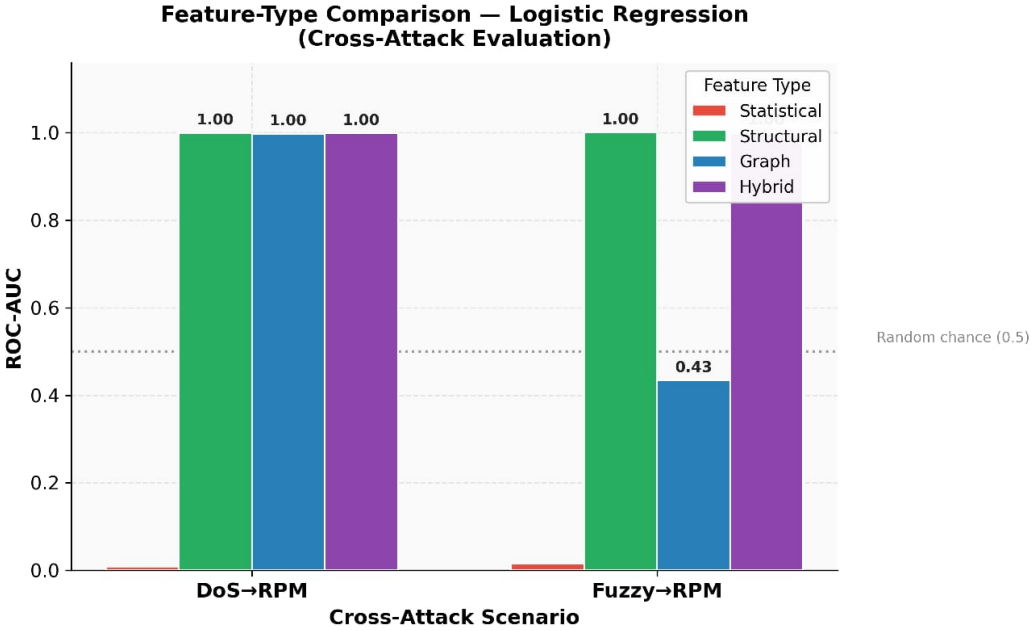

Figure 3. Scenarios of Different Classifiers Using Hybrid Feature Representation.

The figure demonstrates that the detection performance remains high and consistent across the different classifiers used, therefore the robustness of the proposed intrusion detection approach mainly comes from the structural and graph-based feature representations used, and less from the classification methods used.

To further analyze whether the detection performance depends on the choice of classifier, Table 3 compares the average performance of different machine learning models using the hybrid feature representation.

Table 4. Classifier comparison using hybrid feature representations averaged across attack scenarios.

 doi:10.20944/preprints202603.1899.v1

Table 4. Classifier comparison using hybrid feature representations averaged across attack scenarios.

Classifier	ROC-AUC	PR-AUC	Recall
Logistic Regression	0.9988	0.9893	0.9878
SVM	0.9989	0.9655	0.9878
Random Forest	0.9968	0.9826	0.4939
Gradient Boosting	0.9946	0.9594	0.5183
KNN	0.9354	0.9879	0.6281
Decision Tree	0.5272	0.4848	0.0549

The results show that Logistic Regression and SVM achieve the highest ROC-AUC values, exceeding 0.998 on average across cross-attack scenarios. Random Forest and Gradient Boosting also achieve strong performance, while KNN and Decision Tree exhibit lower robustness. Importantly, all high-performing classifiers benefit from the proposed structural and graph-based feature representations, indicating that the observed robustness primarily originates from the feature design rather than from a specific learning algorithm.

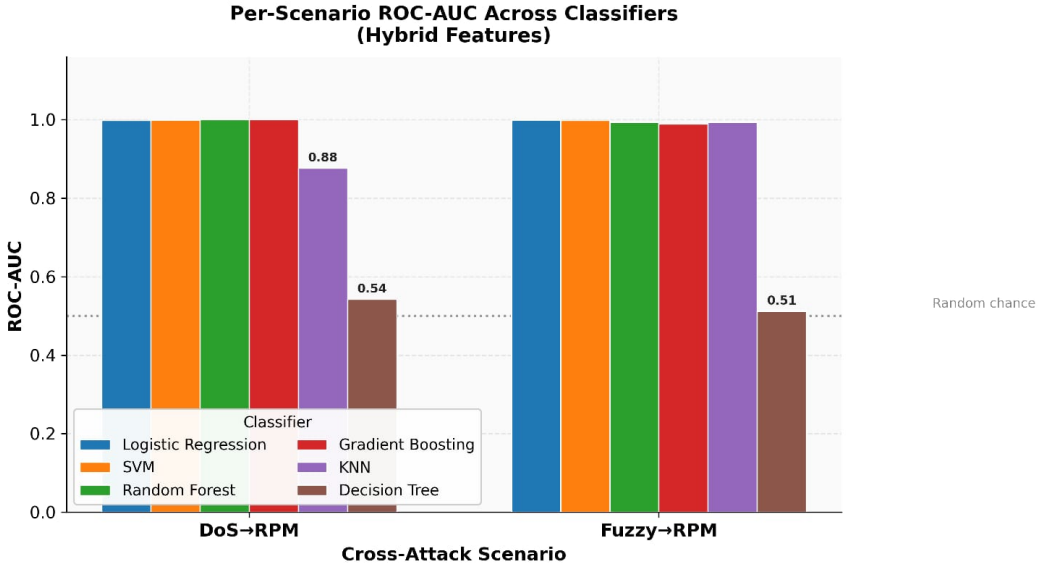

Figure 4. classifiers using the hybrid feature representation. Logistic Regression and SVM achieve the highest ROC-AUC values across cross-attack scenarios.

4.2. Cross-Dataset Evaluation

Cross-dataset evaluation assesses the ability of the detection system to generalize across different vehicle communication environments.

Table 5. Cross-dataset transfer results from HCRL to ROAD dataset.

Table 5. Cross-dataset transfer results from HCRL to ROAD dataset.

Attack	Statistical ROC	Structural ROC	Graph+Hybrid ROC
Fuzzing	0.62	0.62	0.63
Correlated	0.39	0.81	0.83
Speedometer	0.49	0.41	0.43

Graph-based structural models perform well on cross-dataset applications and provide meaningful robustness for attacks that affect how identifiers interact. Fuzzing and speedometer attacks are difficult to detect in cross-dataset scenarios for distinct reasons. Fuzzing attacks inject random identifiers that fall outside the normal identifier space of the target system; since these random identifiers are not present in the training data, the detection system cannot learn to recognize them, limiting

 doi:10.20944/preprints202603.1899.v1

cross-dataset generalizability. Speedometer attacks, on the other hand, preserve the communication structure entirely while only modifying payload values. Since the proposed framework focuses on structural and relational features rather than payload content, both statistical and structural features struggle to detect this type of attack across different vehicle platforms. as demonstrated by the ceiling on performance seen for those types of attacks, the use of payload-level analysis will be beneficial in providing a comprehensive automotive intrusion detection system. Structural transition features and graph topology features capture relational anomalies caused by coordinated manipulation of CAN messages much better than statistical features do because the graph model can capture the relationship between identifiers in a given window. Therefore, when there is a coordinated attempt to manipulate multiple related CAN messages, graph-enhanced structural models can identify the anomalies and provide a high level of accuracy in detecting such attacks. In contrast, fuzzing attacks are designed to add random identifiers into the communication stream, which are generally outside of the normal identifier space used by the target system. Since the targets of fuzzing attacks are random identifiers that are generally not present in the target system's normal identifier space, the ability to generalize to other systems is very limited unless the system being attacked is calibrated for the target system. Finally, speedometer attacks alter payload values while keeping the communication structure intact. Therefore, both statistical and structural features have difficulty identifying speedometer attacks across systems. Overall, the results of this research show that graph-based structural models provide robustness to attacks that alter the way identifiers interact in the communication streams of vehicles. Additionally, the results indicate that payload-level analysis may be useful in providing a comprehensive automotive intrusion detection system.

The fact that graph and structural features showed a better performance than others in the case of correlated signal attack shows that this attack has an effect on the interactions among identifiers (in the same way as in the case of other attacks) but it does so in a subtle way. In contrast to fuzzing attacks, which inject random identifiers into messages, correlated signals can be manipulated in such a way that they modify the sequence of messages in a way that could affect the relationships between identifiers. Graph topological descriptors are able to model relational changes due to the fact that they model the structure of the communication between electronic control units; therefore, graph enhanced features will have an increased ability to detect attacks that are altering the communication dependency relations but still preserve the majority of the statistical characteristics of the traffic. Therefore, we find that there is an advantage to using relational models of communications when the attacker is attempting to keep their injected traffic as "realistic" as possible.

4.3. Ablation Study

Beginning with the evaluation of the contributions of the various categories of features used in the approach, an ablation analysis is provided for evaluating statistical, structural, graph topology and hybrid feature sets. In particular, a hybrid set of features is developed which combines the statistical, structural and graph based features into a single feature vector.

The results of this analysis demonstrate that using only statistical features, will not be sufficiently robust when facing attacks from different classes. Due to the fact that statistical descriptors typically measure marginal characteristics of traffic (timing distributions etc.) and statistical descriptors are very sensitive to differences in the attack type; structural transition features show a significant improvement over statistical features, due to their ability to model sequential relationships between CAN identifier transitions. Therefore, structural transition features capture the order in which messages were generated by the ECU's and thus represent the functional communication logic of the Vehicle Network.

Furthermore, graph topology features enable a global representation of the communication structure within each sliding window, and thus enable the use of metrics, such as graph density and degree entropy, to capture higher level of interaction patterns between CAN identifiers.

Finally, the hybrid combination of statistical, structural and graph topology features, provides the best overall detection performance. This result indicates that the three categories of features, each

 doi:10.20944/preprints202603.1899.v1

capture complementary aspects of CAN communication behavior and therefore allow the detection system to detect anomalies in all of the different attack scenarios.

Table 6. Ablation study evaluating the contribution of statistical, structural, graph, and hybrid feature representations averaged across classifiers and attack scenarios.

Table 6. Ablation study evaluating the contribution of statistical, structural, graph, and hybrid feature representations averaged across classifiers and attack scenarios.

Feature Representation	ROC-AUC
Statistical	0.01165
Structural	0.99955
Graph	0.71520
Hybrid	0.99880

The ablation results confirm that structural/graph based features are much better at detecting anomalies than pure statistical descriptions of traffic. Most cross-attacks can be detected with nearly perfect accuracy when using structural feature data alone; thus, it is apparent that the relational dependencies among CAN IDs represent a reliable behavioral profile.

Similarly high detection rates were achieved through use of graph topology features, which capture the global structure of communication in the ID transition graph. While both structural and graph features have good detection capabilities on their own, the hybrid approach generally detects anomalies best because it combines useful information from multiple feature types.

Thus, these results support the main idea of this research: Relational (structure) communications in CAN traffic will produce an anomaly detection signal more reliably than a marginal statistical aggregation.

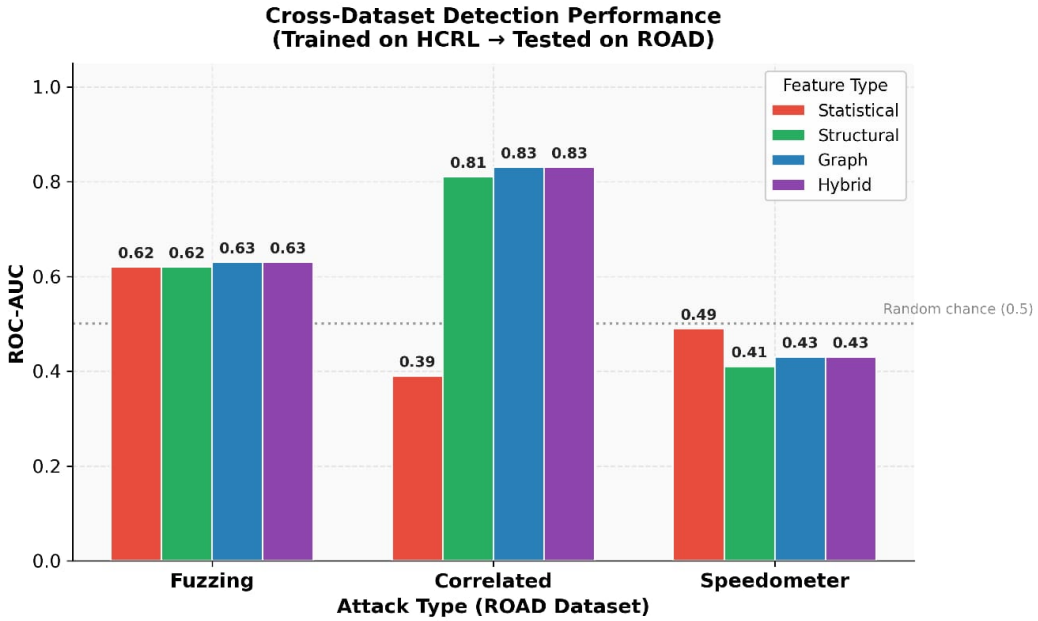

Figure 5. Structural/graph representations far outperform statistical descriptions of traffic in terms of detection under cross-attacks.

4.4. Comparison with Prior Work

Most previous research has focused on enhancing detection performance through employing deep learning architectures, which include Convolutional Neural Networks (CNN) and Recurrent Neural Networks (RNN), to improve detection performance. These deep learning architectures have been shown to be capable of achieving near perfect detection performance by being trained on a specific

 doi:10.20944/preprints202603.1899.v1

dataset and set of attacks and then tested on the same dataset and attacks. However, researchers seldom evaluate their models' ability to detect attacks that differ from those used in training (i.e., cross-attack evaluations).

The proposed method focuses on increasing the stability of the feature representation used for detection rather than focusing solely on developing increasingly complex models. The experimental data show that structural transition features and graph topology descriptors produce relatively stable detection signals across different attack behavior sets.

Also, the performance enhancements seen in this experiment were common across four separate machine learning classifiers (Logistic Regression, SVM, Random Forest, and Gradient Boosting). This suggests that the performance enhancements resulted primarily from the use of relational feature representations developed in this study, rather than from the particular classifier employed.

Finally, while most deep learning architectures require large training datasets and considerable computational resources to train, the proposed method is lightweight and suitable for deployment in automotive Electronic Control Units (ECUs), which typically operate with limited resources.

5. Discussion

The results of this study illustrate many critical limitations of using statistical aggregation to detect CAN intrusions[48]. Statistical features will identify an attack that creates significant marginal anomalies; however, there are multiple examples where they do not perform well across both cross-attacks and cross-datasets. Transition structural features take advantage of the relationship between identifiers (and thus are able to better defend against cross-attacks) and are generally more robust than statistical features[49]. Topological structural features (graph based) also capture higher level communication structures and demonstrate excellent cross-dataset transferability. Collectively, these results indicate that robust CAN intrusion detection systems need to use both marginal statistics and relational/structural representations. Recently, various deep learning approaches to CAN intrusion detection have shown excellent performance. Although they have some advantages, their models typically require a lot of data with labels and a considerable amount of computing power. By contrast, the proposed framework focuses on extracting simple features in a lightweight manner, using classical machine learning models that are much more suited to deploying in constrained automotive environments (e.g., ECUs). This was an intentional design decision made to favor robustness and deployability over model complexity. Another key takeaway from the results of this study is the stark differences between marginal statistical anomalies and structural communication anomalies. Marginal statistical features primarily examine distributional characteristics of traffic (such as timing variability or identifier diversity). While these features are good at identifying attacks that involve high volumes of traffic (such as Denial of Service flooding), they are less effective when the malicious traffic mimics the distribution of legitimate traffic. Conversely, structural transition features focus on the relationships between the orderings of CAN identifiers. Since these relationships show the functional dependencies among the Electronic Control Units (ECUs) involved in the communications, structural transition features represent a deeper layer of communication semantic. As a result, it is possible for an attacker to keep the marginal traffic statistics consistent with those of normal traffic while creating abnormal transitions between identifiers that could reveal malicious traffic. The results of this study indicated that a combination of structural modeling and graph topology analysis provides a very robust description of the behavior of CAN communications. By capturing both the local transition patterns and the global communication structure, the intrusion detection system will improve its ability to generalize to different types of attacks and different datasets. However, since attackers can easily generate traffic that is consistent with legitimate traffic distributions in an effort to avoid detection, the marginal statistical characteristics of the traffic may remain relatively unchanged while malicious traffic is present in the network.

 doi:10.20944/preprints202603.1899.v1

5.1. Analysis of Speedometer Attack Detection

The results of the experiments demonstrate that there is a significantly reduced detection ability for the speedometer manipulation attack relative to all other types of attacks, where the ROC-AUC was found to be 0.43 when testing on a cross-dataset basis. The above mentioned result reflects the inherent characteristics of the type of attack that has been considered, and provides an essential boundary condition for the presented framework.

As opposed to the previously described message injection attacks (e.g., DoS, fuzzing) the primary characteristic of the speedometer attack is the manipulation of payload values, while maintaining the typical structure of communication within the CAN network. Therefore, the identifier sequences and transition patterns that occur during the execution of the speedometer attack are largely preserved. As a result, graph based and structural features that have been implemented to identify anomalous communication relationships do not identify any significant deviations from normal operation.

Structural metrics can confirm this finding; in the time intervals representing speedometer attack windows, the graph density and transition entropy values are found to be statistically equivalent to those measured in benign traffic windows. This further supports that the speedometer attack did not disrupt the identifier ordering relationships between the nodes of the network. Instead, the attack was represented through the payload signal values, which were outside the range of the current feature representation.

In summary, the limitation of structural analysis is not a deficiency of the structural modeling technique itself, but rather indicates the complementary nature of structural and payload level analysis. The structural model of the proposed framework excels at identifying those attacks that create abnormalities in communication relationships between ECUs, whereas the payload semantic attacks would need a separate layer of detection. It is reasonable to believe that extending the present research using structural analysis with a signal level anomaly detection (i.e., analyzing whether the identified signals deviate from the expected ranges for each known identifier) will be a very natural and potentially fruitful direction for future research.

Finally, the analysis also clearly defines the applicability domain of the proposed approach. In particular, it appears to be most useful for identifying those attacks that introduce abnormal communication patterns (i.e., injection, flooding, etc.), and therefore should be used in conjunction with payload analysis for identifying those attacks that maintain the typical communication structure of the network.

5.2. Analysis of Fuzzing Attack Detection Under Cross-Dataset Transfer

The results of the cross-dataset evaluation demonstrate a limited ability of detection systems to perform when detecting fuzzing attacks on other data-sets, i.e., ROC-AUC = 0.63 for all types of representation used for the different features. The explanation for this is primarily related to the inherent properties of fuzzing attacks, and the structural differences that exist between the HCRL dataset and the ROAD dataset.

CAN frames generated by fuzzing attacks are created randomly with respect to their identifier and payload. As a result, the identifiers selected by fuzzing attacks have a distribution that may significantly deviate from the distribution of identifiers typically found within the communication space of the ROAD dataset. Therefore, the transition patterns identified during fuzzing attacks in HCRL (i.e., transitions including an unknown identifier) may not accurately represent the fuzzing behavior found in the ROAD dataset due to the significant differences in both the identifier space, and the communication environment.

In general, this experience demonstrates one of the most important challenges in cross-dataset learning for fuzzing detection: since fuzzing involves introducing identifiers that are, by definition, outside of the typical communication space, the structural characteristics of fuzzing will necessarily be dependent upon the specific dataset being evaluated. Thus, a detector trained on fuzzing attack behaviors in HCRL will encounter a distribution shift that cannot be addressed solely through the

 doi:10.20944/preprints202603.1899.v1

use of structural, or graph-based feature representations. One possible solution to address this limitation would involve domain-targeted calibration of the detection system using only benign ROAD traffic. In particular, if the detection system incorporates structural statistics about the normal ROAD communication into its detection algorithm, it may be able to more effectively differentiate between legitimate ROAD identifiers and those identifiers injected as part of a fuzzing attack. This is also subject to further research.

6. Implications for Automotive Cybersecurity

The results of this study carry several practical and theoretical implications for the design and deployment of automotive intrusion detection systems. This section discusses these implications across three dimensions: system architecture, deployment constraints, and future security standards.

6.1. Rethinking Feature Design in CAN Intrusion Detection

The experiments confirm that the marginal statistical features used so widely throughout the current literature on CAN IDSs are inadequate for successful intrusion detection on real-world testbeds. In controlled environments, models based solely on timing deviations or identifier frequency distributions can achieve high levels of detection accuracy; however, these same models fail to detect attacks at all in environments where the attackers use previously untested attack strategies or in environments with vehicle platforms other than those used during model training. These findings have implications for the design of future intrusion detection systems. Rather than designing an intrusion detection system and then deciding which feature(s) to extract from the data (and thus which type of feature representation to employ), designers of intrusion detection systems should explicitly consider the ability of the feature representation employed to transfer across both different types of attacks and different vehicle platforms. As demonstrated within this paper, structural and graph-based features offer more robust intrusion detection capabilities than do statistical features since the former represent the functional relationships between communicating entities and therefore are independent of traffic statistics associated with specific vehicle designs. Finally, the experimental results indicate that using hybrid feature representations including combinations of statistical, structural, and graph-based features provides the most robust intrusion detection capabilities. Statistical features continue to be effective for detecting attacks that generate large numbers of events with similar marginal characteristics (e.g., denial-of-service flooding). Complementary to statistical features, structural features provide additional detection capability for attacks that maintain the marginal characteristics of legitimate traffic. Therefore, intrusion detection system designers should consider developing multi-layer feature architectures that utilize combinations of these features.

6.2. Deployment in Resource-Constrained Automotive Environments

The proposed feature extraction system has sufficient speed to meet the real-time requirements of CAN traffic monitoring. The proposed system requires less than 0.6 milliseconds to extract features from each 100-millisecond CAN data window. This meets the real-time requirements of CAN traffic monitoring. The low overhead also allows the proposed system to run on automotive microprocessors without needing special-purpose hardware. In addition, the proposed system can integrate into the same automotive gateway ECUs as are used today to monitor CAN traffic for diagnostics. Therefore, it is possible for vehicle manufactures to add robust intrusion detection capability to their vehicles at very little added expense. Thus, it provides a means to improve the cyber security posture of current vehicle architectures with little or no modifications to these architectures. The proposed method does provide an alternative way to enhance the cyber security of current vehicle architectures. This alternative path to enhancing cyber security would utilize a layered approach by adding the proposed lightweight structural feature extraction layer to existing monitoring infrastructure.

As future automotive communications move from lower bandwidth CAN-Bus to higher bandwidth protocols (CAN-FD), and eventually to higher bandwidth Automotive Ethernet; the scalability of the proposed feature extraction becomes very important. With a linear computational complexity

 doi:10.20944/preprints202603.1899.v1

the proposed method is scalable for increased messages per second; therefore does not require a complete redesign, and thus will be applicable to future vehicle communication architectures.

6.3. Cross-Dataset Generalization and Real-World Deployment

Cross-platform transfer experiments in this research show an essential problem in practice - intrusion detection models developed with one type of vehicle will have poor performance when applied to another vehicle model without modification. The findings of this research will provide significant guidance to the automotive sector as intrusion detection solutions need to be able to operate across many types of vehicles and various communication environments.

Cross-platform detection performance will greatly increase by including a small amount of benign target domain calibration which includes samples of normal communications for the target vehicle platform. This form of calibration is very practical in terms of deployment into automobiles as they are able to collect examples of normal operation during an initial training or learning period. At the end of this learning period, the vehicle manufacturer could apply this process to the automobile during the time the vehicle is being commissioned; thus, allowing intrusion detection systems to learn about the unique communication characteristics of each vehicle platform and automatically adjust their behavior to those characteristics, eliminating the necessity for manual configuration.

This finding suggests that using graph-enhanced structural modeling as a method for developing cross-platform features is the best of the methods used in this research; therefore, it appears that using topology based features to develop the communication characteristics of vehicles will result in features that generalize better than other features. Therefore, researchers in this field should use graph-based communication analysis as a standard component of all intrusion detection processes in the automotive sector.

6.4. Broader Security Implications

Beyond this particular context of detecting CAN intrusions, this study's findings also point out a general principle applicable to network anomaly detection systems within safety-critical systems: feature representations that capture relational dependencies between components of a system tend to generalize more robustly than those based on marginal traffic statistics alone. This principle is especially important for connected and autonomous vehicles where the attack surface continues to grow as vehicles have external connectivity through v2x communications over-the-air update mechanisms and cloud-based services. As new communication interfaces provide additional entry points into the vehicle for attackers, intrusion detection systems must be able to generalize across previously unseen attack strategies. Structural and graph-based representations by definition capture fundamental communication dependencies rather than attack-specific statistical signatures; therefore they are a more future-proof foundation for automotive cybersecurity. Finally, the proposed framework could supplement other vehicle security mechanisms such as message authentication codes secure boot processses and intrusion prevention systems. By providing an additional layer of behavioral monitoring capable of detecting anomalous communication pattern even when cryptographic protections are bypassed or not available, structural intrusion detection contributes to a defense-in-depth security architecture suitable for modern connected vehicles.

7. Limitations

Despite the encouraging results shown by this study there are many restrictions that need to be identified. The primary focus of the proposed framework is on the structural aspects of communication, it does not explicitly consider the semantic aspects of payloads. Therefore, if an attacker is able to modify the values of signals while maintaining the communication pattern, such an attack would likely be difficult to identify. In addition, the experiments were performed with two public datasets, these datasets do not represent all possible real world communication environments related to vehicles. Additional vehicle datasets and/or real world driving conditions will be needed to further test the proposed method. Another restriction related to the proposed framework is the choice of the sliding

 doi:10.20944/preprints202603.1899.v1

window size used during feature extraction. A sensitivity analysis was conducted evaluating window sizes of 50, 100, 200, 300, and 500 CAN frames. Results demonstrated that the proposed framework maintains stable detection performance across all tested window sizes, with ROC-AUC values ranging from 0.9989 to 0.9994. However, Recall varied more noticeably, peaking at 0.9964 for a window size of 100 frames and decreasing to 0.9425 at 50 frames, suggesting that smaller windows may limit the structural context available for reliable classification. Based on these results, a window size of 200 frames was selected as the optimal balance between detection stability and computational efficiency.

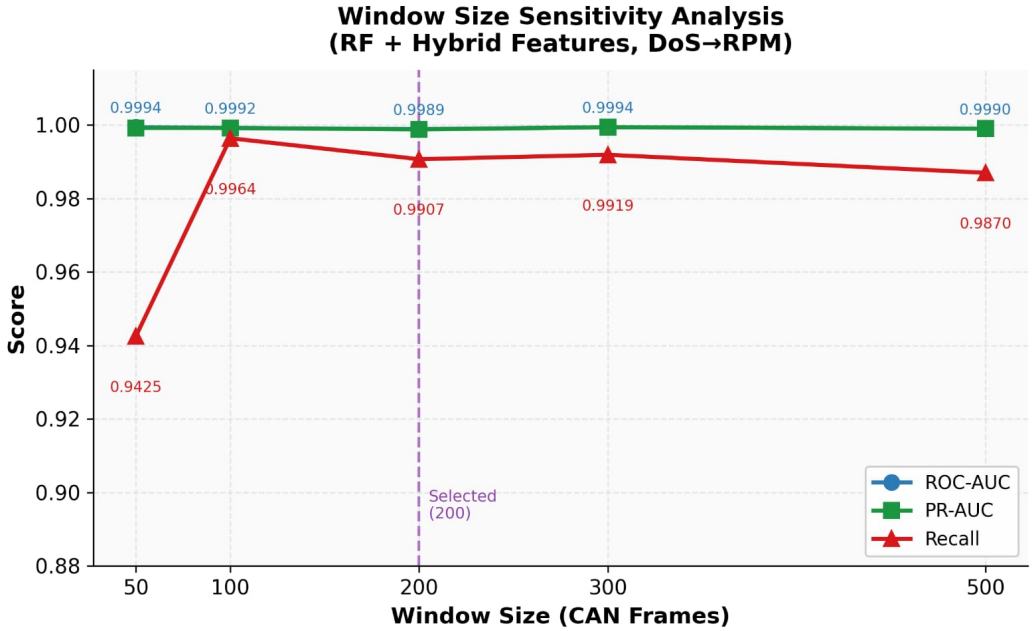

Figure 6. Future work could investigate adaptive windowing strategies that dynamically adjust window size based on real-time vehicle communication characteristics.

8. Threats to Validity

Several threats to validity should be considered when evaluating this experiment.

Threats to internal validity that arise from the design of an experiment are those that result from choices made in the design of the experiment such as choosing window sizes and features for extraction.

external validity threats arise from the diversity of the dataset. While cross-dataset experiments were performed to enhance generalizability across different vehicles manufacturers with additional datasets from other manufacturers would enhance the generalizability even further.

Construct validity threats may arise from labeling strategies used in the datasets.

Reproducibility threats may also arise from differences in preprocessing steps and feature extraction implementations. While the proposed feature extraction pipeline is relatively simple, differences in identifier parsing or preprocessing steps may also affect experimental results. Providing open-source implementations of the proposed framework will allow for easier reproducibility and enable future work to build upon these results.

9. Conclusions

BACKGROUND This research examined the limits of robustness of statistical CAN intrusion detection systems when tested under cross-attack evaluation scenarios. Statistical descriptors of traffic data (e.g., timing anomalies, identifier distributions, and payload characteristics) have been shown to be highly accurate at detecting intrusions in laboratory-controlled test cases; however, the results of this research indicate that these descriptions fail under distribution shifts when testing unseen attack types. In order to overcome these limits, the authors proposed a lightweight intrusion

 doi:10.20944/preprints202603.1899.v1

detection framework for modeling the structural patterns of CAN communication using structural identifier transition graphs and graph topology features. The proposed framework represents CAN traffic as sequential identifier interaction and uses graph-based descriptors to identify relational dependencies between electronic control units (ECU) that cannot be identified via marginal traffic data alone. Using the HCRL Car-Hacking dataset, the authors experimentally demonstrated that statistical descriptor representations produced near-random detection performance in cross-attack scenarios, with ROC-AUC values as low as 0.0088, confirming their inability to generalize across different attack behaviors. However, structural transition descriptor representations produced extremely high detection performance across all of the tested evaluation scenarios. Graph topology representations also demonstrated similar levels of robustness, as they captured the global structure of communication within the identifier transition graph. More importantly, the experimental results showed that this level of robustness was observed regardless of the machine learning classifier used (i.e., Logistic Regression, Support Vector Machine, Random Forest, and Gradient Boosting). Therefore, the authors believe that the primary source of the performance improvements arise from the use of the proposed relational feature representations, rather than being dependent upon a particular classification model. The ablation analysis further demonstrated that combining statistical, structural, and graph-based features would produce the highest overall detection performance. While statistical descriptors alone were found to be inadequate for providing robust detection against cross-attacks, the structural and graph representations demonstrated to be able to extract the stable behavioral dependencies that generalize across all attack strategies. Practically speaking, the proposed framework continues to be very computationally lightweight and therefore suitable for implementation in resource-limited automotive environments. Furthermore, feature extraction is performed in linear time and utilizes interpretable structural metrics rather than the computationally expensive architectures used by deep learning. Therefore, the proposed framework is capable of meeting the requirement for real-time monitoring of automotive electronic control units. Overall, the results of this research clearly illustrate the need to model the relational communication structure in CAN intrusion detection systems. By transitioning from simply analyzing traffic statistically and integrating structural and graph-based representations into the intrusion detection system, the system's robustness can be significantly enhanced against a wide variety of attack behaviors. In addition, unlike many deep learning-based CAN intrusion detection systems, the proposed framework is both highly effective at detecting robustly while maintaining its low computational overhead and interpretability, which make it suitable for deployment on resource limited automotive ECUs. Future research will include expanding the proposed framework to incorporate payload-level semantic analysis, which would enable detection of attacks that preserve communication structure while manipulating signal values. Additionally, evaluating the framework across further automotive datasets and vehicle platforms, along with investigation of adaptive windowing strategies, will help improve cross-domain generalization. Addressing the observed instability of decision tree classifiers within hybrid feature representations also remains an important direction for future work.

References

1. Wang, W.; Guo, K.; Cao, W.; Zhu, H.; Nan, J.; Yu, L. Review of Electrical and Electronic Architectures for Autonomous Vehicles: Topologies, Networking and Simulators. *Automotive Innovation* **2024**, *7*, 82–101. https://doi.org/10.1007/s42154-023-00266-9.
2. Anwar, A.S.; Anwar, A.; Moukahal, L.J.; Zulkernine, M.; Moukahal, L.; Zulkernine, M. Security assessment of in-vehicle communication protocols. *Vehicular Communications* **2023**, *44*, 100639–100639. https://doi.org/10.1016/j.vehcom.2023.100639.
3. Rai, R.; Grover, J.; Sharma, P.; Pareek, A. Securing the CAN bus using deep learning for intrusion detection in vehicles. *Scientific Reports* **2025**, *15*, 13820–13820. https://doi.org/10.1038/s41598-025-98433-x.
4. Rai, R.; Grover, J.; Sharma, P.; Pareek, A. Securing the CAN bus using deep learning for intrusion detection in vehicles. *Scientific Reports* **2025**, *15*, 13820–13820. https://doi.org/10.1038/s41598-025-98433-x.

 doi:10.20944/preprints202603.1899.v1

5. Tron, A.D.F.; Longari, S.; Carminati, M.; Polino, M.; Zanero, S. CANflict. *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security* **2022**, pp. 711–723. https://doi.org/10.1145/3548606.3560618.
6. El-Fatyany, A.; Wang, X.; Li, L.; Ren, K. EF-IDS: An efficient intrusion detection system with enriched features for CAN bus in modern vehicles. *Journal of Systems Architecture* **2025**, *171*, 103646–103646. https://doi.org/10.1016/j.sysarc.2025.103646.
7. Tampa.; Kirichenko, L.; Alghawli, A.S.; Ageyev, D.; Mulesa, O.; Baranovskyi, O.; Ilkov, A.; Kulbachnyi, V.; Bondarenko, O. Statistical and Signature Analysis Methods of Intrusion Detection. *Lecture notes on data engineering and communications technologies* **2022**, pp. 115–131. https://doi.org/10.1007/978-3-030-95161-0_5.
8. Ennaji, S.; De Gaspari, F.; Hitaj, D.; Kbidi, A.; Mancini, L.V. Adversarial Challenges in Network Intrusion Detection Systems: Research Insights and Future Prospects. *IEEE Access* **2025**, *13*, 148613–148645. https://doi.org/10.1109/access.2025.3600984.
9. Jia, H.; Xiong, X.; Luo, H.; Cao, Y. Driveshield: Unmasking stealthy attacks on CAN bus via adversarial spatiotemporal feature learning. *Journal of King Saud University - Computer and Information Sciences* **2025**, *37*. https://doi.org/10.1007/s44443-025-00273-2.
10. Lin, H.; Yu, X.; Chen, Z.; Cao, Y. RAG-HIDS: A multi-relational graph-based hierarchical intrusion detection system for in-vehicle networks. *Ad Hoc Networks* **2025**, *183*, 104108–104108. https://doi.org/10.1016/j.adhoc.2025.104108.
11. Lin, X.; Ma, B.; Wang, X.; Yu, G.; He, Y.; Liu, R.P.; Ni, W. ByCAN: Reverse Engineering Controller Area Network (CAN) Messages From Bit to Byte Level. *IEEE Internet of Things Journal* **2024**, *11*, 35477–35491. https://doi.org/10.1109/jiot.2024.3435833.
12. Lin, H.; Yu, X.; Chen, Z.; Cao, Y. RAG-HIDS: A multi-relational graph-based hierarchical intrusion detection system for in-vehicle networks. *Ad Hoc Networks* **2025**, *183*, 104108–104108. https://doi.org/10.1016/j.adhoc.2025.104108.
13. Donadel, D.; Balasubramanian, K.; Brighente, A.; Cleaveland, R.; Conti, M.; Poovendran, R. CANTXSec: A Deterministic Intrusion Detection and Prevention System for CAN Bus Monitoring ECU Activations. *Lecture notes in computer science* **2025**, pp. 429–458. https://doi.org/10.1007/978-3-031-95764-2_17.
14. Ye, P.; Liang, Y.; Bie, Y.; Qin, G.; Song, J.; Wang, Y.; Liu, W. GDT-IDS: graph-based decision tree intrusion detection system for controller area network. *The Journal of Supercomputing* **2025**, *81*. https://doi.org/10.1007/s11227-025-07116-x.
15. Longari, S.; Cerracchio, P.; Carminati, M.; Zanero, S. Assessing the Resilience of Automotive Intrusion Detection Systems to Adversarial Manipulation. *ACM Transactions on Cyber-Physical Systems* **2025**, *9*, 1–27. https://doi.org/10.1145/3737294.
16. Rimal, Y.; Sharma, N.; Paudel, S.; Alsadoon, A.; Koirala, M.; Gill, S. Comparative analysis of heart disease prediction using logistic regression, SVM, KNN, and random forest with cross-validation for improved accuracy. *Scientific Reports* **2025**, *15*, 13444–13444. https://doi.org/10.1038/s41598-025-93675-1.
17. Liu, C.; Dong, Y.; Xiang, W.; Yang, X.; Su, H.; Zhu, J.; Chen, Y.; He, Y.; Xue, H.; Zheng, S. A Comprehensive Study on Robustness of Image Classification Models: Benchmarking and Rethinking. *International Journal of Computer Vision* **2024**, *133*, 567–589. https://doi.org/10.1007/s11263-024-02196-3.
18. Debicha, I.; Bauwens, R.; Debatty, T.; Dricot, J.; Kenaza, T.; Mees, W. TAD: Transfer learning-based multi-adversarial detection of evasion attacks against network intrusion detection systems. *Future Generation Computer Systems* **2022**, *138*, 185–197. https://doi.org/10.1016/j.future.2022.08.011.
19. Cerezo, M.; Arrasmith, A.; Babbush, R.; Benjamin, S.C.; Endo, S.; Fujii, K.; McClean, J.R.; Mitarai, K.; Yuan, X.; Cincio, L. Variational quantum algorithms. *Nature Reviews Physics* **2021**, *3*, 625–644. https://doi.org/10.1038/s42254-021-00348-9.
20. Saravanan, R.; Balaji, S.; Ganesan, M.; Braveen, M.; Perumal, R.S. Optimal attention deep learning based in-vehicle intrusion detection and classification model on CAN messages. *Scientific Reports* **2025**, *15*, 33952–33952. https://doi.org/10.1038/s41598-025-10637-3.
21. Saravanan, R.; Balaji, S.; Ganesan, M.; Braveen, M.; Perumal, R.S. Optimal attention deep learning based in-vehicle intrusion detection and classification model on CAN messages. *Scientific Reports* **2025**, *15*, 33952–33952. https://doi.org/10.1038/s41598-025-10637-3.
22. Dash, N.; Chakravarty, S.; Rath, A.K.; Giri, N.C.; AboRas, K.M.; Gowtham, N. An optimized LSTM-based deep learning model for anomaly network intrusion detection. *Scientific Reports* **2025**, *15*, 1554–1554. https://doi.org/10.1038/s41598-025-85248-z.

 doi:10.20944/preprints202603.1899.v1

23. Shahriar, M.H.; Xiao, Y.; Moriano, P.; Lou, W.; Hou, Y.T. CANShield: Deep-Learning-Based Intrusion Detection Framework for Controller Area Networks at the Signal Level. *IEEE Internet of Things Journal* **2023**, *10*, 22111–22127. https://doi.org/10.1109/jiot.2023.3303271.
24. Alzubaidi, L.; Zhang, J.; Humaidi, A.J.; Al-Dujaili, A.Q.; Duan, Y.; Al-Shamma, O.; Santamaría, J.; Fadhel, M.A.; Al-Amidie, M.; Farhan, L. Review of deep learning: concepts, CNN architectures, challenges, applications, future directions. *Journal Of Big Data* **2021**, *8*, 53–53. https://doi.org/10.1186/s40537-021-00444-8.
25. Kumar, L.K.S.; Nethi, S.R.; Uyyala, R.; Vurubindi, P.; Narahari, S.C.; Das, A.K.; K, V.B.; Alenazi, M.J.F. Anomaly-based intrusion detection on benchmark datasets for network security: a comprehensive evaluation. *Scientific Reports* **2026**, *16*. https://doi.org/10.1038/s41598-026-38317-w.
26. Yin, T.; Naqvi, S.A.R.; Nandanoori, S.P.; Kundu, S. Advancing Cyber-Attack Detection in Power Systems: A Comparative Study of Machine Learning and Graph Neural Network Approaches. *2024 Resilience Week (RWS)* **2024**, pp. 1–9. https://doi.org/10.1109/rws62797.2024.10799283.
27. Zola, F.; Medina, J.A.; Venturi, A.; Gil, A.; Orduna-Urrutia, R. A Graph Machine Learning Approach for Detecting Topological Patterns in Transactional Graphs. *ArXiv.org* **2025**. https://doi.org/10.48550/arxiv.2509.12730.
28. Lin, H.; Yu, X.; Chen, Z.; Cao, Y. RAG-HIDS: A multi-relational graph-based hierarchical intrusion detection system for in-vehicle networks. *Ad Hoc Networks* **2025**, *183*, 104108–104108. https://doi.org/10.1016/j.adhoc.2025.104108.
29. Ahmed, U.; Nazir, M.; Sarwar, A.; Ali, T.; Aggoune, E.M.; Shahzad, T.; Khan, M.A. Signature-based intrusion detection using machine learning and deep learning approaches empowered with fuzzy clustering. *Scientific Reports* **2025**, *15*, 1726–1726. https://doi.org/10.1038/s41598-025-85866-7.
30. Pollicino, F.; Stabili, D.; Marchetti, M. Performance Comparison of Timing-Based Anomaly Detectors for Controller Area Network: A Reproducible Study. *ACM Transactions on Cyber-Physical Systems* **2023**, *8*, 1–24. https://doi.org/10.1145/3604913.
31. Pollicino, F.; Stabili, D.; Marchetti, M. Performance Comparison of Timing-Based Anomaly Detectors for Controller Area Network: A Reproducible Study. *ACM Transactions on Cyber-Physical Systems* **2023**, *8*, 1–24. https://doi.org/10.1145/3604913.
32. Yu, H.; Yang, W.; Cui, B.; Sui, R.; Wu, X. Renyi entropy-driven network traffic anomaly detection with dynamic threshold. *Cybersecurity* **2024**, *7*. https://doi.org/10.1186/s42400-024-00249-1.
33. de Heij, V.; Niazi, M.U.B.; Ahmed, S.; Johansson, K.H. Distributed Traffic State Estimation in V2X-Enabled Connected Vehicle Networks. *ArXiv.org* **2025**.
34. Redhu, A.; Choudhary, P.; Srinivasan, K.; Das, T.K. Deep learning-powered malware detection in cyberspace: a contemporary review. *Frontiers in Physics* **2024**, *12*. https://doi.org/10.3389/fphy.2024.1349463.
35. Kiflay, A.; Tsokanos, A.; Fazlali, M.; Kirner, R. Network intrusion detection leveraging multimodal features. *Array* **2024**, *22*, 100349–100349. https://doi.org/10.1016/j.array.2024.100349.
36. Alharthi, A.M.; Alaryani, M.; Kaddoura, S. A comparative study of machine learning and deep learning models in binary and multiclass classification for intrusion detection systems. *Array* **2025**, *26*, 100406–100406. https://doi.org/10.1016/j.array.2025.100406.
37. Lampe, B.; Meng, W. can-train-and-test: A curated CAN dataset for automotive intrusion detection. *Computers & Security* **2024**, *140*, 103777–103777. https://doi.org/10.1016/j.cose.2024.103777.
38. Jiang, W.; Zhang, T.; Liu, S.; Ji, W.; Zhang, Z.C.; Xiao, G. Exploring the Physical-World Adversarial Robustness of Vehicle Detection. *Electronics* **2023**, *12*, 3921–3921. https://doi.org/10.3390/electronics12183921.
39. Lampe, B.; Meng, W. can-train-and-test: A curated CAN dataset for automotive intrusion detection. *Computers & Security* **2024**, *140*, 103777–103777. https://doi.org/10.1016/j.cose.2024.103777.
40. Cantone, M.; Marrocco, C.; Bria, A. On the Cross-Dataset Generalization of Machine Learning for Network Intrusion Detection. *arXiv (Cornell University)* **2024**. https://doi.org/10.48550/arxiv.2402.10974.
41. Taneja, A.; Kumar, G. Attention-CNN-LSTM based intrusion detection system (ACL-IDS) for in-vehicle networks. *Soft Computing* **2024**, *28*, 13429–13441. https://doi.org/10.1007/s00500-024-10313-0.
42. Kumar, L.K.S.; Nethi, S.R.; Uyyala, R.; Vurubindi, P.; Narahari, S.C.; Das, A.K.; K, V.B.; Alenazi, M.J.F. Anomaly-based intrusion detection on benchmark datasets for network security: a comprehensive evaluation. *Scientific Reports* **2026**, *16*. https://doi.org/10.1038/s41598-026-38317-w.
43. Zhao, J.; Wu, Y.; Deng, R.; Xu, S.; Gao, J.; Burke, A. A Survey of Autonomous Driving from a Deep Learning Perspective. *ACM Computing Surveys* **2025**, *57*, 1–60. https://doi.org/10.1145/3729420.

 doi:10.20944/preprints202603.1899.v1

44. Munappy, A.R.; Bosch, J.; Olsson, H.H.; Arpteg, A.; Brinne, B. Data management for production quality deep learning models: Challenges and solutions. *Journal of Systems and Software* **2022**, *191*, 111359–111359. https://doi.org/10.1016/j.jss.2022.111359.
45. Saravanan, R.; Balaji, S.; Ganesan, M.; Braveen, M.; Perumal, R.S. Optimal attention deep learning based in-vehicle intrusion detection and classification model on CAN messages. *Scientific Reports* **2025**, *15*, 33952–33952. https://doi.org/10.1038/s41598-025-10637-3.
46. Bhatt, N.; Bhatt, N.; Prajapati, P.; Sorathiya, V.; Alshathri, S.; El-Shafai, W. A Data-Centric Approach to improve performance of deep learning models. *Scientific Reports* **2024**, *14*, 22329–22329. https://doi.org/10.1038/s41598-024-73643-x.
47. Ekundayo, O.; Ezugwu, A.E. Deep learning: Historical overview from inception to actualization, models, applications and future trends. *Applied Soft Computing* **2025**, *181*, 113378–113378. https://doi.org/10.1016/j.asoc.2025.113378.
48. Hassija, V.; Chamola, V.; Mahapatra, A.; Singal, A.; Goel, D.; Huang, K.; Scardapane, S.; Spinelli, I.; Mahmud, M.; Hussain, A. Interpreting Black-Box Models: A Review on Explainable Artificial Intelligence. *Cognitive Computation* **2023**, *16*, 45–74. https://doi.org/10.1007/s12559-023-10179-8.
49. Saeed, N.H.; Hamza, A.A.; Sobh, M.A.; Bahaa-Eldin, A.M. Efficient feature ranked hybrid framework for android Iot malware detection. *Scientific Reports* **2026**, *16*, 3726–3726. https://doi.org/10.1038/s41598-026-35238-6.