

Detecting Masquerade Attacks in Controller Area Networks Using Graph Machine Learning

William Marfo¹, Member, IEEE, Pablo Moriano², Senior Member, IEEE,
Deepak K. Tosh¹, Senior Member, IEEE, and Shirley V. Moore, Member, IEEE

Abstract—Modern vehicles rely on a myriad of electronic control units (ECUs) interconnected via controller area networks (CANs) for critical operations. Despite their ubiquitous use and reliability, CANs are susceptible to sophisticated cyberattacks, particularly masquerade attacks, which inject false data that mimic legitimate messages at the expected frequency. These attacks pose severe risks such as unintended acceleration, brake deactivation, and rogue steering. Traditional intrusion detection systems (IDS) often struggle to detect these subtle intrusions due to their seamless integration into normal traffic. This paper introduces a novel framework for detecting masquerade attacks in the CAN bus using graph machine learning (ML). We hypothesize that the integration of shallow graph embeddings with time series features derived from CAN frames enhances the detection of masquerade attacks. We show that by representing CAN bus frames as message sequence graphs (MSGs) and enriching each node with contextual statistical attributes from time series, we can enhance detection capabilities across various attack patterns compared to using graph-based features only. Our method ensures a comprehensive and dynamic analysis of CAN frame interactions, improving robustness and efficiency. Extensive experiments on the ROAD dataset validate the effectiveness of our approach, demonstrating statistically significant improvements in the detection rates of masquerade attacks compared to a baseline that uses graph-based features only as confirmed by Mann-Whitney U and Kolmogorov-Smirnov tests ($p < 0.05$).

Index Terms—Controller area networks, intrusion detection systems, graph ML, masquerade attacks.

I. INTRODUCTION

CONTROLLER area networks (CANs) have become ubiquitous in modern vehicles and industrial applications, forming the backbone of communication between electronic control units (ECUs) responsible for essential functions such as acceleration, braking, and steering [1], [2], [3]. Renowned for their reliability and efficiency, CANs have become the de facto standard for in-vehicle communication. However, increased integration with external systems—such as diagnostics, firmware updates, and advanced driver assistance systems

(ADAS)—has exposed CANs to a growing range of cyber threats that compromise data integrity and system functionality [4]. CAN attacks are typically grouped into fabrication, suspension, and masquerade categories [5].

Among these threats, masquerade attacks are particularly insidious. These attacks inject malicious CAN frames that emulate legitimate messages, manipulating vehicle behavior while evading conventional intrusion detection systems (IDS) [6], [7]. Unlike fabrication or suspension attacks, which create noticeable disruptions, masquerade attacks mirror the structure and frequency of authentic messages, making them exceptionally difficult to detect [5]. The consequences are severe: a compromised ECU could subtly trigger unintended acceleration, disengage braking mechanisms, or disable safety-critical features—actions that, if left unnoticed, can lead to catastrophic outcomes on the road.

While IDS solutions have been deployed to safeguard CANs, most fall into two categories: signature-based and anomaly-based approaches. Signature-based IDS are effective against known threats but fail against zero-day attacks [8], [9]. Anomaly-based IDS offer broader coverage but struggle with high rates of both false positives and negatives, as well as insufficient context-awareness for detecting stealthy behavior [10], [11]. Existing solutions also tend to rely on either structural or signal-based analysis in isolation, failing to capture the critical relationship between message sequences and signal dynamics—a relationship that is essential for detecting masquerade attacks.

Masquerade attacks are uniquely challenging to detect due to their ability to seamlessly blend into normal traffic. These attacks mimic the frequency and content of legitimate frames, thereby bypassing traditional IDS techniques that rely on surface-level anomalies. Prior studies have shown that detection requires examining multiple consecutive frames, as single-message checks are often insufficient [2]. Moreover, methods that use either graph structure or signal statistics alone often miss subtle patterns that span both temporal and structural domains [12], [13]. For example, a vehicle under attack may gradually exhibit abnormal behavior—such as slowing to a halt—not due to one erroneous message, but due to a sequence of seemingly valid ones. Detecting such behavior demands a multimodal perspective.

To address these limitations, this work proposes a unified framework that integrates graph-based and time series analysis for the detection of masquerade attacks in CANs.

Received 25 April 2025; revised 4 July 2025; accepted 9 November 2025. Date of publication 21 November 2025; date of current version 16 December 2025. The associate editor coordinating the review of this article and approving it for publication was Prof. Kun Sun. (Corresponding author: Pablo Moriano.)

William Marfo, Deepak K. Tosh, and Shirley V. Moore are with the Department of Computer Science, The University of Texas at El Paso, El Paso, TX 79968 USA (e-mail: wmarfo@miners.utep.edu; dktosh@utep.edu; svmoore@utep.edu).

Pablo Moriano is with the Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, TN 37830 USA (e-mail: moriano@ornl.gov).

Digital Object Identifier 10.1109/TIFS.2025.3636019

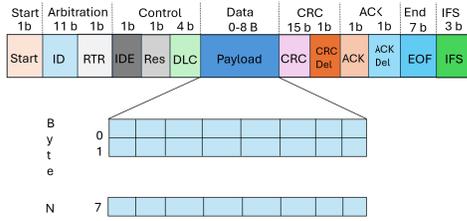


Fig. 1. CAN protocol frame with 8-byte payload. The full structure is presented for comprehensive understanding, with a specific focus on the ID and payload fields for our analysis.

By modeling CAN traffic as MSGs and enriching each node with contextual time series features, our method captures both structural irregularities and signal-level deviations. This dual perspective enables robust detection of subtle intrusions that would otherwise go unnoticed in traditional systems. We also report the testing time per window (TTW) [14] as a proxy for inference time to assess the real-time detection capability of our approach.

The main contributions of this paper can be summarized as follows:

- We propose a comprehensive framework for detecting masquerade attacks in the CAN bus. Our approach uses graph ML by integrating shallow graph embeddings and time series analysis to capture both the structural and temporal aspects of CAN frames. By representing CAN frames as MSGs, we enhance the ability to detect subtle masquerade attacks that traditional IDS might miss.
- We implement a robust node annotation technique within MSGs that significantly enhances the detection process. By incorporating key statistical attributes derived from time series data, such as the mean and standard deviation of signals, into each CAN ID node, our method provides a richer, context-aware analysis. This detailed annotation not only improves the accuracy of masquerade attack detection but also ensures continuous monitoring and dynamic adaptation to changing network behaviors.
- We evaluate our proposed framework through extensive experiments using the ROAD dataset, a benchmark in vehicular network security research. Our results demonstrate significant improvements in detecting masquerade attacks when compared to an approach based graph topology only. The evaluation highlights the scalability and robustness of our approach, ensuring its practical applicability in modern vehicular networks.

We have made the code available to reproduce all the results at [15].

II. BACKGROUND

A. CAN Protocol

CAN is a widely adopted message-based communication protocol designed to enable efficient data exchange among ECUs in vehicles and industrial systems [16]. CAN operates on the physical and data-link layers of the OSI model, using a multi-master, broadcast-based communication model over a two-wire bus, as shown in Fig. 1.

Each CAN frame contains multiple fields; however, in the context of this work, we focus on two: the 11-bit arbitration ID and the up-to-8-byte data field. The ID determines message priority, with lower values corresponding to higher priority. The data field holds signal-level information such as speed or engine temperature. Messages with the same ID are transmitted periodically with a fixed payload structure, enabling the modeling of temporal and structural relationships among signals. For example, a powertrain ECU may broadcast a frame with ID 0×102 at 20 Hz containing real-time values for speed and odometer. These periodic patterns are typically defined using a CAN database (DBC) file, which decodes binary payloads into meaningful physical signals. This representation is particularly useful for downstream ML-based IDS that leverage time series analysis and graph-based representations [17].

Beyond the protocol's operational details, it is crucial to understand the security landscape and the various types of attacks targeting the CAN bus. Our approach targets multiple tiers of CAN bus threats, which are typically classified into three categories based on adversarial intent [5]:

- **Fabrication attacks:** Malicious messages are injected onto the bus by a rogue ECU, without disrupting normal traffic from other ECUs [18].
- **Suspension attacks:** A legitimate ECU is disabled or disconnected, causing its expected messages to vanish temporarily [19].
- **Masquerade attacks:** The most stealthy and severe—these suppress a legitimate ECU and impersonate it by replaying spoofed frames at realistic timing and content patterns [4].

Note that while this study focuses on masquerade attacks due to their stealth and impact, the proposed framework is also applicable to fabrication and suspension attacks, as these introduce structural and temporal deviations captured by MSGs.

B. CAN Message Sequence Graph (MSG)

The MSG is a crucial tool for analyzing the behavior of the CAN [2]. It captures the sequence pattern of messages, reflecting their normal sequence of messages in CAN. Nodes in this graph correspond to unique CAN identifiers (IDs), denoting different ECUs or functions within the vehicle network. Edges map the sequence of messages between these IDs illustrating how information flows within the system. The MSG models regular patterns of communication between ECUs within specified sliding windows, allowing for detailed observation of both normal and potentially malicious activities. Typical sequence patterns found during various vehicle operations, such as parking or acceleration, can be captured by the MSG. These patterns form the basis for detecting deviations that may indicate cybersecurity threats like fabrication and suspension attacks.

Constructing an MSG involves parsing time-stamped CAN messages into a graph structure that captures the sequence and frequency of CAN ID transitions. This highlights communication patterns typical of normal vehicle operations—such as

fuel delivery and speed coordination during acceleration—and enables detection of disruptions indicative of attacks. As shown in Section IV-A, MSGs visualize these dynamics, supporting robust analysis of ECU interactions and vehicular communication integrity.

III. RELATED WORK

We discuss prior work closely related to graph-based CAN IDS (Section III-A) and other prior work related to general CAN IDS for detecting masquerade attacks (Section III-B).

A. Prior Work Closely Related to the Present Study

Islam et al. [20] proposed a four-stage IDS utilizing the chi-squared method to identify both strong and weak cyber-attacks on the CAN bus. Their approach, which is among the first graph-based defense systems for CAN, showed misclassification rates of 5.26% for DoS attacks, 10% for fuzzy attacks, 4.76% for replay attacks, and zero misclassification for spoofing attacks. The dataset used for evaluation was obtained from the Hacking and Countermeasure Research Lab [27]. This study emphasized the need for robust security mechanisms in modern vehicles and demonstrated superior accuracy compared to existing ID sequence-based methods.

Jedh et al. [2] proposed a message injection attack detection solution. Their approach leveraged MSGs and used graph-based analytics and anomaly detection to detect malicious message injections with high accuracy and low detection time. They validated their approach using a dataset collected from a moving vehicle under attack conditions. This study highlighted the importance of addressing the security of in-vehicle communication networks without relying on proprietary ECU information.

Refat et al. [23] presented a lightweight IDS that translates CAN traffic into temporal graphs and applies neighborhood-based graph similarity techniques to detect message injection attacks. The system was evaluated using real vehicle data [28], achieving a detection accuracy of 96.01% for spoofing, fuzzy, and DoS attacks. This work contributed to vehicle security by providing a computationally efficient method without requiring changes to the CAN protocol.

Zhang et al. [17] proposed a CAN bus anomaly detection system using graph neural networks (GNNs) to detect message injection, suspension, and falsification attacks. Their approach involved creating directed attributed graphs from CAN message streams and training a two-stage classifier cascade. The evaluation on a Ford Transit 500 dataset [29] demonstrated the system's efficiency in real-time detection and its ability to handle new anomalies through federated learning training.

Xiao et al. [11] proposed the CAN-GAT model based on graph attention networks for in-vehicle networks. By transforming CAN bus messages into graph structures, the model captured the correlation between traffic bytes, improving detection accuracy and efficiency. The model was evaluated using datasets from the Hacking and Countermeasure Research Lab [27], demonstrating superior performance among compared GNNs.

Meng et al. [26] developed GB-IDS, an IDS leveraging a novel graph structure and a variational autoencoder for

training classifiers without negative samples. Their system, evaluated on the OTIDS dataset [30], achieved high detection success rates for DoS, fuzzing, and impersonation attacks. This study addressed the limitations of traditional IDS, i.e., [23] by avoiding the need for protocol parsing and large training datasets.

Compared to prior works that focus on either graph-based or time series-based CAN IDS, our approach offers a unified framework that addresses their respective limitations. Graph-based methods, such as [2] and [20], use MSGs to detect structural anomalies but lack granularity for subtle signal deviations. Time series-based techniques such as [12] and [13] capture temporal anomalies but often miss intricate message sequence patterns and graph topology that are crucial for uncovering subtle masquerade attacks. These single-modality designs are insufficient for detecting stealthy masquerade attacks. Our framework combines graph embeddings and time series features, enabling detection across both structural and signal domains. Additionally, its flexibility—adjustable via window length and offset—supports fine-tuning in constrained environments. We evaluate this framework on the ROAD dataset [5], demonstrating improved performance over recent approaches [17], [24].

B. Other Prior Work Related to the Present Study

Hanselmann et al. [13] proposed CANet, a neural network architecture for detecting intrusions on the CAN bus. CANet, the first deep learning-based IDS capable of handling the high dimensionality of CAN bus data, was evaluated on real and synthetic CAN data. The method outperformed previous ML-based approaches with a high true negative rate, typically over 0.99, demonstrating robustness in detecting a large number of synthetic masquerade attacks [13].

Moriano et al. [4] focused on detecting masquerade attacks on the CAN bus by analyzing time series extracted from raw CAN frames. Using hierarchical clustering, this study demonstrated that changes in cluster assignments could indicate anomalous behavior. The proposed forensic tool was tested on the ROAD dataset [5], showing significant differences in time series clustering similarity on benign and attack conditions.

Shahriar et al. [3] proposed CANtropy, a feature engineering-based lightweight CAN IDS. CANtropy explores a comprehensive set of features from both temporal and statistical domains utilizing PCA for anomaly detection. Evaluation on the SynCAN dataset [13] showed CANtropy's effectiveness, with an average AUC-ROC score of 0.992, outperforming existing DL-based baselines.

Moriano et al. [12] conducted a benchmark study on four non-deep learning unsupervised online IDS for masquerade attacks in CANs. They controlled streaming data conditions in a sliding window setting and used realistic attacks from the ROAD dataset [5]. They noticed that there is no one-size-fits-all solution for detecting masquerade attacks in CAN in an online setting, i.e., there is no single algorithm that is optimal for detecting each attack type. Notably, a method that detects changes in the hierarchical structure of clusters of time series outperformed others in detecting various attack types at the expense of computational overhead.

IV. PROPOSED DATA-DRIVEN METHODS

Algorithm 1 End-to-End Masquerade Attack Detection

Require: Raw CAN logs, ω_t , δ_t , DBC mappings, $node2vec_{params}$

Ensure: Trained detection model

- 1 $\mathcal{T} \leftarrow \text{ExtractTimeSeries}(CAN, DBC) \triangleright$ Sec. IV-B
- 2 $\mathcal{G} \leftarrow \text{Build MSGs}(CAN, \omega_t, \delta_t) \triangleright$ Alg. 2
- 3 **for** each graph $G_i \in \mathcal{G}$ **do**
- 4 AnnotateNodes(G_i, \mathcal{T}) \triangleright Alg. 3
- 5 $E_i \leftarrow \text{node2vec}(G_i, node2vec_{params}) \triangleright$ Sec. IV-D
- 6 $F_i \leftarrow \text{Combine}(E_i, G_i, \mu, G_i, \sigma) \triangleright$ Feature fusion
- 7 **end for**
- 8 $\mathcal{M} \leftarrow \text{TrainRandomForest}(\{F_1, \dots, F_n\})$ **return** \mathcal{M}

This section presents our approach for detecting masquerade attacks in CAN bus communication using graph ML. We begin by outlining the complete pipeline in Fig. 2, followed by a high-level algorithmic workflow in Algorithm 1. We then provide detailed procedures for MSG construction (Section IV-A), time series extraction (Section IV-B), node annotation (Section IV-C), graph embedding generation (Section IV-D), and supervised learning process (Section IV-E). Evaluation setup and metrics are described in Section IV-F.

Threat model: In our framework, we assume the attacker has gained access to the CAN bus, either physically or remotely [31]. Masquerade attacks, as represented in the ROAD dataset [5], typically involve two compromised ECUs: a *weak attacker* that learns the frequency of messages sent by a targeted ECU and suppresses its transmissions, and a *strong attacker* that injects spoofed messages at the expected frequency using the same ID. The coordinated nature of this attack enables the adversary to mimic legitimate traffic patterns while altering the message content. In this study, our detection framework is evaluated under the assumption that the adversary can play both attacker roles, consistent with the ROAD dataset’s multi-stage masquerade attack model.

Note that to quantify the computational footprint of our end-to-end detection pipeline, we analyze its complexity.

Computational Complexity: The total cost is dominated by:

$$\mathcal{T}_{total} = \underbrace{O\left(\frac{n}{\delta_t} \cdot \omega_t\right)}_{\text{MSG building}} + \underbrace{O(n + k \cdot s_{sig})}_{\text{Time series processing}} + \underbrace{O(r \cdot l \cdot k)}_{\text{Graph embeddings}} + \underbrace{O(t \cdot d \cdot n_{train})}_{\text{ML Training}}$$

where:

- n : Total CAN messages.
- k : Unique CAN IDs.
- s_{sig} : Maximum number of signals per CAN ID (not constant; varies across IDs, typically 2–10, with an upper limit of 16 in the ROAD dataset).
- n_{train} : Number of training samples (approximately the number of sliding windows).
- ω_t, δ_t : Window size and offset for time-based configurations (seconds); $\omega_t \in [2, 15]$, $\delta_t \in [1, \omega_t]$.

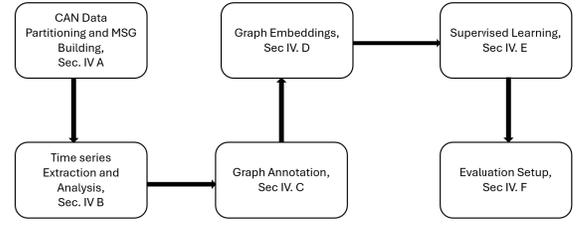


Fig. 2. Proposed data-driven framework for masquerade attack detection in CAN using graph ML.

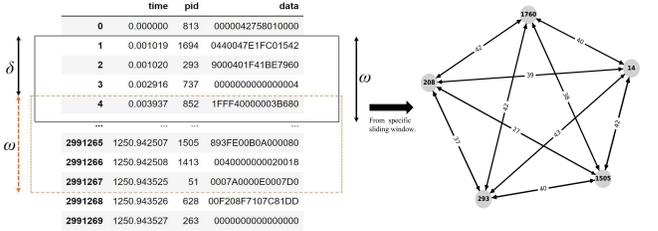


Fig. 3. Sliding window partitioning of CAN messages (left) and resulting MSG subgraph (right). Top five nodes by connectivity are highlighted; edges are weighted.

- ω_s, δ_s : Window size and offset for sample-based configurations (messages); $\omega_s \in \{50, 100, \dots, 400\}$, $\delta_s \in \{50, 100, \dots, \omega_s\}$.
- r : Number of node2vec walks per node (100).
- l : Length of each node2vec walk (15).
- t : Number of trees in the Random Forest classifier (100).
- d : Maximum tree depth in the Random Forest classifier ($\approx \log_2 n_{train}$).

Key Observations:

- Training cost is negligible because:
 - It is done offline (not during vehicle operation)
 - Typically $n_{train} \ll 10^4$ windows
 - Random Forest scales linearly with n_{train}
- Time series extraction is linear in n (CAN-D decoding)
- MSG construction dominates when $\omega_t \gg \delta_t$
- Embedding generation (approximately 70% runtime) scales with k^2 .

A. CAN Data Partitioning and MSG Building

We use sliding windows of fixed size to partition and process the stream of CAN data [32]. The raw CAN data is first transformed into a structured dataframe format where each row represents a message with its timestamp, process ID, and data payload (see Fig. 3). This data is then partitioned into discrete time slices defined by specific window sizes and offsets, enabling a continuous analysis of CAN data.

We now define the key parameters controlling the sliding windows:

- **Window size (ω , either ω_t for time or ω_s for samples):** The length of each window determining the extent of data included in each graph representation.

- **Offset** (δ , either δ_t for time or δ_s for samples): The sliding step between consecutive windows, where δ represents the number of samples separating each window, ensuring no loss of data and the capture of communication patterns that may span multiple windows.

Parameter Rationale—Aligning with Attack Dynamics:

The masquerade attacks in the ROAD dataset [5] are sustained, stealthy intrusions (e.g., `correlated_signal`, `max_speedometer`, `max_engine_coolant_temp`, `reverse_light_off`, `reverse_light_on`) lasting approximately 17–72 seconds. Our parameters were empirically optimized to capture attack duration while maintaining computational feasibility. For example, a 6-second window ensures sufficient temporal context to detect anomalies that unfold gradually, while a 3.2-second offset prevents missed detections near window boundaries through overlapping windows. This offset does not introduce a reaction delay but strategically balances attack coverage and computational feasibility—smaller offsets (e.g., 1s) increase processing load, whereas larger offsets (e.g., 6s) risk missing attacks. The selected window parameters ensure that our detection method aligns with real attack lifecycles.

Motivation for Sliding Window Detection: Masquerade attacks cannot be detected at the message level—they mimic legitimate traffic in short bursts but manifest as cumulative deviations over time. As demonstrated by Jedh et al. [2], stealthy message injections, especially masquerade attacks, necessitate analysis over multiple consecutive frames rather than relying on single-message checks. Our approach aggregates e.g., 600 messages (at 100Hz) within a 6-second window, enabling structural (graph-based) and temporal (statistical) analysis to expose subtle attack patterns.

A MSG is constructed within each defined sliding window to model CAN activity following the procedure introduced by [2] and [20]. Graphs are empirically recognized as effective tools to model relationships among relational data that are too complex for tabular representation or other simpler data structures [33]. The process of building MSGs is crucial as it sets the stage for subsequent analyses, including node annotation and embedding generation. By modeling CAN data using MSG on sliding windows, we enhance our ability to detect and respond to potential masquerade attacks by finetuning the portion of analyzed data. We now detail each of the components of the MSG:

- **Nodes:** Each node in the MSG corresponds to a unique ID in CAN representing the header of the CAN frame. This alignment allows for the detailed representation of each frame’s interactions on streams of CAN data.
- **Edges:** Directed edges are established based on the sequential relationships of messages. That is, an edge is formed when one ID follows another, indicating the flow of communication [20]. This method effectively maps sequences of CAN messages enabling the detection of irregular patterns that may signify an attack. Edge weights are assigned to edges of MSGs patterns to quantify the frequency of observed communication sequences within a sliding window. Specifically, the weight of an edge from node v_i to node v_j is computed as the number of

times the CAN message with ID v_j directly follows the CAN message with ID v_i within the same sliding window. Weights help to identify anomalous recurrent patterns that could indicate security threats.

Algorithm 2 Creation of Message Sequence Graphs (MSGs)

Require: Raw CAN data, ω_t , δ_t
Ensure: MSG representing CAN activity over time

- 1 Initialize an empty directed graph G
- 2 Initialize $current_time = start_time$ of CAN data
- 3 **while** $current_time + \omega_t \leq end_time$ of CAN data **do**
- 4 Define $start_time = current_time$ and $end_time = current_time + \omega_t$
- 5 Extract messages in $[start_time, end_time]$
- 6 Initialize a new subgraph G_{sub}
- 7 **for** each message m in $[start_time, end_time]$ **do**
- 8 Extract id from m
- 9 **if** id not in G_{sub} **then**
- 10 Add id as a node in G_{sub}
- 11 **end if**
- 12 Determine next id in sequence
- 13 **if** edge from current id to next id exists **then**
- 14 Increment weight of edge
- 15 **else**
- 16 Add edge with weight = 1
- 17 **end if**
- 18 **end for**
- 19 Merge G_{sub} into main graph G respecting δ
- 20 Increment $current_time$ by δ_t
- 21 **end while** **return** G

Fig. 3 shows how MSGs are built and their corresponding components. Our modeling framework ensures that both structural and temporal aspects of CAN communications are captured, facilitating a comprehensive analysis of the CAN dynamics. Algorithm 2 describes the creation of MSGs.

To ensure the correctness of Algorithm 2, we establish loop invariants for both the outer windowing loop and inner message processing loop. A loop invariant is a property that holds true before and after each iteration, ensuring algorithmic correctness. We validate these invariants through three steps:

1) **Loop Invariants:**

- **Outer Loop (Lines 3-21):** At iteration i , the main graph G contains all message sequences from windows $\{[t_0, t_0 + \omega_t], \dots, [t_{i-1}, t_{i-1} + \omega_t]\}$ where $t_j = start_time + j \cdot \delta_t$
- **Inner Loop (Lines 7-18):** For message m_k in window $[t_i, t_i + \omega_t]$, the subgraph G_{sub} :
 - Contains all nodes $\{id_1, \dots, id_k\}$ from messages $\{m_1, \dots, m_k\}$
 - Maintains edge weights equal to observed message sequence frequencies

2) **Verification Steps:**

- **Initialization:**
 - Outer: $G = \emptyset$ before first window (trivially satisfies)

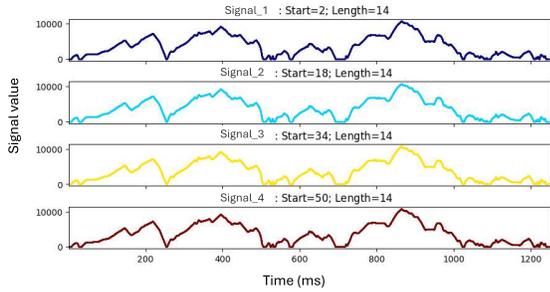


Fig. 4. Speed signals of four wheels encoded in node ID 1760. X-axis: Time (ms); Y-axis: Signal value.

- Inner: $G_{sub} = \emptyset$ before first message (trivially satisfies)
- **Maintenance:**
 - Outer: Each iteration processes one complete window and merges G_{sub} into G
 - Inner: Each message updates G_{sub} 's nodes/edges atomically
- **Termination:**
 - Outer: Terminates when $current_time + \omega_t > end_time$ with all complete windows processed
 - Inner: Processes all messages in each window before continuing

3) Time Complexity:

- **Window processing:**
 - $\lceil \frac{n}{\delta_t} \rceil$ windows processed
 - Each window requires $O(\omega_t)$ operations (message extraction + graph updates)
- **Edge operations:**
 - Average-case $O(1)$ for edge operations (`networkx` default implementation)
 - Each message generates exactly one edge update
 - Worst-case $O(k)$ per operation with hash collisions (k =nodes)
- **Total complexity:**
 - $O\left(\frac{n}{\delta_t} \cdot \omega_t\right)$ for general case
 - $O(n)$ when $\omega_t = \delta_t$ (each message processed once)

Our experimental validation confirms this analysis:

- Time-based windows: $\omega_t \in [2, 15]$ s, $\delta_t \in [1, 15]$ s (all combinations)
- Sample-based windows: $\omega_s = \delta_s \in \{50k \mid k = 1, \dots, 8\}$ messages
- Linear scaling confirmed for $\omega/\delta \in [0.5, 2.0]$

B. Time Series Extraction and Analysis

Concurrently with the creation of the MSGs, raw CAN messages are decoded using a DBC as described in [34]. In doing so, CAN data is transformed into a time series format to capture detailed patterns of the signals contained in every CAN ID per sliding window. We used CAN-D [35] to extract timeseries from CAN logs. Up to the time of this writing, CAN-D is still the state-of-the-art method for CAN reverse

engineering [36]. Fig. 4 shows a decoded representation of the signals representing the four wheels' speed of a vehicle [5]. The plot is generated by translating the CAN frames into a continuous time series format, which captures the network activity over time. This representation facilitates a detailed characterization of masquerade attacks within CANs highlighting variations and trends that may not be evident relying on raw CAN data.

C. Graph Annotation

Algorithm 3 Node Annotation With Statistical Attributes

Require: CAN data \mathcal{D} , sliding windows \mathcal{W} , signal mappings \mathcal{S}
Ensure: Annotated graphs $\{G_t\}$ with statistical node features

- 1 **for** each window $w \in \mathcal{W}$ **do**
- 2 **for** each signal $s \in \mathcal{S}$ **do**
- 3 Extract signal values v_s from \mathcal{D} within window w
- 4 Interpolate missing values in v_s using linear interpolation
- 5 Compute mean μ_s and standard deviation σ_s of v_s
- 6 Annotate corresponding node in G_w with (μ_s, σ_s)
- 7 **end for**
- 8 **end for** **return** Annotated graphs $\{G_t\}$

We enhance each MSG by annotating its nodes with statistical attributes derived from corresponding time series data. For each CAN ID, we compute the mean and standard deviation of its associated signals—features chosen for their simplicity and constant-time complexity, as demonstrated in prior work [3]. This annotation occurs alongside MSG construction, allowing the model to jointly capture structural and temporal dynamics. Fig. 5 illustrates the annotation process, and Algorithm 3 provides implementation details.

D. Graph Embeddings

To model the complex interactions in CANs, we leverage graph-based techniques, specifically graph embeddings, which map network structures into low-dimensional vectors for ML.

Graph embeddings map nodes, edges, or entire graphs into a lower-dimensional vector space, preserving their structural and relational information [37], [38]. In our study, we focus on shallow graph embeddings, specifically the *node2vec* algorithm [37], which is well-suited for capturing the nuances of weighted and directed graphs characteristic of MSGs.

node2vec generates embeddings by simulating random walks on the graph, controlled by two key parameters: the *return parameter* p (likelihood of revisiting a node) and the *in-out parameter* q (preference for exploring local vs. distant neighborhoods). These parameters allow the algorithm to balance the exploration of immediate connectivity and broader network contexts, making it particularly effective for modeling CAN bus data [39].

We configure *node2vec* with 64 dimensions (d) to balance computational efficiency and feature richness [40], a walk length (l) of 15 to capture multi-hop relationships, and 100

walks per node (r) to ensure thorough exploration of the network. Additionally, we set the return parameter (p) to 1.5 and the in-out parameter (q) to 0.5, emphasizing local ECU interactions while retaining global CAN topology. This configuration ensures that the embeddings capture both the immediate communication patterns between ECUs and the broader structure of the CAN.

Unlike prior works that focus solely on node embeddings [2], [20], we construct whole-graph embeddings by averaging node embeddings within sliding windows. This aggregation process provides a unified representation of the network's behavior, capturing both structural and temporal patterns of CAN frames. Specifically, we stack the embeddings per node in the MSGs, labeling each graph to indicate normal or anomalous activity in the CAN. This approach allows us to frame our intrusion detection problem as a graph classification task [41], combining both structural and temporal information for enhanced detection capabilities.

The relatively low embedding dimension (64) has been shown to be sufficient for capturing rich graph-related properties while maintaining computational efficiency [40], which is critical for near real-time detection of masquerade attacks in CANs. By enriching the embeddings with time-series features (e.g., mean, standard deviation) from decoded signals, we ensure a context-aware analysis of masquerade attacks, further enhancing the robustness of our framework.

The result of this process is a data frame that not only encapsulates node embeddings but also includes the augmented node attribute features with labels indicative of normal or anomalous activity states in the CAN. Before classification, we create a comprehensive feature vector for each graph by concatenating the averaged graph embeddings with the time series statistical features (mean and standard deviation of signals). This representation combines both structural (graph) and temporal (time series) information, providing a robust input for ML classifiers. These comprehensive feature vectors serve as the foundation for our supervised learning approach to masquerade attack detection.

E. Supervised Learning

To evaluate the effectiveness of our feature representation and overall anomaly detection framework, we employ supervised learning techniques. We experiment with two ensemble learning methods: Random Forest (RF) [42] and XGBoost [43]. These decision tree-based ensemble learning methods are renowned for their high accuracy, scalability, and robustness, making them exceptionally suitable for dealing with the high-dimensional space characteristic of CAN data. The RF model is instantiated with 100 trees, and depth is capped at 20 to prevent overfitting. A combination of a minimum samples split of 6 and minimum samples leaf of 2 is used to control the growth of the trees further. As the training dataset exhibits a 1:10 ratio between the minority (attack) and majority (normal) classes, we balance the dataset using the synthetic minority over-sampling technique (SMOTE) [44]. A 5-fold stratified k -fold cross-validation is performed to ensure the model's generalizability during training. The XGBoost model yielded similar results to the RF classifier in our experiments. Despite

the comparable performance, we opted to report only results for the RF classifier due to its simpler interpretability and less intensive hyperparameter tuning process, which aligns well with the goals of this study.

F. Evaluation Setup

Evaluation is performed at the sliding window level. In doing so, we verify the proportion of sliding windows coinciding with an attack region in the testing captures, which refers to the recorded CAN bus data used for testing. We found that for the combinations of (ω_t, δ_t) or (ω_s, δ_s) , the proportion of windows containing attacks tended to be balanced, i.e., greater than 40%. To quantitatively assess the performance of our anomaly detection framework, we use the area under the receiver operating characteristic curve (AUC-ROC) [45]. The AUC-ROC is a comprehensive evaluation metric for classification problems across various threshold settings. It reflects the model's ability to distinguish between classes, with higher values indicating better performance. The ROC curve is a plot of the true positive rate (TPR) against the false positive rate (FPR) at different threshold levels. The TPR, also known as sensitivity or recall, measures the proportion of actual positives correctly identified, while the FPR measures the proportion of actual negatives incorrectly identified as positive. Mathematically, the AUC-ROC is defined as: $\text{AUC-ROC} = \int_0^1 \text{TPR}(\text{FPR}^{-1}(x)) dx$.

In practical terms, the AUC-ROC can be interpreted as the probability that the classifier will rank a randomly chosen positive instance higher than a randomly chosen negative instance. This interpretation is particularly useful in our context, where we aim to detect subtle masquerade attacks embedded within normal CAN traffic.

Performance Metrics: We evaluated computational performance using the testing time per window (TTW) metric by Nichelini et al. [14], which quantifies the amount of time required to process each window across attack scenarios. We define TTW as:

$$\text{TTW} = \frac{\text{Total Detection Time}}{\text{Number of Windows}}$$

TTW provides a practical means of evaluating the computational feasibility of our detection framework in real-time settings. All experiments were conducted on a 12th Gen Intel Core i9-12900HK CPU with an NVIDIA RTX 3080 Ti GPU and 32GB RAM.

G. Dataset

Our framework is evaluated on the Real ORNL Automotive Dynamometer (ROAD) dataset [5], collected from a physical vehicle at Oak Ridge National Laboratory (ORNL). This dataset includes physically verified fabrication and simulated masquerade attacks, offering a realistic benchmark for CAN security methods. It is distinguished by its high-quality CAN data and credible attack scenarios across diverse driving conditions. The dataset comprises 3.5 hours of recorded data—3 hours for training and 30 minutes for testing. The test data incorporates five masquerade attacks including correlated signal, max engine coolant, max speedometer, reverse light off/on

TABLE I
COMPARISON OF RELATED WORKS ON INTRUSION DETECTION IN CANS

Reference	Graph-based	Graph embeddings	Time series features	Node annotation	Dataset	Evaluation metrics	Public code/Data
Islam et al. [2020] [20]	✓	✗	✗	✗	HCRL	Misclassification rate	✓
Hanselmann et al. [2020] [13]	✗	✗	✓	✗	Custom	AUC-ROC	✗
Jedh et al. [2021] [2]	✓	✗	✗	✗	Custom	Accuracy, Detection time	✗
Islam et al. [2022] [21]	✓	✗	✗	✗	OpelAstra	Accuracy, Training time	✗
Sreelekshmi et al. [2022] [22]	✓	✗	✗	✗	Car Hacking 2020	Accuracy	✓
Refat et al. [2022] [23]	✓	✗	✗	✗	Real CAN Data	Accuracy	✗
Shahriar et al. [2022] [24]	✗	✗	✓	✗	ROAD, SynCAN	AUC-ROC	✓
Zhang et al. [2023] [17]	✓	✓(GNN)	✗	✓	Ford Transit 500	Accuracy	✗
Park et al. [2023] [25]	✓	✗	✗	✗	Car Hacking 2020	Accuracy	✗
Xiao et al. [2023] [11]	✓	✓(GAT)	✗	✗	HCRL	Accuracy	✗
Meng et al. [2023] [26]	✓	✗	✗	✗	OTIDS	Accuracy	✗
Moriano et al. [2024] [12]	✗	✗	✓	✗	ROAD	AUC-ROC, TTW	✓
Our Work	✓	✓(node2Vec)	✓	✓	ROAD	AUC-ROC, FPR, FNR, TTW	✓

TABLE II
COMPARISON OF RELATED WORKS ON INTRUSION DETECTION IN CANS

Attack Name	Duration (s)	Injection Interval (s)	Description
correlated_signal_1	33.10	[9.19, 30.05]	Injects varying values for wheel speeds, causing the vehicle to halt.
correlated_signal_2	28.23	[6.83, 28.23]	
correlated_signal_3	16.96	[4.32, 16.96]	
max_speedometer_1	88.02	[42.01, 66.45]	Injects the maximum value to be displayed on the speedometer.
max_speedometer_2	59.70	[16.01, 47.41]	
max_speedometer_3	86.77	[9.52, 70.59]	
max_engine_coolant	25.88	[19.98, 24.17]	Injects the maximum value, triggering the coolant warning light.
reverse_light_off_1	28.11	[16.63, 23.35]	Toggles the reverse light irrespective of the actual gear position.
reverse_light_off_2	40.67	[13.17, 36.88]	
reverse_light_off_3	57.88	[16.52, 40.86]	
reverse_light_on_1	54.85	[18.93, 38.84]	Toggles the reverse light irrespective of the actual gear position.
reverse_light_on_2	72.02	[20.41, 57.30]	
reverse_light_on_3	64.26	[23.07, 46.58]	

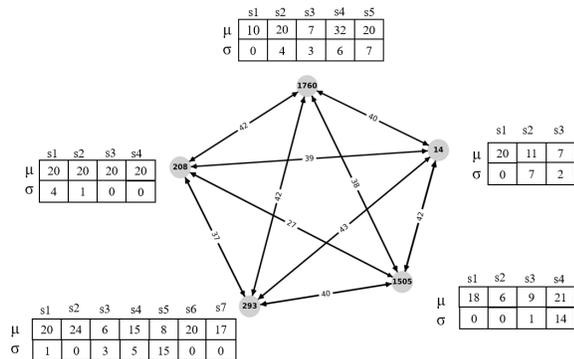


Fig. 5. Illustration of graph annotation with statistical attributes in CAN. Each node represents a CAN ID and is annotated with the mean (μ) and standard deviation (σ) of each signal associated with it. Note that the number of signals may vary per node.

attacks. These attacks are designed to manipulate specific vehicle states. This dataset enables us to conduct a comprehensive evaluation of our framework under a variety of attack conditions. Table II summarizes the durations, injection intervals, and impacts of the masquerade attacks.

V. RESULTS

We evaluate our framework on the ROAD dataset [5] under two settings: (1) a baseline using graph embeddings only and (2) an enhanced configuration combining graph embeddings with time series features. All methods were implemented in Python 3.8.18, using NetworkX for graph operations, sklearn.metrics for evaluation, node2vec 0.4.6 for

embeddings, and timeit for computing TTW. To support reproducibility, our code is publicly available at [15].

Section V-A shows detection results using a baseline model with graph embeddings only, demonstrating significant variation in performance across different ω_t and δ_t combinations using time-based windows (see Fig. 6). Section V-B evidences how incorporating time series features with graph embeddings enhances detection capabilities, improving AUC-ROC values across various attacks (see Fig. 7). Section V-C provides a comparative analysis of detection settings (see Table III) and also presents TTW as a proxy for computational overhead across attack types (see Table V). This analysis supports the feasibility of our approach for time-sensitive applications and highlights trade-offs introduced by different window configurations. Section V-D presents a comparison with state-of-the-art (SOTA) methods, contextualizing our method's performance against existing approaches (see Table VIII). Finally, Section V-E investigates the effect of using more granular windows based on the number of samples. We show that even under these fine-grained configurations, the framework maintains strong detection performance (see Table IX).

A. Baseline Model With Graph Embeddings Only

Our evaluation begins with the establishment of a baseline model that relies solely on graph embeddings, i.e., focusing only on the topology of the MSGs. This baseline model provides a fundamental comparison point and is essential for assessing the effectiveness of incorporating additional features. Fig. 6 shows the performance metrics for this baseline serving as the benchmark for our subsequent enhancements. Specifically, the AUC-ROC values for correlated_signal_attack reach an optimal performance of 0.98 with a ω_t of 8 seconds and an δ_t of 1 second, showcasing strong model effectiveness in this configuration. In contrast, reverse_light_on_attack achieves its highest AUC-ROC of 0.97 with $\omega_t = 7$ seconds and $\delta_t = 4$ seconds, respectively, demonstrating its dependence on extended time frames. In general for the remaining attacks, we notice that high classification results can be obtained for ω_t and δ_t of several seconds. These insights confirm that adjusting the configurations to suit different types of attacks is crucial for maximizing detection capabilities. The variance in performance across different settings shows the importance

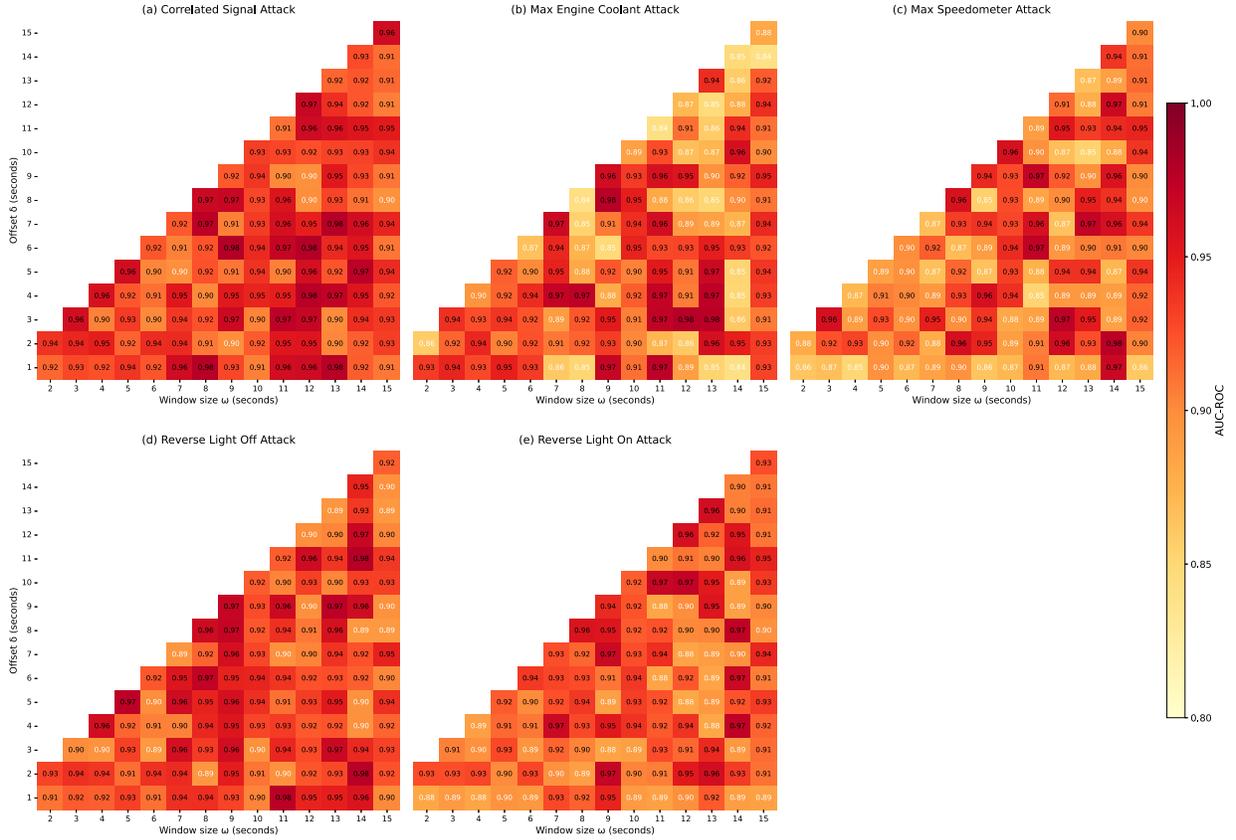


Fig. 6. AUC-ROC across different ω_t and δ_t combinations for masquerade attacks detected by the baseline model using *only* graph embeddings.

of fine-tuning the sliding window parameters to enhance the efficacy of anomaly detection systems in vehicular networks.

B. Enhanced Model With Graph Embeddings and Time Series

Building upon our baseline, we incorporate time series features of CAN signals alongside the graph embeddings. This setting harnesses both the structural patterns and the temporal dynamics of CAN frames aiming to achieve a more comprehensive anomaly detection. Fig. 7 depicts the performance metrics of this enhanced model. The inclusion of time series data significantly improves detection capabilities, as evidenced by the AUC-ROC scores. For *correlated_signal_attack*, the model achieves an AUC-ROC peak of 0.99 at a ω_t of 3 seconds and an δ_t of 3 seconds. This suggests a strong model response to quick changes in CAN signal patterns. Similarly, *reverse_light_on_attack* records its highest AUC-ROC at 0.99 for a ω_t of 7 seconds and an δ_t of 4 seconds, indicating effective monitoring of critical engine parameters. In general, we notice that when integrating time series features, the required time to get optimal performance is lower than when using only graph embeddings. Specifically, across the five masquerade attacks, the average optimal ω_t decreases from 9.8 seconds to 6 seconds, while the average optimal δ_t increases from 3.2 seconds to 4.8 seconds when combining graph embeddings with time series features.

Offset Interpretation and Design Rationale: The offset value is not equivalent to a detection delay—rather, it is a mechanism to ensure continuous attack coverage by allowing consecutive windows to overlap. Without this, attacks occurring near the window boundaries could be missed. Increasing the offset (e.g., to 6s) risks skipping attacks entirely, while reducing it (e.g., to 1s) significantly increases computational overhead. Our configuration optimally balances attack coverage and processing feasibility while ensuring the IDS remains responsive. This 38.8% reduction in ω_t (from 9.8 s to 6 s) coupled with a 50% increase in δ_t (from 3.2 s to 4.8 s) allows for more granular analysis of CAN messages while spacing the analyses further apart, enabling our method to capture detailed patterns more effectively.

C. Metrics Summary

We compare the effectiveness of the two previously defined configurations and include a time series-only variant as an ablation baseline to assess the individual contribution of temporal features.

Table III summarizes results across heatmaps. Attack categories are listed in the columns, and detection settings in the rows. Each cell in this table reports summary statistics from the heatmaps for a specific setting and attack type. We focus on metrics including mean (μ), standard deviation (σ), median (η), minimum (min), and maximum (max). Additionally, we compute the average ($\overline{\max}$) and standard deviation (σ_{\max}) of

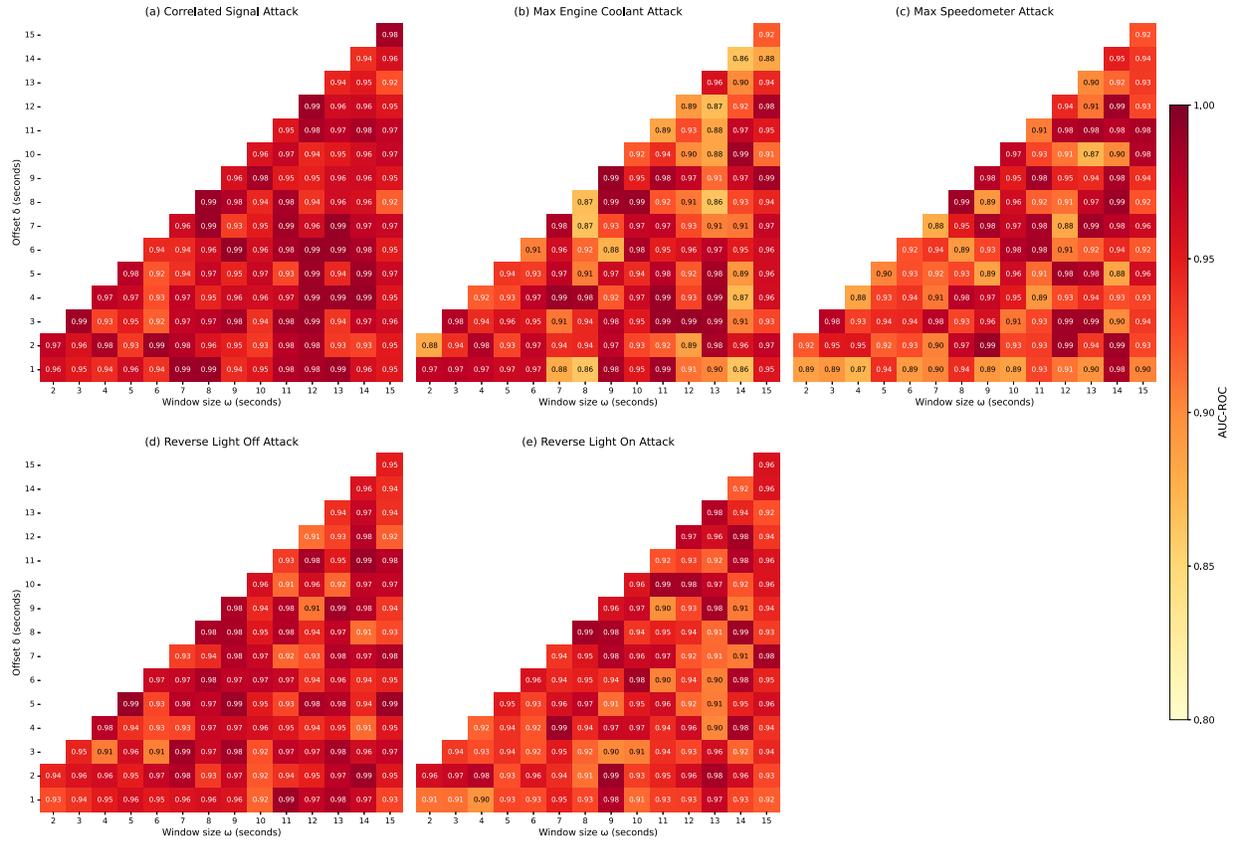


Fig. 7. AUC-ROC across different ω_t and δ_t combinations for masquerade attacks detected by incorporating *both* graph embeddings and time series features.

TABLE III

COMPARISON OF DETECTION SETTINGS BASED ON AUC-ROC. THE TABLE REPORTS THE MEAN (μ), STANDARD DEVIATION (σ), MEDIAN (η), MINIMUM (MIN), AND MAXIMUM (MAX) AUC-ROC VALUES FOR EACH ATTACK SCENARIO. IT ALSO INCLUDES THE AVERAGE ($\overline{\text{max}}$) AND STANDARD DEVIATION (σ_{max}) OF THE MAXIMUM VALUES OBTAINED ACROSS ALL EXPERIMENT TYPES—GRAPH EMBEDDINGS ONLY, TIME SERIES ONLY, AND THE HYBRID APPROACH

Experiment	Correlated signal	Max speedometer	Max engine coolant	Reverse light off	Reverse light on	Experiment max
Graph embeddings only	$\mu = 0.94$ $\sigma = 0.02$ $\eta = 0.93$ min = 0.90 (15,8) max = 0.98 (8,1)	$\mu = 0.91$ $\sigma = 0.03$ $\eta = 0.91$ min = 0.85 (4,1) max = 0.98 (14,2)	$\mu = 0.91$ $\sigma = 0.04$ $\eta = 0.92$ min = 0.84 (15,14) max = 0.98 (9,8)	$\mu = 0.93$ $\sigma = 0.02$ $\eta = 0.93$ min = 0.89 (6,3) max = 0.98 (11,1)	$\mu = 0.92$ $\sigma = 0.03$ $\eta = 0.92$ min = 0.88 (2,1) max = 0.97 (7,4)	$\overline{\text{max}} = 0.98$ $\sigma_{\text{max}} = 0.00$
Time series features only	$\mu = 0.82$ $\sigma = 0.05$ $\eta = 0.83$ min = 0.76 (13,7) max = 0.90 (10,4)	$\mu = 0.85$ $\sigma = 0.04$ $\eta = 0.86$ min = 0.80 (9,6) max = 0.91 (11,5)	$\mu = 0.79$ $\sigma = 0.06$ $\eta = 0.80$ min = 0.72 (14,9) max = 0.89 (12,5)	$\mu = 0.81$ $\sigma = 0.05$ $\eta = 0.83$ min = 0.75 (13,8) max = 0.88 (9,4)	$\mu = 0.84$ $\sigma = 0.03$ $\eta = 0.85$ min = 0.79 (8,5) max = 0.90 (10,5)	$\overline{\text{max}} = 0.90$ $\sigma_{\text{max}} = 0.01$
Graph embeddings + time series features	$\mu = 0.96$ $\sigma = 0.02$ $\eta = 0.96$ min = 0.92 (6,3) max = 0.99 (3,3)	$\mu = 0.94$ $\sigma = 0.03$ $\eta = 0.93$ min = 0.87 (13,10) max = 0.99 (8,8)	$\mu = 0.94$ $\sigma = 0.04$ $\eta = 0.94$ min = 0.86 (8,1) max = 0.99 (7,4)	$\mu = 0.96$ $\sigma = 0.02$ $\eta = 0.96$ min = 0.91 (6,3) max = 0.99 (5,5)	$\mu = 0.94$ $\sigma = 0.03$ $\eta = 0.94$ min = 0.90 (4,1) max = 0.99 (7,4)	$\overline{\text{max}} = 0.99$ $\sigma_{\text{max}} = 0.00$
Attack max	$\overline{\text{max}} = 0.98$ $\sigma_{\text{max}} = 0.01$	$\overline{\text{max}} = 0.98$ $\sigma_{\text{max}} = 0.01$	$\overline{\text{max}} = 0.98$ $\sigma_{\text{max}} = 0.01$	$\overline{\text{max}} = 0.98$ $\sigma_{\text{max}} = 0.01$	$\overline{\text{max}} = 0.98$ $\sigma_{\text{max}} = 0.01$	

the maximum values obtained by each setting and attack type, which are shown in the last row and column of the table.

Graph embeddings combined with time series features consistently outperform both the graph embeddings-only and the time series-only baselines across all attack types. For example, for the *correlated_signal_attack*, the hybrid setting achieves a mean AUC-ROC (μ) of 0.96 with a standard deviation (σ) of 0.02, compared to 0.94 and 0.02 for the graph-only baseline, and 0.82 and 0.05 for the time-series-only

baseline. Similarly, for *max_engine_coolant_attack*, the hybrid setting achieves a mean AUC-ROC of 0.94 and standard deviation of 0.04, versus 0.91 and 0.04 for graph-only, and 0.79 and 0.06 for time-series-only. A comparable improvement is seen across *max_speedometer_attack* (0.94 vs. 0.91 vs. 0.85), *reverse_light_off_attack* (0.94 vs. 0.92 vs. 0.84), with all standard deviations remaining within similar bounds.

TABLE IV

MEAN AND STANDARD DEVIATION OF FPR AND FNR AT THE OPTIMAL ROC-DERIVED THRESHOLD (YOUDEEN’S J INDEX). “GRAPH EMBEDDINGS” REFERS TO THE BASELINE; “+ TIME SERIES FEATURES” DENOTES THE ENHANCED SETTING

Attack Name	Version	FPR (mean \pm std)	FNR (mean \pm std)
correlated_signal	Graph embeddings	0.04 \pm 0.01	0.09 \pm 0.02
	+ time series features	0.02 \pm 0.01	0.07 \pm 0.02
max_speedometer	Graph embeddings	0.01 \pm 0.01	1.20 \pm 0.22
	+ time series features	0.00 \pm 0.01	0.80 \pm 0.17
max_engine_coolant	Graph embeddings	0.02 \pm 0.01	0.27 \pm 0.05
	+ time series features	0.01 \pm 0.01	0.21 \pm 0.05
reverse_light_off	Graph embeddings	0.06 \pm 0.01	0.01 \pm 0.01
	+ time series features	0.04 \pm 0.01	0.01 \pm 0.01
reverse_light_on	Graph embeddings	0.01 \pm 0.01	0.87 \pm 0.19
	+ time series features	0.00 \pm 0.01	0.69 \pm 0.14

TABLE V

TTW SUMMARY ACROSS ATTACK TYPES WITH MIN/MAX CONFIGURATIONS

Attack Name	Mean (s)	Std (s)	Median (s)	Min TTW (s)	Max TTW (s)
correlated_signal	2.36	1.21	2.44	0.024 (12.0, 1.0)	6.84 (10.0, 1.0)
max_speedometer	2.67	1.29	2.65	0.005 (13.0, 4.0)	20.82 (13.0, 4.0)
max_engine_coolant	2.30	1.16	2.30	0.237 (10.0, 5.0)	4.89 (15.0, 9.0)
reverse_light_off	2.60	1.29	2.59	0.027 (10.0, 1.0)	6.22 (15.0, 14.0)
reverse_light_on	2.80	1.36	2.93	0.005 (13.0, 1.0)	8.53 (9.0, 2.0)

TABLE VI

TRADE-OFF BETWEEN DETECTION ACCURACY AND RUNTIME EFFICIENCY FOR SELECTED (ω_t, δ_t) CONFIGURATIONS

Attack Name	ω_t (s)	δ_t (s)	AUC-ROC	TTW (s)
correlated_signal	3	3	0.99	0.44
max_speedometer	8	8	0.99	1.85
max_engine_coolant	7	4	0.99	0.92
reverse_light_off	5	5	0.99	0.73
reverse_light_on	7	4	0.99	1.10
	4	2	0.96	0.18

TABLE VII

MANN-WHITNEY U AND KOLMOGOROV-SMIRNOV TEST RESULTS FOR TIME-BASED WINDOWS

Attack Name	U Statistic	p -value	KS Statistic	p -value
correlated_signal	11337.0	5.55e-16	0.46	3.81e-12
max_speedometer	10082.0	7.98e-09	0.33	2.34e-06
max_engine_coolant	9864.0	8.01e-08	0.33	2.34e-06
reverse_light_off	11033.0	4.98e-14	0.43	1.73e-10
reverse_light_on	10818.0	9.84e-13	0.43	1.73e-10

How do spoofed frames differ in feature domains?

In feature space, spoofed frames differ from benign traffic primarily through anomalies in structural transitions and temporal regularities [12]. Our evaluation shows that the fused feature representation enables near-perfect separation between benign and malicious traffic. Specifically, our hybrid model achieves a mean AUC-ROC of 0.99 across all masquerade attacks, indicating that the combined structural and temporal deviations are sufficient for discrimination, even when content and timing are superficially plausible. Notably, difficult-to-detect attacks such as max_engine_coolant, previously noted for weak signal-level deviation, are still effectively detected by our

TABLE VIII

COMPARATIVE PERFORMANCE OF OUR HYBRID APPROACH VERSUS SOTA METHODS ON THE ROAD DATASET. AUC-ROC AND TTW VALUES CORRESPOND TO EACH METHOD’S BEST-PERFORMING CONFIGURATION PER ATTACK

Attack Name	Metric	Ours	Moriano	CANShield
correlated_signal	AUC-ROC	0.99	0.96	0.93
	TTW (s)	0.44	0.67	N/A
max_speedometer	AUC-ROC	0.99	0.97	0.94
	TTW (s)	1.85	2.01	N/A
max_engine_coolant	AUC-ROC	0.99	0.94	0.91
	TTW (s)	0.92	1.04	N/A
reverse_light_off	AUC-ROC	0.99	0.95	0.92
	TTW (s)	0.73	0.89	N/A
reverse_light_on	AUC-ROC	0.99	0.93	0.90
	TTW (s)	1.10	1.36	N/A
Average	AUC-ROC	0.99	0.95	0.92
	TTW (s)	0.97	1.19	N/A

[†]CANShield does not explicitly define (ω_t, δ_t) values but adopts an equivalent segmentation strategy using window_step_test parameters; TTW is omitted as its framework does not report latency per window.

TABLE IX

MANN-WHITNEY U AND KOLMOGOROV-SMIRNOV TEST RESULTS FOR SAMPLE-BASED WINDOWS

Attack Name	U Statistic	p -value	KS Statistic	p -value
correlated_signal	1069.5	9.84e-07	0.58	2.61e-06
max_speedometer	1085.0	4.20e-07	0.56	9.30e-06
max_engine_coolant	811.5	0.033	0.33	0.018
reverse_light_off	1040.5	4.84e-06	0.50	9.36e-05
reverse_light_on	1112.0	8.35e-08	0.61	6.75e-07

approach. This supports the conclusion that our hybrid features amplify subtle inconsistencies—particularly in message transition structure and short-term temporal signal patterns—that standalone models may miss. Prior CAN intrusion studies also suggest that such composite feature sets yield superior discrimination [46], [47].

We now summarize median and maximum statistics for each configuration. In correlated_signal_attack, the hybrid setting yields a median (η) of 0.96, a minimum (min) of 0.92, and a maximum (max) of 0.99. These statistics capture the central tendency and spread of detection performance across all configurations. Notably, the time-series-only baseline consistently underperforms in terms of both average and peak metrics, which further validates the necessity of combining temporal and structural features for detecting stealthy attacks. Moreover, we observe that for max_engine_coolant_attack, the average maximum AUC-ROC ($\bar{\max}$) reaches 0.99 for the hybrid model, 0.98 for the graph-only baseline, and only 0.89 for the time-series-only variant, with a standard deviation of $(\sigma_{\max}) = 0.00$ or 0.01—highlighting both the performance gap and stability offered by the hybrid method across experimental trials.

To provide clearer insights into the IDS’s classification performance, we computed the false positive rate (FPR) and false negative rate (FNR) at a selected operating threshold derived from the receiver operating characteristic (ROC) curve [45]. Specifically, we selected the threshold that maximizes Youden’s J index ($J = \text{TPR} - \text{FPR}$), a standard method for identifying the optimal balance between true positive and false

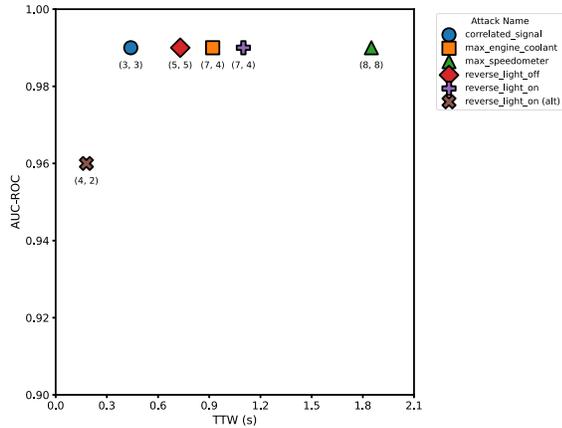


Fig. 8. Trade-off between detection accuracy (AUC-ROC) and runtime cost (TTW) for selected (ω_t, δ_t) configurations. Values in parentheses below each point indicate the ω_t and δ_t used in that configuration. The `reverse_light_on (alt)` point reflects an alternate low-latency configuration for the same attack, included to illustrate the practical trade-off.

positive rates in binary classification tasks [48]. This ensures a meaningful trade-off between sensitivity and specificity, supporting practical deployment. Table IV reports the average FPR and FNR across all attack types and (ω_t, δ_t) combinations, highlighting consistent improvements—especially in FNR—when time series features are integrated.

Table V summarizes TTW results, including mean, median, standard deviation, minimum, and maximum processing times, with corresponding optimal δ_t and ω_t combinations shown in parentheses. Table V demonstrates that our hybrid approach achieves reasonable computational overhead across attack types. For instance, the `correlated_signal_attack` achieves a mean TTW of 2.36s, with a best-case performance as low as 0.024s, highlighting feasibility for real-time application with further optimization.

Table VI presents the AUC-ROC and corresponding TTW for selected (ω_t, δ_t) configurations across all attack types. Each row reflects a configuration that balances high detection accuracy with practical runtime cost. Note that for `reverse_light_on`, the second row highlights a lower-latency alternative with a slight reduction in AUC-ROC, illustrating the effect of window size and offset on real-time feasibility. We show that our hybrid approach consistently achieves near-perfect detection (0.99 AUC-ROC) across all masquerade attacks, while also revealing the varying computational overhead associated with different optimal configurations. For instance, the `correlated_signal` attack is detected with 0.99 AUC-ROC at 0.44s TTW, highlighting its efficiency. This comprehensive view of performance and runtime is critical for deploying IDS in resource-constrained vehicular environments.

Figure 8 visualizes this trade-off by plotting AUC-ROC against TTW, annotated with the corresponding (ω_t, δ_t) values. While most configurations achieve near-identical AUC-ROC scores (0.99), the associated TTW varies, reflecting the computational impact of wider temporal contexts. Note that the alternate `reverse_light_on` configuration, for example, reduces TTW from 1.10s to 0.18s with only a minor drop in

AUC-ROC. We show the importance of selecting (ω_t, δ_t) pairs that offer an optimal balance between detection accuracy and inference latency.

To further validate the differences in detection effectiveness of both settings, we employed the Mann-Whitney U test and the Kolmogorov-Smirnov (KS) test. The Mann-Whitney U test [49] is a non-parametric statistical test that evaluates whether there is a difference between two independent samples. The KS test [50] is another non-parametric test that assesses whether two samples come from the same distribution. In our analysis, we compared the AUC-ROC value distributions of the two settings. Here, the null hypothesis is that the AUC-ROC values for the setting combining embeddings with time series features are less than or equal to those for the setting using only embeddings, while the alternative hypothesis is that the AUC-ROC values for the setting combining embeddings with time series features are greater than those for the setting using only embeddings. A low p -value indicates a significant difference between these two settings. We focused on a significance level of 0.05.

Table VII shows significant differences between the observed distributions for all masquerade attacks. In each of the attacks, the Mann-Whitney U and KS tests produce low p -values, indicating a significant deviation from the expected distribution. This supports our hypothesis of relying on annotated MSGs to effectively detect masquerade attacks at various windows and offset combinations. Consequently, we reject the null hypothesis for all tested masquerade attacks, confirming that the combined setting performs significantly better.

D. Comparison With SOTA Methods

We compare our hybrid approach against two relevant SOTA baselines on the ROAD dataset to further contextualize our method’s performance. Our proposed hybrid method integrates structural graph features with temporal dynamics to enhance anomaly detection. SOTA baselines include:

Moriano et al. [12] — An unsupervised time series approach which is evaluated on the ROAD dataset [5] and reports AUC-ROC and TTW using sliding window parameters (ω_t, δ_t) aligning with our evaluation design. We replicated their best-performing configuration and included results for each masquerade attack.

CANShield [24] — A deep learning-based method that operates on signal-level representations, applied to the ROAD dataset [5]. While effective in detecting various attack classes, it does not report TTW nor explicitly define (ω_t, δ_t) values. Instead, it segments data internally using `window_step_test` parameters. Additionally, due to the high-dimensional signal-level input and deeper architecture, the approach is considerably more computationally intensive on ROAD. As such, TTW is omitted and only AUC-ROC is reported for this baseline.

Table VIII summarizes the attack-wise AUC-ROC and TTW scores. Our approach achieves the highest AUC-ROC (0.99) across all five masquerade attacks while consistently demonstrating lower TTW per window. The optimal (ω_t, δ_t) values for our configuration are presented in Table VI. For

Moriano et al. [12], the values correspond to their best-performing public configuration. We show the effectiveness and practicality of our method for real-time detection, highlighting its robust performance under constrained latency budgets. The code to reproduce these baselines has been added to our GitHub repository [15].

E. Evaluation Based on Sample Windows

We investigate the effect of sample-based windows in our detection framework’s performance, as opposed to time-based windows. Our aim is to understand the impact of varying the number of CAN message samples in each sliding window on masquerade attack detection. In this context, time-based windows refer to windows defined by a fixed duration, whereas sample-based windows are defined by a fixed number of CAN messages. We experiment with different ω_s and δ_s to find the optimal configurations for detecting anomalies.

Due to space constraints, we do not include detailed heatmap results in this version of the paper. However, note that our sample-based results suggest that the integration of time series features with graph embeddings enhances the model’s ability to detect anomalies, even when requiring slightly larger sample windows. The improvement in detection capability comes with a measurable increase in computational requirements, as quantified by the TTW metrics in Section V-C. Further optimization and hardware-specific tuning will be necessary to ensure real-time feasibility in production vehicles.

We also found the distributions of AUC-ROC values to be significantly different between the two settings. The Mann-Whitney U and KS tests confirmed this difference, indicating the superior detection performance of the combined setting. Table IX summarizes the results, showing significant differences across all tested masquerade attacks. For instance, `correlated_signal_attack` yielded a Mann-Whitney U statistic of 1069.5 with a p -value of $9.84e-07$, and a KS statistic of 0.58 with a p -value of $2.61e-06$. Similar trends were observed for other attacks, such as `max_engine_coolant_attack` and `max_speedometer_attack`, confirming the effectiveness of incorporating time series features with graph embeddings. The results for each attack type consistently indicate that the combined setting outperforms the embeddings-only setting.

VI. DISCUSSION

This paper presents a novel approach for detecting masquerade attacks in CAN using graph ML. By combining graph embeddings with time series features, our method significantly improves detection performance across all masquerade attack types in the ROAD dataset. These gains are statistically validated using the Mann-Whitney U and KS tests (see Tables VII and IX). For instance, detection of the `max_engine_coolant_attack`—previously difficult to detect due to weak signal deviation [12]—was notably improved in the enhanced setting, demonstrating stronger anomaly sensitivity. Our evaluation also shows that using sample-based windows—particularly with smaller ω values—enables the detection of subtle anomalies via fine-grained analysis.

What specific deviations does the model capture?

Our hybrid framework captures subtle, cumulative deviations that even sophisticated spoofing cannot perfectly hide. We leverage two complementary domains:

- **Structural deviations (via MSGs and graph embeddings):** Masquerade attacks involve an attacker suppressing a legitimate ECU and impersonating it. Attackers cannot perfectly replicate intricate real-time inter-message dependencies, message orderings, and complex sequential patterns (graph motifs) characteristic of benign ECU communication. Prior work demonstrates motif frequency distributions as sensitive indicators of anomalies [51]. Graph-based approaches effectively capture these structural anomalies in CAN traffic. Our approach leverages `node2vec` to embed these deviations, extracting low-dimensional representations that preserve the complex topology of MSGs.
- **Temporal and content deviations (via timeseries features):** Masquerade attacks can introduce subtle inconsistencies in physical signal values or precise timing patterns, even if overall frequency is matched. Decoded signal values may show anomalies in statistical properties (mean, standard deviation) or inter-signal correlations, which are stable in normal driving but deviate under spoofing [52]. Our node annotation technique enriches CAN ID nodes with these statistics, enabling detection of nuanced temporal and content-level inconsistencies. Similar approaches prove effective for high-frequency signal perturbations in automotive IDS [12], [17].

The strength of our hybrid approach lies in its ability to fuse these structural graph embeddings with the temporal signal statistics. A deviation missed by one modality might be detected by the other, providing a robust, multi-faceted detection capability against attacks that blend seamlessly into normal traffic.

The effectiveness of our approach highlights the value of combining structural and temporal insights in CAN IDS. Graph embeddings capture ECU interactions, while time series features reflect evolving signal patterns. This dual analysis improves detection and allows more efficient sliding window configuration—reducing the optimal offset δ_t by 38.8% (9.8 s to 6 s) and increasing the window length ω_t by 50% (3.2 s to 4.8 s). These adjustments enhance timely detection and emphasize the importance of securing CAN-based systems as vehicle architectures evolve.

While automotive Ethernet adoption grows (projected to reach 11.6 billion US dollars by 2032 [53]), it primarily serves high-bandwidth domains like ADAS and infotainment (38% market share in 2023). However, CAN remains critical for safety subsystems (e.g., engine, braking) due to its role in hybrid architectures and persistent vulnerabilities. Over 90% of modern vehicles use hybrid CAN/Ethernet networks, with CAN handling safety-critical low-speed data (powertrain/chassis) while Ethernet supplements infotainment [53]. CAN’s simplicity and cost ensure its ubiquity in existing fleets, with Ethernet adoption concentrated in high-end models. Despite Ethernet’s growth, CAN lacks encryption/authentication, making it vulnerable to attacks like brake

disablement via OBD-II injection [54]. These risks persist even in hybrid architectures, as compromised Ethernet gateways can inject malicious CAN messages. Our work addresses these gaps with lightweight CAN-specific detection mechanisms. CAN security remains foundational, as even Ethernet-reliant systems interface with CAN for critical functions.

Our work distinguishes itself from prior methods by integrating graph-based and time series analysis to address the subtle nature of masquerade attacks. While existing approaches focus on single-modality detection, our hybrid framework captures both structural and temporal anomalies, enabling the detection of sophisticated intrusions that traditional IDS might miss. This dual approach not only improves detection accuracy but also provides a scalable solution adaptable to the dynamic nature of modern vehicular networks.

Our parameter choices are empirically grounded in the ROAD dataset’s attack profiles, balancing accuracy and practicality. While shorter windows may suit other threat models, masquerade detection necessitates analyzing sustained patterns. For safety-critical functions requiring sub-second detection (e.g., brake disengagement), adaptive parameter tuning and hardware-specific optimizations will be explored in future work. While our approach demonstrates promising detection capabilities, its real-time deployment in resource-constrained vehicular environments requires further validation on embedded platforms (e.g., ARM-based boards). Our current results, based on realistic CAN logs generated on a physical vehicle, focus on validating the detection methodology rather than finalizing an on-vehicle implementation.

While cryptographic measures such as encryption and MACs are effective defenses against masquerade attacks, our anomaly detection framework provides an additional layer of security. By detecting subtle deviations in CAN traffic patterns, our approach complements cryptographic defenses, ensuring robust protection against sophisticated attacks, including those that bypass cryptographic approaches.

While our results are promising, it is important to acknowledge the limitations of this study:

Lack of Validation in a Real Environment: Our study is simulation-based and lacks validation in real vehicle environments, where factors like latency, sensor noise, and varying driving conditions could impact performance. Future work should involve empirical validation on a moving vehicle, including hardware integration and performance tuning.

Evaluation Based Solely on ROAD Dataset: While comprehensive, the ROAD dataset [5] may not capture the full range of real-world attack scenarios. Future studies should incorporate additional datasets with diverse and recent attack patterns to improve robustness and generalizability.

Computational Complexity: The graph embedding phase may introduce computational overhead. Future work should explore algorithmic optimization and hardware acceleration for improved efficiency.

Use of Default Parameter Values in Detection Algorithms: This study relied on default parameters in detection algorithms, which may not yield optimal performance. Future research should involve extensive hyperparameter tuning to maximize detection accuracy in application-specific contexts.

Limits of Detectability: No system is impervious to highly sophisticated attackers. The detectability limits for our framework are as follows:

- **Perfect mimicry:** If an attacker can replicate the full message transition topology and multivariate signal relationships of a legitimate ECU—including physical-layer jitter and temporal correlations—the detection margin narrows. However, this level of mimicry requires access to proprietary ECU timing, bus arbitration, and signal generation logic, which is extremely difficult without insider knowledge [55].
- **Short or bursty attacks:** Transient attacks that operate below the minimum window size or inject very few anomalous transitions may evade detection. This is a known limitation in time-window-based IDS designs [2].
- **Overlapping statistical distributions:** In rare cases, spoofed messages may fall within the natural variability of benign traffic, making them indistinguishable even in a fused feature space. Defense against this class of attack may require incorporating physical-layer fingerprints [46] or entropy-based signal modeling [56].

In summary, our study contributes to the growing body of research on graph-based CAN IDS by demonstrating the effectiveness of combining graph embeddings with time series features. This approach not only enhances detection capabilities but also provides a flexible framework adaptable to various vehicular systems and attack patterns.

VII. CONCLUSION

This paper presents a novel framework for detecting masquerade attacks in CANs by integrating graph ML with time series analysis. By representing CAN frames as MSGs and enriching each node with contextual statistical features, the proposed approach captures both the structural and temporal dynamics of in-vehicle communication. Experiments on the ROAD dataset [5] show that this hybrid model consistently outperforms graph-only baselines across all tested masquerade attack types. These improvements are statistically validated through the Mann-Whitney U and Kolmogorov-Smirnov tests (see Tables VII and IX), confirming the robustness of our detection capabilities across various attack scenarios.

Future work will explore deploying this framework in online detection settings, investigate lightweight graph-based architectures such as those introduced by Zhang et al. [17], and assess its adaptability across diverse datasets. We also aim to explore integration with foundation models for enhancing generalization in evolving vehicular environments.

ACKNOWLEDGMENT

This manuscript has been authored by UT-Battelle, Limited Liability Company (LLC) under Contract No. DE-AC05-00OR22725 with the U.S. Department of Energy (DOE). The publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a non-exclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for U.S. Government purposes. The DOE will provide

public access to these results in accordance with the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan>). This research was sponsored in part by Oak Ridge National Laboratory's Laboratory Directed Research and Development program through the Sustainable Research Pathways (SRP) program and by the U.S. Department of Energy. This work was also partially funded under U.S. Department of Energy Award No. DE-FE0032089. No additional external funding was received. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

REFERENCES

- [1] S. Halder, M. Conti, and S. K. Das, "COIDS: A clock offset based intrusion detection system for controller area networks," in *Proc. 21st Int. Conf. Distrib. Comput. Netw.*, Jan. 2020, pp. 1–10.
- [2] M. Jedh, L. Ben Othmane, N. Ahmed, and B. Bhargava, "Detection of message injection attacks onto the CAN bus using similarities of successive messages-sequence graphs," *IEEE Trans. Inf. Forensics Security*, vol. 16, pp. 4133–4146, 2021.
- [3] M. H. Shahriar, W. Lou, and Y. T. Hou, "CANtropy: Time series feature extraction-based intrusion detection systems for controller area networks," in *Proc. Inaugural Int. Symp. Vehicle Secur. Privacy*, 2023, pp. 1–8.
- [4] P. Moriano, R. A. Bridges, and M. D. Iannacone, "Detecting CAN masquerade attacks with signal clustering similarity," in *Proc. 4th Int. Workshop Automot. Auto. Vehicle Secur.*, 2022, pp. 1–8.
- [5] M. E. Verma et al., "A comprehensive guide to CAN IDS data and introduction of the ROAD dataset," *PLoS ONE*, vol. 19, no. 1, Jan. 2024, Art. no. e0296879.
- [6] C. Miller and C. Valasek, "Remote exploitation of an unaltered passenger vehicle," in *Proc. Black Hat USA*, 2015, pp. 1–91.
- [7] S. Sharmin, H. Mansor, A. F. A. Kadir, and N. A. Aziz, "Comparative evaluation of anomaly-based controller area network IDS," in *Proc. 12th Int. Conf. Softw. Comput. Appl.*, Feb. 2023, pp. 218–226.
- [8] W. Marfo, D. K. Tosh, and S. V. Moore, "Condition monitoring and anomaly detection in cyber-physical systems," in *Proc. 17th Annu. Syst. Syst. Eng. Conf. (SOSE)*, Jun. 2022, pp. 106–111.
- [9] X. Wang, Y. Liu, K. Jiao, P. Liu, X. Luo, and T. Liu, "Intrusion device detection in fieldbus networks based on channel-state group fingerprint," *IEEE Trans. Inf. Forensics Security*, vol. 19, pp. 4012–4027, 2024.
- [10] H. Sun, M. Chen, J. Weng, Z. Liu, and G. Geng, "Anomaly detection for in-vehicle network using CNN-LSTM with attention mechanism," *IEEE Trans. Veh. Technol.*, vol. 70, no. 10, pp. 10880–10893, Oct. 2021.
- [11] J. Xiao, L. Yang, F. Zhong, H. Chen, and X. Li, "Robust anomaly-based intrusion detection system for in-vehicle network by graph neural network framework," *Appl. Intell.*, vol. 53, no. 3, pp. 3183–3206, Feb. 2023.
- [12] P. Moriano, S. C. Hespeler, M. Li, and R. A. Bridges, "Evaluating lightweight unsupervised online IDS for masquerade attacks in CAN," 2024, *arXiv:2406.13778*.
- [13] M. Hanselmann, T. Strauss, K. Dormann, and H. Ulmer, "CANet: An unsupervised intrusion detection system for high dimensional CAN bus data," *IEEE Access*, vol. 8, pp. 58194–58205, 2020.
- [14] A. Nichelimi, C. A. Pozzoli, S. Longari, M. Carminati, and S. Zanero, "CANova: A hybrid intrusion detection framework based on automatic signal classification for CAN," *Comput. Secur.*, vol. 128, May 2023, Art. no. 103166.
- [15] W. Marfo. (2024). *GraphML Controller Area Network Attack Detection*. [Online]. Available: <https://github.com/billmj/GraphML-CONTROLLER-AREA-NETWORK-Attack-Detection>
- [16] M. D. Natale, H. Zeng, P. Giusto, and A. Ghosal, *Understanding and Using the Controller Area Network Communication Protocol: Theory and Practice*. New York, NY, USA: Springer, 2012.
- [17] H. Zhang, K. Zeng, and S. Lin, "Federated graph neural network for fast anomaly detection in controller area networks," *IEEE Trans. Inf. Forensics Security*, vol. 18, pp. 1566–1579, 2023.
- [18] A. Gazdag, R. Ferenc, and L. Buttyán, "CrySys dataset of CAN traffic logs containing fabrication and masquerade attacks," *Sci. Data*, vol. 10, no. 1, p. 903, Dec. 2023.
- [19] S. Rajapaksha, G. Madzudzo, H. Kalutarage, A. Petrovski, and M. O. Al-Kadri, "Can-mirgu: A comprehensive can bus attack dataset from moving vehicles for intrusion detection system evaluation," in *Proc. Symp. Vehicles Secur. Privacy (VehicleSec)*, San Diego, CA, USA, 2024, doi: [10.14722/vehiclesec.2024.23043](https://doi.org/10.14722/vehiclesec.2024.23043).
- [20] R. Islam, R. U. D. Refat, S. M. Yerram, and H. Malik, "Graph-based intrusion detection system for controller area networks," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 3, pp. 1727–1736, Mar. 2022.
- [21] R. Islam, M. K. Devnath, M. D. Samad, and S. M. J. A. Kadry, "GGNB: Graph-based Gaussian naive Bayes intrusion detection system for CAN bus," *Veh. Commun.*, vol. 33, Jan. 2022, Art. no. 100442.
- [22] M. Sreelekshmi and S. Aji, "A graph-based strategy for intrusion detection in connected vehicles," in *Proc. Int. Conf. Inf. Commun. Technol. Competitive Strategies*. Cham, Switzerland: Springer, 2022, pp. 133–143.
- [23] R. U. D. Refat, A. A. Elkhail, and H. Malik, "A lightweight intrusion detection system for CAN protocol using neighborhood similarity," in *Proc. 7th Int. Conf. Data Sci. Mach. Learn. Appl. (CDMA)*, Mar. 2022, pp. 121–126.
- [24] M. H. Shahriar, Y. Xiao, P. Moriano, W. Lou, and Y. T. Hou, "CANShield: Deep-learning-based intrusion detection framework for controller area networks at the signal level," *IEEE Internet Things J.*, vol. 10, no. 24, pp. 22111–22127, Dec. 2023.
- [25] S. B. Park, H. J. Jo, and D. H. Lee, "G-IDCS: Graph-based intrusion detection and classification system for CAN protocol," *IEEE Access*, vol. 11, pp. 39213–39227, 2023.
- [26] Y. Meng, J. Li, F. Liu, S. Li, H. Hu, and H. Zhu, "GB-IDS: An intrusion detection system for CAN bus based on graph analysis," in *Proc. IEEE/CIC Int. Conf. Commun. China (ICCC)*, Aug. 2023, pp. 1–6.
- [27] H. Lee, S. H. Jeong, and H. K. Kim, "CAN intrusion detection dataset," *IEEE DataPort*, 2025, doi: [10.21227/sy7h-9057](https://doi.org/10.21227/sy7h-9057).
- [28] H. M. Song, J. Woo, and H. K. Kim, "In-vehicle network intrusion detection using deep convolutional neural network," *Veh. Commun.*, vol. 21, Jan. 2020, Art. no. 100198.
- [29] L. B. Othmane, L. Dhulipala, M. Abdelkhalik, N. Multari, and M. Govindarasu, "On the performance of detecting injection of fabricated messages into the CAN bus," *IEEE Trans. Dependable Secure Comput.*, vol. 19, no. 1, pp. 468–481, Jan. 2022.
- [30] H. Lee, S. H. Jeong, and H. K. Kim, "OTIDS: A novel intrusion detection system for in-vehicle network by using remote frame," in *Proc. 15th Annu. Conf. Privacy, Secur. Trust (PST)*, Aug. 2017, pp. 57–5709.
- [31] S. Checkoway et al., "Comprehensive experimental analyses of automotive attack surfaces," in *Proc. 20th USENIX Secur. Symp. (USENIX Secur. 11)*, 2011, p. 6.
- [32] T. Qin, Z. Liu, P. Wang, S. Li, X. Guan, and L. Gao, "Symmetry degree measurement and its applications to anomaly detection," *IEEE Trans. Inf. Forensics Security*, vol. 15, pp. 1040–1055, 2020.
- [33] J. Bryan and P. Moriano, "Graph-based machine learning improves just-in-time defect prediction," *PLoS ONE*, vol. 18, no. 4, Apr. 2023, Art. no. e0284077.
- [34] CSS Electronics.(2021). *Understanding DBC Files: Basics for Can Databases*. Accessed: Apr. 23, 2023. [Online]. Available: <https://www.csselectronics.com/pages/can-dbc-file-database-intro>
- [35] M. E. Verma, R. A. Bridges, J. J. Sosnowski, S. C. Hollifield, and M. D. Iannacone, "CAN-D: A modular four-step pipeline for comprehensively decoding controller area network data," *IEEE Trans. Veh. Technol.*, vol. 70, no. 10, pp. 9685–9700, Oct. 2021.
- [36] A. Buscemi, I. Turcanu, G. Castignani, A. Panchenko, T. Engel, and K. G. Shin, "A survey on controller area network reverse engineering," *IEEE Commun. Surveys Tuts.*, vol. 25, no. 3, pp. 1445–1481, 3rd Quart., 2023.
- [37] A. Grover and J. Leskovec, "Node2vec: Scalable feature learning for networks," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, May 2016, pp. 855–864.
- [38] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu, "A comprehensive survey on graph neural networks," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 32, no. 1, pp. 4–24, Jan. 2021.
- [39] Z.-F. Wei, P. Moriano, and R. Kannan, "Robustness of graph embedding methods for community detection," 2024, *arXiv:2405.00636*.
- [40] W. Gu, A. Tandon, Y.-Y. Ahn, and F. Radicchi, "Principled approach to the selection of the embedding dimension of networks," *Nature Commun.*, vol. 12, no. 1, p. 3772, Jun. 2021.
- [41] M. Zhong, M. Lin, C. Zhang, and Z. Xu, "A survey on graph neural networks for intrusion detection systems: Methods, trends and challenges," *Comput. Secur.*, vol. 141, Jun. 2024, Art. no. 103821.
- [42] L. Breiman, "Random forests," *Mach. Learn.*, vol. 45, no. 1, pp. 5–32, 2001.

- [43] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Aug. 2016, pp. 785–794.
- [44] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE: Synthetic minority over-sampling technique," *J. Artif. Intell. Res.*, vol. 16, pp. 321–357, Jun. 2002.
- [45] T. Fawcett, "An introduction to ROC analysis," *Pattern Recognit. Lett.*, vol. 27, no. 8, pp. 861–874, Jun. 2006.
- [46] K.-T. Cho and K. G. Shin, "Fingerprinting electronic control units for vehicle intrusion detection," in *Proc. 25th USENIX Secur. Symp. (USENIX Secur.)*, 2016, pp. 911–927.
- [47] S. Tariq, S. Lee, and S. S. Woo, "CANTransfer: Transfer learning based intrusion detection on a controller area network using convolutional LSTM network," in *Proc. 35th Annu. ACM Symp. Appl. Comput.*, Mar. 2020, pp. 1048–1055.
- [48] L. Yang et al., "UnFlowS: An unsupervised construction scheme of flow spectrum for network traffic detection," *IEEE Trans. Inf. Forensics Security*, vol. 20, pp. 3330–3345, 2025.
- [49] H. B. Mann and D. R. Whitney, "On a test of whether one of two random variables is stochastically larger than the other," *Ann. Math. Statist.*, vol. 18, no. 1, pp. 50–60, Mar. 1947.
- [50] V. W. Berger and Y. Zhou, "Kolmogorov–Smirnov test: Overview," in *Wiley StatsRef: Statistics Reference Online*. Hoboken, NJ, USA: Wiley, 2014.
- [51] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon, "Network motifs: Simple building blocks of complex networks," *Science*, vol. 298, no. 5594, pp. 824–827, Oct. 2002.
- [52] H. Zhang, X. Meng, X. Zhang, and Z. Liu, "CANsec: A practical in-vehicle controller area network security evaluation tool," *Sensors*, vol. 20, no. 17, p. 4900, Aug. 2020.
- [53] Global Market Insights. (Mar. 2024). *Automotive Ethernet Market Size, by Component, Vehicle Type, Application, Bandwidth & Forecast, 2024–2032*. [Online]. Available: <https://www.gminsights.com/industry-analysis/automotive-ethernet-market>
- [54] SealingTech. (Mar. 2024). *Can Bus Vulnerabilities Raise Threats for Cars and Tanks*. [Online]. Available: <https://www.sealingtech.com/2024/03/05/can-bus-vulnerabilities-raise-threats-for-cars-and-tanks/>
- [55] C. Miller and C. Valasek, "A survey of remote automotive attack surfaces," in *Proc. Black Hat USA*, Las Vegas, NV, USA, 2014.
- [56] M. Marchetti and D. Stabili, "Anomaly detection of CAN bus messages through analysis of ID sequences," in *Proc. IEEE Intell. Vehicles Symp. (IV)*, Jun. 2017, pp. 1577–1583.



William Marfo (Member, IEEE) received the Ph.D. degree in computer science from The University of Texas at El Paso, USA. He is currently an Associate Specialist and a Cyberinfrastructure Software Engineer with the Cascadia Region Earthquake Science Center (CRESCENT), University of California, Davis, USA. He develops scalable software and AI/ML models for large-scale seismic data processing and interdisciplinary earthquake science. He previously conducted research with U.S. Department of Energy national laboratories and was named a

2025 Trusted CI Fellow by U.S. National Science Foundation (NSF) Cybersecurity Center of Excellence. His research interests include machine learning, high-performance computing, and cybersecurity for complex computational systems. He is a member of ACM, SIAM, AMS, and ISACA.



Pablo Moriano (Senior Member, IEEE) received the B.S. and M.S. degrees in electrical engineering from Pontificia Universidad Javeriana, Colombia, and the M.S. and Ph.D. degrees in informatics from Indiana University Bloomington, Bloomington, IN, USA. He is a Staff Scientist with the Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, TN, USA. His research interests include data science, network science, and cybersecurity, with a focus on the design of AI and statistical methods for anomaly detection in large-scale networked and cyber-physical systems. Applications of his work include intrusion detection, insider-threat analytics, and Internet-routing security. He is also a Senior Member of the ACM and a member of SIAM.



Deepak K. Tosh (Senior Member, IEEE) is currently an Associate Professor of Computer Science with The University of Texas at El Paso. His research focuses on addressing various multi-disciplinary networking and cybersecurity challenges associated with critical national infrastructures, network virtualization, Industrial the Internet of Things (IIoT), space cyber-physical systems, Blockchain, and tactical networking. He has authored/co-authored more than 90 peer-reviewed conference papers, book chapters, and journal papers. His research has been supported by the Department of Energy (DOE) and National Science Foundation (NSF) grants. Two of his research works on Blockchain and data provenance received the "Top 50 Blockchain Papers in 2018" award at BlockchainConnect Conference, in 2019. He is also a recipient of the prestigious NSF CAREER Award, in 2023.



Shirley V. Moore (Member, IEEE) received the B.A. degree in mathematics and chemistry from Indiana University, the M.Ed. degree from the University of Illinois, the M.S. degree in mathematics from Wichita State University, and the M.S. and Ph.D. degrees in computer sciences from Purdue University. She is currently an Associate Professor with the Computer Science Department, The University of Texas at El Paso, USA. Her focus is on performance optimization methodologies and tools. Her research is in parallel and high-performance

computing, including distributed deep learning. She is a Senior Member of ACM.