Improving CAN anomaly detection with correlation-based signal clustering

Beatrix Koltai, András Gazdag, and Gergely Ács

Abstract—**Communication on the Controller Area Network (CAN) in vehicles is notably lacking in security measures, rendering it susceptible to remote attacks. These cyberattacks can potentially compromise safety-critical vehicle subsystems, and therefore endanger passengers and others around them. Identifying these intrusions could be done by monitoring the CAN traffic and detecting abnormalities in sensor measurements. To achieve this, we propose integrating time-series forecasting and signal correlation analysis to improve the detection accuracy of an onboard intrusion detection system (IDS). We predict sets of correlated signals collectively and report anomaly if their combined prediction error surpasses a predefined threshold. We show that this integrated approach enables the identification of a broader spectrum of attacks and significantly outperforms existing state-of-the-art solutions.**

Index Terms—**CAN, Anomaly Detection, TCN, Correlation.**

I. Introduction

SECURING vehicular communication networks is becoming crucial as the automotive industry rapidly evolves and increasingly adopts connectivity. Applying Intrusion Detection Systems (IDS) in specific domains is becoming essential for identifying and mitigating threats to vehicular networks. One such domain is the vehicles' inner communication on the Controller Area Network (CAN).

The CAN bus is a complex network of Electronic Control Units (ECUs) that collaborate to provide the necessary functions of the vehicle. Cyber attacks targeting these ECUs can have dire consequences for safety-critical subsystems such as brakes, the engine, or the steering wheel. A malfunctioning vehicle not only endangers passengers and others around it but also impacts the VANET (Vehicular Ad-hoc Network). Compromising data used in Vehicle-to-Everything (V2X) communication, an attacker could spread malicious information and alter the behavior of others, which could cause congestion or severe accidents in an urban environment. An attacker can have financial motivation besides deteriorating reliability and driving safety. Gaining control over the vehicle could allow theft, stealing sensitive data, and sabotaging the system.

Laboratory of Cryptography and System Security Department of Networked Systems and Services, Budapest University of Technology and Economics
E-mail: {bkoltai, agazdag, acs}@crysys.hu

Since the CAN protocol does not implement any security measures [1], an attacker can potentially attack the ECUs by making communication inaccessible, injecting new malicious messages, or even modifying valid messages. DoS (Denial-of-Service) attacks disable the benign CAN communication by flooding the network with the highest priority messages. However, this attack can be easily detected because the network load is significantly increased during the attack. Message injection can also affect specific vehicle functions, but these attacks are also easy to detect, with simple statistical methods, as injected messages cause a recognizable change in the regular arrival times.

The most challenging issue is message modification attacks that do not introduce new messages to the network, only the data contents are changed. This attack is the hardest to detect due to the variability in traffic patterns, lack of authentication or encryption, the existence of stealthy attack techniques, and the lack of attack signatures. In general, only the continuously changing message data can be used for identifying anomalies that requires general, accurate methods to differentiate between normal and malicious behavior.

After extracting signals from the message data, the detection of malicious message modifications follow two main approaches: time-series forecasting [2], [3], [4] and signal correlation analysis [5], [6]. In time-series forecasting, a machine learning model is trained per signal that predicts the next, expected signal value. Anomaly is reported when there is a substantial deviation between the prediction and the actual value. Unfortunately, this method is incapable of identifying modifications that fall within the usual, non-anomalous range of signal values, even if they constitute an attack. For instance, this limitation is evident when the speed value is modified, causing it to marginally fall below the speed limit. To overcome this shortcoming, the deviation of the correlation between each pair of signals is checked, where correlation is calculated based on the most recent few minutes' worth of signal data [5], [6]. Indeed, increasing the speed should naturally result in a corresponding increase in the RPM signal; otherwise their correlation would appear anomalous. Consequently, to evade detection, an attacker would need to maintain the original correlation intact and simultaneously modify all correlated signals, which could be prohibitively expensive in practice. Nonetheless, unlike time-series forecasting, this purely correlation-driven approach is

DOI: 10.36244/ICJ.2023.4.3

unable to identify malicious alterations in signals that lack any correlation between them.

Our proposal combines the merits of both time-series forecasting and correlation analysis. We simultaneously forecast multiple correlated signals and flag an anomaly if the cumulative difference between the predicted values and the actual values of all correlated signals exceeds a specified threshold. The underlying idea is that, as a single model forecasts multiple highly correlated signals, any alteration in one signal will inevitably influence the predictions of all other correlated signals. In other words, we leverage signal correlation not only for more accurate prediction, but also to induce detectable deviation of the predicted signals from the actual ones even if only one of them is maliciously modified. For example, the larger the speed the larger the RPM value, which means that increased speed with constant RPM is likely to produce a noticeable cumulative prediction loss over *both* signals if they are predicted jointly by a single model. Furthermore, unlike pure correlation-based approaches, our method is capable of identifying malicious alterations in signals, even those that lack correlation, when their predicted values deviate significantly from their actual values. Additionally, it can detect attacks in which the attacker modifies correlated signals simultaneously without altering their correlation, yet still induces abnormal behavior.

Our contributions in this work are as follows:

- We employ a combination of time-series forecasting and signal correlation analysis to identify anomalies in the vehicular CAN bus. Our unsupervised method relies solely on *unlabeled* CAN traces for training and calibration prior to deployment. It operates by simultaneously predicting correlated signals that allows a more accurate detection of abnormal behaviour.
- We assess the effectiveness of our approach using a dataset comprising eight distinct message modification attack types. Our results demonstrate a substantial performance improvement over the state-of-the-art: we achieve a detection rate of 95% (compared to 68%) with a precision of 80% (versus 30%). Additionally, our method exhibits a minimal average detection delay of just 0.38 seconds.
- Finally, we show that in addition to modification attacks, our solution also effectively identifies injection attacks, allowing the identification of both types of attacks by a single algorithm.

The rest of the paper is organized as follows: Section II briefly covers prior research and developments in anomaly detection in Controller Area Networks. Section III summarizes the relevant background of the CAN bus and vehicular intrusion detection solutions. The attacker model is introduced in Section IV. Section V describes the proposed anomaly detection mechanism, the training process, and the detection process. Section VI evaluates the performance of the method on real-world CAN data. Finally, in Section VII we conclude our paper.

II. RELATED WORK

Intrusion detection systems used in in-vehicle networks differ from those used on the Internet because there are limited known attack signatures. Most research results are based on unsupervised learning, as the available data can only be used appropriately to describe the benign state of the systems. Following this approach, papers have been published on detecting message injection and modification attacks.

IDS systems often rely on measuring and monitoring the timestamp of message arrivals to detect injection attacks. Due to the periodical timing of CAN data messages in a benign state, timing-based detection methods can effectively detect message insertions and drops [7], [8].

Attackers, however, cannot only inject messages into the bus, but it is also possible for them to modify messages, as described in Section IV.

In [9], the proposed method can detect these modification attacks by utilizing the transient state at the beginning of a modification attack. For a short time missing messages could indicate a suspension attack as a preparation step for a modification attack. However, if this phase is not detected in time, the rest of the attack will be successful.

In recent years, many papers have been published on identifying modification attacks based only on the message data contents. Among others, researchers tackled the problem by continuously measuring the relationship between data fields, forecasting future data values and later identifying deviations between the predictions and actual values.

CAN signal correlation analysis is proposed in [5] to identify modification attacks. Even though this approach is robust against attacks that target highly correlated signals, its effectiveness is generally limited. In [6], the authors extend correlation analysis with hierarchical clustering. Their results are demonstrated on a dataset, but it is not compared to other baseline results. As the presented framework can only handle entire traffic logs, it is not applicable as a real-time detector for the CAN bus but only as a forensics tool.

Time series forecasting is also used to predict future values in CAN communication, either on message or signal level. These predictive methods can identify possible modification attacks by measuring deviations between predicted and actual measured values.

Using a neural network for anomaly detection has been proposed in CANet [2]. Although this approach exploits relations between signals for detection, this information is not directly used in the network structure. In [3], the INDRA framework was proposed, which analyzes temporal patterns and behavior of messages using Gated Recurrent Unit (GRU) based recurrent autoencoders. The authors show that INDRA outperforms CANet in accuracy and false positive rate. In [4], the authors introduce a Temporal Convolutional Network based detection system. Their approach separates CAN signals and builds individual predictor models for each signal, similar to CANet and INDRA. However, as TCN networks are smaller and faster than previous neural networks, such as LSTMs, their solution outperforms all previous results. In this paper, we improve on the TCN-based approach by introducing signal clustering to improve detection results while reducing the mechanism's footprint.

III. Background

This section provides an overview of the CAN network's operation within vehicles, and introduces the application of Temporal Convolutional Neural Networks (TCNs) along with signal correlation analysis as part of our proposed anomaly detection approach.

A. CAN

Modern-day vehicles have a complex internal control system comprised of ECUs, each assigned to manage a specific function. These ECUs are interconnected via networks, the most important being the Controller Area Network. While this system has proven reliable over the years, external interfaces have exposed it to potential attacks [10].

On the CAN bus information is transmitted in frames. A CAN frame contains header, payload, and trailer segments. The actual data to be transmitted is in the payload segment.

Within the data part, various digital and analog signals are encoded. Manufacturers do not disclose how the signals are encoded, but they can be reverse-engineered using methods previously proposed in the literature [11].

B. Temporal Convolutional Networks

Convolutional Neural Networks (CNNs) and Temporal Convolutional Networks (TCNs) are deep learning architectures widely used for various tasks, including image recognition and natural language processing. They offer significant benefits when applied to time series data, making them suitable for detecting anomalies in the Controller Area Network (CAN) [4].

CNNs are designed to process grid-like data, such as images, by applying convolutional filters to extract spatial features. In the case of time series data, 1-dimensional causal convolutions can be used to identify local patterns and dependencies within the data.

To process sequences in parallel, TCNs use dilated convolutions, which enable them to capture long-range dependencies efficiently. This ability is critical in identifying anomalies that may occur over extended periods or exhibit complex temporal behaviors. Additionally, TCNs stack multiple layers for hierarchical feature extraction.

TCNs can handle large volumes of data, making them suitable for analyzing extensive CAN message traffic. This architecture can be optimized for real-time processing, allowing immediate anomaly detection and response in safety-critical CAN systems.

IV. Attacker model

This section discusses the attacker model and the attack surface of a CAN network. We describe the capabilities and goals of an attacker and classify the potential attacks that an attacker may perform on CAN messages.

We assume that the attacker can gain access to the vehicle using the most common attack vectors [10]. The goal of the attacker is to send forged data to an ECU, forcing it into a corrupt state. This could cause problems anywhere between showing invalid values on the dashboard to making the vehicle completely unusable or stealing it[1], depending on the target ECU. This goal can be achieved in multiple ways. For example, vehicles with wireless interfaces, such as Bluetooth, WiFi, or a 3G/4G/5G connection, can also be attacked remotely. Once an attacker has the capability to interact with the CAN bus, there are multiple possible attack strategies, including DoS, message injection, and message modification. The latter two are also referred to as a fabrication and a masquerade attack.

We focus on the most challenging problem, which is the message modification attack. During these attacks the repetition times of the messages are unchanged, as there are no new messages introduced to the network. Hence, messages arrive at their expected time but with a modified data content. Carrying out such an attack requires strong technical skills, nevertheless, its feasibility has already been demonstrated in [12]. A practical implementation of such an attack exploits the error handling mechanism of the CAN protocol. If a device detects an error during transmission, an error signal bit can be used to inform the sender about the problem. Repeated error signals can force an ECU into an error state. In this state all further message transmissions are suspended, allowing an attacker to take the place of the ECU in the communication and send modified messages. Therefore, identifying modification attacks based only on meta-data (e.g., the number or timing of CAN messages) is not possible. In this paper, we present a novel anomaly detection mechanism, designed to detect such attacks.

V. Proposed solution

Our solution has three main components: after extracting signals from the raw CAN traffic, (1) correlated signals are grouped together using clustering, (2) a separate and independent supervised forecasting model per group predicts the next value of all correlated signals within a group, and finally (3) an anomaly is reported if at least one of the forecasting model's predictions deviate significantly from the true, observed values of the predicated signals. We detail the operation of each component as follows.

A. Preprocessing of CAN Traffic

All signals from the available CAN messages are extracted using the manufacturer's specification or any state-of-the-art automatic extraction tool [11]. As not all extracted signals are equally useful for anomaly detection, a subset K of all extracted signals are retained while the rest are dropped. Indeed, useless signals are extracted from unused parts of the CAN messages (i.e., there is no device in the vehicle that uses that part of the message), or carry constant values with no

[1] https://arstechnica.com/information-technology/2023/04/crooks-are-stealing-cars-using-previously-unknown-keyless-can-injection-attacks

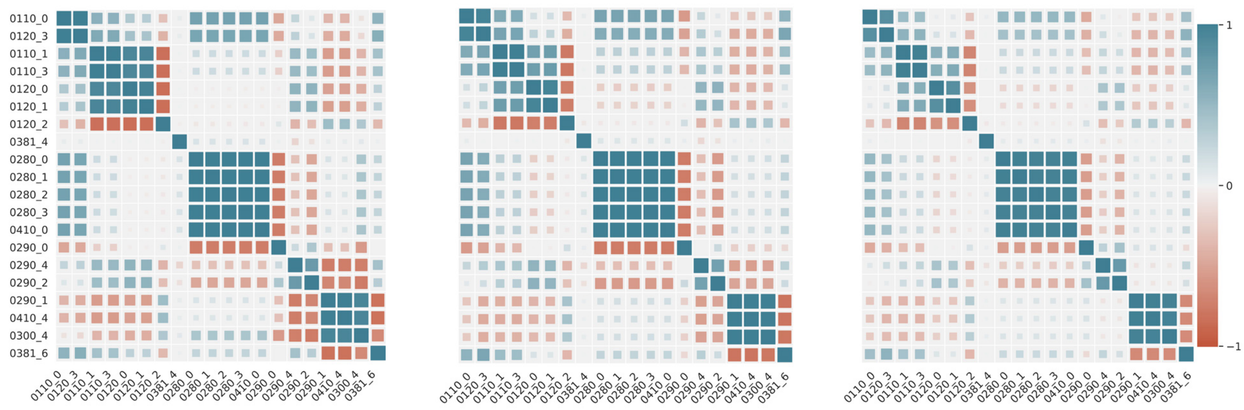

(a) Pearson correlation.

(b) Spearman correlation.

(c) Kendall's Tau correlation.

Fig. 1: Heatmaps of different correlation metrics used to determine similarity between signals (Pearson, Spearman and Kendall correlation form left to right). To better illustrate the magnitude of the correlation, we also varied the size of each point on the heatmaps, which is proportional to the darkness of the color.

predictive power. This filtering process also helps minimize the size of the forecasting model detailed in Section V-C. Finally, all retained signals are normalized by dividing each signal value by their theoretical maximum that is either specified by the manufacturer, or computed as $\lceil 2^s \rceil$ where s is the number bits used to store the signal in the CAN message.

B. Grouping of Correlated Signals

All retained K signals are clustered into C groups based on their pairwise correlation values. Although our approach is not restricted to any specific similarity measure or clustering technique, we show in this section that hierarchical clustering with Pearson correlation is the most effective combination. Specifically, each signal is first assigned to a separate cluster and then the closest clusters are iteratively merged until the number of clusters attains C. The distance of two clusters with centroids c_i, c_j is measured by $1 - |corr(c_i, c_j)|$, where $corr$ denotes the Pearson correlation.

1) Correlation analysis: We have analyzed Pearson, Spearman, and Kendall correlation metrics.

Pearson's correlation coefficient [13] measures the linear relationship between two continuous variables, suitable for typical analog signals on the CAN bus, such as speed, PRM, etc. It is sensitive to outliers, which means that extreme values can significantly influence the correlation value.

Spearman Rank Correlation [13] measures the strength and direction of the monotonic relationship between two variables by calculating the correlation based on the ranks of data points. Other than a monotonic relationship, it does not assume linearity or follow any specific distribution. The ranking property of this metric makes it robust to outliers.

Like Spearman, Kendall's Tau, also known as Kendall's rank correlation coefficient [13], does not assume linearity or follow any specific distribution. Kendall's Tau is often considered more robust than Spearman's.

A heatmap for each correlation method displaying the pairwise correlation between signals is shown in Figure 1. This visualization reveals that all three correlation methods exhibit nearly identical dependencies. However, some significant differences occur in the case of signals `0290_1`, `0410_4`, and `300_4`, which correlate only with signals `0290_4` and `0290_2` according to their Pearson coefficients. As Figure 2 shows, signal `0290_1`, `0410_4`, and `300_4` are indeed more similar, even though their Spearman and Kendall correlation values are significantly smaller.

2) Clustering of signals: We compared four distinct clustering algorithms on our dataset - DBSCAN, Affinity Propagation, Hierarchical Clustering, and Mean Shift Clustering[2]. We chose only clustering techniques that do not require the number of clusters to be specified in advance, as we do not know the optimal number of groups.

DBSCAN (Density-Based Spatial Clustering of Applications with Noise) is a density-based clustering algorithm that groups data points based on their density within the dataset [14]. It can discover clusters of arbitrary shapes and is robust to outliers called noise points.

Affinity Propagation is an exemplar-based clustering algorithm that selects a set of data points as exemplars and assigns the rest of the points to the nearest exemplar [15]. Affinity Propagation can be sensitive to the choice of similarity or distance metric, and the number of exemplars can significantly affect the results.

Hierarchical clustering builds a tree-like hierarchy of clusters, often represented as a dendrogram [14]. It can be agglomerative or divisive. Agglomerative clustering starts with individual data points as clusters. It merges them iteratively, while divisive clustering begins with a single set containing all data points and splits them into smaller groups.

Mean Shift clustering is a mode-seeking clustering algorithm that aims to find the modes or peaks of data density [16]. It is beneficial for finding clusters with non-uniform shapes or densities. Mean Shift is sensitive to the bandwidth parameter, which affects the size and shape of the groups.

[2] https://github.com/CrySyS/CAN-Message-Modification-Detection

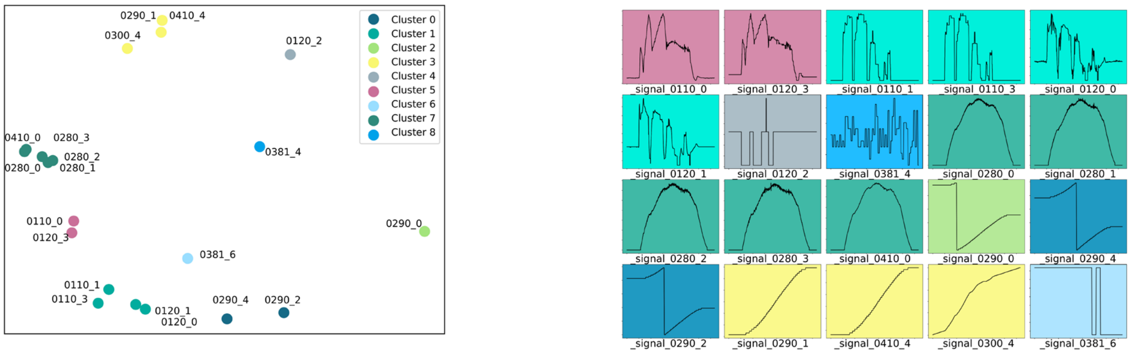

(a) Clustering illustrated in a 2D representation, each point corresponds to a signal.

(b) Result of the clustering, background color indicates the cluster.

Fig. 2: Hierarchical clustering with Pearson correlation.

We visually compared the combinations of all clustering methods with each correlation metrics. In case of DBSCAN, we noticed that some signals are assigned to separate groups, even though they apparently belong together. Moreover, the result was sensitive to the clustering parameters. We also found Affinity Propagation method too sensitive to its parameters, and even with the best settings, it grouped signals that did not belong together. MeanShift and Hierarchical clustering essentially gave the same results. We opted to use the Hierarchical clustering algorithm with Pearson correlation for ease of use. Figure 2 illustrates the result, where the signals are represented in 2D space while preserving their pairwise similarities with Multidimensional Scaling (MDS).

C. Signal Forecasting

We train C supervised models on the clustered CAN data in order to predict the next upcoming signal value: all retained K signals are divided into equally-sized overlapping segments using a sliding window with size w, and each segment serves as input to the forecasting model to predict the subsequent signal value immediately following the segment.

More precisely, let a signal with ID s be represented as a time series $(T_1^s, \ldots, T_n^s)$ after pre-processing, and $\mathbf{O}^G = [(T_1^{g_j}, T_2^{g_j}, \ldots, T_n^{g_j})] \in \mathbb{R}^{|G| \times n}$ denotes the time series of all correlated signals in group G, where $G = \{g_1, \ldots, g_{|G|}\}$ are the set of signal IDs belonging to G. For any signal group G, a forecasting model f_G simultaneously predicts the next element of each signal of the group: given the most recent w signal values $\mathbf{O}_{t-w:t}^G = [(T_{t-w}^{g_j}, T_{t-w+1}^{g_j}, \ldots, T_{t-1}^{g_j})] \in \mathbb{R}^{|G| \times w}$ as input, the forecasting model predicts the next value $\mathbf{O}_{t:t+1}^G = (T_t^{g_1}, T_t^{g_2}, \ldots, T_t^{g_{|G|}})^\top \in \mathbb{R}^{|G|}$ of every signal in G. Before deployment, all forecasting models are trained on CAN data that comes from the same or sufficiently similar distribution as the actual CAN traffic after deployment.

D. Decision

We compare the prediction made by every forecasting model with the actual, observed values of the signals, and report anomaly if the deviation of the prediction is too large for any group.

More precisely, let $\mathbf{O}_{t:t+1}^G$ denote the actual, observed value of the signals at time t in group G after performing the pre-processing steps detailed in Section V-A. The prediction error for group G at time t is defined as

$$\mathsf{err}_G(t) = \frac{1}{|G|} ||f_G(\mathbf{O}_{t-w:t}^G) - \mathbf{O}_{t:t+1}^G||_2^2 \tag{1}$$

which measures the mean squared error (MSE) between the actual signal values and the values predicated by f_G from the last w observed values of the signal. Note that $\mathbf{O}$ denotes the true value of the signal that is observed on-line after the deployment of the trained forecasting model f_G.

A naive method of detection is to directly compare the prediction error with a threshold τ, and report anomaly if $\mathsf{err}_G(t) \geq \tau$ for any group G. However, since the variance of $\mathsf{err}_G(t)$ can be large depending on the accuracy of the forecasting model f_G, this approach can yield large detection error: any value of τ would induce either too many false positives (for smaller τ) or false negatives (for larger τ). To mitigate such effect of forecasting inaccuracy, we rather compare the mean of the last ℓ error values with the threshold, that is, report anomaly if $(1/\ell)\sum_{i=t-\ell}^{t-1} \mathsf{err}_G(i) \geq \tau$ for any group G. This approach also more reliably detects stealthier attacks that span multiple time slots and involve insignificant modification of the signal value per slot, but surpass the threshold when aggregated.

To adjust τ, we follow the standard three-sigma rule and set τ to three times the standard deviation of $(1/\ell)\sum_{i=t-\ell}^{t-1} \mathsf{err}_G(i)$ plus its expected value on normal (attack-free) traffic [17]. The underlying assumption is that, without adversarial manipulation, the cumulative prediction error lies within three standard deviations of its mean that has a probability of 0.9973 if it is normally distributed (which is the case if ℓ is sufficiently large). The three-sigma rule is applicable even without access to attacked traffic before deployment, otherwise an optimal calibration of τ follows from the Neyman-Pearson lemma.

We performed experiments to determine the value of ℓ. We analyzed different values of ℓ ranging from 1 to 500, where $\ell = 1$ means that only the present error value is considered, and 500 corresponds to the mean computed over roughly 0.8 seconds of data. Too large values of ℓ can smooth out shorter attacks potentially increasing false negative rate after deployment. On the other hand, too small values of ℓ yields larger variance of $(1/\ell)\sum_{i=t-\ell}^{t-1} \mathsf{err}_G(i)$ which can increase false positive rate.

Since attacked traces are usually not available during training, the value of ℓ is adjusted to minimize false positives only on benign signals. This is what we will do to evaluate our proposal in Section VI, and set the value of ℓ to 200.

In general, the values of τ and ℓ would depend on the manufacturer's priorities. For instance, a manufacturer may prefer to minimize false positives to detect and respond to attacks quickly or to investigate all suspicious cases. However, this can lead to missing some stealthy and short-duration attacks.

E. Discussion

1) Why Grouping Correlated Signals: The joint forecasting of correlated signals offers several advantages for anomaly detection. First, it allows a single model per group to leverage the inherent interdependencies among group members, resulting in more accurate forecasts for each signal within the group. Second, any malicious modification of a signal is likely to impact the predictions of all group members, thereby increasing the cumulative prediction error as described in Eq. (1). This enhances the detectability of attacks compared to prior methods in the literature, as demonstrated in Section VI. Finally, instead of creating a stand-alone model for each individual signal as in [4], our approach requires the construction of only K forecasting models, rendering it a more appealing choice in resource-constrained environments.

2) Cost Analysis: The cost of our approach is dominated by that of the forecasting models. Apart from the C forecasting models, $K \cdot w$ signal values are stored for forecasting and $K \cdot \ell$ error values for decision purposes. The forecasting models are trained off-line in parallel, and the trained models are deployed in the vehicle. Therefore, the computational cost is dominated by the inference time of the forecasting models, where the inference processes of models are parallelizable.

VI. EVALUATION

A. Dataset

We use two CAN datasets for evaluation: Dataset-1 introduced in [4], and Dataset-2 introduced in [18].

Dataset-1 contains seven short (<1 minute) traces of specific driving and traffic scenarios, and a longer trace (~25 minutes). Dataset-2 contains nine short traces and eleven longer traces.

As the datasets originate from the same vehicle type, both have 20 message IDs and 1-6 signals per ID. Similarly, both datasets contain message injection and message modification attacks. As our main objective is to detect modification attacks, first we only use the corresponding traces.

We evaluate our mechanism on Dataset-1 to compare its performance to the chosen baseline described in Section VI-C. Since the two datasets are based on very similar CAN traffic from the same vehicle type, and most attacks follow the same strategy (only the RANDOM and DELTA attacks are not included in both), we present only the joint results.

The attacks have been performed using 6 different signal modification strategies:

- ADD-DECR - Modify with decrement value: a decrease per message is subtracted from the original value.
- ADD-INCR - Modify with increment: increases the original value by one increment per message.
- CONST - Change to constant: constant value replaces the original value.
- NEG-OFFSET - Modify with delta: a given value is subtracted from the original data value.
- POS-OFFSET - Modify with delta: a given value is added to the original data value.
- REPLAY - Replace the original data value with a previous value.
- DELTA - An attacker chosen value is added to the original value.
- RANDOM - The original value is replaced by a new random value in every attacked message.

B. Model Architecture and Parameters

For evaluation, we instantiate our proposal described in Section V. We create two datasets for training and testing purposes. A total number of 3.2 million CAN messages were used to create a training dataset for signal forecasting and calibrating all parameters of our approach (i.e., K, C, w, ℓ). Our calibrated model is tested on 1.3 million benign and malicious test messages (67 attacked traces and 9 benign traces), each containing one attacked signal. Both datasets undergo the same pre-processing steps with the same parameters that were computed exclusively on the training data.

a) Pre-processing: We use a signal mask based on the bit flip rate to extract relevant signals. We retain $K = 20$ of the $N = 77$ extracted signals that describe the state of the vehicle and likely to have sufficient predictive power for signal forecasting[3]. The retained signals are normalized as described in Section V-A.

b) Signal grouping: We conduct a correlation analysis on the signals and identify groups of correlated signals. We utilize hierarchical clustering with Pearson correlation as a similarity measure, and group linearly dependent signals together accordingly. We identify $C = 9$ clusters of the 20 signals in our dataset.

c) Signal forecasting: For forecasting, we use multi-channel Temporal Convolutional Networks (TCN). We apply an input sliding window of size $w = 1750$, equivalent to roughly 3 seconds, and each TCN has a receptive field with the same size w. Each channel of the multi-channel model corresponds to an individual signal in the group. The output of the TCN layers is then forwarded to a fully connected linear layer which generates the prediction of the upcoming signal values. Each multichannel TCN layer has four dilatation layers with a logarithmic offset of 2 $(1, 2, 4, 8)$. The kernel size is fixed at 16. We train each forecasting model with Adam optimizer and MSE loss using early stopping.

[3] Note that this information is already known to a car manufacturer

The total size of all forecasting models, capable of handling all message IDs together in groups, is approximately 15 MB and contains 4.157 million parameters.

d) Decision: We average the last $\ell = 200$ prediction error values of our forecasting models and compare with threshold τ which is calibrated according to the three-sigma rule on the training data as described in Section V-D. In other words, we do *not* use the attacked traces in our dataset to adjust τ because it is unlikely to have sufficiently representative data about all possible attacks in practice.

C. Comparison with Baselines

The most relevant related works are CANet [2], INDRA [3], and the single TCN (S-TCN) anomaly detector architecture from [4]. To avoid confusion, from now on, we will refer to the Single TCN method (S-TCN), and refer to our proposed solution described in Section VI-B as Correlation-based TCN (C-TCN).

The INDRA framework has been shown to outperform other relevant unsupervised approaches including CANet regarding false positives and detection accuracy. Moreover, according to numerical experiments on two datasets, the SynCAN dataset [2] and Dataset-1, the S-TCN approach has larger accuracy with a significantly lower false positive rate than INDRA. Therefore, it is sufficient to show that our solution outperforms the S-TCN approach, because it has demonstrated superior performance compared to CANet and INDRA [4].

To properly compare the two results, we adapt the S-TCN approach by training one TCN model per signal but keeping the rest of the process, i.e., the data pre-processing, the same as our C-TCN solution. As expected, this adapted approach can reconstruct the expected behavior of CAN signals individually.

D. Evaluation Metrics

We evaluate both the baseline S-TCN and our proposed C-TCN method using standard performance metrics: accuracy, false positive rate, precision, and recall.

Precision and recall are particularly important metrics in this context, since the *testing* dataset is often imbalanced; attacks on the CAN bus are often short, which means that the number of benign instances significantly exceeds the number of attack instances.

In addition, we also measure the time it takes to detect attacks (denoted by T_D), and the fraction of attacked traces that are successfully detected (denoted by R_D):

$$T_D = \frac{\sum_{n=1}^{N_t} (t_{detection} - t_{attack})}{N_t} \tag{2}$$

$$R_D = \frac{\sum_{n=1}^{N_t} \mathbb{1}_{\{\text{trace } n \text{ is detected as anomalous}\}}}{N_t} \tag{3}$$

where N_t is the number of attacked traces, $t_{detection}$ is the time of detection (time of the first message whose signal values trigger anomaly), t_{attack} is the starting time of the attack (time of first attacked message) and $\mathbb{1}$ is the indicator function. Note that, while recall measures the detection performance on individual messages, detection rate measures the recall with respect to the traces. Indeed, both datasets used for evaluation includes short driving scenarios affected by various types of attacks, as described in Section VI-A, and an attacked trace is successfully detected if at least one message belonging to the attacked section of the trace triggers detection.

E. Results

All experiments were done using the TCN implementation in Keras [19].

Table I shows the accuracy and false positive rate for benign and malicious test sets, as well as the precision, recall, detection rate, and detection delay for attacked traces for both message modification and message injection attacks. These metrics are calculated across multiple traces and averaged to provide the overall results shown in the table.

To investigate the use of only one IDS system in a vehicle, we also tested our solution against message injection attacks. Although we do not focus on detecting these attacks, we demonstrate that the solution can be applied to detect message injections as well.

After experimenting, we conclude that correlation-based C-TCN can effectively detect attacks on CAN bus data. Our major findings are as follows:

1) Grouping of CAN signals based on correlation improves the detection performance from 68% to 95% which means that our proposed C-TCN method can detect 95% of all the modification attack scenarios. These attacks are detected with a delay of 0.38 seconds on average.
2) Correlation-based C-TCN significantly outperforms S-TCN on all evaluated metrics, especially regarding precision and recall, where C-TCN achieves 80-83% average performance.
3) In addition to modification attacks, C-TCN also effectively identifies injection attacks, allowing the identification of both types of attacks by a single algorithm.

As Figure 3 shows, S-TCN fails to detect the stealthier ADD-DECR attack, which slowly modifies the original signal message-by-message. It is only detected when the attack abruptly stops, and the signal returns to its original value. In contrast, our C-TCN model can detect the attack earlier when the modification induces a detectable change in the cumulative prediction error.

VII. CONCLUSION

This paper presented a novel approach to intrusion detection on the CAN bus. We mainly aimed at detecting message

TABLE I
COMPARING OVERALL RESULTS OF EVALUATING THE BASELINE S-TCN AND THE PROPOSED CORRELATION-BASED C-TCN ON BENIGN AND MALICIOUS TEST TRACES FROM BOTH DATASET.

	Benign		Message Modification		Message Injection	
	S-TCN	C-TCN	S-TCN	C-TCN	S-TCN	C-TCN
Acc.	0.98	**0.99**	0.93	**0.98**	0.96	**0.99**
FPR	0.03	**0.02**	0.05	**0.04**	0.01	0.01
Prec.	-	-	0.30	**0.80**	0.67	**0.88**
Recall	-	-	0.24	**0.83**	0.28	**0.70**
R_D	-	-	0.68	**0.95**	0.79	**0.94**

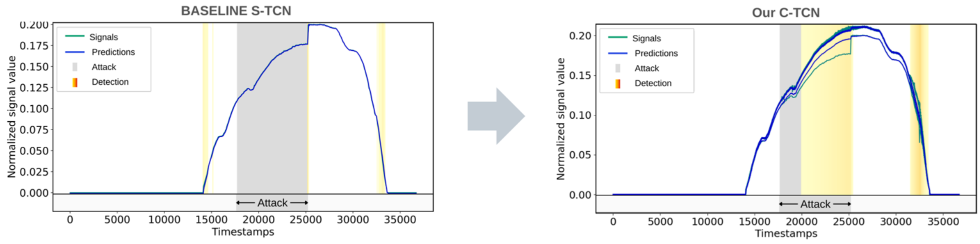

Fig. 3: Comparative evaluation of S-TCN vs. C-TCN on the same attacked trace, containing an ADD-DECR attack. The figure shows the attacked region marked by grey vertical lines and detections marked by yellow to red vertical lines, with the magnitude of the cumulative prediction error indicated by the darkness of the color.

modification attacks, the most complex attack type possible on the CAN bus. We showed that a correlation-based TCN model can efficiently predict the subsequent values of the vehicle signals, which can be used for anomaly detection. Finally, we also presented measurements demonstrating that our approach outperforms the state-of-the-art.

Our main contribution is to combine correlation analysis with time-series forecasting to improve detection accuracy. By grouping signals first based on their correlation, we create models that can predict future values with a high accuracy. During an attack, the forecasting of a group of correlated signals is significantly less accurate, allowing the detection of the anomaly. Furthermore, by grouping the signals, we can use fewer models resulting in a smaller footprint, which is an important factor for embedded systems.

In case an attacker knows which signals are clustered together and understands how the signals usually behave, it may be able to modify all the signals in the group without being detected. This requires maintaining the normal signal behavior including the inter-dependencies between different signals. However, it is unlikely that the attacker have all these capabilities in practice, especially if the groups are sufficiently large and the device running our integrated solution is adequately protected.

In our future work, we plan to analyze correlations in different traffic situations to improve our solution.

VIII. Acknowledgement

The research presented here have been supported by the NRDI Office, Ministry of Innovation and Technology, Hungary, within the framework of the Artificial Intelligence National Laboratory Programme, and the NRDI Fund based on the charter of bolster issued by the NRDI Office.

References

[1] M. Bozdal, M. Samie, S. Aslam, and I. Jennions, "Evaluation of can bus security challenges," *Sensors*, vol. 20, no. 8, 2020. [Online]. Available: https://www.mdpi.com/1424-8220/20/8/2364

[2] 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. 58 194–58 205, 2020.

[3] V. K. Kukkala, S. V. Thiruloga, and S. Pasricha, "Indra: Intrusion detection using recurrent autoencoders in automotive embedded systems," 2020.

[4] I. Chiscop, A. Gazdag, J. Bosman, and G. Biczók, "Detecting message modification attacks on the CAN bus with temporal convolutional networks," in *Proceedings of the 7th International Conference on Vehicle Technology and Intelligent Transport Systems*, 2021. [Online]. Available: DOI: 10.5220/0010445504880496

[5] A. Gazdag, G. Lupták, and L. Buttyán, "Correlation-based anomaly detection for the can bus," in *Euro-CYBERSEC*, Nice, France, 2021.

[6] P. Moriano, R. A. Bridges, and M. D. Iannacone, "Detecting can masquerade attacks with signal clustering similarity," *ArXiv*, vol. abs/2201.02665, 2022. [Online]. Available: https://api.semanticscholar.org/CorpusID:245836760

[7] H. M. Song, H. R. Kim, and H. K. Kim, "Intrusion detection system based on the analysis of time intervals of can messages for in-vehicle network," *2016 International Conference on Information Networking (ICOIN)*, pp. 63–68, 2016. [Online]. Available: https://api.semanticscholar.org/CorpusID:9333718

[8] A. Gazdag, D. Neubrandt, L. Buttyán, and Z. Szalay, "Detection of injection attacks in compressed can traffic logs," in *International Workshop on Cyber Security for Intelligent Transportation Systems, Held in Conjunction with ESORICS 2018*. Springer, 2018.

[9] S. Lee, H. J. Jo, A. Cho, D. H. Lee, and W. Choi, "Ttids: Transmission-resuming time-based intrusion detection system for controller area network (can)," *IEEE Access*, vol. 10, pp. 52 139–52 153, 2022.

[10] S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, S. Savage, K. Koscher, A. Czeskis, F. Roesner, and T. Kohno, "Comprehensive experimental analyses of automotive attack surfaces," in *20th USENIX Security Symposium (USENIX Security 11)*. San Francisco, CA: USENIX Association, aug 2011. [On-line]. Available: https://www.usenix.org/conference/usenix-security-11/comprehensive-experimental-analyses-automotive-attack-surfaces

[11] 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 Transactions on Vehicular Technology*, vol. 70, pp. 9685–9700, 10 2021.

[12] K.-T. Cho and K. G. Shin, "Fingerprinting electronic control units for vehicle intrusion detection," in *Proceedings of the 25th USENIX Security Symposium. USENIX Association*, 2016, pp. 911–927. [Online]. Available: https://www.usenix.org/system/files/conference/usenixsecurity16/sec16_paper_cho.pdf

[13] E. F. El-Hashash and R. H. A. Shiekh, "A comparison of the pearson, spearman rank and kendall tau correlation coefficients using quantitative variables," *Asian Journal of Probability and Statistics*, vol. 20, no. 3, p. 36–48, Oct. 2022. [Online]. Available: https://journalajpas.com/index.php/AJPAS/article/view/425

[14] M. Z. Rodriguez, C. H. Comin, D. Casanova, O. M. Bruno, D. R. Amancio, L. d. F. Costa, and F. A. Rodrigues, "Clustering algorithms: A comparative approach," *PLOS ONE*, vol. 14, pp. 1–34, 01 2019. [Online]. Available: DOI: 10.1371/journal.pone.0210236

[15] B. J. Frey and D. Dueck, "Clustering by passing messages between data points," *Science*, vol. 315, no. 5814, pp. 972–976, 2007. [Online]. Available: https://www.science.org/doi/abs/10.1126/science.1136800

[16] M. A. Carreira-Perpinan, "A review of mean-shift algorithms for clustering," 2015.

[17] M.-C. Dani, F.-X. Jollois, M. Nadif, and C. Freixo, "Adaptive threshold for anomaly detection using time series segmentation," in *Neural Information Processing*, S. Arik, T. Huang, W. K. Lai, and Q. Liu, Eds. Cham: Springer International Publishing, 2015, pp. 82–89.

[18] A. Gazdag, R. Ferenc, and L. Buttyán, "Crysys dataset of can traffic logs containing fabrication and masquerade attacks," *Scientific Data*, 2023.

[19] P. Remy, "Temporal convolutional networks for keras," https://github.com/philipperemy/keras-tcn, 2020.

Beatrix Koltai is an MSc student at the Budapest University of Technology and Economics (BME) in Hungary. Her general research interests include the applicability of sequential data in security, and attack detection using machine learning methods.

András Gazdag assistant professor at Budapest University of Technology and Economics (BME), in Hungary. His research interests are in embedded systems security (with a special focus on vehicle security) and embedded system forensics.

Gergely Ács is an associate professor at Budapest University of Technology and Economics (BME), in Hungary. Before that, he was a post-doc and then research engineer at INRIA, in France. His general research interests include data privacy and security, as well as machine learning in this context.