

Enhancing Cybersecurity of Automotive Ethernet Networks with Deep Learning-based Intrusion Detection Systems

Luigi F. Marques da Luz¹, Paulo Freitas de Araujo-Filho¹, Divanilson R. Campelo¹

¹Centro de Informática – Universidade Federal de Pernambuco (CIn - UFPE)
Av. Jorn. Aníbal Fernandes – s/n – Recife – PE – Brazil

{lfml, pfreitas, dcampelo}@cin.ufpe.br

Abstract. *The growing number of attack surfaces in vehicles, driven by increased connectivity and the demand for automotive Ethernet, a high-bandwidth in-vehicle network technology, emphasizes the need for effective security mechanisms. This dissertation proposes two deep learning-based intrusion detection systems (IDSs) for identifying cyberattacks in automotive Ethernet networks. The first proposal features an IDS based on a multi-criteria optimized convolutional neural network, designed to enhance detection accuracy, speed, and storage efficiency simultaneously. The second proposal introduces a multi-stage deep learning-based IDS, where the initial stage prioritizes fast detection while the second stage focuses on achieving more accurate results and classifying attacks. The main results of this dissertation comprehend the publication of a paper in SBRC 2023 and another in the Ad Hoc Networks journal.*

1. Introduction and motivation

Today's cars contain dozens of electronic control units (ECUs), which are interconnected by a number of in-vehicle networks (IVNs), [Wu et al. 2020]. Recent trends in cars, such as the distribution of automotive functions among different ECUs, higher bandwidth demands from new sensor types (e.g., cameras), and the paradigm shift to a service-oriented architecture have made legacy IVN technologies like controller area network (CAN) unsuitable for supporting the aforementioned tendencies because of its limited bandwidth and limited scalability for future applications. The emergence of Ethernet as a high-bandwidth and flexible IVN solution, especially after the standardization of the IEEE 100BASE-T1 Ethernet, has opened a myriad of opportunities for the introduction of new technologies in vehicles [Matheus and Königseder 2021].

To provide Ethernet with quality of service (QoS) capabilities, the Audio Video Bridging (AVB)/Time Sensitive Networking (TSN) task groups defined several standards that offer time synchronization, low latency, and reliability in switched Ethernet networks [Matheus and Königseder 2021]. For instance, the IEEE 1722-2016 standard defines the audio-video transport protocol (AVTP), which guarantees the reliable transmission of high bandwidth time-sensitive Ethernet traffic such as video frames from automotive infotainment systems [IEEE 2016]. Another example is the generalized precision time protocol (gPTP), which synchronizes the nodes to a common reference time to render the transmitted streams in sync [IEEE 2020]. It is worth noting, however, that Ethernet-based communications in cars must coexist with legacy in-vehicle technologies such as CAN, which is still used for safety-control applications due to its low cost and efficiency.

However, while enhanced connectivity brings new opportunities and capabilities to cars, it also presents security concerns to drivers and passengers [Liu et al. 2017, Jo and Choi 2021, Ghosal and Conti 2020]. [Checkoway et al. 2011] demonstrated the feasibility of remote exploitation of vehicles via connectivity tools, such as Bluetooth and cellular radio. Additionally, it was shown that hijacked wireless communication channels allow long-distance vehicle control and theft. As it follows, [Miller and Valasek 2015] showed the steps to remotely hack a car and turn off its engine on a highway - the vulnerabilities they found resulted in a recall of 1.4 million vehicles. Alongside, [Jeong et al. 2021] demonstrated a replay attack on an automotive Ethernet scenario. This replay attack can manipulate information and mislead the decision-making process of an autonomous vehicle, posing a significant threat to people's lives. So, defending vehicles against security threats is crucial in today's connected cars. Additionally, according to [UN Regulation 155 2021], vehicles manufactured after July 2022 in countries within the United Nations Economic Commission for Europe (UNECE) jurisdiction must be able to detect and report cyberattacks.

Traditional network security mechanisms include encryption and authentication. However, some of these mechanisms have drawbacks when considered for a resource-constrained environment such as IVNs. For example, encryption adds computing and transmission overhead that may not be suitable for IVN timing requirements [Jo and Choi 2021]. On the other hand, intrusion detection systems (IDS) are security mechanisms that work as a second line of defense, triggered when other security measures fail. IDSs monitors devices and networks to identify intrusions and report malicious activities. One of the IDS benefits is that it can be deployed as a separate network node, excluding the necessity of modifying existing nodes [Wu et al. 2020]. Although traditional IDSs can be developed using statistical methods and static rules, machine learning (ML) based IDSs have gained attention because of their detection results and capability of detecting complex attacks in high-dimensional data such as network packets [Freitas de Araujo-Filho et al. 2020]. Moreover, a modern vehicle produces tons of data that represent the behavior of the vehicle components, which could be used to develop IVN IDSs based on ML [Wu et al. 2020]. However, ML-based IDSs usually demand high computational power, often unavailable in IVNs [Bianco et al. 2018].

1.1. Problem statement

Despite the previously mentioned benefits of the use of ML/DL-based IDSs for detecting cyberattacks in automotive environments, there are still open challenges that need to be further addressed and are considered for this dissertation:

- While IDSs for IVNs must accurately detect malicious attacks, they still need a low detection time and a small storage size to be deployed in resource-constrained environments. Thus, it is necessary to propose new detection solutions that aim to optimize the detection accuracy, detection time, and storage size simultaneously, targeting IVN environments;
- While IDSs must report malicious activities quickly, it is important to have a classification of such events to provide information for forensics and future improvements [Wu et al. 2020]. Thus, it is necessary to propose new detection solutions that can accurately detect malicious activity quickly and provide a classification of it.

2. Research objectives and contributions

Our main objective was to improve the security of automotive Ethernet networks by accurately detecting and categorizing malicious activities with a low detection time using DL-based IDSs. To achieve this goal, we have defined the following specific objectives:

- Propose an IDS to detect malicious traffic in an AVTP network that uses a multi-criteria optimization technique that improves detection results, storage size, and detection time;
- Propose a novel multi-stage DL-based IDS, in which the first stage goal is to quickly detect cyberattacks, while the second stage aims to detect and classify the cyberattacks with a lower false positive rate.

3. Related work

Recently, intrusion detection systems for automotive Ethernet networks have been proposed in the literature. The authors in [Jeong et al. 2021] proposed an IDS that uses a 2D-convolutional neural network (2D-CNN) to detect replay attacks in AVTP packets. The authors also released their dataset of replay attacks for public use. The conducted replay attack consists of a specific video frame repeatedly transmitted in the network. In [Carmo et al. 2022], the authors suggested using the XGBoost algorithm to identify replay attacks in AVTP packets. Their proposal detected attacks in 620 μ s/sample using low-cost, CPU-based hardware such as Raspberry Pi and achieved similar detection rates. However, these might indicate that the replay attack presented in [Jeong et al. 2021] might be simple to detect, as the use of XGBoost had a minor impact in the detection metrics.

In [Alkhatib et al. 2022], the authors evaluated the detection time, model size, and detection-related metrics of two autoencoder-based models, a convolutional autoencoder (CAE) and a long short-term memory-based autoencoder (LSTM-AE) to develop an anomaly detector for detecting zero-day cyberattacks in AVTP packets. Despite the ability to detect new attacks, their proposed IDSs had a slight decrease in detection performance when compared to other state-of-the-art works [Jeong et al. 2021, Carmo et al. 2022]. In their other work [Alkhatib et al. 2021], the authors evaluated the performance of deep learning (DL)-based IDSs for detecting cyberattacks in scalable service-oriented middleware over IP (SOME/IP) networks and the IDS was only suitable for offline intrusion detection.

As it follows, the authors of [Han et al. 2023] presented a dataset with novel cyberattacks in a heterogeneous automotive Ethernet network containing packets from the protocols AVTP, CAN over user datagram protocol (UDP), and gPTP. The authors also proposed an IDS that used wavelet transforms to extract features and inputs them in a CNN model. Although their dataset containing different attacks, their IDS was only able to detect the attack and not classify it. Moving on, the authors of [Shibly et al. 2023] presented an IDS that uses Generate Adversarial Networks in order to decrease the need to a large amount of labeled data to train their IDS. However, the obtained results still need further improvement to match the detection of other methods, such as the one proposed in [Han et al. 2023]. At last, the authors of [Jeong et al. 2023] proposed a feature extractor that can gather information from protocol change, payload, and packet timestamps. These features are used as input to an unsupervised deep neural network model to detect cyberattacks. Despite the ability to detect novel attacks, it cannot classify the kind of attack and requires the use of GPU devices to achieve real-time detection.

Targeting the issue of balancing detection time and overall detection metrics, in Section 4 we describe the proposal of an IDS that optimizes detection accuracy, detection time, and model size during the IDS training phase. The main motivation is to achieve an IDS with a low detection time and storage size, enabling it to be deployed in resource-constrained devices. Additionally, Section 5 describes the proposal of an IDS that aims to detect cyberattacks quickly and efficiently. It uses a divide-and-conquer strategy based on using a traditional ML algorithm and a more robust DL algorithm concurrently to identify suspicious events and then accurately classify them among the known cyberattacks.

Table 1 summarizes the works mentioned above, highlighting their methods, datasets, and key characteristics. The timing requirements in the table are regarding the real-time detection threshold mentioned in [Jeong et al. 2021] of $1,735 \mu\text{s}/\text{sample}$ during a replay attack.

Table 1. Related work. AEID and TOW-IDS refer to the datasets proposed in [Jeong et al. 2021] and [Han et al. 2023], respectively. SOME/IP refers to the dataset used in [Alkhateeb et al. 2021]

Reference	Method	Dataset	Supervised	Multi-label	Timing requirements
[Jeong et al. 2021]	2D-CNN	AEID	Yes	No	Yes, with GPU devices
[Alkhateeb et al. 2021]	DL methods	SOME/IP	Yes	No	No, offline IDS
[Carmo et al. 2022]	XGBoost	AEID	Yes	No	Yes
[Alkhateeb et al. 2022]	CAE and LSTMMAE	AEID	No	No	No
[Han et al. 2023]	Wavelet transform feature extractor and customized DCNN	TOW-IDS	Yes	No	Yes, but it is not clear in which device
[Shibly et al. 2023]	Feature-aware semi-supervised learning	TOW-IDS	Partially	No	Not mentioned
[Jeong et al. 2023]	Multimodal feature extractor with a neural network	TOW-IDS	No	No	Yes, with GPU devices
Section 4 proposed system [da Luz et al. 2023]	Pruned and quantized 2D-CNN	AEID	Yes	No	Yes
Section 5 proposed system [da Luz et al. 2024]	Multi-stage IDS	AEID and TOW-IDS	Yes	Yes	Yes

4. Multi-criteria optimized deep learning-based intrusion detection system for detecting cyberattacks in automotive Ethernet networks

In this chapter, we propose a DL-based IDS that uses a multi-criteria loss function that simultaneously improves storage size, detection time, and detection results of a 2D-CNN. This work was published and presented in SBRC 2023 and was selected as one of the best papers of the event [da Luz et al. 2023]. For our threat model, we consider that the attacker has access to the in-vehicle network which is composed AVTP-capable devices. The IDS will distinguish between normal and replay attack frames, where the replay attack consist of the re-injection of pre-captured packets to misguide the network nodes that rely on this information. In this scenario, the re-injected packet may be the information source for the decision making of an autonomous vehicles, which may put people lives in danger, as depicted in Figure 1. In a nutshell, the main contributions of this work are:

- We propose an IDS to detect malicious traffic in an AVTP network that uses a multi-criteria optimization technique that improves detection results, storage size, and detection time. By doing so, we shorten the gap between deploying DL-based IDSs in resource-constrained environments such as an automotive network.
- An experimental comparison between the existing works. This comparison shows a reduction of 900x and 1.43x, respectively, in the storage size and detection time

compared to the method presented in [Jeong et al. 2021] while maintaining similar results regarding the F1-score.

Figure 1. On the left is the original frame, where the people crossing the street are detected by the ADAS. On the right is the frame received during a replay attack, where the vehicle is misguided to see no one crossing the street. Adapted from [Burke 2019].



Table 2 summarizes the F1-score for the test set, detection time, and storage size. Timing experiments were conducted using an Intel(R) Xeon(R) CPU @ 2.20GHz. The detection time and storage size in Table 2 were obtained by reproducing the works from both [Jeong et al. 2021, Carmo et al. 2022]. Our proposed IDS provides well-balanced trade-off metrics, especially regarding storage size. Despite the difference in the detection time when compared to [Carmo et al. 2022] (XGBoost follows lightweight decision-tree branches, while a CNN forwards the input through convolutional layers), we can still detect a cyberattack before a packet is received (according to the $1700 \mu\text{s}/\text{sample}$ real-time threshold specified in [Jeong et al. 2021]). Alongside, our proposed IDS achieved a higher F1-score, indicating a better trade-off in detection time and packet classification, which is significant for a safety-critical scenario such as a vehicle, as false negatives may cause accidents. In summary, we have optimized both model storage size and detection time with a minimal drop of 0.0017 points in the F1-score. This storage size result considers the amount of memory to store the model weights, and indicates that the model could be potentially stored in microcontroller devices with limited memory. Storage size metrics can be further improved with run-time memory profiling in future work.

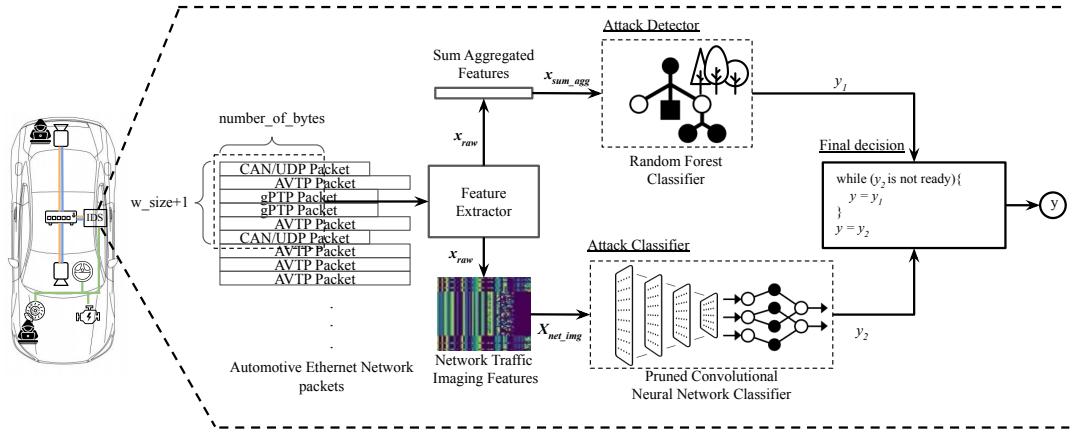
Table 2. Trade-off analysis between detection time, storage size, and F1-score.
The F1-score was obtained by considering the mean value in the test set.

Method	Detection time ($\mu\text{s}/\text{sample}$)	Storage size (KB)	F1-score
Our work	1589	11.7	0.9788
[Jeong et al. 2021]	2273	10617.0	0.9805
[Carmo et al. 2022]	250	10600.0	0.9538

5. Multi-stage deep learning-based intrusion detection system for automotive Ethernet networks

In this chapter, we build upon our previous work and propose a multi-stage deep learning-based IDS to detect several attacks in automotive Ethernet networks (depicted in Figure 2). This work was published in the Ad Hoc Networks Journal special issue on selected

Figure 2. Block diagram of our proposed IDS main components.



best papers from SBRC 2023 [da Luz et al. 2024]. The first stage has the goal of quickly detecting cyberattacks so that it is possible to stop attacks before they can cause damage. The second stage, on the other hand, is responsible for achieving a detection with a lower false positive rate when compared to the first, as well as the cyberattack classification. We aim to distinguish and classify benign and malicious in-vehicle network protocol packets in a heterogeneous automotive network. The considered in-vehicle network contains legitimate and malicious packets from AVTP, gPTP, and CAN protocols, which cover typical applications of video stream transmission, synchronization, and legacy protocol for safety-critical systems. We consider that the attacker has access to the in-vehicle network and is able to conduct attacks in the existing protocols. We extend the considered attacks to include MAC flooding, PTP sync, CAN DoS, CAN replay, and frame injection attacks. In a nutshell, the main contributions of this work are:

- We propose a novel multi-stage deep learning-based IDS, in which the first stage goal is to quickly detect cyberattacks, while the second stage aims to detect and classify the cyberattacks with a lower false positive rate;
- We evaluate our proposed IDS in publicly available automotive Ethernet datasets and compare our experimental results with state-of-the-art automotive Ethernet IDSs regarding their detection metrics and detection time.

In Table 3, we present the detection results achieved by our proposed IDS for the TOW-IDS dataset. In this scenario, our attack classifier achieved the second-best detection results among the compared works, differing only in the third decimal results from the results presented in [Han et al. 2023]. The attack classifier stage detection metrics for the TOW-IDS dataset show that our multi-stage technique provides the expected results since the attack classifier stage's primary goal is to achieve higher detection results. We noticed a performance drop in the attack detector stage when compared to the AEID dataset results. As a result, we took an in-depth analysis of the attack detector results for each attack in the TOW-IDS dataset. In Table 4, we present the number of false and true negatives per attack for our attack detector stage in the TOW-IDS dataset. One can observe that the CAN Replay attack significantly impacts the performance results, mainly due to the number of false negatives it generates. However, the attack detector's performance drop regarding the CAN Replay attack is compensated by the improved detection results of our proposed attack classifier.

Table 3. Test set results for TOW-IDS dataset. The results contain all the attacks present in the dataset. We have used the results presented in the original paper for both [Han et al. 2023] and [Shibly et al. 2023] works.

Method	Accuracy	Precision	Recall	F1-Score
Random Forest (Attack detector)	0.8740	0.9238	0.7778	0.8446
Pruned CNN (Attack classifier) [Han et al. 2023]	<u>0.9962</u>	0.9960	0.9962	<u>0.9960</u>
[Shibly et al. 2023]	0.9965	-	-	0.9974
	0.9700	0.9400	0.9500	0.9500

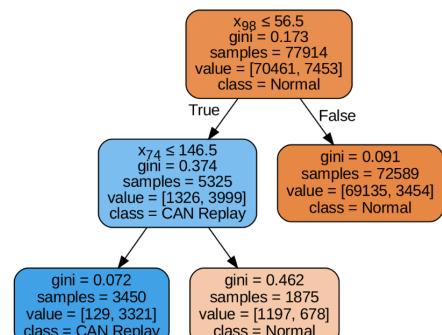
Moving on, we have conducted an explainability analysis to better comprehend the performance of our attack detector. We have used the Trustee Framework [Jacobs et al. 2022], which focuses on generating a high-fidelity and easy-to-interpret decision tree from the training data and a black box model. We have conducted this analysis only with normal and CAN replay attack samples from the training set, as it was the attack with higher false negatives. We have used the pruned tree analysis for its easier interpretation.

In Figure 3, we present one of the pruned decision trees from the Trustee framework. Based on the resulting decision tree, we can observe that the features considered for splitting the nodes are related to the CAN payload and ID fields. However, the split that uses feature 74 shows a high value of gini impurity, indicating a misclassification probability of approximately 40% in this branch. The observed high gini impurity values indicate an uncertain decision boundary between normal and CAN replay attack samples, making it hard to find the best split. This result corroborates with the information that the CAN replay attack is more complex than other attacks as they have valid CAN IDs and payload values [Jeong et al. 2023, Han et al. 2023]. Due to the higher complexity nature of the CAN replay attack and the TOW-IDS dataset heterogeneity, some factors may contribute to out-of-distribution (o.o.d) samples in the test set. These reasons include but are not limited to, a high cardinality in the values of the features coming from different CAN IDs and payloads. Another possible reason is the variable number of CAN frames in the sampling window considered for the feature extractor.

Table 4. Attack Detector Random Forest Classifier false negatives and true negatives per attack for the TOW-IDS dataset.

Attack	False Negatives (% Total)	True Negatives (% Total)
CAN DoS	6931 (9.11)	79790 (29.32)
CAN Replay	62593 (82.26)	40497 (14.88)
MAC Flooding	38 (0.05)	50708 (18.63)
PTP Sync	99 (0.13)	76137 (27.98)
Frame Injection	6433 (8.45)	25004 (9.19)
Total	76094 (100.00)	272136 (100.00)

Figure 3. Pruned decision tree obtained using Trustee framework with normal and CAN replay attack samples.



6. Conclusions and future work

In Section 4 (Chapter 5), we proposed a DL-based IDS to detect replay attacks in an automotive Ethernet network. Unlike previous works, our IDS optimizes detection results, detection time, and storage size simultaneously during the training step by applying the LilNetX framework introduced in [Girish et al. 2022]. The optimization process relies on adding two terms to the loss function: the storage size and the detection time. This optimization takes into consideration that the future deployment of the proposed IDS is a microcontroller device, which has memory limitations. We have also analyzed the trade-off between detection results, detection time, and storage size. The balance of the evaluated metrics is essential to design and deploy an IDS in a resource-constrained environment such as an IVN. We have also compared our results with other state-of-the-art intrusion detection systems for automotive Ethernet networks.

Secondly, in Section 5 (Chapter 6), we proposed a novel multi-stage deep learning approach that consists of two stages. The attack detector stage aims to ensure a fast detection time, while the attack classifier stage focuses on achieving the most accurate results. We have used a Random Forest classifier in the attack detector stage and a Pruned CNN, obtained in our previous work (presented in Section 4), in the attack classifier stage. We have evaluated our proposed IDS in two publicly available automotive Ethernet datasets: the AEID dataset [Jeong et al. 2021] and the TOW-IDS dataset [Han et al. 2023]. We have also compared our experimental results with state-of-the-art works, and our attack classifier and attack detector stages obtained the best results among the works that used the AEID dataset, with F1-Score greater than 0.995. For the TOW-IDS dataset, our detection results only differed in the third decimal digit compared to IDS proposed in [Han et al. 2023]. Furthermore, our proposed IDS presented a significant improvement in the detection time, obtaining an overall detection time result of 116 microseconds per sample for the AEID dataset and 774 microseconds per sample for the TOW-IDS dataset, being able to fulfill the real-time detection threshold of 1,000 μ s/sample proposed in [Jeong et al. 2021].

We have identified future work opportunities that could be utilized to advance the field of IDSs for automotive Ethernet networks, which are discussed below:

- Employ a feedback loop between the stages to improve the first-stage results based on the second-stage more accurate detections;
- Incorporate interpretability and explainability techniques in the design of ML and DL-based IDSs to enhance trust in their outputs, improve their robustness, and assist forensics teams in tracing cyberattacks;
- Utilizing unsupervised ML and DL methods enables the detection of zero-day attacks, which pose challenges to supervised approaches;
- Utilize online machine learning and reinforcement learning techniques to develop IDSs that are network-agnostic and capable of adapting to continuously changing environments;
- Use of in-vehicle network simulation environments to provide a lower-risk test scenario and increase the flexibility of evaluating different scenarios. For instance, the CARLA simulator can be used alongside network simulation tools such as socketCAN or OMNeT++ to have vehicle-meaningful data being transmitted over a simulated network;

- Use of automotive-grade hardware to conduct novel cyberattacks and contribute to the generation of more diverse datasets. For instance, we are conducting research at CIn-UFPE that uses TSN-capable devices to simulate real in-vehicle network traffic for cybersecurity studies;
- Implement the proposed IDSs in low-cost devices like microcontrollers and evaluate their performance in a testbed that closely resembles in-vehicle network environments.

References

Alkhatib, N., Ghauch, H., and Danger, J.-L. (2021). SOME/IP intrusion detection using deep learning-based sequential models in automotive ethernet networks. In *2021 IEEE 12th Annual Information Technology, Electronics and Mobile Communication Conference (IEMCON)*, pages 0954–0962.

Alkhatib, N., Mushtaq, M., Ghauch, H., and Danger, J.-L. (2022). Unsupervised network intrusion detection system for avtp in automotive ethernet networks. In *2022 IEEE Intelligent Vehicles Symposium (IV)*, page 1731–1738. IEEE Press.

Bianco, S., Cadene, R., Celona, L., and Napoletano, P. (2018). Benchmark analysis of representative deep neural network architectures. *IEEE Access*, 6:64270–64277.

Burke, K. (2019). How does a self-driving car see? <https://blogs.nvidia.com/blog/2019/04/15/how-does-a-self-driving-car-see/>. Accessed: 2022-12-30.

Carmo, P., Freitas de Araujo-Filho, P., Campelo, D., Freitas, E., Filho, A. O., and Sadok, D. (2022). Machine learning-based intrusion detection system for automotive ethernet: Detecting cyber-attacks with a low-cost platform. In *Anais do XL Simpósio Brasileiro de Redes de Computadores e Sistemas Distribuídos*, pages 196–209, Porto Alegre, RS, Brasil. SBC.

Checkoway, S., McCoy, D., Kantor, B., Anderson, D., Shacham, H., Savage, S., Koscher, K., Czeskis, A., Roesner, F., and Kohno, T. (2011). Comprehensive experimental analyses of automotive attack surfaces. In *20th USENIX Security Symposium (USENIX Security 11)*, San Francisco, CA. USENIX Association.

da Luz, L., Freitas de Araujo-Filho, P., and Campelo, D. (2023). Multi-criteria optimized deep learning-based intrusion detection system for detecting cyberattacks in automotive ethernet networks. In *Anais do XLI Simpósio Brasileiro de Redes de Computadores e Sistemas Distribuídos*, pages 197–210, Porto Alegre, RS, Brasil. SBC.

da Luz, L. F. M., Freitas de Araujo-Filho, P., and Campelo, D. R. (2024). Multi-stage deep learning-based intrusion detection system for automotive ethernet networks. *Ad Hoc Networks*, 162:103548.

Freitas de Araujo-Filho, P., Kaddoum, G., Campelo, D. R., Santos, A. G., Macêdo, D., and Zanchettin, C. (2020). Intrusion detection for cyber–physical systems using generative adversarial networks in fog environment. *IEEE Internet of Things Journal*, 8(8):6247–6256.

Ghosal, A. and Conti, M. (2020). Security issues and challenges in V2X : A Survey. *Computer Networks*, 169:107093.

Girish, S., Gupta, K., Singh, S., and Shrivastava, A. (2022). Lilnetx: Lightweight networks with extreme model compression and structured sparsification.

Han, M. L., Kwak, B. I., and Kim, H. K. (2023). TOW-IDS: Intrusion detection system based on three overlapped wavelets for automotive ethernet. *IEEE Transactions on Information Forensics and Security*, 18:411–422.

IEEE (2016). IEEE standard for a transport protocol for time-sensitive applications in bridged local area networks. *IEEE Std 1722-2016 (Revision of IEEE Std 1722-2011)*, pages 1–233.

IEEE (2020). IEEE standard for local and metropolitan area networks—timing and synchronization for time-sensitive applications. *IEEE Std 802.1AS-2020 (Revision of IEEE Std 802.1AS-2011)*, pages 1–421.

Jacobs, A. S., Beltiukov, R., Willinger, W., Ferreira, R. A., Gupta, A., and Granville, L. Z. (2022). Ai/ml for network security: The emperor has no clothes. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, CCS '22, page 1537–1551, New York, NY, USA. Association for Computing Machinery.

Jeong, S., Jeon, B., Chung, B., and Kim, H. K. (2021). Convolutional neural network-based intrusion detection system for AVTP streams in automotive ethernet-based networks. *Vehicular Communications*, 29:100338.

Jeong, S., Kim, H. K., Han, M. L., and Kwak, B. I. (2023). AERO: Automotive ethernet real-time observer for anomaly detection in in-vehicle networks. *IEEE Transactions on Industrial Informatics*, pages 1–12.

Jo, H. J. and Choi, W. (2021). A Survey of Attacks on Controller Area Networks and Corresponding Countermeasures. *IEEE Transactions on Intelligent Transportation Systems*, pages 1–19.

Liu, J., Zhang, S., Sun, W., and Shi, Y. (2017). In-vehicle network attacks and countermeasures: Challenges and future directions. *IEEE Network*, 31(5):50–58.

Matheus, K. and Königseder, T. (2021). *Automotive Ethernet*. Cambridge University Press.

Miller, C. and Valasek, C. (2015). Remote Exploitation of an Unaltered Passenger Vehicle. *Defcon 23*, 2015:1–91.

Shibly, K. H., Hossain, M. D., Inoue, H., Taenaka, Y., and Kadobayashi, Y. (2023). A feature-aware semi-supervised learning approach for automotive ethernet. In *2023 IEEE International Conference on Cyber Security and Resilience (CSR)*, pages 426–431.

UN Regulation 155 (2021). UN Regulation No. 155 - cyber security and cyber security management system.
<https://unece.org/transport/documents/2021/03/standards/un-regulation-no-155-cyber-security-and-cyber-security>.

Wu, W., Li, R., Xie, G., An, J., Bai, Y., Zhou, J., and Li, K. (2020). A survey of intrusion detection for in-vehicle networks. *IEEE Transactions on Intelligent Transportation Systems*, 21(3):919–933.