CISTER
Research Centre in
Real-Time & Embedded
Computing Systems

Conference Paper

Adaptive Intrusion Mitigation in Software-Defined Vehicles Using Deep Reinforcement Learning

Harrison Kurunathan*
Hazem Ali
Gowhar Javanmardi*
Mohamed Eldefrawy
Miguel Gutiérrez Gaitán*
Ramiro Robles*
Patrick Meumeu Yomsi*
Eduardo Tovar*

*CISTER Research Centre

CISTER-TR-250301

2025

Adaptive Intrusion Mitigation in Software-Defined Vehicles Using Deep Reinforcement Learning

Harrison Kurunathan*, Hazem Ali, Gowhar Javanmardi*, Mohamed Eldefrawy, Miguel Gutiérrez Gaitán*, Ramiro Robles*, Patrick Meumeu Yomsi*, Eduardo Tovar*

*CISTER Research Centre

Polytechnic Institute of Porto (ISEP P.Porto)

Rua Dr. António Bernardino de Almeida, 431

4200-072 Porto

Portugal

Tel.: +351.22.8340509, Fax: +351.22.8321159

E-mail: jhk@isep.ipp.pt, haali@isep.ipp.pt, gowha@isep.ipp.pt, mjggt@isep.ipp.pt, rsr@isep.ipp.pt, pmy@isep.ipp.pt, emt@isep.ipp.pt

https://www.cister-labs.pt

Abstract

Software-defined vehicles (SDVs) leverage vehicle-to-everything (V2X) communication to enable advanced connectivity and autonomous driving capabilities. However, this increased interconnectivity also exposes them to cyber threats such as spoofing, denial-of-service attacks, and data manipulation, making intrusion detection systems (IDS) essential for ensuring SDV security and reliability. In this work, we propose a novel intrusion mitigation approach that integrates Advantage Actor-Critic (A2C) reinforcement learning with a Long Short-Term Memory (LSTM) network to detect anomalies and intrusions in V2X communications. The LSTM component captures temporal dependencies in V2X data, enhancing the model 19s ability to identify emerging attack patterns, while the A2C framework dynamically adjusts defensive actions, including flagging, blocking or monitoring traffic, based on evolving threat levels. Experimental results demonstrate the model 19s effectiveness, achieving high detection accuracy and sensitivity. Additionally, we analyze how the system adapts over time, becoming more confident in its decision-making and optimizing security enforcement. This work enhances SDV cybersecurity by introducing a learning-based adaptive intrusion response system aiming at mitigating threats in highly dynamic vehicular networks.

Adaptive Intrusion Mitigation in Software-Defined Vehicles Using Deep Reinforcement Learning

Harrison Kurunathan+, Hazem Ismail Ali*, Gowhar Javanmardi+, Mohamed Eldefrawy*, Miguel Gutiérrez Gaitán**, Ramiro Robles+, Patrick Yomsi+, Eduardo Tovar+
+(jhk,gowha,rasro,pmy,emt)@isep.ipp.pt,*(hazem.ali,mohamed.eldefrawy)@hh.se,**miguel.gutierrez@uc.cl
+ CISTER-ISEP, Porto, Portugal, * CERES, ITE, Halmstad University, Halmstad, Sweden
** Pontificia Universidad Católica de Chile, Chile

Abstract

Software-defined vehicles (SDVs) leverage vehicle-to-everything (V2X) communication to enable advanced connectivity and autonomous driving capabilities. However, this increased interconnectivity also exposes them to cyber threats such as spoofing, denial-of-service attacks, and data manipulation, making intrusion detection systems (IDS) essential for ensuring SDV security and reliability. In this work, we propose a novel intrusion mitigation approach that integrates Advantage Actor-Critic (A2C) reinforcement learning with a Long Short-Term Memory (LSTM) network to detect anomalies and intrusions in V2X communications. The LSTM component captures temporal dependencies in V2X data, enhancing the model's ability to identify emerging attack patterns, while the A2C framework dynamically adjusts defensive actions, including flagging, blocking or monitoring traffic, based on evolving threat levels. Experimental results demonstrate the model's effectiveness, achieving high detection accuracy and sensitivity. Additionally, we analyze how the system adapts over time, becoming more confident in its decision-making and optimizing security enforcement. This work enhances SDV cybersecurity by introducing a learning-based adaptive intrusion response system aiming at mitigating threats in highly dynamic vehicular networks.

CCS Concepts

• **Computer systems organization** → **Reliability**; • **Security and privacy** → **Malware and its mitigation**.

Keywords

Software-defined vehicles, intrusion mitigation, deep reinforcement learning

ACM Reference Format:
Harrison Kurunathan+, Hazem Ismail Ali*, Gowhar Javanmardi+, Mohamed Eldefrawy*,, Miguel Gutiérrez Gaitán**, Ramiro Robles+, Patrick Yomsi+, Eduardo Tovar+. 2025. Adaptive Intrusion Mitigation in Software-Defined Vehicles Using Deep Reinforcement Learning. In *The 4th International Workshop on Real-time and IntelliGent Edge computing (RAGE '25), May 6–9, 2025, Irvine, CA, USA*. ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3722567.3727848

ACM ISBN 979-8-4007-1611-9/2025/05
https://doi.org/10.1145/3722567.3727848

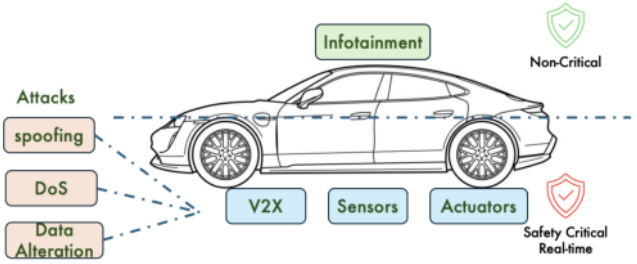

Figure 1: Cyberattacks on the V2X Architecture in SDVs.

1 Introduction

Over the past decade, Vehicle-to-Everything (V2X) communication has become a cornerstone in the evolution of Software-Defined Vehicles (SDVs), enabling applications such as autonomous driving, traffic management [1], and connected platooning [2] [3]. SDVs focus on three critical aspects of modern vehicles: connectivity, automation, and safety [4]. By leveraging software-driven reconfiguration, SDVs enable scalable and adaptable automotive architectures that continuously update and optimize vehicle operations [9]. However, V2X connectivity also introduces significant security challenges, particularly in control, data exchange, and decision-making processes, which are crucial for safe vehicle operation. As illustrated in Figure 1, the SDV ecosystem consists of safety-critical components (e.g., sensors and actuators responsible for vehicle control) and non-critical components (e.g., infotainment systems). The interconnected nature of these subsystems makes them susceptible to cyber threats, including spoofing, denial-of-service (DoS) attacks, and data manipulation, which can target not only V2X communications but also onboard vehicular networks. Mitigating these risks requires robust intrusion detection mechanisms that can safeguard SDVs without compromising system performance.

Traditional Intrusion Detection Systems (IDS) rely on predefined rules and signatures to identify malicious activities, such as abrupt deviations from normal patterns or unexpected timestamp variations. While effective against known threats, these rule-based approaches struggle to recognize novel and sophisticated cyberattacks, particularly zero-day attacks [6]. Given the highly dynamic nature of SDV networks, more adaptive and intelligent solutions are required to detect anomalies within short response times. To address these limitations, researchers have increasingly turned to machine learning-based IDS [8], leveraging reinforcement learning and deep learning techniques to dynamically learn attack patterns and analyze temporal variations in V2X traffic. These approaches

improve detection accuracy and responsiveness, enhancing the security of cyber-physical systems against evolving threats.

To address these challenges, we propose a combinatorial (or hybrid) learning framework that combines Advantage Actor-Critic (A2C) reinforcement learning and Long Short-Term Memory (LSTM) networks to detect intrusions in SDV communications and dynamically respond with mitigation actions. Given the need for sequential pattern recognition and adaptive decision-making, A2C-LSTM [7] provides a robust framework for processing dynamic data streams in SDV networks. The proposed approach analyzes several packet-level features, including packet size, arrival rate, source and destination IDs, signal strength, channel utilization, timestamps, error rates, sequence numbers, and protocol types, to detect anomalous patterns that may indicate security threats.

Irregular packet sizes or erratic arrival rates may indicate flooding attacks, while discrepancies in sequence numbers can be a sign of replay attacks [9]. By analyzing temporal patterns in vehicular data traffic, the LSTM layer captures sequential dependencies, allowing the A2C-LSTM model to adaptively adjust its actions for intrusion detection and mitigation. In this preliminary work, we define four actions based on the model's confidence in detected anomalies: (i) allowing communication, (ii) flagging suspicious packets, (iii) alerting and monitoring, and (iv) blocking. These actions reflect different levels of intervention criticality. Minor irregularities, such as slight deviations in packet size, lead to flagging without interrupting traffic. Major anomalies, such as inconsistencies in sequence numbers, high error rates, or sudden signal strength drops, trigger immediate blocking to prevent further threats. This adaptive decision-making enables SDVs to handle a variety of attack types under fluctuating conditions. By dynamically adjusting responses, the A2C-LSTM model enhances both the reliability and resilience of SDV security operations.

The contributions of this work are as follows:

- We introduce a hybrid approach combining A2C and LSTM for intrusion detection in SDV communications by leveraging temporal dependencies.
- We propose a reinforcement learning framework that enables the model to learn dynamic attack scenarios and adaptively execute mitigation actions, improving resilience against a range of threats without manual intervention.
- We analyze the model's detection accuracy, error rates, and sensitivity, providing insights into its effectiveness in identifying cyber threats.

The rest of the paper is organized as follows: Section II presents the state of the art on machine learning-aided intrusion detection systems. Section III describes the system model, detailing the proposed algorithm and its preliminaries. Section IV discusses the experimental results, followed by conclusions and future research directions.

2 State of Art

The security in autonomous vehicles (AVs) is a major concern for research due to their reliance on interconnected systems, AI-driven decision-making, and communication networks. One of the recent surveys that addresses that concern in terms of communication is [14] where the authors examine the security challenges associated with the Internet of Autonomous Vehicles (IoAV). They identify various vulnerabilities in IoAV systems and emphasize the necessity of adopting lightweight security protocols. These protocols are designed to provide strong protection while maintaining system performance. Additionally, the authors explore the implementation of fast, lightweight security protocols on heterogeneous embedded, low-power, high-performance computing platforms as an effective approach to addressing these security challenges. Other surveys show the need for Machine learning to enable IDSs to further classify between benign and malicious behaviours, and its several challenges, such as lack of quality training data and high false-positive rates, have been discussed through various surveys in the literature [10, 11].

Table 1: Summary of Related IDS

Ref	Method
[10]	Explores the use of Generative Machine Learning Models (GMLMs) in intrusion detection, presenting a systematic study of existing research and a detailed review with insights for future advancements.
[11]	Examines intrusion detection technologies, comparing deep learning models, datasets, feature extraction, classifiers, and preprocessing techniques. It discusses challenges, experimental design, and key issues, concluding with recommendations for researchers.
[12]	Proposes a load balancing algorithm that maps sensor data, vehicles, and data centers performing tasks.
[13]	Proposes an intrusion detection approach combining Federated Learning (FL) with the BERT model, referred to as FL-BERT to mitigate data manipulation.
[15]	Introduces DRL-KeyAgree, a channel-based key generation technique using combinatorial deep reinforcement learning (DRL) to mitigate spoofing attacks.
[16]	Extreme gradient boost and K-near neighbour algorithms to classify intrusion in CAN systems of vehicles
[18]	context-sensitive IDS that detects vehicle intrusions by analyzing sensor data, including AI and human drivers, using hyperparameter-optimized detection.
[19]	Anomaly detection method for zero-day attacks using state vector machines by distinguishing high-dimensional abnormal patterns.
[20]	Introduces an anomaly detection method that serves as a NIDS for V2X cyber-attacks, utilizing IP and payload whitelisting, frequency analysis, and payload inspection.

Example of these IDSs is [20], the authors proposed an anomaly detection method for V2X systems using a Network Intrusion Detection System (NIDS). While it detects abnormal behaviour, it may miss differentiation between attacks or operational issues, leading to possible drops of legitimate packets. To improve accuracy, complementary solutions should analyze NIDS alerts and provide context, reducing false positives.

Another IDS [21] where the authors propose an anomaly detection method for securing Connected and Automated Vehicles

(CAVs) in V2X networks using 5G. By learning normal traffic patterns with deep auto-encoders, their approach detects deviations that may signal cyberattacks or faults, including zero-day attacks, without requiring labeled data. This method provides a scalable and privacy-preserving solution for V2X environments.

Also, novel methods like Federated Learning (FL) capabilities in conjunction with the BERT model for sequence classification (FL-BERT) have been used for intrusion detection [13]. In this work, the authors train their models locally and then send the model's weights to the server. After aggregation, the server aggregates the weights from all clients to update a global model. After aggregation, the global model's weights are shared with the clients. In [15], the combination of A2C and LSTM was used to capture the temporal characteristics of V2V communication and generate secret keys. In this work, we capture more temporal characteristics and enable A2C to define an appropriate action based on the models' ability to define the intrusions in the system.

Machine learning classifiers have been used in the literature to identify cyber assaults in CAN systems of vehicular networks. The authors of [16] use extreme gradient boost and K-nearest neighbour algorithms to classify such attacks. In this work, a modified metaheuristic optimizer is used to classify with higher accuracy. The pivotal role of AI and ML in threat prediction and mitigation in line with vehicular networks has been discussed in [17]. These methods in training data generation, model optimization, and real-time decision-making. In our proposed work, we utilize deep reinforcement learning to mitigate imminent threats to the integrity of the vehicular network.

Researchers in [18] introduce CAHOOTv2, a context-sensitive intrusion detection system to protect the vehicular network against intrusions such as denial of service, replay, spoofing, additive and selective attacks. The authors use a modified version of the MetaDrive simulator to generate data of the intruder that contradicts with the driver data. In our work, we use a synthetic dataset that varies based on attack patterns such as denial of service (DoS).

The authors in [19] propose an anomaly detection method, "ZeroCAN," to address the security threats posed by zero-day attacks in vehicular networks. Zero-day attacks exploit undiscovered vulnerabilities at both the software and hardware levels, making their detection crucial for ensuring system security. To tackle this issue, the authors leverage Support Vector Machines (SVM) for anomaly detection, as SVMs effectively handle high-dimensional data and distinguish between normal and abnormal patterns in complex environments. Their approach involves modelling the behaviour of each electronic control unit (ECU) on the network using separate SVMs, along with a set of high-level features that capture the timing and data payload characteristics of CANbus traffic. The proposed method achieves an anomaly detection rate exceeding 99% and maintains a false positive rate below 0.01% in most normal operation scenarios. In this work, we use A2C to mitigate such attacks so that the integrity of communication will not be compromised.

Packet-level intrusion detection model has been proposed for SDVs in [12]. The authors in this work initially propose a load-balancing algorithm to map sensor data, vehicles, and data centers performing tasks. This load balancing can track the attack on the network by learning network configuration, routing, network characteristics, and system load factors. In our work, we also consider several packet-level data to enrich the environment to make adaptive decisions.

3 System Model

We present the system model of our combinatorial mitigation system in Figure 2. In this model, the actor learns the policy by selecting actions, while the critic evaluates these actions by estimating the value of the states. The feedback from the critic through the advantage function helps the actor refine its policy, which results in better action and eventually better rewards. LSTM layers model the temporal nature of V2X data, as intrusions that are influenced by any prior communication patterns. The LSTM cells process sequences of states $\{\boldsymbol{s}_{t-k}, \ldots, \boldsymbol{s}_t\}$ and capture dependencies over time k and generate hidden representations that inform the policy and value functions. The state space represents inputs to the LSTM model at each time step. The state S_t at time t consists of packet-level features relevant to network communication in the V2X environment, capturing essential data for intrusion detection. The state is defined as follows:

$$S_t = \{\mathrm{PS}_t, \mathrm{AR}_t, \mathrm{S\_ID}_t, \mathrm{D\_ID}_t, \mathrm{RSS}_t, \mathrm{TS}_t, \mathrm{ER}_t, \mathrm{SN}_t\} \tag{1}$$

where, PS_t is the size of the packet at time t, AR_t is the arrival rate of the packets, $\mathrm{S\_ID}_t$ and $\mathrm{D\_ID}_t$ are the source and the destination IDs, RSS_t is the signal strength, TS_t is the timestamp, and, SN_t is the sequence number of the packets.

The action A_t at time t is the decision made by the agent based on the observed state S_t. Define A_t as:

$$A_t \in \{a_{\text{allow}}, a_{\text{flag}}, a_{\text{monitor}}, a_{\text{block}}\} \tag{2}$$

where the actions correspond to *allowing* the communication to proceed, *flagging* suspicious communication, *monitoring* the communication and setting it up for further inspection, and *blocking* the communication due to high suspicion. The Actor network, parameterized by θ, represents the policy $\pi(A_t|S_t;\theta)$ and outputs a probability distribution over actions given the current state. The policy function with LSTM output is:

$$\pi(A_t|S_t;\theta) = \mathrm{softmax}(f_{\mathrm{LSTM}}(S_t;\theta)) \tag{3}$$

where $f_{\mathrm{LSTM}}(S_t;\theta)$ denotes the output of the LSTM network with parameters θ. The Critic network, parameterized by ϕ, estimates the value function $V(S_t;\phi)$, which provides a measure of the expected future reward from state S_t under the current policy:

$$V(S_t;\phi) = f_{\mathrm{LSTM}}(S_t;\phi) \tag{4}$$

The critic networks Value Estimate of the expected reward assesses the quality of the actions in terms of improving detection accuracy. This is done by learning the temporal difference (through LSTM), and defining which actions will result in favorable outcomes for intrusion detection. The reward R_t is designed to incentivize actions that accurately identify and respond to intrusions, based on detection accuracy, false positives/negatives, and resource usage:

$$R_t = \alpha \cdot \mathrm{TPR} - \beta \cdot \mathrm{FPR} - \gamma \cdot \text{Resource Usage} \tag{5}$$

where α, β, and γ are weighting parameters and, TPR and FPR denote the true positive and the false positive rates. The resource

Harrison Kurunathan+, Hazem Ismail Ali*, Gowhar Javanmardi+, Mohamed Eldefrawy*,
Miguel Gutiérrez Gaitán**, Ramiro Robles+, Patrick Yomsi+, Eduardo Tovar+

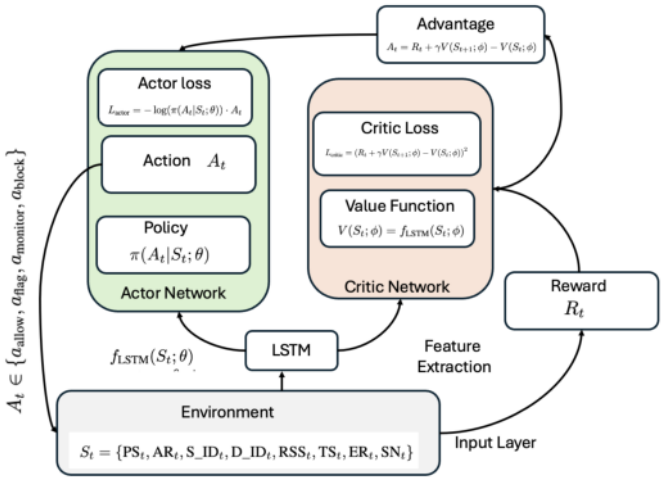

Figure 2: Combinatorial model for intrusion mitigation system where the environment is enriched through LSTM and actions such as allow, block, monitor, and flag are taken based on detection accuracy and positive rates

usage penalizes resource-intensive actions such as the monitor and block, so that they are utilized more carefully, thus preserving bandwidth and resources for essential functions. Using Advantage Actor-Critic (A2C), the network is updated through temporal difference learning. The advantage A_D is computed as:

$$A_D = R_t + \gamma V(S_{t+1};\phi) - V(S_t;\phi) \tag{6}$$

The actor and critic are updated by minimizing the actor loss L_{actor} and critic loss L_{critic} as follows:

$$L_{\text{actor}} = -\log(\pi(A_t|S_t;\theta)) \cdot A_t \tag{7}$$

Based on the probability of selecting the correct action given the state, L_{actor} optimizes the actor model to improve its decision-making in choosing appropriate actions.

$$L_{\text{critic}} = (R_t + \gamma V(S_{t+1};\phi) - V(S_t;\phi))^2 \tag{8}$$

L_{critic} presents the difference between predicted value and observed reward. This loss function helps in improving the critic's value estimation. Finally, a feedback mechanism adjusts the A2C policy in an iterative fashion and adapts the value estimates in line with the real-time interactions from the environment. This reinforces the actions that enhance intrusion detection accuracy.

4 Results

The experimental setup involves simulating a sequence of packet-level feature states over 500 episodes. Every episode taken in this dataset contains between 100 to 1000 packets of data representing packet-level attributes like packet size, arrival rate, source and destination IDs, and signal strength. These features are used for training the actor-critic network, which decides a set of four potential actions representing a response in the form of a mitigation strategy. In order to mimic a real-life vehicular networking scenario, the packet size follows an exponential distribution, arrival rates follow a Poisson process, signal strength varies according to a normal distribution, and channel utilization follows a CSMA/CA back-off model. The error rates are influenced by interference, signal strength fluctuations, and network congestion. The dataset comprises two types of normal and attack traffic, with 30% of the total network traffic containing attack scenarios. Normal traffic exhibits standard TCP/IP and UDP patterns, while attacks include DoS attacks, which lead to a spike in arrival rates and channel utilization. The reinforcement learning model learns an optimal intrusion mitigation strategy by analyzing the feature evolution.

4.1 Adapatability to detect intrusion

To understand how well the system adapts to intrusion detection over time, we first measure the model sensitivity to this application. Model sensitivity measures the ability to accurately identify any irregularities within the V2X communication system. It represents the *proportion of actual attacks that the model correctly classified as intrusions*. When the sensitivity is high, it denotes that the model has high reliability in defining and mitigating the threat. From Figure 3, it can be noticed that the model sensitivity rises from 80% to almost 100% around 500 episodes. This denotes the capability of the model to accurately identify the attacks.

Furthermore, to add validation to the sensitivity results, we plot the detection accuracy and its range along the episodes. Detection accuracy defines the effectiveness in correctly identifying both attacks and secure states. It is the *proportion of all correct predictions made by the model*, combining true positives (correctly detected attacks) and true negatives (correctly identified secure states) relative to the total number of cases. From Figure 4, it can be noticed the detection accuracy of the intrusion mitigation system increases steadily with learning and reaches almost 95-100 % around 500 episodes. The improvement in both the model sensitivity and the detection accuracy can be attributed to the LSTM capturing the temporal data across the episodes and enriching the environment with hidden information. This improves the models' capability to improve the accuracy and the eventual reliability.

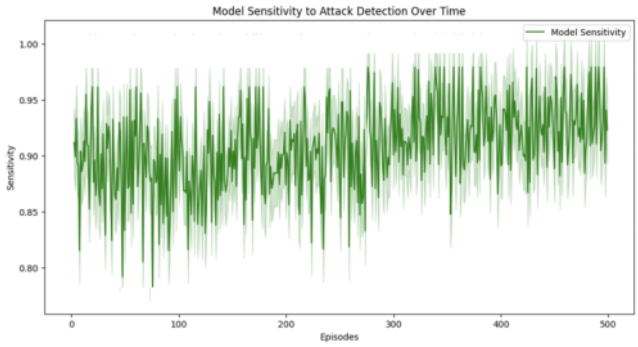

Figure 3: Model sensitivity defines the proportion of actual attacks that the model correctly classified as intrusions.

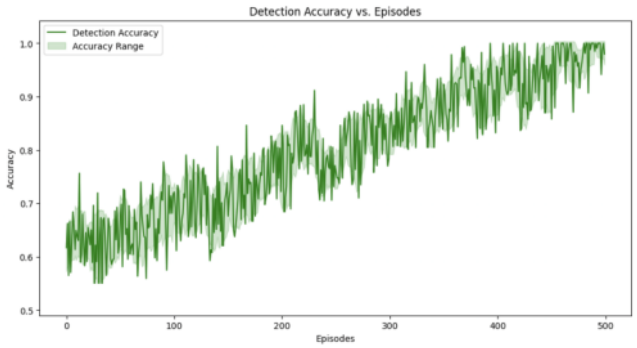

Figure 4: Detection accuracy increases steadily with learning and reaches almost 95-100 % around 500 episodes

Figure 5: FPR is steadily below 0.20, and FNR under 0.1indicating that the model is accurate

Figure 6: Heatmap of allow, flag, block, and monitor communications defined by the A2C network across 100 episodes

4.2 Intrusion mitigation through learning

Having defined the accuracy of the model, we then plot the false positive and negative rates that show the negative aspects of the model. From our results in Figure 5, we witness a slightly higher False Positive Rate (FPR) compared to the False Negative Rate (FNR). This behaviour suggests that the system is more sensitive to potential threats, resulting in occasional "false alarms" where benign communications are flagged as suspicious. However, it must be noticed that FPR is steadily below 0.20, and FNR under 0.10. These lower values indicate that the model is reasonably accurate in classifying malicious activity. The higher FPR also implies a conservative approach that prioritizes the detection of potential threats, even at the cost of some benign packets being wrongly flagged as malicious. This design can be a strong contender for safety-security-critical applications of SDV communication where missing a true threat (a false negative) could be very catastrophic.

Finally, in Figure 6, we present the heatmap of the four actions: allow, flag, block, and monitor communications defined by the A2C network across 100 episodes. Early episodes show more diversity in actions; this is due to the exploratory behaviour of the model. The model tries various responses to understand their outcomes. As the episodes progress, we can see the actions becoming more consistent. This implies that the model is shifting toward exploitation, and the model relies on actions that provide better rewards. We also can identify that this model is more conservative initially as it prefers defensive responses such as blocking or flagging due to positive detections. However, it is also equally adaptive by providing long stretches of allow and flag as the model becomes confident in detecting non-threatening states.

4.3 Comparison to state-of-the-art techniques

To compare our method with benchmarks, we test it against DQN and the novel dual-DQN methods. DQN is a DRL that approximates the Q-value function, which in turn estimates the expected future rewards for every action in a given state. Dual-DQN reduces the overestimation bias that occurs in the action-value estimates. This improves the stability of this algorithm. Akin to the Q-value, our method uses LSTM to estimate and learn from the temporal dependencies and improve the reward at every turn of the learning process. As these algorithms are much more suitable for complex and sequential decision processes, we compare them against each other in this experiment. The integration of LSTM helps in handling temporal and sequential dependencies much more efficiently than the DQN and Dual-DQN. DQN relies on feed-forward networks and results in Q-value overestimation, whereas the A2C-LSTM, with its recurrent structure, retains the past experiences that result in convergence much earlier than its counterparts. As noticeable in

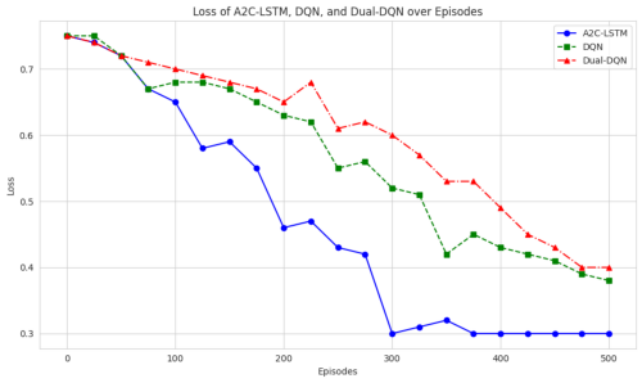

Figure 7: Comparison of loss against DQN and Dual-DQN, where A2C-LSTM gets early convergence

Figure 7, the loss drops around 300 episodes and reaches convergence, whereas DQN and Dual DQN have a constant difference of about 20-50% against the proposed A2C-LSTM-based intrusion detection mechanism.

5 Future Scope

In this work, we propose a novel intrusion mitigation approach that leverages the functionalities of the Advantage Actor-Critic (A2C) framework with a Long Short-Term Memory (LSTM) network to detect anomalies and intrusions in V2X communications. From our preliminary analyses, we were able to confirm that the proposed model provides a detection accuracy of almost 95-100% and the actions taken are well adapted based on the sensitivity of the model. As a future scope of this work, we aim to explore other reinforcement learning techniques like DDPG and other methods like CNN that can capture temporal variances similar to the LSTM. Another interesting take would be to integrate federated learning approaches that enable distributed learning across SDVs while preserving data privacy. Adapting these countermeasures to the SDV architecture can enable countermeasures such as dynamic adaptive vehicle behaviour in response to detected threats, thus improving the overall security and resilience of the system.

Acknowledgments

This work was supported by the CISTER Research Unit (UIDP/UIDB/04234/2020), financed by National Funds through FCT/MCTES (Portuguese Foundation for Science and Technology); by HAL4SDV funding within the Chips Joint Undertaking (Chips JU) - the Public-Private Partnership for research, development and innovation under Horizon Europe – and National Authorities under grant agreement n° 101139789. In Addition, it is funded by Halmstad University (Halmstad i Högskolan) research base funding. Also by the Chilean Research Agency ANID through FONDECYT project No. 11241221.

References

[1] Reddy, R., Almeida, L., Kurunathan, H., Gaitán, M.G., Santos, P.M. and Tovar, E., 2024. Worst-case response time of mixed vehicles at complex intersections. in *IEEE Open Journal of Intelligent Transportation Systems*, 5, pp.186-201.
[2] H. Kurunathan, J. Santos, D. Moreira, and P. M. Santos, "Towards Safe Cooperative Autonomous Platoon Systems Using COTS Equipment," in *Proc. 2023 IEEE 24th Int. Symp. on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, pp. 464–469, 2023.
[3] P. Duarte, M. Soyturk, R. Robles, M. Araújo, B. Yaman, A. Goes, B. Mendes, G. Javanmardi, and M. G. Gaitán, "5G Network as Key-Enabler for Vehicular Platooning," in *Proc. 2023 20th Annual IEEE Int. Conf. on Sensing, Communication, and Networking (SECON)*, pp. 369–371, 2023.
[4] P. Veloso Teixeira, D. Raposo, R. Lopes, and S. Sargento, "Software Defined Vehicles for Development of Deterministic Services," *arXiv e-prints*, p. arXiv-2407, 2024.
[5] K. Yadav, S. Shinde, V. Chebrolu, A. Kumar, V. Ramakrishna, and P. Rajalakshmi, "Design Considerations and Framework Analysis for Software-Defined Autonomous Vehicles," in *Proc. 2024 IEEE 99th Vehicular Technology Conference (VTC2024-Spring)*, pp. 1–5, 2024.
[6] R. Ahmad, I. Alsmadi, W. Alhamdani, and L. Tawalbeh, "Zero-day Attack Detection: A Systematic Literature Review," *Artificial Intelligence Review*, vol. 56, no. 10, pp. 10733–10811, 2023.
[7] H. Kurunathan, K. Li, W. Ni, N. Li, E. Tovar, and M. Guizani, "Exploring LSTM-Assisted A2C for Physical Layer Security in Vehicular Cyber-Physical Systems," in *Proc. 2024 International Wireless Communications and Mobile Computing (IWCMC)*, pp. 1177–1182, 2024.
[8] F. Sangoleye, J. Johnson, and E. E. Tsiropoulou, "Intrusion Detection in Industrial Control Systems Based on Deep Reinforcement Learning," *IEEE Access*, 2024.
[9] K. H. M. Gularte, J. A. R. Vargas, J. P. J. Da Costa, A. A. S. Da Silva, G. A. Santos, Y. Wang, C. A. Müller, C. Lipps, R. T. S. Júnior, W. B. Vidal Filho, *et al.*, "Safeguarding the V2X Pathways: Exploring the Cybersecurity Landscape through Systematic Literature Review," *IEEE Access*, 2024.
[10] J. Halvorsen, C. Izurieta, H. Cai, and A. Gebremedhin, "Applying Generative Machine Learning to Intrusion Detection: A Systematic Mapping Study and Review," *ACM Computing Surveys*, vol. 56, no. 10, pp. 1–33, 2024.
[11] H. Liao, M. Z. Murah, M. K. Hasan, A. H. M. Aman, J. Fang, X. Hu, and A. U. R. Khan, "A Survey of Deep Learning Technologies for Intrusion Detection in Internet of Things," *IEEE Access*, 2024.
[12] U. Ahmed, J. C.-W. Lin, G. Srivastava, U. Yun, and A. K. Singh, "Deep Active Learning Intrusion Detection and Load Balancing in Software-Defined Vehicular Networks," *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, no. 1, pp. 953–961, 2022.
[13] S. I. Ahsan, P. Legg, and S. M. Alam, "Privacy-Preserving Intrusion Detection in Software-Defined VANET Using Federated Learning with BERT," *arXiv preprint arXiv:2401.07343*, 2024.
[14] H. I. Ali, H. Kurunathan, M. H. Eldefrawy, F. Gruian and M. Jonsson, "Navigating the Challenges and Opportunities of Securing Internet of Autonomous Vehicles With Lightweight Authentication," in *IEEE Access*, vol. 13, pp. 24207-24222, 2025, doi: 10.1109/ACCESS.2025.3537800.
[15] H. Kurunathan, K. Li, E. Tovar, A. M. Jorge, W. Ni, and A. Jamalipour, "DRL-KeyAgree: An Intelligent Combinatorial Deep Reinforcement Learning-Based Vehicular Platooning Secret Key Generation," *IEEE Transactions on Intelligent Transportation Systems*, 2024.
[16] P. Dakic, M. Zivkovic, L. Jovanovic, N. Bacanin, M. Antonijevic, J. Kaljevic, and V. Simic, "Intrusion Detection Using Metaheuristic Optimization Within IoT/IIoT Systems and Software of Autonomous Vehicles," *Scientific Reports*, vol. 14, no. 1, p. 22884, 2024.
[17] B. J. Asaju, "Advancements in Intrusion Detection Systems for V2X: Leveraging AI and ML for Real-Time Cyber Threat Mitigation," *Journal of Computational Intelligence and Robotics*, vol. 4, no. 1, pp. 33–50, 2024.
[18] D. Micale, I. Matteucci, F. Fenzl, R. Rieke, and G. Patanè, "A Context-Aware On-Board Intrusion Detection System for Smart Vehicles," *International Journal of Information Security*, vol. 23, no. 3, pp. 2203–2223, 2024.
[19] W. Balte, J. Rendal, H. Kurunathan, H. I. Ali, A. dos Santos Roque, W. de Morais and M. Fazeli, "ZeroCAN: Anomaly-Based Zero-Day Attack Detection in Vehicular CAN Bus Networks," 2025 33rd Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP), Turin, Italy, 2025.
[20] A. Lekidis, "Anomaly Detection Mechanisms for In-Vehicle and V2X Systems," in *Proc. 19th Int. Conf. on Availability, Reliability and Security*, pp. 1–8, 2024.
[21] A. A. Korba, A. Boualouache, B. Brik, R. Rahal, Y. Ghamri-Doudane, and S. M. Senouci, "Federated Learning for Zero-Day Attack Detection in 5G and Beyond V2X Networks," in *Proc. ICC 2023 - IEEE Int. Conf. on Communications*, pp. 1137–1142, 2023.

Received 10 February 2025; accepted X March 2025