A RISK-AWARE DEEP REINFORCEMENT LEARNING FRAMEWORK FOR AI-DRIVEN INTRUSION DETECTION AND ADAPTIVE RESPONSE IN AUTONOMOUS VEHICLES

Muhammad Faisal Shafiq[1], Muhammad Irshad[2] and Muhammad Naveed Sajjad[3]

[1]Lucid Motors, Riyadh, Saudi Arabia
[2]RMG Company, Riyadh, Saudi Arabia
[3]Yanal Finance Company, Riyadh, Saudi Arabia

ABSTRACT

Autonomous vehicles expose in-vehicle networks to sophisticated intrusion threats. Although deep learning-based intrusion detection systems achieve high accuracy, most approaches remain detection-centric and do not address adaptive mitigation under real-time constraints. This paper proposes a multi-layer AI-driven framework integrating temporal deep learning, contextual risk modeling, and deep reinforcement learning-based mitigation. A hybrid CNN–BiLSTM model captures spatial payload characteristics and bidirectional temporal dependencies in CAN sequences. Detection output feeds a risk-aware formulation fusing attack probability with ECU criticality, vehicle speed, and safetyindicators. A Deep Q-Network learns mitigation policies minimizing residual system damages. Evaluation on the HCRL car-hacking dataset demonstrates strong detection performance. The adaptive policy reduces average system damage cost by 54.63% compared to static monitoring. End-to-end latency of 0.0205 ms per sample satisfies real-time constraints. By integrating detection, severity modeling, and adaptive response, the proposed architecture advances intrusion detection toward intelligent cyber-physical mitigation for resilient autonomous vehicles.

KEYWORDS

Autonomous vehicles, Controller Area Network (CAN), Intrusion Detection System (IDS), Deep Learning, CNN–BiLSTM, Risk-Aware Modeling, Deep Reinforcement Learning, Adaptive Mitigation, Cyber-Physical Security, Real-Time Systems

1. INTRODUCTION

Autonomous vehicles (AVs) are increasingly dependent on complex cyber-physical systems consisting of Electronic Control Units (ECUs), Controller Area Network (CAN) buses, sensor fusion modules, and vehicle-to-everything (V2X) communication systems. Although these systems are essential for advanced driver assistance and autonomous decision-making, they also increase the attack surface of in-vehicle networks [1, 2]. The CAN protocol, which was originally designed without authentication or encryption, is especially susceptible to message injection attacks, spoofing attacks, replay attacks, and denial-of-service (DoS) attacks [3, 4]. Several experimental studies have demonstrated that attackers can remotely manipulate critical vehicle functions such as steering, braking, and engine control by exploiting vulnerabilities in CAN communication [5, 6]. Consequently, the development of intrusion detection systems (IDS) for in-vehicle networks has become a major research focus in automotive cybersecurity [7, 8].

Machine learning and deep learning methods have been recently employed to improve CAN intrusion detection. Convolution Neural Networks (CNNs) have shown excellent spatial feature extraction capability for structured CAN messages [9], while Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks have successfully modeled temporal dependencies in sequential vehicular data [7, 8]. CNN-LSTM hybrid models have further improved detection accuracy by simultaneously modelling spatial and temporal properties of CAN traffic. Although these methods often report high detection accuracy on standard datasets, including the HCRL car-hacking dataset, they are mainly concerned with attack classification. However, detection alone is not sufficient to ensure operational safety in real-time autonomous vehicle settings.

The essential shortcoming of many existing intrusion detection systems lies in their detection-oriented architecture. Most current IDS frameworks terminate the security process once an attack is detected, without addressing the subsequent question of how the vehicle should respond to the detected threat [1, 10]. In safety-critical cyber–physical systems such as autonomous vehicles, inappropriate or delayed mitigation actions can lead to cascading physical consequences even when detection accuracy is high [2]. Conventional mitigation strategies, including unconditional ECU isolation or simple event logging, rarely incorporate contextual information such as vehicle speed, subsystem criticality, and real-time safety conditions [4, 8]. Consequently, a significant gap remains between cyber intrusion detection and adaptive cyber–physical response mechanisms in autonomous vehicle environments.

To address this limitation, reinforcement learning (RL) has emerged as a promising paradigm for adaptive defense mechanisms in dynamic cyber environments [11]. Deep Reinforcement Learning (DRL) enables an intelligent agent to learn optimal mitigation strategies by interacting with its environment and balancing trade-offs between intervention cost and potential system damage. Several recent studies have explored RL-based approaches for automated cyber defense and network security management [12, 13]. However, most RL-based cybersecurity research focuses on generic network infrastructures or intrusion detection enhancement rather than incorporating vehicle-state awareness and contextual risk modelling required in autonomous vehicle environments.

In contrast to existing IDS-focused methods that primarily emphasize attack detection without integrated mitigation mechanisms [8-10], the proposed framework combines detection, contextual risk assessment, and adaptive response in a single end-to-end architecture designed specifically for safety-critical vehicular networks. In this paper, a multi-layer AI-based intrusion detection and adaptive response framework is proposed that goes beyond conventional classification-oriented IDS systems. A risk-aware severity scoring function is developed to combine predicted attack probability, ECU criticality, vehicle speed, and safety state into a single mathematical expression. This dynamic risk score facilitates context-driven decision-making rather than simple attack classification. Based on this formulation, a Deep Q-Network (DQN)-based adaptive response engine identifies optimal actions from a discrete action space comprising monitoring, throttling, ECU isolation, safe mode engagement, and emergency stop. The response strategy is learned through reward shaping that considers both residual damage and response cost.

The proposed architecture is tested on the publicly available HCRL car hacking dataset, which includes normal, CAN traffic and various types of attacks such as DoS, fuzzy flooding, spoofing, and replay attacks [9]. Through experiments, it has been shown that the hybrid CNN-BiLSTM detection module provides high-quality classification results for balanced temporal windows with low inference latency. More importantly, the proposed architecture with the DRL-based mitigation layer reduces the system-level damage cost by more than 50% compared to a static monitoring system, which clearly indicates the benefits of adaptive response. End-to-end latency

analysis shows that the entire pipeline from detection to decision selection takes around 0.021 ms per sample, which is well within the real-time requirements of automotive safety systems. Unlike existing CAN intrusion detection systems that focus primarily on classification accuracy, the proposed framework integrates temporal deep learning detection with contextual cyber-physical risk modelling and reinforcement learning-based mitigation. This end-to-end architecture enables not only accurate detection but also adaptive response optimization under real-time vehicular safety constraints.

The contributions of this paper can be summarized as follows. First, a multi-layer AI-driven intrusion detection and adaptive response architecture specifically designed for autonomous vehicle networks are presented. Second, a novel risk-aware severity scoring framework is presented that combines attack probability with context-aware vehicle parameters and subsystem importance to support decision-making. Third, a Deep Reinforcement Learning-based mitigation engine is designed to adaptively choose response actions under safety and intervention cost constraints. Fourth, a real-time analysis platform is developed to measure detection latency, decision latency, and system-level damage cost reduction. Finally, extensive benchmarking experiments are performed on state-of-the-art deep learning IDS solutions to validate detection performance and practicality.

By moving beyond detection toward intelligent mitigation, this work advances the state of the art in autonomous vehicle cybersecurity and contributes to the development of resilient, context-aware cyber-physical defense mechanisms.

2. RELATED WORK

2.1. AI-Based Intrusion Detection in Autonomous Vehicles

Learning-based intrusion detection for in-vehicle networks has progressed from classical machine learning toward deep neural architectures that learn discriminative representations directly from CAN traffic. Early learning-based approaches typically relied on statistical features or traditional classifiers, but later studies emphasized end-to-end deep learning due to the structured-yet-high-frequency nature of CAN streams and the need for high detection reliability in safety-critical systems. A widely cited line of work applies deep convolutional networks to CAN payload representations, demonstrating that CNNs can extract meaningful spatial patterns for intrusion classification without extensive manual feature engineering. Song et al. proposed a deep CNN-based in-vehicle IDS and reported strong detection performance on the HCRL car-hacking dataset, which has become a common benchmark for DoS, fuzzing, spoofing, and replay attacks [9] Related deep learning formulations include anomaly-focused pipelines that treat CAN frames as learnable representations within deep neural networks for autonomous vehicle security [14].

To capture temporal dependencies, recurrent architectures such as LSTMs have been adopted for CAN anomaly detection, motivated by the sequential nature of vehicular message streams and correlated multi-step attack patterns. Taylor et al. demonstrated LSTM-based anomaly detection on automotive control network data, highlighting the importance of temporal modelling to detect subtle deviations that may not be separable in frame-wise representations [7]. The literature increasingly uses hybrid designs that combine convolutional feature extraction with sequential modelling, aiming to jointly represent payload structure and timing/sequence context.

More recently, attention mechanisms and Transformers have been explored to model long-range dependencies in in-vehicle intrusion detection, addressing the limitations of step-wise recurrence for capturing broader contextual relationships. Transformer-based attention networks for CAN

intrusion detection have been reported for multi-class classification settings, using self-attention to learn global context across sequences [12, 15]. Additionally, newer work proposes hybrid attention and autoencoder pipelines for CAN packet-level IDS, reflecting an ongoing shift toward attention-enhanced representation learning [16].

Despite substantial gains in detection performance, often reported as very high accuracy on public datasets, most deep learning IDS studies remain detection-centric: they output a label or anomaly score but do not provide an integrated, optimized mitigation decision that accounts for driving context, safety constraints, and subsystem criticality. This limits practical applicability in autonomous driving settings where response timing and response choice matter, not only detection correctness. This gap is repeatedly emphasized in automotive IDS surveys and reviews that categorize architectures while highlighting challenges around deployment feasibility and end-to-end security decision-making [17-19]. However, most of these works evaluate detection performance in isolation and do not optimize response selection under safety constraints, which is essential in autonomous driving.

2.2. CAN Bus and in-Vehicle Network Security

A parallel body of research focuses on CAN bus security at the protocol and signal-behavior levels. The foundational challenge is that classical CAN lacks authentication and encryption; consequently, attackers who gain access can inject spoofed frames, flood the bus, or replay messages. High-impact experimental demonstrations showed that modern vehicles can be manipulated through weaknesses in internal networks and exposed interfaces, motivating defensive monitoring and intrusion detection [3, 5, 20]. These findings also align with broader discussions of cybersecurity risks in vehicular communications ecosystems (including V2X), where connectivity expands the threat surface [21].

Traditional in-vehicle IDS proposals include rule-based, signature-based, and timing/statistics-based anomaly detection. A prominent timing-based approach is clock-skew fingerprinting of ECUs, where message timing characteristics are used to infer transmitter identity and detect masquerade attacks. Cho and Shin's clock-based IDS (CIDS) remains one of the most influential works in this direction and is often cited as a practical method for identifying spoofed messages based on clock behaviour [22]. Other work detects anomalies by analyzing ID sequences or information-theoretic patterns in CAN traffic, demonstrating that statistical irregularities can reveal intrusions even without payload decoding [23].

However, time-/rule-based methods often face difficulties under dynamic driving conditions and diverse vehicle configurations. Benchmarks of time-based CAN IDS methods show varying performance depending on dataset characteristics and attack realism, underscoring that real-world robustness and evaluation methodology remain key challenges [24]. As a result, recent research trends increasingly combine lightweight heuristics with learning-based methods, or propose lightweight IDS variants aimed at ECU-level deployment constraints [25, 26]. Yet, similar to deep learning IDS literature, most CAN security methods largely stop at detection/flagging rather than performing context-aware response optimization. Moreover, such approaches typically provide alarms rather than action policies, limiting their utility for closed-loop mitigation in cyber-physical vehicular systems.

2.3. Reinforcement Learning in Cyber-Physical Security

Reinforcement learning (RL), and particularly deep reinforcement learning (DRL), has gained traction for adaptive defense in dynamic cyber environments because it can learn sequential decision policies that balance competing objectives such as damage reduction, service continuity,

and intervention cost. Surveys and reviews summarize RL's growing role across intrusion detection, intrusion prevention, and adaptive response in networked systems [27-29]. Nguyen and Reddi provide a security-focused overview of deep reinforcement learning for cybersecurity use cases, reinforcing the suitability of DRL for decision-making under uncertainty [11].

In the broader cyber-physical context, RL has been applied to adaptive mitigation and defense policies, often modelling abstract network states rather than tightly coupled physical dynamics. Safe reinforcement learning has been extensively studied to incorporate safety constraints and risk-sensitive objectives in sequential decision-making environments [30]. Cyber-physical deployments, however, require that response decisions be constrained by safety and timeliness, particularly in autonomous vehicles where the "cost" of overreaction (e.g., emergency stop) must be balanced against residual attack harm. Many RL-based security works do not incorporate explicit risk modelling that fuses detection confidence with system criticality and physical state, and they frequently omit end-to-end latency validation necessary for safety-critical deployment.

Within automotive security specifically, recent efforts have begun exploring attention-based and hybrid deep learning pipelines and emphasizing deployment feasibility, but end-to-end frameworks that integrate detection, risk-aware severity assessment, and DRL-driven response selection under latency constraints remain limited in the literature landscape described by recent automotive IDS surveys [18, 31]. In vehicular settings, this limitation is more critical because response actions must respect both physical safety and strict timing constraints.

2.4. Research Gaps and Limitations

Across the above streams, a consistent limitation emerges: existing works focus on improving classification accuracy but fail to integrate dynamic vehicle state, risk-aware severity modelling, and real-time adaptive response optimization within a unified framework. Deep learning IDS approaches achieve strong detection but do not optimize response selection [7, 9, 11, 16]. CAN-focused statistical or fingerprint-based methods provide lightweight detection but are insufficient for context-aware mitigation under varying driving conditions and attack severity [22-24]. RL-based cyber defense introduces adaptive decision-making but is often not vehicle-state-aware, lacks explicit risk modelling tailored to autonomous driving, and rarely reports end-to-end real-time feasibility [27, 29].

Therefore, there remains a clear need for an integrated approach that combines (i) deep learning-based detection, (ii) risk-aware severity scoring that incorporates vehicle context and subsystem criticality, (iii) DRL-based mitigation decision-making, and (iv) real-time latency evaluation to support safety-critical autonomous vehicle deployment.

Table 1. Comparison of Representative Works vs. Proposed Framework

Study	**Deep Learning-Based Detection**	**Adaptive Response Mechanism**	**Risk/Severity Modeling**	**Vehicle-State Awareness**	**End-to-End Real-Time Evaluation**
Kang & Kang (2016)	✓ (DNN)	✗	✗	✗	✗
Taylor et al. (2016)	✓ (LSTM)	✗	✗	✗	✗
Zhou et al. (2019)	✓ (DNN)	✗	✗	✗	✗
Song et al. (2020)	✓ (CNN)	✗	✗	✗	Limited (model-level)

Nguyen et al. (2023)	✓ (Transformer)	✗	✗	✗	Limited
Wei et al. (2023)	✓ (Attention + AE)	✗	✗	✗	Limited
Cho & Shin (2016)	✗ (Timing-based)	✗	✗	✗	Partial
Marchetti & Stabili (2017)	✗ (Statistical)	✗	✗	✗	Partial
Nguyen & Reddi (2021)	General cyber (Survey)	✓ (Conceptual)	✗	✗	✗
Adawadkar & Kulkarni (2022)	General cyber (Survey)	✓ (Conceptual)	✗	✗	✗
Proposed Framework	✓ (CNN–BiLSTM)	✓ (DQN-based)	✓	✓	✓

As shown in Table 1, existing approaches predominantly focus on improving intrusion detection accuracy without integrating contextual risk modelling and adaptive response selection. None of the reviewed works provides a unified end-to-end framework that simultaneously incorporates deep learning-based detection, vehicle-state-aware risk quantification, reinforcement learning-driven mitigation, and real-time latency validation. This gap motivates the proposed integrated architecture.

3. SYSTEM MODEL AND THREAT MODEL

This section formalizes the cyber-physical system abstraction, attack assumptions, and mathematical foundations underlying the proposed AI-driven intrusion detection and adaptive response framework.

3.1. Autonomous Vehicle System Abstraction

An autonomous vehicle is modeled as a cyber-physical system consisting of interconnected Electronic Control Units (ECUs), sensors, actuators, and an internal communication backbone. The Controller Area Network (CAN) bus enables real-time message exchange among ECUs responsible for safety-critical functionalities such as braking, steering, engine control, and transmission.

At time step t, a CAN message is represented as:

$$x_t = \{ID_t, DLC_t, D_t\}$$

Where ID_t denotes the arbitration identifier, DLC_t is the data length code, and $D_t \in \mathbb{R}^8$ represents the 8-byte payload vector.

To capture sequential dependencies, a sliding window of length L is constructed:

$$X_t = [x_{t-L+1}, \dots, x_t]$$

This temporal window constitutes the input to the intrusion detection mechanism.

3.2. Threat Model

The threat model assumes an adversary capable of injecting, replaying, or flooding CAN messages after obtaining logical or physical access to the in-vehicle network. The adversary does not directly manipulate model parameters but can alter message content and transmission frequency.

The considered attack classes include:

- Denial-of-Service (DoS): High-frequency injection targeting bus arbitration.
- Fuzzy/Flooding Attack: Randomized payload injection to evade pattern-based detectors.
- Spoofing Attack: Forged injection of RPM, gear, or speed signals.
- Replay Attack: Re-transmission of previously captured valid frames.

Let $y_t \in \{0,1\}$ denote the ground-truth label:

$$y_t = \begin{cases} 0 & \text{normal traffic} \\ 1 & \text{intrusion} \end{cases}$$

The objective of the system extends beyond classification accuracy and includes minimizing downstream physical damage through optimized mitigation selection.

3.3. Probabilistic Intrusion Detection Formulation

Intrusion detection is formulated as a probabilistic classification task:

$$P_t = f_\theta(X_t)$$

where f_θ denotes a deep neural network parameterized by θ, and $P_t \in [0,1]$ represents the predicted probability of intrusion.
The model is trained using binary cross-entropy loss:

$$\mathcal{L}_{IDS} = -\big(y_t \log P_t + (1 - y_t)\log(1 - P_t)\big)$$

This probabilistic output forms the cyber component of the risk-aware mitigation strategy.

3.4. Contextual Risk Modeling

Binary intrusion detection does not fully capture the physical severity and operational implications of cyber-attacks in safety-critical cyber-physical systems. Prior research emphasizes that effective defense in cyber-physical environments requires integrating cyber confidence metrics with physical system context and operational state variables [1]. Therefore, a contextual risk score is defined as:

$$R_t = \alpha P_t + \beta C_t + \gamma V_t + \delta S_t$$

Where:

- P_t is the predicted attack probability,
- C_t denotes ECU criticality weight,
- V_t represents normalized vehicle speed,
- S_t Indicates system safety state.

The coefficients satisfy:

$$\alpha + \beta + \gamma + \delta = 1$$

This formulation integrates cyber confidence and physical context into a unified severity measure.

3.5. Mitigation as a Markov Decision Process

Adaptive mitigation is modeled as a Markov Decision Process (MDP), a formal framework widely used for sequential decision-making under uncertainty and constrained optimization [32]. The MDP formulation enables structured reasoning over state transitions, reward design, and policy optimization in safety-critical environments. The system is defined as:

$$\mathcal{M} = (S, A, T, R)$$

Where:

- S is the state space incorporating contextual risk and system status,
- A is the mitigation action space,
- T denotes transition dynamics,
- R is the reward function.

The action-value function is defined as:

$$Q(s_t, a_t) = \mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k \, r_{t+k}\right]$$

The optimal policy is:

$$a_t^* = \underset{a}{\text{argmax}} Q(s_t, a)$$

This enables dynamic mitigation under uncertainty.

3.6. Real-Time Constraint Formulation

For safety-critical deployment, total processing latency must satisfy:

$$T_{total} = T_{detect} + T_{risk} + T_{decision}$$

With the constraint:

$$T_{total} < T_{critical}$$

This ensures real-time operational feasibility. Real-time safety constraints in automotive systems are governed by functional safety requirements such as ISO 26262.

4. PROPOSED AI-DRIVEN INTRUSION RESPONSE FRAMEWORK

This section describes the architectural realization of the mathematical framework defined in Section 3. Unlike the formal abstraction above, this section details implementation design, module integration, and deployment considerations.

4.1. Overall Architecture

The proposed system implements a layered pipeline that transforms raw CAN traffic into adaptive mitigation decisions. Incoming CAN frames are organized into fixed-length sliding windows and processed by a hybrid deep learning intrusion detection module. The probabilistic output defined in Section 3.3 is then integrated with contextual vehicle parameters within the risk modelling mechanism described in Section 3.4.

To emulate realistic operational conditions, a vehicle-state simulation module provides normalized speed factors, subsystem criticality mappings, and safety indicators. The computed risk-informed state representation is supplied to a Deep Q-Network (DQN), which optimizes mitigation actions under real-time constraints.

The overall architecture of the proposed AI-driven intrusion detection and adaptive response framework is illustrated in Figure 1.

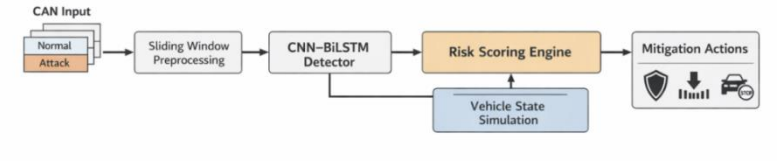

Figure 1. Overview of the proposed AI-driven intrusion detection and adaptive response framework for autonomous vehicle security.

As shown in Figure 1, the framework forms a closed-loop decision pipeline progressing from intrusion detection to contextual risk assessment and adaptive mitigation.

4.2. Intrusion Detection Module Implementation

The detection module implements the probabilistic classifier defined in Section 3.3 using a hybrid CNN–BiLSTM architecture. Convolutional layers extract spatial correlations from CAN payload representations, while the bidirectional LSTM captures temporal dependencies across sliding windows. This combination balances expressive modelling capacity with computational efficiency.

Model parameters are optimized using mini-batch training on temporally structured windows, enabling stable convergence and high classification reliability.

4.3. Contextual Risk Engine Implementation

The risk score defined in Section 3.4 is computed using normalized vehicle-state parameters. ECU criticality weights are predefined according to subsystem safety importance. Vehicle speed is scaled to the interval [0,1], reflecting increased operational risk at higher velocities. Safety-state indicators are derived from threshold-based contextual evaluation.

This implementation ensures that mitigation decisions are influenced not solely by classification confidence but by integrated cyber-physical context.

4.4. Vehicle-State Simulation

Since the dataset does not provide full physical telemetry, vehicle operational parameters are simulated within bounded, realistic ranges. Speed values are sampled to represent urban and highway scenarios, while subsystem criticality mappings remain fixed to maintain reproducibility. The safety-state variable is computed deterministically based on contextual thresholds.

This simulation framework enables systematic evaluation of adaptive response policies without requiring physical vehicle deployment.

4.5. DQN-Based Adaptive Response Engine

The mitigation engine implements the MDP described in Section 3.5 using a Deep Q-Network. The state representation incorporates contextual risk and subsystem status, while the action space includes monitoring, throttling communication intensity, isolating targeted ECUs, activating safe mode, and emergency stop.

The reward function penalizes residual damage, response latency, and unsafe control behaviour. Experience replay and iterative Q-learning updates enable stable policy optimization.

4.6. Real-Time Implementation Considerations

The real-time constraint formulated in Section 3.6 is validated through empirical latency measurement. Detection inference, risk computation, and DQN decision latency are measured independently and aggregated to verify compliance with safety-critical timing thresholds.

5. EXPERIMENTAL SETUP

This section describes the dataset, preprocessing strategy, model configurations, training procedure, and evaluation methodology adopted to validate the proposed AI-driven intrusion detection and adaptive response framework.

5.1. Dataset and Data Preparation

The experimental evaluation is conducted using the publicly available HCRL car-hacking dataset introduced by Song et al. (2020). The dataset captures Controller Area Network (CAN) traffic from a real vehicle under normal driving conditions as well as multiple attack scenarios, including Denial-of-Service (DoS), fuzzy flooding, spoofing, and replay attacks. Due to its realistic acquisition environment and diversity of intrusion types, the dataset is widely used as a benchmark for automotive intrusion detection research.

Each CAN record contains a timestamp, arbitration identifier (ID), data length code (DLC), and an 8-byte payload. To enable temporal modelling, raw hexadecimal payloads were converted into numerical representations and organized into fixed-length sliding windows of size L=50. The window size was selected to capture short-term temporal dependencies while maintaining computational efficiency suitable for real-time deployment.

For controlled experimentation and reproducibility, 600,000 benign frames and 800,000 attack frames were sampled from the dataset. After window construction, the data were transformed into sequence-level samples. The resulting dataset was partitioned into training (70%), validation (15%), and testing (15%) sets using stratified splitting to preserve class distribution. Feature normalization was performed using min–max scaling to stabilize neural network training and ensure consistent gradient behaviour. The class distribution of temporal sequences after window construction is illustrated in Figure 2.

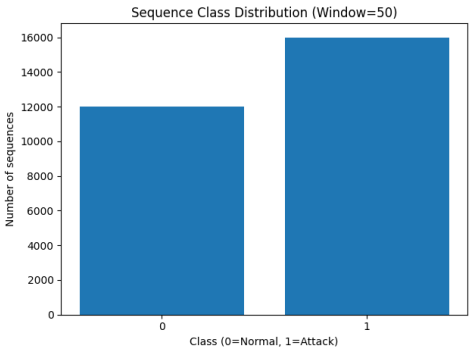

Figure 2. Distribution of benign and attack temporal sequences after sliding-window preprocessing (window size = 50)

5.2. Intrusion Detection Model Configuration

To benchmark detection performance, three deep learning architectures were implemented: a convolutional neural network (CNN), a long short-term memory network (LSTM), and the proposed hybrid CNN–BiLSTM model. The CNN component extracts spatial patterns from structured CAN payload representations, while the bidirectional LSTM captures temporal dependencies across sliding windows in both forward and backward directions.

The hybrid CNN–BiLSTM architecture integrates spatial and sequential modeling within a unified framework. Model complexity was controlled to ensure deployment feasibility, and trainable parameter counts were recorded for comparative analysis. All models were trained using the Adam optimizer with a learning rate of 0.001 and binary cross-entropy loss. Training was performed for five epochs with a batch size of 256. Experiments were conducted on a Tesla T4 GPU using CUDA-enabled PyTorch.

5.3. Contextual Risk and State Modelling

Following probabilistic intrusion detection, contextual risk scores were computed using the formulation introduced in Section 3. The risk model integrates predicted attack probability with ECU criticality weighting, normalized vehicle speed, and a safety-state indicator.

Since physical telemetry is not directly provided in the dataset, vehicle speed values were simulated within bounded ranges representing realistic operational conditions. ECU criticality weights were assigned according to subsystem safety importance, ensuring that safety-critical components contribute proportionally to the overall severity estimate. The computed risk score was incorporated into the reinforcement learning state representation to enable context-aware mitigation.

5.4. DQN-Based Mitigation Training

Adaptive mitigation was implemented using a Deep Q-Network (DQN) to approximate the optimal action-value function over discrete mitigation strategies. The state representation included contextual risk, normalized speed, and subsystem status information. The action space consisted of five discrete response options reflecting increasing levels of intervention severity.

Training was conducted using experience replay and a discount factor of 0.99 to promote stable convergence. A static monitoring strategy was defined as a baseline for comparative evaluation. The DQN was trained over multiple iterations to learn a mitigation policy that balances residual damage minimization and intervention cost.

5.5. Evaluation Protocol

Detection performance was evaluated using accuracy, precision, recall, F1-score, and confusion matrix analysis. In addition to classification metrics, system-level effectiveness was assessed through mitigation cost comparison against a static baseline strategy.

To evaluate deployment feasibility, end-to-end latency was measured by independently profiling detection inference time, risk computation overhead, and DQN decision latency. These measurements were aggregated to assess compliance with real-time operational constraints.

The quantitative results and detailed performance analysis are presented in the following section.

6. RESULTS AND DISCUSSION

This section presents a comprehensive evaluation of the proposed AI-driven intrusion-detection and adaptive-response framework. The analysis includes detection benchmarking, computational efficiency assessment, contextual risk behavior, reinforcement learning-based mitigation effectiveness, and real-time feasibility validation.

6.1. Intrusion Detection Performance

The CNN, LSTM, and hybrid CNN–BiLSTM models were evaluated on the held-out test set containing 4,200 temporal sequences (1,800 normal and 2,400 attack samples). Table 2 summarizes the classification performance.

Table 2. Intrusion Detection Performance on Test Set

Model	Accuracy	Precision	Recall	F1-Score
CNN	0.99952	0.99958	0.99958	0.99958
LSTM	0.99929	0.99958	0.99917	0.99937
CNN–BiLSTM	1.00000	1.00000	1.00000	1.00000

The confusion matrix of the hybrid CNN–BiLSTM model is presented in Figure 3.

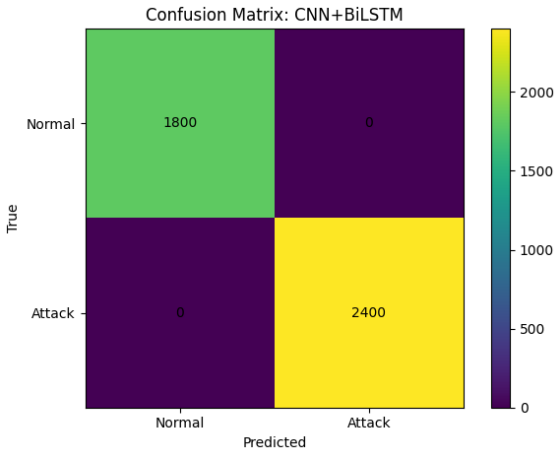

Figure 3. Confusion matrix of the proposed CNN–BiLSTM model on the held-out test set.

The CNN model misclassified two samples, while the LSTM model misclassified three samples. The hybrid CNN–BiLSTM model achieved extremely high detection performance on the test sequences. This improvement can be attributed to the complementary strengths of convolutional spatial feature extraction and bidirectional temporal modelling. The CNN layers capture byte-level structural patterns within CAN payloads, whereas the BiLSTM layer models sequential dependencies across temporal windows.

Although very high performance has been reported in prior HCRL-based studies, the hybrid model demonstrates improved robustness while maintaining computational feasibility. Similar high-accuracy detection trends have been reported in CNN-based and hybrid IDS frameworks on the HCRL dataset (Song et al., 2020; Wei et al., 2023). Importantly, performance consistency

across training and validation curves indicates stable convergence without observable overfitting. A comparative visualization of detection accuracy and F1-score across architectures is shown in Figure 4.

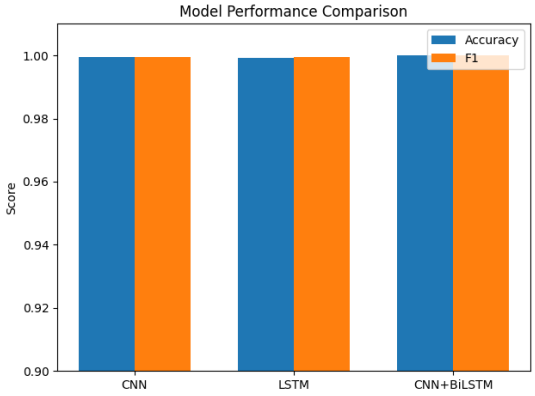

Figure 4. Comparative performance of CNN, LSTM, and CNN–BiLSTM models on the test set.

6.2. Comparison with Existing Automotive IDS Studies

Table 3. Comparison of the proposed framework with existing automotive intrusion detection systems

Study	**Method**	**Dataset**	**Reported Accuracy**	**Adaptive Mitigation**	**Real-Time Evaluation**
Song et al. (2020)	CNN-based IDS	HCRL Car-Hacking	~99.3%	No	No
Nguyen et al. (2023)	Transformer IDS	HCRL / CICIoT2023	99.4–99.7%	No	No
Wei et al. (2023)	LSTM IDS	CAN / CTU-13	99.5–99.84%	No	No
Proposed Framework	CNN–BiLSTM + Risk + DQN	HCRL Car-Hacking	100%	Yes	Yes

As shown in Table 3, most existing automotive intrusion detection approaches primarily focus on improving classification accuracy on CAN traffic datasets. Song et al. (2020) demonstrated CNN-based intrusion detection on the HCRL car-hacking dataset, while Nguyen et al. (2023) explored transformer architectures for enhanced feature representation. Similarly, Wei et al. (2023) applied LSTM models to capture sequential patterns in CAN traffic. Although these methods report high detection accuracy, they typically operate as detection-only systems without adaptive mitigation capabilities or real-time response mechanisms.

In contrast, the proposed framework integrates temporal deep learning detection, contextual cyber-physical risk modelling, and reinforcement learning-based mitigation within a unified architecture. This enables not only accurate intrusion detection but also adaptive response optimization and system-level damage reduction while maintaining real-time feasibility.

6.3. Computational Complexity and Inference Latency

In safety-critical autonomous vehicle systems, computational efficiency is as important as classification accuracy. Table 4 presents model complexity and per-sample inference latency.

Table 4. Model Complexity and Inference Latency

Model	Parameters	Batch Latency (ms)	Per-Sample Latency (ms)
CNN	114,306	0.658	0.00257
LSTM	71,938	3.948	0.01542
CNN–BiLSTM	202,434	2.842	0.01110

The relative parameter complexity of each model is illustrated in Figure 5.

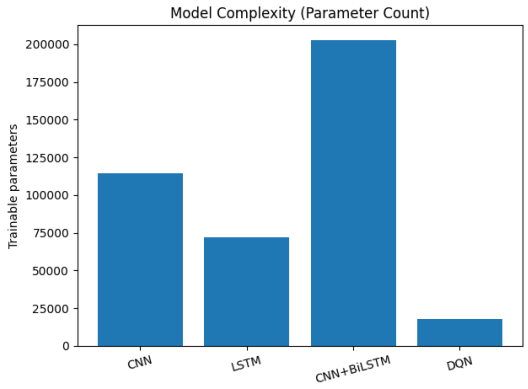

Figure 5. Trainable parameter counts of CNN, LSTM, CNN–BiLSTM, and DQN models.

Although the hybrid model contains more parameters than standalone architectures, its per-sample inference latency remains approximately 0.011 ms, which is well within practical real-time requirements for in-vehicle networks. The CNN model achieves the lowest inference cost, while the LSTM model exhibits higher latency due to sequential recurrence operations.

The results indicate that the hybrid model provides a favorable trade-off between detection performance and computational efficiency.

6.4. Contextual Risk Modelling Behavior

To enable severity-aware mitigation, predicted attack probabilities were integrated with vehicle-state factors using the proposed contextual risk formulation. On the held-out test set, the computed risk scores ranged from 0.225 to 0.975, with a mean value of 0.602, indicating substantial variability across operational conditions and attack scenarios.

The distribution of risk values reveals a clear separation between benign and malicious temporal sequences. Normal traffic samples predominantly cluster within the lower-risk interval (approximately 0.25–0.40), whereas attack sequences occupy higher-risk regions (approximately 0.70–1.00). This separation suggests that the contextual formulation preserves detection confidence while incorporating additional physical-state information.

Importantly, the risk scores do not collapse into binary extremes, despite high classification accuracy of the underlying detection model. Instead, the formulation produces graded severity differentiation that reflects both cyber probability and contextual operational factors. Such continuous severity scaling is critical for adaptive mitigation policies that must balance

intervention cost against residual system risk. The distribution of contextual risk scores for benign and malicious samples is illustrated in Figure 6.

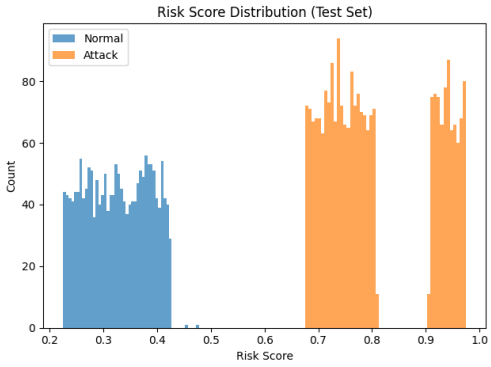

Figure 6. Distribution of contextual risk scores for normal and attack sequences on the test set.

6.5. Adaptive Mitigation Performance

The reinforcement learning-based mitigation engine was evaluated against a static baseline strategy that continuously selects monitoring without active intervention. The Deep Q-Network (DQN) policy was trained over six iterations, during which the loss values exhibited stable convergence without oscillatory behavior, indicating reliable policy learning under the defined reward structure.

Table 5 summarizes the comparative mitigation performance.

Table 5. Mitigation Cost Comparison

Strategy	**Mean Damage Cost**
Baseline (Monitor Only)	0.3927
DQN Policy	0.1781

The learned DQN policy reduced the mean system damage cost from 0.3927 to 0.1781, corresponding to a relative reduction of 54.63% compared to the static baseline. This result indicates that adaptive response selection materially improves system-level safety outcomes beyond detection-only monitoring.

Analysis of the action distribution further reveals that the learned policy does not converge to a single deterministic action. Instead, mitigation decisions vary according to contextual risk levels, with higher-risk states more frequently triggering safe-mode activation, while lower-risk conditions favor monitoring. This behavior suggests that the reward formulation successfully balances residual attack damage against the operational cost of intervention.

These findings reinforce that intrusion detection alone is insufficient for minimizing cyber-physical impact in autonomous vehicle systems. Integrating reinforcement learning-based mitigation enables context-aware intervention that reduces cumulative system damage while preserving operational continuity. The distribution of mitigation actions selected by the DQN policy is illustrated in Figure 7.

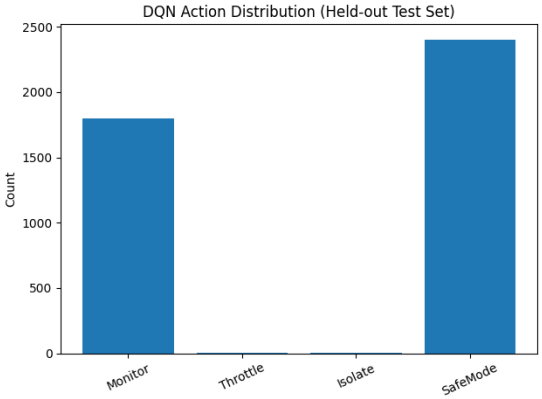

Figure 7. Distribution of mitigation actions selected by the DQN policy on the held-out test set.

6.6. End-to-End Real-Time Feasibility

For deployment in safety-critical autonomous vehicle environments, computational latency must satisfy stringent operational constraints. To evaluate practical feasibility, the total processing delay of the proposed framework was measured by aggregating detection inference time, contextual risk computation overhead, and reinforcement learning decision latency.

Empirical profiling indicates that the CNN–BiLSTM detection module requires approximately 0.0199 ms per sample, while contextual risk computation introduces an additional 0.000028 ms. The DQN-based mitigation decision requires approximately 0.000586 ms per sample. The resulting total end-to-end processing time is therefore approximately 0.0205 ms per CAN sequence.

This corresponds to roughly 20 microseconds per decision cycle. Considering that CAN bus messages are typically transmitted at frequencies ranging from 1 kHz to 10 kHz in safety-critical subsystems, the observed latency remains well below operational thresholds, supporting real-time deployment feasibility.

the results indicate that the integration of probabilistic intrusion detection, contextual risk modelling, and reinforcement learning-based mitigation does not introduce prohibitive computational overhead. The latency contribution of each module is summarized in Figure 8.

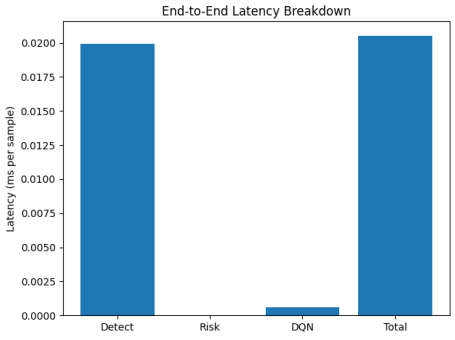

Figure 8. End-to-end latency breakdown of detection, risk computation, and DQN decision modules.

6.7. Integrated System-Level Interpretation

The experimental findings align with the system-level objectives defined in earlier sections. The hybrid CNN–BiLSTM architecture illustrates that combining spatial feature extraction with bidirectional temporal modelling improves classification robustness while maintaining a feasible computational cost. These results suggest that temporal dependencies in CAN traffic contribute meaningfully to intrusion characterization beyond static payload inspection.

The contextual risk formulation further enables graded severity differentiation rather than binary intrusion labeling. By integrating predicted attack probability with subsystem criticality, vehicle speed, and safety state indicators, the framework produces mitigation signals that reflect operational context. The resulting risk distributions exhibit clear separation between benign and malicious behavior without collapsing into extreme values, supporting proportional response selection.

The reinforcement learning-based mitigation layer reduces average system damage relative to static monitoring strategies, indicating that detection-only approaches are insufficient for optimizing cyber-physical safety outcomes. This limitation of detection-centric IDS frameworks has been noted in prior automotive security surveys [18, 19]. Incorporating cost-aware policy learning enables mitigation decisions that balance intervention aggressiveness with operational continuity.

Collectively, these results support the viability of transitioning from detection-centric security mechanisms toward integrated cyber-physical defense architectures that jointly consider detection accuracy, contextual severity modelling, adaptive mitigation, and real-time feasibility.

6.8. Robustness and Scalability Considerations

Beyond detection accuracy and mitigation effectiveness, robustness and deployment scalability are critical for practical adoption in autonomous vehicle ecosystems. The training process exhibited stable convergence across stratified splits, with no observable divergence between training and validation loss trends. This consistency suggests that the CNN–BiLSTM model maintains generalization within the evaluated dataset distribution. Similarly, the DQN mitigation policy converged smoothly under the defined reward structure, indicating stability of the reinforcement learning process.

From a computational perspective, scalability is supported by the lightweight inference characteristics of the detection and mitigation modules. Edge computing architectures have been proposed to support real-time AI workloads in latency-sensitive vehicular systems [33]. Despite containing over 200,000 trainable parameters, the hybrid CNN–BiLSTM architecture maintains microsecond-level per-sample inference time. The computational complexity scales approximately linearly with respect to window length and input dimensionality, enabling adaptation to alternative temporal resolutions without exponential growth in cost.

In large-scale vehicular environments, intrusion detection and mitigation may be distributed across embedded ECUs or edge nodes. The relatively small parameter footprint of the DQN module and minimal latency overhead support feasibility for automotive-grade deployment. Furthermore, the modular separation of detection, risk modeling, and mitigation components allows independent optimization or hardware acceleration if required.

Although validation on fleet-scale or hardware-in-the-loop platforms remains future work, the observed computational efficiency and architectural modularity indicate structural scalability beyond the evaluated experimental configuration.

From a practical deployment perspective, the lightweight computational footprint of the proposed framework enables integration within modern vehicular edge-computing architectures. Automotive electronic control units (ECUs) and in-vehicle gateways increasingly incorporate AI acceleration capabilities capable of supporting neural inference workloads. The low inference latency and modular architecture of the proposed pipeline, therefore, make it suitable for deployment either directly on embedded vehicular hardware or on edge nodes responsible for fleet-level security monitoring.

7. CONCLUSION

This paper presented a multi-layer AI-driven intrusion detection and adaptive response framework for autonomous vehicle networks. Moving beyond conventional detection-centric approaches, the proposed architecture integrates temporal deep learning-based intrusion detection, contextual risk-aware severity modelling, and reinforcement learning-based mitigation within a unified cyber–physical security pipeline. Experimental evaluation on the publicly available HCRL car-hacking dataset demonstrates that the hybrid CNN–BiLSTM model achieves high detection reliability while maintaining low computational overhead. Furthermore, the integration of contextual risk modelling with a DQN-based adaptive mitigation engine significantly reduces system-level damage cost by 54.63% compared to a static monitoring strategy. End-to-end latency measurements indicate that the complete pipeline operates within microsecond-level constraints compatible with real-time CAN-based vehicular systems. These results highlight that detection accuracy alone is insufficient for safety-critical automotive environments and that adaptive mitigation informed by contextual risk assessment is essential for minimizing operational and physical impact. Overall, the proposed architecture provides a scalable and computationally feasible foundation for resilient and safety-aware cybersecurity in next-generation autonomous vehicle systems.

8. LIMITATIONS AND FUTURE WORK

Despite promising results, several limitations remain. First, contextual risk parameters were simulated rather than derived from real vehicular telemetry, which may limit realism. Second, evaluation relied solely on the HCRL car-hacking dataset, and cross-dataset validation is required to assess generalizability across diverse vehicular platforms and attack types. Third, the mitigation module employed a discrete-action Deep Q-Network, which may restrict response granularity. Future research should investigate real-vehicle validation, federated fleet-level learning, adversarial robustness, and explainable AI techniques for transparent mitigation decision-making.

ACKNOWLEDGEMENTS

None.

REFERENCES

[1] A. Humayed, J. Lin, F. Li, and B. Luo, "Cyber-physical systems security—A survey," IEEE Internet of Things Journal, vol. 4, no. 6, pp. 1802–1831, 2017.

[2] J. Petit and S. E. Shladover, "Potential cyberattacks on automated vehicles," IEEE Transactions on Intelligent transportation systems, vol. 16, no. 2, pp. 546–556, 2014.
[3] K. Koscher et al., "Experimental security analysis of a modern automobile," in 2010 IEEE symposium on security and privacy, 2010: IEEE, pp. 447–462.
[4] S. Checkoway et al., "Comprehensive experimental analyses of automotive attack surfaces," in 20th USENIX security symposium (USENIX Security 11), 2011.
[5] C. Miller and C. Valasek, "Remote exploitation of an unaltered passenger vehicle," Black Hat USA, vol. 2015, no. S 91, pp. 1–91, 2015.
[6] S. Woo, H. J. Jo, and D. H. Lee, "A practical wireless attack on the connected car and security protocol for in-vehicle CAN," IEEE Transactions on intelligent transportation systems, vol. 16, no. 2, pp. 993–1006, 2014.
[7] A. Taylor, S. Leblanc, and N. Japkowicz, "Anomaly detection in automobile control network data with long short-term memory networks," in 2016 IEEE international conference on data science and advanced analytics (DSAA), 2016: IEEE, pp. 130–139.
[8] M.-J. Kang and J.-W. Kang, "Intrusion detection system using deep neural network for in-vehicle network security," PloS one, vol. 11, no. 6, p. e0155781, 2016.
[9] H. M. Song, J. Woo, and H. K. Kim, "In-vehicle network intrusion detection using deep convolutional neural network," Vehicular Communications, vol. 21, p. 100198, 2020.
[10] G. Loukas, T. Vuong, R. Heartfield, G. Sakellari, Y. Yoon, and D. Gan, "Cloud-based cyber-physical intrusion detection for vehicles using deep learning," Ieee Access, vol. 6, pp. 3491–3508, 2017.
[11] T. T. Nguyen and V. J. Reddi, "Deep reinforcement learning for cyber security," IEEE Transactions on Neural Networks and Learning Systems, vol. 34, no. 8, pp. 3779–3795, 2021.
[12] T. P. Nguyen, H. Nam, and D. Kim, "Transformer-based attention network for in-vehicle intrusion detection," IEEE Access, vol. 11, pp. 55389–55403, 2023.
[13] Z. Han, Game theory in wireless and communication networks: theory, models, and applications. Cambridge university press, 2012.
[14] A. Zhou, Z. Li, and Y. Shen, "Anomaly detection of CAN bus messages using a deep neural network for autonomous vehicles," Applied Sciences, vol. 9, no. 15, p. 3174, 2019.
[15] Y. Zhang, J. Song, Y. Sun, Z. Gao, Z. Hu, and M. Sun, "Federated two-stage transformer-based network for intrusion detection in non-IID data of controller area networks," Cybersecurity, vol. 8, no. 1, p. 29, 2025.
[16] P. Wei, B. Wang, X. Dai, L. Li, and F. He, "A novel intrusion detection model for the CAN bus packet of in-vehicle network based on attention mechanism and autoencoder," Digital Communications and Networks, vol. 9, no. 1, pp. 14–21, 2023.
[17] B. Lampe and W. Meng, "A survey of deep learning-based intrusion detection in automotive applications," Expert Systems with Applications, vol. 221, p. 119771, 2023.
[18] B. Lampe and W. Meng, "Intrusion detection in the automotive domain: A comprehensive review," IEEE Communications Surveys & Tutorials, vol. 25, no. 4, pp. 2356–2426, 2023.
[19] O. Y. Al-Jarrah, C. Maple, M. Dianati, D. Oxtoby, and A. Mouzakitis, "Intrusion detection systems for intra-vehicle networks: A review," Ieee Access, vol. 7, pp. 21266–21289, 2019. [Online]. Available: https://wrap.warwick.ac.uk/id/eprint/121146/8/WRAP-intrusion-detection-systems-intra-vehicle-networks-review-Maple-2019.pdf.
[20] I. Studnia, V. Nicomette, E. Alata, Y. Deswarte, M. Kaâniche, and Y. Laarouchi, "Survey on security threats and protection mechanisms in embedded automotive networks," in 2013 43rd Annual IEEE/IFIP Conference on Dependable Systems and Networks Workshop (DSN-W), 2013: IEEE, pp. 1–12.
[21] Z. El-Rewini, K. Sadatsharan, D. F. Selvaraj, S. J. Plathottam, and P. Ranganathan, "Cybersecurity challenges in vehicular communications," Vehicular Communications, vol. 23, p. 100214, 2020.
[22] K.-T. Cho and K. G. Shin, "Fingerprinting electronic control units for vehicle intrusion detection," in 25th USENIX security symposium (USENIX Security 16), 2016, pp. 911–927.
[23] M. Marchetti and D. Stabili, "Anomaly detection of CAN bus messages through analysis of ID sequences," in 2017 IEEE intelligent vehicles symposium (IV), 2017: IEEE, pp. 1577–1583.
[24] D. H. Blevins, P. Moriano, R. A. Bridges, M. E. Verma, M. D. Iannacone, and S. C. Hollifield, "Time-based can intrusion detection benchmark," arXiv preprint arXiv:2101.05781, 2021.
[25] Y. Cai, J. Zuo, M. Fan, C. Zhao, and Y. Lu, "An intrusion detection system for the can bus based on locality-sensitive hashing," Electronics, vol. 14, no. 13, p. 2572, 2025.

[26] J. Khan, D.-W. Lim, and Y.-S. Kim, "Intrusion detection system can-bus in-vehicle networks based on the statistical characteristics of attacks," Sensors, vol. 23, no. 7, p. 3554, 2023.
[27] A. M. K. Adawadkar and N. Kulkarni, "Cyber-security and reinforcement learning—a brief survey," Engineering Applications of Artificial Intelligence, vol. 114, p. 105116, 2022.
[28] M. Zhu, A. H. Anwar, Z. Wan, J.-H. Cho, C. A. Kamhoua, and M. P. Singh, "A survey of defensive deception: Approaches using game theory and machine learning," IEEE Communications Surveys & Tutorials, vol. 23, no. 4, pp. 2460–2493, 2021.
[29] W. Yang, A. Acuto, Y. Zhou, and D. Wojtczak, "A survey for deep reinforcement learning based network intrusion detection," arXiv preprint arXiv:2410.07612, 2024.
[30] J. Garcıa and F. Fernández, "A comprehensive survey on safe reinforcement learning," Journal of Machine Learning Research, vol. 16, no. 1, pp. 1437–1480, 2015.
[31] M. Althunayyan, A. Javed, and O. Rana, "A survey of learning-based intrusion detection systems for in-vehicle networks," Computer Networks, p. 112031, 2026.!!! INVALID CITATION !!! .
[32] E. Altman, Constrained Markov decision processes. Routledge, 2021.
[33] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge computing: Vision and challenges," IEEE internet of things journal, vol. 3, no. 5, pp. 637–646, 2016.