ITM Web of Conferences **79**, 01031 (2025) https://doi.org/10.1051/itmconf/20257901031
*KEIS-2025*

# Rule-Based Reasoning in Dynamic Knowledge Transformer-based Intrusion Detection System for Autonomous Vehicles

*Parutagouda* S. Khanagoudar[1], *Savita* Muchakhandi[2], *Sangamesh*. H[3*], *Kushal* Badgujar[4], and *Sarala*. D V[5]

[1]Department of Computer Science and Engineering, Sambhram Institute of Technology, Bengaluru, India
[2]Department of Electronics and Communication Engineering, Nitte Meenakshi Institute of Technology, Nitte (Deemed to be University), Bengaluru, India
[3]Department of Electronics and Communication Engineering, H.K.E Society's Sir M Visvesvaraya College of Engineering, Raichur, India
[4]Department of Electronics Engineering, COEP Technological University, Pune, India
[5]Department of Computer Science and Engineering, BMS College of Engineering, Bengaluru, India

**Abstract.** In recent years, Autonomous vehicles (AVs) require robust reasoning mechanisms to ensure safety and efficiency dynamic traffic environments. However, existing Hybrid Intrusion Detection System using Deep Reinforcement Learning (HIDS-DRL) model limited by static rule hierarchies, complications in resolving rule conflicts, and lacking generalization to unseen situations. To overcome these limitations, this research proposes the Rule-based Reasoning in Dynamic Knowledge Transformer for AV (RDK-AV) framework that combines symbolic reasoning and data-driven learning. Initially, a dynamic knowledge base encodes traffic laws, contextual constraints, and safety priorities, updated continuously with environmental perception. After that, a rule-based reasoning layer filters candidate actions and resolves conflicts adaptively, while a knowledge-infused transformer captures temporal traffic patterns and aligns them with rule embeddings to support contextual reasoning. Then, a decision fusion module integrates both outputs to generate safe, explainable and adaptive driving policies with increased robustness for autonomous navigation in real-world scenarios. Experimental results demonstrates that the proposed RDK-AV model achieved fairly good performance in optimizing the whole traffic system with Average Cumulative Delay (ACD) per vehicle of 2.6%, when compared with existing HIDS-DRL model.

## 1 Introduction

In the past few years, Autonomous Vehicles (AVs) have changed the landscape of intelligent transportation systems by utilizing sensors, machine-learning algorithms, and vehicle-to-everything (V2X) connectivity to operate with decreased human intervention [1]. The intrinsic connection between AVs and vast cyber-physical infrastructures–all of the dependencies that exist, from in-vehicle networks to cloud- or edge-services– have exposed AVs to sophisticated cyberattacks [2]. Intrusion Detection Systems (IDS) are extremely important to protect AVs from malicious threats and build their safety and trustworthiness [3]. The challenge here is the diverse and layered networks of AVs, where abbreviated decision-making is crucial, little tolerance for delay exists, and attacks can operate through multiple vectors, such as perception sensors, control modules, and V2X channels [4]. The more traditional signature and rule-based IDS systems generally cannot detect novel or zero-day attacks in dynamic environments. There have been recent advances in areas of Deep-Learning (DL) based IDS models, specifically supervised anomaly detection and hybrid DL architectures applicable to vehicular contexts [5].

These models are effective at detecting novel and known attacks, but still present issues with scalability, explainability, and overcoming adversarial attacks. Motivated by the safety-critical nature of AV operations, there is a growing need for adaptive, lightweight, and explainable IDS solutions [6]. It dynamically learns from evolving attack patterns, interoperate across heterogeneous vehicular platforms, and provide actionable intelligence to both on-board decision systems and remote security operation centers [7]. Federated learning and edge-assisted IDSs have also emerged as promising directions to support collaborative training across vehicles while taking advantage of data privacy but still have missed the issues of synchronization and communication overhead [8]. In addition, state-of-the-art IDS frameworks increasingly adopt multimodal fusion by integrating data from V2X messages to improve detection robustness, yet this integration raises challenges of data heterogeneity and alignment [9]. The design of IDS for autonomous vehicles must balance detection accuracy, computational efficiency, and real-time responsiveness, thereby ensuring secure and resilient autonomous

* Corresponding author: sangamesh.hosgurmath@gmail.com

mobility. Moreover, resource constraints in embedded automotive hardware being resource-constrained and ultra-low latency being required, it is difficult to transfer traditional enterprise IDS solutions [10].

The contribution of this research include:

- Rule-based Reasoning in Dynamic Knowledge Transformer for AV (RDK-AV) integrates a dynamic knowledge base with a knowledge-infused transformer which enables the autonomous vehicle to make decisions that are both legally compliant and contextually adaptive in dynamic traffic environments.
- Through the combination of symbolic action scores, predictions based on a transformer, and real-time risk assessment, the framework produces safety aware and interpretable driving policies allowing traceability of how each component contributes to the final decision.
- Adaptive and real-time decision making ensures temporal adaptivity, allowing RDK-AV to respond to evolving traffic patterns, unforeseen obstacles, and environmental changes while maintaining robust autonomous navigation.

The remainder of this paper is organized as follows: Section 2 establishes the literature review, Section 3 demonstrates the proposed methodology, Section 4 specifies the results of the proposed model with existing models as well as discussion, and Section 5 summarizes the overall conclusion of the research.

## 2 Literature review

Cynthia Anthony et al. [11] suggested Non-Tree based Machine Learning (NT-ML) algorithm for intrusion detection system in autonomous vehicles. Initially, the RDK-AV framework constructed a dynamic knowledge base that encoded traffic rules, contextual constraints, and safety priorities, which was continuously updated through real-time environmental perception from LiDAR, radar, and cameras. Then, a rule-based reasoning layer evaluated candidate driving actions against the knowledge base, filtering out unsafe maneuvers and assigning scores based on feasibility and priority. After that, a knowledge-infused transformer captured temporal traffic patterns and contextual dependencies by aligning sequential sensory inputs with embedded rules, producing context-aware representations of dynamic traffic conditions. However, this framework required high computational resources for real-time updates and transformer operations.

Lukas Westhofen et al. [12] introduced Ontologies formalization and recognition of criticality for automated driving. Initially, the ontology-based framework formalized criticality phenomena by constructing a T-Box that encoded regulatory knowledge, traffic rules, and scenario semantics derived from expert-defined catalogs. Then, the system transformed heterogeneous traffic data, collected from multiple sources such as drones, vehicles, and cameras, into ontology-compliant A-Boxes that represented entities, roles, and relationships in a structured manner. After that, reasoning engines using description logic and rule-based inference evaluated the A-Boxes against the formalized ontology to detect and classify safety-relevant criticality events. However, the high computational cost of ontology reasoning, especially with large A-Boxes, and the sensitivity of inference accuracy to incomplete or uncertain sensor data, which could limit real-time deployment in highly dynamic traffic scenarios.

Xiaohui Zhang et al. [13] represented Multi-Lane Autonomous Driving (MLAD) using Reinforcement Learning (RL) with Traffic Flow Adaptive (TFA) respectively. Initially, the hierarchical framework for multi-lane autonomous driving was designed by combining deep reinforcement learning with rule-based methods to overcome the limitations of purely data-driven or purely symbolic approaches. Then, a high-level deep reinforcement learning agent that produced the Instantaneous Desired Speed (IDS) which was an intermediate decision variable that showed the vehicle's desire to accelerate or lane-change under different traffic states. After that, low-level rule-based car-following models and lane-changing models did the action based on the provided IDS, while upholding safety and interpretability, and remaining aware of their dynamic environment for adaptive actions. However, the computational demand during training and reliance on accurate parameter calibration, which could affect real-time performance in highly congested or uncertain environments.

Majd Alkorabi et al. [14] presented rule-based RL for autonomous vehicle communications with modern intelligent transportation. Initially, deep learning algorithms for autonomous vehicles were designed to improve decision-making, object detection, and trajectory planning by leveraging high-dimensional sensory data from LiDAR, radar, and cameras. Then, reinforcement learning strategies such as Deep Q-Networks and Proximal Policy Optimization were applied to model adaptive driving behaviors, enabling vehicles to handle dynamic environments like lane-changing, overtaking, and collision avoidance. After that, hybrid approaches combining rule-based constraints with machine learning were developed to balance interpretability and flexibility. However, the high computational demand of real-time training and inference with reduced generalization under adverse weather and unseen conditions, which limited deployment reliability.

Guanzhi Xiong et al. [15] demonstrated Large Language Model (LLM) decision making for AV across multiple driving scenarios.Initially, the Dilu framework to integrate large language models into autonomous vehicle decision-making, treating driving as a language-based reasoning task. Then, a reasoning engine processed environmental descriptions through scene encoding, memory recall, and prompt generation, allowing the LLM to reason step by step before producing driving decisions. After that, a reflection engine analyzed successful and failed experiences to improve adaptability, while a memory storage module maintained both successful and reflective knowledge for future decision support. It also provided improved interpretability by exposing reasoning chains. However, its high computational cost, latency in real-time decision-making, and performance instability at higher speeds, which limited direct deployment in safety-critical real-world driving contexts

## 3 Methodology

The proposed RDK-AV framework integrates symbolic reasoning and data-driven learning to achieve safe, adaptive, and explainable autonomous navigation. It begins with a dynamic knowledge base that encodes traffic rules, contextual constraints, and safety priorities, continuously updated with real-time sensor observations. A rule-based reasoning layer evaluates candidate actions using hard and soft constraints, resolving conflicts and incorporating probabilistic confidence measures. Parallelly, a knowledge-infused transformer captures temporal traffic patterns and aligns them with rule embeddings. Finally, a decision fusion module combines symbolic and learned predictions with risk-aware adjustment to generate contextually adaptive, interpretable, and safe driving policies suitable for dynamic real-world scenarios. The overall performance of this framework is shown in Figure 1.

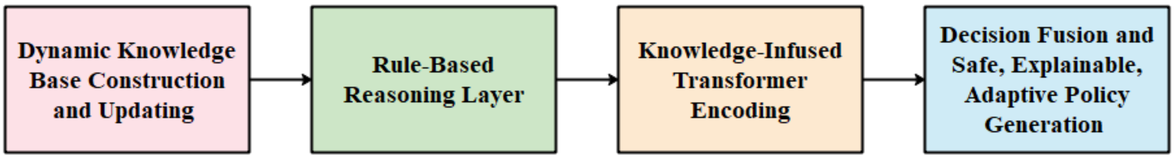


**Fig. 1.** Block diagram of the proposed RDK-AV framework.

### 3.1 Dynamic knowledge base construction and updating

In this research, the foundational layer of the RDK-AV framework, as it provides a structured repository for encoding traffic rules, contextual constraints, and safety priorities necessary for autonomous vehicle's navigation. Static rules models are not sufficient for navigation in dynamic driving environments in which traffic conditions, road layouts, and events are constantly changing. In contrast, the constraints for a vehicle operating in a dynamic environment requires a dynamic knowledge base that keep both some symbolic domain knowledge and real-time environmental context consistent. Formally, the DKB at time $t$ is defined as a tuple of facts, rules, and confidence measures defined in Equation (1):

$$\mathcal{K}_t = \{(f_i, r_i, w_i) \mid i = 1,2, \dots, N\} \tag{1}$$

Where, $f_i$ represents individual facts such as road types, lane markings, traffic signals, or vehicle state observations; $r_i$ denotes rule statements encoding relationships between facts and permissible actions ("if red light, then stop"). Similarly, $w_i \in [0,1]$ represents a confidence weight which indicates the reliability of the fact-rule pair and is dynamically updated depending on the sensor perception and the historic continuity of the traffic environment. In order to account for real time perception, environmental observations that are made by the LiDAR, camera and radar systems, collectively denoted as $\mathcal{E}_t$, are utilized to update the knowledge base. The operative update is expressed in the following Equation (2):

$$\mathcal{K}_{t+1} = \mathcal{U}(\mathcal{K}_t, \mathcal{E}_t) \tag{2}$$

Where, $\mathcal{U}$ denotes the knowledge update operator, which revises the observable based on existing weights and new rules or facts to add to the existing knowledge base. For a fact $f_i$, observed at time t, the weight update is shown as following Equation (3):

$$w_i^{t+1} = \lambda w_i^t + (1-\lambda) conf(f_i, \mathcal{E}_t) \tag{3}$$

Where, $\lambda \in [0,1]$ is a forgetting factor that regulates the weight of old knowledge, and $\mathrm{conf}(f_i, \mathcal{E}_t)$ represents the confidence of $f_i$ based on sensor evidence, derived from probabilistic models and Bayesian confidence levels. Also, if $r_i$ previously unencoded patterns are observed, new rules are dynamically added, shown as in Equation (4):

$$\mathcal{K}_{t+1} \leftarrow \mathcal{K}_{t+1} \cup \{(f_j, r_j, w_j)\} \tag{4}$$

Through this dynamic updating the knowledge base is never out of sync with real-world observations, ensuring accuracy and completeness. By uniting formal, structured symbolic knowledge with perceptual feedback, allows the AV system to record domain knowledge but be supple towards changing traffic conditions. The final DKB is the main input to the rule-based reasoning layer, which ensures each candidate action now abides by both the pre-defined rules and situational context, thereby establishing a safe and explainable foundation for autonomous driving decisions.

### 3.2 Rule-based reasoning layer

Following the constructure, the core symbolic inference layer of the RDK-AV framework---the reasoning layer. It is responsible for transmuting knowledge loaded in the dynamic knowledge base into actionable and safe decisions. The DKB provides a structured collection of facts, rules, and associated confidence weights, but knowledge in and of itself cannot drive an autonomous vehicle through complex, dynamic traffic environments. Therefore, a rule-based reasoning layer is implemented to keep track of candidate actions, enforce laws of the road, resolve conflicting situations of competing rules, and maintain higher priority safety-critical actions. Formally, let the candidate action set at time be $\mathcal{A}_t = \{a_1, a_2, \dots, a_m\}$, where each $a_j$ candidate action represents an intended vehicle maneuver. The reasoning process evaluates each action against the knowledge base $\mathcal{K}_t$ through a symbolic inference function, which is represented as Equation (5):

$$s_j = R(\mathcal{K}_t, a_j) \tag{5}$$

Where, $s_j$ indicates the feasibility and priority score assigned to action $a_j$ respectively, then the function $R(.)$ is broken into two critique aspects: hard constraints such as legal traffic rules and soft constraints including safety margins and comfort levels. Hard constraints are dealt with logically and inferences are enforced as predicates: If $r_i(f_i)$ breaks the hard constraints, then $s_j = 0$, where actions can not break fundamental laws of the road like running a red light. Soft constraints distill scores into reasonable weighted penalties defined in Equation (6):

$$s_j = s_j - \sum_k \beta_k \phi_k(f_i, a_j) \tag{6}$$

Where, $\phi_k$ is the distance from safety or comfort violation, and $\beta_k$ is a weight factor reflecting the relative importance of the constraint. By scoring each candidate action without changing its meaning, the RBR layer is able to resolve conflicts of competing rules, and select actions that maximized the critical with the best remaining preferences. A lane-change action is penalized if the preceding action was a high-speed maneuver without adequate spacing. Mathematically, the temporal penalty can be represented as Equation (7):

$$s_j^t = s_j - \gamma\psi(a_{t-1}, a_j) \tag{7}$$

Where, $\psi(.)$ captures risk induced by consecutive maneuvers and $\gamma$ refers a tunable parameter. Furthermore, to account for uncertainties in perception and partial knowledge, the RBR layer uses probabilistic reasoning by treating the fact weights $w_i$ from the DKB as confidence measures. The candidate action score will become as Equation (8):

$$s_j^{final} = s_j \cdot \prod_{i \in \mathcal{F}_j} w_i \tag{8}$$

Where, $\mathcal{F}_j \subseteq \mathcal{K}_t$ is the set of facts relevant to action $a_j$, and the probabilistic reasoning allows the RBR layer to flexibly balance rule compliance with the uncertainty in the environment. Filtering out infeasible actions, resolving conflicts, and ordering actions, the RBR layer guarantees that every candidate behavior is compliant with traffic laws at their point and time of execution (including the real time, uncertainty laden safety constraints). The output of the layer, a prioritized action set $\mathcal{A}_t^{RBR}$ with associated scores, will serve as one of the primary inputs to the decision fusion module. Moreover, the symbolic reasoning predictions from the RBR layer are processed with the data-driven predictions performed by the transformer-based encoder to produce final adaptive decision policies for the vehicle. Next, This complementary fusion of structured symbolic reasoning with probabilistic confidence measures makes for a solid, trustworthy platform for how to autonomously navigate.

### 3.3 Knowledge-Infused transformer encoding

In this step, the data-driven reasoning component of the RDK-AV framework complementing the symbolic outputs from rule-based reasoning layer by capturing complex contextual dependencies and temporal traffic patterns. The rule-based reasoning ensures compliance with traffic laws, it does not model high-dimensional correlations in dynamic environments, such as interactions among multiple vehicles, pedestrians, and traffic signals over time. To overcome these limitations, this framework employes a transformer-based architecture infused with knowledge embeddings derived from the dynamic knowledge based, allowing for context-aware temporal modeling. Additionally, consider a sequence of observations over a temporal window $L$ represented as $X_{t-L:t} = [x_{t-L}, \dots, x_t]$, where $x_\tau$ includes vehicle states, sensor inputs, and traffic features at time $\tau$. Each fact-rule pair $(f_i, r_i)$ from the dynamic knowledge based is embedded into vector space using an embedding function $E_r$, which is represented as Equation (9):

$$E_r(f_i, r_i) \in \mathbb{R}^{d_c} \tag{9}$$

Where, $d_e$ denotes the embedding dimension, the knowledge-infused transformer incorporates both sequential observations and rule embeddings as input to a multi-head self-attention mechanism. If consider one head, the first- order attention scores are computed as in Equation (10):

$$Attention(Q, K, V) = softmax(\frac{QK^\top}{\sqrt{d_k}})V \tag{10}$$

Where queries Q, keys K, and values V are projections from the concatenated sensory embeddings and rule embeddings: $H_\tau = [x_\tau; E_r(f_i, r_i)]$. By directly embedding knowledge into the attention computation of the transformer, latent representations of temporal events are placed in opposition to symbolic constraints in such a way during the learning process that learned features respect traffic rules but represent nuance in the environment. The multi-head self-attention outputs a context-specific representation as expressed in Equation (11):

$$H_t = \mathcal{T}_\theta(X_{t-L:t}, E_r(\mathcal{K}_t)) \in \mathbb{R}^{L \times d_k} \tag{11}$$

Where, $\mathcal{T}_\theta$ is the transformer parameterized by $\theta$ and $d_h$ is the hidden dimension, these embeddings encode both temporal dependencies and rule-oriented context. Also, enabling future downstream modules to reason over sequential structures to predict safe and adaptive maneuvers in reconfiguring traffic flows. Positional encodings are also added to preserve the order of the sequence, while the multi-headed attention allows for parity of modeling the diverse relational dependencies between agents and environmental features. The knowledge-infused transformer enables a hybrid of symbolic reasoning and deep learning applications through embeddings that respect contextual relevance to actions and adherence to rules of traffic. The generated rule-relevant embedding context formed the principal inputs to the decision fusion module comprised of queueing the symbolic action scores from the rule - based reasoning layer with the action predictions generated with the transformer mechanism.

### 3.4 Decision fusion and safe, explainable, adaptive policy generation

The final step describes the integrative and execution stages of the RDK-AV framework, combining outputs from the symbolic image reasoning and predictive knowledge-infused transformer decision-makers into rational, safe, and interpretable driving actions. The rule-based reasoning layer attaches a score to candidate actions according to the encoded traffic rules and safety constraints contained in this symbolic knowledge base. It also attaches a score to candidate actions based on the transformer-based traffic trajectory predictions that capture thoughtful interpretations of temporal traffic patterns while accounting for complex environmental dependencies. The scoring of candidate actions is fused in an Equation (12):

$$S(a_j) = \alpha s_j^{RBR} + (1-\alpha)\hat{s}_j^T \quad (12)$$

Where, α weights the relative influence of the symbolic rules and the data-driven learning because the environment is constantly changing, such as when surrounding vehicles change speed or appear with no warning, the symbol-based rules do not capture every source of risk. Therefore, a real-time risk metric extracted from the sensor using Equation (13):

$$S^{safe}(a_j) = S(a_j) - \lambda R(a_j, \mathcal{E}_t) \quad (13)$$

Where, $R(a_j, \mathcal{E}_t)$ assesses the degree of risk based on possible threats like nearby vehicles, pedestrians, or unsafe lane positions, and λ denotes the degree of sensitivity to the risk. The final driving action is selected according to which candidate action has the highest safety adjusted score shown as in Equation (14):

$$a_t^* = arg \max_{a_j \in \mathcal{A}_t} S(a_j) \quad (14)$$

Therefore, this selection process provides a systematic means to balance compliance with traffic laws, the ability to be responsive to the changing dynamics of traffic, and accounting for current risk. Explainability is embedded in this framework because each component of the final score-rule influence, data-driven prediction, and risk adjustment is traced and analyzed individually, which is formulated as Equation (15):

$$S(a_t^{safe}) = \underbrace{\alpha s_j^{RBR}}_{Rule\ Influence} + \underbrace{(1-\alpha)\hat{s}_j^T}_{Data\text{-}Driven\ Influence} - \underbrace{\lambda R(a_j, \mathcal{E}_t)}_{Risk\ Adjustment} \quad (15)$$

Finally, adaptivity is achieved through continuous updates of both the dynamic knowledge base and the transformer embeddings, allowing the AV system to adjust its policies in real-time in response to evolving traffic conditions, unexpected obstacles, or changes in road topology. This adaptive, fused policy can be formally represented as a time-dependent mapping $\pi_t: (\mathcal{K}_t, H_t, \mathcal{E}_t) \mapsto a_t^{safe}$, which captures the complete decision-making process. By integrating symbolic reasoning, data-driven predictions, and risk-aware adjustments in a single framework, RDK-AV produces safe, explainable, and contextually adaptive driving actions, suitable for deployment in complex, dynamic, real-world traffic scenarios. Next, the performance of proposed RDK-AV model is illustrated in below section.

## 4 Experimental results

The Experimental results were obtained in a high-performance computing environment to ensure efficient model training and validation. This environment uses multicore processors such as Intel i7, at least 16 GiB of random-access memory (RAM), and a GPU such as NVIDIA GTX 1080 to compute DL training. In particular, the software environment includes Python and DL libraries such as TensorFlow and PyTorch. The proposed RDK-AV framework integrates a dynamically updated knowledge base, rule-based reasoning, and knowledge-infused transformer encoding to generate safe, adaptive, and explainable driving policies. Symbolic rules ensure compliance with traffic laws, where transformer embeddings capture temporal traffic patterns. A decision fusion module combines these outputs with real-time risk assessment, producing robust, context-aware actions for autonomous navigation in complex, dynamic environments. The evaluation metrics like Average Cumulative Delay (ACD) per vehicle is utilized, which are given in below Equation (16):

$$ACD = \frac{1}{N}\sum_i^N D_i \quad (16)$$

Here, $N$ denotes total number of vehicles observed and $D_i$ indicates total delay experienced by the $i$th vehicle, which is the difference between actual travel time and ideal travel time.

### 4.1 Performance analysis

**Table 1.** Performance analysis of proposed RDK-AV.

| Performance Models | ACD (%) |
|---|---|
| EBRB | 6.1 |
| LLMs | 5.3 |
| Proposed RDK-AV | 2.6 |

The performance of the proposed RDK-AV model is compared with that of various existing methods such as EBRB and LLMs, using evaluation metrics. Existing approaches are listed in Table 1.

Table 1 presents the performance of the proposed RDK-AV framework compared with the existing methods. The RDK-AV model achieved better results than the conventional approaches, including ACD of 2.6% per vehicle, outperforming the EBRB model in terms of ACD (97.14%) per vehicle. Similarly, the LLMs model had ACD (97.94%) per vehicle. The RDK-AV integrates a dynamic knowledge base, rule-based reasoning, and knowledge-infused transformer encoding, fused with real-time risk assessment, to generate safe, adaptive, and explainable autonomous driving policies in dynamic traffic environments.

### 4.2 Comparative analysis

The proposed RDK-AV framework is compared with the existing HIDS-DRL [13] model using comprehensive evaluation metrics. An evaluation of the proposed RDK-AV method is presented in Table 2.

**Table 2.** Comparative analysis of proposed RDK-AV.

| Comparative Models | ACD (%) |
|---|---|
| HIDS-DRL [13] | 4.9 |
| Proposed RDK-AV | 2.6 |

Table 2 shows that the proposed RDK-AV framework obtained better results than the existing HIDS-DRL [13] model in terms of ACD of 2.6% per vehicle respectively. The existing HIDS-DRL [13] model obtained minimal results in terms of ACD of 4.9% per vehicle, which represent significant improvements in model performance.

### 4.3 Discussion

The proposed RDK-AV countermeasure framework is designed to fill the long-standing gaps not covered by previous approaches such as EBRB systems, LLMs, and HIDS-DRL [13] approaches. LLM-based defenses improve contextual reasoning at scale and demonstrably use natural language-based threat intelligence requirements gathering from unstructured reports of the attack. However, due to their computational overhead, uncertainty in the decision-making rationale, and inherent susceptibility to adversarial prompt manipulation. The EBRB relies primarily on static event-condition-action rules that make it less resilient to threats presented by novel or evolving adversaries. The HIDS-DRL [13] model offers a semi-automated, host-based intrusion detection system integrated with reinforcement learning to facilitate the learning of autonomous policies that optimize according to dynamic threat conditions. However, due to their training costs, delayed convergence, and vulnerability to reward exploitation, HIDS-DRLs are insufficiently operationally reliable in time-critical environments. To overcome these challenges, this research proposed the RDK-AV model which combines strong deep knowledge representation and adversarial validation capabilities, leveraged knowledge graphs for semantic-based inference modeling while employing adversarial trained policies to resist evasion attempts. This overlapping architecture affords proactive detection of emergent and novel attack vectors, adaptive reconfiguration of non-static defensible applications, and subject to knowledge models, verifiable explainable decision support.

## 5 Conclusion

In this research, the proposed RDK-AV framework combines symbolic rule-based reasoning and knowledge-infused transformer learning in an attempt to create safe, explainable, adaptive driving policies for autonomous vehicles. The framework is designed and implemented with a dynamic knowledge base capable of constantly encoding traffic laws including static and contextualized constraints and real-time observations from the environment. The symbolic rule-based reasoning layer assesses candidate maneuvers, resolves conflicts if necessary, and assures safety. The knowledge-infused transformer captures temporal patterns of traffic and complex interactions between multiple agents as they manifest in decreasingly flexible contexts. The symbolic-safe and knowledge-based insights inform the final critical risk-based decision module, resulting in interpretable action scores that weigh symbolically compliant reasoning, predictive reasoning, and real-time mitigation of hazards. Experimental evaluations demonstrated that RDK-AV achieves minimized ACD per vehicle of 2.6%, when compared with existing HIDS-DRL model. Future work will include multi-agent coordination and probabilistic risk modelling to facilitate decision making in robust manner in high density urban and dynamically uncertain traffic environments.

## References

1. B. Lampe, W. Meng, Intrusion detection in the automotive domain: A comprehensive review. IEEE Commun. Surv. Tutor. **25**, 2356-2426 (2023). 10.1109/COMST.2023.3309864
2. M. Almehdhar, A. Albaseer, M.A. Khan, M. Abdallah, H. Menouar, S. Al-Kuwari, A. Al-Fuqaha, Deep learning in the fast lane: A survey on advanced intrusion detection systems for intelligent vehicle networks. IEEE Open J. Veh. Technol. **5**, 869-906 (2024). 10.1109/OJVT.2024.3422253
3. E.E. Abdallah, A. Aloqaily, H. Fayez, Identifying intrusion attempts on connected and autonomous vehicles: A survey. Procedia Comput. Sci. **220**, 307-314 (2023). https://doi.org/10.1016/j.procs.2023.03.040
4. J. Nagarajan, P. Mansourian, M.A. Shahid, A. Jaekel, I. Saini, N. Zhang, M. Kneppers, Machine learning based intrusion detection systems for connected autonomous vehicles: A survey. Peer-to-Peer Netw. Appl. **16**, 2153-2185 (2023). https://doi.org/10.1007/s12083-023-01508-7

5. S. Anbalagan, G. Raja, S. Gurumoorthy, R.D. Suresh, K. Dev, IIDS: Intelligent intrusion detection system for sustainable development in autonomous vehicles. IEEE Trans. Intell. Transp. Syst. **24**, 15866-15875 (2023). 10.1109/TITS.2023.3271768
6. M. Sorkhpour, A. Yazdinejad, A. Dehghantanha, Auto-CIDS: An autonomous intrusion detection system for vehicular networks. In Proceeding of the Workshop on Autonomous Cybersecurity, Association for Computing Machinery, New York, NY, United States, November 07 (2024), 45-55
7. S. Anbalagan, G. Raja, S. Gurumoorthy, K. Ayyakannu, Blockchain assisted hybrid intrusion detection system in autonomous vehicles for industry 5.0. IEEE Trans. Consum. Electron. **69**, 881-889 (2023). 10.1109/TCE.2023.3320282
8. L. Wang, X. Zhang, D. Li, H. Liu, Multi-sensors space and time dimension based intrusion detection system in automated vehicles. IEEE Trans. Veh. Technol. **73**, 200-215 (2023). 10.1109/TVT.2023.3306345
9. M. Moradi, M. Kordestani, M. Jalali, M. Rezamand, M. Mousavi, A. Chaibakhsh, M. Saif, Sensor and decision fusion-based intrusion detection and mitigation approach for connected autonomous vehicles. IEEE Sens. J. **24**, 20908-20919 (2024). 10.1109/JSEN.2024.3397966
10. P. Mansourian, N. Zhang, A. Jaekel, M. Kneppers, Deep learning-based anomaly detection for connected autonomous vehicles using spatiotemporal information. IEEE Trans. Intell. Transp. Syst. **24**, 16006-16017 (2023). 10.1109/TITS.2023.3286611
11. C. Anthony, W. Elgenaidi, M. Rao, Intrusion detection system for autonomous vehicles using non-tree based machine learning algorithms. Electronics **13**, 809 (2024). https://doi.org/10.3390/electronics13050809
12. L. Westhofen, C. Neurohr, M. Butz, M. Scholtes, M. Schuldes, Using ontologies for the formalization and recognition of criticality for automated driving. IEEE Open J. Intell. Transp. Syst. **3**, 519-538 (2022). https://doi.org/10.1109/OJITS.2022.3187247
13. X. Zhang, J. Sun, Y. Wang, J. Sun, A hierarchical framework for multi-lane autonomous driving based on reinforcement learning. IEEE Open J. Intell. Transp. Syst. **4**, 626-638 (2023). 10.1109/OJITS.2023.3300748
14. M. Alkorabi, A. Souri, N. İnanç, Deep learning algorithms for autonomous vehicle communications: Technical insights and open challenges. Concurr. Comput. Pract. Exp. **37,** e70218 (2025). https://doi.org/10.1002/cpe.70218
15. G. Xiong, S. Liu, Y. Yan, Q. Li, H. Li, Efficacy of autonomous vehicle's adaptive decision-making based on large language models across multiple driving scenarios. IEEE Access. **13**, 108076-108092 (2025). https://doi.org/10.1109/ACCESS.2025.3580641