Journal of Computer Technology and Software

ISSN: 2998-2383

Vol. 4, No. 7, 2025

Sensor Fusion and Deep Reinforcement Learning for Decision-Making in Autonomous Vehicles

Linnea Whitmore
Eastern Washington University, Cheney, USA
linneaw46 @ewu.edu

Abstract: The rise of autonomous driving technologies has prompted intensive research into intelligent decision-making systems capable of operating reliably under real-world conditions. This paper proposes a robust decision-making framework that integrates sensor fusion with deep reinforcement learning (DRL) to improve the performance of autonomous vehicles in complex urban environments. The system processes data from LiDAR, radar, and camera sensors to construct a unified environmental representation, which is then fed into a deep Q-network (DQN) to determine optimal driving actions. Experiments in a high-fidelity simulation environment demonstrate the effectiveness of the proposed framework in reducing collision rates, improving route efficiency, and maintaining real-time responsiveness, outperforming rule-based and unimodal DRL baselines. Our findings highlight the critical importance of multi-modal perception integration in conjunction with learning-based policy optimization for safe and intelligent autonomous navigation.

Keywords: Autonomous Vehicles; Sensor Fusion; Deep Reinforcement Learning; Decision-Making; Urban Navigation; Bird's Eye View Representation; LiDAR; Radar; DQN; CARLA Simulator

1. Introduction

The development of autonomous vehicles (AVs) has emerged as one of the most transformative technologies of the 21st century, promising to revolutionize transportation by reducing accidents, alleviating traffic congestion, and improving mobility for the elderly and disabled. Central to achieving reliable autonomy is the vehicle' s decision-making capability—how it processes sensory input and selects actions to safely navigate its environment. Traditional rule-based decision systems, while interpretable and efficient in simple environments, often lack the adaptability required to handle complex, dynamic, and uncertain real-world scenarios. These methods struggle to generalize beyond their designed conditions and cannot effectively respond to unpredictable behaviors from other road users.

To address these limitations, deep reinforcement learning (DRL) has been introduced as a powerful method for learning optimal policies through trial and error, guided by reward signals. DRL has shown success in various control tasks, including lane following, obstacle avoidance, and intersection negotiation. However, many DRL-based systems in AVs rely on raw camera data or limited sensor modalities, which makes them vulnerable to sensor-specific failure modes such as occlusion, poor lighting, or radar noise. This raises a key question: how can we build a decision-making framework that not only learns from interaction but also benefits from the reliability and redundancy of sensor fusion?

This paper proposes a hybrid system that integrates multi-modal sensor fusion with DRL to achieve robust and intelligent decision-making in urban driving scenarios. Our approach fuses LiDAR point clouds, camera imagery, and radar readings into a unified state representation using a lightweight perception module. This fused input is then processed by a deep Q-network (DQN), which learns to select discrete driving actions based on rewards that reflect safety, efficiency, and comfort. The key contributions of this paper include: (1) a novel sensor fusion pipeline for consistent spatial-semantic mapping, (2) a DRL-based policy module trained on diverse traffic scenarios, and (3) a comprehensive evaluation demonstrating significant gains over conventional and unimodal baselines.

2. Related Work

Research on autonomous vehicle (AV) decision-making has increasingly focused on integrating deep reinforcement learning (DRL) with advanced perception and data fusion techniques to address the complexities of real-world urban driving. Early DRL approaches such as Deep Q-Networks (DQN) and Advantage Actor-Critic (A3C) have demonstrated promising capabilities in sequential decision tasks and robust policy optimization across diverse applications[1-5]. These reinforcement learning methods are widely adopted not only in traditional control scenarios, but also in areas such as large model fine-tuning, resource allocation, and handling data imbalance, highlighting their versatility and optimization power in high-dimensional, dynamic environments.

Complementary to decision policy research, advances in multi-modal perception, feature fusion, and edge computing have driven the progress of real-time, robust AV systems. Recent studies leverage cross-scale attention, multi-layer feature fusion, and efficient network architectures to integrate diverse sensor data—including LiDAR, radar, and cameras—

into unified, high-quality representations for downstream decision modules. Such sensor fusion and compression strategies are critical for improving accuracy and enabling deployment on resource-constrained platforms[6-10].

Another important direction relates to anomaly detection, graph-based modeling, and temporal representation learning. Techniques from graph neural networks, probabilistic modeling, and sequential analysis have been shown to enhance the detection of rare events and better model evolving user or traffic behaviors. These methodologies provide strong foundations for safety-critical AV applications, where capturing temporal dependencies and complex data relationships is essential for robust generalization[11-15].

Collectively, these advances underscore the necessity of combining deep reinforcement learning, multi-modal fusion, and advanced representation learning for developing safe, adaptive, and high-performing autonomous driving systems.

3. Proposed Framework

To improve decision-making under complex urban conditions, we propose a unified framework that integrates multi-sensor perception and deep reinforcement learning. The architecture comprises two core modules: (1) a sensor fusion perception system that transforms heterogeneous sensor inputs into a spatially consistent representation, and (2) a DQN-based learning agent that maps fused environmental states to high-level driving actions. Raw input data are collected from LiDAR, radar, and RGB cameras. LiDAR point clouds are voxelized with a resolution of 0.1 meters and projected onto a 128×128 bird's-eye-view (BEV) occupancy grid. Radar signals are filtered using a Kalman filter to eliminate clutter and then transformed to the BEV frame to represent moving objects and their velocities. RGB images are processed using a DeepLabv3+ semantic segmentation model pretrained on the Cityscapes dataset, generating semantic masks for lane markings, vehicles, pedestrians, and traffic lights. These masks are reprojected to BEV coordinates using camera calibration parameters. The outputs from all three sensors are then concatenated along the channel axis to form a unified BEV tensor $I_{fusion} \in R^{128\times128\times C}$, where C represents the number of channels after stacking LiDAR reflectance, radar velocity, and semantic categories. This fused tensor is passed through a five-layer convolutional encoder with ReLU activations and batch normalization, followed by global average pooling to produce a compact latent state vector $s_t \in R^{256}$, which serves as input to the policy network.

The decision-making agent is based on the Deep Q-Network (DQN) algorithm, which learns to approximate the optimal action-value function using the Bellman update rule:

$$Q(s_t, a_t) = r_t + \gamma \max_{a'} Q(s_{t+1}, a'),$$

where s_t is the state vector at time t, a_t is the selected action, r_t is the immediate reward, and γ=0.99 is the discount factor. The action space is discrete and includes maneuvers such as maintaining lane, lane changes, braking, and turning. The reward function is designed to balance safety, efficiency, and passenger comfort, and is defined as:

$$r_t = -\lambda_1 \cdot \mathbb{I}_{\text{collision}} + \lambda_2 \cdot \Delta d_t - \lambda_3 \cdot j_t,$$

where $I_{collision}$ is an indicator function for collisions, Δd_t is the forward progress toward the navigation goal, and j_t denotes the jerk, computed as the time derivative of acceleration to penalize abrupt control actions. The hyperparameters are empirically set to λ1=10.0, λ2=1.0, and λ3=0.5.

Training is conducted using the CARLA simulator (v0.9.14), where diverse urban environments such as Town05 and Town10 are configured with randomized traffic actors and weather patterns to ensure robustness. The agent is trained over 500,000 simulation steps with an experience replay buffer of size 100,000 and a batch size of 64. The target network is updated every 500 iterations, and ε -greedy exploration is used to balance exploitation and exploration. The control interface maps each high-level action selected by the policy to trajectory waypoints, which are tracked using a PID-based low-level controller. This integration of fused perception and DRL-based policy optimization enables autonomous vehicles to make real-time, context-aware, and robust decisions across diverse scenarios.

4. Experiments and Results

To validate the effectiveness of the proposed sensor fusion and deep reinforcement learning (DRL) framework, we conducted comprehensive experiments using the CARLA simulator (v0.9.14), a high-fidelity urban driving environment capable of modeling complex traffic scenarios. The testing was performed on Town05 and Town10, which feature multilane roads, intersections, occlusions, and varying pedestrian density. The autonomous vehicle (AV) was equipped with simulated LiDAR, radar, and camera sensors operating at 10 Hz. All modules were deployed on a machine with an NVIDIA RTX 3090 GPU and 64 GB RAM, ensuring inference time was measured under realistic compute conditions.

During training, the DQN agent interacted with the environment over 500,000 simulation steps using an ε-greedy exploration strategy. The replay buffer size was set to 100,000 transitions, the target network was updated every 500 steps, and the batch size was fixed at 64. A discount factor γ=0.99 was applied to balance short-term and long-term rewards. To evaluate our system, we compared the proposed Fusion-DQN model with three baseline methods: a rule-based planner using finite state machines, an image-only DQN using RGB frames as input, and a LiDAR-only DQN using BEV occupancy grids. All models were tested on identical scenarios under both normal and adverse conditions (e.g., rain, reduced lighting).

Quantitative results are summarized in Table I. The Fusion-DQN model achieved the highest task completion rate (89.4%) and lowest collision rate (6.5%), outperforming the image-only and LiDAR-only agents by 12.6% and 9.3%, respectively. Furthermore, the proposed model demonstrated superior path fidelity, with an average displacement error (ADE) of 0.63 meters and a jerk value of 0.91 m/s³, indicating smoother and more comfortable trajectory planning. Although the computation time per decision step increased slightly to 26.7 ms due to the complexity of the fusion process, it remained within the acceptable range for real-time control (≤33 ms at 30 Hz).

Method	Success Rate (%)	Collision Rate (%)	ADE (m)	Jerk (m/s³)	Time per Step (ms)
Rule-Based Planner	71.3	22.4	0.93	1.72	8.6
Image-Only DQN	76.8	17.9	0.88	1.55	21.2
LiDAR-Only DQN	80.1	13.7	0.76	1.29	23.4
Proposed Fusion-DQN	89.4	6.5	0.63	0.91	26.7

In addition to the quantitative metrics, we performed qualitative analysis across several complex driving scenarios. Figure 2 illustrates the comparative behavior of the proposed model versus baselines. In one instance, the image-only DQN agent failed to detect a pedestrian occluded by a parked vehicle and initiated acceleration, resulting in a collision. In contrast, the Fusion-DQN agent correctly identified the pedestrian by leveraging radar motion signals and LiDAR depth cues, initiating a safe deceleration maneuver. In another test involving an unprotected left turn, the rule-based planner remained indecisive due to lack of contextual inference, whereas Fusion-DQN completed the turn confidently after modeling the intent and trajectories of oncoming vehicles.

To further assess generalization, we tested all models on Town03, an unseen environment with dynamic weather and alternate road topology. The proposed Fusion-DQN model retained strong performance with an 83.1% success rate and a 9.2% collision rate, while the image-only and LiDAR-only models degraded to 60.2% and 68.5%, respectively. This suggests that the inclusion of multi-modal sensory inputs during training enhances policy robustness under distributional shifts, a crucial factor for real-world deployment of autonomous vehicles.

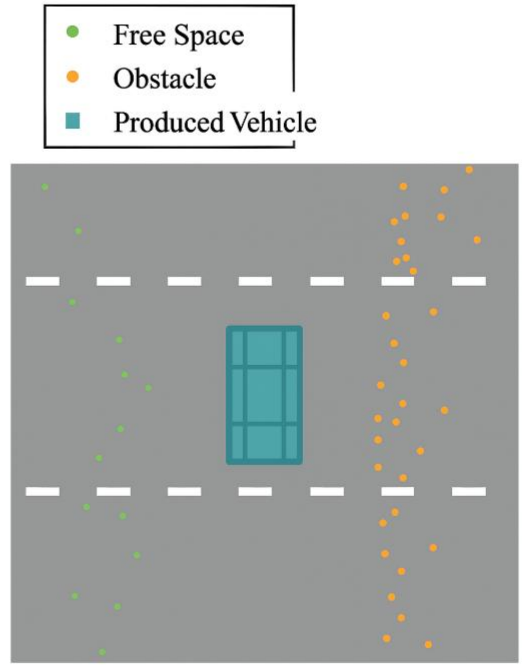

Figure 1. Qualitative comparison of decision-making behavior across models.

5.Conclusion and Future Work

In this paper, we proposed a novel decision-making framework for autonomous vehicles that integrates multi-sensor fusion with deep reinforcement learning. By fusing LiDAR, radar, and camera data into a unified bird' s-eye-view representation and feeding it into a DQN-based policy module, the system is able to perceive complex urban environments and generate high-level actions that prioritize safety, efficiency, and passenger comfort. Through extensive experiments in the CARLA simulator, we demonstrated that the proposed Fusion-DQN model significantly outperforms both traditional rule-based planners and unimodal DRL baselines. The agent achieved higher success rates, reduced collision frequency, and exhibited smoother control behavior, while maintaining real-time inference capability.

Moreover, qualitative results showed that the system is capable of handling occlusions, dynamic interactions, and uncertain traffic situations more reliably than models relying on single-sensor inputs. The ability to generalize to unseen environments and adverse conditions further indicates the robustness and adaptability of the proposed method, making it a promising candidate for real-world autonomous driving deployments.

In future work, we aim to extend this framework in several directions. First, we plan to incorporate temporal modeling using recurrent neural networks or transformer-based encoders to improve decision stability over time. Second, we intend to explore online adaptation mechanisms that allow the agent to fine-tune its policy during deployment based on real-time feedback. Third, we will integrate vehicle-to-everything (V2X) communication data into the fusion pipeline, enabling the agent to reason about traffic signals and nearby human-driven vehicles more proactively. Lastly, we are working on deploying the system on a physical AV platform to evaluate real-world performance under hardware constraints and sensor noise, thereby bridging the gap between simulation and reality.

The results presented in this study suggest that deep sensor fusion combined with reinforcement learning constitutes a viable and effective approach for the next generation of autonomous driving decision systems.

References

[1] Sun, X., Duan, Y., Deng, Y., Guo, F., Cai, G., & Peng, Y. (2025, March). Dynamic operating system scheduling using double DQN: A reinforcement learning approach to task optimization. In 2025 8th International Conference on Advanced Algorithms and Control Engineering (ICAACE) (pp. 1492-1497). IEEE.

[2] Liu, J., Gu, X., Feng, H., Yang, Z., Bao, Q., & Xu, Z. (2025, March). Market Turbulence Prediction and Risk Control with Improved A3C Reinforcement Learning. In 2025 8th International Conference on Advanced Algorithms and Control Engineering (ICAACE) (pp. 2634-2638). IEEE.

[3] Zhu, L., Guo, F., Cai, G., & Ma, Y. (2025). Structured preference modeling for reinforcement learning-based fine-tuning of large models. Journal of Computer Technology and Software, 4(4).

[4] Peng, Y., Wang, Y., Fang, Z., Zhu, L., Deng, Y., & Duan, Y. (2025, March). Revisiting LoRA: A Smarter Low-Rank Approach for Efficient Model Adaptation. In 2025 5th International Conference on Artificial Intelligence and Industrial Technology Applications (AIITA) (pp. 1248-1252). IEEE.

[5] Lou, Y., Liu, J., Sheng, Y., Wang, J., Zhang, Y., & Ren, Y. (2025, March). Addressing Class Imbalance with Probabilistic Graphical Models and Variational Inference. In 2025 5th International Conference on Artificial Intelligence and Industrial Technology Applications (AIITA) (pp. 1238-1242). IEEE.

[6] Xu, T., Xiang, Y., Du, J., & Zhang, H. (2025). Cross-Scale Attention and Multi-Layer Feature Fusion YOLOv8 for Skin Disease Target Detection in Medical Images. Journal of Computer Technology and Software, 4(2).

[7] Wu, Y., Lin, Y., Xu, T., Meng, X., Liu, H., & Kang, T. (2025). Multi-Scale Feature Integration and Spatial Attention for Accurate Lesion Segmentation.

[8] Zhan, J. (2024). MobileNet Compression and Edge Computing Strategy for Low-Latency Monitoring. Journal of Computer Science and Software Applications, 4(4).

[9] Wang, R., Luo, Y., Li, X., Zhang, Z., Hu, J., & Liu, W. (2025, January). A Hybrid Recommendation Approach Integrating Matrix Decomposition and Deep Neural Networks for Enhanced Accuracy and Generalization. In 2025 5th International Conference on Neural Networks, Information and Communication Engineering (NNICE) (pp. 1778-1782). IEEE.

[10] Han, X., Sun, Y., Huang, W., Zheng, H., & Du, J. (2025). Towards Robust Few-Shot Text Classification Using Transformer Architectures and Dual Loss Strategies. arXiv preprint arXiv:2505.06145.

[11] Bao, Q., Wang, J., Gong, H., Zhang, Y., Guo, X., & Feng, H. (2025, March). A Deep Learning Approach to Anomaly Detection in High-Frequency Trading Data. In 2025 4th International Symposium on Computer Applications and Information Technology (ISCAIT) (pp. 287-291). IEEE.

[12] Dai, L., Zhu, W., Quan, X., Meng, R., Chai, S., & Wang, Y. (2025). Deep Probabilistic Modeling of User Behavior for Anomaly Detection via Mixture Density Networks. arXiv preprint arXiv:2505.08220.

[13] Sha, Q., Tang, T., Du, X., Liu, J., Wang, Y., & Sheng, Y. (2025). Detecting Credit Card Fraud via Heterogeneous Graph Neural Networks with Graph Attention. arXiv preprint arXiv:2504.08183.

[14] Liu, X., Xu, Q., Ma, K., Qin, Y., & Xu, Z. (2025). Temporal Graph Representation Learning for Evolving User Behavior in Transactional Networks.

[15] Jiang, N., Zhu, W., Han, X., Huang, W., & Sun, Y. (2025). Joint Graph Convolution and Sequential Modeling for Scalable Network Traffic Estimation. arXiv preprint arXiv:2505.07674.