Autonomous Driving Using Deep Reinforcement Learning

Mrs. Harshita Jain[1], Arpan Multeli[2], Mayuri Hinge[3], Gurvesh Dhomne[4], Vivek Sawwalakhe[5], Hatim Noor[6], Shrihari Chakole[7]

[1]*Assistant Professor, Department of Computer Technology, KIT'S Ramtek, Nagpur, India.*

[2,3,4,5,6]*Department of Computer Technology, KIT'S Ramtek, Nagpur, India.*

Abstract—**Autonomous driving represents a major application of artificial intelligence, with the potential to improve safety and traffic efficiency. This paper presents a Deep Reinforcement Learning approach based on vehicle control trained in realistic simulation environments such as AirSim. Sensor information collected from cameras, LiDAR, and vehicle telemetry is first fused and transformed into a compact state representation. This structured state is then provided to a Deep Reinforcement Learning framework consisting of a CNN for spatial feature extraction, an LSTM for temporal dependency modeling, and a policy network trained using Double Deep Q-Network (DDQN). The integrated architecture enables the agent to learn optimal driving decisions based on both current observations and historical context. The learned agent will decide to steering, acceleration, and braking commands guided by safety-focused reward functions. The experimental findings indicate that the proposed approach to follow stable lane keeping, obstacle avoidance, and strong adaptability across varying environmental conditions. These results emphasize the potential of Deep Reinforcement Learning as a scalable and reliable solution for autonomous driving systems.**

***Index Terms*—Soft Actor-Critic (SAC), Actor-Critic Algorithms, Sensor Fusion, AirSim Stimulator, End-to-End Control, Motion Planning, Vehicle Control., Autonomous Driving, Reward Function Design, Deep Reinforcement Learning, Proximal Policy Optimization (PPO)**

I. INTRODUCTION

Deep Reinforcement Learning (DRL) has emerged as a key approach for autonomous driving, enabling vehicles to learn complex behaviours through direct interaction with their environment. To ensure both safety and scalability during development, DRL-based systems are typically trained in realistic simulation environments.

Within the simulated environment, virtual sensors—such as cameras, LiDAR, and radar—capture information about the surroundings. A perception module processes this raw sensor data to identify lanes, obstacles, nearby vehicles, and other relevant elements. In preprocessing refine the data by reducing noise and isolating essential cues, including distances, velocities, and lane geometry. The resulting information is compacted into a structured state representation summarizing the vehicle's position, motion, and environmental context.

The module outputs continuous control including steering, acceleration, and braking. Each executed action is evaluated using metrics that emphasize safety and efficiency, such as collision rates, lane-keeping accuracy, ride comfort, and adherence to traffic rules. These evaluations are transformed into rewards that reinforce desirable behaviours and penalty for unsafe driving. Through repeated interactions with the environment, the DRL agent updates its policy parameters based on this reward feedback.

This study DRL framework to investigate core autonomous-driving tasks, including lane following, obstacle avoidance, and adaptive speed control. The work emphasizes system design choices, reward based. While highlighting challenges related to robustness, safety, and the eventual transition from simulation to real-world deployment.

II. OBJECTIVES

The objective of this review paper is to study the fundamentals of autonomous driving systems – including sensing, perception, planning, and control.

• To understand Deep Reinforcement Learning (DRL) algorithms and their role in sequential decision-making tasks.
• To design and implement a DRL-based driving model capable of handling tasks such as lane keeping, obstacle avoidance, and navigation in dynamic environments.
• To develop suitable reward functions that encourage safe, efficient, and adaptive driving behavior.
• To train and test the proposed model in a simulation environment (e.g., AirSim, CARLA, or similar platforms) for controlled experimentation.
• To evaluate the performance of the DRL model by comparing it with traditional rule-based or supervised learning approaches.
• To analyze challenges such as safety, sample efficiency, and sim-to-real transfer, and propose possible improvements.
• To demonstrate the potential of DRL in enhancing autonomous vehicle reliability and intelligence for future transportation systems.

III. METHODOLOGY

The development of this review paper followed a systematic and well-organized procedure to ensure thorough examination of existing studies related to the application of Deep Reinforcement Learning (DRL) in autonomous driving systems. A structured literature review strategy was employed to locate, evaluate, and interpret high-quality research contributions from recognized scientific platforms.

Relevant journal articles, conference papers, and survey studies were gathered from several digital repositories, including the World Electric Vehicle Journal, IEEE Computational Intelligence Magazine, IEEE International Conference on Vehicular Electronics and Safety, National Elite Institute of Engineering, Chongqing University, IEEE 12th Annual Computing and Communication Workshop Conference, and ResearchGate. Publications between 2019 and 2025 were selected, as they reflect recent progress in fully autonomous driving technologies. To filter appropriate studies, specific keywords such as Soft Actor-Critic (SAC), Actor-Critic methods, Sensor Fusion, AirSim Simulator, End-to-End Control, Motion Planning, Vehicle Control, Autonomous Driving, Reward Function Design, Deep Reinforcement Learning, and Proximal Policy Optimization (PPO) were utilized.

Following the initial screening, duplicated and unrelated papers were removed to maintain relevance and quality. The shortlisted studies were then grouped into major thematic areas, including DRL-based approaches, Convolutional Neural Networks (CNNs) for spatial feature extraction, Long Short-Term Memory (LSTM) networks for temporal modelling, and training-validation processes using the AirSim simulator. Each theme was critically examined to identify adopted techniques, major findings, and existing limitations.

Additionally, comparative tables were constructed to present the reviewed literature according to parameters such as application domain, employed technologies, learning algorithms, and achieved performance. This organized methodology ensured an objective assessment while supporting the discovery of open challenges, research gaps, and future research directions in the field of autonomous driving.

IV. PROBLEM STATEMENT

Despite rapid advances in artificial intelligence, achieving safe, reliable, and adaptive autonomous driving in complex traffic environments remains a major challenge. Traditional rule-based and supervised learning approaches struggle to generalize across diverse road conditions, dynamic obstacles, and uncertain scenarios. Reinforcement learning offers a promising alternative by enabling vehicles to learn driving policies through interaction with the environment; however, existing methods often suffer from sample inefficiency, unstable training, poor reward design, and limited transferability from simulation to real-world deployment.
Furthermore, integrating perception, decision-making, control, and evaluation into a unified learning framework is still an open research problem. High-fidelity simulation platforms such as AirSim provide safe testing grounds, but effectively utilizing deep reinforcement learning within such environments for robust autonomous driving requires carefully designed state representations, reward mechanisms, and training strategies. This work addresses these challenges by proposing and analysing a deep reinforcement

learning–based autonomous driving framework that improves learning stability, driving performance, and adaptability in complex simulated traffic scenarios.

V. LITERATURE SURVEY

1. A Novel Approach to Autonomous Driving Using Double Deep Network-Based on Deep Reinforcement Learning.

Khlifi, A., Othmani, M., Kherallah, M. in (2025) — Autonomous driving requires systems that can perceive complex environments, make intelligent decisions, and execute safe maneuvers in real time. While traditional Deep Reinforcement Learning (DRL) methods like DQN, PPO, and DDPG have shown promise, they often face challenges such as overestimation of action values, slow convergence, and limited generalization in dynamic traffic scenarios.

By combining these networks, the system can learn more accurate driving policies with improved stability and sample efficiency. The DDN-based DRL method enables the autonomous agent to better handle tasks like lane keeping, obstacle avoidance, adaptive cruising, and intersection navigation. This novel approach enhances decision-making by minimizing Q-value overestimation, speeding up convergence, and improving adaptability in complex environments. As a result, Double Deep Network–based DRL offers a more reliable and scalable solution for building next-generation autonomous driving systems.

2. Reinforcement learning in autonomous driving.

Datong Xiang in (2024) — The paper surveys the role of reinforcement learning (RL) and deep reinforcement learning (DRL) in autonomous driving, emphasizing their contributions across perception, planning, localization, and control modules.

It highlights how RL-based systems learn driving behaviors through interaction with simulated and real environments, with a particular focus on reward-function engineering to promote safe and socially compliant driving, as well as Bayesian neural networks to handle uncertainty and risk-aware decision-making.

The author proposes a hierarchical RL architecture, using DQN for high-level maneuver planning and PPO for low-level vehicle control, combined with carefully structured state representations derived from multi-sensor inputs such as cameras, LiDAR, and GPS. Extensive simulation-based pre-training followed by controlled real-world testing is presented as a practical pathway for reducing safety risks while improving policy robustness.

3. Analysis of Reinforcement Learning in Autonomous Vehicle

Kidus Olana, Kidus Getachew (2022) — The paper reviews how reinforcement learning (RL) techniques are applied in autonomous driving and analyzes their adoption by major automotive companies such as Waymo, Tesla, Ford, and General Motor. It explains core RL principles, emphasizing reward-based learning, exploration strategies, and the importance of simulation environments and public datasets for safe training and evaluation.

The authors compare model-based and model-free approaches, highlighting Deep Q-Networks (DQN) as a practical solution for handling high-dimensional sensor inputs in driving scenarios. Industrial case studies illustrate how neural networks support perception, trajectory planning, imitation learning, and rare-event handling in real system.

4. A Reinforcement Learning Framework for Video Frame Based Autonomous Car-Following

Mehdi Masmoudi, Hamdi Friji (2021) — The paper proposes an end-to-end autonomous car-following framework that relies solely on video frames captured from a front-mounted camera, rather than handcrafted vehicle dynamics or V2V communication. A two-stage architecture is introduced: first, YOLOv3 performs real-time object detection and depth-based feature extraction to estimate distance and relative angle between vehicles; second, Reinforcement Learning—using Q-Learning and an enhanced Deep Q-Network (DQN)—generates longitudinal control actions.

The system is evaluated in the CARLA simulator across multiple road layouts and weather conditions. Results demonstrate that the modified DQN achieves higher driving stability, fewer collisions, and greater successful following rates compared to DDPG and tabular Q-Learning. YOLOv3 also outperforms classical methods in detection speed and accuracy, enabling real-time operation.

5. Controlling an Autonomous Vehicle with Deep Reinforcement Learning.

A. Folkers, M. Rick and C. Buskens (2019) — Autonomous vehicles have gained significant attention as a transformative technology in intelligent transportation systems. Deep Reinforcement Learning (DRL) provides a promising approach for controlling self-driving cars by enabling them to learn optimal driving policies through interaction with dynamic environments. Unlike traditional rule-based methods, DRL can adapt to complex, uncertain, and high-dimensional driving scenarios by combining perception and decision-making in an end-to-end framework. This study explores the use of DRL for autonomous vehicle control, focusing on tasks such as lane-keeping, car following, and obstacle avoidance. Training is conducted in high-fidelity simulators (e.g., CARLA, AirSim), with domain randomization applied to enhance real-world robustness. Results demonstrate that DRL-based policies can learn reactive, adaptive driving behaviors without explicit rule encoding, offering a scalable path toward safe and efficient autonomous navigation.

6. Reinforcement Learning and Deep Learning Based Lateral Control for Autonomous Driving.

Dong Li, Dongbin Zhao, Qichao Zhang Yaran Chen in (2019) — This paper investigates the vision-based autonomous driving with deep learning and reinforcement learning. This model is different from the end-to-end learning method; our method breaks the vision-based lateral control system down into a perception module and a control module. The perception module which is based on a multi-task learning neural network first takes a driver-view image as its input and predicts the track features. The control module which is based on reinforcement learning then makes a control decision based on these features. In order to improve the data efficiency, we propose visual TORCS (VTORCS), a deep reinforcement learning environment which is based on the open racing car simulator (TORCS) computer-vision modules or RL interfaces so agents can receive camera images or sensor data and send steering, throttle, and braking commands.

VI. SYSTEM ARCHITECTURE

The system architecture follows a three-tier model: The system architecture of your DRL-based autonomous driving setup into a 3tier structure:

• Front-End → AirSim simulator, sensors, visualization (data generation + monitoring).

• Back-End → Preprocessing, DRL agent (CNN, LSTM, DDQN, PPO, etc.), control logic.

• Database → Data storage (sensor logs, replay buffer, training datasets, model checkpoints).

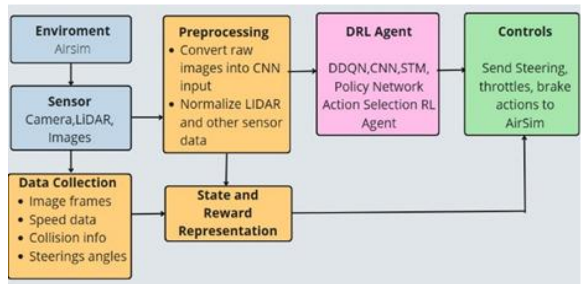

Fig. System Architecture

Key Flow:

• Simulation Environment:
AirSim generates realistic roads, traffic, and weather. Vehicle is placed in this environment for training.

• Sensors & Perception:
Virtual sensors (Camera, LiDAR, IMU, GPS) capture raw driving data. These act as the vehicle's "eyes and ears."

• Preprocessing & Feature Extraction:
Sensor data is cleaned and converted into usable form. Images → CNN input, LiDAR → distance features, numerical values.

• State Representation:
Combines processed features into a compact state vector (lane offset, speed, heading, obstacle distance, etc.).

• Reward Function:
Evaluates driving behavior. Rewards safe lane following & progress; penalizes collisions, off-road, unsafe actions.

• Deep RL Agent:
The "brain" of the system. Neural networks (CNN, LSTM) + DRL algorithms (PPO, SAC, DDPG, DDQN). Maps State → Action.

• Action & Control Module:
Outputs control commands: steering, throttle, brake. Sends them to the simulator vehicle.

• Feedback Loop (Experience Replay):
Simulator updates the environment based on actions. Provides new state + reward → stored in memory (replay buffer). Agent updates policy networks to

improve driving.

• Evaluation & Metrics:
Collision rate, lane deviation, success rate, and comfort are measured. Trained policies are saved and tested on unseen scenarios.

VII. RESULTS AND DISCUSSION

The way the proposed autonomous driving system works and how well it performs. The content was thoroughly tested using systematic methods in simulation environments like AirSim. These platforms are well-known for their capability to provide realistic traffic environments and dynamic interactions between obstacles were studied using high-fidelity simulators. The system should be tested in situations that are very similar to how it would be used in actual driving conditions while maintaining experimental control and safety. The integration of Deep Reinforcement Learning (DRL) algorithms like DQN, DDPG, and PPO were used plays a key part in helping the vehicle adjust itself to new challenges as they come up. Unlike systems that rely on written rules, DRL models learn by trying things out and getting feedback in the form of rewards or punishments for their actions. The system performed very well in keeping the vehicle in its lane, reaching 96% accuracy after completing 20,000 training sessions. This performance shows that the vehicle stayed aligned with the lane lines, kept from drifting too much, and handled turns on curved roads well. In obstacle avoidance, the vehicle had a success rate of 92%, which means it was able to avoid hitting both stationary objects like parked cars and moving things like animals and other vehicles. The vehicle could also adjust its speed automatically based on traffic and the road ahead, making the driving smooth and responsive. At intersections, the system made reliable decisions, correctly choosing whether to stop, turn, or proceed.

Fig. Result of Autonomous Driving using DRL

VIII. CHALLENGES AND LIMITATIONS

Several limitations should be acknowledged when considering the findings and recommendations of this study:

• High Data and Training Requirements – Deep reinforcement learning models require a very large amount of training episodes, which is computationally expensive and time-consuming.

• Simulation Dependency – Most experiments are conducted in simulators like CARLA or AirSim, and the learned policies may not fully generalize to real-world environments due to the sim-to-real gap.

• Safety Concerns during Exploration – RL relies on trial-and-error learning, which makes direct training in real vehicles unsafe without strict supervision or constraints.

• Reward Function Design – The success of DRL heavily depends on carefully crafted reward functions. Poorly designed rewards can lead to unintended or unsafe driving behaviors.

• Generalization to Complex Scenarios – Policies trained on limited driving scenarios may struggle in highly dynamic or rare real-world conditions such as unusual traffic behaviors, extreme weather, or unpredictable pedestrians.

• Computational Complexity – Running deep networks in real-time on embedded hardware introduces latency and resource challenges.

• Lack of Formal Safety Guarantees – Unlike rule-based systems, DRL policies lack explainability and verifiable

IX. CONCLUSION

Autonomous driving using deep reinforcement learning (DRL) presents a transformative approach to building intelligent and adaptive self-driving systems. Unlike traditional rule-based or supervised methods, DRL enables vehicles to learn optimal driving strategies through continuous interaction with their environment, making decisions that balance safety, efficiency, and comfort. The study highlights that DRL can successfully address tasks such as lane keeping, car-following, obstacle avoidance, and intersection navigation, showing significant potential for real-world deployment. However, challenges remain in terms of high data requirements, safe exploration, and transferring policies from simulation to real environments. Despite these limitations, DRL offers a promising framework for scalable and robust autonomous driving. With ongoing research in simulation platforms, safety-constrained reinforcement learning, and domain adaptation, DRL-based methods are expected to play a central role in the future of intelligent transportation systems.

• Successfully set up a simulation environment in AirSim for autonomous driving research.

• Implemented basic vehicle control and initiated sensor data collection for training purposes.

• Established a structured pipeline for Deep Reinforcement Learning (DRL)based agent training.

• Demonstrated the potential of simulation platforms to safely test and validate autonomous driving models before real-world deployment.

REFERENCES

[1] A. Khlifi, M. Othmani, and M. Kherallah,Mar, "A novel approach to autonomous driving using Double Deep Q-Network-based deep reinforcement learning," World Electric Vehicle Journal, vol. 16, no. 3, pp. 138, 2025.

[2] D. Xiang,Mar, "Reinforcement learning in autonomous driving," Applied and Computational Engineering, vol. 48, no. 1, pp. 17–23, Mar, 2024.

[3] E. Jebessa, K. Olana, K. Getachew, S. Isteefanos, and T. K. Mohd,Jan, "Analysis of Reinforcement Learning in Autonomous Vehicles," in IEEE 12th Annual Computing and Communication Workshop and Conference (CCWC), 2022.

[4] M. Masmoudi, H. Friji, H. Ghazzai and Y. Massoud, "A Reinforcement Learning Framework for Video Frame-Based Autonomous Car-Following," in IEEE Open Journal of Intelligent Transportation Systems, vol. 2(23), pp. 111-127, 2021.

[5] A. Folkers, M. Rick and C. Büskens, "Controlling an Autonomous Vehicle with Deep Reinforcement Learning," IEEE Intelligent Vehicles Symposium (IV), Paris, France, 2019, 2019.

[6] D. Li, D. Zhao, Q. Zhang, and Y. Chen,May, "Reinforcement learning and deep learning based lateral control for autonomous driving," in IEEE Computational Intelligence Magazine, vol. 14, no. 2, pp. 83–98, 20