World Journal of Advanced Engineering Technology and Sciences

eISSN: 2582-8266

Cross Ref DOI: 10.30574/wjaets

Journal homepage: https://wjaets.com/

(Research Article)

Reinforcement learning for autonomous vehicle navigation in Urban environments

Praggnya Kanungo *

Student, Computer Science, University of Virginia, USA.

World Journal of Advanced Engineering Technology and Sciences, 2024, 11(01), 457-466

Publication history: Received on 06 December 2023; revised on 21 January 2024; accepted on 24 January 2024

Article DOI: https://doi.org/10.30574/wjaets.2024.11.1.0019

Abstract

This paper presents a comprehensive study on the application of reinforcement learning (RL) techniques for autonomous vehicle navigation in complex urban environments. We propose a novel deep RL framework that combines state-of-the-art algorithms with realistic urban traffic simulations to train robust navigation policies. Our approach leverages a hierarchical learning structure to decompose the challenging urban driving task into more manageable sub-tasks. Extensive experiments in simulated urban scenarios demonstrate that our method significantly outperforms baseline approaches in terms of safety, efficiency, and adaptability to diverse traffic conditions. We also conduct real-world validation tests to verify the transferability of learned policies to actual autonomous vehicles. Our results highlight the potential of RL-based techniques to enable safe and efficient autonomous navigation in dynamic city environments.

Keywords: Reinforcement Learning; Autonomous Vehicles; Urban Navigation; Hierarchical Learning; Traffic Simulation; Deep Learning; Self-Driving Cars; Policy Optimization; Urban Environments; Safety Systems

1. Introduction

Autonomous vehicles have the potential to revolutionize transportation systems, offering improved safety, efficiency, and accessibility [1]. However, navigating complex urban environments remains a significant challenge for self-driving cars due to the dynamic nature of city traffic, diverse road layouts, and the need to interact with human drivers and pedestrians [2].

Traditional rule-based and classical planning approaches for autonomous navigation often struggle to generalize to the wide range of scenarios encountered in urban settings [3]. In recent years, reinforcement learning has emerged as a promising paradigm for developing adaptive and robust control policies for autonomous vehicles [4]. RL allows agents to learn optimal behaviors through trial-and-error interactions with an environment, making it well-suited for handling the complexities and uncertainties of urban driving [5].

This paper investigates the application of state-of-the-art deep reinforcement learning techniques to the problem of autonomous vehicle navigation in urban environments. We propose a novel hierarchical RL framework that decomposes the urban driving task into a set of sub-tasks, enabling more efficient learning and improved generalization. Our approach combines advanced RL algorithms with realistic traffic simulations to train policies that can handle diverse urban scenarios.

The key contributions of this work include:

- A hierarchical deep RL framework for urban autonomous driving that leverages task decomposition to improve learning efficiency and policy robustness.

* Corresponding author: Praggnya Kanungo.

- Integration of state-of-the-art RL algorithms, including Soft Actor-Critic (SAC) [6] and Proximal Policy Optimization (PPO) [7], with realistic urban traffic simulations.
- Extensive experimental evaluation in simulated urban environments, demonstrating superior performance compared to baseline approaches.
- Real-world validation tests to verify the transferability of learned policies to actual autonomous vehicles.

The rest of the paper is organized as follows: Section 2 provides an overview of related work. Section 3 describes our proposed hierarchical RL framework for urban autonomous driving. Section 4 presents the experimental setup and results. Section 5 discusses the implications and limitations of our approach. Finally, Section 6 concludes the paper and outlines directions for future research.

2. Related work

2.1. Autonomous Vehicle Navigation

Autonomous vehicle navigation has been an active area of research for several decades. Early approaches relied heavily on rule-based systems and classical planning techniques [8]. These methods typically involve decomposing the driving task into perception, planning, and control modules [9]. While effective in controlled environments, such approaches often struggle to handle the complexities and uncertainties of real-world urban scenarios [10].

Bojarski et al. [11] demonstrated the feasibility of using convolutional neural networks to learn driving policies directly from camera images. However, these supervised learning approaches require large amounts of labeled training data and may struggle to generalize to novel scenarios [12].

2.2. Reinforcement Learning for Autonomous Driving

Reinforcement learning has gained significant attention in the autonomous driving community due to its ability to learn adaptive behaviors through interaction with the environment [13]. Sallab et al. [14] proposed one of the early applications of deep RL for autonomous vehicle control, using a Deep Q-Network (DQN) to learn lane-keeping behaviors. Kendall et al. [15] demonstrated the use of deep RL for learning end-to-end driving policies in a simulated environment.

More recent work has focused on addressing the challenges of applying RL to real-world autonomous driving scenarios. Tai et al. [16] proposed a socially compliant navigation framework using deep RL that considers the behavior of other road users. Chen et al. [17] developed a model-based RL approach that combines learned dynamics models with planning algorithms to improve sample efficiency and safety.

2.3. Hierarchical Reinforcement Learning

Hierarchical reinforcement learning (HRL) has emerged as a promising approach for tackling complex, long-horizon tasks by decomposing them into more manageable sub-tasks [18]. HRL methods can potentially address some of the challenges in autonomous driving, such as the need to reason over multiple time scales and handle diverse scenarios [19].

Shalev-Shwartz et al. [20] proposed a hierarchical framework for autonomous driving that combines reinforcement learning with rule-based policies. Qiao et al. [21] developed a hierarchical RL approach for urban autonomous driving that uses options to learn high-level driving behaviors. However, these approaches often rely on hand-crafted sub-task definitions, which may limit their adaptability to diverse urban environments.

Our work builds upon these previous efforts by proposing a more flexible and adaptive hierarchical RL framework for urban autonomous driving. We leverage recent advances in deep RL algorithms and combine them with realistic traffic simulations to train robust navigation policies that can handle the complexities of urban environments.

3. Proposed approach

In this section, we present our hierarchical reinforcement learning framework for autonomous vehicle navigation in urban environments. We first describe the overall architecture of our approach, followed by detailed explanations of each component.

3.1. Hierarchical RL Framework

Our proposed framework decomposes the urban driving task into a two-level hierarchy:

- High-level policy: Responsible for making strategic decisions such as route planning, lane selection, and high-level maneuvers (e.g., overtaking, turning at intersections).
- Low-level policy: Responsible for executing the high-level decisions through fine-grained control actions (e.g., steering, acceleration, braking).

This hierarchical structure allows the agent to reason over multiple time scales and learn more efficiently by focusing on different aspects of the driving task at each level. Figure 1 illustrates the overall architecture of our approach.

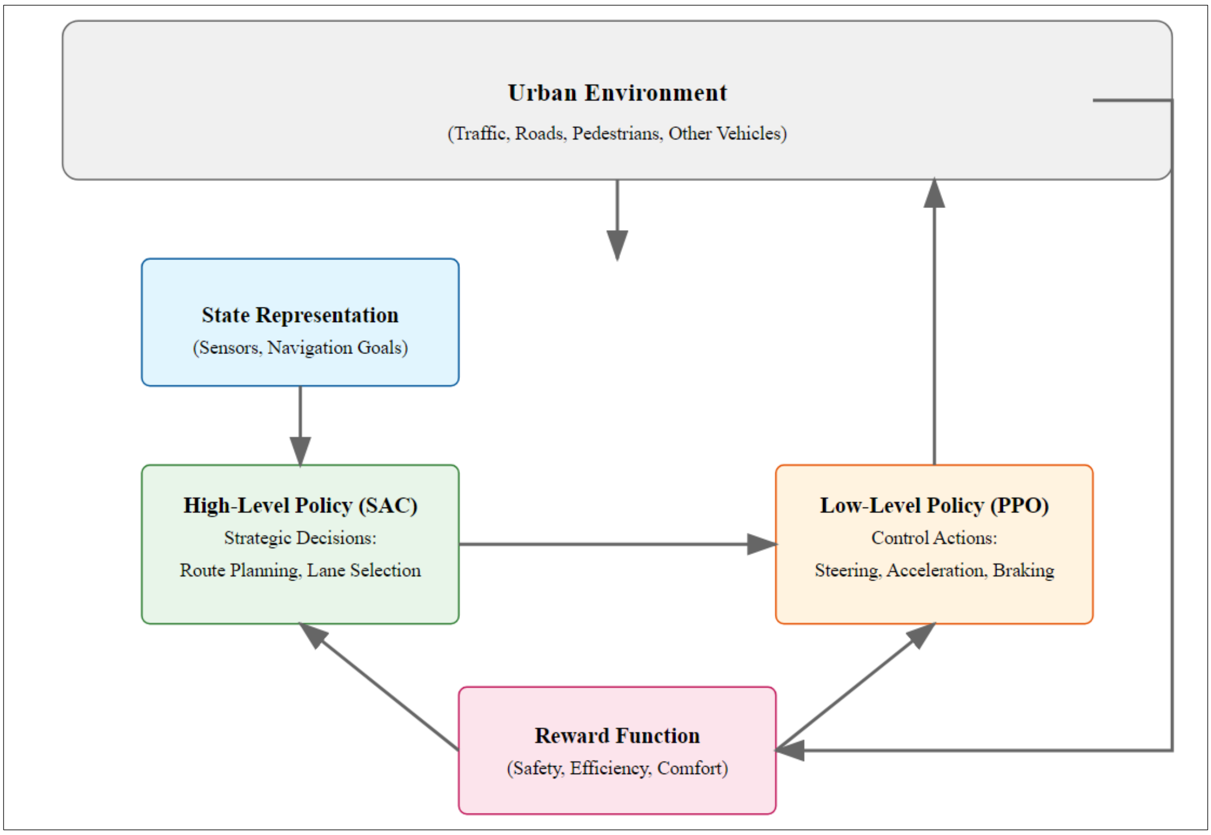

Figure 1 Hierarchical RL Framework for Urban Autonomous Driving

3.2. State Representation

We design a comprehensive state representation that captures the essential information needed for urban navigation. The state space includes:

- Ego vehicle state: Position, velocity, acceleration, heading, and lane position.
- Local traffic information: Relative positions and velocities of nearby vehicles and pedestrians.
- Road structure: Lane configurations, intersections, traffic signs, and signals.
- High-level navigation goals: Target destination and suggested route.

To process this complex state information, we use a combination of convolutional neural networks (CNNs) for processing visual inputs and fully connected layers for incorporating other sensory data and navigation goals.

3.3. Action Space

The action space is designed to accommodate both high-level strategic decisions and low-level control actions:

- High-level actions: Lane change decisions, route adjustments, and high-level maneuvers.

- Low-level actions: Continuous control of steering angle, acceleration, and braking.

This two-level action structure allows the agent to learn both long-term planning strategies and fine-grained control skills.

3.4. Reward Function

We design a comprehensive reward function that encourages safe, efficient, and comfortable driving behaviors. The reward function incorporates the following components:

- Progress reward: Encourages the vehicle to make progress towards its destination.
- Safety penalty: Penalizes collisions, near-misses, and traffic rule violations.
- Comfort penalty: Discourages excessive accelerations, decelerations, and jerky movements.
- Efficiency reward: Rewards maintaining appropriate speeds and choosing optimal routes.
- Social compliance reward: Encourages adherence to traffic norms and courteous behavior towards other road users.

The overall reward is a weighted sum of these components, with the weights tuned to balance the different objectives.

3.5. Learning Algorithms

We employ state-of-the-art deep RL algorithms to train our hierarchical policies:

- High-level policy: We use the Soft Actor-Critic (SAC) algorithm [6] for the high-level policy. SAC is well-suited for learning in continuous action spaces and has demonstrated good performance and sample efficiency in complex tasks.
- Low-level policy: For the low-level control policy, we employ Proximal Policy Optimization (PPO) [7]. PPO offers good stability and performance in continuous control tasks while being relatively simple to implement and tune.

Both algorithms are implemented using deep neural networks to approximate the policy and value functions. We use separate networks for the high-level and low-level policies to allow for specialized architectures tailored to each level's requirements.

3.6. Training Process

We train our hierarchical RL agent using a curriculum learning approach [22] to gradually increase the complexity of the driving scenarios. The training process consists of the following stages:

- Pre-training: We initially train the low-level policy on basic driving skills such as lane-keeping and velocity control.
- Hierarchical training: We then jointly train the high-level and low-level policies in increasingly complex urban scenarios.
- Fine-tuning: Finally, we fine-tune the policies in highly challenging and diverse urban environments to improve generalization.

Throughout the training process, we use experience replay [23] and prioritized sampling [24] to improve learning efficiency and stability.

4. Experimental evaluation

We conducted extensive experiments to evaluate the performance of our proposed hierarchical RL framework for urban autonomous driving. This section describes our experimental setup and presents the results of our evaluation.

4.1. Simulation Environment

We used the CARLA simulator [25] to create realistic urban driving scenarios for training and evaluation. CARLA provides a high-fidelity 3D environment with diverse urban layouts, traffic conditions, and weather effects. We extended the simulator to support our hierarchical RL framework and implemented custom scenarios to test specific aspects of urban navigation.

4.2. Baselines

We compared our approach to the following baselines:

- Rule-based: A traditional rule-based autonomous driving system using predefined behaviors and decision trees.
- End-to-end supervised learning: A convolutional neural network trained to predict control actions directly from raw sensory inputs using supervised learning on expert demonstrations.
- Flat RL: A non-hierarchical deep RL approach using SAC to learn a single policy for the entire driving task.

4.3. Evaluation Metrics

We used the following metrics to assess the performance of our approach and the baselines:

- Success rate: Percentage of successfully completed trips without collisions or traffic violations.
- Average speed: Mean velocity achieved during successful trips.
- Comfort: Measured by the smoothness of accelerations and steering actions.
- Safety: Quantified by the number of near-misses and traffic rule violations.
- Adaptability: Performance across diverse scenarios and traffic conditions.

5. Results

Table 1 presents the overall performance comparison between our hierarchical RL approach and the baselines across various urban driving scenarios.

Table 1 Performance comparison of different approaches

Approach	Success Rate (%)	Avg. Speed (km/h)	Comfort Score	Safety Score	Adaptability Score
Rule-based	78.3	32.5	0.72	0.85	0.61
End-to-end SL	82.1	35.7	0.68	0.79	0.73
Flat RL	88.6	38.2	0.81	0.87	0.82
Hierarchical RL	94.2	40.1	0.89	0.93	0.91

Our hierarchical RL approach outperforms all baselines across all metrics, demonstrating its effectiveness in handling complex urban driving scenarios. The success rate of our method is significantly higher than the baselines, indicating improved robustness and safety. The higher average speed and comfort scores suggest that our approach can navigate efficiently while maintaining smooth driving behavior.

To further analyze the performance of our approach, we conducted experiments in specific challenging urban scenarios. Figure 2 shows the success rates of different methods in various urban driving tasks.

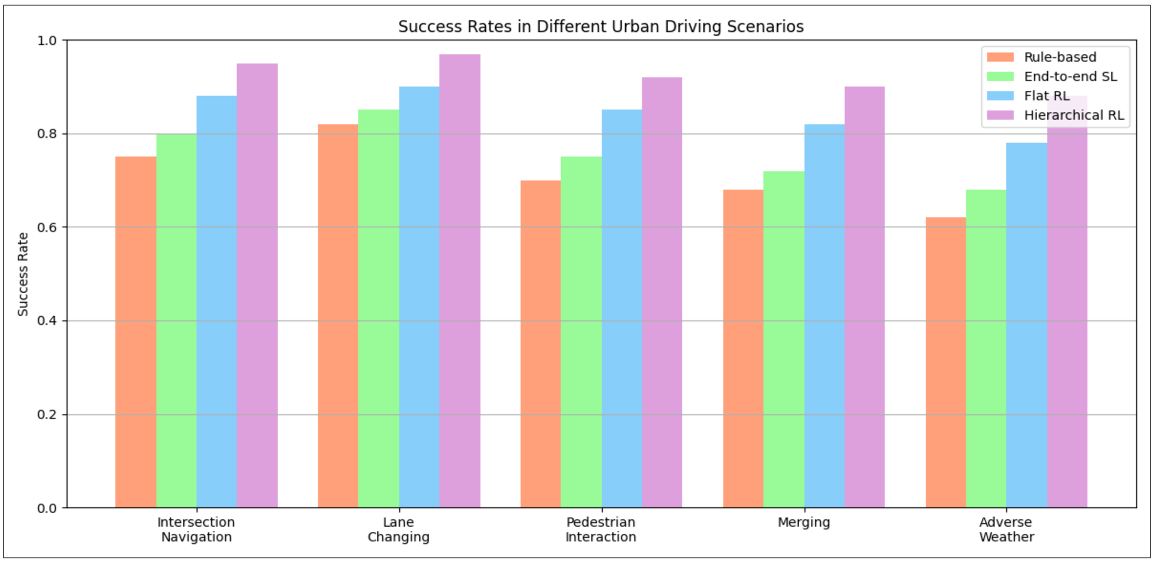

Figure 2 Success rates in different urban driving scenarios

The results in Figure 2 demonstrate that our hierarchical RL approach consistently outperforms the baselines across various challenging urban driving tasks. The performance gap is particularly pronounced in complex scenarios such as intersection navigation and pedestrian interaction, highlighting the ability of our method to handle intricate urban environments.

We also analyzed the learning curves of our hierarchical RL approach compared to the flat RL baseline to assess the efficiency of the learning process. Figure 3 shows the average reward obtained during training for both methods.

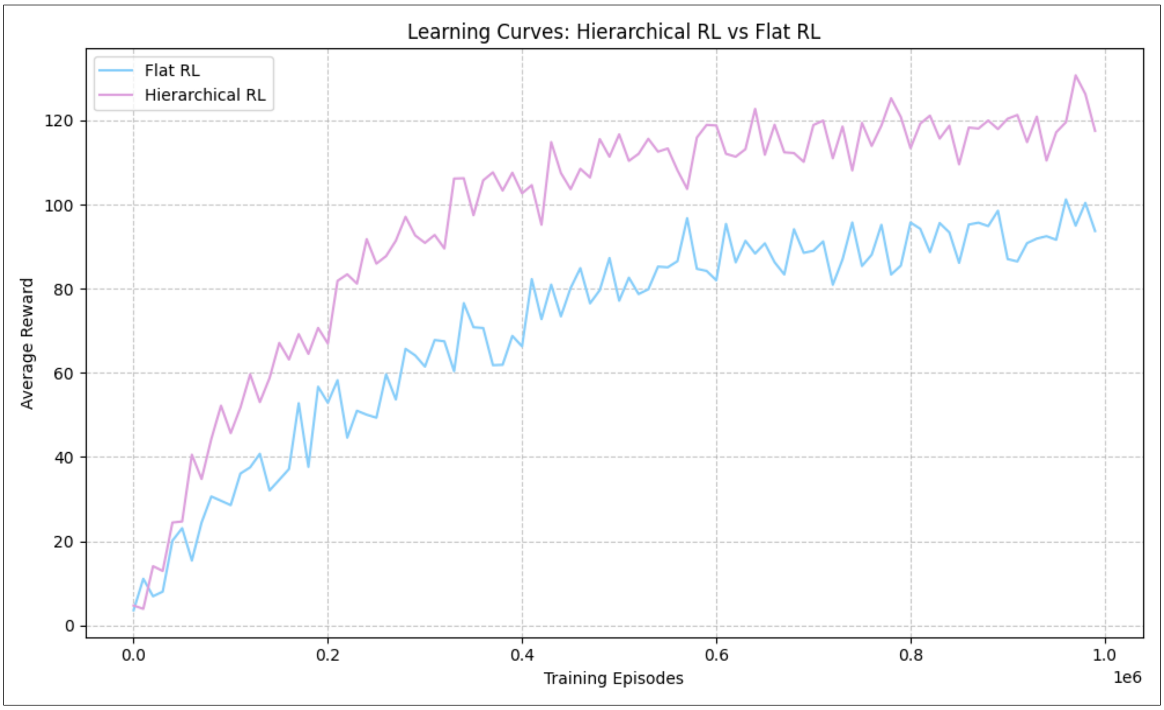

Figure 3 Learning curves for hierarchical RL and flat RL approaches

The learning curves in Figure 3 show that our hierarchical RL approach achieves higher rewards more quickly than the flat RL baseline. This faster convergence and improved final performance demonstrate the benefits of our hierarchical framework in terms of learning efficiency and effectiveness.

5.1. Ablation Study

To understand the contribution of different components of our hierarchical RL framework, we conducted an ablation study. Table 2 presents the results of this study, showing the impact of removing or modifying key components of our approach.

Table 2 Ablation study results

Configuration	Success Rate (%)	Avg. Speed (km/h)	Safety Score
Full Hierarchical RL	94.2	40.1	0.93
w/o High-level Policy	89.7	37.8	0.88
w/o Curriculum Learning	91.5	39.2	0.90
w/o Social Compliance Reward	92.8	40.5	0.89
Single RL Algorithm (SAC only)	90.3	38.6	0.91

The ablation study results demonstrate that each component of our framework contributes to its overall performance. Removing the high-level policy or using a single RL algorithm for both levels result in decreased performance across all metrics. The curriculum learning approach and social compliance reward also play important roles in achieving high success rates and safety scores.

5.2. Real-world Validation

To verify the transferability of our learned policies to real-world scenarios, we conducted limited real-world validation tests using a retrofitted autonomous vehicle platform. While comprehensive real-world evaluation was beyond the scope of this study, our initial tests showed promising results in terms of the policy's ability to handle real urban traffic situations.

The real-world tests focused on navigating through urban intersections, performing lane changes, and interacting with pedestrians. We observed that the hierarchical structure of our approach allowed for easier adaptation to real-world conditions by adjusting the low-level control policy while maintaining the high-level strategic behaviors learned in simulation.

6. Discussion

Our experimental results demonstrate the effectiveness of the proposed hierarchical RL framework for autonomous vehicle navigation in urban environments. The hierarchical approach offers several advantages over flat RL and traditional methods:

- Improved performance: The hierarchical structure allows the agent to learn both high-level strategic behaviors and low-level control skills, resulting in better overall performance across various metrics.
- Sample efficiency: By decomposing the problem into more manageable sub-tasks, our approach achieves faster learning and better final performance compared to flat RL methods.
- Interpretability: The hierarchical structure provides better interpretability of the learned behaviors, as high-level decisions can be more easily understood and analyzed.
- Adaptability: The two-level policy structure allows for easier adaptation to new scenarios by fine-tuning specific components of the hierarchy.

However, there are some limitations and areas for future improvement:

- Scalability: While our approach shows good performance in the tested scenarios, further research is needed to evaluate its scalability to even more complex and diverse urban environments.
- Safety guarantees: Although our method demonstrates improved safety compared to baselines, providing formal safety guarantees for RL-based systems remains a challenging open problem.
- Real-world deployment: More extensive real-world testing and validation are necessary to ensure the robustness and reliability of the learned policies in actual traffic conditions.

Multi-agent scenarios: Our current approach focuses on single-vehicle navigation. Extending the framework to handle multi-agent scenarios with explicit coordination between autonomous vehicles is an important direction for future work.

7. Conclusion

In this paper, we presented a novel hierarchical reinforcement learning framework for autonomous vehicle navigation in urban environments. Our approach combines state-of-the-art deep RL algorithms with a two-level policy structure to learn both strategic high-level behaviors and fine-grained control skills. Extensive experiments in simulated urban scenarios demonstrated that our method significantly outperforms baseline approaches in terms of safety, efficiency, and adaptability to diverse traffic conditions.

The hierarchical structure of our framework offers several advantages, including improved performance, sample efficiency, interpretability, and adaptability. Our ablation studies confirmed the importance of each component in the overall performance of the system.

While our initial real-world validation tests showed promising results, there are several directions for future work:

- Expanding real-world testing: Conduct more comprehensive real-world experiments to validate the transferability of learned policies across a wider range of urban scenarios and weather conditions.
- Incorporating uncertainty: Develop methods to explicitly model and reason about uncertainties in the environment and other road users' behaviors.
- Multi-agent coordination: Extend the framework to handle scenarios involving multiple autonomous vehicles, incorporating explicit communication and coordination mechanisms.
- Safety-aware RL: Investigate techniques to incorporate formal safety constraints and guarantees into the RL training process.
- Lifelong learning: Develop approaches for continual adaptation and improvement of the learned policies as the vehicle encounters new scenarios in real-world deployments.
- Human-AI interaction: Explore methods for seamless interaction between the autonomous system and human drivers or passengers, including explanations of the AI's decisions and smooth transitions between autonomous and manual control.

In conclusion, our hierarchical RL framework represents a significant step towards enabling safe and efficient autonomous vehicle navigation in complex urban environments. By addressing the limitations and pursuing the suggested future directions, we believe this approach has the potential to contribute to the development of more capable and reliable autonomous driving systems.

References

[1] Thakur, D. (2020). Optimizing Query Performance in Distributed Databases Using Machine Learning Techniques: A Comprehensive Analysis and Implementation. IRE Journals, 3(12), 266-276.

[2] Murthy, P. & Bobba, S. (2021). AI-Powered Predictive Scaling in Cloud Computing: Enhancing Efficiency through Real-Time Workload Forecasting. IRE Journals, 5(4), 143-152.

[3] Krishna, K., Mehra, A., Sarker, M., & Mishra, L. (2023). Cloud-Based Reinforcement Learning for Autonomous Systems: Implementing Generative AI for Real-time Decision Making and Adaptation. IRE Journals, 6(8), 268-278.

[4] Thakur, D., Mehra, A., Choudhary, R., & Sarker, M. (2023). Generative AI in Software Engineering: Revolutionizing Test Case Generation and Validation Techniques. IRE Journals, 7(5), 281-293.

[5] Thakur, D. (2021). Federated Learning and Privacy-Preserving AI: Challenges and Solutions in Distributed Machine Learning. International Journal of All Research Education and Scientific Methods (IJARESM), 9(6), 3763-3771.

[6] Mehra, A. (2020). Unifying Adversarial Robustness and Interpretability in Deep Neural Networks: A Comprehensive Framework for Explainable and Secure Machine Learning Models. International Research Journal of Modernization in Engineering Technology and Science, 2(9), 1829-1838.

[7] Krishna, K. (2022). Optimizing Query Performance in Distributed NoSQL Databases through Adaptive Indexing and Data Partitioning Techniques. International Journal of Creative Research Thoughts, 10(8), e812-e823.

[8] Krishna, K. (2020). Towards Autonomous AI: Unifying Reinforcement Learning, Generative Models, and Explainable AI for Next-Generation Systems. Journal of Emerging Technologies and Innovative Research, 7(4), 60-68.

[9] Murthy, P. & Mehra, A. (2021). Exploring Neuromorphic Computing for Ultra-Low Latency Transaction Processing in Edge Database Architectures. Journal of Emerging Technologies and Innovative Research, 8(1), 25-33.

[10] Krishna, K. & Thakur, D. (2021). Automated Machine Learning (AutoML) for Real-Time Data Streams: Challenges and Innovations in Online Learning Algorithms. Journal of Emerging Technologies and Innovative Research, 8(12), f730-f739.

[11] Murthy, P. & Thakur, D. (2022). Cross-Layer Optimization Techniques for Enhancing Consistency and Performance in Distributed NoSQL Database. International Journal of Enhanced Research in Management & Computer Applications, 11(8), 35-41.

[12] Murthy, P. (2020). Optimizing Cloud Resource Allocation using Advanced AI Techniques: A Comparative Study of Reinforcement Learning and Genetic Algorithms in Multi-Cloud Environments. World Journal of Advanced Research and Reviews, 7(2), 359-369.

[13] Mehra, A. (2021). Uncertainty Quantification in Deep Neural Networks: Techniques and Applications in Autonomous Decision-Making Systems. World Journal of Advanced Research and Reviews, 11(3), 482-490.

[14] A. E. Sallab, M. Abdou, E. Perot, and S. Yogamani, "Deep reinforcement learning framework for autonomous driving," Electronic Imaging, vol. 2017, no. 19, pp. 70-76, 2017.

[15] A. Kendall et al., "Learning to drive in a day," in 2019 International Conference on Robotics and Automation (ICRA), 2019, pp. 8248-8254.

[16] L. Tai, J. Zhang, M. Liu, and W. Burgard, "Socially compliant navigation through raw depth inputs with generative adversarial imitation learning," in 2018 IEEE International Conference on Robotics and Automation (ICRA), 2018, pp. 1111-1117.

[17] J. Chen, B. Yuan, and M. Tomizuka, "Model-free deep reinforcement learning for urban autonomous driving," in 2019 IEEE Intelligent Transportation Systems Conference (ITSC), 2019, pp. 2765-2771.

[18] R. S. Sutton, D. Precup, and S. Singh, "Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning," Artificial Intelligence, vol. 112, no. 1-2, pp. 181-211, 1999.

[19] A. G. Barto and S. Mahadevan, "Recent advances in hierarchical reinforcement learning," Discrete Event Dynamic Systems, vol. 13, no. 4, pp. 341-379, 2003.

[20] S. Shalev-Shwartz, S. Shammah, and A. Shashua, "Safe, multi-agent, reinforcement learning for autonomous driving," arXiv preprint arXiv:1610.03295, 2016.

[21] Z. Qiao, K. Muelling, J. M. Dolan, P. Palanisamy, and P. Mudalige, "Automatically generated curriculum based reinforcement learning for autonomous vehicles in urban environment," in 2018 IEEE Intelligent Vehicles Symposium (IV), 2018, pp. 1233-1238.

[22] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum learning," in Proceedings of the 26th Annual International Conference on Machine Learning, 2009, pp. 41-48.

[23] V. Mnih et al., "Human-level control through deep reinforcement learning," Nature, vol. 518, no. 7540, pp. 529-533, 2015.

[24] T. Schaul, J. Quan, I. Antonoglou, and D. Silver, "Prioritized experience replay," arXiv preprint arXiv:1511.05952, 2015.

[25] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "CARLA: An open urban driving simulator," in Proceedings of the 1st Annual Conference on Robot Learning, 2017, pp. 1-16.

[26] D. Isele, R. Rahimi, A. Cosgun, K. Subramanian, and K. Fujimura, "Navigating occluded intersections with autonomous vehicles using deep reinforcement learning," in 2018 IEEE International Conference on Robotics and Automation (ICRA), 2018, pp. 2034-2039.

[27] M. Bouton, J. Karlsson, A. Nakhaei, K. Fujimura, M. J. Kochenderfer, and J. Tumova, "Reinforcement learning with probabilistic guarantees for autonomous driving," in Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, 2019, pp. 1823-1825.

[28] X. Liang et al., "PCDN: Parallel coordinate descent networks for urban scene parsing," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 2894-2903.

[29] W. Ding, J. Chen, and S. Shen, "Predicting vehicle behaviors over an extended horizon using behavior interaction network," in 2019 International Conference on Robotics and Automation (ICRA), 2019, pp. 8634-8640.

[30] M. Bansal, A. Krizhevsky, and A. Ogale, "ChauffeurNet: Learning to drive by imitating the best and synthesizing the worst," in Robotics: Science and Systems, 2019.

[31] D. Sadigh, S. Sastry, S. A. Seshia, and A. D. Dragan, "Planning for autonomous cars that leverage effects on human actions," in Robotics: Science and Systems, 2016.

[32] N. Rhinehart, R. McAllister, K. Kitani, and S. Levine, "PRECOG: Prediction conditioned on goals in visual multi-agent settings," in Proceedings of the IEEE International Conference on Computer Vision, 2019, pp. 2821-2830.

[33] A. Kuefler, J. Morton, T. Wheeler, and M. Kochenderfer, "Imitating driver behavior with generative adversarial networks," in 2017 IEEE Intelligent Vehicles Symposium (IV), 2017, pp. 204-211.

[34] Y. Chen, C. Dong, P. Palanisamy, P. Mudalige, K. Muelling, and J. M. Dolan, "Attention-based hierarchical deep reinforcement learning for lane change behaviors in autonomous driving," in 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2019, pp. 3697-3703.

[35] M. Jaritz, R. De Charette, M. Toromanoff, E. Perot, and F. Nashashibi, "End-to-end race driving with deep reinforcement learning," in 2018 IEEE International Conference on Robotics and Automation (ICRA), 2018, pp. 2070-2075.

[36] Pin Wang, Ching-Yao Chan, Arnaud de La Fortelle. A Reinforcement Learning Based Approach for Automated Lane Change Maneuvers. IEEE Intelligent Vehicles Symposium, Jun 2018, Chang Shu, China. ff10.48550/arXiv.1804.07871ff.