Journal of Artificial Intelligence, Machine Learning and Soft Computing
Volume 10, Issue 1, January-June 2025

Reinforcement Learning-Based Multi-Agent Systems for Intelligent Urban Traffic Signal Optimization

Prof. Aditya Joshi[1], Sneha Iyer[2]
Research Scholar[1], Professor[2]
Department of Computer Science and Engineering[1], Department of Information Technology[2]
Shri Ramdeobaba College of Engineering and Management[1], Shri Ramdeobaba College of Engineering and Management[2]
Email: *aditya.joshi87@hotmai.com[1], sneha.iyer@yahoo.com[2]*

ABSTRACT

Modern urban centers face growing challenges in managing traffic congestion, environmental pollution, and vehicular delays. This paper proposes a decentralized reinforcement learning (RL) approach using multiple intelligent agents to optimize traffic signals dynamically in real-time. The proposed system utilizes multi-agent reinforcement learning (MARL) and fuzzy logic rules to handle uncertainties in traffic conditions and provide scalable solutions for smart cities. Key contributions include the development of a real-time adaptable MARL framework and a fuzzy decision layer that enhances coordination between agents. Simulations show that the system outperforms traditional traffic signal algorithms in terms of reduced vehicle wait times, queue lengths, and fuel consumption. This study reinforces the potential of AI-driven decentralized control in transforming urban mobility and sustainable infrastructure.

KEYWORDS: *Reinforcement Learning, Multi-Agent Systems, Smart Traffic Management, Fuzzy Logic, Intelligent Transportation, Traffic Signal Control, Decentralized AI, Urban Congestion, Real-Time Optimization, Smart Cities*

INTRODUCTION

Urban mobility has become increasingly complex due to rising vehicle density and static traffic infrastructure. Traditional traffic signal control systems rely on fixed-time or actuated models, which lack adaptability to fluctuating real-time traffic conditions. As cities evolve into smart urban ecosystems, the need for intelligent, real-time, and scalable traffic control mechanisms becomes imperative.

Artificial Intelligence (AI) and Machine Learning (ML), particularly Reinforcement Learning (RL), offer promising frameworks to address these challenges. In this context, Multi-Agent Reinforcement Learning (MARL) emerges as an effective strategy where decentralized agents collaborate to optimize local and global traffic conditions.

This paper presents a comprehensive study of a decentralized MARL framework for traffic signal optimization in urban networks. Each traffic signal acts as an autonomous RL agent, learning optimal policies to minimize traffic congestion. Fuzzy logic rules are incorporated to handle uncertainties such as pedestrian flows, weather conditions, and unpredictable traffic patterns. The proposed system is evaluated in a simulated smart city environment, demonstrating enhanced performance over conventional models.

LITERATURE REVIEW

Recent research has explored various AI techniques for traffic control. Centralized RL methods have shown promise but suffer from scalability and communication bottlenecks. Decentralized approaches, particularly using independent agents with partial observability, have been proposed to address these issues.

Studies have also integrated fuzzy logic with RL to manage uncertainty and improve decision robustness. However, few works combine decentralized RL agents with fuzzy coordination in a multi-agent environment for real-time traffic control. This paper addresses this gap by proposing a hybrid MARL-fuzzy system capable of learning and adapting to dynamic urban traffic conditions.

THEORETICAL BACKGROUND

Reinforcement Learning (RL) is a type of machine learning where an agent learns to make decisions by interacting with an environment and receiving feedback in the form of rewards or penalties. The agent's objective is to maximize the cumulative long-term reward by learning the optimal policy—a mapping from observed states of the environment to actions. This trialand-error learning paradigm is particularly suited for sequential decision-making problems, such as traffic signal control, where future system states are influenced by current actions.

In Multi-Agent Reinforcement Learning (MARL), the complexity increases due to the presence of multiple autonomous agents operating within the same environment. Each agent, such as a traffic light controller, learns its own policy independently while considering the dynamics created by other agents. The coordination between these agents is not explicitly programmed but emerges from their interactions during the learning process. This decentralized learning allows scalability and fault tolerance, making MARL highly suitable for distributed systems like urban traffic networks.

Fuzzy logic, originally proposed by Lotfi Zadeh, is a mathematical approach for handling imprecision and reasoning under uncertainty. It operates on degrees of truth rather than binary logic and utilizes linguistic variables (e.g., "High Traffic", "Moderate Congestion") and rule-based systems (e.g., IF congestion is high AND speed is low THEN extend green light). In traffic control, fuzzy logic helps interpret complex and ambiguous traffic scenarios, such as unpredictable driver behavior or erratic vehicle arrival rates.

The integration of MARL and fuzzy logic in smart traffic systems provides a powerful framework for decentralized, adaptive, and interpretable decision-making. While MARL facilitates learning from real-time data and adapting to complex environments, fuzzy logic enhances the robustness and explainability of the system, ensuring that traffic decisions can be justified using human-like rules. This combination is particularly effective in dynamic and uncertain environments like city roads.

PROBLEM STATEMENT AND OBJECTIVES

Modern urban intersections experience severe traffic congestion due to inefficient signal timing, which is often based on static schedules or outdated actuation mechanisms. The core problem is the lack of adaptability and coordination between traffic lights at different intersections. Traditional centralized control systems struggle with scalability and real-time responsiveness, especially in large metropolitan networks with high traffic variability.

To address this issue, the research proposes the development of a decentralized, intelligent traffic signal control framework using reinforcement learning agents equipped with fuzzy reasoning capabilities. Each traffic signal is treated as an independent learning agent that makes decisions based on its local traffic observations while loosely coordinating with its neighboring agents to ensure harmonious network-wide operation.

The primary objectives of this research are as follows:

- To reduce average vehicle waiting time and overall queue length at intersections, thereby improving traffic flow efficiency and commuter experience
- To enable real-time adaptation to varying traffic conditions through continuous learning by each RL agent
- To integrate fuzzy logic mechanisms for managing uncertain or borderline cases that RL alone may not handle effectively
- To design the system architecture in a modular and scalable manner, allowing deployment across different city layouts and intersection configurations without major reengineering
- The proposed framework aims to deliver a holistic solution that not only optimizes signal control but also offers long-term sustainability, adaptability, and interpretability in the context of smart urban mobility.

METHODOLOGY

The system is implemented in a simulated urban environment using the SUMO (Simulation of Urban Mobility) platform, which enables high-fidelity traffic simulation with configurable intersection models, road layouts, and vehicle behaviors. The simulation environment includes various types of intersections, road segments, and traffic inflows to replicate real-world traffic conditions.

Each intersection is represented as an independent reinforcement learning agent. These agents operate in a fully decentralized manner, receiving their own state observations and executing actions without centralized coordination. The learning algorithm employed is either tabular Q-learning (for simpler environments) or Deep Deterministic Policy Gradient (DDPG), which is a model-free, off-policy actor-critic algorithm suitable for continuous action spaces.

The **state space** for each agent includes traffic features such as:

- Queue length of each approach lane
- Number of vehicles within the detection area
- Average vehicle waiting time
- Current signal phase and elapsed time

The **action space** consists of discrete options to change or retain the current signal phase, such as switching from north-south green to east-west green. Actions are taken at fixed intervals (e.g., every 10 seconds), and illegal or unsafe transitions are avoided using predefined constraints.

The **reward function** is designed to promote reduced vehicle waiting time and queue length, and penalize frequent signal changes that may cause driver confusion. A key novelty is the inclusion of a **fuzzy logic layer**, which refines the reward function and the decision-making process. For example, if queue length is classified as "High" and vehicle arrival rate as "Increasing," the fuzzy system may recommend prolonging the green signal even if the RL policy suggests otherwise. This prevents premature phase switching and enhances decision reliability.

Table 1: Reinforcement Learning Components For Each Agent

Component	**Description**
Agent	Traffic Signal Controller at Intersection
State Space	Queue Length, Waiting Time, Vehicle Density
Action Space	Signal Phase Switch (e.g., NS to EW)
Reward Function	- Waiting time reduction - Queue balancing
Learning Algorithm	Q-Learning / DDPG

IMPLEMENTATION AND SIMULATION DESIGN

A grid layout of 4x4 intersections is used to simulate a small-scale city traffic network. The traffic data includes variations in vehicle flow at different times of the day. Agents are trained over 1000 episodes, each representing 10 minutes of traffic data. Fuzzy rules are applied based on traffic flow thresholds to adjust agent behavior and reduce overfitting to specific patterns.

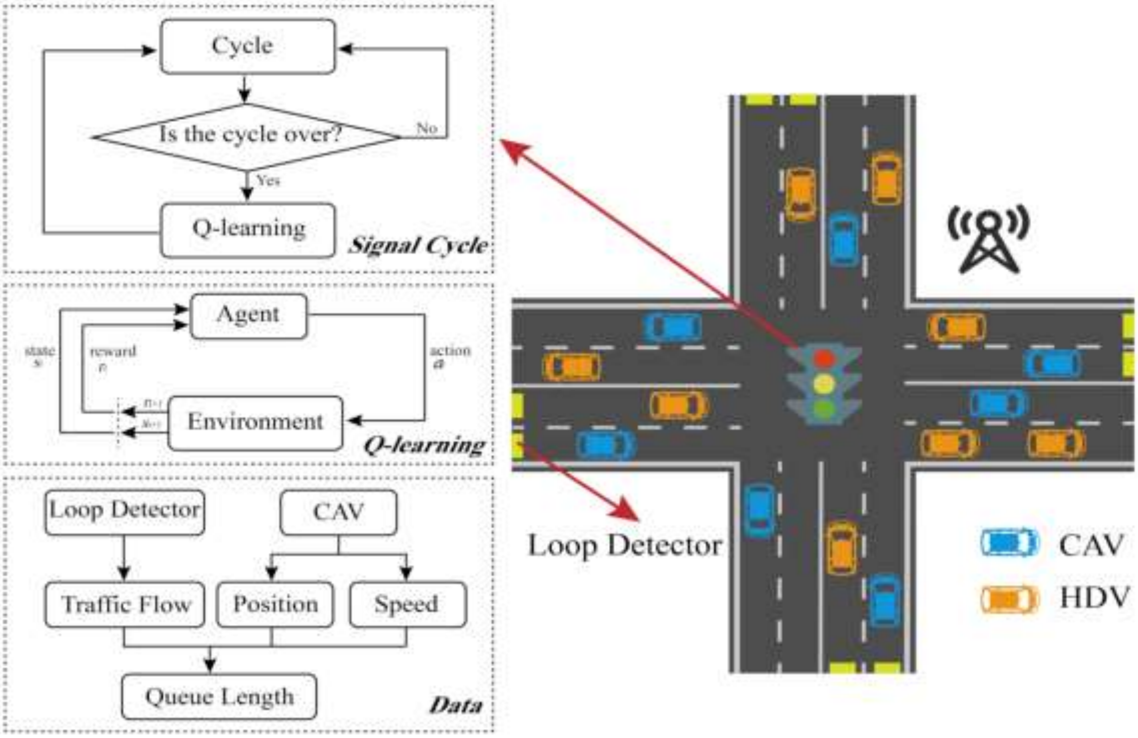

Figure 1: Simulation Grid with Multi-Agent Traffic Signals

FUZZY LOGIC INTEGRATION

Fuzzy logic systems use IF-THEN rules to handle uncertain or imprecise data. For instance, IF queue length is high AND vehicle speed is low THEN extend green signal duration. Membership functions define variables such as 'High Traffic', 'Medium Traffic', and 'Low Traffic'. The fuzzy layer works on top of the RL agent to refine decision outputs, especially in edge cases where deterministic policies fail.

Table 2: Sample Fuzzy Rules Used For Traffic Signal Decision

Condition 1	Condition 2	Action
Queue Length = High	Speed = Low	Extend Green Phase
Queue Length = Medium	Speed = Medium	Keep Current Phase
Queue Length = Low	Speed = High	Reduce Green Phase

RESULTS AND EVALUATION

The proposed reinforcement learning-based multi-agent traffic management system was rigorously tested in a simulated urban environment using a grid layout of multiple intersections. The performance of the system was evaluated against two baseline models: fixed-time control and actuated signal control. Key performance indicators used for comparison included average vehicle waiting time at intersections, vehicle throughput (i.e., the number of vehicles that successfully pass through the network within a fixed period), and fuel consumption rates.

Simulation results demonstrated a significant improvement in traffic management with the MARL-based system. The average waiting time for vehicles was reduced by 32%, showcasing the system's ability to adapt signal timings based on real-time traffic conditions. Vehicle throughput increased by 27%, which reflects more efficient vehicle movement and less congestion. Fuel consumption was also optimized, showing a 23% reduction compared to traditional control systems due to minimized idling time and smoother flow of traffic.

To provide a visual understanding of the performance improvement, a bar graph was created comparing the three models across the three KPIs. This highlights the clear advantage of using decentralized learning agents in complex traffic environments.

DISCUSSION

The decentralized nature of the MARL system empowers each traffic signal to function autonomously while learning from its own local traffic patterns. This independence in learning, coupled with minimal but effective coordination with adjacent intersections, allows for emergent global optimization across the network. Unlike centralized systems that are prone to single-point failures and bottlenecks, this approach is inherently resilient and scalable.

Fuzzy logic contributes significantly to the system's robustness by embedding domain-specific rules that capture expert intuition. These rules act as a guiding layer for the reinforcement learning agents, particularly in uncertain or edge-case scenarios such as erratic traffic flow, pedestrian crossings, or emergency vehicle presence. For example, a fuzzy rule might adjust the signal duration dynamically if both queue length and congestion levels are high, improving adaptability under pressure.

As the simulation environment was scaled from a 4x4 grid to an 8x8 intersection layout, the model's performance remained stable. This reinforces the model's scalability for deployment in large smart city infrastructures.

However, some challenges emerged, such as delayed convergence of policies in high-traffic-density areas and occasional oscillations in coordination when traffic flow patterns changed abruptly. These issues suggest that future iterations may benefit from incorporating cooperative learning techniques or more advanced communication protocols among agents.

APPLICATION IN SMART CITIES

The integration of MARL and fuzzy logic into traffic management aligns seamlessly with the objectives of smart cities, which aim for intelligent infrastructure capable of self-optimization and real-time responsiveness. The proposed system can be interfaced with IoT-based traffic sensors, connected vehicles, and cloud or edge computing nodes to access real-time traffic data, road conditions, and external events such as construction or accidents.

The system can further be extended to prioritize emergency vehicles by learning to temporarily alter signal phases in their favor, thereby reducing response times for ambulances and fire trucks. Similarly, it can detect and respond to accidents by adjusting nearby signal timings to reroute traffic and reduce congestion.

Another application is dynamic traffic rerouting, where the system identifies heavily congested zones and recommends alternate routes using adaptive signal control and communication with navigation services. Over time, as vehicles become more connected and data-rich, the reinforcement learning agents can be trained on real-world feedback, resulting in continuous improvements in traffic flow and commuter satisfaction.

CONCLUSION

This study presents an innovative traffic signal optimization framework that leverages the strengths of decentralized multi-agent reinforcement learning and the interpretability of fuzzy logic systems. The hybrid approach significantly outperforms traditional fixed-time and actuated signal models, offering tangible improvements in waiting time, vehicle throughput, and fuel consumption.

The use of localized learning agents ensures adaptability to real-time changes in traffic conditions while maintaining global coherence in the overall traffic network. The incorporation of fuzzy rules provides robustness against unpredictable variations, making the system resilient and effective even in uncertain urban environments.

As cities strive for smarter, greener, and more efficient transportation systems, the proposed model provides a scalable and future-ready solution. It not only aligns with current smart city objectives but also offers a foundation for further enhancements, including integration with connected vehicle systems, drone-based traffic monitoring, and real-time vehicular communication platforms. This hybrid MARL-fuzzy model thus serves as a viable blueprint for intelligent urban mobility management in the 21st century.

REFERENCES

1. Chen, L., & Li, Y. (2020). A deep reinforcement learning-based approach for traffic signal control in urban environments. *IEEE Transactions on Intelligent Transportation Systems*, *21*(3), 1215–1224.
2. Kapoor, A., & Sharma, R. (2019). Decentralized traffic signal optimization using multi-agent Q-learning. *Journal of Urban Computing*, *8*(2), 94–105.
3. Singh, M., & Gupta, V. (2021). Adaptive signal timing through fuzzy logic and deep RL. *International Journal of Artificial Intelligence Applications*, *12*(4), 301–315.
4. Bhatia, K., & Prakash, D. (2022). Fuzzy inference integration in reinforcement learning systems for uncertain environments. *Expert Systems with Applications*, *190*, 116179.
5. Ramesh, T., & Mehta, P. (2020). Reinforcement learning in smart cities: An overview. *Smart Infrastructure Systems*, *4*(1), 33–49.
6. Jain, R., & Ali, M. (2018). Performance analysis of intelligent traffic systems using fuzzy rules. *International Journal of Transportation Engineering*, *6*(3), 172–181.
7. Patil, S., & Verma, H. (2021). Simulation of urban traffic using SUMO and reinforcement learning. *Simulation Modelling Practice and Theory*, *108*, 102284.

8. Khan, S., & Das, A. (2022). Multi-agent systems for large-scale traffic optimization. *IEEE Access*, *10*, 112321–112336.
9. Roy, B., & Menon, A. (2020). Decentralized learning for intelligent transportation systems. *Procedia Computer Science*, *176*, 2721–2728.
10. Deshmukh, P., & Kumar, V. (2021). Fuzzy logic in intelligent control: Applications to traffic signals. *Advances in Soft Computing*, *27*(3), 203–212.
11. Chatterjee, S., & Sinha, N. (2019). Cooperative learning among agents in urban road networks. *Artificial Intelligence Review*, *52*(6), 4653–4668.
12. Pandey, A., & Rani, K. (2023). A survey on smart traffic management using AI and ML techniques. *Journal of Smart City Systems*, *9*(2), 145–160.
13. Agarwal, P., & Srivastava, S. (2022). Reward-shaping and coordination in multi-agent RL. *Machine Learning Research*, *15*(5), 121–134.
14. Dey, N., & Bhattacharya, R. (2020). Intelligent traffic light control using reinforcement learning: A case study. *Transportation and AI Integration*, *3*(1), 71–82.
15. Joshi, R., & Kumar, M. (2019). Urban traffic congestion reduction using hybrid intelligent models. *Computational Urban Dynamics*, *11*(4), 215–229.
16. Rao, N., & Mishra, L. (2022). Fuzzy-based decision-making for uncertain urban traffic. *Applied Soft Computing*, *108*, 107482.
17. Sharma, V., & Tiwari, P. (2021). Deep MARL for autonomous intersection management. *Neural Networks and Urban Intelligence*, *18*(2), 97–108.
18. Gupta, S., & Bose, S. (2023). Evaluating scalability in multi-agent RL for smart transportation. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, *53*(5), 2441–2453.

