iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 **Issue:** III **Month of publication:** March 2025

DOI: https://doi.org/10.22214/ijraset.2025.67570

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 13 Issue III Mar 2025- Available at www.ijraset.com

Adaptive Traffic Management System Using Reinforcement Learning

Dr Kanjana G[1], B R Neha[2], Nayana V Nair[3], Pooja A J[4], Varsha P S[5]
[1]*Assistant Professor, Dept. of Electronics and Communication*
[2, 3, 4, 5]*Dept. of Applied Electronics and Instrumentation*

Abstract: ***Traffic congestion is a significant issue in urban areas, leading to increased travel delays, fuel consumption, and environmental pollution. Traditional traffic management systems, which use fixed-timer signals, rule-based controls, manual intervention by traffic police, and electronic sensor-based methods, often struggle to adapt to dynamic traffic conditions. To address these challenges an Adaptive Traffic Management System (ATMS) using Reinforcement Learning (RL) is proposed to optimize signal timings and improve traffic flow. The system dynamically adjusts traffic signal timings based on real-time traffic flow, ensuring efficient traffic movement. Key features of this approach include prioritization of emergency vehicles, reduced waiting time for emergency vehicles and zero waiting time for emergency vehicles. Additionally, our system minimizes pollution and prevents potential collisions by optimizing traffic flow. The system leverages the SUMO (Simulation of Urban MObility) platform to simulate real-world traffic scenarios and evaluate performance. Reinforcement learning enables the system to learn from traffic patterns and make real-time decisions to minimize congestion. The model is trained using deep Qlearning to optimize signal control at intersections. The adaptive system reduces average waiting time, travel delays, and fuel consumption compared to conventional traffic light control methods. Through iterative learning, the model adjusts signal phases dynamically, ensuring smooth and efficient vehicle movement. The simulation results demonstrate significant improvements in road network efficiency and reduced congestion. The proposed system can also handle unpredictable traffic disruptions such as accidents or roadblocks. Performance is evaluated based on key metrics, including average vehicle delay, queue length, and throughput. The flexibility of reinforcement learning allows the system to adapt to varying traffic conditions without manual intervention. Implementing AI-driven adaptive traffic control can enhance urban mobility and commuter experiences. By optimizing traffic signal control, this approach contributes to sustainable and efficient transportation networks. This research highlights the potential of reinforcement learning in addressing traffic congestion through intelligent decision-making.***

Keywords: ***Adaptive Traffic Management, Deep Q Learning, Intelligent Transportation Systems, Queue Length, Real-Time Traffic Control, Reinforcement Learning, SUMO, Traffic Congestion, Traffic Signal Optimization, Traffic Signal Control, Vehicle Delay.***

I. INTRODUCTION

Traffic congestion poses a serious challenge in urban areas, causing extended travel delays, greater fuel consumption, and heightened environmental pollution. As cities continue to grow, the demand for efficient traffic management solutions has become more critical. Traditional traffic control systems, which rely on fixed signal timings or pre-defined rules, often fail to adapt to changing traffic conditions. These limitations result in inefficient traffic flow, longer waiting times, and increased carbon emissions. A critical issue in traffic management is the delay faced by emergency vehicles, such as ambulances which require immediate passage through intersections. Fixed timer traffic signals and conventional rule-based systems often fail to prioritize emergency vehicles, leading to dangerous delays that can impact lifesaving services. Ensuring rapid and uninterrupted movement of emergency vehicles is essential for public safety and effective traffic management.

To address these challenges, adaptive traffic management systems (ATMS) have emerged as a promising solution. By utilizing Reinforcement Learning (RL), traffic signals can dynamically adjust based on real-time traffic conditions. RL-based models continuously learn from traffic patterns and optimize signal timings to reduce congestion. Unlike traditional approaches, reinforcement learning enables a data-driven and self-improving decision-making process, making the system more efficient and responsive. This project proposes an adaptive traffic management system using reinforcement learning, implemented and evaluated using the SUMO (Simulation of Urban MObility) platform. SUMO provides a realistic traffic simulation environment, allowing for the testing and validation of AI-driven traffic control strategies. The proposed system is trained using deep Q-learning to optimize signal control at intersections.

©IJRASET: All Rights are Reserved

The goal is to minimize vehicle waiting time, improve traffic throughput, and enhance overall road network efficiency. By leveraging reinforcement learning, this system aims to create an intelligent traffic management framework that can respond dynamically to varying traffic conditions. The paper also discusses ethical issues and difficulties, with a focus on the significance of precise data analysis and protecting client privacy.

II. METHODOLOGY

1) Traffic Simulation Environment

For evaluating the effectiveness of Reinforcement Learning in Adaptive Traffic Control, Simulation of Urban MObility (SUMO) is used as the simulation platform. SUMO provides a realistic traffic environment with different types of vehicles including emergency vehicles.

2) Problem Formulation

The objective of the system is to improve traffic signal control to reduce vehicle waiting time while ensuring that emergency vehicle almost achieve zero waiting time. This is achieved by dynamically adjusting traffic signals based on real time traffic density and vehicle priority.

3) Reinforcement Learning

Reinforcement Learning is used to design an intelligent traffic signal control that learns from experience to optimize traffic flow. The traffic signal control problem is formulated as a Markov Decision Process (MDP) where:

- State (S): Represents real time traffic conditions including vehicle densities and emergency vehicle presence.
- Action (A): Possible changes in traffic signal durations and phase sequences.
- Reward (R): Feedback mechanism where shorter waiting times for regular vehicles and immediate passage for emergency vehicles yield higher rewards.
- Policy (π): Learning strategy that determines the best actions based on observed traffic conditions.

To train the system, Deep Reinforcement Learning techniques are used which utilize neural networks to approximate optimal policies.

4) Deep Q Network

By using deep learning, DQN can efficiently process high-dimensional inputs, such as real-time traffic data from multiple intersections, vehicle movement patterns, or sensor readings. This capability enables the system to handle complex traffic situations more effectively than traditional Q-learning. Additionally, techniques like experience replay (where past experiences are stored and replayed to stabilize training) and target networks (which help prevent instability during learning) make DQN more robust and reliable for real-world applications.DQN extends Q-learning by leveraging deep neural networks to approximate Q-values, making it well-suited for large, complex environments like adaptive traffic management, where handling real-time decision-making with high-dimensional data is critical.

5) Bellman's Equation

The learning process follows Bellman's equation, which iteratively updates the Q values.

$$Q(s,a)=R+\gamma \max Q(s',a')$$

- Q (s,a) represents the expected reward for taking action **a** in state **s**.
- R is the immediate reward obtained.
- γ is the discount factor that account for future rewards.
- max Q (s', a') is the highest expected reward from the next state.

6) System Implementation

The program is implemented in Python which utilizes Reinforcement Learning frameworks such as TensorFlow

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

a) Data Collection – Traffic data is collected using the Simulation of Urban Mobility (SUMO), an open-source traffic simulation tool. The TraCI (Traffic Control Interface) module is used to extract real-time traffic data, including vehicle count, waiting times, speeds, and queue lengths. This data serves as input for the reinforcement learning model, providing a detailed representation of the traffic environment.

b) State Representation -To enable the RL model to make informed decisions, raw traffic data must be transformed into numerical state representations. These include traffic density, signal phase information, and vehicle waiting times. Proper state representation ensures that the model effectively interprets traffic conditions, allowing it to learn optimal signal control policies

c) Training the model – The Deep Q-Network (DQN) algorithm is used to train the RL model for adaptive traffic signal control. The model selects actions (traffic signal changes) based on observed traffic states and learns from rewards, which are designed to minimize vehicle waiting times and congestion. Through iterative training over multiple episodes, the model improves its decision-making capabilities.

d) Simulation And Evaluation – Once trained, the model is tested within SUMO to compare its performance against traditional fixed-time traffic signals. Key performance metrics include vehicle waiting time(prioritizing emergency vehicles), queue length, and overall traffic flow efficiency. The expectation is that the trained RL model will dynamically adjust signal timings to improve traffic conditions more effectively than fixed-time systems

7) Training

During training, the agent explores different signal timing strategies and updates its policy using a reward function that prioritizes reduced waiting time for both regular and emergency vehicles. As the agent interacts with the environment, it balances exploration and exploitation to refine its policy. Initially, it explores various traffic signal configurations to gather diverse experiences. Over time, the shift towards exploiting the learned strategies those yield higher rewards.

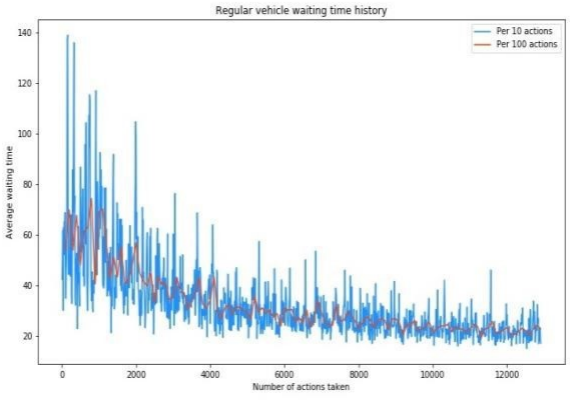


Figure 1: Regular Vehicle Waiting Time History

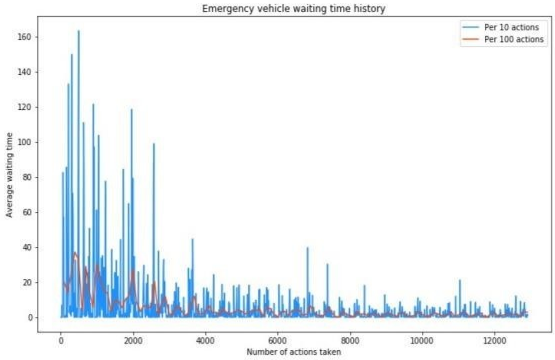


Figure 2: Emergency Vehicle Waiting Time History

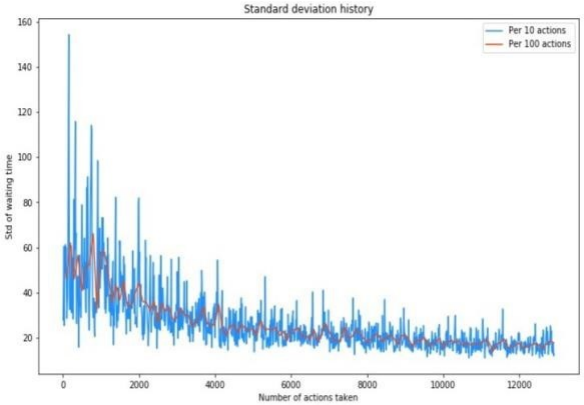


Figure 3: Regular Vehicle Waiting Time Standard Deviation

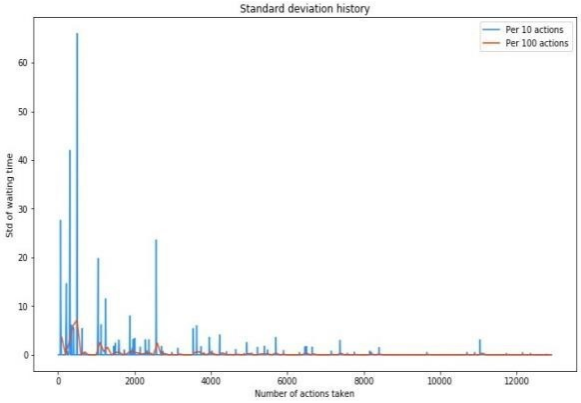


Figure 4: Emergency Vehicle Waiting Standard Deviation

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

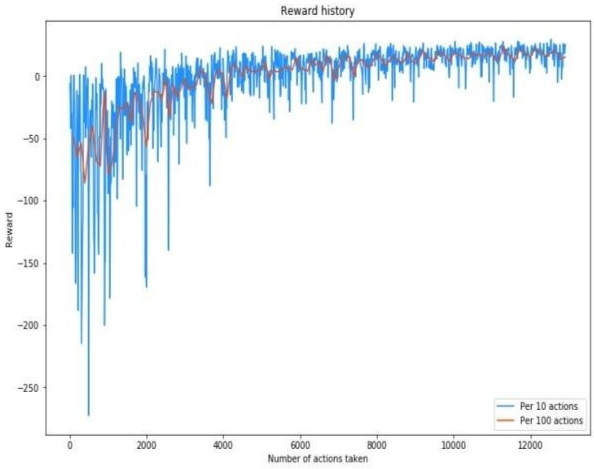


Figure 5: Reward History

In Figure 1, initially the waiting time fluctuates significantly, indicating high variability as the reinforcement learning model explores different control strategies. Over time, both curves show a decreasing trend, suggesting that the model learns to reduce waiting times and optimize regular vehicle flow. The smoother orange line (per 100 actions) shows a clearer trend, confirming that the system stabilizes waiting times as training progress.

In Figure 2, the emergency vehicle waiting time follows a similar pattern to the regular vehicles, with sharp fluctuations at the beginning and a gradual decrease as learning improves. The high peaks at the start indicate scenarios where emergency vehicles experienced long delays before the system optimized their priority .As training continues, the waiting time stabilizes at lower values, meaning the model effectively reduces delays for emergency vehicles.

In Figure 3, initially the standard deviation is high, indicating large fluctuations in waiting times across different traffic scenarios. As training progresses, both curves decrease, signifying that the system is stabilizing waiting times for regular vehicles. The orange line provides a smoother trend, confirming that the system is reducing inconsistencies over time.

In Figure 4, the standard deviation for emergency vehicles starts high but decreases significantly over time. The sharp initial spikes indicate inconsistent waiting times, possibly due to unpredictable early-stage learning policies. Over time, fluctuations diminish, showing that the model has successfully learned to prioritize emergency vehicles consistently. The orange line remains almost flat towards the end, meaning that waiting times for emergency vehicles have become highly stable.

In Figure 5, as training progresses, the reward steadily increases, suggesting that the model is learning to optimize traffic control strategies effectively. Towards the later stages, the reward stabilizes at high value, with minimal fluctuations, indicating that agent has converged to an optimal policy. The orange line provides a smoother trend, showing the long term improvement in policy performance.

8) Performance Evaluation

The effectiveness of this approach is measured based on:

- Reduction in average vehicle waiting time compared to traditional methods.
- Almost zero waiting time for emergency vehicles.
- Overall improvement in traffic flow efficiency.

III. SYSTEM REQUIREMENTS

A. Simulation of Urban Mobility (SUMO)

SUMO is used to simulate traffic scenarios. SUMO is a well-known open-source traffic simulation tool that provides APIs and a Graphical user interface (GUI) for efficiently modelling and managing complex road networks. It supports dynamic vehicle routing based on right-side driving rules at intersections and allows users to create road networks in different grid formats through its visual interface. Additionally, SUMO enables users to control traffic signal at intersections using custom policies. It also offers the capability to capture simulation snapshots at each step helping us extract state information.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |



Figure 6: SUMO Simulation

B. Anaconda

Anaconda is an open-source distribution that makes coding in Python easier and is especially used for data science, machine learning and deep learning–based applications. Anaconda includes a package manager (conda) which makes installation and management easy for dependency free libraries. It has everything for code development, which includes vital utilities like Jupyter Notebook, Spyder and VS Code and experimentation. Furthermore, Anaconda easily integrates with frameworks such as TensorFlow, PyTorch, and OpenAI Gym, hence, it is preferred for reinforcement learning-oriented traffic simulations

IV. RESULT AND DISCUSSION

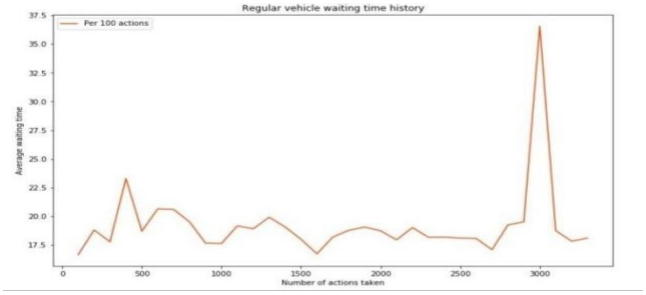


Figure 7: Emergency Vehicle Time Per 100 actions

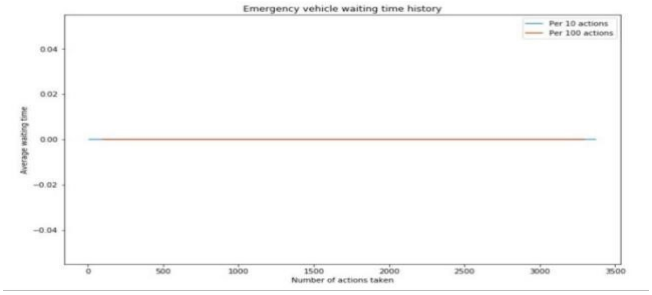


Figure 8: Standard Deviation of waiting time Per 100 actions

In Figure 7, it represents the average waiting time history of regular vehicles over the training phase, measured per 100 actions. It shows fluctuations initially but gradually stabilizes with a peak at one point, possibly due to a significant traffic surge or policy adjustment.

In Figure 8, it represents the waiting time history of emergency vehicles, where the waiting time remains nearly zero throughout indicating that the reinforcement learning agent has successfully prioritized emergency vehicles, ensuring minimal delays.

The analysis of the graphs provides a comprehensive evaluation of the RL-based adaptive traffic management system. The result highlights significant improvements in traffic efficiency, emergency vehicle prioritization, and real time adaptability compared to the traditional fixed-signal systems. The RL-based system significantly reduces the waiting time of regular vehicles over multiple training steps, ensuring smoother traffic flow. Also, emergency vehicles experiences near zero waiting time, demonstrating the system's ability to dynamically prioritize critical traffic without disrupting overall efficiency. The standard deviation analysis shows that waiting time for both emergency and regular vehicles becomes more stable as the RL model learns and adapts. This shows that the system not only improves over time but also ensures predictability and consistency in traffic control decisions. The reward function analysis confirms that the RL model continuously improves by learning optimal traffic control strategies. Over time the model converges towards an efficient policy, ensuring that traffic signals operates with maximum effectiveness. The SUMO simulation results confirm that RL-based traffic control outperforms traditional methods in all tested scenarios. The system proves to be scalable and adaptable, making it a viable solution for urban traffic management.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

V. CONCLUSION

This research successfully demonstrates how reinforcement learning, specifically Deep Q-Networks (DQN), can be applied to adaptive traffic management to optimize signal control. The primary goal is to minimize vehicle waiting times by dynamically adjusting traffic signals based on real-time traffic density while ensuring that emergency vehicles experience little to no delay. Using Simulation of Urban MObility (SUMO) as the testing environment real world traffic conditions were simulated to evaluate the performance of our system compared to traditional fixed-time traffic signal control methods. The findings indicate that reinforcement learning provides a significant improvement in traffic flow efficiency. The Deep Reinforcement Learning model (DRL) based model, trained using Bellman's equation, continuously learns and adapts to changing traffic patterns, making real-time decisions that outperform static signal timing strategies. The system prioritizes emergency vehicles, enabling them to navigate through intersections without unnecessary halts, which can be lifesaving in critical situations. By implementing Python-based reinforcement learning techniques, the decision-making process of traffic lights was successfully optimized. The simulation results show a clear reduction in average waiting time for regular vehicles while achieving near-zero delay for emergency vehicles. It has broader applications which include:

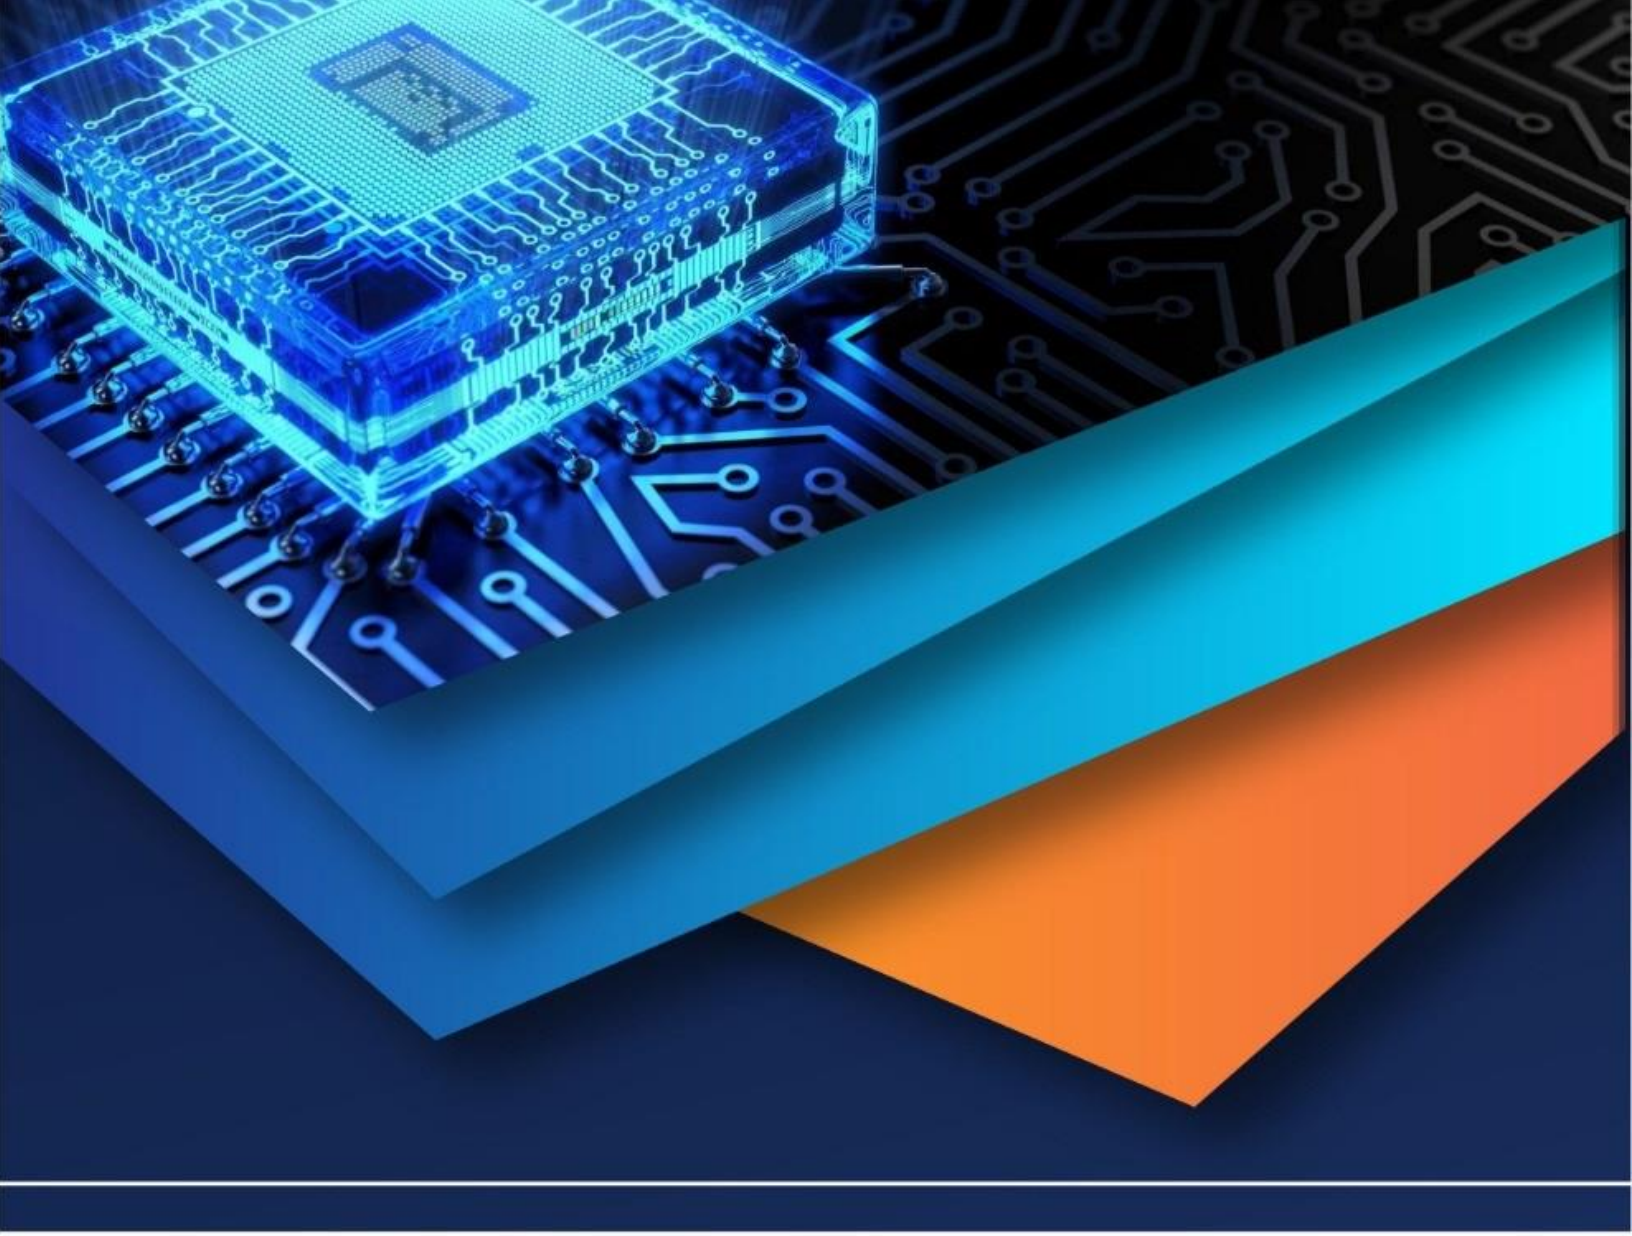
- Reduced fuel consumption.
- Lower greenhouse gas emissions.
- Improved public safety

Real world deployment would require addressing additional complexities such as pedestrian crossings, unpredictable traffic incidents, varying weather conditions and multimodal transportation systems. In summary, this research demonstrates the potential of AI-powered adaptive traffic management in developing more intelligent and efficient urban transportation networks. Continuous refinement and expansion of this approach can help cities transition toward advanced traffic systems that improve mobility, alleviate congestion, and enhance overall infrastructure efficiency.

REFERENCES

[1] "Adaptive Traffic Signaling Control Using SUMO Simulator", Taha Abdulwahid, Muzamil Eltejani Mohammed, IJISRT, January 2025.
[2] "AI Based Traffic Management System" Bibek Shrestha, Bishal Singh, Gaurav Darlami, Nabin Upadhayaya, IJEPE, 2025.
[3] "Machine Learning Framework for sustainable traffic management and safety", Ali Louati, December 2024.
[4] "Pri-DDQN: Learning Adaptive Traffic Signal Control Strategy through a hybrid agent, Yanliu Zheng, Juan Luo, Han Gao", November 2024.
[5] "Adaptive Traffic Signal's Safety and Efficiency Improvement by Multi-Objective Deep Reinforcement Learning Approach", Shahin Mirbakhsh, Mahdi Azizi, IJIRME, July 2024.
[6] "A Deep Reinforcement Learning Approach for Adaptive Traffic Routing in Next-gen Networks", Akshita Abrol, Purnima Murali Mohan, Tram Truong Hou, A*Star, February 2024.
[7] "Adaptive Traffic Lights Based on Traffic Flow Prediction using Machine Learning Models", Idriss Moumen, Jaafar Abouchabaka, IJECE, October 2023.
[8] "Traffic Light Control with Reinforcement Learning", Taoyu Pan, Cambridge A Level Centre, August 2023.
[9] "Attacking Deep Reinforcement Learning Based Traffic Signal Control Systems with Colluding Vehicles", Ao Qu, Yihong Tang, Wei Ma, AAAI, November 2021.
[10] Reinforcement Learning forLogistics and Supply Chain Management:Methodologies, State of the Art,and Future Opportunities, Yimo Yan, Chin Pang Ho, Yong-Hong Kuo, 2022.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |



10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 (24*7 Support on Whatsapp)