Project 2322 | June 2025

SJSU SAN JOSÉ STATE UNIVERSITY

MINETA TRANSPORTATION INSTITUTE

Addressing Urban Traffic Congestion: A Deep Reinforcement Learning-Based Approach

Tairan Liu, PhD

California State University Transportation Consortium

CALIFORNIA STATE UNIVERSITY LONG BEACH

MINETA TRANSPORTATION INSTITUTE

Founded in 1991, the Mineta Transportation Institute (MTI), an organized research and training unit in partnership with the Lucas College and Graduate School of Business at San José State University (SJSU), increases mobility for all by improving the safety, efficiency, accessibility, and convenience of our nation's transportation system. Through research, education, workforce development, and technology transfer, we help create a connected world. MTI leads the Mineta Consortium for Emerging, Efficient, and Safe Transportation (MCEEST) funded by the U.S. Department of Transportation, the California State University Transportation Consortium (CSUTC) funded by the State of California through Senate Bill 1 and the Climate Change and Extreme Events Training and Research (CCEETR) Program funded by the Federal Railroad Administration. MTI focuses on three primary responsibilities:

Research

MTI conducts multi-disciplinary research focused on surface transportation that contributes to effective decision making. Research areas include: active transportation; planning and policy; security and counterterrorism; sustainable transportation and land use; transit and passenger rail; transportation engineering; transportation finance; transportation technology; and workforce and labor. MTI research publications undergo expert peer review to ensure the quality of the research.

Education and Workforce Development

To ensure the efficient movement of people and goods, we must prepare the next generation of skilled transportation professionals who can lead a thriving, forward-thinking transportation industry for a more connected world. To help achieve this, MTI sponsors a suite of workforce development and education opportunities. The Institute supports educational programs offered by the Lucas Graduate School of Business: a Master of Science in Transportation Management, plus graduate certificates that include High-Speed and Intercity Rail Management and Transportation Security Management. These flexible programs offer live online classes so that working transportation professionals can pursue an advanced degree regardless of their location.

Information and Technology Transfer

MTI utilizes a diverse array of dissemination methods and media to ensure research results reach those responsible for managing change. These methods include publication, seminars, workshops, websites, social media, webinars, and other technology transfer mechanisms. Additionally, MTI promotes the availability of completed research to professional organizations and works to integrate the research findings into the graduate education program. MTI's extensive collection of transportation-related publications is integrated into San José State University's world-class Martin Luther King, Jr. Library.

Disclaimer

The contents of this report reflect the views of the authors, who are responsible for the facts and accuracy of the information presented herein. This document is disseminated in the interest of information exchange. MTI's research is funded, partially or entirely, by grants from the U.S. Department of Transportation, the U.S. Department of Homeland Security, the California Department of Transportation, and the California State University Office of the Chancellor, whom assume no liability for the contents or use thereof. This report does not constitute a standard specification, design standard, or regulation.

Report 23-13

Addressing Urban Traffic Congestion: A Deep Reinforcement Learning-Based Approach

Tairan Liu, PhD

June 2025

A publication of the
Mineta Transportation Institute
Created by Congress in 1991

College of Business
San José State University
San José, CA 95192-0219

TECHNICAL REPORT DOCUMENTATION PAGE

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
25-13		

4. Title and Subtitle Addressing Urban Traffic Congestion: A Deep Reinforcement Learning-Based Approach	**5. Report Date** June 2025 **6. Performing Organization Code**
7. Authors Tairan Liu, PhD	**8. Performing Organization Report** CA-MTI-2322
9. Performing Organization Name and Address Mineta Transportation Institute College of Business San José State University San José, CA 95192-0219	**10. Work Unit No.** **11. Contract or Grant No.** SB1-SJAUX_2023-26
12. Sponsoring Agency Name and Address State of California SB1 2017/2018 Trustees of the California State University Sponsored Programs Administration 401 Golden Shore, 5th Floor Long Beach, CA 90802	**13. Type of Report and Period Covered** **14. Sponsoring Agency Code**

15. Supplemental Notes
10.31979/mti.2025.2322

16. Abstract
In an innovative venture, the research team embarked on a mission to redefine urban traffic flow by introducing an automated way to manage traffic light timings. This project integrates two critical technologies, Deep Q-Networks (DQN) and Auto-encoders, into reinforcement learning, with the goal of making traffic smoother and reducing the all-too-common road congestion in simulated city environments. Deep Q-Networks (DQN) are a form of reinforcement learning algorithms that learns the best actions to take in various situations through trial and error. Auto-encoders, on the other hand, are tools that help simplify complex data, making it easier for the DQN to understand and make decisions. To enhance the accuracy of these decisions, the research team chose average vehicle speed as a crucial indicator of traffic flow and employed HyperOPT, a method for fine-tuning the system's hyper-parameters. The team put their method to the test in three different traffic scenarios: controlling a single intersection, managing multiple intersections, and overseeing protected left-turn signals. The results were clear and promising. The innovative system significantly improved traffic conditions by either reducing the average wait time at lights or increasing the overall speed of vehicles passing through intersections. This research not only presents a leap forward in traffic management but also offers a glimpse into a future where road congestion could be significantly alleviated. By employing cutting-edge AI and data processing techniques, the project stands as a testament to the potential for smart cities where traffic flow is optimized, making commutes faster and safer for everyone.

17. Key Words	18. Distribution Statement
Traffic signal timing, left turns, advanced traffic management systems, optimization, vehicle to everything communications.	No restrictions. This document is available to the public through The National Technical Information Service, Springfield, VA 22161.

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
Unclassified	Unclassified	51	

Form DOT F 1700.7 (8-72)

Copyright © 2025

by **Mineta Transportation Institute**

All rights reserved.

DOI: 10.31979/mti.2025.2322

Mineta Transportation Institute
College of Business
San José State University
San José, CA 95192-0219

Tel: (408) 924-7560
Fax: (408) 924-7565
Email: mineta-institute@sjsu.edu

transweb.sjsu.edu/research/2322

ACKNOWLEDGMENTS

This study was supported by the CSU Transportation Consortium and the California State University, Long Beach (CSULB) TRANSPORT 2023 program. Any opinions, findings, conclusions, and recommendations expressed in this material are those of the author and do not necessarily reflect the views of these institutes. The author would like to thank his students, Anurag Balakrishnan, Satyam Pathak, and Pedro Herrera, from the Department of Computer Engineering and Computer Science; and Avi Jagdish, from the Department of Mechanical and Aerospace Engineering, all at CSULB, for working on this research project, implementing and testing the proposed methods, and collecting all the data. The author would also like to thank the Mineta Transportation Institute's staff, Executive Director Dr. Karen Philbrick, and Deputy Executive Director Dr. Hilary Nixon.

Contents

Acknowledgments vi

List of Figures ix

List of Tables x

Executive Summary 1

1. Introduction 3

    1.1 Background and Motivation 3

    1.2 Traditional Approaches to Alleviating Traffic Congestion 5

    1.3 Role of Traffic Management System in Alleviating Traffic Congestion 7

    1.4 Role of V2X Technology in Alleviating Traffic Congestion 8

    1.5 Research Objectives 9

    1.6 Scope of Work 9

    1.7 Overview of the Report 10

2. Methodologies 11

    2.1 Fundamental Methods 11

3. Case Studies: Applying Proposed Method on Traffic Signal Timing Optimization 16

    3.1 Scenario 1: Single-Intersection Traffic Light Control 16

    3.2 Scenario 2: Multi-Intersection Traffic Light Control 20

    3.3 Scenario 3: Dynamic Signal Timing Optimization for Left-Turn Maneuvers via Vehicle-Infrastructure Communication 24

4. Summary and Conclusions 34

Bibliography 35

About the Author 40

LIST OF FIGURES

Figure 1. An Illustrative Image of a Road Intersection 5

Figure 2. Workflow of the Proposed Method 15

Figure 3. Two-Way Single-Intersection Road Network 17

Figure 4. Comparative Analysis Between the Baseline Method and Our DQN + Auto-Encoder Approach 18

Figure 5. A Four by Four Road Network Used in Simulation 20

Figure 6. The System Mean Waiting Time Quickly Converges in the Training Process of the Multiple-Intersection Scenario 21

Figure 7. System's Mean Waiting Time With Fixed-Timing Controller 22

Figure 8. System's Mean Speed with Fixed-Timing Controller 22

Figure 9. System's Mean Waiting Time with Our Trained Controller 23

Figure 10. System's Mean Speed with Fixed-Timing Controller 23

Figure 11. Concept of a Platoon of Automated Vehicles Relaying Information 26

Figure 12. Diagram of the Deep Q-Network Architecture for Left-Turn Maneuvers 27

Figure 13. Sketch of the Two-Way Single-Intersection Used in Protected-Left Turn Simulation 31

Figure 14. System Mean Speed of Each Episode in the Learning Process 32

Figure 15. Reward Function of the Learning Process 32

Figure 16. System Mean Speed of Non-Algorithmic Method and Our Method 33

List of Tables

Table 1. Average Waiting Time at a Two-way Single Intersection for Different Methods... 19

Table 2. Configuration of Traffic Flows in the Simulation.. 31

Executive Summary

California has been in a transportation crisis for a long while. In this crisis, the economy is significantly affected due to the substantial impact on goods transportation and residents' hindered mobility and safety. Two urgent issues in the transportation crisis are congestion and traffic safety. A common situation is that congestion leads to traffic accidents, and traffic accidents, in turn, cause more severe congestion. The dynamic nature of traffic requires traffic management solutions that can monitor, analyze, and intervene in real time. The objection of this project is to tackle transportation challenges with the help of ad hoc networks of vehicles and roadside infrastructure so as to achieve efficient, real-time, intelligent traffic management..

The research team proposed an approach to controlling traffic signals that makes them dynamically adaptable to varying traffic conditions such that traffic congestion can be significantly alleviated. This system utilizes the data collected from existing infrastructures, including sensors and cameras, to make informed decisions about traffic signal phases. By utilizing deep reinforcement learning algorithms, especially deep Q-network (DQN) algorithms, coupled with an auto-encoder, the system can dynamically adjust signal timings in real time to improve traffic flow and efficiency.

As a result of this work, the research team has made the following discoveries:

- The deep reinforcement learning method can significantly improve traffic flow in multiple evaluation metrics.
- The proposed DQN + auto-encoder approach with the HyperOPT optimizer can quickly find the optimal solution for controlling traffic signals.
- The selection of the reward function for the deep Q-network algorithm can significantly affect its performance and impact multiple evaluation metrics.
- Coordinating multiple intersections that utilize the DQN algorithm can achieve more improvements than a single intersection case.

This project has led to three peer-reviewed papers included in conference proceedings:

- Anurag Balakrishnan, Satyam Pathak, Pedro Herrera, and Tairan Liu. "Addressing Urban Traffic Congestion: A Hybrid DQN-Autoencoder Model with HyperOPT Tuning." In *International Conference on Transportation and Development 2024*, pp. 739–749. 2024.
- Pedro Herrera, Avi Jagdish, and Tairan Liu. "Dynamic Signal Timing Optimization for Left-Turn Maneuvers via V2I and V2V Communication." IEEE International Conference on Green Energy and Smart Systems.

- Avi Jagdish and Tairan Liu. “Impact of Reward Function Selection on DQN-Based Traffic Signal Control.” *IEEE International Conference on Green Energy and Smart Systems.*

1. Introduction

1.1 Background and Motivation

Urban traffic congestion remains one of the most pressing challenges facing cities worldwide, profoundly impacting the economy, environment, public health, and overall quality of urban life. The persistent traffic jams on city roads not only symbolize the complexity of urban mobility but also spotlight the intricate relationship between transportation infrastructure and urban development. City planners, traffic engineers, and technological innovators have long struggled with congestion, seeking effective strategies to untangle the knots of stalled traffic that frustrate commuters, impede emergency services, and pollute city environments. Despite advances in technology and urban design, the problem persists, underscoring the need for innovative and integrated solutions that can adapt to the dynamic nature of urban growth and mobility patterns.

The adverse effects of traffic congestion extend beyond mere inconvenience. Economically, it represents a significant drain on productivity, as countless hours that could be spent on work or leisure are lost to the inertia of gridlock. The cost of congestion, therefore, is not just measured in the time wasted but also in the increased operational costs for businesses whose deliveries and services are delayed. Small businesses, in particular, suffer from delayed supplies and reduced customer services, further amplifying the economic toll. Environmentally, the implications are equally severe. Congested lines of vehicles emit a vast number of pollutants, contributing to poor air quality and exacerbating public health issues. The health impacts are not limited to respiratory problems alone; the stress and frustration associated with prolonged commutes are also responsible for decreases in mental well-being. Long-term exposure to high levels of pollution can lead to chronic health conditions, placing a further burden on healthcare systems.

Moreover, the challenge of traffic congestion is intertwined with broader urban planning and societal issues. It reflects and exacerbates inequalities, with less affluent communities often bearing the brunt of pollution and having limited access to efficient transportation alternatives. These communities frequently reside in areas with higher traffic volumes and fewer green spaces, compounding the health risks associated with poor air quality. The sprawling expansion of urban areas, driven by the quest for more livable space, further complicates the scenario, stretching the capacity of existing road networks and public transportation systems. This sprawl often leads to a car-dependent culture, increasing the number of vehicles on the road and making it harder to implement public transport solutions that are both efficient and economically viable. Therefore, addressing traffic congestion requires a holistic approach that considers urban design, environmental sustainability, and social equity to create cities that are not only more navigable but also more livable for all residents.

Traditional responses to congestion, while well-intentioned, often address the symptoms rather than the root causes of the problem. Expanding road infrastructure, for example, can lead to induced demand, where increasing road capacity encourages more car use, thus returning congestion levels to their original state or worse. This phenomenon highlights a paradox where efforts to alleviate congestion through road expansion inadvertently exacerbate the problem. Similarly, while enhancing public transportation is crucial, it requires substantial investment and coordination across various stakeholders to increase coverage and service quality to levels that can effectively attract commuters away from private car use. Investments in public transportation need to be strategic, ensuring accessibility, reliability, and affordability, to genuinely compete with the convenience of private vehicles.

These traditional methods, constrained by physical, financial, and societal limitations, underscore the need for innovative solutions that can adapt to and address the dynamic nature of urban traffic. The rapid growth of urban populations and the evolving demands on urban infrastructure call for a shift in how we conceive and manage urban mobility. This situation has set the stage for the exploration of intelligent traffic management systems, which promise a more adaptive, efficient, and sustainable approach to alleviating urban congestion. By leveraging advanced technologies and data analytics, these systems aim to optimize the flow of traffic in real time, reducing bottlenecks and improving the overall efficiency of the urban transportation network. Intelligent traffic management systems can include synchronized traffic signals, real-time traffic monitoring, and dynamic toll pricing, all of which help manage traffic flow and reduce congestion.

As cities continue to grow and evolve, the quest to mitigate traffic congestion becomes increasingly urgent. It is clear that solving this issue requires a multifaceted approach that not only incorporates technological innovation but also rethinks urban planning practices and transportation policies. Urban planners need to consider mixed-use developments that reduce the need for long commutes, promote walking and cycling, and create efficient transit-oriented developments. Policies should encourage carpooling, the use of electric vehicles, and the implementation of congestion pricing to manage demand during peak hours. Furthermore, public engagement and behavior change initiatives are essential to shift cultural attitudes about car ownership and use.

The pursuit of solutions to traffic congestion, therefore, is not just about easing commutes but about envisioning and creating cities that are more livable, equitable, and sustainable for all their residents. This vision encompasses a future where urban environments are designed for people, not just cars. It includes green spaces, clean air, and accessible transportation options that cater to diverse populations. By addressing traffic congestion comprehensively, cities can improve the quality of life for their residents, enhance economic productivity, and contribute to environmental sustainability. Ultimately, the challenge of urban congestion provides an opportunity to reimagine urban living, fostering cities that prioritize human well-being and environmental health.

Figure 1. An Illustrative Image of a Road Intersection

1.2 Traditional Approaches to Alleviating Traffic Congestion

Traffic congestion in urban areas has been a persistent challenge. Many researchers have delved into various methodologies and techniques to address this challenge in the past decades.

One of the most intuitive responses to congestion has been to increase road capacity. This approach is grounded in the basic principle of supply and demand (Lindsey et al., 2001)—expanding roads should theoretically accommodate more vehicles and reduce congestion. However, Duranton and Turner's work (2011) challenges this notion, demonstrating that increased capacity often leads to higher demand, thus failing to alleviate congestion in the long term. Their analysis of American cities underscores the limitations of road expansions as a sustainable solution to traffic congestion.

Enhancing public transportation systems has been widely advocated as a means to reduce reliance on private vehicles and, by extension, congestion. Studies have shown that investments in bus rapid transit (BRT), light rail, and metro systems can significantly shift commuting patterns. Cervero and Guerra (2011) highlight the success of BRT systems in several American cities in improving mobility and reducing traffic volume on roads. However, the effectiveness of public transportation improvements depends heavily on integration with existing urban infrastructure and the overall attractiveness of the services offered.

Congestion pricing involves charging drivers a fee to use certain roads during peak hours, aiming to discourage unnecessary trips and redistribute traffic load (AbuLibdeh, 2017). This approach has been implemented in cities such as Singapore (Li, 2002), London (Litman, 2006), and Stockholm (Börjesson et al., 2012), with varying degrees of success. Eliasson (2014) provides a comprehensive evaluation of Stockholm's congestion charging scheme, noting significant reductions in traffic volumes and improvements in air quality. Nonetheless, the success of congestion pricing is contingent upon public acceptance and the availability of viable transportation alternatives (Selmoune et al., 2020).

Travel demand management (TDM) strategies seek to reduce the need for car travel by promoting alternative modes of transportation, flexible work schedules, and telecommuting (Bianco, 2020; Gärling, 2007). Calthrop et al. (2000) explore various TDM measures, emphasizing the importance of policy incentives and employer participation in reducing peak-hour congestion. While TDM can be effective in certain contexts, its success is highly dependent on cultural factors and the availability of practical alternatives to driving.

Traffic flow management techniques, such as signal timing optimization and the use of intelligent transportation systems (ITS), aim to improve the efficiency of existing road networks. Studies have highlighted the potential of ITS to reduce congestion through real-time traffic monitoring and adaptive signal control (Ren et al., 2016; Mannion et al., 2016). The approaches include but are not limited to traffic adaptive control (Spall, 1997), genetic algorithms (Teklu, 2007; Park, 1999), max pressure strategies (Varaiya, 2013), and adaptive traffic-responsive strategies (Baldi et al., 2015). The use of reinforcement learning (RL) algorithms in traffic management has gained a lot of attention in the past two decades, with pioneering work utilizing RL to minimize overall waiting times by adapting dynamically to fluctuating traffic scenarios (Wiering, 2000). A major leap forward occurred in 2017 when deep learning was integrated with RL algorithms in traffic control, enhancing the capabilities of traffic data processing and signal management (Gao et al., 2017). This integration, especially with advanced technologies such as graph convolutional neural networks, has facilitated a more in-depth understanding of traffic dynamics across various intersections. A graph convolutional neural network approach was developed to automatically extract features of multi-intersection road networks instead of manually obtaining these features (Nishi, 2018). Lin (2018) introduced an RL algorithm that is capable of tackling complex traffic problems at multiple intersections, which makes substantial progress towards practical traffic management problems. A method combining speed guidance systems with RL-based traffic signal control, tailored for future traffic conditions, was introduced to minimize the total queue length and reduce traffic congestion (Maadi, 2022). Tunc (2023) introduced a novel approach, which blends fuzzy logic with deep Q-learning, merging the adaptability of RL with the robustness of fuzzy logic, to address the challenges of complexity and uncertainty in traffic networks (Tunc, 2023). A cooperative multi-objective architecture was recently proposed to address the challenges (e.g., suboptimal, unstable) in RL, as well as link carbon emissions and global traffic throughput (Tang, 2023). In summary, RL-based algorithms have been widely studied, which provide a significant advantage over traditional approaches in terms of handing complex traffic problems.

Vehicle-to-Everything (V2X) communication technology represents a significant leap forward in the quest to mitigate urban traffic congestion and enhance road safety (Zhang et al., 2018). The recent progress in sensing, communication, and computing technologies has paved the way for a viable solution to the challenge of vehicle coordination. V2X technology enables vehicles to interact with a variety of entities in their surroundings, making this coordination feasible. Key components of V2X technology encompass vehicle-to-vehicle (V2V), vehicle-to-infrastructure (V2I), vehicle-to-device (V2D), and vehicle-to-cloud (V2C) (Zhang, 2018). By gathering data from different road elements, vehicles can make better-informed choices, such as modifying speed for traffic optimization, reacting to changes in traffic signals, and maneuvering safely in complex road scenarios. V2X is instrumental in enhancing road safety, offering preemptive alerts about risks such as abrupt traffic deceleration or nearby emergency vehicles. Additionally, by promoting smoother traffic flow and improved control of traffic lights, V2X can effectively alleviate congestion and enhance travel efficiency. While V2X technology plays a pivotal role in the evolution of intelligent transportation systems, the challenge of efficiently utilizing collected data for vehicle coordination is complex. This complexity stems from the necessity to process and interpret large volumes of data in real time for precise and prompt decision-making. Additionally, achieving compatibility and integration of V2X systems across diverse vehicle models and infrastructures presents a considerable challenge. The technology must be adaptable to different traffic conditions and environments, demanding sophisticated algorithms and robust system designs.

1.3 Role of Traffic Management System in Alleviating Traffic Congestion

Traffic Management Systems (TMS) embody a modern approach to resolving the long-existing problem of urban traffic congestion. By integrating state-of-the-art sensory technology, real-time data analytics, and strategic traffic engineering, TMS offers a comprehensive solution to improve traffic flow, reduce bottlenecks, and enhance the overall efficiency of urban transportation networks. The core of TMS's strategy in combating congestion lies in its ability to monitor, analyze, and adapt to traffic conditions on the fly.

One primary method by which TMS alleviates congestion is through real-time traffic monitoring. Utilizing an array of sensors, cameras, and GPS data, TMS can paint a detailed picture of traffic patterns across an urban area, identifying congestion points as they arise. This capability allows for immediate responses to traffic fluctuations, such as rerouting traffic, adjusting signal timings, or dispatching traffic management teams to address issues directly.

Adaptive signal control technology stands out as a hallmark feature of TMS. It allows traffic lights to adjust dynamically based on actual traffic demand rather than on preset schedules. This adaptability significantly reduces unnecessary waiting at intersections, one of the chief contributors to urban congestion. By optimizing signal timings, TMS ensures that traffic moves more smoothly and efficiently, decreasing the overall time vehicles spend idling and contributing to better air quality.

Moreover, TMS employs predictive analytics to forecast traffic conditions, enabling pre-emptive measures to counteract anticipated congestion. For events known to increase traffic volume, such as concerts, sports games, or major public gatherings, TMS can adjust the entire network to accommodate increased demand, thereby smoothing traffic flow and preventing congestion from taking hold.

Incident management is another crucial aspect of TMS. Quick detection and resolution of traffic incidents, such as accidents or breakdowns, are vital in maintaining steady traffic flow. TMS facilitates rapid emergency response and communicates with drivers in real time to advise on alternative routes, greatly minimizing the congestion that accidents can cause.

In essence, TMS leverages sensory technology and data to create a more responsive and efficient urban traffic network. By addressing congestion through real-time monitoring, adaptive signal control, predictive planning, and effective incident management, TMS plays a pivotal role in enhancing urban mobility around the globe.

1.4 Role of V2X Technology in Alleviating Traffic Congestion

V2X significantly contributes to alleviating traffic congestion by facilitating the real-time exchange of traffic information among vehicles and infrastructure. For example, V2I communication allows traffic signals to interact with approaching vehicles, adjusting signal timings based on actual traffic conditions and reducing unnecessary stops. This dynamic adjustment of traffic lights according to real-time demand ensures that intersections, often chokepoints for congestion, become areas of fluid movement.

Moreover, V2V communication can dramatically improve traffic efficiency by enabling vehicles to maintain optimal speeds and safe distances from each other, effectively reducing the stop-and-go driving patterns that exacerbate congestion. By sharing information about speed, direction, and intended actions, vehicles can synchronize their movements, allowing for smoother flows and more cohesive traffic patterns.

V2X also enhances traffic management by providing precise, real-time data on traffic conditions, surpassing the capabilities of traditional sensor-based systems. This rich data environment allows for more accurate traffic forecasting and modeling, enabling city planners and traffic managers to implement more effective congestion mitigation strategies.

Additionally, V2X plays a critical role in incident detection and management. By immediately broadcasting alerts about accidents or road hazards, V2X can swiftly redirect traffic away from affected areas, minimizing the cascading effect of congestion that such incidents often trigger.

As urban areas continue to grow and evolve, the implementation of V2X communication technology stands as a beacon of hope for resolving traffic congestion. Its ability to connect all elements of the urban transportation ecosystem offers a path towards not only more efficient and less congested roads but also safer and more sustainable urban environments. Through the innovative use of V2X, the future of urban mobility looks brighter, promising a smoother journey for all road users.

1.5 Research Objectives

The overarching goal of our research is to explore and quantify the effectiveness of Traffic Management Systems (TMS) and Vehicle-to-Everything (V2X) communication technologies in mitigating urban traffic congestion. By delving into these advanced transportation solutions, the research aims to provide a comprehensive understanding of their potential impacts, identify better practices for implementation, and offer actionable insights for urban planners and policymakers. The specific objectives of this research are as follows:

1. Evaluate the Current State of Traffic Congestion: The first objective is to assess the extent and nature of traffic congestion in selected urban areas. This involves collecting and analyzing data on traffic volume, travel times, peak congestion periods, and identifying major congestion hotspots. Understanding the current traffic congestion landscape is crucial for benchmarking and later evaluating the impact of TMS and V2X interventions.

2. Investigate the Effectiveness of TMS in Reducing Congestion: This objective focuses on analyzing how various components of TMS—including adaptive signal control, real-time traffic monitoring, and incident management strategies—contribute to alleviating congestion. The research will examine case studies where TMS has been implemented, evaluating changes in traffic flow, reductions in travel times, and improvements in air quality.

3. Explore the Potential of V2X Communication Technologies: The research aims to explore the role of V2X technologies in enhancing traffic efficiency and safety. This includes examining how V2V, V2I, V2D, and V2C communications can reduce congestion, facilitate smoother traffic flow, and improve emergency response times. Special attention will be given to the integration of V2X with autonomous vehicles and its potential to transform urban mobility.

1.6 Scope of Work

The scope of this research encompasses a comprehensive examination of TMS and V2X communication technologies as innovative solutions to urban traffic congestion. The work is structured to cover a broad range of aspects related to these technologies, from their theoretical underpinnings and technical mechanisms to real-world applications, challenges, and outcomes. Here is a summary of the key areas that will be covered within the scope of this research:

- TMS, especially traffic signal control;
- V2X communication integrated into traffic signal control in solving the congestion challenge.

1.7 Overview of the Report

The report is organized as follows. In Section 1, we present an introduction to the work in the report. The methodology of the work is presented in Section 2. In Section 3, we show the effectiveness of our methodology in three case studies: single-intersection traffic light control, multiple-intersection traffic light control, and protected left-turn traffic light control. Section 4 summarizes and concludes the work in this report.

2. Methodologies

2.1 Fundamental Methods

This report focuses on using reinforcement learning algorithms to optimize traffic signal timings. The aim is to dynamically determine the optimal duration for each phase of the traffic light (green, yellow, red) based on real-time observations, improving vehicular flow and throughput in urban areas. Incorporating average speed into the observation space marks a distinct strategy compared to previous studies. This addition provides the reinforcement learning model with an expanded set of data to learn from and observe. Consequently, this enhancement empowers the model to make more informed and effective decisions relevant to the specific environment it is tasked with in the context of reinforcement learning. HyperOPT was utilized to navigate the hyperparameter space effectively to find the optimal solution. This approach is effective in identifying one set of the settings for our reinforcement learning algorithm, thereby enhancing our capability to refine and advance the traffic signal control system.

Reinforcement Learning

Reinforcement learning is a type of machine learning where an agent learns decision-making by interacting with its environment, executing actions, and receiving feedback in the form of rewards . The central goal of RL is to derive an optimal strategy, termed a policy, instructing the agent on actions to maximize cumulative rewards over time. This process involves crucial components such as the agent (decision-maker), the environment (external influence), states (representations of the environment), actions (agent decisions impacting the environment), rewards (scalar feedback for actions), and the policy (strategic guide for decision-making). The interplay between these elements characterizes the dynamics of reinforcement learning.

Q-Learning

Q-learning aims to learn the optimal action-value function $Q^*(s,a)$, which represents the expected return of taking an action a in the state s and following the optimal policy thereafter. The update rule for Q-learning is: $Q(s,a) \leftarrow Q(s,a) + \alpha[r + \gamma\, max_{a'} Q(s',a') - Q(s,a)]$, where α is the learning rate, r is the immediate reward after taking the action a, γ is the discount factor, s' is the subsequent state after the current state s, and a' is the subsequent action.

Deep Q-Network

We employed the Deep Q-Network (DQN) , an adaptation of the traditional Q-learning algorithm, for our study's objective. DQN seamlessly integrates deep neural networks to approximate Q-values, offering enhanced capabilities in handling large-scale and intricate challenges such as optimizing traffic light control in urban settings. Noteworthy features of DQN include its utilization of a neural network as a function approximator, adeptly generating Q-values for diverse actions in each state. Experience replay is employed, allowing the algorithm to break correlations between consecutive experiences by randomly sampling batches from a replay buffer. DQN further enhances stability by incorporating a target network, periodically updated with the primary network's weights. The training process involves minimizing the squared difference between predicted and target Q-values, encapsulated in the concise loss function $L(\theta) = \mathrm{E}_{(s,ar,s')\sim U(D)}[(r + \gamma\, max_{a'}\, Q(s', a'; \theta^{-}) - Q(s, a; \theta))^2]$. These features make DQN a robust choice for addressing complex problems in reinforcement learning.

Auto-Encoder

Auto-encoders , a neural network variant, specialize in unsupervised learning tasks, notably excelling in data compression and de-noising. Their primary objective is to learn efficient encodings of input data, often employed for dimensionality reduction. The auto-encoder comprises key components: the encoder, responsible for compressing input into a reduced-dimensional latent-space representation; the decoder, which reconstructs the input from this representation; and the bottleneck, representing the minimum dimensionality. The network's performance is evaluated by the reconstruction loss, typically measured using mean squared error, assessing the disparity between the original and reconstructed data. The latent space representation encapsulates the compressed data produced by the encoder. Activation functions, such as ReLU or Sigmoid, introduce the non-linearity crucial for handling complex data structures in encoding and decoding processes.

Our Approach

The foundation of our RL system was formed by combining the capabilities of DQN architecture with auto-encoders. This integration allows the auto-encoder to effectively represent and compress the state space, ensuring that the DQN operates optimally with the environment's high-dimensional data. With this framework established, our focus shifts to adjustments (tunings) to adapt the DQN and auto-encoder model to the specific challenges of our traffic signal control problem.

Observation Space

Each traffic signal agent's observation is stacked as a vector:

$$obs = [p_o, m_g, m_v, \rho_n, Q_n]$$

where:

- p_o is a one-hot encoded vector highlighting the currently active green phase.

- m_g is a binary variable that indicates if the minimum green time duration has been surpassed for the ongoing phase.
- m_v is the average speed of all the vehicles at the intersection.
- ρ_n represents the density of vehicles in an incoming lane n; it is calculated as the ratio of the total number of vehicles in the lane to its maximum capacity.
- Q_n indicates the queue density in an incoming lane n, determined by the proportion of vehicles with a speed below 0.1 m/s to the lane's total capacity.

Action Space

The action space for this reinforcement learning task is discrete and corresponds to selecting the next green phase that will be open for a specified duration. The action space is represented as an integer, where each value from 0 to the number of green phases minus 1 corresponds to a specific green phase. This is used as an index to keep track of the current green phase of the traffic signal. In traffic signal systems, different phases represent different sets of lights being green to control traffic flow. For example, $A = \{1, 2, 3\}$ indicates the set of all traffic signal phases (i.e., 1, 2, and 3 indicating green, yellow, and red, respectively). Taking an action "$a \in A$" means selecting an appropriate traffic signal phase for the next time step, which will be "Keep" or "Switch," such that if the agent chooses a set of non-conflicting phases to be assigned the green light, then it will be "Keep"; otherwise, a mandatory yellow phase will be enforced before the "Switch." This discrete action space allows the agent to control the traffic signal's behavior by determining which direction of traffic should have the right of way at any given moment.

Reward Function

The agent's learning and decision-making are molded by the reward functions:

1. Diff-Waiting-Time Reward: A function evaluating the disparities in waiting times across traffic phases. The aim is to ensure uniform waiting times, thus facilitating smooth traffic flow. The reward is given by:

$$\Delta t(s, a, s') = \sum w_{v_i}^{t-1} - \sum w_{v_i}^{t}$$

Here, $w_{v_i}^{t}$ denotes the waiting time of a vehicle, which is the number of seconds the vehicle has a speed less than 0.1 m/s at an intersection. As the delay increases over time, the reward is always a negative value as we are looking to decrease Δt.

2. Average-Speed Reward: This reward emphasizes the average speed of vehicles, giving preference to scenarios where vehicles maintain a decent pace. This reward is expressed as:

$$m(s, a, s') = \frac{\sum_i v_i}{N}$$

Here, v_i refers to the speed of the vehicles, and N represents the number of lanes.

3. Queue Reward: This reward focuses on curtailing the number of halted vehicles, signaling that traffic is flowing without extended stops, which is given by:

$$Q(s, a, s') = -Q$$

Here, Q represents the number of vehicles that are waiting at the intersection.

4. Pressure Reward: This reward quantifies the difference between outgoing and incoming vehicles, advocating for a balanced distribution of vehicles across lanes. It is expressed in the following form:

$$p(s, a, s') = \sum_{out} v_{out} - \sum_{in} v_{in}$$

Here, v_{in} represents the number of vehicles coming into an intersection, and v_{out} represents the number of vehicles going out of the intersection.

5. Average-Speed-Maximization Reward: This function is intricately tied to the vehicles' average speed, which is given by:

$$\text{max\_}v = \frac{\sum_i v_i}{L}$$

A higher mean speed leads to a better reward, symbolizing efficient traffic movement.

Epsilon-Greedy Strategy with DQN + Auto-Encoder

The intricate combination of our DQN and auto-encoders, augmented by an epsilon-greedy strategy, forms the core of our learning mechanism. Central to this framework is the auto-encoder, which is pivotal in efficiently compressing and representing the state space effectively, equipping the DQN to process high-dimensional data. Upon observing the current state, s, the agent either opts for a random action a with a likelihood ε, or, drawing from its learned experiences, selects the action that maximizes its Q-value: $a^* = argmax_a Q(s, a; \theta)$.

Following the action and transitioning to the new state, s, the agent logs the reward r and refines the DQN parameters: $Q(s, a; \theta) \leftarrow r + \gamma \; max_{a'} Q(s', a'; \theta)$.

The tuning of reward functions, the integration of the epsilon-greedy strategy, and the DQN-auto-encoder architecture together form a robust and adaptable reinforcement learning framework.

In the illustrated framework (as shown in Figure 2), we show the architecture of the RL-based traffic signal control system. The model integrates a DQN with an auto-encoder, a neural network structured to encode and reconstruct high-dimensional input data into a lower-dimensional encoded format. Situated in a standard urban traffic intersection, the agent leverages the auto-encoder to transform complex traffic parameters such as lane density, queue length, and average speed into simplified representations. The DQN then processes this condensed representation, guiding the agent to make informed decisions about adjusting traffic light timings for optimized traffic flow. Upon executing a decision, the agent receives feedback from a reward system that quantifies the effectiveness of its actions based on metrics such as waiting times and vehicle speeds. An auxiliary HyperOPT module fine-tunes the system's hyperparameters, ensuring its precision and adaptability.

Figure 2. Workflow of the Proposed Method

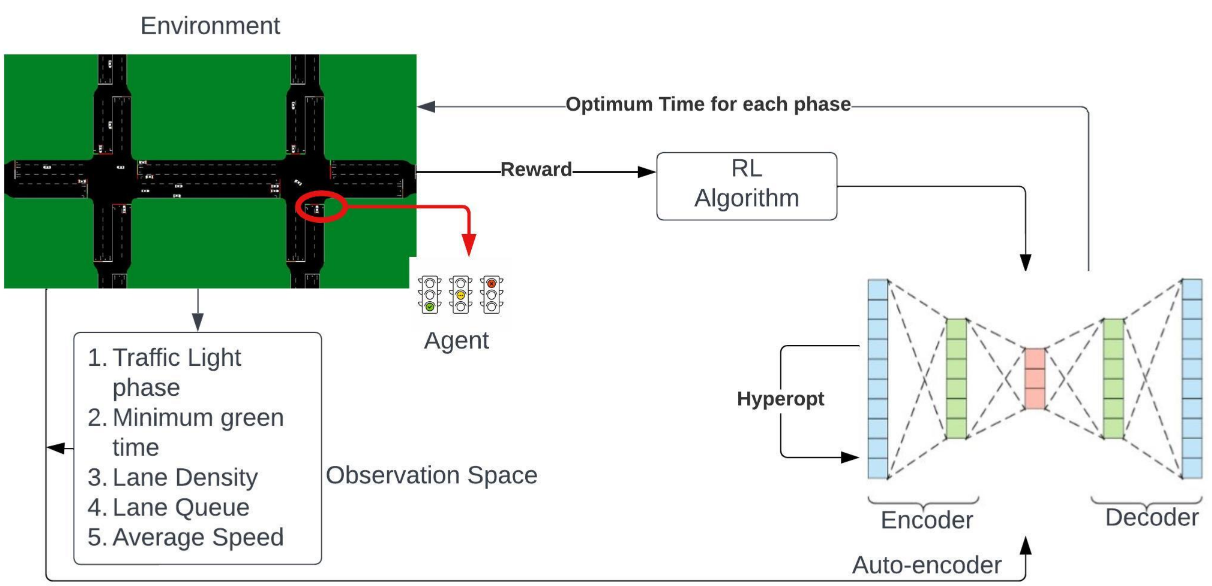

3. Case Studies: Applying the Proposed Method on Traffic Signal Timing Optimization

3.1 Scenario 1: Single-Intersection Traffic Light Control

This section presents a case study of our method as applied to a single intersection scenario to validate the effectiveness of this method.

3.1.1 Simulation Environment Configuration

The simulations for this case study were conducted on a MacBook Pro featuring an M1 chip with an 8-core CPU, 14-core GPU, and a 16-core Neural Engine, coupled with 16 GB of memory. In the simulation, we employed the Simulation of Urban MObility (SUMO) to model a two-way single intersection scenario as shown in Figure 3. With the road network, the flow of vehicles is configured as follows:

- From north to south and from north to west, 300 vehicles per hour were generated.
- From north to east, the flow is slightly lower as 250 vehicles per hour were generated.

For the deep learning aspect, we utilized the capabilities of PyTorch and implemented the DQN algorithm, an established RL algorithm.

To capture an extended temporal perspective, we set the simulation time to 40,000 seconds, allowing us to observe and evaluate traffic dynamics effectively over a prolonged period. In terms of training hyperparameters, we adjusted the learning rate to 0.01, optimizing the learning process for our DQN model.

Figure 3. Two-Way Single-Intersection Road Network

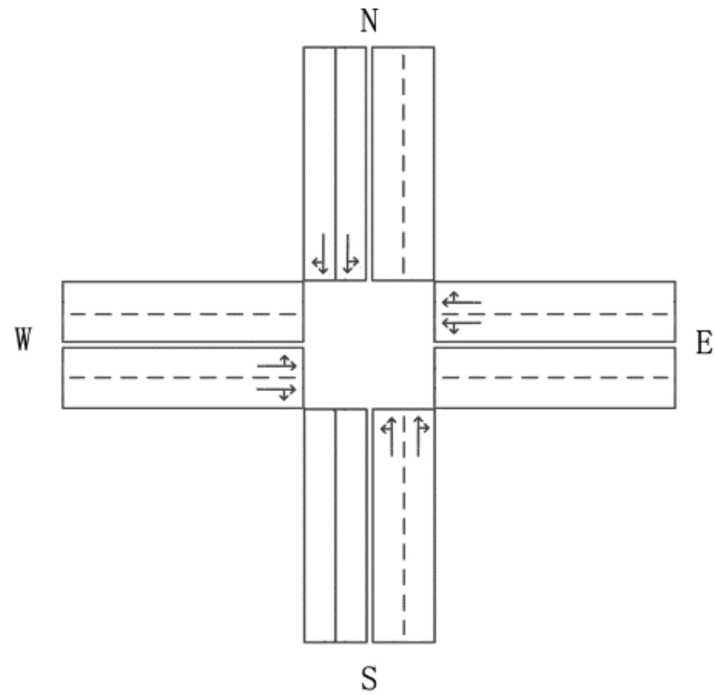

3.1.2 Results and Comparative Analysis

Our extensive analysis, conducted over 400 episodes using the SUMO traffic simulation, focused on evaluating the effectiveness of three distinct methodologies in managing traffic congestion, as measured by the average waiting time of vehicles. These methodologies include a pre-existing implementation of traffic management (DQN), our innovative approach that employs the DQN with an auto-encoder, and a naïve method devoid of deep learning techniques, wherein the traffic signals have been programmed with a fixed time interval for each of the phases. The timings have been set based on real-world traffic signals, where the signals change every 30 seconds.

In our research, we carried out two types of simulation tests using SUMO. The first type of test looked at how well two different traffic control methods improved as they were trained over time. The second type of test checked how reliable these methods were. We ran each method five times in simulations that lasted 40,000 seconds each. This was done in a virtual environment that mimics a two-way road with a single intersection. By doing this, we wanted to make sure that the results we obtained were not just a one-time thing but could be repeated reliably.

Convergence Test

As shown in Figure 4, the baseline method shows a gradual decrease in mean waiting time as the number of episodes increases, suggesting that the method is learning and improving over time. However, the improvement is slow, and the waiting times remain relatively high, averaging around 300 seconds even after 400 episodes. In our test, we noticed that the computational speed of the training process became very slow after the initial convergence. This indicates that while the baseline method is making progress, it is not very efficient. The decision to stop the baseline method training due to slow progress in training the DQN model, with each episode needs several hours. Even though with longer training, the DQN could potentially reach a smaller value of the average waiting time, the slow training speed would request an extended period of time.

On the other hand, the DQN + auto-encoder method shows a very different pattern. Initially, there is a significant spike in mean waiting time, but this quickly drops and begins to stabilize. After this initial fluctuation, the waiting times reduce dramatically to about 186 seconds within 50 episodes, suggesting that the method starts to converge to a more effective strategy much faster than the baseline method. Despite some variability, the general trend for the DQN + auto-encoder method is towards lower waiting times, with episodes leveling out at a much lower mean waiting time compared to the baseline method. This indicates a more efficient learning process and a quicker convergence towards an optimized solution for traffic management.

The sharp peak in the beginning for the DQN + auto-encoder method could be due to the initial exploration phase, which is common in reinforcement learning algorithms where the system tries out different strategies to learn which one works better. After this phase, the algorithm starts to exploit the best-known strategies, leading to the observed improvement and stabilization in performance.

Figure 4. Comparative Analysis Between the Baseline Method and Our DQN + Auto-Encoder Approach

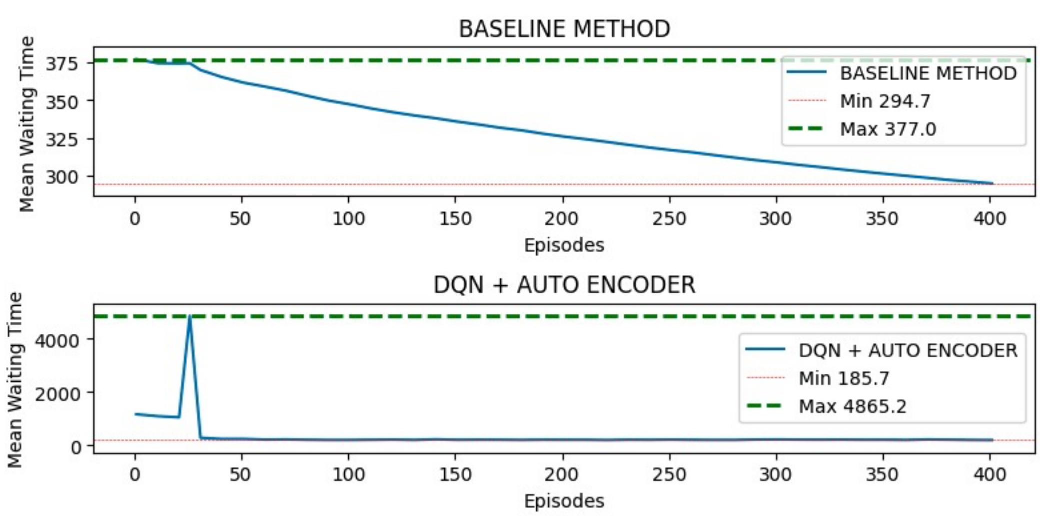

Repeatability Test

The assessment was also carried out through a repeatability test, each conducted 5 times and lasting 40,000 seconds based on the policy of the 400th episode, in a two-way single intersection scenario. The result is shown in Table 1.

Table 1. Average Waiting Time at a Two-way Single Intersection for Different Methods

Method	Mean (seconds)	Standard Deviation (seconds)
Naïve	1132	135
Baseline	340	37
DQN + Auto-Encoder	212	12

The naïve approach registered the highest mean duration at 1,132 seconds, accompanied by a substantial standard deviation of 135 seconds. This significant deviation underscores a variability in performance and shows potential inconsistencies in the application of this method.

The baseline method exhibits a notable enhancement in efficiency, providing a mean duration of 340 seconds, noticeably lower than the naïve approach. Its standard deviation, positioned at 37 seconds, also indicates improved consistency in performance compared to the naïve method.

Our method, the integration of DQN with auto-encoders, outperformed both preceding methods. This approach achieved an impressive mean duration of 212 seconds, significantly reducing the waiting time. Furthermore, it maintained a standard deviation of 12 seconds, the lowest among the evaluated methods. This minimal deviation is indicative of a high level of consistent performance, reinforcing the method's reliability.

3.1.3 Conclusion and Discussion

This work shows that the DQN + auto-encoder approach is effective in reducing the average waiting time of vehicles at intersections. The integration of DQN and auto-encoder not only accelerates the processing time but also ensures consistent performance, proving to be an optimal solution for traffic management in a two-way single intersection environment within SUMO.

The inclusion of average speed into the observation space distinguishes the approach from previous approaches. This integration of a DQN with an auto-encoder represents a considerable advancement in the reinforcement learning models' ability to react, anticipate, and adapt to emerging patterns in traffic congestion management. The model resulted in reduced average waiting times and improved consistency. This method surpassed traditional DQN implementations and significantly outperformed the naïve method.

For future research, the focus could shift to include vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) communications. V2V and V2I communications implemented in the traffic management systems can help create a more connected and efficient transportation network. By enabling vehicles to share important information such as their location, speed, and planned routes in real time, a more cooperative traffic network can be developed. Upcoming studies could also work on creating new algorithms that use V2V and V2I communications to improve how vehicles change lanes and merge on ramps. These improvements are crucial to helping vehicles work together better when merging, which can help keep traffic flowing smoothly and reduce road congestion. Research in this area has the potential to greatly improve the safety and effectiveness of our highway systems. It could lead to better coordination between vehicles, making overall traffic movement smoother and reducing delays.

3.2 Scenario 2: Multi-Intersection Traffic Light Control

This section's objective is to test our method on a multiple-intersection scenario and validate its effectiveness.

3.2.1 Simulation Environment

The simulation environment used in this scenario is a four-by-four road network, as shown in Figure 5.

Figure 5. A Four by Four Road Network Used in Simulation

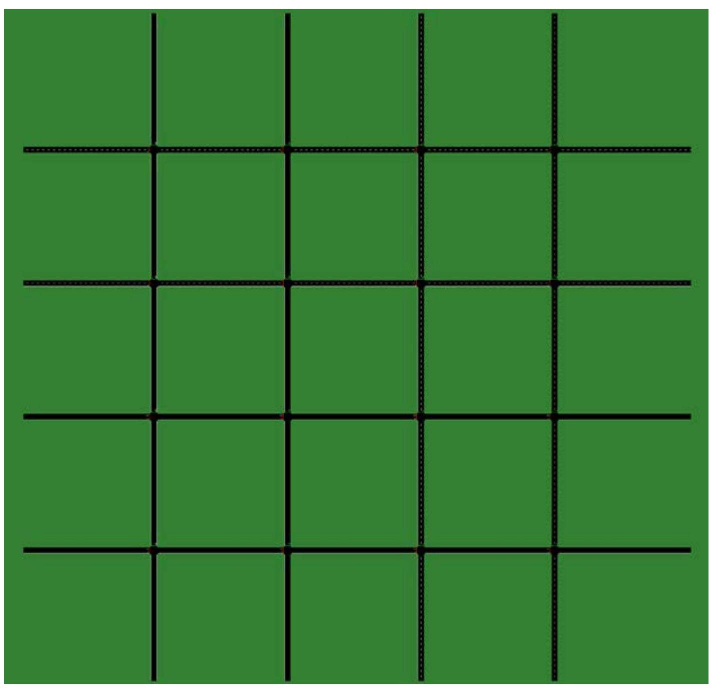

3.2.2 Results and Analysis

Our algorithm, though it was designed for single intersections, was also applied to this four-by-four road network, and we observed a significant improvement in vehicles' waiting time and speed. Besides, the method can also lead to a fast convergence in the training process. As shown in Figure 6, the system's mean waiting time quickly converges with very small fluctuations.

Figure 6. The System Mean Waiting Time Quickly Converges in the Training Process of the Multiple-Intersection Scenario

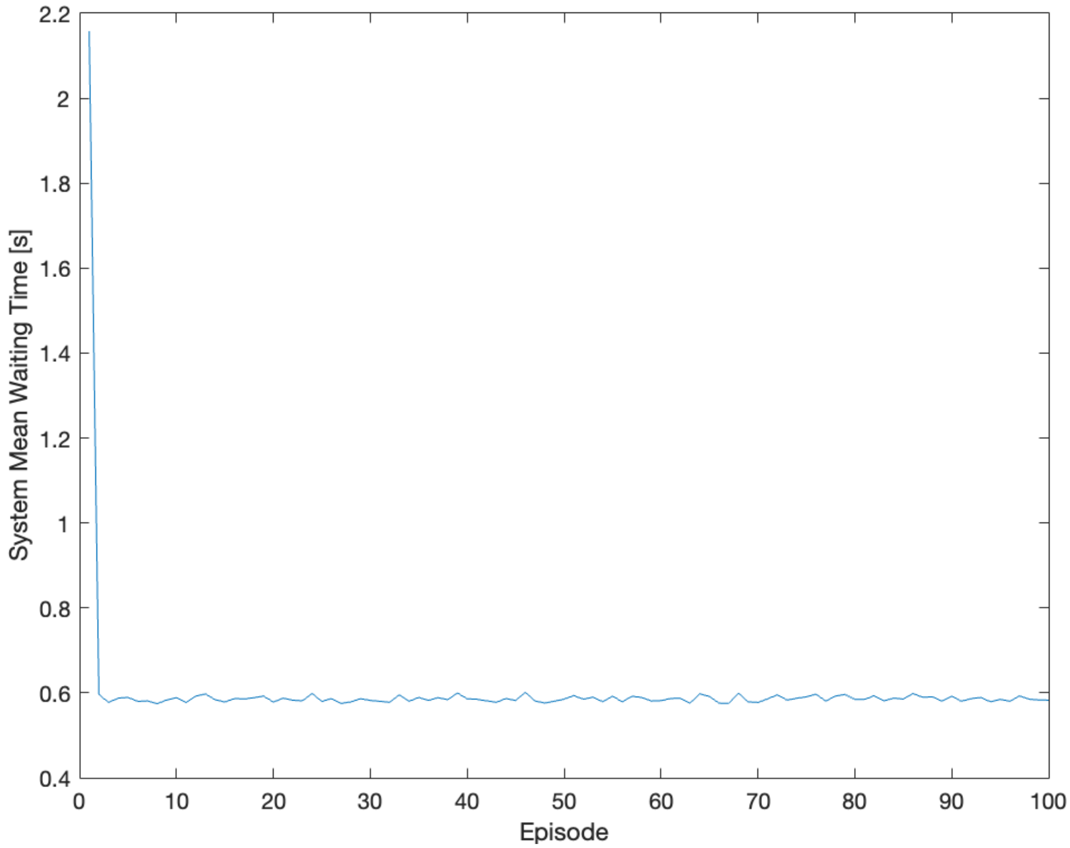

We also compared the performance of the trained controller with a fixed timing controller (a pre-timed traffic signal controller where each phase has a predetermined length) in terms of vehicles' average waiting time and speed. As shown in Figures 7 and 8, vehicles' average waiting time and speed are mostly in the range of 0 to 16 seconds and 4 to 8 m/s, respectively, with the fixed-timing controller. The average value is 7.01 seconds and 6.20 m/s, respectively. With the trained controller, the vehicles' average waiting time and speed can be significantly improved, where the average waiting time was reduced to 0.59 seconds, and the average speed was increased to 7.77 m/s (as shown in Figures 9 and 10). These results indicate that the proposed method also works for multiple intersection scenarios and can achieve better performance compared to the traditional fixed-timing traffic signal controllers.

Figure 7. System's Mean Waiting Time with Fixed-Timing Controller
The Dashed Line is the Average Value

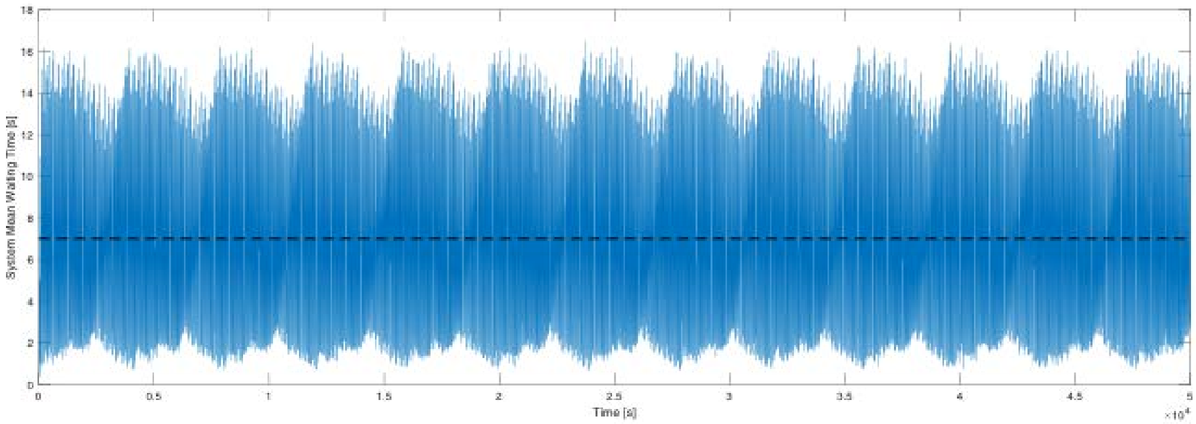

Figure 8. System's Mean Speed with Fixed-Timing Controller

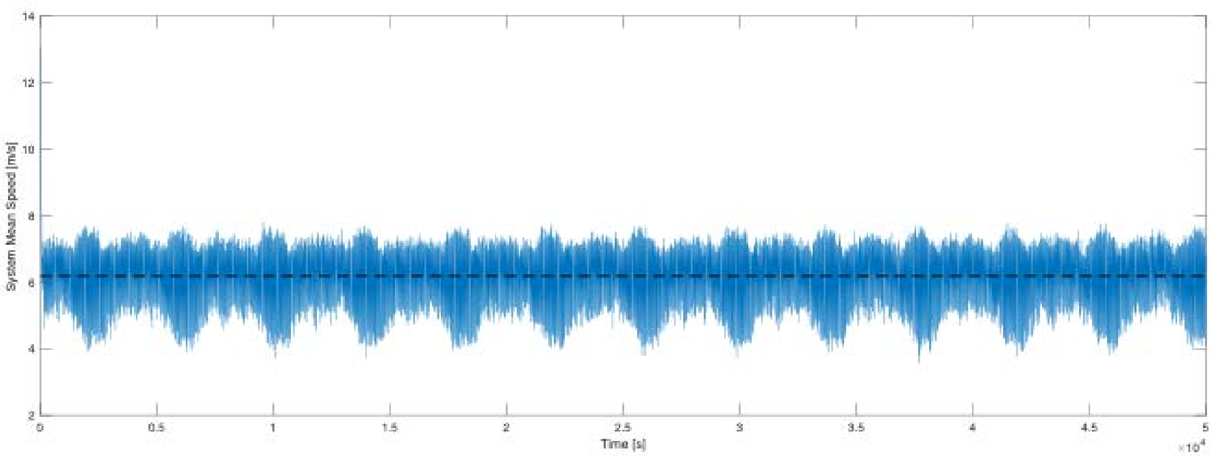

Figure 9. System's Mean Waiting Time with Our Trained Controller
The dashed line is the average value.

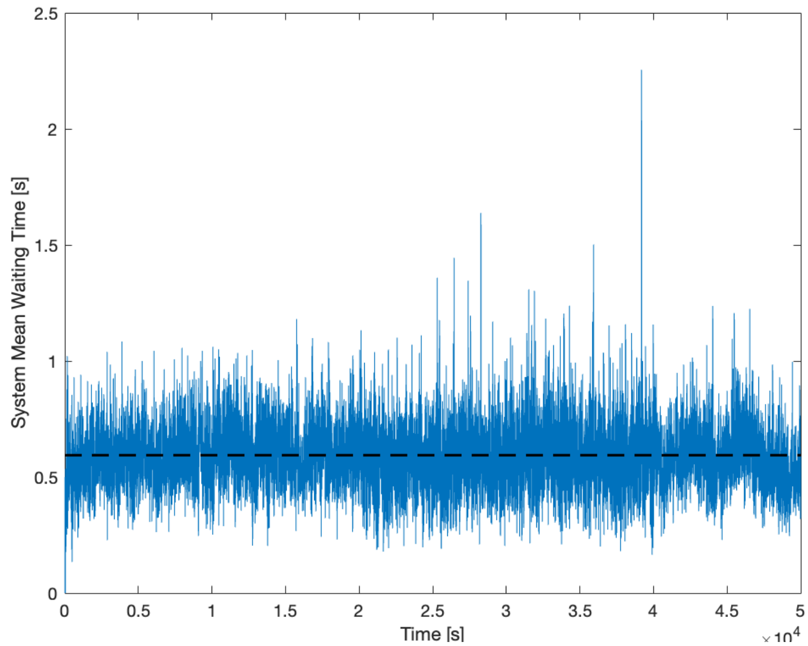

Figure 10. System's Mean Speed with Fixed-Timing Controller
The dashed line is the average value.

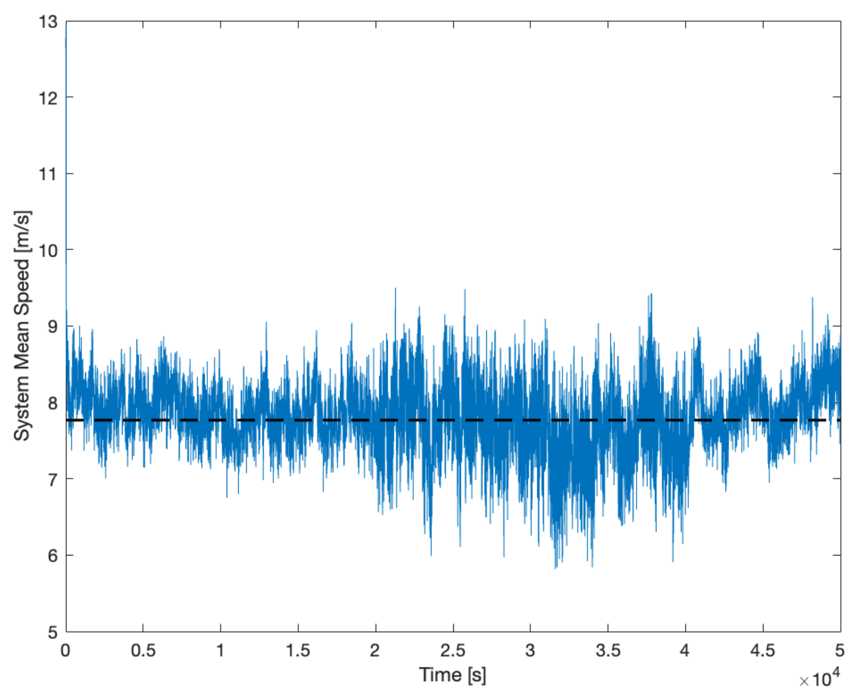

3.3 Scenario 3: Dynamic Signal Timing Optimization for Left-Turn Maneuvers via Vehicle-Infrastructure Communication

3.3.1 Problem Statement of Left Turn Maneuvers

In urban areas with high population density, left-turn movements at intersections present a significant challenge since they cross opposing traffic lanes. This requires drivers to judge the speed and distance of oncoming vehicles and make decisions about the appropriate timing of their turns. Left-turn maneuvers impact the flow of both vehicles and pedestrians, often leading to unsafe situations. Adding the unpredictability of human factors further complicates this scenario. Some examples include drivers under time stress due to traffic signal timing (TST) (Pawar, 2022), impaired visibility, and judgment errors. For instance, challenges arising from impaired visibility account for 20% of the critical reasons for accidents, while judgment errors, such as failing to estimate a safe gap distance and making false predictions about others' actions, account for approximately 34% (National Motor Vehicle Crash Causation Survey, 2008). These factors contribute to the complexity of the decision-making process that results in the execution of a left turn maneuver. Given these factors, it becomes increasingly essential to devise a strategy that effectively coordinates vehicles, particularly in scenarios where automated vehicles (AVs) and human-driven vehicles coexist.

While V2X technology plays a pivotal role in the evolution of intelligent transportation systems, the challenge of efficiently utilizing collected data for vehicle coordination is complex. This complexity stems from the necessity to process and interpret large volumes of data in real time for precise and prompt decision-making. Additionally, achieving compatibility and integration of V2X systems across diverse vehicle models and infrastructures presents a considerable challenge. The technology must be adaptable to different traffic conditions and environments, demanding algorithms and robust system designs.

Several studies have been conducted to address these challenges. Lu (2023) undertook an empirical study using deep learning models to improve the reliability of autonomous driving systems through V2X, particularly in high-speed and dynamic environments. Jiang (2020) studied an improved layout for intersections integrating real-time V2X communication to optimize traffic signal timing and improve traffic flow for left-turn intersections in unsaturated traffic conditions. Zhang (2022b) studied cooperative control methods in intelligent connected vehicle environments to address the challenges in heterogeneous traffic flow. Vehicles equipped with a stack of perception technology use sensor measurements to establish a vehicle's state and environmental information, which is crucial for the safe navigation of intersections (Emamifar, 2023). Regarding traffic safety, Xiang (2022) studied V2X communication-aided message dissemination to resolve the delay of warning information received by vehicles. To tackle the potential privacy concerns in sharing data in V2X communication, Song (2023) proposed a pipeline for applying federated learning for intelligent transportation systems. Du (2023) explored the optimization of left-turn maneuvers with vehicles equipped with safety measuring sensors, seeking to reduce potential collisions among opposing vehicles in the network.

This study proposes an innovative approach that employs V2I and V2V communication to dynamically optimize signal timings for intersections featuring both protected and unprotected left turns. Our work seeks to utilize the information provided by advanced technologies, such as AVs equipped with sensing technology and cooperative communication devices (e.g., cameras, lidar systems, and dedicated short-range communications (DSRC)), to improve visibility and coordination between vehicles. The contributions of this work are as follows.

1. We propose an approach for the implementation of protected left turns by incorporating reinforcement learning algorithms in traffic signal timing for V2I/V2V traffic networks.

2. Our approach improves the efficiency and safety of protected left-turn signal phases while improving traffic flow for all signal phases.

3. We show the effectiveness and applicability of reinforcement learning algorithms in dynamic TST adjustments for improved traffic flow (higher speed).

3.3.2 System Architecture for Detection and Communication in Mixed Traffic Environments

The design and operation of the system are tailored to facilitate interaction and connectivity among infrastructures, connected and automated vehicles (CAVs), and manually operated vehicles in mixed-traffic situations. At the heart of the system's architecture lies the communication module, tasked with managing V2X communications. DSRC components are utilized to transmit information crucial for trajectory planning. This enables CAVs to merge and slow down effectively when approaching intersections. Employing this method of communication has demonstrated enhancements in the performance of individual vehicles and has positively impacted on the overall traffic flow within the network.

Recognizing that conventional vehicles will remain prevalent (Mahdinia, 2021), our system design incorporates these human-driven vehicles into a unified communication network with CAVs. When in proximity to a CAV platoon, human-operated vehicles can engage in the communication chain, using the DSRC module to share traffic information, thereby enhancing collaborative driving dynamics. Our system's architecture is geared towards refining existing frameworks (Eskandarian, 2019) by focusing on detection, communication, and adaptive traffic management strategies. This approach aims to bolster safety and efficiency in an era characterized by a mix of autonomous and human-driven vehicles. Traditionally, platoons were formed by synchronizing traffic signal timings. However, with the advent of AVs equipped with DSRC modules, a more dynamic organization of vehicle batches at intersections is possible, facilitating smoother collective crossing (Gholamhosseinian, 2022). Moreover, the concept of vehicle platooning extends beyond intersections. Organized traffic flows have been observed to alleviate congestion across various road networks, including highways. Additionally, these systems can be tailored to prioritize emergency vehicles, further enhancing traffic management and response times (Michaud, 2006).

Figure 11. Concept of a Platoon of Automated Vehicles Relaying Information

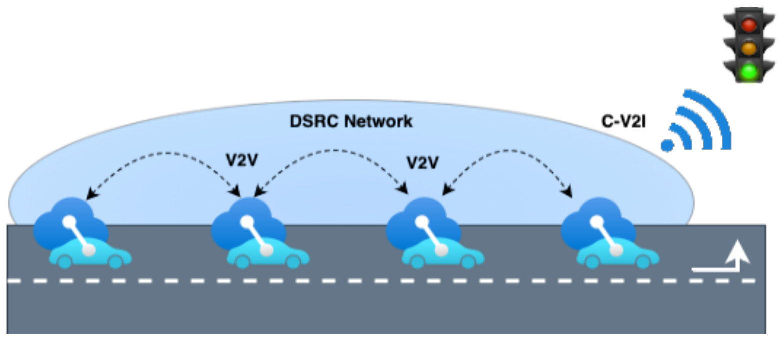

The primary control objective of each vehicle in a platoon is to maintain a safe following distance from the leading vehicle (Timmerman, 2021). A platoon of automated vehicles refers to a group of vehicles, typically equipped with automated driving systems and V2V communication, that travel closely together in a coordinated, cooperative manner, usually in a single lane and often at highway speeds. The concept of a platoon of automated vehicles relaying information is shown in Figure 11.By using the platoon-leading vehicle as a reference, the velocity is set and tracked by the following vehicles, creating a group of vehicles in a platoon of cooperative cruise control (Badnava, 2021). The proposed algorithms synchronize the formation of platooning AVs and adjust individual vehicle behaviors based on real-time traffic conditions, improving traffic congestion, and enhancing the safety of the unprotected left-turn maneuver.

Platoon Merging and Coordination

We propose algorithms that synchronize the formation of platooning autonomous vehicles. The formation of a platoon is defined in the simulation by four types of vehicles. The leading vehicle is defined with a minimum gap distance of 1.5 meters between the trailing vehicles. The three following vehicles are defined with a minimum gap distance of 0.7 m and an increased rate of permitted acceleration by 20%. With varying velocity rates, which are dependent on the vehicle's position in the platoon, the behavior results in a smooth flow of traffic.

This study develops an algorithm that enables the compatibility of vehicles equipped with a DSRC module to merge multi-lane platoons into a single lane consisting of platoons of left-turning vehicles. AVs seeking to join the platoon from an adjacent lane ease congestion by using all available lanes during periods of high congestion. The vehicles seeking to merge abide by protocols set by the algorithm to execute the merging maneuver. The algorithm defines a specific gap between vehicles. This merging scenario of AVs into the turning lane mitigates congestion by using all three available lanes, as opposed to forming a lengthy queue waiting to turn left. In this merging scenario, a zipper merger is executed with multiple entry points along the left-turning lane. Human-driven vehicles in adjacent lanes benefit from the reduction of bottlenecks typical of a first-in-first-out traffic system (Woo, 2021), thereby reaching the intersection more efficiently.

The individual vehicle behaviors of the vehicles in platoons adjust based on real-time traffic conditions. Through broadcasted vehicle data (vehicle speed and vehicle position), vehicle arrivals at intersections arrange themselves through relative position and speed, with speed adjustment as the key feature (Ghoul & Sayed, 2021). Through the training and adjustment, emergent behaviors of the platoons occurred concerning queuing and batching for the crossing maneuver at the intersection. Safe utilization of road space is prioritized, while overall waiting time was reduced.

3.3.3 Modified Method

In this case study, we modified the method proposed in Section 2 specifically for both unprotected and protected left-turning vehicles in a synthetic network road environment using SUMO. The model responds to varying traffic conditions, helping autonomous vehicles in platoons and human-driven vehicles by increasing throughput during protected and unprotected left turn signal phases. The architecture of the Deep Q-network is shown in Figure 12.

Figure 12. Diagram of the Deep Q-Network Architecture for Left-Turn Maneuvers

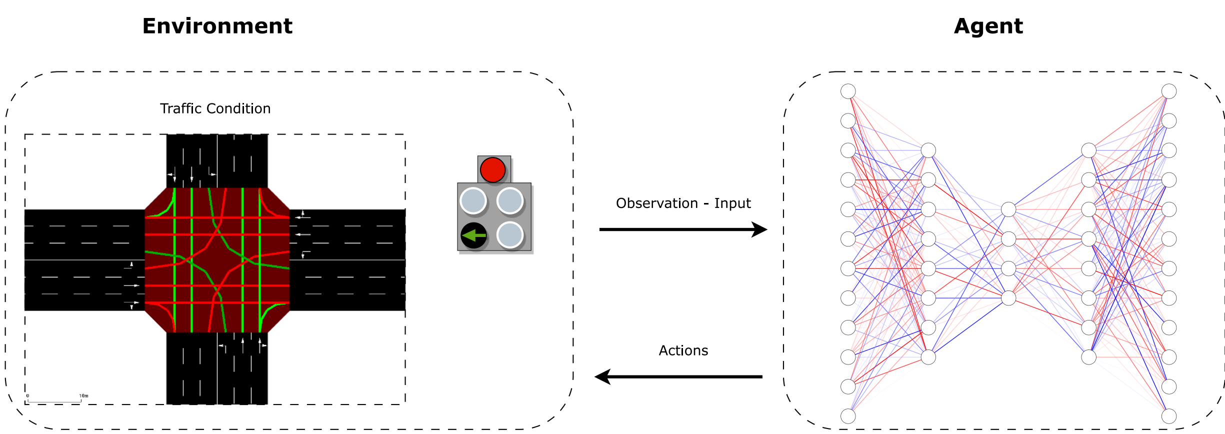

The DQN model continuously analyzes real-time traffic data, taking inputs on the presence and characteristics of the vehicles that are part of a platoon. Using V2X communication, platoons are detected and then prioritized during the unprotected signal phase. Platoon movement is streamlined, leading to improvements in traffic flow and time metrics. The DQN's memory feature enables adaptations over time, fine-tuning the decision-making process based on traffic data and time metrics. The improvements in traffic pattern recognition optimized the management of traffic signal timings.

State Space

The state space represents the different parameters and conditions of the traffic environment at a given time. This involves analyzing the positions and velocities of individual vehicles, the state and timing of traffic lights, tracking the presence of vehicle platoons, and additional measurements such as vehicle count, density, and the flow rate of traffic segments traveling in different directions. Each state of space S per traffic simulation episode can be defined as the set of all possible states. Each state $s \in S$ is a vector representing different aspects and parameters of the traffic environment at a given time. A simplified model, where states are defined by the positions and velocity of each vehicle, can be represented as (x, y, v), where x and y are the coordinates of each vehicle and v is the speed. The complete state space is the cartesian product of all individual vehicle states represented as $S = \{(x_1, y_1, v_1), (x_2, y_2, v_2), \dots, (x_n, y_n, v_n)\}$, where n is the number of vehicles in the simulation.

Action Space

In the SUMO environment, a component utilized by the reinforcement agent decision-making process is the action space. The space defines the set of actions the algorithm can take and adjust to influence the traffic environment. The agent can decide the duration of each signal timing (green, red, or yellow), adjusting based on the network traffic conditions. Algorithm 1 (next page) determines this action by creating custom signal phases prioritizing platooning vehicles, improving efficiency and safety for the platoon by creating a longer permissive turning phase for the vehicles. By broadcasting data to nearby vehicles, movements are coordinated by the provided updates. By adjusting parameters for vehicle platoons, behaviors of individual vehicles are altered for route choice.

Reward Function

The reward function is given by (Alegre, 2019)

$$r_{\boldsymbol{t}} = -(\Delta D_t) = D_t - D_{t+1}$$

where $r_{\boldsymbol{t}}$ is the reward at the time step t. ΔD_t is the change in cumulative vehicle delay. The reward at each time step is based on the change in cumulative delay relative to the previous time step $t - 1$. $\Delta D_t = D_{t+1} - D_t$. This approach rewards actions that reduce the overall waiting time of vehicles in the network. If the cumulative delay decreases (i.e., traffic flow improves), the agent receives a positive reward, encouraging it to take similar actions in future time steps. Conversely, if the cumulative delay increases, a negative reward is received, discouraging it from taking similar actions in the future. The agent's policy is continuously updated based on these rewards to optimize traffic flow over time.

Algorithm 1. Pseudocode for TST Control for Platoons

Function adjust_traffic_signals_for_platoons(vehicle_ids)

Initialize an empty list: platooning_vehicles_for_left_turn

For Each vehicle_id **in** vehicle_ids

If is_platooning_vehicle(vehicle_id) and is_making_left_turn(vehicle_id)

Add vehicle_id to platooning_vehicles_for_left_turn

If platooning_vehicles_for_left_turn **is not empty**

// Set a custom traffic signal phase for platooning vehicles

// This phase gives protected left turns to north and south directions

Return *"rrrGrrrrrrrGrrrr"* # 'G' for green in specific lanes, *'r'* for red in others

Else

// Set the default traffic signal phase

// This phase allows regular traffic flow with unprotected left turns

Return *"GGGgrrrrGGGgrrrr"* # Regular traffic signal phase

End Function

Neural Network Layers

The neural network structure was defined with an auto-encoder for dimensionality reduction of traffic data. This unconventional method has two main parts: an encoder and a decoder. The input captures non-linear features of the traffic data, such as traffic flow patterns, and congestion dynamics.

Encoder Layer

$$E(x) = f_a(f_{BN}(W_2 f_a(f_{BN}(W_1 x + b_1)) + b_2)),$$

where W_1 and W_2 are the weights of the neural network, b_1 and b_2 are the biases, f_{BN} represents batch normalization, and f_a is the ReLU activation function.

Decoder Layer

$$D(E(x)) = f_a(f_{BN}(W_4 f_a(f_{BN}(W_3 E(x) + b_3)) + b_4)),$$

where W_3 and W_4 are the weights of the neural network, and b_3 and b_4 are the biases.

Batch normalization helps normalize the input layer by adjusting and scaling the activations. The formula is as follows: $\hat{x}_i \leftarrow \frac{x_i - \mu_{\mathrm{B}}}{\sqrt{\sigma_{\mathrm{B}}^2 + \varepsilon}}$, where x_i is the input to a neuron, μ_{B} is the mean of a batch, σ_{B}^2 is the variance of the batch, and ε is a small number to prevent division by 0.

3.3.4 Configuration of Simulation Environment

In the simulated road network (as shown in Figure 13), we configured vehicle flow to emulate a congested city by adjusting probabilities of vehicle generation over a set period of 5,000 seconds. The probability of the routes each vehicle took is shown in Table 2. Vehicles with a 20% probability of being generated travel horizontally on the route and do not turn left. The route directions with the protected and unprotected left turn maneuver travel from north to east and south to west. The most congested portion of the road network is the lane traveling from south to west, consisting of vehicles in platoons turning left and vehicles in adjacent lanes merging into the turning lane. 30% of vehicles in the simulation are part of a platoon, and the other 70% use a human-driven model for control. The platoons use a specialized flow where the leader and following vehicles are defined independently. This configuration was designed to serve as a stress test for the algorithms, with the objective of providing insights into real-world applicability.

Figure 13. Sketch of the Two-Way Single-Intersection Used in Protected-Left Turn Simulation

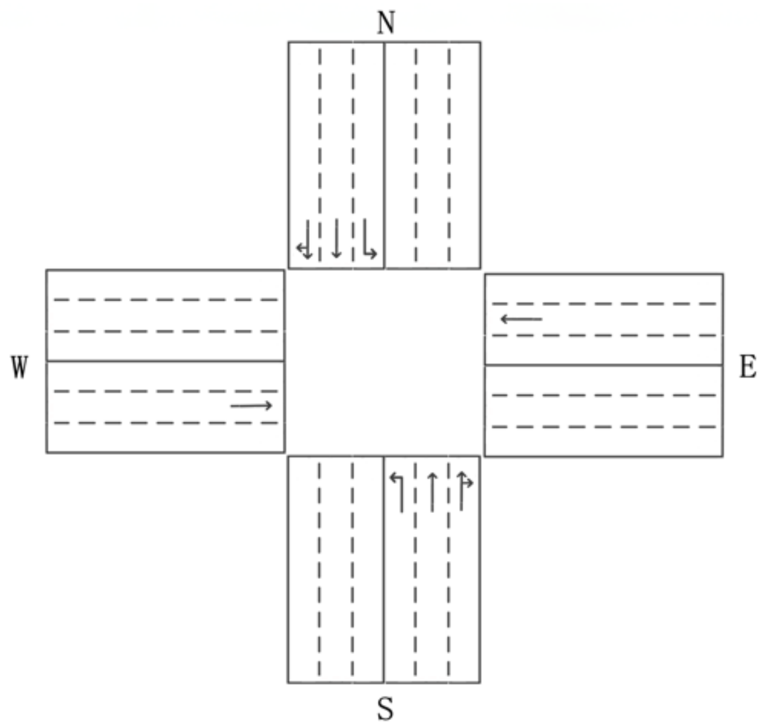

Table 2. Configuration of Traffic Flows in the Simulation

Type	Probability (%)	Route
Passenger	10%	East – West
Passenger	10%	West – East
Passenger	10%	South – East (right turn)
Passenger	10%	North – West (right turn)
Passenger	10%	South – North
Passenger	10%	North – South
Passenger	10%	North – East (left turn)
Platoons	30%	South – West (left turn)

3.3.5 Data Representation and Interpretation

After the network undergoes optimization through reinforcement learning, the efficiency of vehicle platoons traversing the intersection improves, resulting in higher speeds and throughput. As depicted in Figure 14, there is a notable increase in the system's mean speed (the average speed of all vehicles in the simulation) around the 50th episode of the learning process. Following this, the system's mean speed fluctuates based on traffic conditions but does not show significant variation. The reward function, illustrated in Figure 15, also indicates stabilization after the 50th episode. A comparison of our trained model (Figure 16(b)) with a simulation without any algorithmic intervention (Figure 16(a)) reveals a clear boost in terms of the system's mean speed. In saturated traffic conditions, our model achieved a system mean speed of 3.39±0.06 m/s, compared to 2.49±0.03 m/s in the non-algorithmic scenario, representing a 36% increase in the system mean speed.

Figure 14. System Mean Speed of Each Episode in the Learning Process

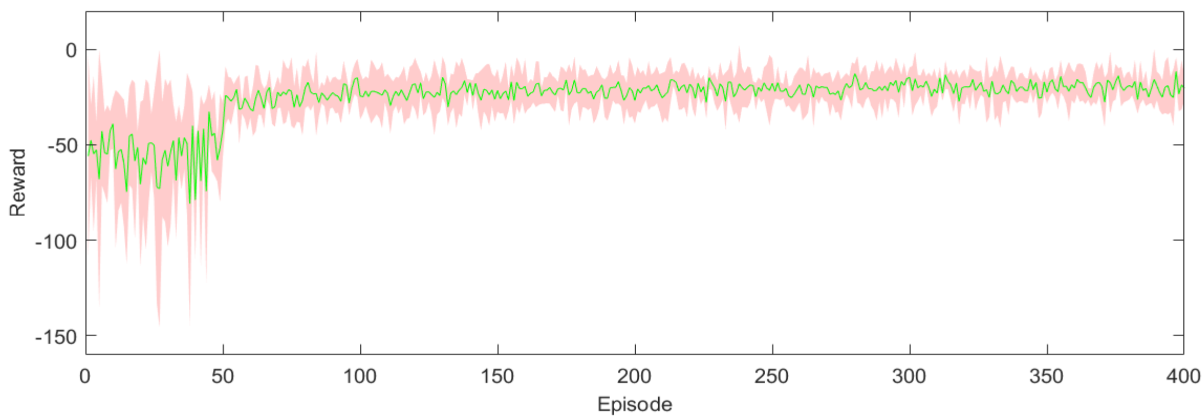

Figure 15. Reward Function of the Learning Process

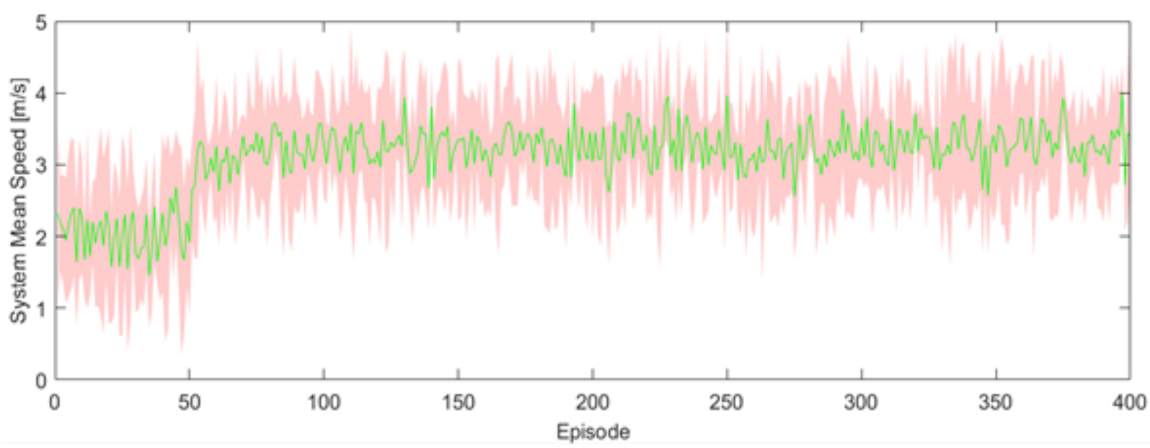

Figure 16. System Mean Speed of Non-Algorithmic Method and Our Method

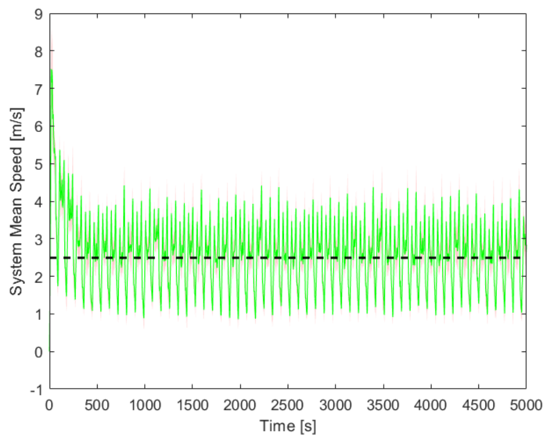

(a) Non-algorithmic

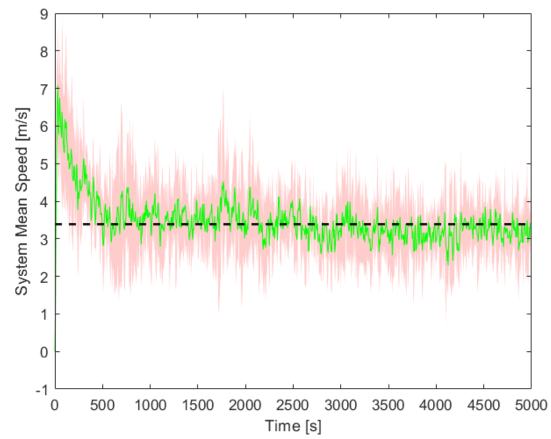

(b) Our Method

The green line represents the mean value; the pink region is the confidence interval; the black dashed line is the mean value for the time between 1,000 and 5,000 seconds.

3.3.6 Conclusions and Discussions

The unprotected left-turn maneuver is often a source of stress for drivers, demanding rapid decision-making and sharp perception skills, as highlighted by Park et al. (2017). The integration of reinforcement learning techniques into congested road environments presents a promising avenue to alleviate the risks and difficulties associated with these maneuvers. Our findings demonstrate a substantial improvement in the system's mean speed, showing a 36% increase compared to intersections operating without any algorithmic assistance. This enhancement is not only reflected in faster traffic flow but also suggests a potential reduction in the stress and complexity involved in executing unprotected left turns, thereby contributing to more efficient road conditions.

In our research, the traffic generated was defined in a certain direction of the network to create a mixed traffic scenario. Further investigation will be conducted into the introduction of human-driven vehicles into platoons and the introduction of more uncertainty in traffic flows. Further experimentation with autonomous vehicles traveling in different directions on the network will also be explored. Autonomous vehicles are often praised for their ability to perform well in different driving maneuvers, however, recent legislative actions in California have cast doubt on their readiness for real-world challenges (Tabarrok, 2023). The legislation has barred companies such as GM's Cruise and Google's Waymo from operating in the state, citing safety concerns. This development underscores the need for researchers to consider real-world applicability further in the future.

4. Summary and Conclusions

In this project, we proposed a deep reinforcement learning-based approach to the traffic congestion problem. Utilizing one or more traffic signals being controlled by a DQN algorithm coupled with an auto-encoder, we were able to achieve significant performance improvement in terms of vehicles' average waiting time or speed. Our simulation results demonstrate that the proposed approach outperformed traditional traffic signal control methods. Future work may involve exploring other RL-based frameworks or algorithms for improved performance and robustness, testing the algorithm in more complex traffic scenarios and integrating it with real-time traffic data, and incorporating additional environmental factors, such as weather conditions and pedestrian movements, could further enhance the robustness of the system. Collaboration with city planners and transportation authorities will be essential to facilitate real-world implementation and to ensure that the system can adapt to varying urban layouts and traffic patterns. Another promising direction for future research is the integration of vehicle-to-everything (V2X) communication technologies. By enabling direct communication between vehicles or between vehicles and traffic infrastructures, we can achieve even more responsive and efficient traffic management. This could include real-time adjustments to signal timings based on the density and flow of traffic, potentially reducing congestion and improving overall traffic flow.

In summary, our deep reinforcement learning-based approach to traffic signal control has shown great potential in reducing traffic congestion and improving vehicle flow. However, continued research and collaboration with various stakeholders will be crucial for the successful real-world implementation of this technology.

Bibliography

AbuLibdeh, A. (2017). Traffic congestion pricing: Methodologies and equity implications. *Urban Transp. Syst.*, 203–227.

Alegre, L. N. (2019). SUMO-RL, https://github.com/LucasAlegre/sumo-rl

Badnava, S., Meskin, N., Gastli, A., Al-Hitmi, M. A., Ghommam, J., Mesbah, M., & Mnif, F. (2021). Platoon transitional maneuver control system: A review. *IEEE Access*, *9*, 88327–88347.

Baldi, S., Michailidis, I., Ntampasi, V., Kosmatopoulos, E. B., Papamichail, I., & Papageorgiou, M. (2015). Simulation-based synthesis for approximately optimal urban traffic light management. In *2015 American Control Conference (ACC)* (pp. 868–873). IEEE.

Bianco, M. J. (2000). Effective transportation demand management: Combining parking pricing, transit incentives, and transportation management in a commercial district of Portland, Oregon. *Transportation Research Record*, *1711*(1), 46–54.

Börjesson, M., Eliasson, J., Hugosson, M. B., & Brundell-Freij, K. (2012). The Stockholm congestion charges—5 years on. Effects, acceptability and lessons learnt. *Transport Policy*, *20*, 1–12.

Calthrop, E., Proost, S., & Van Dender, K. (2000). Parking policies and road pricing. *Urban studies*, *37*(1), 63–76.

Cervero, R., & Guerra, E. (2011). Urban densities and transit: A multi-dimensional perspective. *UC Berkeley: Center for Future Urban Transport*.

Du, W., Ye, J., Gu, J., Li, J., Wei, H., & Wang, G. (2023). Safelight: A reinforcement learning method toward collision-free traffic signal control. In *Proceedings of the AAAI conference on artificial intelligence* (Vol. 37, No. 12, pp. 14801–14810).

Duranton, G., & Turner, M. A. (2011). The fundamental law of road congestion: Evidence from US cities. *American Economic Review*, *101*(6), 2616–2652.

Eliasson, J. (2014). *The Stockholm congestion charges: An overview* (p. 42). Centre for Transport Studies.

Eskandarian, A., Wu, C., & Sun, C. (2019). Research advances and challenges of autonomous and connected ground vehicles. *IEEE Transactions on Intelligent Transportation Systems*, *22*(2), 683–711.

Emamifar, M., & Ghoreishi, S. F. (2023). Modeling and state estimation of autonomous vehicles in signalized intersections. In *International Conference on Transportation and Development 2023* (pp. 10–21).

Gao, J., Shen, Y., Liu, J., Ito, M., & Shiratori, N. (2017). Adaptive traffic signal control: Deep reinforcement learning algorithm with experience replay and target network. *arXiv preprint* arXiv:1705.02755.

Gärling, T., & Schuitema, G. (2007). Travel demand management targeting reduced private car use: Effectiveness, public acceptability and political feasibility. *Journal of Social Issues*, *63*(1), 139–153.

Gholamhosseinian, A., & Seitz, J. (2022). A comprehensive survey on cooperative intersection management for heterogeneous connected vehicles. *IEEE Access*, *10*, 7937–7972.

Ghoul, T., & Sayed, T. (2021). Real-time signal-vehicle coupled control: An application of connected vehicle data to improve intersection safety. *Accident Analysis & Prevention*, *162*, 106389.

Jiang, X., & Gao, S. (2020). Signal control method and performance evaluation of an improved displaced left-turn intersection design in unsaturated traffic conditions. *Transportmetrica B: Transport Dynamics*, *8*(1), 264–289.

Li, M. Z. (2002). The role of speed–flow relationship in congestion pricing implementation with an application to Singapore. *Transportation Research Part B: Methodological*, *36*(8), 731–754.

Lin, Y., Dai, X., Li, L., & Wang, F. Y. (2018). An efficient deep reinforcement learning model for urban traffic control. *arXiv preprint* arXiv:1808.01876.

Lindsey, R., & Verhoef, E. (2001). Traffic congestion and congestion pricing. In *Handbook of transport systems and traffic control* (pp. 77–105). Emerald Group Publishing Limited.

Litman, T. (2006). London congestion pricing. *Implications for other cities*, 1–13.

Lu, Y., Zhang, Y., Shi, T., Wang, J., Wu, J. M., & Liu, B. (2023). Empirical study and signal intensity prediction for cellular vehicle-to-everything (C-V2X). In *2023 IEEE 98th Vehicular Technology Conference (VTC2023-Fall)* (pp. 1–6). IEEE.

Maadi, S., Stein, S., Hong, J., & Murray-Smith, R. (2022). Real-time adaptive traffic signal control in a connected and automated vehicle environment: Optimisation of signal planning with reinforcement learning under vehicle speed guidance. *Sensors, 22*(19), 7501.

Mahdinia, I., Mohammadnazar, A., Arvin, R., & Khattak, A. J. (2021). Integration of automated vehicles in mixed traffic: Evaluating changes in performance of following human-driven vehicles. *Accident Analysis & Prevention*, *152*, 106006.

Markakis, M. G., Talluri, K., & Tikhonenko, D. (2022). Managing lane-changing of algorithm-assisted drivers. *Transportation Research Part C: Emerging Technologies*, *138*, 103586.

Mannion, P., Duggan, J., & Howley, E. (2016). An experimental review of reinforcement learning algorithms for adaptive traffic signal control. *Autonomic road transport support systems*, 47–66.

Michaud, F., Lepage, P., Frenette, P., Letourneau, D., & Gaubert, N. (2006). Coordinated maneuvering of automated vehicles in platoons. *IEEE Transactions on Intelligent Transportation Systems*, *7*(4), 437–447.

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., & Bellemare, M. G. (2015). Human-level control through deep reinforcement learning. *Nature*, *518*(7540), 529–533.

National Highway Traffic Safety Administration. (2008). National motor vehicle crash causation survey: Report to congress. *National Highway Traffic Safety Administration Technical Report DOT HS*, *811*, 059.

Nishi, T., Otaki, K., Hayakawa, K., & Yoshimura, T. (2018). Traffic signal control based on reinforcement learning with graph convolutional neural nets. In *2018 21st International conference on intelligent transportation systems (ITSC)* (pp. 877–883). IEEE.

Park, B., Messer, C. J., & Urbanik, T. (1999). Traffic signal optimization program for oversaturated conditions: Genetic algorithm approach. *Transportation Research Record*, *1683*(1), 133–142.

Park, S. J., Subramaniyam, M., Kim, S. E., Hong, S., Lee, J. H., & Jo, C. M. (2017). Older driver's physiological response under risky driving conditions—Overtaking, unprotected left turn. In *Advances in Applied Digital Human Modeling and Simulation: Proceedings of the AHFE 2016 International Conference on Digital Human Modeling and Simulation, July 27–31, 2016, Walt Disney World, Florida, USA* (pp. 107–114). Springer International Publishing.

Pawar, N. M., Velaga, N. R., & Mishra, S. (2022). Impact of time pressure on acceleration behavior and crossing decision at the onset of yellow signal. *Transportation research part F: traffic psychology and behaviour*, *87*, 1–18.

Ren, Y., Wang, Y., Yu, G., Liu, H., & Xiao, L. (2016). An adaptive signal control scheme to prevent intersection traffic blockage. *IEEE Transactions on Intelligent Transportation Systems*, *18*(6), 1519–1528.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by backpropagating errors. *Nature*, *323*(6088), 533–536.

Selmoune, A., Cheng, Q., Wang, L., & Liu, Z. (2020). Influencing factors in congestion pricing acceptability: A literature review. *Journal of Advanced Transportation*, *2020*(1), 4242964.

Song, R., Lyu, L., Jiang, W., Festag, A., & Knoll, A. (2023). V2X-boosted federated learning for cooperative intelligent transportation systems with contextual client selection. *arXiv preprint arXiv:2305.11654*.

Spall, J. C., & Chin, D. C. (1997). Traffic-responsive signal timing for system-wide traffic control. *Transportation Research Part C: Emerging Technologies*, *5*(3–4), 153–163.

Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT Press.

Tabarrok, A. (2023). Driverless cars may already be safer than human drivers. https://policycommons.net/artifacts/4782643/driverless-cars-may-already-be-safer-than-human-drivers/5618936/. Accessed: Dec. 1, 2023.

Tang, C. R., Hsieh, J. W., & Teng, S. Y. (2023). Cooperative multi-objective reinforcement learning for traffic signal control and carbon emission reduction. *arXiv preprint* arXiv:2306.09662.

Teklu, F., Sumalee, A., & Watling, D. (2007). A genetic algorithm approach for optimizing traffic control signals considering routing. *Computer-Aided Civil and Infrastructure Engineering*, *22*(1), 31–43.

Timmerman, R. W., & Boon, M. A. (2021). Platoon forming algorithms for intelligent street intersections. *Transportmetrica A: transport science*, *17*(3), 278–307.

Tunc, I., & Soylemez, M. T. (2023). Fuzzy logic and deep Q learning based control for traffic lights. *Alexandria Engineering Journal*, *67*, 343–359.

Varaiya, P. (2013). Max pressure control of a network of signalized intersections. *Transportation Research Part C: Emerging Technologies*, *36*, 177–195.

Wang, T., Cao, J., & Hussain, A. (2021). Adaptive traffic signal control for large-scale scenario with cooperative group-based multi-agent reinforcement learning. *Transportation research part C: emerging technologies*, *125*, 103046.

Wiering, M. A. (2000). Multi-agent reinforcement learning for traffic light control. In *Machine Learning: Proceedings of the Seventeenth International Conference (ICML'2000)* (pp. 1151–1158).

Woo, S., & Skabardonis, A. (2021). Flow-aware platoon formation of connected automated vehicles in a mixed traffic with human-driven vehicles. *Transportation research part C: emerging technologies*, *133*, 103442.

Xiang, X., Fan, B., Dai, M., Wu, Y., & Xu, C. Z. (2022, June). V2X communication aided emergency message dissemination in intelligent transportation systems. In *2022 IEEE 23rd International Conference on High Performance Switching and Routing (HPSR)* (pp. 35–40). IEEE.

Xu, Y., Wang, Y., & Liu, C. (2022). Training a reinforcement learning agent with AutoRL for traffic signal control. In *2022 Euro-Asia Conference on Frontiers of Computer Science and Information Technology (FCSIT)* (pp. 51–55). IEEE.

Zhang, S., Chen, J., Lyu, F., Cheng, N., Shi, W., & Shen, X. (2018). Vehicular communication networks in the automated driving era. *IEEE Communications Magazine*, *56*(9), 26–32.

Zhang, X., Mo, H., Ma, H., & Luo, Q. (2022a). Deep recurrent Q networks for urban traffic signal control. In *CICTP 2022* (pp. 318–325).

Zhang, L., Wang, Y., & Zhu, H. (2022b). Theory and experiment of cooperative control at multi-intersections in intelligent connected vehicle environment: Review and perspectives. *Sustainability*, *14*(3), 1542.

About the Author

Tairan Liu, PhD

Dr. Tairan Liu received his B.Sc. degree from the University of Science and Technology of China, Hefei, China, in 2012, and his Ph.D. in Mechanical Engineering from Louisiana State University in 2020. He held postdoctoral appointments at the University of Georgia (2020–2021) and North Carolina State University (2021–2022). He is currently an Assistant Professor in the Department of Mechanical and Aerospace Engineering at California State University, Long Beach. His research interests include control theory, networked systems, and robotics.

MTI FOUNDER

Hon. Norman Y. Mineta

MTI BOARD OF TRUSTEES

Founder, Honorable Norman Mineta***
Secretary (ret.),
US Department of Transportation

Chair, Jeff Morales
Managing Principal
InfraStrategies, LLC

Vice Chair, Donna DeMartino
Retired Transportation Executive

Executive Director, Karen Philbrick, PhD*
Mineta Transportation Institute
San José State University

Rashidi Barnes
CEO
Tri Delta Transit

David Castagnetti
Partner
Dentons Global Advisors

Kristin Decas
CEO & Port Director
Port of Hueneme

Stephen J. Gardner*
President & CEO
Amtrak

Kimberly Haynes-Slaughter
Executive Consultant
Olivier, Inc.

Ian Jefferies*
President & CEO
Association of American Railroads

Diane Woodend Jones
Principal & Chair of Board
Lea + Elliott, Inc.

Priya Kannan, PhD*
Dean
Lucas College and
Graduate School of Business
San José State University

Will Kempton**
Retired Transportation Executive

David S. Kim
Senior Vice President
Principal, National Transportation
Policy and Multimodal Strategy
WSP

Therese McMillan
Retired Executive Director
Metropolitan Transportation
Commission (MTC)

Abbas Mohaddes
Chairman of the Board
Umovity

Stephen Morrissey
Vice President – Regulatory and
Policy
United Airlines

Toks Omishakin*
Secretary
California State Transportation
Agency (CALSTA)

Sachie Oshima, MD
Chair & CEO
Allied Telesis

April Rai
President & CEO
Conference of Minority
Transportation Officials (COMTO)

Greg Regan*
President
Transportation Trades Department,
AFL-CIO

Paul Skoutelas*
President & CEO
American Public Transportation
Association (APTA)

Rodney Slater
Partner
Squire Patton Boggs

Tony Tavares*
Director
California Department of
Transportation (Caltrans)

Lynda Tran
CEO
Lincoln Room Strategies

Jim Tymon*
Executive Director
American Association of
State Highway and Transportation
Officials (AASHTO)

Josue Vaglienty
Senior Program Manager
Orange County Transportation
Authority (OCTA)

* = Ex-Officio
** = Past Chair, Board of Trustees
*** = Deceased

Directors

Karen Philbrick, PhD
Executive Director

Hilary Nixon, PhD
Deputy Executive Director

Asha Weinstein Agrawal, PhD
Education Director
National Transportation Finance Center Director

Brian Michael Jenkins
Allied Telesis National Transportation Security Center

MTI MINETA TRANSPORTATION INSTITUTE