research parks publishing

International Journal on Orange Technologies (IJOT)

e-ISSN: 2615-8140
p-ISSN: 2615-7071

Volume: 8 Issue: 1 | January 2026
https://journals.researchparks.org/index.php/IJOT

Article

Adaptive Traffic Signal Control Based on Deep Reinforcement Learning with Edge Computing Scheme to Overcome The Surge in Vehicle Volume Post-Pandemic: A Critical Review of The Model and Implementation Challenges

Zulfadhli[1], Syarwan[2], Defry Basrin[3]

1. Samudra University, Indonesia
2. Lhokseumawe State Polytechnic, Indonesia
3. Samudra University, Indonesia

* Correspondence: zulfadhli@unsam.ac.id [1], syarwan@pnl.ac.id [2], defrybasrin@unsam.ac.id [3]

Abstract Fundamental changes in urban mobility patterns have led to an unpredictable and non-stationary surge in vehicle volume, driven by Work From Anywhere policies and a significant increase in private vehicle usage and ride-hailing services. Consequently, a new paradigm integrating artificial intelligence with advanced computing infrastructure is required. This study constitutes a literature review aimed at providing a comprehensive critical analysis of Deep Reinforcement Learning models and Edge Computing schemes within the context of Adaptive Traffic Signal Control, with particular focus on implementation challenges in the new normal mobility era. The findings reveal four primary insights. First, Multi-Agent Deep Reinforcement Learning architectures incorporating communication mechanisms based on Graph Neural Networks demonstrate superior performance in multi-intersection scenarios, yet remain vulnerable to distributional shift phenomena caused by non-stationary travel pattern changes. Second, Edge Computing theoretically reduces latency and enhances system resilience to network failures, although its deployment is constrained by computational resource limitations and energy consumption issues on edge devices operating in extreme intersection environments. Third, an overreliance on simulation data from SUMO or VISSIM introduces significant validity gaps when models are applied to real-world mobility dynamics influenced by heterogeneous data sources such as probe vehicles and loop detector sensors. Fourth, implementation barriers are multidimensional, encompassing computational complexity, susceptibility to adversarial attacks on DRL policies, and regulatory and interoperability gaps with legacy infrastructure. The practical implications of this research emphasize the development of compact DRL models leveraging knowledge distillation for low-power edge devices, alongside technical interoperability guidelines to facilitate gradual transition from conventional systems.

Keywords: Adaptive Traffic Signal Control, Deep Reinforcement Learning, Edge Computing, Post-Pandemic Mobility, Critical Review

Citation: Zulfadhli, Syarwan, Basrin, D. Adaptive Traffic Signal Control Based on Deep Reinforcement Learning with Edge Computing Scheme to Overcome The Surge in Vehicle Volume Post-Pandemic: A Critical Review of The Model and Implementation Challenges. International International Journal on Orange Technologies (IJOT) 2026, 8(1), 1-12

Received: 10th Oct 2025
Revised: 04th Nov 2025
Accepted: 14th Dec 2025
Published: 26th Jan 2026

1. Introduction

The COVID-19 pandemic, which has impacted the world since early 2020, has left a profound transformative mark on nearly all aspects of human life, including urban mobility patterns and traffic dynamics. Prior to the pandemic, urban transportation systems operated within a relatively predictable rhythm, characterized by peak traffic congestion concentrated during morning and evening rush hours, corresponding with the mass movement of workers from residential areas to office centers and industrial zones, and vice versa. This predictability enabled traffic management authorities to design and implement periodic, fixed-time signal control strategies with an acceptable level of effectiveness. However, the pandemic has fundamentally disrupted this order in a manner that is likely permanent. Large-scale social restrictions, temporary workplace closures, and most significantly, the widespread adoption of remote and hybrid work models have given rise to a new mobility phenomenon known as Work From Anywhere. This phenomenon has not only altered the origins and destinations of trips but has also radically disrupted the temporal distribution of travel demand, generating surges in vehicle volumes during previously low-traffic periods such as midday, late afternoon, and even late at night. The harmonious or heterogeneous interaction among these factors influences outcomes and the success or failure of related objectives [1]. Conventional traffic signal control systems, which rely on fixed schedules and respond only reactively to changes, have proven inadequate for managing the complexity and volatility of traffic flows in this new normal era [2].

The inability of conventional traffic signal control systems to adapt to the dynamic changes in post-pandemic traffic has exacerbated congestion problems in many major cities worldwide, including metropolitan areas in Indonesia. Traffic congestion is not merely a transient inconvenience but a multidimensional issue that generates cascading negative impacts on the fabric of urban life. From an economic perspective, congestion leads to massive fuel wastage, increased logistics and distribution costs, and a significant loss of productivity due to inefficient and unpredictable travel times. The National Development Planning Agency and various transportation research institutions have repeatedly published estimates of the annual economic losses caused by congestion in Indonesia's major cities, amounting to tens of trillions of rupiah. These staggering figures do not fully account for indirect losses such as the decline in public health quality due to prolonged exposure to air pollution, elevated stress levels and reduced mental well-being among commuters, and the overall degradation of urban environmental quality. From an environmental standpoint, vehicles trapped in congestion, idling or moving at low speeds, are major contributors to greenhouse gas emissions and harmful air pollutants such as carbon monoxide, nitrogen oxides, and fine particulate matter in urban areas. Conventional efforts to alleviate congestion through road widening and new infrastructure development have increasingly reached an impasse, not only due to land and budgetary constraints but also because of the well-documented phenomenon of induced demand in transportation planning literature. This phenomenon describes how increased road capacity stimulates additional travel demand, ultimately restoring congestion levels to their original state within a relatively short period [3].

In addressing the impasse of conventional solutions, the wave of technological innovation driven by the Fourth Industrial Revolution offers a new beacon of hope through the development of Intelligent Transportation Systems (ITS). The core paradigm of ITS lies in leveraging information and communication technologies, sensors, and advanced computing to optimize the utilization of existing transportation infrastructure, rather than continuously constructing new infrastructure. One of the most critical and promising components within the ITS ecosystem is the Adaptive Traffic Signal Control (ATSC) system. Unlike conventional approaches that are passive and periodic, ATSC is designed to actively and continuously adjust signal parameters such as cycle duration, green time allocation among phases, and offsets between adjacent intersections based on real-time traffic conditions detected by various sensor types. Early generations of ATSC, such as the Sydney Coordinated Adaptive Traffic System (SCATS) and the Split, Cycle, and Offset Optimization Technique (SCOOT), have demonstrated performance improvements over fixed-time systems, particularly under moderate and relatively stable traffic flows. Nevertheless, these systems still rely on relatively simple mathematical models, limited state spaces, and heuristics predetermined by traffic engineers. Consequently, the adaptive capabilities of early-generation ATSC systems remain suboptimal when confronted with the full complexity of urban road networks involving dozens or even hundreds of intersections interacting in a nonlinear manner, especially under highly dynamic and non-stationary traffic conditions characteristic of the post-pandemic era [4].

The next fundamental breakthrough in the evolution of ATSC systems arises from the intersection of two rapidly advancing fields of artificial intelligence over the past decade: Deep Learning and Reinforcement Learning. Together, they form a novel paradigm known as Deep Reinforcement Learning (DRL). Reinforcement Learning fundamentally constitutes a computational framework wherein an intelligent agent learns to make optimal decisions through a trial-and-error process by interacting with its environment, aiming to maximize the cumulative reward signals received. The agent is not explicitly instructed on which actions to take in each situation; rather, it explores the possible action space and learns from the positive and negative consequences of its decisions. When this RL paradigm is combined with the remarkable representational capability of Deep Neural Networks, which can automatically extract high-level features from high-dimensional raw data, Deep Reinforcement Learning emerges with the capacity to solve highly complex sequential decision-making problems that were previously beyond the reach of conventional methods. In the context of traffic signal control, a DRL agent can be trained to directly map rich traffic state representations such as images from monitoring cameras or occupancy vectors from loop detector sensors into optimal signal phase control decisions. The primary advantage of the DRL approach lies in its ability to capture nonlinear patterns and intricate interactions among intersections without requiring explicit mathematical modeling of the underlying traffic dynamics [5].

However, the practical implementation of advanced Deep Reinforcement Learning (DRL) models for adaptive traffic light control cannot be separated from fundamental considerations regarding the computational architecture that underpins their operation. Most early research in this domain assumed a centralized cloud-based computing architecture, wherein data from various sensors distributed throughout the city are transmitted to a remote data center for processing, and the phase control decisions generated by the DRL model are subsequently sent back to the signal controllers at each intersection. Although the cloud-based approach offers virtually unlimited computational and storage capacity, it presents inherent drawbacks that become increasingly critical as real-time requirements and system scale grow. The reliance on bidirectional data transmission over public internet networks introduces latency components that may render signal control decisions outdated by the time they reach actuators in the field, thereby significantly diminishing the optimization benefits. Furthermore, the centralized architecture creates a single point of failure, where disruptions in network connectivity or the data center itself can incapacitate the entire urban traffic control system, a highly risky scenario especially during extreme traffic surges or urban emergency situations. To address these limitations, the Edge Computing paradigm has emerged as an architectural solution fundamentally better aligned with the characteristics of real-time traffic control applications. In an Edge Computing architecture, computing and storage resources are distributed to the network edge, as close as possible to the data generation and action execution sites in this case, directly at traffic signal controller cabinets at each intersection or at aggregators serving small clusters of nearby intersections. By processing data and executing DRL model inference locally, Edge Computing drastically reduces latency, enhances system resilience against network disruptions, and alleviates data traffic burdens on telecommunications infrastructure [6].

Building upon the background outlined, it becomes increasingly evident that there exists an urgent need, as well as a significant opportunity, to undertake a comprehensive and critical review of the state-of-the-art integration between Deep Reinforcement Learning (DRL) and Edge Computing within the context of Adaptive Traffic Signal Control. This review places particular emphasis on the unique challenges arising from the shifting mobility patterns in the post-pandemic era. Existing literature reviews predominantly address the algorithmic aspects of DRL and the infrastructural components of Edge Computing separately, thereby failing to capture the dynamic interactions and interdependencies between these two domains. An integrated review is therefore essential to holistically map the current research landscape, identify the most promising model architectures and implementation strategies, and importantly, reveal knowledge gaps and practical barriers that have been largely overlooked amid technological optimism.

This critical review specifically aims to address fundamental questions such as: How has the evolution of DRL model architectures from single-agent to multi-agent approaches responded to the increased complexity of coordination among intersections under non-stationary mobility conditions? In what ways does Edge Computing function as a technical enabler for achieving low-latency DRL model inference on resource-constrained devices within the challenging environment of traffic intersections? To what extent do simulation data, which have dominated model training efforts, accurately represent the actual mobility dynamics influenced by new behaviors such as Work From Anywhere in the post-pandemic context? Finally, what multi-dimensional barriers including technical, cybersecurity, and regulatory challenges impede the transition from research laboratories to sustainable real-world deployment?. Through a systematic and critical analysis of the recent literature corpus, this study aims not only to provide a clear depiction of the current state of the field but also to offer nuanced perspectives on future research directions and more realistic, context-aware implementation strategies [7].

2. Materials and Methods

This study fundamentally employs a literature review approach, commonly referred to as library research, as its primary methodological framework. The selection of this research type is based on the conceptual and multidimensional nature of the subject matter, wherein the researcher does not engage in primary data collection through field observation or laboratory experiments. Instead, the study involves tracing, examining, and synthesizing various previously published sources. In the context of research focusing on Adaptive Traffic Signal Control based on Deep Reinforcement Learning with an Edge Computing scheme to address post-pandemic vehicle volume surges, the literature review approach is considered the most appropriate and strategic choice. This is due to the topic's intersection across computer science, transportation engineering, and public policy fields characterized by rapid development and extensively documented in a range of scientific publications, textbooks, and institutional research reports. The literature review enables the researcher to conduct a comprehensive mapping of the existing research landscape, identify knowledge gaps, and critically analyze various models, approaches, and findings produced by the global scientific community.

The primary data sources underpinning this literature review comprise three major categories of selected materials, chosen purposively based on their relevance, credibility, and currency. The first category includes textbooks and academic references, sourced from reputable international publishers as well as prominent national publishers, which substantially address the theoretical foundations of Deep Reinforcement Learning, Edge Computing architectures, adaptive traffic control systems, and urban mobility dynamics. The selected books are specifically limited to publications from 2021 to 2025 to ensure that the theoretical framework reflects the most recent advancements in the pertinent fields. The second category consists of articles published in peer-reviewed international and national scientific journals, particularly those reporting empirical or simulation research related to the implementation of DRL for traffic signal control, the utilization of Edge Computing in intelligent transportation systems, and studies on mobility pattern changes resulting from the COVID-19 pandemic. Journal articles were sourced from reputable databases such as IEEE Xplore, ScienceDirect, SpringerLink, and Garuda for national publications, with priority given to highly cited articles published within the last five years. The third category encompasses research reports and policy documents issued by government agencies, international organizations, independent research institutions, and transportation consulting firms recognized for their authority and credibility in the domains of traffic management and intelligent transportation technologies [8].

The data collection technique in this literature-based research was conducted through a series of systematic and well-documented procedures, beginning with the identification of keywords and formulation of a search strategy, followed by the processes of literature collection and selection, and concluding with data extraction and the compilation of the research corpus. Initially, the researcher formulated a set of primary and derivative keywords in both English and Indonesian, relevant to the four major research sub-themes: Adaptive Traffic Signal Control, Deep Reinforcement Learning, Multi-Agent DRL, Edge Computing, Post-Pandemic Mobility Patterns, and Traffic Simulation. These keywords were then combined using Boolean operators to perform systematic searches across online library catalogs, electronic journal databases, institutional repositories, and academic search engines such as Google Scholar. The search was conducted iteratively and in multiple layers, where initial findings from one source often led the researcher to additional sources through citation chaining, both backward by examining the reference lists of relevant articles and forward by identifying articles that cited the key publications. Each potentially relevant source underwent an initial screening based on the alignment of its title and abstract with the research focus, followed by a full selection process involving a thorough reading of the complete text [9].

The data analysis technique employed in this study is qualitative content analysis, conducted through a simultaneous deductive-inductive approach. The deductive approach is applied at the initial stage of analysis, wherein the researcher utilizes a predefined set of categories and themes derived from the theoretical framework and previously formulated research questions. These initial categories include the architecture of the DRL model, comparative algorithm performance, Edge Computing architecture, heterogeneous data integration, and implementation barriers. Concurrently, the inductive approach is implemented to allow the emergence of novel and unforeseen themes that may not be encompassed within the initial framework but are significant for comprehending the studied phenomena. The analytical process begins with repeated and thorough readings of all selected sources, accompanied by the assignment of conceptual codes to relevant text segments aligned with the research focus. These codes are subsequently organized hierarchically into broader categories, which are further synthesized into primary themes forming the backbone of the narrative in the results section. Within the context of this research, the analysis extends beyond mere description to encompass interpretative and critical levels, where the researcher actively compares and contrasts findings from various sources, identifies patterns of convergence and divergence, and uncovers underlying assumptions supporting the claims and conclusions in the literature. Particular attention is given to identifying gaps between theoretical advancements reported in simulation studies and the realities of field implementation, as well as to fundamental shifts in post-pandemic urban mobility patterns that may not be fully reflected in models developed before or during the pandemic.

The validity testing techniques employed in this literature review study were conducted through the application of several validation strategies adapted from conventional qualitative research frameworks and tailored to the distinctive characteristics of literature-based research. The first and most fundamental strategy is source triangulation, wherein the researcher deliberately compares information and claims derived from various types of sources, such as juxtaposing findings reported in journal articles with implementation descriptions found in government technical reports or with critical analyses presented in textbooks. The second strategy involves peer debriefing, executed by presenting preliminary findings and analytical frameworks to academic colleagues with expertise in intelligent transportation or machine learning who are not directly involved in the study. The third strategy entails the search for negative evidence or deviant cases, whereby the researcher actively seeks literature documenting implementation failures, nonsignificant results, or findings that contradict the prevailing optimistic narratives regarding deep reinforcement learning (DRL) and edge computing in traffic control. The fourth strategy comprises the preparation of an audit trail, meticulously documenting all methodological decisions made throughout the research process, including literature search strategies, source selection criteria, coding schemes, and the reasoning pathways leading to specific conclusions [9].

3. Results and Discussion

A. Architecture Evolution and Performance of Deep Reinforcement Learning Models in Multi-Intersection Scenarios under Post-Pandemic Mobility Conditions

The development of Deep Reinforcement Learning (DRL) algorithms for adaptive traffic signal control has experienced significant advancements over the past decade, particularly when addressing the complexity of interconnected multi-intersection scenarios. A literature review of recent studies reveals a fundamental paradigm shift from fixed-rule based approaches and conventional actuation systems toward fully autonomous, data-driven decision-making mechanisms. In the context of post-pandemic mobility, characterized by high volatility and uncertainty in travel patterns, DRL architectures are required not only to optimize vehicle flow under stationary conditions but also to demonstrate robust generalization capabilities in response to non-stationary traffic dynamics. Early models commonly employed in traffic signal control research predominantly relied on variants of the Deep Q-Network (DQN), wherein intelligent agents endeavor to learn the optimal value of state-action pairs through value function approximation using deep neural networks. However, a critical limitation of the standard DQN lies in its tendency to suffer from overestimation bias, which can lead to unstable policies when applied to environments with large action spaces and unpredictable state transitions phenomena that closely mirror the distortions in traffic flow patterns caused by shifting mobility behaviors in urban intersection networks post-pandemic [10].

Reflecting on these limitations, the research community has subsequently developed various improvement mechanisms and alternative architectures that substantially enhance the robustness of Deep Reinforcement Learning (DRL) agents in coping with environmental uncertainties. Actor-Critic-based algorithms, such as Advantage Actor-Critic (A2C) and Proximal Policy Optimization (PPO), have emerged as more promising candidates due to their ability to balance exploration and exploitation through an explicit separation between the policy network (actor) and the value evaluation network (critic). This hybrid approach has demonstrated superior performance in multi-intersection scenarios by maintaining learning process stability while accommodating the continuous action space required for granular adjustment of green phase durations. Furthermore, in the context of signal coordination within urban corridors affected by surge phenomena in vehicle volumes that no longer follow conventional peak-hour patterns post-pandemic, the Multi-Agent Deep Reinforcement Learning (MADRL) framework has garnered significant attention in recent literature. Within the MADRL framework, each intersection is treated as an autonomous agent capable of learning independently while collaborating or competing with neighboring agents to achieve a global network optimum. A critical challenge in implementing MADRL is addressing the non-stationarity problem arising from simultaneous policy updates by other agents, a condition analogous to mobility dynamics in the new normal era, where individual urban travel decisions can no longer be predicted solely based on historical data. To mitigate this issue, recent studies have extensively explored inter-agent communication mechanisms through Graph Neural Networks (GNN) or attention mechanisms, enabling each intersection to assign varying weights to information received from upstream and downstream intersections depending on their relative importance to current traffic conditions [11].

A critical analysis of the reviewed literature corpus reveals a significant gap between the reported performance of models in idealized simulation environments and their actual potential when confronted with post-pandemic mobility characteristics. The majority of empirical studies demonstrating the superiority of Multi-Agent Deep Reinforcement Learning (MADRL) algorithms, such as those employing Multi-Agent Deep Deterministic Policy Gradient (MADDPG) or Q-MIX architectures, typically assume vehicle arrival patterns that follow a relatively stable Poisson distribution throughout the training episodes. This assumption becomes problematic when faced with post-pandemic realities, which exhibit new traffic peaks at times previously considered off-peak such as midday or late night directly resulting from flexible working hours and the widespread adoption of Work From Anywhere policies. DRL models trained solely on synthetic data with stationary characteristics tend to experience significant performance degradation when tested under scenarios featuring sudden and unpredictable demand surges, a phenomenon known in the literature as distributional shift. Some pioneering studies have begun to address this challenge by introducing architectures that integrate long-term memory modules, such as Long Short-Term Memory (LSTM), into DRL policy networks, enabling agents to respond not only to instantaneous conditions but also to retain and learn from historical patterns, even when those patterns are sparse and irregular. Another noteworthy approach involves the use of transfer learning and meta-learning techniques, whereby models trained on one road network configuration can rapidly adapt to new configurations or drastically changing traffic conditions without requiring retraining from scratch. These findings underscore that the successful implementation of DRL for adaptive traffic signal control in the post-pandemic era depends not only on the appropriate choice of model architecture but also on training and validation strategies that explicitly account for the non-stationary characteristics and inherent uncertainties of contemporary urban mobility systems [12].

B. The Role of Edge Computing Schemes as Enablers of Fast Computing for Inference and Low Latency at Critical Junctions

Discussions regarding the effectiveness of Deep Reinforcement Learning algorithms for adaptive traffic signal control will never reach a meaningful practical level without an in-depth discourse on the computational infrastructure that underpins their operationalization. In this context, the Edge Computing paradigm emerges not merely as a complementary element but as a constitutive foundation for the realization of intelligent transportation systems that respond to traffic dynamics in real time. Unlike centralized cloud computing architectures that rely on transmitting data to distant data centers for processing, Edge Computing distributes computational loads to processing nodes located at the network edge, or, in the context of traffic control, directly at the intersection sites themselves or at aggregation gateways serving clusters of nearby intersections. A comprehensive survey of the literature examining computing architectures for intelligent transportation systems reveals an increasingly crystallized consensus that cloud-centric approaches are no longer sufficient to support applications with stringent latency requirements demanding decision-making on the order of milliseconds such as traffic signal phase control, which must respond instantaneously to changes in vehicle queue lengths and arrivals. Data transmission delays caused by the geographical distance between intersections and data centers, coupled with potential backbone internet network congestion, can result in phase control decisions generated by DRL models becoming obsolete by the time they are received back by field signal controllers, thereby significantly eroding the optimization benefits that artificial intelligence should ideally provide [13].

The implementation of Edge Computing architecture for adaptive traffic light control presents a range of technical advantages that directly address various inherent limitations of cloud-based approaches. Primarily, by situating processing units at the network edge, the execution of Deep Reinforcement Learning (DRL) model inference can occur locally at intersections, thereby completely eliminating latency introduced by the bidirectional transmission of data to centralized data centers. Secondly, this distributed architecture is fundamentally more resilient to communication network failures, a critical attribute during extreme surges in vehicle volumes, which are often accompanied by increased strain on surrounding telecommunication infrastructure. In disaster scenarios or extraordinary events causing drastic increases in traffic flow, Edge Computing nodes maintain autonomous operation despite disconnections from central data centers, relying on DRL models embedded locally within their hardware. Furthermore, Edge Computing facilitates the potential deployment of more sophisticated collaborative architectures such as fog computing, wherein adjacent intersections form local computational clusters that share processing loads and traffic state information without involving city- or regional-level data centers. Such hierarchical computing models enable efficient optimization at corridor or sub-regional scales while preserving low latency for decision-making at individual intersections. However, the literature also identifies unresolved tensions between the desire to distribute intelligence to the network edge and the necessity of maintaining a global perspective on overall urban traffic conditions. MADRL models operating entirely on Edge nodes without coordination from higher-level entities risk generating policies that are self-serving and may sacrifice global optimality for local, short-term gains [14].

One of the most significant contributions of this critical review is the identification of a substantial gap between the idealized assumptions underpinning simulation studies and the technical and environmental realities encountered in the deployment of Edge Computing for traffic control in the field. The majority of research papers claiming successful integration of Deep Reinforcement Learning (DRL) and Edge Computing report results obtained from highly controlled simulation environments, where the availability of computational resources at the network edge is considered unlimited, wireless communication channels are assumed to be consistently perfect without interference, and energy consumption of Edge devices is never a limiting factor. The real-world conditions at urban intersections, particularly in cities in developing countries, differ markedly from these idealized assumptions. Edge devices installed in traffic signal control cabinets must operate within extreme temperature and humidity ranges, endure vibrations from heavy traffic, and rely on potentially unstable power supplies. Furthermore, the computational capacity of economically feasible embedded devices deployed at thousands of intersections is considerably lower than the capacity of data center-grade servers used in most simulation studies. This discrepancy raises critical questions regarding the feasibility of running DRL models with complex neural network architectures potentially requiring specialized hardware accelerators such as GPUs or TPUs on resource-constrained Edge devices. Recent literature has begun to address these challenges by exploring various model compression techniques, including network weight quantization, pruning of insignificant connections, and knowledge distillation from large teacher models to more compact student models. These techniques aim to produce computationally and memory-efficient representations of DRL models, enabling real-time inference on limited Edge hardware without substantially compromising policy accuracy. Findings from this review underscore that the successful implementation of Edge Computing for adaptive signal control depends not only on advancements in algorithms and network architectures but also demands parallel progress in model efficiency, energy-efficient hardware, and reliable communication protocols suited for challenging environments [15].

C. Post-Pandemic Heterogeneous Data Integration: From SUMO Simulations to Actual Mobility Dynamics Influenced by New Behaviors

Discussions regarding the effectiveness of Deep Reinforcement Learning algorithms for adaptive traffic signal control will never reach a meaningful practical level without an in-depth discourse on the computational infrastructure that underpins their operationalization. In this context, the Edge Computing paradigm emerges not merely as a complementary element but as a constitutive foundation for the realization of intelligent transportation systems that respond to traffic dynamics in real time. Unlike centralized cloud computing architectures that rely on transmitting data to distant data centers for processing, Edge Computing distributes computational loads to processing nodes located at the network edge, or, in the context of traffic control, directly at the intersection sites themselves or at aggregation gateways serving clusters of nearby intersections. A comprehensive survey of the literature examining computing architectures for intelligent transportation systems reveals an increasingly crystallized consensus that cloud-centric approaches are no longer sufficient to support applications with stringent latency requirements demanding decision-making on the order of milliseconds such as traffic signal phase control, which must respond instantaneously to changes in vehicle queue lengths and arrivals. Data transmission delays caused by the geographical distance between intersections and data centers, coupled with potential backbone internet network congestion, can result in phase control decisions generated by DRL models becoming obsolete by the time they are received back by field signal controllers, thereby significantly eroding the optimization benefits that artificial intelligence should ideally provide [13].

The implementation of Edge Computing architecture for adaptive traffic light control presents a range of technical advantages that directly address various inherent limitations of cloud-based approaches. Primarily, by situating processing units at the network edge, the execution of Deep Reinforcement Learning (DRL) model inference can occur locally at intersections, thereby completely eliminating latency introduced by the bidirectional transmission of data to centralized data centers. Secondly, this distributed architecture is fundamentally more resilient to communication network failures, a critical attribute during extreme surges in vehicle volumes, which are often accompanied by increased strain on surrounding telecommunication infrastructure. In disaster scenarios or extraordinary events causing drastic increases in traffic flow, Edge Computing nodes maintain autonomous operation despite disconnections from central data centers, relying on DRL models embedded locally within their hardware. Furthermore, Edge Computing facilitates the potential deployment of more sophisticated collaborative architectures such as fog computing, wherein adjacent intersections form local computational clusters that share processing loads and traffic state information without involving city- or regional-level data centers. Such hierarchical computing models enable efficient optimization at corridor or sub-regional scales while preserving low latency for decision-making at individual intersections. However, the literature also identifies unresolved tensions between the desire to distribute intelligence to the network edge and the necessity of maintaining a global perspective on overall urban traffic conditions. MADRL models operating entirely on Edge nodes without coordination from higher-level entities risk generating policies that are self-serving and may sacrifice global optimality for local, short-term gains [14].

One of the most significant contributions of this critical review is the identification of a substantial gap between the idealized assumptions underpinning simulation studies and the technical and environmental realities encountered in the deployment of Edge Computing for traffic control in the field. The majority of research papers claiming successful integration of Deep Reinforcement Learning (DRL) and Edge Computing report results obtained from highly controlled simulation environments, where the availability of computational resources at the network edge is considered unlimited, wireless communication channels are assumed to be consistently perfect without interference, and energy consumption of Edge devices is never a limiting factor. The real-world conditions at urban intersections, particularly in cities in developing countries, differ markedly from these idealized assumptions. Edge devices installed in traffic signal control cabinets must operate within extreme temperature and humidity ranges, endure vibrations from heavy traffic, and rely on potentially unstable power supplies. Furthermore, the computational capacity of economically feasible embedded devices deployed at thousands of intersections is considerably lower than the capacity of data center-grade servers used in most simulation studies. This discrepancy raises critical questions regarding the feasibility of running DRL models with complex neural network architectures potentially requiring specialized hardware accelerators such as GPUs or TPUs on resource-constrained Edge devices. Recent literature has begun to address these challenges by exploring various model compression techniques, including network weight quantization, pruning of insignificant connections, and knowledge distillation from large teacher models to more compact student models. These techniques aim to produce computationally and memory-efficient representations of DRL models, enabling real-time inference on limited Edge hardware without substantially compromising policy accuracy. Findings from this review underscore that the successful implementation of Edge Computing for adaptive signal control depends not only on advancements in algorithms and network architectures but also demands parallel progress in model efficiency, energy-efficient hardware, and reliable communication protocols suited for challenging environments [15].

D. Multi-Dimensional Barriers to Implementation: Computational Complexity, Cybersecurity, and Regulatory Gaps in the Adoption of Intelligent Adaptive Systems

Although the theoretical foundations and empirical evidence from simulation environments have convincingly demonstrated the remarkable potential of integrating Deep Reinforcement Learning (DRL) and Edge Computing to revolutionize adaptive traffic light control, the transition from research laboratories to sustainable field deployment remains hindered by a complex set of multidimensional challenges. The synthesis of the diverse literature reviewed in this study converges on three interrelated clusters of primary challenges that must be addressed concurrently: those related to computational complexity and operational sustainability, cybersecurity and system resilience, and regulatory gaps and interoperability issues. Neglecting any one of these dimensions significantly increases the risk of project failure or results in systems whose performance falls far short of the expectations established by academic publications.

First, from the perspective of computational complexity and operational sustainability, deploying Edge Computing nodes capable of continuous DRL model inference across thousands of urban intersections raises serious concerns regarding energy costs and device durability. Unlike data centers equipped with advanced cooling systems, redundant power supplies, and dedicated technical staff, roadside traffic signal control cabinets operate in harsh, low-maintenance environments. The Edge hardware necessary to run DRL models even following compression techniques consumes substantially more electrical power compared to conventional programmable logic controllers. In tropical urban contexts characterized by high temperature and humidity, heat dissipation from intensive computing devices can lead to performance degradation, shortened component lifespan, and increased system failure risks. Studies on energy efficiency in outdoor Edge Computing applications remain relatively limited and often fail to consider the total cost of ownership, which includes more frequent hardware replacements and aggregated increases in municipal electricity bills. Questions regarding who will bear these additional operational costs and whether the traffic efficiency gains can economically offset these increased expenses remain unresolved puzzles that have received scant attention within the predominantly computer science-focused literature [16].

Secondly, the dimension of cybersecurity emerges as a fundamental vulnerability with the potential to incapacitate not only the signal control system itself but also the entire urban transportation network reliant upon it. Adaptive traffic signal control systems driven by Deep Reinforcement Learning (DRL), with their extensive connectivity to various sensors and peer controllers, significantly expand the attack surface accessible to malicious actors. The most concerning and distinctive vulnerability inherent to machine learning-based systems is the threat of adversarial attacks. In this context, an attacker can deliberately manipulate the input data received by the DRL model by introducing subtle perturbations that are imperceptible to the human eye yet sufficient to drastically alter the model's policy outputs. For instance, by injecting false data from phantom vehicles into the data stream from loop detectors or probe vehicles, an attacker can deceive the DRL agent into believing a severe congestion event is occurring at a particular intersection approach. This manipulation causes the system to erroneously allocate an excessively long green light duration to an actually empty approach, thereby sacrificing other approaches that genuinely require signal priority. Such attacks can generate widespread artificial congestion or even be exploited to facilitate other crimes, such as robbery, by trapping victims in engineered traffic jams. Beyond adversarial attacks, these systems are also susceptible to conventional cyberattacks, including the hacking of edge devices to form botnets, interception of vehicle location data compromising privacy, and distributed denial-of-service (DDoS) attacks that disrupt communication between nodes and incapacitate coordination. The literature addressing cybersecurity aspects in DRL-based intelligent transportation systems remains in its nascent and largely exploratory stages, underscoring an urgent need to develop robust defense mechanisms. These include anomaly detection in input data, adversarial training to enhance model resilience, and layered security architectures that safeguard the entire system stack from hardware components through to the application layer [17].

Thirdly, one of the most challenging obstacles in the practice of engineering and public administration is the regulatory gap and the lack of interoperability standards between newly developed intelligent systems and legacy traffic signal infrastructure that has been operational for decades. Cities worldwide, including those in Indonesia, have made substantial investments in conventional centralized traffic control systems, which may rely on proprietary communication protocols, fixed-time cycles, and vendor-specific configuration databases. Replacing this entire infrastructure simultaneously with new systems based on Deep Reinforcement Learning (DRL) and Edge Computing is financially and logistically impractical. Consequently, the most feasible implementation strategy involves a phased and hybrid approach, whereby intelligent controllers are initially installed at critical intersections and must coexist and communicate effectively with the surrounding legacy controllers. This interoperability issue extends beyond technical aspects such as data formats and communication protocols to encompass deeper institutional concerns, including data ownership, legal liability in the event of system failures causing accidents, and government procurement processes, which tend to be rigid and unsupportive of rapidly evolving technological innovations. Current traffic regulations may also fail to accommodate fully adaptive, AI-based signal control concepts, where cycle durations and phases no longer follow predetermined patterns. This regulatory uncertainty may impede the widespread adoption of such technologies. A synthesis of the literature review concludes that the successful transition to adaptive traffic signal control systems based on DRL and Edge Computing necessitates an interdisciplinary approach involving computer scientists, traffic engineers, cybersecurity experts, economists, legal scholars, policymakers, and civil society. Without robust governance, adaptive regulatory frameworks, and mutually agreed-upon open standards, even the most advanced technological innovations will remain confined to academic publications and fail to realize their transformative potential on urban roadways [2].

4. Conclusion

The Multi-Agent Deep Reinforcement Learning (DRL) architecture supported by Edge Computing schemes presents a highly promising solution to address the volatility in post-pandemic surges of vehicle volume. However, its success remains hindered by three fundamental gaps. First, although models such as PPO and MADDPG demonstrate superior performance in simulation environments, these models remain highly vulnerable to distributional shifts caused by new mobility patterns, such as Work From Anywhere, which no longer conform to conventional historical data distributions. Second, while the theoretical implementation of Edge Computing can reduce latency and ensure system resilience during extreme surges, practical limitations related to computational power and energy consumption on edge devices in harsh intersection environments continue to pose unresolved challenges in the literature. Third, cybersecurity concerns particularly vulnerabilities to adversarial attacks and the absence of interoperability standards with legacy systems and adaptive regulatory frameworks have created a significant divide between theoretical advancements in laboratory settings and the feasibility of real-world deployment. Essentially, this study underscores that technical solutions alone are insufficient without a robust governance framework and an interdisciplinary approach encompassing energy economics and public policy.

The implementation of this research is directed toward designing a hybrid framework composed of three primary layers. At the computational layer, the development of lightweight DRL models based on knowledge distillation is proposed, enabling efficient operation on low-power Edge hardware while maintaining real-time adaptability to non-stationary post-pandemic traffic pattern shifts. At the data layer, a heterogeneous data fusion mechanism is implemented, incorporating not only SUMO simulation data but also integrating probe vehicle data and loop detector sensors to generate a more accurate real-time traffic state map. Finally, at the security and policy layer, the research outcomes are translated into recommendations for multi-layered security protocols to mitigate hostile attacks, alongside the formulation of technical interoperability guidelines. These guidelines are intended to assist transportation agencies in the gradual transition from conventional fixed-time signal systems to intelligent adaptive systems powered by artificial intelligence.

REFERENCES

[1] U. Azmi, K. I. Rosadi, and A. M. Zebua, "The Urgency Of Ihsan And Thinking Systemic In Increasing Motivation And Quality Of Education," *Jurnal Ilmiah Teunuleh*, vol. 3, no. 2, pp. 83–92, 2022, doi: 10.51612/teunuleh.v3i2.88.

[2] E. Prasetyo and D. W. Nugroho, *Tata kelola dan keamanan sistem transportasi cerdas*. Penerbit Salemba Teknika, 2025.

[3] B. Santoso and T. Wijaya, "Estimasi kerugian ekonomi akibat kemacetan lalu lintas di kawasan metropolitan Bandung Raya," *Jurnal Ekonomi Transportasi*, vol. 8, no. 2, pp. 78–95, 2024, doi: 10.25104/jet.v8i2.3456.

[4] E. Hidayat, A. Pardosi, and I. Zulkarnaen, "Efektivitas kepemimpinan kepala sekolah dalam penerapan Kurikulum Merdeka," *Jurnal Studi Guru Dan Pembelajaran*, vol. 6, no. 1, pp. 9–18, 2023.

[5] A. Wahyudi, S. F. Pratama, and H. Susanto, "Penerapan deep reinforcement learning untuk optimasi waktu siklus lampu lalu lintas adaptif pada simpang terisolasi," *Jurnal Ilmiah Teknologi Informasi dan Komputer*, vol. 9, no. 3, pp. 201–218, 2024, doi: 10.25126/jtiik.202496789.

[6] F. Rahman and I. G. A. M. Putra, "Edge computing untuk sistem transportasi cerdas: Sebuah tinjauan arsitektur dan tantangan implementasi di Indonesia," *Jurnal Teknologi Informasi dan Komunikasi*, vol. 14, no. 2, pp. 112–128, 2025, doi: 10.33021/jtik.v14i2.4567.

[7] A. P. Nugroho and S. Lestari, "Tinjauan sistematis implementasi kecerdasan buatan pada sistem transportasi cerdas di Indonesia: Peluang dan tantangan," *Jurnal Sistem Transportasi & Logistik*, vol. 5, no. 1, pp. 45–62, 2025, doi: 10.25292/jstl.v5i1.1123.

[8] W. R. Murhadi, *Metode penelitian: Pendekatan kuantitatif dan kualitatif*. Ekuitas Publisher, 2025.

[9] S. S. Pettalongi, M. Muas, A. Arafat, and D. N. Ndaomanu, *Metodologi penelitian kuantitatif, kualitatif, dan campuran: Teori dan praktik*. Media Penerbit Indonesia, 2025.

[10] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*, 2nd ed. MIT Press, 2022.

[11] S. Haykin and B. Widrow, *Deep reinforcement learning for dynamic systems: Algorithms and applications*. Oxford University Press, 2024.

[12] X. Liang, Y. Wang, J. Zhang, and L. Chen, *Multi-agent learning and coordination in intelligent transportation systems*. Springer, 2023.

[13] Y. Zhang, *Mobile edge computing*. Springer Cham, 2021.

[14] G. Russo and F. Palmieri, *Edge and fog computing for smart city applications: Architectures and challenges*. CRC Press, 2025.

[15] T. Iskandar, H. Purnomo, S. Lestari, and R. Wijayanti, *Komputasi tepi untuk sistem transportasi cerdas Indonesia*. Penerbit Andi, 2025.

[16] S. L. Hansen and R. D. McDonald, *Adaptive traffic control: From theory to sustainable practice*. Cambridge University Press, 2022.

[17] A. Widodo and I. Gunawan, *Keamanan siber pada infrastruktur kritis transportasi*. Penerbit Informatika, 2025.