

Article

Inclusive and Adaptive Traffic Management for Smart Cities: A Framework Combining Emergency Response and Machine Learning Optimization

Ioana-Miruna Vlasceanu ^{1,*} , João Sarraipa ² , Ioan Sacala ¹ , Janetta Culita ¹  and Mircea Segarceanu ¹

¹ Automatic Control and Systems Engineering Department, Faculty of Automatic Control and Computers, National University of Science and Technology Politehnica of Bucharest, Splaiul Independentei, No. 313, 060042 Bucharest, Romania; ioan.sacala@upb.ro (I.S.); janetta.culita@upb.ro (J.C.); mircea.segarceanu@upb.ro (M.S.)

² Faculty of Sciences and Technology & UNINOVA-CTS, Nova University of Lisbon, 2829-516 Caparica, Portugal; jfss@uninova.pt

* Correspondence: ioana.vlasceanu1512@upb.ro

Abstract

Smart control technologies that can manage the complexity of urban traffic while also reducing response times for emergency vehicles are necessary. This article proposes AETM (Adaptive and Equitable Traffic Management), an adaptive and equitable traffic management system that integrates contextual methods for handling emergencies with traffic light control based on reinforcement learning. The system uses Q-learning to optimize traffic light phases under normal traffic conditions and integrates a dedicated emergency vehicle module, which includes detection, dynamic rerouting and contextual preemption functions. The system adaptively optimizes traffic light phases under normal traffic conditions and integrates a specialized module for emergency vehicles, which ensures their detection, dynamic rerouting and contextual preemption. The priority level is evaluated through an auxiliary fuzzy mechanism, based on interpretable rules, which takes into account local conditions without influencing the learning process. The performance of the framework is evaluated in a microscopic simulation environment by comparing classical control, adaptive control, and the full AETM configuration. The results highlight significant reductions in travel times and stops for emergency vehicles while maintaining overall traffic stability.

Keywords: emergency prioritization; adaptive signal control; traffic automation; reinforcement learning; intelligent transport systems



Academic Editor: Isam Shahrour

Received: 3 December 2025

Revised: 13 January 2026

Accepted: 17 January 2026

Published: 2 February 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article

distributed under the terms and

conditions of the [Creative Commons](#)

[Attribution \(CC BY\)](#) license.

1. Introduction

As cities grow and people move around more, managing traffic has become harder. Keeping roads safe, efficient, and fair is now a major challenge. Old systems that use fixed traffic light schedules or only react to problems cannot keep up with changing traffic patterns. This leads to more congestion, longer wait times, and wasted resources in cities. These issues get even worse during emergencies, when first responders and ambulances are already delayed by traffic. Research shows that emergency vehicles are often held up because there is not enough smart coordination or dynamic prioritizing at intersections, making it harder to respond quickly.

Intelligent transport systems (ITS) have become a promising technological framework for improving urban mobility in recent years. They can support real-time monitoring,

adaptive traffic light control, and communication between vehicles and infrastructure. Recent methodologies advocate for the amalgamation of VANET (Vehicular Ad-hoc Network), V2X (Vehicle to Everything), or IoT (Internet of Things) as instruments for traffic surveillance and ambulance prioritization, underscoring substantial decreases in travel durations within simulated environments [1]. However, most of these solutions only focus on local preemptions at intersections or isolated routing. They don't have a consistent way to combine global network optimization with the needs of emergency vehicles. Many implementations also have structural problems, such as being dependent on costly infrastructure, not being able to quickly adapt to changes in traffic, or not having strong mechanisms for coordinating multiple intersections, as recent studies [2,3] have shown. Figure 1 provides a conceptual illustration of an ITS-enabled urban intersection, emphasizing the interaction between connected vehicles, smart traffic signals, and pedestrians through V2X and IoT-based communication.



Figure 1. Conceptual representation of an Intelligent Transportation System [4].

On the other hand, advances in artificial intelligence have opened new directions for traffic light control, especially through the use of Reinforcement Learning (RL) algorithms. Models based on Q-learning or deep-RL variants have demonstrated the ability to learn optimal policies in complex road networks, reducing waiting times and improving traffic flow [5]. Although these techniques are increasingly used for general optimization, their integration with mechanisms dedicated to emergency vehicles is still insufficiently explored, and interoperability with real-time routing remains an open challenge.

New directions in traffic management combine adaptive control methods with modern monitoring and communication systems. Some studies use learning from experience to recognize unusual situations in the network and to strengthen the exchange of information between vehicles and infrastructure, thus contributing to safer operation [6]. Regarding emergency vehicles, solutions based on VANET communications show that the direct connection between ambulances and traffic lights can significantly shorten the travel time to the destination [7]. At the same time, studies using multiple decision agents, sometimes supported by video analysis, show that closer coordination between intersections can lead to more fluent traffic in variable situations [8].

At the same time, fuzzy logic continues to represent an appropriate solution for scenarios in which uncertainty, variability of flows and the lack of precise traffic models require flexible and interpretable control. Studies applied to traffic management demonstrate that fuzzy systems can alleviate local overloads and generate smooth transitions between

traffic light states, avoiding sudden behaviors that affect network stability [3]. However, most of the fuzzy implementations analyzed in the recent literature do not include advanced prioritization mechanisms for ambulances and are not tested in realistic severe congestion scenarios.

In this context, an integrated methodological framework is needed, capable of combining emergency vehicle detection, optimal routing, adaptive intersection control and mechanisms to ensure equity of access to mobility. Although some works propose optimization strategies inspired by IoT, PIC-based control (Programmable Interface Controller) or sensor-based systems for ambulance management, these are often limited to specific hardware configurations, without providing an extensible, scalable and compatible approach with modern ITS architectures. Figure 2 illustrates the urban intersection context considered in this work, highlighting the coexistence of multiple traffic participants under signalized control.




Figure 2. Conceptual illustration of a signalized urban intersection.

2. State of Art

2.1. Adaptive Traffic Optimization Through Machine Learning and IoT Architectures

Recent developments in intelligent traffic management have increasingly depended on machine learning and Internet of Things technologies to enable adaptive signal control and real-time congestion forecasting. Lilhore et al. [9] propose an IoT-based adaptive traffic-management system in which streaming sensor data is clustered using the DBSCAN algorithm to identify anomalous congestion patterns, and signal timings are adaptively modified according to detected density levels. The approach demonstrates a notable reduction in vehicle waiting times and queue lengths when compared with fixed-time control, using MATLAB simulations and microcontroller-based hardware prototypes. Similarly, Sahil and Sood [10] introduce an edge–cloud framework that integrates vehicular IoT nodes and fog computing layers to predict inflow using J48 decision-tree models and to dynamically redistribute green-phase durations across multiple intersections. Their architecture is modular—composed of sensing, edge, cloud, and application layers—allowing distributed inference near the roadside units to minimize latency.

De Luca et al. [11] extend the optimization perspective by combining traffic control (TC) with route guidance (RG) in a hybrid within-day and day-to-day dynamic traffic-assignment model calibrated for the Naples urban network. Their results show a measurable improvement in travel times and pollutant emissions, validating the benefit of combining control and behavioral adaptation in one loop. Despite their methodological

diversity—clustering, supervised ML, hybrid DTA (Dynamic Traffic-Assignment)—all of these studies share a strong emphasis on efficiency metrics such as queue length, delay, or throughput. They remain validated mainly in simulation environments and seldom integrate modules for emergency-vehicle prioritization or for assessing equity among different road-user groups. Thus, the optimization focus is predominantly technical, with limited consideration of social or operational inclusivity. Further advancements have increasingly integrated Deep Learning (DL) and computer vision technologies into traffic surveillance systems. Although the model illustrated in Figure 3 originates from large-scale traffic management and route guidance studies, it is referenced here solely to illustrate the general interaction between route choice and control mechanisms, not to suggest a highway-oriented application.




Figure 3. Day-to-Day route choice and compliance model for user response to route guidance [11].

Study [12] proposes a CNN-based classifier (CNN - Convolutional Neural Network) combined with spatio-temporal feature extraction to detect crash-like events from camera feeds. The pipeline uses convolutional layers to capture localized visual patterns and LSTM-like (Long Short-Term Memory) temporal modeling to distinguish between continuous flow and sudden disruptions, achieving high accuracy in controlled datasets. However, despite strong detection performance, its integration with real-time control actions—such as signal preemption or diversion strategies—is absent, limiting system usefulness to mere alerting rather than active mitigation. There is related research on sensor-based intelligent traffic management that employs the Internet of Things and traditional machine learning techniques [13] to assess congestion levels in specific areas. Although this approach is more computationally manageable and practical compared to deep learning systems, it primarily emphasizes classification instead of adaptive decision-making.

In parallel with these directions based on supervised learning and data analysis, intelligent traffic light control through reinforcement learning has become a central research direction in urban traffic management. Recent works [14–16] have explored both formulations for individual intersections and multi-agent extensions for intersection networks, demonstrating the ability of RL to learn control regulations promptly from the interaction with the environment. Network-level cooperation mechanisms and control based on the concurrency of traffic light phases are approaches that allow capturing spatio-temporal dependencies and distributing commands between adjacent intersections. The results from the mentioned works show notable reductions in waiting times and queue lengths, especially in dense urban networks or under conditions of variable demand.

However, most of these solutions predominantly focus on optimizing aggregate traffic performance, using objectives such as average delay or throughput, and involve complex models with high training and coordination requirements between agents. There are, however, a number of actions that are not taken into account much, such as clear operational constraints, the integration of emergency vehicle prioritization or equity considerations. As these aspects are not often addressed in the same system, the resulting

decision mechanisms are often more difficult to implement in concrete, real cases. The need for architectures that ensure robustness, transparency and compatibility in the case of urban adaptive control systems still remains a point of interest. For each situation, a pragmatic and integrable solution must be implemented, as agreeable as possible in terms of complexity and that manages both general traffic and explicit constraints such as emergency situations.

2.2. Emergency and Incident-Response Integration

Another stream of research examines the management of non-recurrent congestion caused by accidents or emergencies. Farrag et al. [17] developed the Smart Traffic Incident Management Framework (STIMF), which combines a fuzzy logic decision-support system with a microsimulation environment (VISSIM) to assist control center operators in determining the most effective responses to traffic incidents. The framework evaluates candidate measures—dynamic hard shoulder running, variable speed limits, ramp metering, and V2X communication—through multi-criteria indicators (mobility and safety), generates fuzzy rules from simulation data, and then recommends ranked strategies in real time via MATLAB's Fuzzy Logic Toolbox. Applied to a 20 km section of the Muscat Expressway, STIMF achieved clear mobility and safety benefits by automating decision support, yet it relies on expert-defined rule bases and does not fully embed adaptive learning or automated coordination with the city's wider traffic-optimization systems.

Comparable decision-support approaches often focus on isolated subsystems—either incident detection or signal preemption—without linking them to network-level flow control. Consequently, even when response times improve locally, overall traffic equilibrium or fairness among corridors is not considered. This indicates a persistent fragmentation between emergency-response management and global traffic optimization. A second cluster of works advances hybrid ML + classical optimization for traffic prediction, short-term forecasting, and adaptive signal timing. The work Intelligent Traffic Management using IoT and Machine Learning [18] employs a combination of KNN regression (K-Nearest Neighbors), Naïve Bayes, and linear models to forecast densities and inform green-phase adjustments. Despite providing a lightweight model appropriate for embedded systems, the architecture lacks consideration for special traffic classes such as emergency vehicles and does not include fairness-aware priority rules.

The paper of Hu et al. [19] introduces a method to partition large-scale road networks using graph cuts, spectral clustering, and multi-level partitioning heuristics to improve distributed congestion control. The framework supports faster routing and decentralized signal optimization by creating dynamically reconfigurable sub-networks. However, the partitioning criterion is strictly based on flow and topology, without introducing emergency accessibility or equity-of-access requirements as constraints. Altogether, these hybrid works enhance predictive accuracy and scalability, yet still treat all vehicles uniformly, ignoring the operational asymmetry introduced by EMS (Emergency Medical Services) and vulnerable road users.

Another group of works addresses the prioritization of emergency vehicles through dedicated traffic light preemption mechanisms. Dynamic right-of-way approaches use information about the position and trajectory of emergency vehicles to activate coordinated interventions on traffic lights in the network, reducing travel times on critical routes [20]. Another contribution on this subject is presented in [21] where solutions for improving traffic detection and preemption decisions through V2I (Vehicle to Infrastructure) communication are explored. Pai et al. [22] suggest adaptive systems that change traffic light cycles based on how far ambulances are from intersections and current traffic conditions. However, these works mainly aim at local performance regarding emergencies, prioritizing

a separate mechanism for traffic control. This exclusion of integration into a large-scale urban architecture and the impact on the entire network limits their applicability in reality.

2.3. Integrated Modeling Frameworks for Active Traffic Management

At the macroscopic scale, the integration of control and demand has been examined through agent-based models and dynamic traffic-assignment approaches. Xiong et al. [23] put forward an Integrated Modeling Framework for Active Traffic Management, which merges the Agent-Based Behavior Model (AgBM) with DTA-Lite. This framework simulates how travelers respond to real-time control measures, including variable speed limits (VSL) and ramp metering (RM). It operates in a feedback loop—collecting data, updating behaviors, and reaching system convergence—to effectively capture the interactions between driver adaptation and network performance. As a result, it achieves high-resolution predictions across the Washington D.C. area. Figure 4 illustrates the framework proposed by these authors for active traffic management at a macroscopic scale, in which DTA is coupled with an agent-based behavioral model. Through a feedback loop, the system identifies user paths and applies control measures such as variable speed limitation and on-ramp adjustment, iteratively updating network conditions and traveler behavior until convergence.




Figure 4. Integrated framework for real-time active traffic management in AgBM-DTALite [23].

The paper Autonomous Driving in the Framework of Three-Phase Traffic Theory [24] provides a theoretical investigation of how ACC (Adaptive Cruise Control) and automated-driving behaviors alter traffic stability, breakdown probability, and jam-formation mechanisms. Using three-phase theory dynamics, the authors show that poorly tuned ACC controllers can increase breakdown probability, even while smoothing individual trajectories. This is critical for emergency scenarios: automated vehicles might unintentionally amplify congestion when yielding to preemption corridors unless coordinated in a city-wide system. While rich in theory and simulation-based insights, this research direction remains decoupled from multi-objective traffic management and does not embed EMS-aware control policies.

The strength of these frameworks lies in their ability to realistically depict driver behavior and evaluate various active traffic management strategies under changing conditions. However, they primarily focus on recurrent congestion and improving efficiency and do not address issues related to emergency logistics or social inclusivity. Importantly, there are no specific modules for emergency-vehicle preemption, and fairness or accessibility

metrics are absent from the optimization goals. In short, these frameworks are technically integrated but socially blind.

2.4. Inclusivity, Equity, and Governance Metrics in Smart Mobility

Parallel to the algorithmic progress, a smaller but growing body of research emphasizes equitable and inclusive urban mobility. Kaparias et al. [25] extend the CONDUITS_DST (City-On-Network Data and Urban Indicators for Transport Sustainability—Decision Support Tool) framework to introduce Key Performance Indicators (KPIs) spanning efficiency, safety, environmental impact, and social inclusion. The tool aggregates microscopic simulation data to evaluate policy scenarios, enabling decision-makers to balance technical performance with social sustainability. Similarly, Alsaawy et al. [26] present an IoT- and data-analytics-based four-layer framework (sensing, fog, cloud, service) for congestion monitoring that incorporates governance and interoperability considerations.

While these works successfully operationalize “equity” as an assessment dimension—using accessibility and exposure indicators—they typically treat it as a post-hoc evaluation metric rather than as an active control variable within the real-time management process. None embeds equity safeguards, privacy-by-design, or fairness constraints directly into the traffic-control logic, leaving a gap between ethical intent and algorithmic implementation.

The article in [27] focuses on linking traffic management with environmental KPIs, such as CO₂, NO_x, delay-based fuel metrics, and eco-efficiency. It provides system-level evaluation of strategies such as VSL, dynamic routing, and lane-management, validating them against sustainability metrics via microsimulation. However, just like CONDUITS_DST, the environmental KPIs are merely evaluative, not embedded in the real-time decision-making logic. Moreover, emergency operations and equity considerations remain outside the scope, showing once again the gap between assessment and actionable control.

Antipova et al. [28] highlight the link between transport network configuration and differential access to opportunities, arguing for the need to include equity criteria from the design phase of mobility policies. Similarly, Pereira et al. [29] show that the results of equity assessments depend significantly on the choice of indicators and time thresholds, even in situations where global performance indicators improve. The paper referenced in [30] offers a more recent approach to the use of intelligent algorithms in traffic management. Without considering distributional impacts, these techniques can worsen preexisting systemic imbalances, as pointed out by the authors. Nevertheless, the necessity of operational frameworks that are capable of converting these principles into judgments that are both automated and transparent is still being brought to light.

2.5. Secure and Distributed Architectures for Traffic Control

Regarding the reliability and trustworthiness of intelligent transportation infrastructures, a series of research studies propose the integration of fog computing mechanisms with blockchain technologies, to ensure secure and low-latency communication between sensors, control centers and vehicles. A relevant contribution in this regard [31] describes a layered IoT–Fog–Cloud architecture, in which blockchain ledgers are used for authentication and guaranteeing the immutability of data distributed between nodes. The framework aims to enhance network transparency and fault tolerance through the implementation of smart contracts and distributed consensus mechanisms. This approach is intended to facilitate collaboration among vehicles in densely populated networks. While these designs address significant issues related to data integrity and latency, they still fall short in linking security protocols with adaptive traffic control systems and emergency prioritization strategies. The blockchain component functions primarily as a middleware for trust, not as an enabler of responsive decision-making. As a result, existing secure frameworks strengthen

the infrastructure layer but remain loosely connected to higher-level control policies and fairness auditing—areas increasingly critical in inclusive smart-city contexts.

Sefati et al. [32] propose a framework to improve the resilience of smart infrastructures. Their approach uses clustering with GANs (Generative Adversarial Network) and adaptive routing with Ant Colony Optimization to control congestion in wireless sensor networks. The system addresses congestion detection, notification, and mitigation, aiming to keep quality of service high by balancing loads, choosing strong routes, and reducing energy use. Although this system significantly improves the reliability of the sensor layer and illustrates how ML techniques can prevent communication degradation under high load conditions, it does not define a direct link between WSN (Wireless Sensor Network) resilience and adaptive traffic control strategies or preemption mechanisms for emergency vehicles. Consequently, the solution reinforces the information base of the ITS infrastructure, but remains, like other similar approaches, disconnected from the higher level of decision-making and operational continuity policies in critical scenarios.

2.6. Routing and Multi-Objective Optimization Approaches

In the field of multi-objective routing, Paricio and López-Carmona [33] propose the Traffic-Weighted Multimap (TWM) model, which addresses the problem of route allocation by defining distinct weight maps for different vehicle categories. Each type of fleet is associated with a specific cost function, and the routes are determined by A* and Dijkstra algorithms adapted for multi-criteria optimization, so that the result simultaneously reflects travel time, congestion level and penalties related to the use of overloaded segments. This structure allows for the efficient management of interactions between fleets and leads to a more uniform distribution of flows, preventing the overuse of certain arteries of the network. However, the framework does not include a prioritization mechanism for emergency vehicles and does not integrate constraints related to fair access or non-discriminatory treatment of users, objectives that remain external to its optimization function. Although Traffic Wave Management (TWM) provides a methodologically coherent framework for distributing traffic and alleviating congestion on certain segments, its focus is primarily on optimizing network efficiency. Without dedicated prioritization mechanisms for emergency vehicles or clearly defined constraints regarding access equity, its effectiveness in an inclusive urban setting is somewhat limited.

The contribution in Applied Soft Computing [34] introduces a multi-objective optimization architecture integrating ML predictors and reinforcement elements for controlling signals under multiple competing objectives (delay minimization, throughput maximization, stability). The authors utilize heuristic search and ML-based state estimation; however, the optimization does not include EMS priority handling or any fairness constraints. Similarly, paper [35] explores adaptive strategies that combine optimization algorithms with real-time sensing for signal phasing. Despite strong optimization performance, both frameworks share a focus on aggregate flow efficiency, reinforcing the broader pattern that emergency logistics and mobility equity remain unaddressed.

Another important area of research looks at using reward learning to add clear rules for safety, risk, or key functions. The synthesis study carried out in [36] highlights a clear transition from classic paradigms oriented exclusively on performance (waiting time, throughput) to multi-objective approaches. The approach consists of introducing further restrictions either at the reward function level or employing constrained optimization mechanisms. However, the constraints addressed are treated in an implicit manner and do not explicitly distinguish between classes of users or types of vehicles.

The work [37] has as its main objective the assimilation of percentile risk criteria or CVaR (Conditional Value at Risk). In terms of scalability, reference [38] recommends

distributing the learning process across numerous intersections, decreasing processing burden and facilitating network-level coordination. The processing cost of such approaches makes them unsuitable in the short term, but they provide a strong mathematical model and are very resilient; nonetheless, this is irrelevant to traffic fairness. Figure 5 illustrates a representative deep reinforcement learning control loop for urban traffic signals, included here to highlight the complexity and learning paradigm of federated DRL (Deep Reinforcement Learning) approaches, in contrast to the lightweight and interpretable control strategy adopted in AETM.




Figure 5. Schematic representation of a deep reinforcement learning-based traffic signal control framework from [38].

Therefore, these investigations demonstrate that limited and multi-objective RL can contribute to improved traffic management safety and stability. They also highlight the disconnect between complex theory and practical, easy-to-understand solutions. The AETM architecture combines a simple RL mechanism for optimization with a clear fuzzy module. This setup manages priorities and contextual constraints in a way that is practical and easy to audit.

2.7. Synthesis and Identified Research Gaps

Overall, the theoretical and methodological framework provided by these works is robust, especially in terms of their ability to realistically represent the behavior of road users and to evaluate a wide range of active traffic management strategies under dynamic conditions. The models analyzed are well-suited to study the trade-offs between travel time, congestion level and network capacity utilization, which makes them valuable in the design and testing of control scenarios. However, their application remains partial from the perspective of a complete urban system, as dimensions such as emergency response logistics and social inclusion are generally either marginally addressed or completely absent from optimization formulations. Explicit integration of emergency vehicle preemption modules, as well as equity and accessibility indicators directly into the objective functions, would represent a natural extension of these models, transforming them from predominantly efficiency-oriented tools into components of more comprehensive traffic management solutions, capable of responding to both the operational and social requirements of urban transport networks.

Large-scale modelling approaches—such as agent-based behavioral models, dynamic traffic assignment, and network partitioning—improve realism and scalability, but they do not explicitly incorporate emergency-vehicle prioritization or constraints related to accessibility and equity. At the same time, works addressing sustainability, governance, or social inclusion introduce relevant evaluation metrics but do not embed these metrics

into real-time control logic. Security- and reliability-focused architectures strengthen the underlying infrastructure, yet they remain disconnected from operational decision-making, especially in emergency contexts. Similarly, deep-learning-based incident-detection systems achieve high detection accuracy but stop short of linking detection to automated control actions such as preemption or rerouting.

3. Materials and Methods

3.1. General Considerations and Theoretical Foundations of the Proposed Architecture

The AETM (Adaptive and Equitable Traffic Management) architecture is built as an integrative framework designed to ensure adaptive, efficient and equitable urban traffic management, with the ability to simultaneously respond to operational requirements, emergency situations and social objectives of modern mobility. In the context of urban road networks, where traffic dynamics are influenced by rapid demand variations, heterogeneous driver behaviors and external disturbances, the AETM architecture introduces a flexible decision-making mechanism based on real-time analysis, fuzzy inference and contextual prioritization.

At the heart of the framework is the idea of multi-objective optimization, in which each traffic signal decision represents a calculated trade-off between the overall network performance, the response time for emergency vehicles and maintaining an adequate level of operational equity for all users. This relationship can be conceptualized by an aggregate cost function:

$$J = \alpha C_{\text{traffic}} + \beta C_{\text{EMS}} + \gamma C_{\text{equity}} \quad (1)$$

The coefficients α , β and γ are established through urban policies and define the relative importance given to each component. The formula expresses the general structure of the decision, without imposing an exhaustive mathematical modeling, and reflects the fact that AETM does not pursue a single unilateral objective, but seeks the balance between efficiency, rapid response and fairness.

The architecture abandons a strictly formal description of the network state, preferring instead a conceptual model in which traffic is interpreted through the relationships between density, flow speed, emergency vehicle position and intersection demand. In this framework, fuzzy inference plays a central role, as it allows the robust interpretation of noisy and incomplete data, using flexible linguistic categories and rules. This mechanism provides a natural way of dealing with uncertainty, a fundamental characteristic of the real urban environment.

Figure 6 shows the AETM framework which aims to manage urban traffic using a layered, closed-loop control system. Each layer has a clear task, such as collecting data, optimizing traffic flow, handling emergencies, or carrying out decisions. Dividing the system in this way helps it assess situations more effectively and makes future improvements easier to add.

The first layer collects and combines data from sources like sensors, connected infrastructure, and the locations of emergency vehicles. It filters and organizes this information over time to give a clear view of current traffic conditions. The adaptive optimization layer is at the core of the system. It uses Q-learning to create a traffic light control policy. This layer monitors each intersection, tries different actions based on a reward function, and updates its strategy as traffic and congestion change.




Figure 6. Layered architecture of the Adaptive and Equitable Traffic Management Framework.

The emergency response layer supports EMS vehicles using two main tools. First, it dynamically reroutes ambulances by recalculating their routes with a shortest-path algorithm that uses real-time traffic data. Second, it prioritizes ambulances at intersections with a rule-based module that adjusts their priority based on local factors like distance and congestion. This module does not learn or optimize for the whole system, but instead fine-tunes its response for each situation. In real time, the operational execution layer alters the settings of the traffic lights in accordance with safety and stability requirements, putting decisions into action—even when EMS vehicles are given precedence.

The architecture also includes components for operational equity, governance, and transparency. These parts monitor the effects of interventions and help audit decisions. They do not directly control the system, but they help prevent unfair impacts on specific groups of users.

3.2. Layer 1—Urban Data Acquisition and Integration

The first layer of the AETM architecture is responsible for the acquisition, unification and filtering of data describing the state of the traffic network. In a contemporary urban context, characterized by high density, pronounced variability and multiple sources of information, the quality and coherence of data are essential prerequisites for any adaptive control mechanism. This layer functions as an interface between the physical infrastructure of the city and the higher levels of the architecture, ensuring the semantic coherence of the observations and their availability in real time.

In the urban environment, data is generated by a variety of devices and subsystems, including traditional road sensors, radar equipment, video cameras with intelligent processing, weather stations, connected vehicles and, especially, emergency vehicles. Each of these sources provides information with different granularity, frequency and accuracy, which is why the AETM system must operate in the presence of a significant degree of heterogeneity. To express in a general way the relationship between this data and the real state of traffic, the acquisition layer uses the standard model of noisy observations, in which what is measured represents an approximation of the real state, affected by perception and transmission errors:

$$Z(t) = H(X(t)) + \eta(t) \quad (2)$$

In this conceptual framework, $Z(t)$ represents the data actually collected by the sensors, $X(t)$ denotes the real state of the network (which cannot be directly observed in its entirety), $H()$ is an observability function that reflects the limitations of each type of sensor, and $\eta(t)$ represents the noise associated with the measurement. The formula is used for describing in a unified way the fact that any urban infrastructure produces imperfect observations, which require pre-processing to be used in higher-level analysis modules.

An important aspect of this layer is the process of temporal and semantic data harmonization. Since the transmission rate differs from one sensor to another, the AETM architecture uses fixed temporal windows in which it aggregates the available information, building a coherent representation of the network state at that moment. In parallel, statistical filtering procedures are applied to eliminate outliers, as well as interpolation algorithms to fill in any gaps in the data streams. This step ensures that the final value transmitted to the analytical layer is internally consistent and free from artificial fluctuations that could destabilize the fuzzy inference mechanism.

The acquisition layer also treats data from emergency vehicles in a distinct way, as these are not only traffic observations, but also direct indicators of operational intent. The position, speed and direction of the ambulance are updated at short intervals and form the basis for the proximal assessment in the analysis layer. The AETM architecture thus introduces a logical separation between the general network data and those of critical vehicles, without treating them in isolation, but integrating them into a unified decision model.

Finally, this layer plays the role of guarantor of information quality. Through data filtering, aggregation, and standardization, it provides a robust view of the network state. Because of this, the system is able to function well in actual settings. The whole architecture's performance would suffer if this connecting step wasn't taken; higher-level modules such as preemption, fuzzy analysis, and fairness mechanisms would be forced to handle data that is inconsistent or contradictory.

3.3. Layer 2—Prediction, Analysis and Machine Learning

Q-learning is a reinforcement learning method that improves its decisions by interacting with the environment many times. The agent does not use a set model of traffic. It looks at the current state of the intersection, picks a traffic light action, and gets a reward based on how much it improves things like waiting time or congestion. Over time, by updating its value function $Q(s,a)$, the agent learns a stable way to manage traffic that fits local conditions. In AETM, Layer 2 serves as the primary component responsible for managing optimization. Q-learning optimizes the timing of traffic lights as necessary. In the event of an emergency, Layer 3 has the capability to modify or override decisions within that region; however, the fundamental function of reinforcement learning in managing routine traffic remains unchanged.

In the implementation, the state of each traffic light is discretized into a compact 4-component tuple $(q_{bin}, w_{bin}, p_{bin}, t_{bin})$, built from: the difference/coherence of the queues on the competing directions, the waiting time level, the current phase group and the time spent in the phase. This discretization reduces the state space and allows robust online learning.

The action space is intentionally minimal and interpretable: $A = \{\text{HOLD}, \text{SWITCH}\}$, maintaining the current phase or switching to the opposite group. Decisions are not applied at each step, but at a fixed interval (in this case it was set to 4 steps), to avoid oscillations. In addition to online updating of the Q values associated with the current transition, the agent also uses an experience replay mechanism. Replay-based updates are performed periodically, every five simulation steps, in order to increase the stability of the learning process and improve its efficiency.

The reward is formulated to discourage congestion and delays by introducing “pressure” terms associated with waiting gaps and waiting time variations, while incentivizing traffic flow dissipation through metrics such as throughput and average speed. In addition, when the emergency system influences the operation of a TLS, the reward is adjusted slightly to reflect the fact that the reinforcement learning decision has been conditioned or partially replaced by the EMS logic.

The process by which the control policy is learned is based on the progressive updating of the state-action value function $Q(s, a)$, which reflects how advantageous it is to apply an action a in a given state s . In the proposed implementation, this function is not updated directly, but through a recursive relationship, equivalent to the classical Q-learning rule, presented below.

$$Q_{t+1}(s, a) = (1 - \alpha) Q_t(s, a) + \alpha \left[r_t + \gamma \max_{a'} Q_t(s', a') \right] \quad (3)$$

where the terms have the following meaning:

- $Q_t(s, a)$ represents the estimated value, at time t , of action a in state s , according to the current policy;
- $Q_{t+1}(s, a)$ is the updated value of the same state-action pair after observing the transition;
- $\alpha \in (0, 1]$ is the learning rate, which controls how much newly obtained information influences the current estimate; higher values allow for rapid adaptation, while lower values favor stability;
- r_t is the immediate reward received after applying the action, calculated based on traffic performance (e.g., latency, congestion, or throughput);
- $\gamma \in [0, 1)$ is the discount factor, which expresses the importance of future rewards over immediate ones; a γ close to 1 favors long-term oriented policies;
- s' is the resulting state after applying action a ;
- $\max_{a'} Q_t(s', a')$ represents the best estimate for the next state, according to the current policy.

This formulation highlights the fact that the value of Q is not replaced directly, but is gradually adjusted by combining the existing estimate with a newly calculated learning target, based on the reward obtained and on the information provided by the next state. In the implementation, the estimation error is limited (clipped) to avoid sudden variations and to maintain the numerical stability of the learning process, especially in very dynamic traffic contexts.

The Q-learning algorithm is trained directly during the traffic simulation, without including a separate offline training stage. This option is consistent with the main goal of the AETM framework, which aims at continuous adaptation to variable and difficult-to-predict traffic conditions typical of the real urban environment.

The learning process is structured in two stages. In the first stage, the warm-up, the agent explores the environment and collects state–action–reward transitions, without directly influencing the control strategy. The role of this phase is to ensure the accumulation of a sufficiently varied initial experience and to avoid premature updates of the Q function, which could be affected by a small number of observations.

Subsequently, the policy update is not activated suddenly, but introduced gradually. The agent follows an ϵ -greedy exploration–exploitation strategy, in which, with a probability ϵ , an action is chosen at random, and in the remaining cases the action that maximizes the value of Q is selected. During the simulation, the value of ϵ is gradually reduced, ensuring a progressive and controlled transition from exploration to exploitation of the learned policy.

For clarity and reproducibility, Algorithm 1 summarizes the control logic of the reinforcement learning layer, highlighting the interaction between state observation, action selection, reward computation, and policy update.

Algorithm 1 Adaptive traffic light control based on Q-learning (Layer 2), with handover to EMS

Require: Set of controlled intersections TLS – *Traffic Light Signal*, hyperparameters α, γ, ϵ , decision interval K , warm-up W , replay update frequency U

Ensure: Action-value table $Q(s, a)$ and induced greedy policy $\pi(s) = \arg \max_a Q(s, a)$

```

1: Initialize  $Q(s, a) \leftarrow 0$  for all states  $s$  and actions  $a$ 
2: Initialize ReplayBuffer with prioritization (PER - Prioritized Experience Replay)
3: Set  $\epsilon \leftarrow \epsilon_{\text{start}}$ 
4: for  $t \leftarrow 1$  to  $T$  do
5:   (Optional) update the EMS system (detection, rerouting, pre-emption)
6:   for all  $tls \in TLS$  do
7:     Observe the current state  $s_{tls} \leftarrow discretize\_state(tls)$ 
8:     if  $t \bmod K = 0$  then
9:       if  $EMS\_influence(tls)$  then
10:         $a_{tls} \leftarrow \text{HOLD}$ 
11:       else
12:        if  $rand() < \epsilon$  then
13:           $a_{tls} \leftarrow \text{random action from } \{\text{HOLD}, \text{SWITCH}\}$ 
14:        else
15:           $a_{tls} \leftarrow \arg \max_a Q(s_{tls}, a)$ 
16:        end if
17:      end if
18:      Apply action  $a_{tls}$  to the traffic light, respecting operational constraints
19:    end if
20:  end for
21:  Run a simulation step (advance environment)
22:  for all  $tls \in TLS$  do
23:    Observe the next state  $s'_{tls} \leftarrow discretize\_state(tls)$ 
24:    Calculate the reward  $r_{tls} \leftarrow reward(tls)$ 
25:     $\delta_{tls} \leftarrow r_{tls} + \gamma \max_{a'} Q(s'_{tls}, a') - Q(s_{tls}, a_{tls})$ 
26:     $Q(s_{tls}, a_{tls}) \leftarrow Q(s_{tls}, a_{tls}) + \alpha \cdot clip(\delta_{tls})$ 
27:    Store  $(s_{tls}, a_{tls}, r_{tls}, s'_{tls})$  in ReplayBuffer with priority  $\propto |\delta_{tls}|$ 
28:     $\epsilon \leftarrow \max(\epsilon_{\min}, \epsilon \cdot \epsilon_{\text{decay}})$ 
29:  end for
30:  if  $t \geq W$  and  $t \bmod U = 0$  then
31:    Run additional updates from ReplayBuffer (PER)
32:  end if
33: end for
34: return  $Q$  (policy  $\pi$  is induced as  $\pi(s) = \arg \max_a Q(s, a)$ )

```

3.4. Layer 3—Emergency Response Module and Adaptive Prioritization

The emergency response module represents the operational center of the AETM architecture, since it is in it that the information interpreted in the analytical layer is converted into control decisions with a direct effect on the traffic light infrastructure. This module manages situations in which an emergency vehicle, such as an ambulance, approaches an intersection and needs a priority route to reduce the intervention time. Unlike traditional systems that grant preemption mechanically, regardless of the context, AETM implements an adaptive decision-making mechanism, in which the real need for preemption is rigorously evaluated, and its effects on the network are monitored in real time.

Within this module, preemption is activated only when there is a clear operational benefit. This evaluation is based on how long the emergency car would need to go the whole route without any help, compared to how long it would need to go the whole route with a better traffic signal. This idea can be expressed by the simplified formula:

$$\Delta T_{\text{EMS}} = T_{\text{normal}} - T_{\text{preempt}} \quad (4)$$

where ΔT_{EMS} represents the time gain, T_{normal} the estimated crossing time under normal conditions, and T_{preempt} the time resulting from the traffic light adjustment. The formula is not used for complex calculations, but functions as a conceptual expression that justifies the fact that preemption should be activated exclusively in situations where it produces a net advantage. If this difference is negligible or null, the architecture avoids the introduction of preemption, so as not to generate unjustified delays for other network users. In practice, however, the decision to activate and modulate preemption also depends on contextual factors that cannot be captured by a single scalar criterion. The AETM architecture includes a fuzzy auxiliary module to support contextual decision-making for emergency vehicle prioritization. This module estimates the priority level for emergency vehicles at intersections using a limited set of contextual variables.

The fuzzy system is built around two continuous input variables: the distance of the ambulance from the intersection and the local level of congestion. Each of these is modeled by Gaussian membership functions, which ensure smooth transitions between the different value regions. The distance is divided into the categories Near, Medium and Far, and the level of congestion into Low, Medium and High. The output variable corresponds to the priority level given to the emergency vehicle—Low, Medium or High—which is in turn represented by Gaussian membership functions, defined on a discretized scale. The set of these functions is illustrated in Figure 7.

The rule base (Table 1) is deliberately kept small and easy to follow, being designed to reflect intuitive operational situations. For example, when the ambulance is in the immediate vicinity of the intersection and the level of congestion is high, the system assigns a high level of priority; on the other hand, for longer distances or in conditions of flowing traffic, the priority granted is lower. Such a formulation allows for continuous and gradual adjustment of preemption decisions, reducing the risk of abrupt changes in traffic light phases. Finally, this fuzzy mechanism has a strict interpretable decision support role for contextual prioritization of emergency vehicles, without intervening in the reward learning process or in the convergence of the adaptive control based on Q-learning described in Layer 2.

Table 1. Rule base for fuzzy EMS prioritization module.

Rule	Distance	Congestion	Priority Level
1	Near	High	High
2	Near	Medium	High
3	Near	Low	Medium
4	Medium	High	Medium
5	Medium	Medium	Medium
6	Medium	Low	Low
7	Far	High	Low
8	Far	Medium	Low
9	Far	Low	Low




Figure 7. Gaussian Membership Functions for distance, congestion and priority.

Figure 7 explicitly illustrates the Gaussian membership functions used in the fuzzy decision support module for emergency vehicle prioritization. This module uses a Mamdani-type fuzzy inference system, which is chosen because it is easy to interpret and works well for decisions in uncertain urban settings. The input variables—the distance of the emergency vehicle from the intersection and the local congestion level—are fuzzified by Gaussian membership functions, and the activation of the rules is performed using min–max operators. The inference result consists of a gradually expressed priority level (e.g., no priority, moderate priority, or high priority), obtained by aggregating and normalizing the activation degrees of the rules. The fuzzy module functions as an auxiliary, transparent and independent mechanism that contextually adjusts the preemption decisions for emergency vehicles, without interfering with the learning process or the convergence of the Q-learning algorithm in the adaptive optimization layer.

3.5. Layer 4—Operational Decision and Actuation on Infrastructure

The decisions made by the proposed framework in the analytical stages are being transmitted through the help of the traffic signal infrastructure to the operational level. As the name suggests, its main task is that of maintaining a smooth traffic flow and the prioritization of emergency vehicles when is necessary, by directly managing system operations and applying real-time adjustments. This fourth layer allows gradual and controlled changes of the duration of the traffic light phases, while being aware of the current operational context. It provides calculated decisions without altering the main structure of the traffic light sequence, unlike traditional strategies based on fixed cycles.

At this level, the system uses the outputs of the adaptive control mechanisms, together with those generated by the emergency vehicle management module, to adjust the duration of the green phases at the intersection level. The adjustment takes place through incremental changes to the duration of the phase in progress, changes that can be expressed, at a conceptual level, by the corresponding relationship:

$$G'(t) = G(t) + \Delta u(t) \quad (5)$$

In this formulation, $G(t)$ denotes the duration of the green phase before the intervention, and $G'(t)$ its duration after the update. The prioritization module provides the term $\Delta u(t)$, which expresses the amplitude of the adjustment required either to facilitate the crossing of an emergency vehicle or to alleviate a congestion state identified at the control point level. The relationship has a predominantly conceptual role and highlights how the decision obtained through fuzzy inference is materialized in a physical intervention on the functioning of the traffic lights.

With the help of a well-defined set of operational constraints according to road safety requirements, traffic light adjustments can be performed stably, without affecting the

network. When the request to extend the green period of the traffic light in one direction of the road is detected in the event of an emergency vehicle approaching that point, the overriding of the initial mechanism is limited so as not to disproportionately affect other critical phases, such as those intended for pedestrian crossings or public transport. Thus, the framework has both upper and lower limitations, precisely to avoid the emergence of significant imbalances between competing flows. In this sense, an operational benefit of the intervention is estimated to justify, or not, the penalization of other traffic participants.

Intersection coordination is another essential element of this operational layer. When an emergency vehicle crosses a sequence of intersections, the system does not apply control decisions in isolation at a local level, but synchronizes them spatio-temporally in a coherent manner. To this end, AETM uses estimates of the emergency vehicle's speed, as well as information on the remaining distance to downstream intersections, to determine the optimal moments for initiating interventions. Thus, the vehicle encounters favorable phases upon arrival, delays are reduced, and the risk of secondary congestion decreases, allowing the configuration of a dynamic and adaptive "green lane".

3.6. The Transversal Module of Operational Fairness and User Protection

The AETM framework is built to be expandable, which means that it may be used to add non-technical parts of urban traffic management over time. These parts include fairness, safety for vulnerable users, and openness in the decision-making process. In this context, a transversal module of operational equity is introduced, integrated at both the conceptual and architectural levels, which enhances the adaptive control and emergency vehicle prioritization mechanisms.

The goal of this module is not to directly improve traffic performance or get in the way of the reward-based learning process. Instead, it is to create a logical framework that allows for the monitoring and, when necessary, the limiting of the cumulative impacts of interventions. The module is meant to test how control decisions affect different types of users and traffic flows, with a focus on situations where emergency vehicle preemption systems are often used.

To give a better understanding of this dimension, a purely conceptual fairness indicator is introduced, which aims to highlight how delays are distributed between competing traffic flows. In a simplified formulation, this indicator can be represented by the following relationship:

$$\theta = 1 - \frac{\sigma_{\max}}{\sigma} \quad (6)$$

where σ stands for the range of observed delays and $\sigma_{\max}$ is a reference value set by urban policies or functional guidelines. This indicator is just meant to show how a more fair distribution of intervention effects would work; it is not used in software implementation or to measure outcomes quantitatively.

Functionally, the fairness module is designed to work with emergency vehicle prioritization mechanisms by moderating the duration or frequency of preemption when the cumulative effects on other road users grow considerable. This interaction is intended as a flexible operational constraint, not a hard-and-fast rule, so as to preserve priority in urgent scenarios so as to avoid the accumulation of persistent imbalances in the network.

In this context, one can consider, for example, the discrepancies between the waiting times of the flows competing for the green light, the differences in the accumulated delays of the various modes of transport—such as private vehicles, public transport or pedestrians—as well as how often and for how long preemptive interventions are applied to the same directions. It is also relevant to observe situations in which certain intersections or corridors are repeatedly subject to penalties, indicating an uneven distribution of the effects

of control decisions. Such metrics can highlight situations in which the benefits obtained via prioritizing emergency vehicles are concentrated locally, to the detriment of certain categories of users or areas of the network.

From an operational perspective, the goal of the fairness module is not to achieve complete uniformity of performance among all users, but rather to avert the emergence of structural imbalances in the medium and long run. For example, the system can aim for delays imposed on public transport or pedestrian flows to remain below certain thresholds over longer time periods, or for preemption applied repeatedly in an area to be followed by a gradual return to standard adaptive control regimes. Through such mechanisms, priority is maintained when facing critical events, without operational exceptions becoming permanent network states.

The fairness module is currently only described in terms of its concept and architecture; it has not been fully developed or tested with simulations to see how well it works. This approach highlights how difficult it is to define and measure equity in real urban traffic infrastructures, as well as the need to correlate such mechanisms with municipal policies and regulatory frameworks adapted to each local context. In this regard, the full integration of the module—including the establishment of operational indicators and the analysis of their impact on the overall performance of the system—represents an important direction for future research.

3.7. Transversal Governance, Audit and Transparency Module

Operational transparency in this case refers to the ability to easily clarify the control decisions that were made, as to present, analyze and assess them from an institutional point of view when needed. As such, in this part of the paper it is described the proposed mechanism that has the purpose of supporting operational transparency, managerial responsibility and compatibility with the civil governance regulations. With this purpose, a specialized design that builds up and strengthens public trust, that aims to avoid a blurred functioning of the system, is integrated in the AETM architecture via a particular module for documenting and evaluating decisions.

To begin with, several valuable features such as network status, intervention type and the criteria that substantiated it, will be verified in every interference applied to the traffic infrastructure and recorded by the system for operational context. Retrospective analysis of the system's performance and coherence, as well as justification for the taken maneuvers, are all available through this decision-trail mechanism. Thus, this module's main goal of ensuring the traceability of control decisions can be achieved, aiding in the transition to an auditable and comprehensible algorithmic method.

The system can enhance transparency by aggregating this information over extended time intervals, in the form of operational reports. Such reports can include synthetic indicators on the frequency of interventions, the spatial distribution of preemptions or the cumulative effects of adjustments on traffic. Authorities can use this information as support in the decision-making process, both for the evaluation of mobility policies and for the adaptation of operational parameters according to established urban objectives.

The governance module also supports institutional calibration processes of the architecture. Periodic analysis of aggregated data can highlight situations in which certain intervention patterns generate unwanted effects, such as the repeated penalization of certain areas or categories of users. In such cases, administrative factors can adjust the configuration parameters of the decision-making layers, in order to maintain an appropriate balance between system performance and the social impact of the interventions.

Data protection is another essential dimension integrated into this module. Information from connected vehicles, sensors or monitoring systems may include sensitive data,

which is why the architecture allows the use of aggregation and anonymization mechanisms, in accordance with current data protection regulations. This approach facilitates the exploitation of operational information without compromising user privacy.

All of these functionalities that the AETM architecture embeds are essential parts of a robust system that aspires to facilitate the evaluation, explanation and adjustment of traffic management decisions. This module includes a complementary component for traffic control and optimization mechanisms as an institutional support element.

3.8. Interoperability and Scalability in the AETM Architecture

Interoperability and scalability are essential features of the AETM architecture, which allow its integration into complex and constantly evolving urban networks without requiring major structural changes. In the context of smart cities, the transport infrastructure includes heterogeneous combinations of sensors, existing equipment, local control systems and emerging communication technologies. In this situation, an adaptable traffic control system needs to be able to work in a variety of settings and grow gradually as needed.

The modular architecture of AETM is designed to support this progressive expansion. Each layer of the framework, from data collection to decision and actuation, can be extended or replicated depending on the size and complexity of the network. An important role in this process is played by the data acquisition layer, which ensures the standardization of information coming from various sources and their delivery in a coherent format to the analytical components, regardless of the technology or equipment used. Thanks to this structure, AETM can be implemented incrementally, starting with a small number of intersections or critical nodes and later expanded to the neighborhood or complete urban network level.

Each decision can be coordinated through simple allocation mechanisms between intersections, multiple computing units responsible for different analytical components, as the density of the infrastructure and the volume of data increase. This approach allows for the management of higher levels of complexity without compromising operational stability. Emerging technologies in the area of urban mobility, like connected cars and advanced communications systems, impose a need to allow progressive integration, an aspect that the proposed architecture takes into account. The framework can adapt to such technologies and leverage data through interfacing mechanisms and operational strategies, all without requiring structural changes to the decision logic.

Existing examples of technologies integrated into urban traffic to enhance various performances and facilitate surveillance can be seen in the following. The Sydney Coordinated Adaptive Traffic System (SCATS) [39] involves an intelligent method of monitoring and optimizing traffic by installing sensors under the asphalt at each intersection, and pedestrians are detected by pressing buttons at pedestrian crossings. All measurements are transmitted to regional servers, which in turn are linked to the central control and optimization system. The solution proposed by Hikvision [40] consists of installing surveillance cameras on roads and intersections with the role of detecting road violations, reducing significant traffic load based on regulation algorithms, thus raising the level of safety in the city compared to simply installing video cameras. Other solutions involve technologies such as V2X, a good example being the pilot system implemented in Georgia (USA) by SMARTER Center [41], a project that uses C-V2X (Cellular V2X) for communication between traffic lights and vehicles, including prioritization for emergency vehicles and public transport.

Despite these advances, gaps still exist, as most cities do not yet have large-scale deployments that allow seamless modifications and can connect vehicles, road users and infrastructure. In this context, AETM aims to facilitate the calibration of the system between the regulatory and technical requirements of a city. By enabling informed decisions, perfor-

mance monitoring and incremental deployment, the proposed architecture can contribute to the continuously evolving nature of real-world traffic infrastructures.

4. Results

To validate the operation of the proposed architecture and to robustly evaluate the contribution of each module of the system—reinforcement learning-based control, emergency-vehicle preemption mechanism, and dynamic routing—we conducted an extensive series of experiments in the SUMO (Simulation of Urban Mobility) microscopic environment. This stage aims to answer both global performance questions, regarding the stability of the network under variable traffic conditions, and operational questions, regarding the behavior of ambulances when interacting with the control system.

The evaluation is built around the same traffic-demand file, which automatically generates alternations between low, moderate, heavy, and decreasing flow periods, thus allowing for the direct comparison of four experimental configurations: standard SUMO control, isolated preemption, Q-learning-based control, and the proposed method that combines all components in a unified framework.

The structure of this chapter starts from the description of the configuration and traffic scenarios, continues with the presentation of the evaluation methodology, and ends with the analysis of the results obtained, including both network-level measures and detailed metrics regarding the performance of emergency vehicles.

4.1. Experimental Setup and Traffic Scenarios

The simulations were performed in the microscopic SUMO environment, using an urban network of medium complexity, composed of several directional corridors and diagonal connections that allow the natural formation of congestion zones. The network used in the simulations represents a medium-complexity urban environment, consisting of a combination of signalized cross intersections and Y-intersections, connecting residential, main and arterial roads. This mixed structure was chosen to reproduce realistic interactions between local and transit traffic and to generate varied congestion contexts at the controlled intersections. The signalized intersections are located on main corridors and connecting nodes, allowing the evaluation of both local adaptive control and coordination over intersection sequences.

The traffic demand was generated based on a single route file, which remained unchanged throughout all experiments, to ensure comparability. Within it, flows composed of cars, heavy vehicles, and public transport are defined, with variable temporal distributions that induce, in a controlled manner, alternations between periods of low, medium, and high load. During periods of low traffic, the network operates in fluid conditions, with approximately 120–180 vehicles present simultaneously. As demand increases, a moderate traffic regime is entered, characterized by approximately 300–450 vehicles active on the network and the temporary formation of queues on the main corridors. During high traffic phases, determined by maximum overlap of flows, the network operates in a regime close to saturation, with approximately 650–900 vehicles present simultaneously, allowing the evaluation of the system's behavior under operational stress conditions. All simulations start with a stabilization phase, after which the traffic level gradually increases, reaches a saturation zone, and finally decreases. This natural variation allows the evaluation of the proposed system in very different traffic contexts, without modifying the network or route sets.

Emergency vehicles are introduced into the network at distinct moments of the same time cycle, so that some crossings take place in fluid conditions, and others in intervals dominated by heavy congestion. Origin and destination points are established for each

ambulance, and the movement is initially guided by an automatically generated route. During the simulation, the active dynamic routing mechanism can adapt the route followed, depending on the observed traffic conditions, in order to reduce the total travel time to the destination. Emergency vehicle detection is performed automatically based on the operational class, which allows the prioritization module to be activated at the right time.

To observe the performance of the proposed procedure, three operational regimes were analyzed, evaluated individually, by activating the control mode specific to each one. The first of them, Baseline Control, was used as a reference, using only the static logic, default, from the SUMO simulation environment, without including additional optimization mechanisms. The second scenario, called RL-Based Adaptive Control, additionally introduces a configuration that allows adaptive control of traffic lights, but without making exceptions in traffic depending on the type of vehicle. The last scenario, the one proposed by the paper (AETM—Adaptive and Emergency-Aware Control), also adds the detection of emergency vehicles and their dynamic routing, allowing them to make temporal changes to the traffic light cycle and redirect when necessary. Also, the adaptive learning part for traffic optimization in the second scenario is active in parallel. This approach reproduces realistic situations where regular traffic management must perform well alongside rapid priority for emergency vehicles.

From a computational point of view, the Q-learning-based control mechanism uses a discrete state representation and a reduced action space, with local updates at the intersection level. The learning process is performed online, during the simulation, without separate offline training steps. The computational effort associated with the learning component is negligible compared to the cost of microscopic simulation, which makes the approach suitable for real-time control scenarios or for use in digital twin environments.

All simulations were run for a total duration of 20,000 simulation steps, corresponding to approximately 5 to 6 h of real urban traffic. Although this interval does not cover a full 24-h day, it was chosen as a representative of a normal urban day, able to capture the essential traffic dynamics, including the increase in demand, reaching a peak level and the subsequent recovery phase. Such a concentrated temporal window is frequently used in microscopic traffic studies, as it allows the evaluation of the behavior of control systems under relevant conditions, while maintaining computational feasibility.

The difference between scenarios 2 and 3 does not consist exclusively in granting priority to traffic lights, but also in integrating the spatio-temporal coordination of interventions on multi-intersection routes, while maintaining the overall stability of the network. Using the same flows, the same network and the same insertion times for emergency vehicles ensures that the differences observed in the results are strictly due to the control strategies and not to the particularities of the scenarios. Thus, the proposed architecture can be assessed under steady-state and operational-stress conditions, facilitating a thorough analysis of the trade-offs between overall mobility and emergency response performance.

The setup used in the SUMO simulation environment was meant to show a generic medium-complexity urban situation, which made it possible to test how the system would work under different demand conditions without being affected by the specifics of a real site. Arranged on important corridors and linking streets, the network has both traditional traffic signalized crossings and Y-type crossroads, creating a variety of congestion and interaction scenarios. This structure allows testing both local adaptive control and coordination on intersection sequences, an essential aspect for the analysis of emergency vehicle prioritization.

Although the simulation environment provides a controlled and reproducible framework for validating the proposed mechanisms, it is important to emphasize that the numerical values obtained may vary in real implementations, depending on the geometry

of the intersections, the number of lanes, the existing traffic light phases, the presence of pedestrian or public transport flows, and the location of the detection infrastructure. In this sense, the presented results should be interpreted as comparative trends between control strategies, and the practical implementation of the AETM architecture would require local calibration of parameters and integration with existing urban systems.

4.2. Metrics and Evaluation Methodology

The evaluation of the proposed system is based on a limited but expressive set of indicators, chosen to capture both the overall behavior of the network and the concrete impact on emergency vehicles. For the network level, the analysis compares the three scenarios from the perspective of average waiting time, total queue length and average speed. Average waiting time directly reflects the penalties to which vehicles are subjected at traffic lights or in congestion, being sensitive to the onset and dissipation of congestion. Total queue length provides a quantitative picture of the accumulations at intersections and along the main corridors, allowing the identification of periods in which the network is approaching saturation. The average speed at the network level complements this information, indicating to what extent flows remain fluid or, on the contrary, are slowed down as demand increases or when prioritization mechanisms are applied.

In scenarios that include emergency vehicles, the evaluation is completed by a set of metrics dedicated to them. Each ambulance is tracked throughout its journey, and the total travel time is aggregated as an average travel time, separately for each scenario, to highlight the variations introduced by adaptive and preemptive control. In addition to the duration, the number of stops per route and the average number of stops per trip are also monitored; these values provide additional information on how fragmented the progression of ambulances through the network is and to what extent they are forced to stop or slow down severely near intersections. To assess efficiency from a relative perspective, independent of the length of the route, the travel time of each ambulance is compared to a reference time under free conditions, thus defining an EV (Emergency Vehicle) Travel Time Index, which quantifies how much real traffic degrades the access of the emergency vehicle to its destination.

By combining these network indicators with metrics dedicated to emergency vehicles, the evaluation methodology allows for a coherent analysis of the trade-offs introduced by each scenario: how much the ambulance response time improves when routing and preemption are activated and, in parallel, what effect these mechanisms have on the delays and congestion perceived by the rest of the traffic participants.

Figure 8 illustrates, in a qualitative manner, the differences in the behavior of an intersection under similar traffic conditions, with and without the presence of an emergency vehicle, in the initial phases of increasing demand. In the baseline scenario, without EMS interventions, the intersection operates according to the standard regime, and queue formation occurs gradually in the directions with higher demand. In case (b) however, once the ambulance approached close enough to the intersection, the prioritization mechanism was activated, thus temporarily changing the traffic light lighting cycle. By comparing the two cases, one can observe the application of the preemption module in the context of the simulation, without the appearance of disturbances at the intersection level or blockages within the competing flows.




Figure 8. Visual comparison of intersection behavior during early-stage traffic buildup: (a) baseline flows without emergency vehicles, under moderate demand; (b) EMS approach scenario in the same traffic conditions, with the ambulance initiating adaptive preemption under the AETM controller.

4.3. Comparative Analysis of Control Scenarios

The network performance and the mobility of emergency vehicles were analyzed comparatively in the three scenarios defined previously, using the same traffic flows, the same ambulance insertion times, and the same total simulation duration. The differences that appear in the evolution of the indicators can be attributed exclusively to the control methods introduced progressively, which allows a coherent interpretation of the impact of each component.

The first indicator analyzed is the average vehicle waiting time, shown in Figure 9. In scenario 1 (left), which represents the basic operation of the network with the default traffic light schedules, the waiting time increases sharply with the increase in demand, reaching high values in the peak congestion interval. This steep increase shows the vulnerability of the fixed regime to the overlapping of flows. In scenario 2 (center), the activation of adaptive control leads to a much more stable evolution: although a peak appears in the peak demand area, the amplitudes are significantly lower, and after passing the peak, the network quickly returns to low values. In scenario 3 (right), in which dynamic routing and preemption are also activated, the waiting time profile remains comparable to that of scenario 2, which indicates that the interventions dedicated to ambulances do not introduce additional delays for the rest of the network.




Figure 9. Average network waiting time across simulation scenarios.

Figure 10 shows how the total queue length changes over time. This shows that the three scenarios are very different from each other. In scenario 1, the rise in demand causes vehicles to gradually and steadily pile up at intersections. The curve reaches high values, over 300 vehicles, when the network becomes very congested. After this point, the changes stay big, which means that the network stays in a saturated area for a long

time. In scenario 2, with the activation of adaptive control, the queue level decreases significantly: the accumulations are shorter, and the maximum peak—although still present in the high load interval—is much lower than in the reference scenario. After this peak, the network quickly returns to a stable level, between 20 and 40 vehicles, indicating an increased capacity to dissipate traffic. Scenario 3 maintains the same general trend: although a distinct peak occurs when emergency vehicles cross the network, its amplitude is comparable to that in the adaptive scenario, and the accumulations are promptly reduced after the event. This behavior confirms that the dedicated ambulance mechanisms do not introduce persistent additional congestion and that the network remains stable even under priority intervention conditions.




Figure 10. Average queue length across the evaluated scenarios.

Figure 11 shows the average speed in the network, which backs up what was seen before and gives more information about how each control strategy deals with the slow rise in demand. In scenario 1, the speed drops steadily as the network gets busier, reaching speeds of about 11.5 m/s in the area of highest congestion. This progressive reduction reflects the loss of flow absorption capacity when the traffic lights operate in fixed mode. In scenario 2, the activation of adaptive control considerably stabilizes the speed profile: although a minimum occurs in the same critical interval, the amplitude of the decrease is smaller, and the network recovers quickly, returning to values above 14 m/s at the end of the simulation. Scenario 3 shows the same general dynamics as the adaptive scenario, including a pronounced minimum in the congestion area, but the subsequent recovery is just as fast, and the average speed stabilizes at a level comparable to that in scenario 2. This behavior shows that the introduction of dynamic routing and preemption dedicated to emergency vehicles does not compromise the overall performance of the network and does not prevent its recovery after periods of high demand.




Figure 11. Average network speed under baseline, adaptive, and AETM conditions.

Three main indicators were monitored to assess the performance of emergency vehicles. The first is the average travel time, which represents the total duration of the route traveled by an ambulance between the moment of insertion into the network and arrival at the destination, as reported by SUMO for each trip. The second indicator is the average number of stops along the route, calculated by identifying the sequences in which the vehicle remains at a speed below the movement threshold, each such sequence being considered a stop; the values are aggregated and averaged at the scenario level. The third indicator is the EV Travel Time Index (TTI), used to report the relative delay of the ambulance compared to ideal free-flow conditions. To figure out this index, the scenario first figures out the free-flow time by pretending to be the ambulance on the fastest route possible in the network (without any stops, delays, or conflicts). The TTI is then calculated by dividing the actual travel time by the estimated ideal time. Values close to 1 mean that traffic doesn't have much of an effect on the trip, while high values mean that there are long delays because of traffic jams or having to stop and start again.

The performance of emergency vehicles shows much more pronounced differences between the three scenarios than in the case of network-level indicators. The values summarized in Table 2 show a progressive but substantial improvement of all metrics with the introduction of adaptive control and, subsequently, of mechanisms dedicated to ambulance prioritization. In scenario 1 (Baseline Control), ambulances record an average travel time of 307.3 s, associated with 3.3 stops along the route and an EV Travel Time Index of 7.31, which reflects their direct exposure to the increased congestion of the network.

Table 2. Performance indicators for emergency vehicles.

Scenario	EV Travel Time (Mean Value)	Average EV Stops per Journey	Average EV Travel Time Index
Baseline Control	307.3 s	3.3	7.31
RL-Based Adaptive Control	145.2 s	2.2	3.79
RL + Emergency Prioritization	68.6 s	1.6	1.90

The activation of adaptive control in scenario 2 significantly reduces these values: the average time drops to 145.2 s, the average number of stops decreases to 2.2, and the index reaches 3.79, signaling a much less fragmented progression and a reduction of delays by almost half compared to the baseline scenario. Because of preemption and dynamic routing, the third scenario allows ambulances to travel with minimal disruption from typical traffic. Almost ideal free-flowing conditions are met: an EV Travel Time Index of 1.90, an average number of stops of 1.6, and a decrease in travel time to 68.6 s. The results demonstrate that the network remains robust and that the response times and route flow for emergency vehicles are significantly improved with the addition of priority mechanisms.

5. Conclusions

This paper proposes and analyzes an integrated framework for intelligent urban traffic management, with the main objective of reducing emergency vehicle response times while maintaining the stability and overall performance of the traffic network. The AETM architecture combines an adaptive traffic light control mechanism based on reward learning with dedicated emergency vehicle detection, rerouting and contextual prioritization modules, and is designed to operate efficiently in heterogeneous urban infrastructures.

The results obtained in the simulation environment highlight the efficiency of the adaptive control component. Compared to classical control strategies, the Q-learning-based

approach contributes to traffic stabilization during peak periods, limiting the increase in waiting times and queue lengths. The integration of emergency vehicle management mechanisms does not negatively affect the overall network performance; on the contrary, the complete configuration, which includes rerouting and contextual preemption, maintains values comparable to those obtained by pure adaptive control. At the same time, emergency vehicles benefit from substantial improvements, reflected in significant reductions in travel times, a lower number of stops and values of travel indices close to free-flow conditions. In highly saturated networks or in scenarios with frequent EMS interventions, local delays for competing flows may still occur, representing an inherent trade-off between emergency response efficiency and general traffic performance.

Beyond the quantitative results, the proposed framework integrates an operational equity dimension at the architectural level. The system achieves the prioritization of emergency vehicles through local, temporary and proportional interventions, avoiding the occurrence of cumulative or discriminatory effects on certain flows or categories of users. Prioritization decisions are context-dependent and easily interpretable thanks to the fuzzy and audit components, without affecting the learning process of the adaptive control layer. Thus, equity is supported from the perspective of inclusive mobility, without currently implementing a specific objective for it within the current framework.

Extension of this framework is planned for future research directions. These include calibration of prioritization parameters carried out by testing various optimization methods and integration of advanced congestion prediction models. The exploitation of additional data sources, such as connected vehicles and urban infrastructure, is also envisaged. In addition, evaluation in real implementations or in digital twin environments would allow analyzing the robustness of the proposed solution in diverse urban contexts.

The results presented in this paper are obtained exclusively in a microscopic simulation environment, used to evaluate the behavior of the proposed architecture in a controlled and reproducible setting. Although the observed trends indicate clear benefits in terms of reducing emergency vehicle travel times and maintaining general traffic stability, the numerical values should not be interpreted as guaranteed performances in real implementations, as these are dependent on the local characteristics of the road network and traffic signal infrastructure.

Integrating with current urban traffic management systems, localizing the calibration of prioritizing factors and adaptive control mechanisms, and a basic infrastructure for monitoring traffic and emergency vehicles are all necessary for the proper implementation of the AETM design. A natural intermediate step is the validation in a digital twin environment or in a small-scale pilot project, which would allow for progressive adjustment of the system before a possible expansion at the urban level.

Overall, the proposed AETM architecture offers, following the presented results, a stable basis for further developments in terms of smart and inclusive urban mobility solutions. From the current stage, where its adaptive and transparent nature for emergency vehicle management, integrated without compromising network optimization performance, has been presented, the components currently remaining at the conceptual level can be easily implemented in future stages of research.

Author Contributions: Conceptualization, J.S. and I.-M.V.; methodology, J.S.; software, I.-M.V. and I.S.; validation, I.-M.V., I.S. and J.C.; formal analysis J.C.; investigation, I.-M.V.; resources, M.S.; data curation, I.-M.V.; writing—original draft preparation, I.-M.V. and M.S.; writing—review and editing, I.-M.V. and M.S.; visualization, J.C.; supervision, I.S.; project administration, J.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: Dataset available on request from the authors.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AETM	Adaptive and Equitable Traffic Management system
ITS	Intelligent Transport Systems
VANET	Vehicular Ad-hoc Network
V2X	Vehicle-to-Everything
IoT	Internet of Things
RL	Reinforcement Learning
PIC	Programmable Interface Controller
ML	Machine Learning
TC	Traffic Control
RG	Route Guidance
DTA	Dynamic Traffic-Assignment
DL	Deep Learning
CNN	Convolutional Neural Network
LSTM	Long Short-Term Memory
KNN	K-Nearest Neighbors
EMS	Emergency Medical Services
V2I	Vehicle to Infrastructure
VSL	Variable Speed Limits
RM	Ramp Metering
ACC	Adaptive Cruise Control
KPI	Key Performance Indicators
GAN	Generative Adversarial Network
WSN	Wireless Sensor Network
CVaR	Conditional Value at Risk
DRL	Deep Reinforcement Learning
PER	Prioritized Experience Replay
TLS	Traffic Light Signal
SUMO	Simulation of Urban Mobility
EV	Emergency Vehicle
TTI	Travel Time Index

References

1. Saïda, A.; Rahul, C.; Meghana, C.; Akanksha, G.; Sathvika, A. IoT Based Traffic Light Control System with Ambulance Designation. In Proceedings of the 3rd International Conference Advancement in Electronics & Communication Engineering (AECE), Ghaziabad, India, 23–24 November 2023; pp. 98–101.
2. Rajesh, R.; Kathir, M.N.; Senthilkumar, H.; Parthiban, S.K. Traffic Optimization and Automated Emergency Vehicle Response System. In Proceedings of the 2025 6th International Conference on Recent Advances in Information Technology (RAIT), Dhanbad, India, 6–8 March 2025; pp. 1–6.
3. Munasinghe, K.D.S.A.; Waththegedara, T.D.; Wickramasinghe, I.R.; Herath, H.M.O.K.; Logeeshan, V. Smart Traffic Light Control System Based on Traffic Density and Emergency Vehicle Detection. In Proceedings of the 2022 Moratuwa Engineering Research Conference (MERCon), Moratuwa, Sri Lanka, 27–29 July 2022; pp. 1–6.
4. Jamal, H. Physical Components and Devices in ITS. Available online: <https://www.aboutcivil.org/ITS-physical-components-devices.html> (accessed on 23 January 2017).
5. Ashokkumar, C.; Kumari, D.A.; Gopikumar, S.; Anuradha, N.; Krishnan, R.S.; Sakthidevi, I. Urban Traffic Management for Reduced Emissions: AI-based Adaptive Traffic Signal Control. In Proceedings of the 2024 2nd International Conference on Sustainable Computing and Smart Systems, Coimbatore, India, 10–12 July 2024; pp. 1609–1615.

6. Agarwal, I.; Singh, A.; Agarwal, A.; Mishra, S.; Satapathy, S.K.; Cho, S.-B.; Prusty, M.R.; Mohanty, S.N. Enhancing Road Safety and Cybersecurity in Traffic Management Systems: Leveraging the Potential of Reinforcement Learning. *IEEE Access* **2024**, *12*, 9963–9975. [\[CrossRef\]](#)
7. Bairi, P.; Swain, S.; Bandyopadhyay, A.; Aurangzeb, K.; Alhussein, M.; Mallik, S. Intelligent VANET-based Traffic Signal Control System for Emergency Vehicle Prioritization and Improved Traffic Management. *Egypt. Inform. J.* **2025**, *30*, 100700. [\[CrossRef\]](#)
8. Cui, Z.; Henrickson, K.; Ke, R.; Wang, Y. Traffic Graph Convolutional Recurrent Neural Network: A Deep Learning Framework for Network-Scale Traffic Learning and Forecasting. *IEEE Trans. Intell. Transp. Syst.* **2019**, *21*, 4883–4894. [\[CrossRef\]](#)
9. Lilhore, U.K.; Imoize, A.L.; Li, C.-T.; Simaiya, S.; Pani, S.K.; Goyal, N.; Kumar, A.; Lee, C.-C. Design and Implementation of an ML and IoT Based Adaptive Traffic-Management System for Smart Cities. *Sensors* **2022**, *22*, 2908. [\[CrossRef\]](#)
10. Sahil, N.; Sood, S.K. Smart Vehicular Traffic Management: An Edge Cloud Centric IoT Based Framework. *Internet Things* **2021**, *14*, 100140. [\[CrossRef\]](#)
11. de Luca, S.; Di Pace, R.; Memoli, S.; Pariota, L. Sustainable Traffic Management in an Urban Area: An Integrated Framework for Real-Time Traffic Control and Route Guidance Design. *Sustainability* **2020**, *12*, 726. [\[CrossRef\]](#)
12. Han, X. Traffic Incident Detection: A Deep Learning Framework. In Proceedings of the 2019 IEEE International Conference Mobile Data Management (MDM), Hong Kong, China, 10–13 June 2019; pp. 379–380.
13. Muntean, M.V. Multi-Agent System for Intelligent Urban Traffic Management Using Wireless Sensor Networks Data. *Sensors* **2022**, *22*, 208. [\[CrossRef\]](#)
14. El-Tantawy, S.; Abdulhai, B. Multi-Agent Reinforcement Learning for Integrated Network of Adaptive Traffic Signal Controllers (MARLIN-ATSC). In Proceedings of the IEEE Conference on Intelligent Transportation Systems (ITSC), Anchorage, AK, USA, 16–19 September 2012; pp. 319–326.
15. Wei, H.; Xu, N.; Zhang, H.; Zheng, G.; Zang, X.; Chen, C.; Zhang, W.; Zhu, Y.; Xu, K.; Li, Z. CoLight: Learning Network-Level Cooperation for Traffic Signal Control. *arXiv* **2019**, arXiv:1905.05717.
16. Zheng, G.; Xiong, Y.; Zang, X.; Feng, J.; Wei, H.; Zhang, H.; Li, Y.; Xu, K.; Li, Z. Learning Phase Competition for Traffic Signal Control. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD), Anchorage, AK, USA, 4–8 August 2019; pp. 1963–1972.
17. Farrag, S.G.; Sahli, N.; El-Hansali, Y.; Shakshuki, E.M.; Yasar, A.; Malik, H. STIMF: A Smart Traffic Incident Management Framework. *J. Ambient. Intell. Humaniz. Comput.* **2021**, *12*, 85–101. [\[CrossRef\]](#)
18. Atassi, R.; Sharma, A. Intelligent Traffic Management Using IoT and Machine Learning. *J. Intell. Syst. Internet Things* **2023**, *8*, 8–20. [\[CrossRef\]](#)
19. Hu, C.; Tang, J.; Wang, Y.; Li, Z.; Dai, G. A Flexible Road Network Partitioning Framework for Traffic Management via Graph Contrastive Learning and Multi-Objective Optimization. *Comput.-Aided Civ. Infrastruct. Eng.* **2025**, *40*, 2671–2694. [\[CrossRef\]](#)
20. Kuang, Z.; Zhao, X.; Feng, L. Priority of Emergency Vehicle Dynamic Right-Of-Way Control Method in Networked Environment. *Appl. Sci.* **2023**, *13*, 5883. [\[CrossRef\]](#)
21. Goodall, N.; Smith, B.; Park, B. Traffic Signal Control with Connected Vehicles. *Transp. Res. Rec. J. Transp. Res. Board* **2013**, *2381*, 65–72. [\[CrossRef\]](#)
22. Pai, B.; Shaikh, M.; Vartak, M.; Shaikh, M.; Lade, S. Adaptive Traffic Signal System for Emergency Vehicle. *Int. J. Creat. Res. Thoughts* **2023**, *11*, 306.
23. Xiong, C.; Yang, X.T.; Zhang, L.; Lee, M.; Zhou, W.; Raqib, M. An Integrated Modeling Framework for Active Traffic Management and Its Applications in the Washington, DC Area. *J. Intell. Transp. Syst.* **2021**, *25*, 609–625. [\[CrossRef\]](#)
24. Kerner, B. Autonomous Driving in Framework of Three-Phase Traffic Theory. *Procedia Comput. Sci.* **2018**, *130*, 785–790. [\[CrossRef\]](#)
25. Kaparias, I.; Liu, P.; Tsakarestos, A.; Eden, N.; Schmitz, P.; Hoadley, S.; Hauptmann, S. Predictive Road Safety Impact Assessment of Traffic Management Policies and Measures. *Case Stud. Transp. Policy* **2020**, *8*, 508–516. [\[CrossRef\]](#)
26. Alsaawy, Y.; Alkhodre, A.; Abi Sen, A.; Alshanqiti, A.; Bhat, W.A.; Bahboub, N.M. A Comprehensive and Effective Framework for Traffic Congestion Problem Based on the Integration of IoT and Data Analytics. *Appl. Sci.* **2022**, *12*, 2043. [\[CrossRef\]](#)
27. Musa, A.A.; Malami, S.I.; Alanazi, F.; Ounaies, W.; Alshammari, M.; Haruna, S.I. Sustainable Traffic Management for Smart Cities Using IoT-Oriented Intelligent Transportation Systems (ITS): Challenges and Recommendations. *Sustainability* **2023**, *15*, 9859. [\[CrossRef\]](#)
28. Antipova, A.; Sultana, S.; Hu, Y.; Rhudy, J.P., Jr. Accessibility and Transportation Equity. *Sustainability* **2020**, *12*, 3611. [\[CrossRef\]](#)
29. Pereira, R. Future Accessibility Impacts of Transport Policy Scenarios: Equity and Sensitivity to Travel Time Thresholds for Bus Rapid Transit Expansion in Rio de Janeiro. *J. Transp. Geogr.* **2018**, *74*, 321–332. [\[CrossRef\]](#)
30. Sanchez, T.W.; Qian, Y.; Yan, X. AI Applications in Transportation and Equity: A Survey of U.S. Transportation Professionals. *Future Transp.* **2024**, *4*, 1161–1176. [\[CrossRef\]](#)
31. Nirmala, S. Blockchain Based Secured Framework for Road Traffic Management Using Fog Computing. *Int. J. Comput. Intell. Control* **2021**, *13*, 331–337.

32. Sefati, S.S.; Arasteh, B.; Craciunescu, R.; Comsa, C.-R. Intelligent Congestion Control in Wireless Sensor Networks (WSN) Based on Generative Adversarial Networks (GANs) and Optimization Algorithms. *Mathematics* **2025**, *13*, 597. [[CrossRef](#)]
33. Paricio, A.; López-Carmona, M. Multimap Routing for Road Traffic Management. In *Lecture Notes in Computer Science*; Springer: Berlin/Heidelberg, Germany, 2019; pp. 1–12.
34. Almukhalifi, H.; Noor, A.; Noor, T.H. Traffic Management Approaches Using Machine Learning and Deep Learning Techniques: A Survey. *Eng. Appl. Artif. Intell.* **2024**, *133*, 108147. [[CrossRef](#)]
35. Louati, A. Machine Learning Framework for Sustainable Traffic Management and Safety in AlKharij City. *Sustain. Futures* **2025**, *133*, 100407. [[CrossRef](#)]
36. Michailidis, P.; Michailidis, I.; Lazaridis, C.R.; Kosmatopoulos, E. Traffic Signal Control via Reinforcement Learning: A Review on Applications and Innovations. *Infrastructures* **2025**, *10*, 114. [[CrossRef](#)]
37. Tessler, C.; Mankowitz, D.J.; Mannor, S. Reward Constrained Policy Optimization. In Proceedings of the International Conference on Learning Representations (ICLR), New Orleans, LA, USA, 6–9 May 2019. Available online: <https://openreview.net/forum?id=SkfrvsA9FX> (accessed on 5 September 2025).
38. Li, M.; Pan, X.; Liu, C.; Li, Z. Federated Deep Reinforcement Learning-Based Urban Traffic Signal Optimal Control. *Sci. Rep.* **2025**, *15*, 11724. [[CrossRef](#)]
39. Sydney Coordinated Adaptive Traffic System (SCATS). Available online: <https://www.scats.nsw.gov.au/home> (accessed on 5 September 2025).
40. Hikvision. Available online: <https://www.hikvision.com/en/> (accessed on 5 September 2025).
41. SMARTER Center. Connected Vehicle Deployments in the United States. Available online: <https://smartercenter.org/media/cv-deployments-in-the-us/> (accessed on 5 September 2025).

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.