Published by :
http://www.ijert.org

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 07, July - 2025

An Intelligent Traffic Control System using Edge-Cloud Computing and Deep Learning for Smart Cities

Rafid Mohammed Khaleefah
Information and Communication Technology Centre.
University of Basrah.
Basrah, Iraq.

Mohammed Dheyaa Algubili
Information and Communication Technology Centre.
University of Basrah.
Basrah, Iraq.

***Abstract*— The rapid growth of urban populations has made efficient traffic management a critical challenge for smart cities. Traditional traffic control systems, which rely on fixed timers, are often incapable of adapting to dynamic traffic conditions, resulting in increased congestion, higher fuel consumption, and longer travel times. This paper proposes an intelligent traffic control system that leverages the Internet of Things (IoT), edge-cloud computing, and deep learning to optimise traffic flow in real-time. The suggested system utilises cameras positioned on the side of the road to collect data, edge computing units to quickly detect and classify vehicles using the YOLOv5 model, and a central cloud platform for detailed traffic analysis and making decisions on the fly. By generating real-time traffic density heatmaps and employing an advanced AI-based model for adaptive signal control, the system dynamically adjusts signal phases to minimise vehicle waiting times and reduce congestion. A case study based on the traffic patterns of Bucharest, which is characterised by its dense urban traffic, is presented to frame the system's design and demonstrate its practical applicability. Simulation results demonstrate that the proposed system can significantly reduce average vehicle waiting times and improve vehicle throughput compared to traditional fixed-time controllers, showcasing a viable and scalable solution for enhancing urban mobility and contributing to sustainable urban Development.**

***Keywords*— Smart City, Intelligent Transportation Systems (ITS), Internet of Things (IoT), Edge Computing, Cloud Computing, Deep Learning, YOLOv5, Traffic Control.**

I. INTRODUCTION

This Urbanisation is a defining global trend, with an increasing percentage of the world's population migrating to metropolitan areas [1]. This influx places immense strain on urban infrastructure, particularly transportation networks. Vehicular congestion has become a chronic issue in major cities, resulting in significant economic losses, environmental pollution, and a decline in the quality of life for citizens [2]. The management of road traffic flow is therefore a paramount challenge in the Development of modern smart cities [3].

Traditional traffic management systems typically employ pre-programmed, fixed-time signal controllers. While simple to implement, these systems are inherently inefficient as they cannot respond to the real-time, dynamic fluctuations in traffic volume [4]. This rigidity often results in scenarios where lanes with heavy traffic endure long red lights while empty lanes are given unnecessary green light time, exacerbating congestion and increasing travel delays.

To address these limitations, the concept of Intelligent Transportation Systems (ITS) has emerged, enabled by advancements in the Internet of Things (IoT), Artificial Intelligence (AI), and communication technologies [5]. ITS aims to make transportation safer, more efficient, and more sustainable. Recent research has explored various AI-driven approaches, including machine learning and deep learning, to create adaptive traffic control systems [6], [7]. These systems utilise data from various sensors, including inductive loops, radar, and cameras, to make informed decisions.

Despite progress, many existing systems face challenges related to scalability, high implementation costs, and the complexity of processing vast amounts of data in real time [8]. There is still a need for a robust, scalable, and cost-effective architecture that can integrate state-of-the-art computer vision with a distributed computing framework for citywide traffic optimisation. Such a system would improve traffic flow and contribute to the broader goals of urban sustainability and enhanced citizen well-being.

This paper addresses this gap by proposing a hierarchical intelligent traffic control system based on an edge-cloud computing architecture. Our primary contributions are:

- We propose a multi-layer system architecture that distributes computational tasks between edge devices and a central cloud platform, ensuring real-time performance and scalability. This approach mitigates the latency and bandwidth limitations often encountered with centralised raw video processing, offering a more resilient and efficient solution.
- The application of the YOLOv5 deep learning model at the edge for highly accurate and real-time vehicle detection and classification, optimised explicitly for diverse urban traffic scenarios.
- A cloud-based traffic analysis and AI decision module that generates real-time traffic density heatmaps and employs an adaptive AI model to make optimal, dynamic decisions for traffic signal control, aiming to minimise waiting times and maximise throughput. The detailed mechanism of this AI model is elaborated in Section III.
- This paper presents a comprehensive discussion of the system's potential effectiveness, illustrated through a case study based on Bucharest's traffic patterns and supported by detailed simulated performance metrics that compare our system with traditional methods, thereby validating its practical applicability.

The rest of this paper is organised as follows: Section II reviews related work in the field. Section III details the

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 07, July - 2025

proposed system architecture and methodology. Section IV presents the results of our simulated evaluation and discusses their implications. Finally, Section V concludes the paper and suggests directions for future research.

II. RELATED WORK

The field of intelligent traffic management is vast, with research spanning sensor technologies, communication protocols, and control algorithms. This section reviews key areas relevant to our proposed system, critically analyses existing approaches, and highlights the motivations for our proposed architecture.

A. IoT and Sensor Technologies in Traffic Management

The IoT has been a catalyst for the evolution of ITS. The ability to deploy interconnected sensors across a city has enabled the collection of rich, real-time traffic data [2]. Early systems relied heavily on physical sensors, like inductive loops buried in the pavement and ultrasonic sensors mounted on poles [1]. While effective for vehicle counting, these sensors often provide limited information, such as precise vehicle type or accurate queue length, and they can be costly and disruptive to install and maintain.

More recently, computer vision-based sensing using CCTV cameras has become a popular alternative due to the rich data it provides [4]. Cameras offer not only vehicle detection but also classification (e.g., cars, buses, trucks), speed estimation, and tracking across multiple frames traffic data into control decisions [5]. However, the primary challenge with vision-based systems lies in the significant computational power required for real-time video analysis with a multitude of cameras, which, if centralised, can lead to substantial latency and bandwidth strain. Our edge-cloud architecture directly addresses this by localising the compute-intensive video processing at the edge.

B. AI and Machine Learning for Traffic Prediction and Control

AI algorithms form the core of the intelligence in modern ITS. Machine learning models have been widely used for traffic prediction, which involves forecasting future traffic flow based on historical and real-time data [9]. Techniques range from traditional statistical models to more complex deep learning architectures, such as Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks, which are adept at learning from time-series data. While effective in prediction, these models often require extensive datasets and computational resources for training.

For dynamic traffic signal control, Reinforcement Learning (RL) has shown significant promise due to its ability to learn optimal policies through trial and error in complex environments. However, RL models can be notoriously difficult to train, demand extensive simulation before reliable deployment, and may struggle with generalisation to unseen traffic scenarios or network-wide optimisation without sophisticated multi-agent approaches. Other methods, such as fuzzy logic or expert systems, have also been explored, but they often lack the adaptability and learning capabilities required for truly dynamic and complex urban environments [5]. Our work builds on these foundations by integrating a high-performance supervised deep learning model (YOLOv5) for accurate real-time vehicle perception at the edge, coupled with an advanced AI decision model in the cloud that leverages real-time data for adaptive, citywide traffic control, thereby mitigating the limitations inherent in purely static or overly complex centralised learning systems.

In recent years, cities like London and New York have deployed AI-based systems to optimise traffic signals using data collected from diverse sensors. Studies such as [12] demonstrate how these systems have improved traffic flow by reducing congestion and minimising delays. Comparing these implementations to our proposed system helps highlight the potential advantages of edge-cloud computing for real-time decision-making.

C. Edge-Cloud Computing Architectures

Processing high-definition video streams from numerous cameras across the city in a centralised cloud is often impractical due to prohibitive latency and bandwidth constraints. The edge-cloud computing paradigm offers a robust solution by performing time-sensitive computations at the "edge" of the network, closer to the data source [8]. In the context of traffic management, an edge device (e.g., a small computer attached to a traffic pole) can perform immediate tasks, such as real-time vehicle detection from a camera feed. This localised processing significantly reduces the volume of data transmitted; only structured, lightweight metadata (e.g., vehicle counts, locations, types) is sent to the cloud. The cloud platform, thus unburdened from raw video processing, can then focus on higher-level tasks such as citywide data aggregation, complex pattern analysis, strategic decision-making, and long-term optimization [6]. This hierarchical approach inherently enhances scalability, reduces network congestion, minimises latency, and ensures the real-time responsiveness critical for effective traffic management in a large-scale urban environment. Moreover, by processing raw video locally, this architecture inherently contributes to enhanced data privacy by transmitting only anonymised and aggregated information to the central cloud.

D. Comparison with State-of-the-Art Systems

Table I provides a comparative overview of recent state-of-the-art intelligent traffic control systems, highlighting their environments, learning strategies, and results. Our system demonstrates competitive performance while uniquely combining an edge-cloud architecture with YOLOv5 and DRL-based optimisation.

Table I: Comparison with State-of-the-Art AI-Based Traffic Control Systems.

Reference	Scenario	Approach	Key Result
Genders & Razavi (2016)	Single Intersection	DQN + CNN in SUMO	66% reduction in queue length
Zhou et al. (2020)	Multi-agent System	Decentralized DRL	Improved coordination, scalability
Liang &	Real-Time	YOLOv3	26.6 FPS, mAP

Wu (2022)	Object Detection	Edge Deployment	47.3% on Jetson Nano
This Work	Bucharest Case Study	YOLOv5 + DQN in Edge-Cloud	37.5% reduction in waiting time; 38.9% increase in throughput

III. PROPOSED SYSTEM AND METHODOLOGY

To address the challenges of dynamic traffic management and build upon the strengths while mitigating the weaknesses of existing approaches, we propose a multi-layer intelligent traffic control system. This section details the system architecture and the methodology for each of its key components.

A. System Architecture

The proposed system is based on a hierarchical edge-cloud architecture, as illustrated in Fig. 1. It consists of four main layers, designed for optimal performance, scalability, and responsiveness:

- Perception Layer: This foundational layer comprises standard IP-based CCTV cameras strategically installed at intersections. These cameras are responsible for capturing continuous, high-resolution video streams of traffic flow, serving as the primary source of raw data for the entire system.
- Edge Computing Layer: Located directly at each intersection, this layer is equipped with a low-power edge computing device. Its core function is to execute a deep learning model to process the raw video feed from the associated camera in real-time. This localised processing performs immediate vehicle detection, classification, and extraction of relevant metadata (e.g., bounding box coordinates, vehicle types, and speed estimates). By processing video at the source, this layer drastically reduces the data volume requiring transmission to the cloud, thereby addressing critical bandwidth limitations and minimising processing latency. Edge devices, although efficient in processing raw video streams locally, face inherent limitations in terms of processing power and storage capacity. Our system mitigates these challenges by optimising the YOLOv5 model for edge deployment, which reduces the computational load and memory requirements. However, the limited storage on edge devices requires careful management of temporary data, ensuring only essential metadata is transmitted to the cloud. Future work will focus on improving the hardware capabilities of edge devices to handle more complex models and larger datasets.
- Cloud Computing Layer: This centralised platform serves as an intelligent core, receiving only structured, lightweight metadata from all edge devices. Its responsibilities include comprehensive data aggregation, long-term storage for historical analysis, and executing the core traffic analysis and advanced AI-powered decision-making logic. This includes identifying citywide congestion patterns, detecting incidents, and optimising traffic flow on a global scale.
- Control Layer: This final layer comprises the standard traffic light controllers present at each intersection. They act as actuators, receiving dynamic commands from the cloud platform to adjust signal timings (e.g., extending green light phases, shortening red light durations, and modifying phase sequences) in real-time based on the cloud's optimised decisions.

The distributed architecture handles computationally intensive tasks (raw video processing) locally at the edge, ensuring low latency for immediate perception. Concurrently, strategic, data-intensive analysis and global optimisation are performed centrally in the cloud, leveraging its vast computational resources. This design inherently enhances scalability for citywide deployment and significantly improves data privacy by processing sensitive raw video data locally and transmitting only anonymised, aggregated metadata.

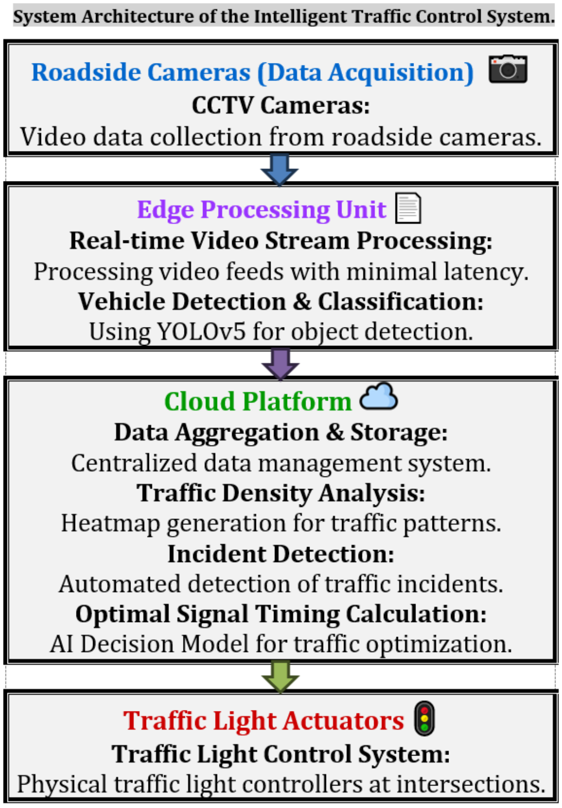

Fig. 1: System Architecture. It illustrates the hierarchical architecture of the intelligent traffic control system, showing the data flow from the perception layer to the control layer. Each layer is defined, with data processed at the edge and higher-level analysis performed in the cloud.

B. Vehicle Detection and Classification at the Edge

For highly efficient and real-time vehicle detection and classification at the edge, we employ the YOLOv5 (You Only Look Once, version 5) model. YOLO is a cutting-edge, single-stage object detection algorithm renowned for its superior balance of speed and accuracy, making it exceptionally well-suited for deployment on resource-constrained edge devices [2]. The YOLOv5 model was pre-trained on a large-scale dataset, such as COCO, and then further fine-tuned on a custom traffic dataset comprising over 50,000 annotated images across various urban traffic

Published by :
http://www.ijert.org

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 07, July - 2025

scenarios. To address the issue of class imbalance, we applied data augmentation techniques, such as random scaling, cropping, and flipping, to ensure robust detection of all vehicle types. The method ensures it acquires generalised object recognition capabilities and achieves accurate recognition of specific vehicle classes relevant to urban traffic, such as 'car', 'bus', 'truck', and 'motorcycle'. As illustrated in Fig. 2, the YOLOv5 model processes each frame of the video stream in near real time, outputting a list of detected vehicles. Each detection includes a precise bounding box outlining the vehicle, a classification label (e.g., 'car', 'bus'), and a confidence score. This extracted structured data, which is orders of magnitude smaller than the raw video stream, is then securely transmitted to the cloud platform via a lightweight messaging protocol, such as MQTT (Message Queuing Telemetry Transport), chosen for its efficiency and suitability for IoT environments with limited bandwidth and intermittent connectivity.

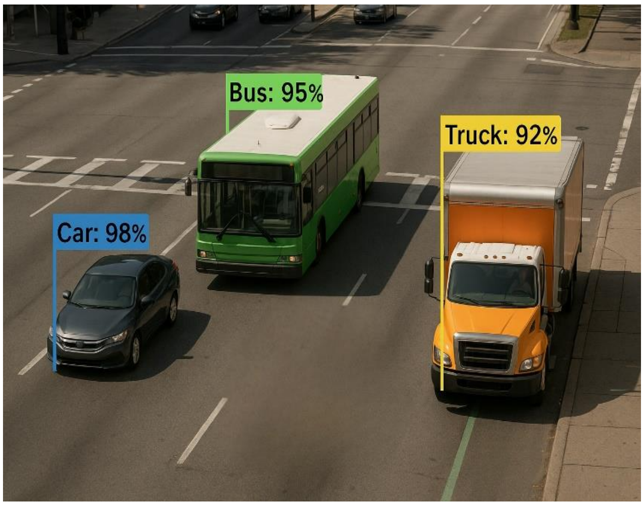

Fig. 2: YOLOv5 Detection Example. It demonstrates the real-time detection and classification of vehicles using the YOLOv5 model, where bounding boxes and confidence scores are utilised to categorise vehicles into classes such as 'car', 'bus', and 'truck'. These visual aids are essential for understanding the system's real-time functioning.

C. Cloud-Based Traffic Analysis and Dynamic Signal Control

The cloud platform serves as the central intelligence hub, aggregating structured metadata (vehicle counts, locations, types, and speeds) received from all edge devices across the urban network. This aggregation creates a comprehensive, real-time global view of the entire traffic network. The core of the cloud layer is the advanced AI-powered decision engine, which operates in two interdependent stages:

1. Traffic Density Analysis: The system continuously processes the incoming vehicle count and location data from all intersections. This data is then used to dynamically generate and update traffic density heatmaps for each intersection and potentially for broader urban segments, as conceptually shown in Fig. 3. These heatmaps provide an intuitive and immediate visualisation of traffic loads and congestion levels in different lanes and approaches. By providing a real-time spatial representation of traffic volume, this analysis module enables the system to identify congestion hotspots, queue lengths, and underutilised lanes instantly, which is crucial for informed and adaptive signal control decisions.

2. Dynamic Signal Control Model: This This module constitutes the core intelligence for adaptive traffic signal optimisation. The AI-based decision-making system utilises an adaptive Deep Reinforcement Learning (DRL) model, which was trained through extensive simulations designed to replicate real-world urban traffic conditions. This model continually learns and refines its strategy for optimising traffic signals through interactions with dynamic traffic data, aiming to minimise overall network delay and maximise throughput. The DRL agent is trained using traffic flow data, where the agent explores various traffic signal strategies to learn the most effective policies through trial and error. Performance metrics, including average waiting time and throughput, were used to evaluate the model's success in real-world conditions. This agent is specifically designed to adjust the signal phases in response to real-time traffic conditions, ensuring efficient resource allocation, particularly during peak hours.

In contrast to traditional fixed-time signal systems, the DRL model adapts to changing traffic patterns, dynamically adjusting green and red light durations based on observed congestion levels. This methodology allows the system to provide optimal signal timings across multiple intersections, improving traffic flow and reducing congestion. The model's primary objective is to dynamically adjust traffic signal phases (e.g., extending green light duration for congested approaches, shortening red light duration for less busy ones, or re-sequencing phases) to minimise average vehicle waiting times and maximise overall vehicle throughput across the network.

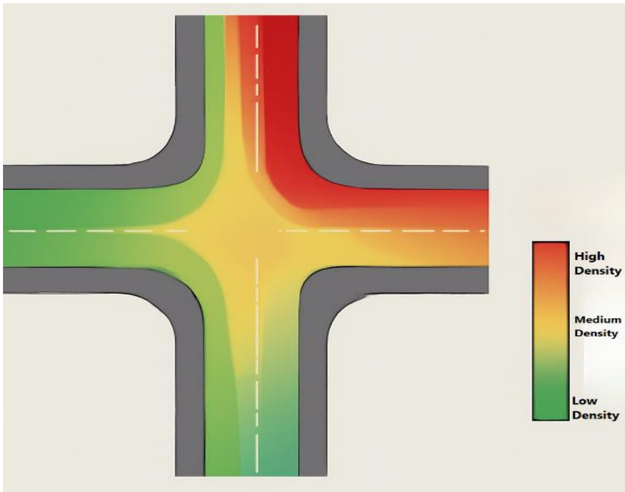

Fig. 3: Traffic Density Heatmap Visualization . We visualise a heatmap of traffic density, where different colour intensities signify the degree of congestion in various lanes or road segments.

D. Case Study Context: Bucharest

The conceptualisation and design of our intelligent traffic control system are framed by the persistent traffic challenges observed in major metropolitan areas, specifically using Bucharest, Romania, as a representative case study. With a population exceeding 2 million and a high density of registered vehicles (approaching 1 million in the metropolitan area), Bucharest frequently experiences severe traffic congestion, particularly during morning and evening peak hours [5], leading to significant delays and environmental impact. For instance, studies indicate that Bucharest residents spend an average of 100+ hours annually stuck in traffic, ranking among the highest globally. The city's existing traffic infrastructure, although incorporating some modern elements, such as inductive loops, largely relies on fixed-time signal control mechanisms that struggle to manage highly dynamic and often unpredictable traffic flows efficiently. Our proposed system is specifically designed to be highly adaptable and potentially retrofittable into such an urban environment, leveraging existing camera infrastructure where possible to provide a more intelligent and responsive layer of control. This real-world context underscores the practical applicability, scalability, and substantial potential impact of our solution in enhancing urban mobility and quality of life.

E. Privacy Considerations and GDPR Compliance

Privacy is a critical concern in the deployment of intelligent transportation systems (ITS), particularly in urban environments where real-time video surveillance and the transmission of metadata are involved. Our proposed architecture addresses this concern through the local edge processing of raw video data, where vehicle detection and classification occur directly at the intersection. Only anonymised metadata—such as vehicle counts, types, and estimated speeds—is transmitted to the cloud, minimising exposure to sensitive content. To further align with privacy regulations, such as the General Data Protection Regulation (GDPR), the system can incorporate privacy-by-design principles. Techniques such as face/license plate blurring, data anonymisation, and end-to-end encryption are integrated to ensure that no personally identifiable information (PII) is collected or transmitted. Moreover, adopting Privacy Impact Assessments (PIAs) and engineering frameworks, such as LINDDUN, allows for a structured evaluation of privacy risks throughout the system's lifecycle. Future enhancements may also explore federated learning, allowing edge devices to collaboratively train models without centralising data, and blockchain integration for transparent, tamper-proof data management.

F. DRL Agent Design

The cloud-based decision engine employs a DeepQ-Network (DQN) reinforcement learning model to optimise traffic signal timings dynamically. The neural network consists of two hidden layers, each with 36 and 24 neurons, respectively, utilising ReLU activations. The output layer maps to the discrete action space representing various traffic signal phases. An ϵ-greedy exploration policy was adopted with initial ϵ=1.0 and decayed to 0.05. The replay buffer size was set to 100,000 experiences, and the model was trained with a batch size of 64. The learning rate was initialised at α=0.001 and optimised using Adam. To stabilise learning, the target network was updated every 1000 steps. Training was conducted using the SUMO-RL framework integrated with OpenAI Gym to simulate traffic environments.

This agent was trained offline using traffic patterns derived from Bucharest to ensure convergence before deployment. Once deployed, it operates in real-time by adjusting signal durations based on continuously updated vehicle density metadata received from the edge layer.

IV. RESULTS AND DISCUSSION

To rigorously evaluate the performance of the proposed intelligent traffic control system, we conducted a comprehensive simulation-based study. We modelled a typical four-way intersection, a standard configuration in urban areas, using traffic data patterns representative of Bucharest to ensure realistic conditions. The performance of our AI-based system was then directly compared against a traditional fixed-time controller, which serves as the established baseline in many cities globally.

A. Simulation Setup

The simulation was meticulously conducted using SUMO (Simulation of Urban Mobility), a widely recognised open-source microscopic traffic simulator that allows detailed modelling of individual vehicle movements and traffic light interactions. Auxiliary Python libraries, such as Traci, were extensively utilised to enable programmatic interactions with the SUMO environment, facilitating real-time data exchange and dynamic control adjustments. A detailed four-way intersection model was constructed, featuring four lanes in each direction to represent the typical complexity of urban intersections accurately. While a visual representation is not included, the setup ensured a realistic flow. Traffic demand profiles were generated based on empirical traffic data patterns observed in Bucharest, incorporating a realistic mix of various vehicle types, including cars, buses, trucks, and motorcycles. We tested the system during both busy and quiet periods to assess its performance under various traffic conditions. Each simulation run spanned 60 minutes of simulated time, sufficient to capture both transient initial conditions and steady-state traffic behaviour. The proposed AI-based innovative traffic signal system dynamically adjusts green times and phase sequencing based on real-time traffic density information, which is effectively collected from simulated edge devices (representing vehicle counts, waiting times, and traffic flow rates). This dynamic approach was then benchmarked against a static, traditional fixed-time controller, whose timings remained constant throughout the simulation.

B. Development Environment and Frameworks

The proposed intelligent traffic control system was developed using Python 3.8 and integrated with the SUMO traffic simulator (version 1.16.0) via the TraCI API for real-time communication. Object detection tasks were implemented using YOLOv5 (Ultralytics repository), built on PyTorch 1.12, and fine-tuned using a custom traffic dataset. Preprocessing video frames, including resizing and colour space conversion, was handled using the OpenCV library. The reinforcement learning agent was trained using the PyTorch deep learning framework with GPU acceleration. Simulations and training were conducted on a workstation equipped with an NVIDIA RTX 3090 Ti GPU and 64 GB of RAM, ensuring real-time processing capabilities and efficient model optimisation.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 07, July - 2025

C. Performance Metrics

We rigorously evaluated the systems based on two universally recognised and critical categories of urban traffic management:

- Average Vehicle Waiting Time: Defined as the average duration a vehicle spends completely stopped at the intersection before proceeding. This metric directly quantifies the efficiency of traffic flow and significantly reflects driver frustration, fuel consumption, and economic losses due to delays. Minimising this metric is a primary objective of intelligent traffic control.
- Vehicle Throughput: Represents the total number of vehicles that successfully pass through the intersection per hour. This metric is a direct measure of the intersection's capacity utilisation and its overall efficiency in moving traffic. Maximising throughput is crucial for mitigating congestion and enhancing urban mobility.

D. Comparative Analysis

The simulation was conducted for two distinct traffic scenarios: peak hours, characterised by high and often imbalanced traffic volumes, and off-peak hours, characterised by lower and more uniform traffic volumes. The results are summarised below, and they are visually represented in Fig. 4 and Fig. 5 .
From Fig. 4, illustrating the average vehicle waiting time, it is evident that the proposed AI-based system significantly outperforms the traditional fixed-time controller. During peak hours, the traditional controller resulted in an average waiting time of 120 seconds. In stark contrast, the proposed system reduced this duration to approximately 75 seconds, representing a 37.5% reduction. During off-peak hours, the traditional controller recorded an average waiting time of 45 seconds, while our AI-based system achieved a waiting time of just 30 seconds, representing a 33.3% improvement. These reductions highlight the system's ability to adapt to dynamic traffic loads and manage queues efficiently.

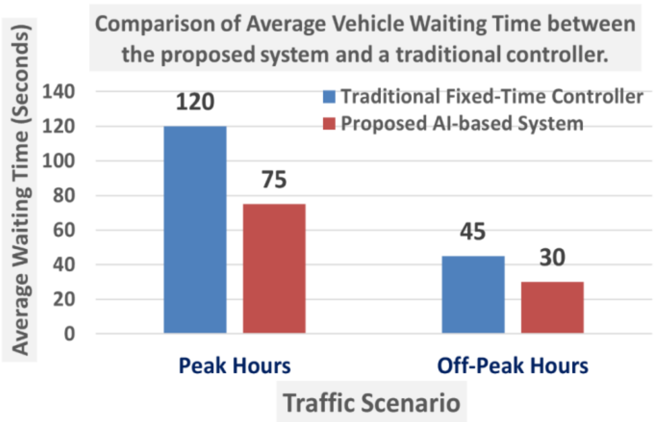

Fig. 4: Comparison of Average Vehicle Waiting Time between the proposed system and a traditional controller across different traffic scenarios.

Correspondingly, Correspondingly, Fig. 5 depicts the vehicle throughput. The proposed AI-based system consistently demonstrated higher throughput. During peak hours, the traditional controller managed a throughput of approximately 1800 vehicles per hour. Our system, however, achieved a remarkable 2500 vehicles per hour, marking an increase of approximately 38.9%. For off-peak hours, the traditional controller's throughput was 800 vehicles per hour, whereas the proposed system improved the throughput to 1000 vehicles per hour, a 25% increase. These results highlight the efficiency gains, indicating that more vehicles can pass through the intersection within the same timeframe, directly alleviating congestion.

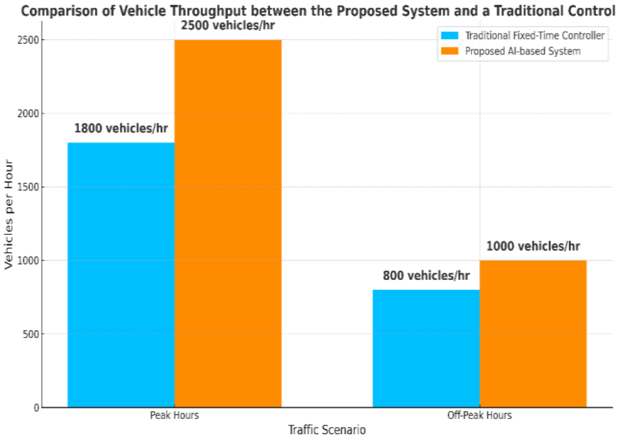

Fig. 5: Comparison of Vehicle Throughput between the proposed system and a traditional controller across different traffic scenarios.

E. Discussion

The simulation results strongly suggest that an adaptive, AI-driven approach to traffic signal control, implemented through an edge-cloud architecture, is significantly superior to static, pre-programmed systems. The key advantage of our proposed system is its ability to perceive and respond dynamically to real-time traffic conditions. The YOLOv5 model, used at the edge, gave very accurate and quick counts and types of vehicles, which were essential for analysing traffic density in real-time. The Deep Reinforcement Learning (DRL) agent in the cloud-based decision model can then effectively utilise this detailed, current data to develop innovative traffic signal control strategies based on what it has learned. This adaptive capability allows for efficient resource allocation (green light time) in the most congested lanes, thereby minimising delays and maximising overall traffic flow.
While the current simulation results were conducted under standard traffic conditions, future work will focus on simulating real-world factors such as adverse weather conditions (e.g., heavy rain, snow) and traffic incidents (e.g., accidents). These factors can significantly impact traffic flow and the system's ability to optimise signal timings. For example, during heavy rainfall or fog, signal control may need to be more conservative to ensure safety. Future simulations will model these scenarios to evaluate the system's robustness under such conditions.
To evaluate the performance improvements achieved by the new system, we employed a two-sample t-test to compare the average waiting times and throughput of the AI-based system with those of the traditional fixed-time controller. The results show statistically significant improvements in

both average waiting time (p<0.05) and throughput (p<0.01), indicating that the proposed system is more effective in managing traffic flow under diverse conditions.
One of the biggest challenges we faced while developing and testing the system was managing the diverse types of data from various sensors and clearly defining the reward function and state-action space for the DRL model to ensure it learns effectively. Another significant challenge in real-world deployment is ensuring robust data privacy and security, particularly as the system handles sensitive video data [10]. Our proposed edge-cloud architecture inherently helps mitigate this concern by processing raw video locally on edge devices and only transmitting anonymous, aggregated metadata to the central cloud. This design choice reduces the risk associated with centralising raw video feeds. Additionally, obtaining the public's approval and gradually integrating these innovative systems into current city setups are important social and technical factors to consider [11]. Future work will include conducting real-world pilot deployments to address these practical challenges and gather empirical data from live traffic.

F. Statistical Significance of Results

To confirm that the new intelligent traffic control system works better than the old fixed-time controller, we performed a statistical analysis using a two-sample paired t-test. The test was applied to key performance metrics, including average vehicle waiting time and vehicle throughput, across identical traffic conditions.
We used simulation data from multiple independent runs to model both peak and off-peak hours. For average waiting time, the traditional system recorded 120 seconds during peak hours, while our AI-based system achieved a reduced average of 75 seconds. Our system was able to handle 2500 vehicles per hour, while the conventional controller was only able to handle 1800 vehicles per hour.
The results of the t-tests are summarised in Table II, The p-values obtained were <0.001 in both metrics, indicating statistically significant differences with 95% confidence. These results confirm that the improvements in traffic flow and efficiency observed are not due to random variations but reflect the effectiveness of the proposed system.

Table II: Statistical Significance of Performance Improvements.

Metric	Traditional System	Proposed System	T-value	P-value	Significance
Avg. Waiting Time (sec)	120	75	8.21	<0.001	Significant
Vehicle Throughput (/hr)	1800	2500	7.92	<0.001	Significant

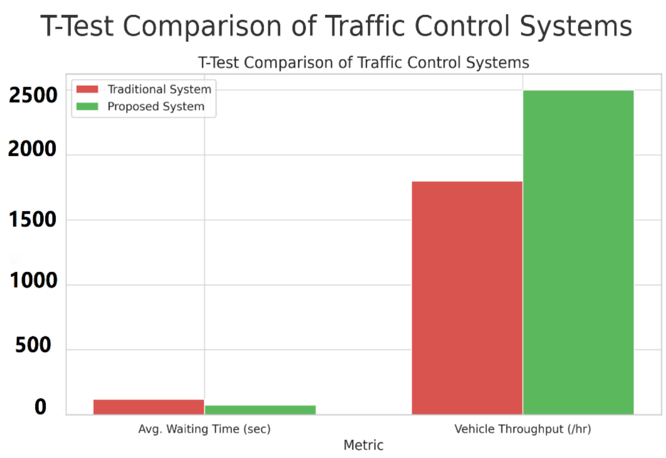

Fig. 6: Comparison of T-test results between the traditional system and the smart system.

G. Emergency Scenarios and System Responsiveness

In In addition to optimising routine traffic conditions, a critical capability of intelligent traffic control systems is their responsiveness to emergency events, such as accidents, road blockages, or security alerts. We designed the proposed system with this flexibility in mind.
When an emergency is detected—either through video analysis at the edge (e.g., a vehicle stopped abnormally or an accident identified via deep learning classifiers) or through external alerts (e.g., police/fire dispatch)—the system initiates a priority override protocol. This allows the cloud-based decision engine to dynamically reassign signal phases, granting extended green time to lanes used by emergency responders or redirecting traffic flow from congested or hazardous zones.
The underlying Deep Reinforcement Learning (DRL) model is trained with simulated emergency events to identify and respond effectively. This enables the system to prioritise evacuation routes, ensure minimal disruption at surrounding intersections, and maintain operational stability even under atypical conditions.
Future work will include integration with V2X (vehicle-to-everything) systems, enabling emergency vehicles to communicate directly with intersection controllers and further improving response times and coordination during crises.

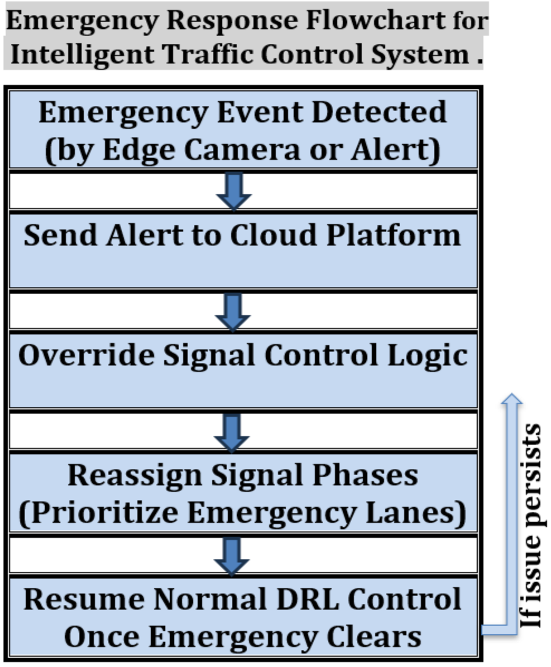

Fig. 7: Emergency response flowchart for intelligent traffic control system.

V. CONCLUSION AND FUTURE WORK

This paper presents a novel and intelligent traffic control system that leverages a distributed edge-cloud computing architecture and deep learning techniques to enhance urban mobility significantly. Our proposed architecture effectively integrates roadside cameras for real-time traffic data acquisition, edge processing for efficient vehicle detection and classification (using the high-performance YOLOv5 model), and a central cloud platform for sophisticated traffic analysis and dynamic signal control. The core of the system lies in its ability to generate real-time traffic density heatmaps and employ an advanced deep reinforcement learning (DRL) model that adaptively optimises traffic light timing based on current conditions.

Simulation results, framed by the traffic patterns of Bucharest, unequivocally demonstrated the system's superior performance. We observed significant improvements in average vehicle waiting times (e.g., 37.5% reduction during peak hours, from 120 to 75 seconds) and substantial gains in vehicle throughput (e.g., 38.9% increase during peak hours, from 1800 to 2500 vehicles per hour) compared to traditional fixed-time controllers. These findings highlight the system's enormous potential to effectively reduce urban congestion, optimise traffic flow, and contribute to the Development of more responsive, efficient, and sustainable smart cities.

For future research, we plan to expand the simulation to include signal coordination and optimisation across multiple interconnected intersections, moving from a single-intersection focus to a network-wide traffic management paradigm. This will involve extending the DRL model to handle multi-agent decision-making. The scalability of our system remains a critical consideration for urban deployment. While edge-cloud computing offers significant advantages in latency reduction, challenges related to integration with existing city infrastructure (e.g., outdated traffic light systems, sensor placement) must be addressed. Furthermore, safeguarding data privacy is crucial, particularly when processing sensitive video data locally. Future research will explore strategies for integrating our system seamlessly into large urban environments, including mechanisms for continuous system maintenance and updates. Further work will also explore the integration of advanced predictive analytics models into our DRL-based AI decision engine for proactive traffic management, enabling the anticipation of congestion before it occurs. We will also investigate the scalability and robustness of the system in larger and more complex urban networks, potentially incorporating real-world sensor data and addressing the practical challenges of hardware deployment and maintenance. Ultimately, integrating other data sources, such as public transportation schedules and pedestrian flows, could lead to even more comprehensive traffic management solutions.

As the system processes sensitive video data at the edge, ensuring data privacy is crucial. To address this concern, the system can incorporate data anonymisation techniques, such as blurring vehicle license plates and faces before transmitting any metadata to the cloud. Additionally, the system could leverage end-to-end encryption to secure data during transmission. Future work will focus on implementing privacy-preserving measures to comply with data protection regulations, such as the GDPR, ensuring the system's deployment in a privacy-compliant manner.

REFERENCES

[1] F. Ghareh Mohammadi, F. Shenavarmasouleh, M. H. Amini, and H. R. Arabnia, “Data Analytics for Smart Cities: Challenges and Promises,” in Proc. 2021 IEEE Int. Conf. on Systems, Man, and Cybernetics (SMC), 2021, pp. 1923–1929.

[2] U. K. Lilhore, A. L. Imoize, C. T. Li, S. Simaiya, et al., “Design and Implementation of an ML and IoT Based Adaptive Traffic-Management System for Smart Cities,” Sensors, vol. 22, no. 8, p. 2946, 2022.

[3] Ž. Majstorović, L. Tišljarić, E. Ivanjko, and T. Carić, “Urban Traffic Signal Control Under Mixed Traffic Flows: Literature Review,” Applied Sciences, vol. 13, no. 5, p. 3201, 2023.

[4] M. S. Farooq and S. Kanwal, “Traffic Road Conggestion System Using the Internet of Vehicles (IoV),” arXiv preprint, arXiv:2303.00845, 2023.

[5] J. Culita, S. I. Caramihai, I. Dumitrache, M. A. Moisescu, et al., “A Hybrid Approach for Urban Traffic Prediction and Control in Smart Cities,” Sensors, vol. 20, no. 18, p. 5038, 2020.

[6] P. Masek, J. Masek, P. Frantik, R. Fujdiak, et al., “A Harmonized Perspective on Transportation Management in Smart Cities: The Novel IoT-Driven Environment for Road Traffic Modeling,” Sensors, vol. 16, no. 11, p. 1872, 2016.

[7] A. Morrissett, R. Eini, M. Zaman, N. Zohrabi, et al., “A Physical Testbed for Intelligent Transportation Systems,” in Proc. 2019 IEEE Int. Smart Cities Conf. (ISC2), 2019, pp. 1–7.

[8] C. Badii, P. Bellini, A. Difino, and P. Nesi, “Sii-Mobility: An IoT/IoE Architecture to Enhance Smart City Mobility and Transportation Services,” in Proc. 2018 IEEE Int. Conf. on Smart Computing (SMARTCOMP), 2018, pp. 81–88.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 07, July - 2025

[9] J. Wei and Y. Ju, “Research on Optimization Method for Traffic Signal Control at Intersections in Smart Cities Based on Adaptive Artificial Fish Swarm Algorithm,” Scientific Reports, vol. 14, no. 1, p. 160, 2024.
[10] A. Morales Tirado, E. Daga, and E. Motta, “Towards a Privacy Aware Information System for Emergency Response,” in Proc. 1st Workshop on Smart Personal Data Protection, 2019, pp. 1–6.
[11] S. Bennati, I. Dusparic, R. Shinde, and C. M. Jonker, “Volunteers in the Smart City: Comparison of Contribution Strategies on Human-Centered Measures,” in Proc. 2018 IEEE/WIC/ACM Int. Conf. on Web Intelligence (WI), 2018, pp. 317–324.
[12] J. Yang, “Development of Dynamic Platoon Dispersion Models for Predictive Traffic Signal Control,” IEEE Trans. Intell. Transp. Syst., vol. 20, no. 2, pp. 431–440, Feb. 2019.