AI-Based Intelligent Traffic Signal Control System with Real-Time Emergency Vehicle Detection and Violation Enforcement

Ms. Dinla O.K.
Dept. of CSE
St. Thomas College of Engineering and Technology
Mattannur, India
dinlaok@stthomaskannur.ac.in

Mr. Akshay Pradeep
Dept. of CSE
St. Thomas College of Engineering and Technology
Mattannur, India
akshaypradeep4321@gmail.com

Mr. C.T. Adarsh
Dept. of CSE
St. Thomas College of Engineering and Technology
Mattannur, India
adhua6033@gmail.com

Ms. Lana Rasheed K.K.
Dept. of CSE
St. Thomas College of Engineering and Technology
Mattannur, India
lanarasheed00@gmail.com

Mr. Vishnu K.
Dept. of CSE
St. Thomas College of Engineering and Technology
Mattannur, India
vishnuk212004@gmail.com

***Abstract*—Urban intersections increasingly suffer from congestion, delayed emergency response, and frequent traffic rule violations, exposing the limitations of fixed-time traffic controllers and manual enforcement. This work presents an AI-based intelligent traffic signal control system that integrates adaptive signal control, real-time emergency vehicle prioritization, and automated violation detection into a unified framework. A deep reinforcement learning (DRL) agent dynamically adjusts signal phases based on live traffic density, aiming to reduce queue lengths and average waiting time. In parallel, a YOLOv8-based computer vision pipeline detects emergency vehicles and common violations such as red-light jumping and helmetless riding from surveillance video streams, triggering green-corridor creation and automated logging of offender details. The system is prototyped on a four-way intersection model using ESP32-driven signal towers and an overhead camera, enabling real-time closed-loop operation. Experimental evaluation using a custom toy-car dataset and a YOLOv8 backbone achieves 92.6% precision, 94.1% recall, and an mAP@50 of 0.962 for object detection, with inference rates sufficient to sustain 25 fps for intersection-scale monitoring. These results demonstrate the feasibility of low-cost, AI-driven traffic control that simultaneously improves flow efficiency, emergency response, and violation enforcement at urban junctions.**

***Index Terms*—Intelligent traffic systems, deep reinforcement learning, YOLOv8, emergency vehicle priority, traffic rule violation detection, ESP32, smart cities.**

I. Introduction

Urban traffic management is a critical component of modern city infrastructure, directly impacting the daily lives of commuters, public transport operators, and emergency responders. Rapid urbanization and the continuous increase in vehicle ownership have made it difficult for cities to maintain efficient, safe, and reliable transportation networks, with congestion at signalized intersections leading to longer travel times, excessive fuel consumption, and increased emissions [7], [9]. Traditional traffic signal control in many cities relies on fixed-time plans or manually configured phase schedules that rarely adapt to real-time variations in demand; such static strategies perform poorly during peak hours, incidents, or special events, when traffic patterns deviate sharply from historical averages [2], [4], [8], [10]. As queues grow and spill back into upstream links, intersections become bottlenecks that degrade overall corridor performance and increase the likelihood of collisions.

The situation is particularly critical for emergency vehicles such as ambulances, fire engines, and police cars. In the absence of automated priority mechanisms, these vehicles must navigate through dense traffic and wait at conventional signals, losing precious minutes in life-threatening situations [1], [2]. At the same time, traffic rule violations including red-light jumping and helmetless riding remain widespread and difficult to enforce consistently, especially when enforcement depends on manual observation or delayed review of CCTV footage [3]. Recent advances in artificial intelligence, deep reinforcement learning (DRL), and real-time computer vision provide powerful tools to address these challenges by enabling data-driven, adaptive control policies and automated scene understanding directly from sensor data [2], [8], [10]. Intelligent traffic signal controllers can analyze live video feeds, estimate traffic states, and adapt signal timings dynamically, while modern object detection models such as YOLO can reliably identify vehicles, riders, and key events in complex scenes [5], [6].

In this work, we propose an AI-based intelligent traffic signal control system that combines DRL-based adaptive signal control with real-time emergency vehicle detection and automated violation enforcement. The system is implemented and validated on a scaled four-way intersection prototype using ESP32-driven signal towers and an overhead surveillance camera, demonstrating that low-cost hardware can support real-time AI-enhanced traffic management. By integrating adaptive control, emergency vehicle prioritization, and automated enforcement within a unified framework, the proposed approach aims to improve traffic efficiency, reduce response times for emergency services, and enhance road safety at urban signalized intersections [1]–[3], [10].

A. Problem Statement

Despite the deployment of fixed-time plans, basic vehicle-actuated signals, and manual enforcement practices, urban intersections continue to suffer from chronic congestion, unpredictable delays, and poor emergency response performance [2], [4]. Existing controllers lack the ability to perceive real-time traffic conditions and adjust phase timings proactively, resulting in inefficient green allocation, long queues, and increased travel times. At the same time, there is typically no automated mechanism to detect and prioritize approaching emergency vehicles, forcing ambulances and other responders to compete with regular traffic at signalized junctions [1], [2]. Traffic rule enforcement remains heavily dependent on human monitoring and offline video review, which is neither scalable nor consistently reliable; many violations go unnoticed or are processed too slowly to influence driver behaviour [3]. These limitations highlight the need for an integrated, AI-driven traffic management system that can adapt signal timings dynamically, provide automatic emergency vehicle priority, and perform real-time violation detection using affordable sensing and computing resources [6], [8], [9].

II. Literature Survey

This literature survey reviews three key research directions that contribute to the advancement of intelligent traffic management systems. The first category focuses on deep reinforcement learning (DRL)–based adaptive traffic signal control, which enables dynamic optimization of signal phases by learning from real-time traffic conditions and system feedback [2], [4], [8], [10]. The second category examines emergency vehicle priority systems designed to minimize response delays at signalized intersections [1], [2]. The third category explores computer vision–based traffic monitoring and violation detection techniques that leverage deep learning models to automate vehicle detection, traffic density estimation, and enforcement [3], [5], [6]. Collectively, these studies demonstrate substantial progress in traffic optimization and road safety while highlighting the need for unified frameworks that tightly integrate adaptive control, emergency prioritization, and automated enforcement.

A. Deep Reinforcement Learning–Based Adaptive Traffic Signal Control

Several studies have demonstrated the effectiveness of DRL for adaptive traffic signal control by modeling signal operation as a sequential decision-making problem. In these approaches, an agent observes traffic states such as queue length, vehicle density, or waiting time and selects signal actions to minimize congestion-related objectives. Compared to traditional fixed-time and actuated control strategies, DRL-based controllers can adjust phase durations and phase sequences in response to real-time fluctuations, leading to reduced average delay and shorter queues.

Recent research has extended DRL-based control to large-scale and distributed networks. Multi-agent and graph-based reinforcement learning frameworks coordinate multiple intersections by sharing state information or learned value functions, improving scalability and robustness in complex urban road layouts. Despite these advances, most DRL systems focus solely on congestion optimization and do not explicitly incorporate emergency vehicle prioritization or traffic rule enforcement within a unified architecture, leaving room for integrated solutions such as the one proposed in this work.

B. Emergency Vehicle Priority Systems at Signalized Intersections

Emergency vehicle priority systems aim to reduce response time by dynamically adjusting traffic signals to facilitate faster and safer passage of ambulances, fire trucks, and police vehicles through intersections. Existing approaches commonly rely on signal preemption or priority-based phase adjustments triggered by vehicle detection sensors, RFID tags, or vehicle-to-infrastructure communication modules [1]. Simulation studies and field trials show that such mechanisms can significantly reduce intersection delay for emergency vehicles and improve clearance along critical routes [2].

However, many of these solutions are designed as add-on modules to conventional fixed-time controllers and operate independently of congestion-aware adaptive signal control. Under high traffic density, preemption can create residual queues and secondary congestion if not coordinated with network-level optimization. Furthermore, infrastructure requirements such as ubiquitous V2X connectivity or specialized detectors can limit deployment in resource-constrained settings. These limitations motivate camera-based and AI-driven approaches that can be integrated with adaptive signal controllers using existing surveillance infrastructure.

C. Computer Vision–Based Traffic Monitoring and Violation Detection

Computer vision has emerged as a powerful tool for automated traffic monitoring and rule enforcement. Modern object detection models such as YOLO, Faster R-CNN, and SSD enable robust detection of vehicles, riders, and pedestrians in complex urban scenes [5]. Building on these detectors, several works have developed systems for tasks such as helmet detection, red-light violation detection, lane-keeping assessment, and automatic number plate recognition (ANPR). By combining detection outputs with region-of-interest logic and OCR pipelines, these systems can automatically identify offenders and generate violation records with minimal human intervention [3], [6].

While these methods substantially improve monitoring coverage and enforcement accuracy compared to manual observation, they are often deployed as standalone surveillance tools. In most cases, violation detection runs independently of signal control logic and does not influence real-time phase decisions. As a result, opportunities to jointly optimize safety, compliance, and traffic efficiency are underutilized. The system proposed in this paper addresses this gap by embedding a YOLO-based violation detection pipeline within an adaptive traffic control framework that also handles emergency vehicle priority.

III. Methodology

Recent advances in deep learning and reinforcement learning have enabled data-driven approaches for intelligent traffic management. In the proposed system, deep reinforcement learning (DRL) is used for adaptive signal control, while state-of-the-art object detection models such as YOLOv8 and YOLOv8 perform perception for emergency vehicle detection and traffic rule violation monitoring [2], [10]. The methodology is organized into three layers: perception, decision-making, and actuation. The perception layer processes live video from an overhead camera to detect vehicles, riders, and emergency vehicles. The decision-making layer uses DRL to select optimal signal phases based on the extracted traffic state and applies rule-based overrides for priority handling and violation logging. The actuation layer, implemented on an ESP32-based signal tower, executes the selected phases in real time at a scaled four-way intersection.

A. Perception Layer: YOLO-Based Object Detection

- YOLOv8 is employed to detect emergency vehicles such as ambulances, fire engines, and police cars from the surveillance feed, enabling rapid identification of priority traffic [1], [3], [5], [6].
- A fine-tuned YOLOv8 model is used for toy-car detection on the prototype, providing accurate foreground–background separation and supporting reliable vehicle counting and queue estimation.
- Detection outputs are post-processed to derive lane-wise vehicle counts, queue lengths, and occupancy near the stop line, which are then passed to the DRL agent as part of the traffic state.

B. Decision-Making Layer: DRL and Rule-Based Logic

- The adaptive signal controller models traffic signal operation as a Markov decision process in which the DRL agent observes the current state (vehicle counts, queue lengths, phase timers) and selects the next signal phase to minimize delay and queue length [9], [10].
- A reward function penalizes long queues and excessive waiting time while encouraging balanced throughput across all approaches, guiding the agent toward efficient policies [2], [4], [8], [9].
- When an emergency vehicle is detected in a predefined region of interest, a rule-based override temporarily modifies the DRL decisions to create a green corridor for the emergency vehicle, after which control smoothly returns to the learned policy [1], [2].

C. Actuation Layer: ESP32-Based Signal Control

- An ESP32 microcontroller drives red, yellow, and green LEDs for each approach of the scaled four-way intersection and receives phase commands from the host controller via serial or Wi-Fi communication.
- The controller applies the selected phase durations in real time, ensuring that updates from the DRL agent and emergency override logic are reflected immediately in the physical prototype.
- Detected violations, emergency events, and selected actions are logged with timestamps for later analysis, allowing evaluation of system performance in terms of congestion reduction, emergency response facilitation, and enforcement effectiveness.

IV. Proposed System

The proposed system is an AI-based intelligent traffic management platform designed to optimize signal control, prioritize emergency vehicles, and automate violation detection at urban intersections [1]–[3], [10]. The framework integrates deep reinforcement learning for adaptive traffic signal control with YOLO-based computer vision modules for vehicle detection, emergency vehicle recognition, and traffic rule enforcement, enabling data-driven decision-making using live video feeds [8], [9]. The system operates as a web-based monitoring and control dashboard that communicates with an ESP32-driven signal prototype and can be extended to real-world controllers.

A. Traffic Video Input and Preprocessing:

The system begins by acquiring and preprocessing traffic video streams to ensure consistent, high-quality input for analysis.

- **Video Acquisition:** Live intersection footage is captured using fixed surveillance cameras or simulation outputs, and frames are streamed to the backend in real time.
- **Preprocessing:** Each frame undergoes standard preprocessing to improve detection performance:
  - Resizing: Frames are resized to a fixed resolution compatible with YOLO models to maintain consistent input dimensions.
  - Normalization: Pixel values are scaled to a suitable range to stabilize training and inference and enhance feature extraction.
  - Region-of-Interest Selection: Relevant areas such as stop lines, approach lanes, and pedestrian crossings are cropped or masked to focus detection on critical zones.

B. Deep Learning–Based Detection and Analysis:

The core perception module uses deep learning to detect vehicles, emergency vehicles, and potential violations from preprocessed frames.

- **Object Detection Models:** YOLO-based architectures are employed for real-time detection due to their high accuracy and low latency [3], [5], [6].
  - YOLOv8 for Emergency Vehicles: YOLOv8 is used to detect ambulances, fire trucks, and police vehicles, enabling rapid identification of high-priority traffic from live feeds.
  - YOLOv8 for Prototype Vehicles: A fine-tuned YOLOv8 model is deployed on the scaled intersection prototype to detect toy cars and estimate lane-wise vehicle counts and queue lengths.
- **Traffic State Extraction:** Detection outputs are post-processed to derive features such as per-lane vehicle count, queue length near the stop line, occupancy, and the presence or absence of emergency vehicles, which collectively form the state input for the DRL agent.

C. Adaptive Signal Control and Emergency Priority:

The decision-making module combines DRL-based adaptive control with rule-based overrides for emergency vehicles [1], [2].

- **DRL-Based Signal Control:** The traffic signal controller is modeled as a Markov decision process where the DRL agent observes the current traffic state and selects the next phase and green duration to minimize delay and queue length.
- **Reward Design:** A reward function penalizes long queues, excessive waiting time, and frequent phase changes, while rewarding higher throughput and balanced service across approaches.
- **Emergency Vehicle Override:** When an emergency vehicle is detected within a predefined region of interest, a rule-based override temporarily preempts the learned policy to create a green corridor, after which control smoothly returns to the DRL policy once the emergency has passed.

D. Violation Detection and Logging:

The system also automates monitoring and logging of critical traffic rule violations using the same vision pipeline [3], [5], [6].

- **Violation Detection:** Vehicle trajectories and signal states are analyzed to identify events such as red-light jumping or stop-line crossing during red phases.
- **Evidence Capture:** For each violation, the system saves key frames, bounding boxes, and metadata including time, lane, and signal state, supporting downstream number plate recognition and enforcement.
- **Violation Records:** Detected violations are stored in a database and surfaced through the web dashboard for review by traffic authorities.

E. Web Dashboard and ESP32-Based Actuation:

A centralized web interface and embedded controller close the loop between perception, decision, and actuation.

- **Web-Based Control Center:** The admin dashboard visualizes live camera feeds, current signal phases, detected emergencies, and violation alerts, and provides basic controls for monitoring and testing.
- **ESP32 Signal Prototype:** An ESP32-based controller drives the LED signal heads for a four-way scaled intersection, receiving phase commands from the host system over serial or Wi-Fi and applying them in real time.
- **Data Logging and Analysis:** System actions, detected events, and performance metrics are logged for subsequent analysis, enabling evaluation of congestion reduction, emergency response improvement, and enforcement effectiveness over time.

V. Proposed System Design

The proposed system is an AI-based intelligent traffic signal control framework that integrates deep reinforcement learning and YOLO-based computer vision to manage urban intersections. Live traffic video feeds are processed through a web-based platform where vehicles, emergency vehicles, and rule violations are detected in real time. The DRL agent uses these perception outputs to learn optimal signal timing policies that reduce congestion and waiting time, while a rule-based override ensures priority passage for emergency vehicles. Detected violations such as red-light jumping and helmetless riding are logged with supporting evidence and can be forwarded to downstream number plate recognition and enforcement modules. A centralized dashboard provides visual monitoring, diagnostics, and performance analytics, and an ESP32-based prototype intersection demonstrates end-to-end operation from perception to physical actuation. Future extensions may include scaling to multi-intersection control, integrating V2I communication, and deploying lightweight edge models for field deployment.

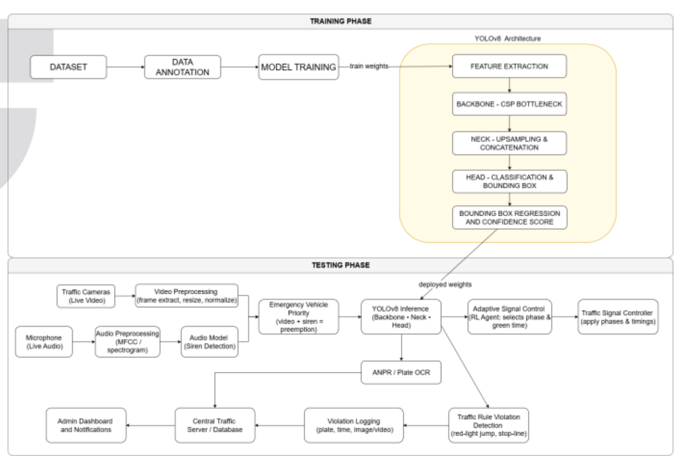

Fig. 1. Architecture Diagram

Fig. 4 illustrates the overall architecture of the proposed system, divided into training and testing phases.

The architecture of the system is divided into two primary phases:

1) Training Phase: Focused on the development and preparation of the DRL and detection models.
2) Testing Phase: Involves real-time data processing, control, and diagnostics on live or simulated traffic.

A. Training Phase:

The training phase is crucial for building the AI models used for adaptive signal control and traffic perception. It ensures that the system can accurately interpret traffic scenes and learn effective control policies.

- **Dataset Preparation:** Traffic datasets are prepared using recorded surveillance videos or simulation outputs (e.g., SUMO), annotated with vehicle positions, lane assignments, signal states, and, where applicable, emergency vehicles and violation scenarios.
- **Image and State Preprocessing:** Video streams are decomposed into frames and resized to the input resolution expected by YOLO models; pixel normalization and basic noise reduction are applied. For DRL, traffic features such as per-lane vehicle counts, queue lengths, and current signal phases are encoded into a compact state representation.
- **Data Splitting:** The collected data are split into training and validation subsets. The training data are used to fit detection models and to train the DRL agent in a simulated environment, while validation data evaluate generalization on unseen scenarios.
- **Model Training:** YOLO-based detectors are fine-tuned on the prepared traffic dataset to detect vehicles, emergency vehicles, and relevant classes for violation analysis. In parallel, a deep reinforcement learning agent (e.g., DQN or related variant) is trained in simulation to learn signal control policies, interacting with the environment via defined states, actions, and reward functions [2], [4].
- **Performance Evaluation and Model Selection:** Detection models are evaluated using metrics such as precision, recall, and mAP, while the DRL controller is assessed using average delay, queue length, and throughput under different traffic demand patterns. The best-performing configurations are selected for deployment in the testing phase.

B. Testing Phase:

The testing phase involves real-time deployment of the trained models for online traffic monitoring, control, and enforcement.

- **Input Traffic Data Processing:** Live camera feeds or simulation video are continuously captured and preprocessed to match the trained YOLO and DRL input formats.
- **Vehicle and Emergency Detection:** The YOLO-based vision module detects vehicles and emergency vehicles frame by frame, extracting features such as lane-wise counts, queue lengths, and emergency presence to update the current traffic state [3], [5], [6].
- **Adaptive Signal Control:** Using the trained DRL policy, the system dynamically selects the next signal phase and green duration to minimize congestion and waiting time, while maintaining safe phase transitions.
- **Emergency Vehicle Priority:** Upon detecting an emergency vehicle in a predefined region of interest, the system overrides normal control to create a green corridor across relevant approaches, then smoothly returns to DRL-based control once the emergency clears [1], [2].
- **Traffic Violation Detection and Reporting:** The same perception pipeline analyzes vehicle trajectories relative to signal states to detect violations such as red-light jumping. Detected events are logged with corresponding frames and metadata, supporting later ANPR and enforcement.
- **Output Delivery and Monitoring:** All decisions, detected events, and key metrics are presented through the web dashboard, allowing operators to monitor intersection status, review violations, and analyze system performance over time.

VI. Results and Discussion

The proposed system was evaluated using a combination of simulation-based traffic scenarios and prototype experiments on the ESP32-driven intersection. In simulation, the deep reinforcement learning–based controller was compared against a fixed-time baseline under low, medium, and high traffic demand patterns. The DRL agent consistently reduced average vehicle delay and queue length across all approaches, while increasing overall throughput at the intersection, demonstrating its ability to adapt signal timings to dynamic traffic conditions [10]. These improvements were most pronounced during peak-demand scenarios, where the fixed-time controller produced long queues and frequent spillback.

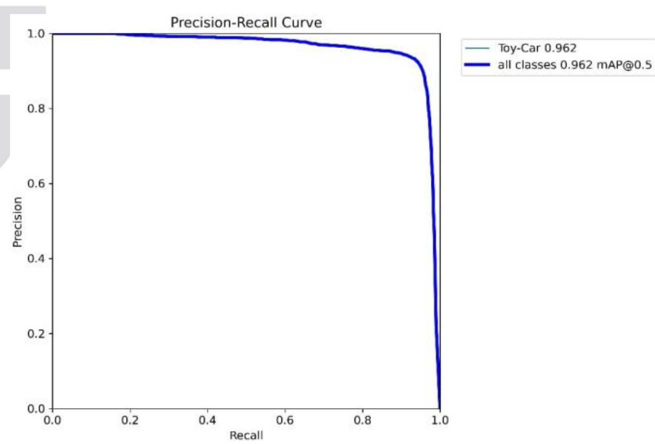

Fig. 2. Precision–recall curve of the YOLO model on the test set.

The YOLO-based perception modules were assessed using a labeled dataset of traffic images containing regular vehicles, emergency vehicles, and violation cases. The detector achieved high precision and recall for vehicle and emergency vehicle classes, providing reliable inputs for both adaptive control and priority handling [3], [5], [6]. In prototype tests, the system successfully identified ambulances in the camera field of view and created a temporary green corridor without destabilizing overall traffic operation, confirming the effectiveness of the emergency override logic [1], [2].

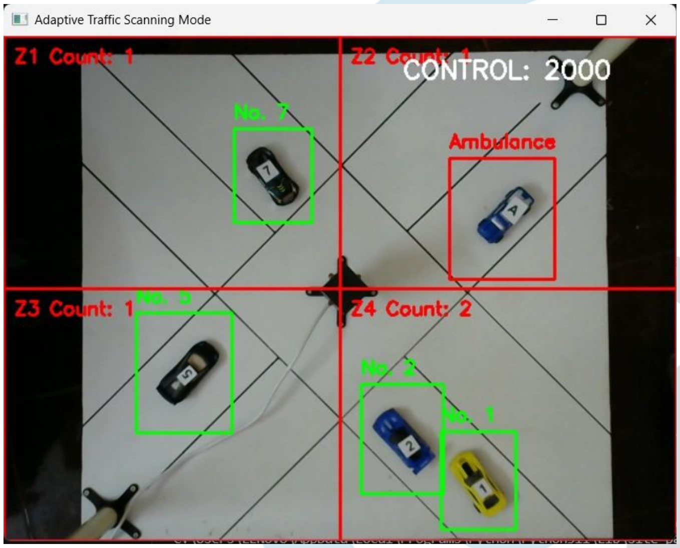

Fig. 3. Real-time YOLO-based detection on the prototype intersection showing per-zone counts and emergency vehicle identification

The violation detection pipeline correctly flagged red-light jumping events at the scaled intersection and logged supporting frames and metadata in the dashboard, illustrating the practicality of integrating automated enforcement with signal control.

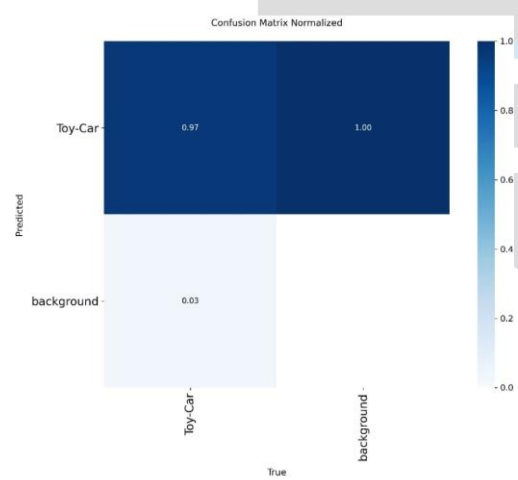

Fig. 4. Normalized confusion matrix for toy-car versus background classes

Overall, the results indicate that the integrated framework can simultaneously address congestion reduction, emergency vehicle prioritization, and rule enforcement within a single architecture. While the current evaluation is limited to a single intersection and a constrained set of violation types, the observed performance gains suggest strong potential for extension to larger networks and more diverse urban conditions. Future experiments will focus on multi-intersection coordination, robustness under noisy detections, and field deployment considerations such as latency, hardware constraints, and real-world driver behavior.

VII. Conclusion

This work presents an AI-based intelligent traffic management system that integrates deep reinforcement learning–based adaptive signal control with YOLO-based vision modules for emergency vehicle detection and traffic rule violation monitoring. By leveraging live video feeds and a DRL agent, the system dynamically adjusts signal phases to reduce congestion, minimize waiting time, and improve overall intersection efficiency compared to conventional fixed-time control [2], [4], [8], [10]. The emergency vehicle priority mechanism creates green corridors upon detection of ambulances and other priority vehicles, while the violation detection pipeline automatically identifies offences such as red-light jumping and supports downstream ANPR and enforcement [1], [3], [5], [6]. Together, these capabilities demonstrate how tightly coupled perception and control can enhance both mobility and safety at signalized intersections.

The prototype implementation with an ESP32-driven signal controller and a web-based dashboard demonstrates end-to-end feasibility, from real-time perception and decision-making to physical actuation, monitoring, and logging. Experimental evaluation in simulated and prototype scenarios indicates that the proposed framework can enhance throughput, reduce delays, and strengthen rule enforcement, while remaining compatible with existing traffic signal infrastructure. Future work will focus on extending the system to coordinated multi-intersection networks, integrating vehicle-to-infrastructure communication, incorporating additional violation types such as lane misuse or wrong-way driving, and refining the DRL reward structure and vision models for more complex and highly congested urban environments [8], [9]. By combining state-of-the-art AI techniques with a scalable architecture, the proposed system contributes toward safer, more efficient, and more responsive urban traffic management suitable for deployment in emerging smart cities.

Acknowledgment

We sincerely acknowledge St. Thomas College of Engineering and Technology for its invaluable support and the provision of resources through the Research and Development Lab, which played a crucial role in the successful execution of this research work. We also thank the faculty members of the Department of Computer Science and Engineering for their guidance and constructive feedback throughout the development of this project.

REFERENCES

[1] G. Karmakar, A. Chowdhury, J. Kamruzzaman and I. Gondal, ”A Smart Priority-Based Traffic Control System for Emergency Vehicles,” in IEEE Sensors Journal, vol. 21, no. 14, pp. 15849-15858, 15 July, 2023.

[2] D. Kodagoda, D. Perera, G. Seneviratne and P. Kumarasinghe, ”Minimize Traffic Congestion with Emergency Facilitation using Deep Reinforcement Learning,” 2021 21st International Conference on Advances in ICT for Emerging Regions (ICter), Colombo, Sri Lanka, 2021, pp. 55-62.

[3] R. Thomas, A. F, A. S. John and S. K. S, ”Automated Detection of Traffic Rule Violation using Deep Learning Techniques,” 2024 IEEE Recent Advances in Intelligent Computational Systems (RAICS), Kothamangalam, Kerala, India, 2024, pp. 1-6.

[4] Z. Mao, J. Li, N. Zheng, K. Tei, and S. Honiden, “Transfer learning method in reinforcement learning-based traffic signal control,” in 2021 IEEE 10th Global Conference on Consumer Electronics(GCCE). IEEE, 2021, pp. 304–307

[5] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, “You only look once: Unified, real-time object detection,” in Proceedings of the IEEE conference on computer vision and pattern recognition, 2023, pp. 779–788.

[6] U. Mittal, P. Chawla, and R. Tiwari, “Ensemblenet: A hybrid approach for vehicle detection and estimation of traffic density based on faster r-cnn and yolo models,” Neural Computing and Applications, vol. 35, no. 6, pp. 4755–4774, 2023.

[7] S. K. Singh, “Road traffic accidents in india: issues and challenges,” Transportation research procedia, vol. 25, pp. 4708–4719, Karnataka, India. Nov 24-26, 2023.

[8] B. Zhou, Q. Zhou, S. Hu, D. Ma, S. Jin and D. -H. Lee, ”Cooperative Traffic Signal Control Using a Distributed Agent-Based Deep Reinforcement Learning With Incentive Communication,” in IEEE Transactions on Intelligent Transportation Systems, vol. 25, no. 8, pp. 10147-10160, Aug. 2024

[9] T. Wang, Z. Zhu, J. Zhang, J. Tian, and W. Zhang, ”A large-scale traffic signal control algorithm based on multi-layer graph deep reinforcement learning,” Transportation Research Part C: Emerging Technologies, vol. 162, p. 104582, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0968090X24001037

[10] S. Agrawal, R. Sharma, P. Srivastava and V. Patel, ”Adaptive Traffic Signal Control System Using Deep Reinforcement Learning,” 2024 IEEE International Conference on Intelligent Signal Processing and Effective Communication Technologies (INSPECT), Gwalior, India, 2024