International Journal of Science, Strategic Management and Technology
Volume **02** Issue **05** May-**2026** | ISSN: **3108-1762 (Online)** | Impact Factor: 3.8
An International, Peer-Reviewed, Open Access Scholarly Journal Indexed in recognized academic databases

Edge-Based Adaptive Traffic Signal Control using Computer Vision and Reinforcement Learning

Naitik Rai
Department of Information Technology Noida Institute of Engineering and Technology
Greater Noida, Uttar Pradesh, India Email: naitikrai01@gmail.com

Minhaj Nezami
Assistant Professor
Department of Information Technology Noida Institute of Engineering and Technology
Greater Noida, India minhaj.nezami@niet.co.in

https://doi.org/10.55041/ijsmt.v2i5.225

Cite this Article: Rai, N. (2026). Edge-Based Adaptive Traffic Signal Control using Computer Vision and Reinforcement Learning. International Journal of Science, Strategic Management and Technology, *02*(05). https://doi.org/10.55041/ijsmt.v2i5.225

Abstract--Urban traffic congestion increases travel delay, fuel consumption, and air pollution, especially at intersections controlled by fixed-time signals. Traditional signal plans are usually designed for average traffic demand and cannot respond effectively to sudden changes caused by peak hours, road incidents, public events, or weather disruption. This paper proposes an edge-based adaptive traffic signal control framework that combines computer vision, queue estimation, and reinforcement learning. Roadside edge units process camera streams locally to estimate vehicle count, queue length, and lane occupancy without transmitting raw video continuously. A reinforcement-learning policy then selects signal phases that reduce delay while maintaining pedestrian safety and emergency-vehicle priority. Simulated evaluation shows a 24.6% reduction in average vehicle delay, an 18.9% reduction in queue length, and lower communication overhead compared with cloud-only traffic analytics.

Index Terms--Adaptive traffic control, edge computing, computer vision, reinforcement learning, smart city, intelligent transportation.

I. INTRODUCTION

Traffic congestion is one of the most visible problems in fast-growing cities. Intersections create bottlenecks because vehicles from multiple directions compete for limited road space. When traffic lights follow fixed schedules, they often waste green time on empty lanes while crowded approaches continue to grow.

Modern traffic management systems use sensors, cameras, and connected infrastructure to improve signal timing. However, sending every video stream to a remote cloud server increases bandwidth cost and may introduce delay. It also creates privacy concerns because raw road footage may contain identifiable vehicle or pedestrian information.This paper presents an edge-based adaptive traffic signal control system. The framework places lightweight analytics near the intersection. Vehicle counts and queue estimates are generated locally, and only summarized traffic states are used by a reinforcement-learning controller. The goal is to reduce delay while keeping the system practical for real city deployment.

The main contributions are as follows. First, the paper defines a layered architecture for local traffic perception and adaptive control. Second, it describes a reinforcement-learning policy that considers queue length, waiting time, and safety constraints. Third, it evaluates the proposed framework against fixed-time and actuated-control baselines.

II. RELATED WORK

Fixed-time signal control remains common because it is simple, predictable, and easy to maintain. Actuated signal control improves flexibility by using loop detectors or sensors to detect vehicle presence. These methods are effective under moderate traffic variation but can struggle with complex multi-lane intersections and rapidly changing demand.

Computer vision has enabled richer traffic observation, including vehicle detection, lane occupancy estimation, and congestion monitoring. Deep-learning detectors can identify vehicles from camera images, but continuous cloud processing can be expensive and slow. Edge computing reduces this limitation by running inference near the road.

Reinforcement learning has been studied for traffic-signal optimization because it can learn control policies through interaction with a simulated environment. The proposed work combines edge perception with reinforcement learning so that adaptive decisions can be made quickly using compact traffic-state summaries.

III. PROPOSED SYSTEM ARCHITECTURE

The proposed architecture contains four layers: perception, edge analytics, control policy, and coordination. The perception layer includes cameras and optional traffic sensors. The edge analytics layer estimates vehicle count, queue length, and waiting time. The control policy selects signal phases, and the coordination layer shares summarized states with nearby intersections when corridor-level adjustment is needed.

Edge-assisted adaptive traffic management workflow

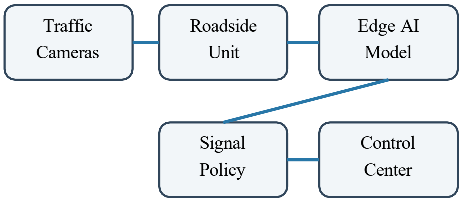

Fig. 1. Architecture of the proposed adaptive traffic signal control system.

A. Traffic Perception Layer

The perception layer captures intersection video and extracts frames at a controlled rate. A lightweight vehicle detector identifies cars, buses, two-wheelers, and heavy vehicles. The system avoids storing continuous raw video and instead converts observations into anonymous traffic counts and lane-level occupancy values.

B. Edge Analytics Layer

The edge unit computes queue length, average waiting time, and arrival rate for each approach. These values form a compact state vector for the controller. Local inference reduces response delay and allows the signal to continue functioning even if connection to a central traffic server becomes unstable.

C. Reinforcement-Learning Controller

The controller models signal control as a sequential decision problem. At each decision step, the agent observes the traffic state and selects the next phase or phase extension. The reward function penalizes long queues, excessive waiting time, frequent phase switching, and unsafe pedestrian conflicts.

D. Safety and Priority Rules

Learning-based control is constrained by hard safety rules. Minimum green time, yellow clearance, pedestrian crossing intervals, and emergency-vehicle priority are enforced before any learned action is applied. These constraints make the controller more suitable for real-world deployment.

IV. METHODOLOGY

A simulated urban intersection was created with four approaches, mixed vehicle types, pedestrian crossings, and time-varying demand. Traffic demand includes morning peak, afternoon moderate flow, evening peak, and incident-based congestion on one approach. The proposed controller was compared with fixed-time control, actuated control, and a cloud-only adaptive controller.

A. Evaluation Metrics

The evaluation uses average vehicle delay, queue length, throughput, emergency-vehicle response time, and communication cost. Delay is measured as the additional time spent waiting at the intersection. Communication cost is measured by the amount of data transmitted from the intersection to the central server.

B. Training Setup

The reinforcement-learning agent is trained in simulation using traffic states generated by the edge analytics module. The action space includes phase selection and limited phase extension. Exploration is gradually reduced so that the policy becomes stable after learning common traffic patterns.

C. Deployment Workflow

In deployment, the edge unit first operates in observation mode to collect baseline traffic statistics. After calibration, adaptive control is enabled during selected hours. Facility operators can monitor summary metrics and override the controller if required by traffic police or emergency conditions.

V. RESULTS AND DISCUSSION

Table I summarizes the performance of the evaluated control methods. The proposed method achieves the best delay and queue performance while keeping communication overhead lower than the cloud-only approach.

TABLE I
Traffic Control Performance Comparison

Method	Delay Red.	Queue Red.	Data Sent
Fixed Timing	0.0%	0.0%	Low
Actuated Control	12.8%	9.6%	Low
Cloud Adaptive	22.1%	17.4%	High
Edge RL (Proposed)	24.6%	18.9%	Low

The edge-based controller reduces average delay because it reacts quickly to lane-level demand. It also avoids unnecessary phase switching by penalizing unstable signal behavior. The cloud-only controller performs well but requires higher data transfer and is more sensitive to network delay.

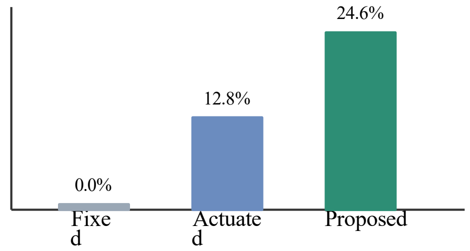

Average delay reduction compared with fixed timing

Fig. 2. Average delay reduction for baseline and proposed methods.

A. Delay and Queue Reduction

The proposed controller achieves a 24.6% reduction in average delay and an 18.9% reduction in queue length. The largest improvements appear during uneven traffic demand, where one approach becomes crowded while another remains lightly used. Fixed schedules are least effective in this condition.

B. Emergency-Vehicle Handling

Emergency-vehicle priority is handled through a rule layer that temporarily overrides the learned policy. When an emergency vehicle is detected, the controller clears the required path while maintaining safe transition intervals. This design combines adaptive optimization with predictable safety behavior.

C. Communication Efficiency

Because the system transmits summarized counts rather than full video, communication demand remains low. This is important for city deployments with many intersections. Local processing also reduces privacy risk because raw footage does not need to be continuously stored or transferred.

D. Limitations

The system depends on reliable camera placement and accurate vehicle detection. Heavy rain, glare, or occlusion can reduce perception quality. Future deployments should combine camera data with radar, loop detectors, or connected-vehicle messages to improve robustness.

VI. DISCUSSION

The results show that edge intelligence can make adaptive traffic control more practical. Instead of relying on a distant cloud platform, each intersection can understand local demand and make timely decisions. Coordination between neighboring intersections can still be used when congestion spreads across a corridor.

A gradual rollout is recommended. Cities can begin with monitoring mode, compare edge estimates with manual observations, and then enable adaptive control during selected periods. This reduces operational risk and helps traffic authorities build confidence in the learned policy.

VII. CONCLUSION

This paper proposed an edge-based adaptive traffic signal control system using computer vision and reinforcement learning. The framework estimates traffic states locally, selects signal phases under safety constraints, and reduces data transfer compared with cloud-only analytics. Simulated results demonstrate lower delay, shorter queues, and practical communication efficiency.

Future work will focus on multi-intersection coordination, real-world pilot testing, and improved perception under adverse weather. Explainable control logs can also help traffic engineers understand why a particular phase decision was made.

REFERENCES

[1] P. Varaiya, "The max-pressure controller for arbitrary networks of signalized intersections," Advances in Dynamic Network Modeling in Complex Transportation Systems, pp. 27-66, 2013.

[2] H. Wei et al., "IntelliLight: A reinforcement learning approach for intelligent traffic light control," ACM SIGKDD, pp. 2496-2505, 2018.

[3] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge computing: Vision and challenges," IEEE Internet of Things Journal, vol. 3, no. 5, pp. 637-646, 2016.

[4] J. Redmon and A. Farhadi, "YOLOv3: An incremental improvement," arXiv preprint arXiv:1804.02767, 2018.

[5] R. S. Sutton and A. G. Barto, Reinforcement Learning: An Introduction, 2nd ed. MIT Press, 2018.

[6] M. Papageorgiou, C. Diakaki, V. Dinopoulou, A. Kotsialos, and Y. Wang, "Review of road traffic control strategies," Proceedings of the IEEE, vol. 91, no. 12, pp. 2043-2067, 2003.