BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

AI-BASED TRAFFIC CONTROLLING SYSTEM : *A Literature Review of Computer Vision, Reinforcement Learning, and IoT Approaches*

[1]Tejas Pavali, [2]Prof. Supriya Survase

[1]Student Researcher; [2]Assistant Professor and Project Guide
[1],[2]Department of Artificial Intelligence and Data Science
[1],[2]Ajeenkya DY Patil School of Engineering, Pune, India
[1]Email: tejaspavali567@gmail.com
[2]Email: supriyasurvase@dypic.in

Abstract

Traffic jams become a frequent issue in rapidly growing cities; however, the time-fixed lights do not adjust properly to changing car traffic. In this paper, different approaches using the following elements: artificial intelligence, computer vision, reinforcement learning, deep learning, and Internet of Things have been discussed in the field of traffic control. With these methods, traffic can be detected and adjusted for in real time. However, there are several issues which should be resolved in the future, such as bad weather conditions, different traffic, high computational requirements, and others.

Keywords: Artificial intelligence, traffic light control, computer vision, reinforcement learning, deep reinforcement learning, Internet of Things, YOLO.

Introduction

The intersections in urban areas have to cope with the traffic demand that varies depending on the time of day, the orientation of the road, weather conditions, incidents, special events, and adjacent land use. In fixed-time control, the green times are set in advance, which does not allow the system to react appropriately to an unexpected queue or underused lane. The outcome is excess waiting, increased time in motion, fuel waste, pollution, and emergency response delay.

Artificial intelligence traffic control has two related abilities. The first one is perception when the cameras, loop detectors, radars, connected vehicles, and IoT devices evaluate the number of vehicles, queue length, speed, occupancy rate, and pedestrians' need for crossing. The other one is the decision when optimization, machine learning, or reinforcement learning selects the appropriate signal phase and green times based on the estimated states. Object detectors like YOLO prove that visual detection is possible in real-time mode, while advanced algorithms enhance the trade-off between time and accuracy (Redmon et al., 2016; Bochkovskiy et al., 2020). RL methods also consider signal control as a decision-making task where the policy uses traffic state, action, and reward as inputs (Wei et al., 2018, 2019a, 2019b).

This work is a literature survey and does not introduce any newly developed traffic controller, data set, model, experiment, or numerical result. The goal of this paper is to discuss the key technology groups,

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

compare ten selected papers, and list the engineering challenges that have to be overcome for a new adaptive traffic controller to be safe and useful in practice.

Research Questions and Objectives

The review aims at addressing the following questions:

1. How do the AI-driven systems perceive the current traffic state?

2. What is the use of RL, DRL, federated learning, and language-model reasoning for signal action selection?

3. In what way does IoT, edge/cloud computing, and connected infrastructure help in traffic management?

4. Why are the limitations such that simulation-driven strategies cannot be directly applied to real intersections?

The objectives are to outline the traffic-control task, provide a practical classification of the approaches, compare the chosen research without creating an artificial set of performance metrics, differentiate perception and control, and point out future directions for adaptive urban traffic systems.

Intelligent Traffic-Control Problem

From Fixed Timing to Adaptive Control

The traffic signal controller chooses a phase and decides on how long this phase stays active. The fixed timing is easy and straightforward, but it operates under historical assumptions. Actuated control is more responsive, as it utilizes detectors, and adaptive control makes its decisions on updated timing basing on current observations. An AI controller takes this one step further by learning relations between the current traffic situation and the action necessary to achieve the desired objective.

Minimizing the delay would not necessarily guarantee pedestrians protection, giving priority to emergency vehicles, or avoiding starvation of a low traffic approach. Therefore, the reward and constraints should be of equal importance to the learning algorithm itself.

Traffic State and Sensors

A traffic state may be described based on measures like vehicle counts, lane occupancy, queue lengths, waiting times, speeds, phase status, turns and turning movements, or messages from connected vehicles. Cameras provide extensive coverage and semantic sensing; however, their operation may be impaired by occlusion, glare, rain, fog, night conditions, camera vibrations, or dense mixed traffic.

Perception Is Not Complete Optimization

Vehicle detection and counting may be necessary inputs but are not sufficient for traffic control. A comprehensive traffic system has to translate perception into a safe signal phase transition, conflict resolution, minimum green and clearance periods, protection of vulnerable road users, and missing or erroneous data handling.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

Method of Review

A focused narrative review was carried out using ten relevant articles from the literature made available during the seminar readings. These consist of articles on real-time object detection, RL for a single intersection, pressure-based control, cooperation through graphs, motion-centric DRL, federated DRL, survey of coordinated RL, and LLM-based controller. Since the articles use different datasets, road networks, simulators, sensors, rewards, and evaluation metrics, any attempt to generate an aggregate ranking score was not considered.

Figure 1 below illustrates the method used in the review.

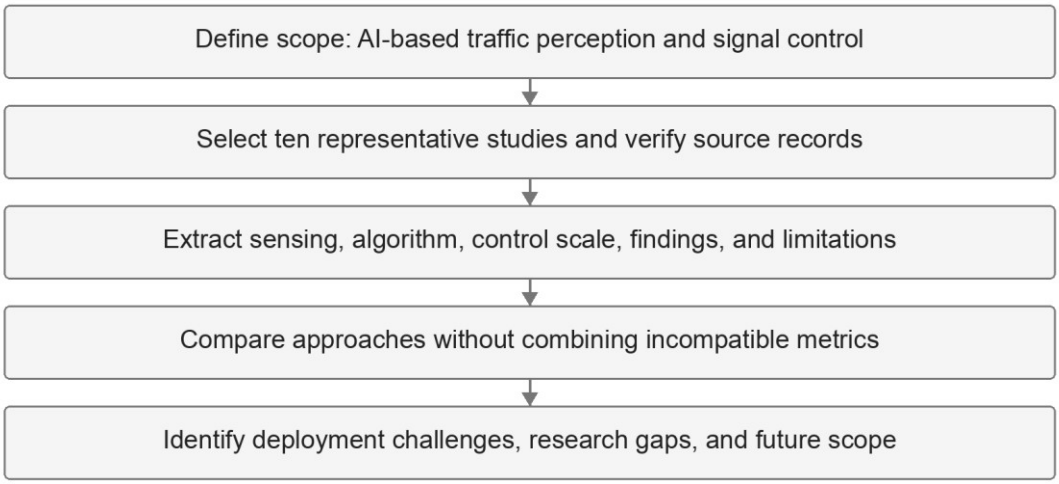


Figure 1: Focused literature-review workflow used in this paper.

Classification of Existing Methods

The existing literature could be classified into four intersecting categories, as shown in Figure 2 below.

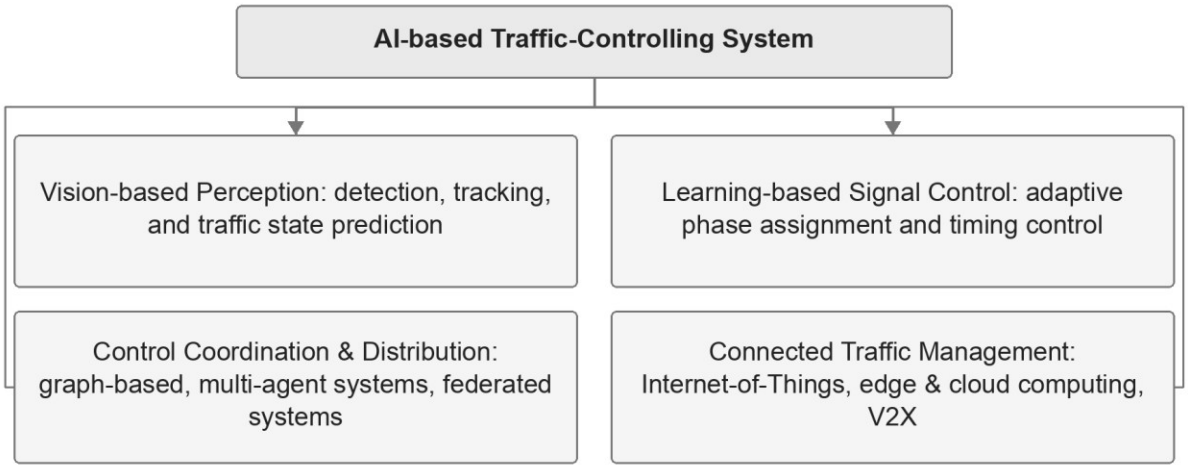


Figure 2: Classification of AI-based traffic-control approaches by function.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

AI & IoT Techniques

Computer Vision & Object Detection

YOLO detects objects using a single regression that calculates the bounding box coordinates along with class probabilities for an input image (Redmon et al., 2016). The latest model in this family, YOLOv4, seeks to create a better balance between detection accuracy and speed through architectural and training enhancements (Bochkovskiy et al., 2020). The models could be used in traffic surveillance applications such as vehicle counting, class estimation, lane occupancy, and queue estimation.

A detector trained on clean daylight images cannot be expected to achieve the same level of performance in low-light conditions, rainy weather, or dense Indian traffic scenarios where motorcycles, auto-rickshaws, buses, automobiles, pedestrians, and lane encroachments exist. Moreover, tracking is needed to estimate motion while avoiding counting the same vehicle again and again.

Reinforcement Learning and Deep Reinforcement Learning

RL treats traffic management as interactions between an agent and its environment where the agent sees a traffic state, selects an action like phases and durations, gets a reward, and updates its policy. IntelliLight provides a DRL framework for traffic-light management at an intersection (Wei et al., 2018). PressLight combines learning and max-pressure control through the use of pressure-based data to control traffic lights (Wei et al., 2019b). In a recent work involving DQN and DDPG methods, the emergence of a green wave is described in both single-intersection and networked traffic scenarios (“Deep Reinforcement Learning,” 2023).

A reward might optimize throughput without considering fairness, safety, or pedestrian management. Training usually involves extensive interactions; hence simulations are widely used. Transfer problems due to differences between simulated and real-life behavior of policies arise without proper restrictions, calibration, and validation.

Network Level vs. Movement Based Control

CoLight makes use of graph attention for representation of the process of communication and coordination between intersections (Wei et al., 2019a). MoveLight is focused on the movements and lane-based control and implements a movement-based DRL approach for various network sizes (Shao et al., 2024). It relies upon promptness of state information and on the representation of the real-world road network through the network graph, movement pattern, and communication.

Federated Learning and LLM-Based Control

In federated DRL, multiple traffic control environments can be used to train a shared algorithm without having all observations available in a single centralized data set. Li et al. (2025) suggest federated PPO for cross-domain traffic-signal control and discuss the problem of generalization of models between intersections. Heterogeneity in demand, hardware, state space, and communication capabilities can pose challenges to distributed learning.

In Traffic-R1, another approach is investigated where a 3 billion parameter language model gets reinforced with expert knowledge in a simulated traffic environment (Zou et al., 2025). Its reasoning-focused interface

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

might make the system easier to understand and adapt to unexpected situations, although computation costs, latency, verification, and safety concerns are still important issues.

IoT, Edge, Cloud, and V2X

There are cameras and sensors producing the data; edge devices that can process videos and do inference with low latency; cloud systems can aggregate historical information and do analytics; and V2X communications can send messages from the vehicles and the infrastructure. Being connected does not mean being smart: drift of sensors, dropped packets, cyber attacks, synchronization problems, and delayed commands may render an otherwise correct algorithm wrong.

Table 1: Functional roles in an AI traffic-control system

Technology	**Primary role**	**Important limitation**
Cameras and detectors	Estimate vehicle count, class, queue, and movements	Occlusion, weather, lighting, privacy, and calibration influence data quality.
RL/DRL controller	Choose adaptive phases or timings	Difficult to design rewards, ensure safety, and achieve simulation-to-reality transfer.
Graph/network/multi-agent approaches	Control neighboring intersections	Mobility becomes an issue due to communication and model complexities.
Federated Learning	Allows improvement of shared models without centralizing all raw data	Non-identical local data and unreliable connectivity could hinder training.
IoT, Edge, and Cloud	Data acquisition, processing, communication, storage, and visualization	Depend on security, latency, availability, and maintenance issues.
V2X Communication	Provide information for connected vehicles and infrastructure	Depends on the penetration rate, standards, and coverage.

Comparing Methods of Control

Table 2 offers a comparative overview that is purposely lacking an accuracy metric due to the differing outputs used by object detection, queue control, network cooperation, and surveys.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

Table 2: Qualitative comparison of major approaches

Approach	Strengths	Deployment considerations
Vision-based monitoring	High-quality, non-intrusive sensing and re-use of camera infrastructure	Takes detection, tracking, privacy safeguards and sensing backups into account.
Single-agent reinforcement learning	Learns a strategy in response to changing demand at one intersection	Could overfit the simulation environment or reward function.
Multi-agent / graph-based reinforcement learning	Model the cooperation across a road network	Complexity of training and communication increases with size.
Federated DRL	Enables learning across domains using decentralized raw data	Variety of sensors and distributions are possible at each site.
LLM control support	Allows for reasoning with expert help	Could have issues with computing, latency, determinism, verification, and safety.

Workflow for Conceptual System Architecture

The workflow for the conceptual system architecture, as illustrated in Figure 3, is based on the generalized understanding of the reviewed literature. This is not an architectural prototype proposed in this paper; it is a conceptual architecture.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

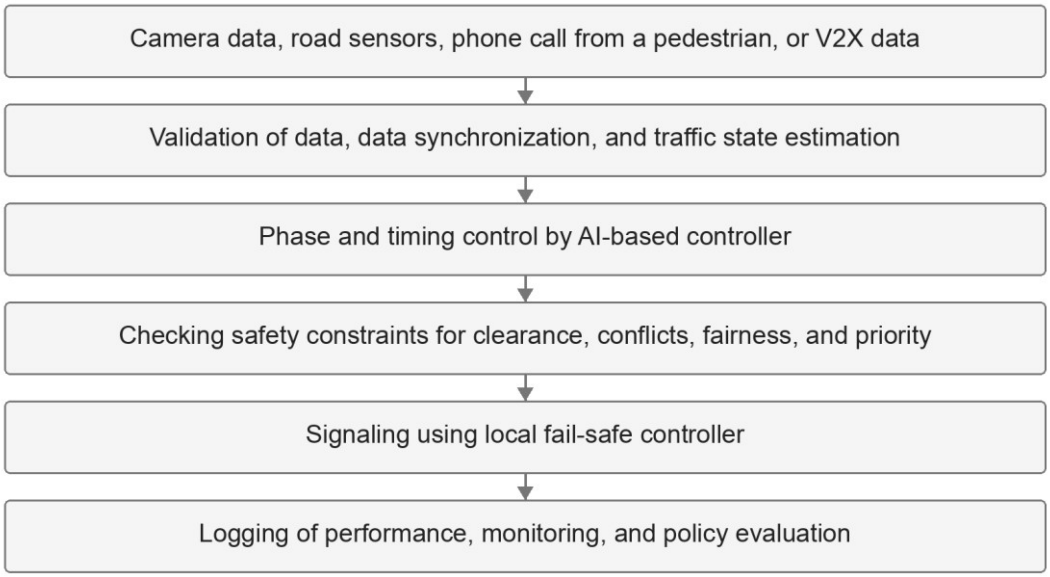


Figure 3: Conceptual workflow for AI-assisted traffic control with safety checks and local fail-safe operation.

Literature Survey

Table 3 summarizes the ten selected studies. The observations describe each publication's main contribution to this review; they do not imply that all works were tested on the same network or can be ranked directly.

Table 3: Comparison of ten representative studies

Study	Study emphasis	Method	Main contribution	Limitation relevant to deployment
Redmon et al. (2016)	Real-time object detection	YOLO single network detector	Develops fast end-to-end visual object detection	General detector but traffic counting and adverse condition testing need further development.
Bochkovskiy et al. (2020)	Balance between speed and accuracy in detection	YOLOv4	Improve real-time visual perception alternatives	Performance depends on computing platform, camera angles, weather conditions, and local training data.
Wei et al. (2018)	Intersection intelligent control	Deep reinforcement learning	Learn traffic light control through states and rewards	Assumptions for evaluation and reward design are not applicable to all real junctions.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

Study	Study emphasis	Method	Main contribution	Limitation relevant to deployment
Wei et al. (2019b)	Arterial coordination using pressure-based approach	RL with max-pressure concept	Couples the learned control with a traffic theory signal	Relies on the accuracy of lane and pressure estimates.
Wei et al. (2019a)	Coordinated control of multiple signals	Graph attention and multi-agent RL	Simulates communication between adjacent intersections	Implementation is influenced by network scale, communication requirements, and network topology changes.
Shao et al. (2024)	Movement-based adaptive control	Lane-level DRL based on movement information	Capable of representing turning movements at several network scales	Needs timely movement level traffic data.
"Deep Reinforcement Learning" (2023)	Adaptive control in single and grid networks	DQN and DDPG	Analyzes threshold policies and green-wave formation in simulations	The findings are simulation-based and dependent on environment configuration.
Li et al. (2025)	Cross-domain collaborative learning	Federated PPO/DRL	Collaborates policy learning in domains without centralized all observations	Multiplicity of heterogeneous locations and communication can affect implementation.
Saadi et al. (2025)	Coordination research and taxonomy	Survey on RL and DRL	Classifies techniques, issues, and directions in intelligent lighting control	More work is needed to adapt the survey results to practical engineering.
Zou et al. (2025)	Traffic signal control by reasoning	Reinforced 3B-parameter language model	Expert knowledge, reasoning, and traffic signal control learning	Computational, latency, verification, and real-time safety are still unresolved problems.

Comparative Discussion

Methods operating with one intersection are easier to explain and evaluate, though they can contribute to the further development of congestion. Network-level approaches imply intersection coordination explicitly, yet, they need more information and learning. Reward based on pressure provides transportation-

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

related sense to the process of reinforcement, while graph attention determines which intersections should cooperate in the communication. Flow control takes into account the flow of turns, not just the number of vehicles.

Federated learning and control using LLM are the areas pushing the boundaries of research. Despite the fact that federated learning could contribute to more cooperation and privacy, raw data decentralization cannot address problems like model poisoning, privacy leaks, or communication issues. Even though LLM-based decision-making may result in easier explanations of decisions, explanation accuracy is no substitute for safety testing in terms of determinism. In either case, the use of a traditional signal controller and fail-safe algorithm is essential.

A comparison of the given improvements expressed as percentages with the existing literature would be incorrect due to the difference in simulators, traffic demands, baseline methods, time horizons, rewards, and road configurations used by researchers.

Proposed Evaluation Framework

The methods described in this paper serve to solve different problems and thus are not comparable according to a single measure of accuracy. Below we provide evaluation criteria which are not claimed to be applied in this paper.

Precision and recall	Performance of traffic control
vehicle count estimation Queue estimates and tracking error different lighting conditions, weather, occlusion and types of traffic	average and high-percentile waiting time queue length, number of stops, throughput, and travel time queues further downstream
Safety and equity	**Latency and Resilience**
permissible phases, minimum green, and clearance intervals Pedestrian service, emergency priority, maximum waiting time, and worst served approaches	latency from sensing to decision-making delayed or missed measurements, communications failures fallback

Procedure for Comparison. For future simulations of an actuated/pressure model against the fixed-time model and the AI controller, the comparison can use the same geometry of the road, demand parameters, control signal policies, and available sensors' information. Multiple simulation runs with the recorded random seed and separate testing data would contribute to estimating the variance. The perception inaccuracies should be incorporated.

If any field test is done, it will have to be authorized by the traffic authority, run in shadow mode, monitored by humans, and with an approved backup plan.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

Deployment Challenges

Unmanned and Unorganized Traffic

It may contain motorcycles, buses, auto-rickshaws, bicycles, pedestrians, parked cars, and unconventional maneuvers to create occlusion and motions that are unlikely to appear in the benchmark data sets. Thus, it is necessary to have local training and calibration for each intersection.

Adverse Conditions

Weather conditions, including rain, fog, lighting, and any other condition that will reduce visibility, can reduce the efficiency of the model. Sensory fusion can help but it will increase the price of the system.

Transfer Learning

It is possible for an RL agent to take actions that are both unsafe and inefficient. This makes the deployment of the policy online at the intersection unacceptable.

Reward Design and Fairness

The worst case wait times, pedestrian crossings, accessibility, and emergency priority services should be modeled explicitly. The reporting should include the tail delay and worst-served population, rather than average metrics.

Scalability, Latency, and Reliability

In the context of a large-scale controller that needs to process many video streams and produce control signals in a timely fashion, centralized control may become a bottleneck; while in the context of distributed control, synchronization and consistency problems will arise. The solution should tolerate any failures of cameras, network connections, model, and cloud services.

Security, Privacy, and Governance

The video streams and the connected vehicle data may contain private data. Appropriate access controls, encryption, retention policies, logging, security patching, network segmentation are required. All changes to the models must be tracked and approved, and a procedure for rollback and investigation of incidents must be in place.

Research Gap

The key gap in research is the difference between potentially useful simulation experiments and actual traffic control systems. The current benchmarks do not take into account the mixture of traffic, changes in weather, accidents, failures of sensors, demand for pedestrians, priority of emergency vehicles, or model drifting.

The second gap is related to end-to-end evaluation. Just because the vehicles can be detected properly does not necessarily mean the formation of queue is adequate, and the decrease of delay achieved by simulation does not mean correct functioning of the camera-based detection system in the real world. Everything from detecting and estimating state to making actions and observing performance should be evaluated.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

Practical Implications

From the perspective of traffic management agencies, AI needs to be deployed as a supervised decision support system initially, prior to any control being granted directly to AI. The most suitable controller is likely the more simple controller that remains interpretable and stable rather than the one that gives the best results in simulation.

From a researcher's point of view, reproducibility calls for clearly stating all of the following parameters: detector readings, simulator version, road layout, traffic generation, signal phases, reward, seeds, baselines, safety constraints, and train-test split.

Future Scope

Any future research should consider the link between the quality of observation and the safety and feasibility of signal actions. Every type of sensing has to be compared with respect to the installation, calibration, maintenance cost and privacy issues.

Another way to proceed is through the constrained learning guided by the traffic engineering criteria. The proposed learning procedure can suggest some action, however, this action will be tested for conflicts, clearance times, pedestrian service, and waiting time limit by a separate safety system. The special handling of emergency vehicles and public transport should not rely on the assumption that this will result from the reward based solely on average delays.



In the long run, validation must address drift in the model, controller workload, maintenance costs, and restoration from failure.

Limitations of This Review

This paper is a literature review of ten selected papers, not a systematic review or a meta-analysis. Several works are preprints or simulation experiments, thus their findings cannot be treated as independent validation. No novel traffic dataset, no hardware implementation, no trained model and no experimental performance results have been developed in the course of this review.

Conclusion

In this paper, AI-based traffic controlling systems have been considered with special focus on computer vision, reinforcement learning, coordination control, federated learning, language-model reasoning and Internet-of-Things infrastructure. Nevertheless, accurate perception, good reward design, fairness, latency, cybersecurity, fail-safety and validation remain challenging tasks. The major direction of future research is developing an end-to-end validated traffic controlling system.

BODHIVRUKSHA JOURNAL OF DIVERSE DISCIPLINE
E-ISSN : 3139-1486
An Open Access, Peer-Reviewed, Multidisciplinary Scholarly Journal
Bi-monthly | Multilingual (English, Hindi, Marathi) | Academic Research
@ Volume 2, Issue 5 (September-October 2026) (Special Issue 1) ISSN: 3139-1486 (Online)

References

1. Bochkovskiy, A., Wang, C.-Y., & Liao, H.-Y. M. (2020). YOLOv4: Optimal speed and accuracy of object detection. arXiv preprint arXiv:2004.10934. https://arxiv.org/abs/2004.10934
2. Deep reinforcement learning for traffic light control in intelligent transportation systems. (2023; revised 2025). arXiv preprint arXiv:2302.03669. https://arxiv.org/abs/2302.03669
3. Li, M., Pan, X., Liu, C., et al. (2025). Federated deep reinforcement learning-based urban traffic signal optimal control. Scientific Reports, 15, Article 11724. https://doi.org/10.1038/s41598-025-91966-1
4. Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). You only look once: Unified, real-time object detection. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (pp. 779–788). https://doi.org/10.1109/CVPR.2016.91
5. Saadi, A., Abghour, N., Chiba, Z., et al. (2025). A survey of reinforcement and deep reinforcement learning for coordination in intelligent traffic light control. Journal of Big Data, 12, Article 84. https://doi.org/10.1186/s40537-025-01104-x
6. Shao, J., Zheng, C., Chen, Y., Huang, Y., & Zhang, R. (2024). MoveLight: Enhancing traffic signal control through movement-centric deep reinforcement learning. arXiv preprint arXiv:2407.17303. https://arxiv.org/abs/2407.17303
7. Wei, H., Zheng, G., Yao, H., & Li, Z. (2018). IntelliLight: A reinforcement learning approach for intelligent traffic light control. In Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 2496–2505). https://doi.org/10.1145/3219819.3220096
8. Wei, H., et al. (2019a). CoLight: Learning network-level cooperation for traffic signal control. In Proceedings of the 28th ACM International Conference on Information and Knowledge Management (pp. 1913–1922). https://doi.org/10.1145/3357384.3357902
9. Wei, H., et al. (2019b). PressLight: Learning max pressure control to coordinate traffic signals in arterial network. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 1290–1298). https://doi.org/10.1145/3292500.3330949
10. Zou, X., et al. (2025). Traffic-R1: Reinforced LLMs bring human-like reasoning to traffic signal control systems. arXiv preprint arXiv:2508.02344. Accepted to ACL 2026. https://arxiv.org/abs/2508.02344