

Original Research Article

AI-based dynamic optimization of smart traffic signals for improved mobility efficiency

Jiangeng Long^{1,2}, M. Kazem Chamran^{1*}¹ Faculty of Information Technology, City University Malaysia, Petaling Jaya, Selangor Darul Ehsan, 46100, Malaysia² Jiangxi Vocational & Technical College of Information Application, Nanchang, Jiangxi, 330043, China

Abstract: Artificial intelligence (AI) can improve urban mobility by adapting traffic signals to changing demand, incidents, and multimodal priorities. We propose a fieldable framework combining real-time state estimation, short-horizon prediction, and constraint-aware reinforcement learning. The policy suggests phase, split, and offset updates, while a safety and fairness layer enforces minimum green, clearance, pedestrian service, and side-street protection. Neighborhood messaging enables corridor coordination, and deployment uses monitoring and rollback. Illustrative studies show reduced delay and queues with higher throughput.

Keywords: artificial intelligence; intelligent transportation systems; adaptive signal control; reinforcement learning; multi-agent control; traffic efficiency

1. Introduction

Urban intersections are bottlenecks where small timing errors amplify delay, emissions, and unreliability. Recent reviews show that fixed-time, actuated, and adaptive signal plans remain widely used in practice, yet their performance can deteriorate when demand becomes highly nonstationary and corridor interactions intensify^[1,2]. This motivates data-driven control strategies that can react to changing arrivals, incidents, and spatial spillback more effectively.

Recent reinforcement-learning studies report encouraging results for signal optimization, including benchmarking against commercial systems^[3], arterial coordination with multi-agent attention-based soft actor-critic control^[4], hierarchical graph-based multi-agent learning^[5], fairness-aware visual perception and control^[6], and multi-objective optimization that jointly considers efficiency, safety, and emissions^[7]. These advances support the use of AI-driven signal control while also highlighting the need for deployable constraints, monitoring, and fallback mechanisms.

2. Methodology

This paper presents an implementable architecture for AI-enabled dynamic signal optimization. Rather than replacing engineering practice, it links learning-based decisions to the constraints, interfaces, and evaluation checks needed for field deployment. We focus on measurable mobility indicators, explicit safety/fairness constraints, and continuous monitoring for rollback and auditability.

2.1. Perception and state construction

Smart intersections use loops/radar, video analytics, and sometimes connected-vehicle messages. Because sources differ in rate and noise, a lightweight fusion layer time-aligns data and outputs per-approach features: arrivals, queue estimates, saturation-flow proxies, and spillback flags from stop-line occupancy.

To mitigate partial observability, we add pressure-like features comparing inbound/outbound congestion and short-horizon (30–90 s) arrival predictions for upstream platoons. Prediction can be simple smoothing, or learned sequence models when trajectories are available.

2.2. Control policy and operational constraints

At each decision interval (e.g., 5–10 s), the policy proposes the next phase, incremental green-split updates,

and optional offset nudges. Any learning method is acceptable, but outputs must map to deployable controller commands (phase/split/offset) rather than arbitrary continuous signals.

A constraint layer enforces operational rules. Hard limits cover min/max green, clearance, pedestrian walk/FDW, and spillback protection; soft limits bound split changes and caps on repeated deferral. Infeasible actions are projected to the nearest feasible plan so safety is never overridden.

2.3. Training and validation

Policies are trained offline in microscopic simulation with randomized demand, incidents, and sensor noise to improve robustness. Calibration prioritizes queue formation and discharge dynamics, which largely determine delay and spillback.

Validation should cover average and stress cases (oversaturation, blockages, detector dropouts). Beyond mean travel time, report tail metrics (e.g., 90th-percentile delay) and worst-approach queues to avoid hiding extreme failures.

2.4. Field integration and governance

Deployment starts in shadow mode, comparing AI recommendations to the incumbent plan. Activation is time-boxed and guarded (e.g., delay cannot exceed baseline by a margin for several intervals). Governance records objective weights, model versions, and approval roles, enabling accountable updates and rollback.

2.5. Real-time implementation considerations

Real-time control must satisfy tight latency budgets; sensing-to-command delay should remain below the decision interval. Edge inference limits network dependence, while learning remains offline and auditable for simpler certification.

3. AI-enabled signal optimization loop

Figure 1 outlines the closed-loop pipeline from sensing and prediction to decision, constraint projection, and monitoring. Learning proposes actions; operations and safety define non-negotiable boundaries.

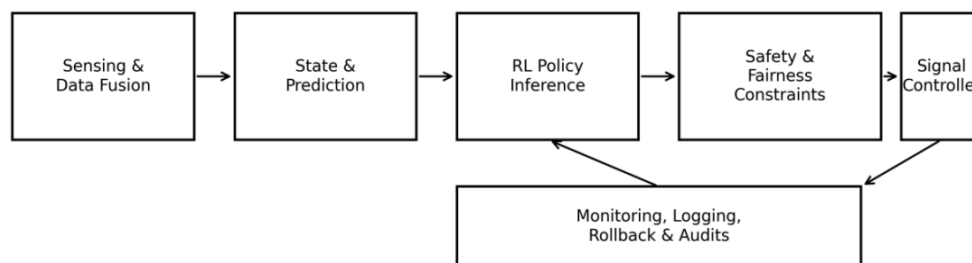


Figure 1. Closed-loop architecture for AI-enabled dynamic traffic signal control.

3.1. Objective and reward design

The controller optimizes a weighted objective for mobility and stability. A typical reward is $r_t = -(w_d * D_t + w_q * Q_t + w_s * S_t) + w_f * F_t$, where D_t is delay, Q_t is queue/spillback risk, S_t is stop rate, and F_t rewards served flow. Weights reflect policy goals (e.g., peak delay reduction, pedestrian regularity) and should be reviewed with stakeholders.

3.2. Coordination across intersections

Coordination uses neighborhood messaging: intersections share queues, predicted arrivals, and phase intent. Simple aggregation or graph attention helps align responses to platoons and spillback; under packet loss, a stable local policy degrades gracefully.

3.3. Monitoring and fallback operation

A supervisory layer translates decisions into standard phase/split commands compatible with common controllers. A baseline plan remains available for automatic fallback. Dashboards track mean and tail delay, queue occupancy, pedestrian compliance, and faults; logs support incident review and controlled retraining.

4. Discussion

This section interprets the framework in operational terms and illustrates how agencies can evaluate benefits and risks with compact indicators.

4.1. Evidence and evaluation indicators

Agencies need a small indicator set capturing mobility and operational stability. Table I provides illustrative pre/post peak-hour metrics, and **Figure 2** shows relative improvements. All measures can be derived from routine detector and controller logs.

Combine intersection-level metrics (approach delay, maximum queue, split violations) with corridor outcomes (travel time reliability, progression). Include equity views by comparing benefits across approaches and time periods so side streets are not systematically degraded.

Table 1. Illustrative corridor-level performance indicators (peak hour).

Metric	Baseline	AI-Dynamic	Relative change
Average delay (s/veh)	68.0	42.0	-38.2%
Mean queue (veh)	14.5	9.1	-37.2%
Throughput (veh/h)	1650	1830	+10.9%
Stop rate (stops/veh)	1.85	1.32	-28.6%




Figure 2. Relative mobility improvements of AI-dynamic control vs. baseline (illustrative).

4.2. Operationalizing risks as design constraints

Operational risks should be encoded as constraints. Safety (clearance, pedestrian service, spillback) is enforced as hard limits; fairness can bound maximum red or penalize repeated deferral. Robustness requires fault detection and bounded actions under abnormal demand; cybersecurity needs authenticated channels and auditable updates.

5. Conclusion

We presented a practical framework for AI-based dynamic traffic signal optimization linking sensing, prediction, constraint-aware learning, and monitoring to measurable mobility outcomes. By pairing learned policies with hard operational constraints and corridor coordination, agencies can reduce delay and queues without sacrificing safety or interpretability. Future work includes richer multimodal objectives (transit/bikes), better transfer across cities, and using connected-vehicle trajectories for improved prediction and online validation. For deployment, the architecture assumes detector/controller logs and a supervisory interface, which lowers integration cost. The explicit constraint projection and fallback plan provide predictable behavior during sensor faults or unusual demand, supporting agency acceptance and routine operations, and makes post-hoc audits straightforward for safety review processes.

About the author

*Corresponding author: M. Kazem Chamran.

References

- [1] P. Michailidis, I. Michailidis, C. R. Lazaridis, and E. Kosmatopoulos, "Traffic Signal Control via Reinforcement Learning: A Review on Applications and Innovations," *Infrastructures*, vol. 10, no. 5, Art. no. 114, 2025, doi: 10.3390/infrastructures10050114.
- [2] M. Noaen, A. Naik, L. Goodman, J. Crebo, T. Abrar, Z. Shakeri Hossein Abad, A. L. C. Bazzan, and B. Far, "Reinforcement learning in urban network traffic signal control: A systematic literature review," *Expert Systems with Applications*, vol. 199, Art. no. 116830, 2022, doi: 10.1016/j.eswa.2022.116830.
- [3] A. Cabrejas-Egea, R. Zhang, and N. Walton, "Reinforcement Learning for Traffic Signal Control: Comparison with Commercial Systems," *Transportation Research Procedia*, vol. 58, pp. 638-645, 2021, doi: 10.1016/j.trpro.2021.11.084.
- [4] F. Mao, Z. Li, Y. Lin, and L. Li, "Mastering Arterial Traffic Signal Control With Multi-Agent Attention-Based Soft Actor-Critic Model," *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, no. 3, pp. 3129-3144, 2023, doi: 10.1109/TITS.2022.3229477.
- [5] S. Yang, "Hierarchical graph multi-agent reinforcement learning for traffic signal control," *Information Sciences*, vol. 634, pp. 55-72, 2023, doi: 10.1016/j.ins.2023.03.087.
- [6] W. Fang, X. Zhao, and C. Zhang, "Fairness-aware multi-agent reinforcement learning and visual perception for adaptive traffic signal control," *Optoelectronics Letters*, vol. 20, pp. 764-768, 2024, doi: 10.1007/s11801-024-3267-2.
- [7] M. Elharoun, S. M. El-Badawy, E. A.-E. Shwaly, and U. Elrawy Shahdah, "Adaptive traffic signal control using deep reinforcement learning: A multi-objective approach for single and multi-intersection scenarios," *IATSS Research*, vol. 49, no. 4, pp. 481-492, 2025, doi: 10.1016/j.iatssr.2025.10.004.