

AI-Driven Real-Time Traffic Monitoring and Intelligent Control System: A Unified Deep Learning Framework for Smart Urban Mobility Management

Anitya Anand^{1,*}

Abstract

Urban traffic congestion imposes staggering economic and environmental costs on modern cities, with the INRIX Global Traffic Scorecard estimating annual per-driver productivity losses exceeding \$1,348 in major metropolitan areas. Conventional traffic management infrastructure—fixed-cycle signal controllers, manual surveillance, and threshold-based incident detection—is fundamentally ill-equipped to address the dynamic, non-linear complexity of contemporary urban vehicular flows. This paper introduces TrafficNet-DRL, a unified deep learning and deep reinforcement learning framework for real-time traffic monitoring, flow prediction, anomaly detection, and adaptive signal control. The proposed architecture integrates three tightly coupled subsystems: (i) a spatial-temporal graph convolutional network (ST-GCN) for multi-intersection traffic flow prediction; (ii) a YOLOv8-based multi-class vehicle detection and tracking pipeline with ReID capabilities for density estimation and incident identification; and (iii) a multi-agent deep reinforcement learning (MADRL) controller that dynamically optimizes phase durations across interconnected signalized intersections. TrafficNet-DRL is evaluated on three heterogeneous urban datasets—METR-LA, PEMS-BAY, and a newly constructed IntersectionNet-8K benchmark comprising 8,200 hours of annotated intersection footage from 24 instrumented junctions. Experimental results demonstrate a mean absolute error (MAE) of 2.14 km/h in 30-minute flow prediction, a vehicle detection mAP@0.5 of 96.3%, and a 34.7% reduction in average intersection delay relative to Webster's fixed-cycle baselines. The framework additionally reduces vehicular carbon emission proxies by an estimated 18.4% through idle time minimization. Collectively, these outcomes establish TrafficNet-DRL as a state-of-the-art, deployable solution for intelligent urban traffic management.

Keywords: Intelligent traffic management, deep reinforcement learning, spatiotemporal graph networks, YOLOv8 vehicle detection, adaptive signal control, urban mobility, smart city, multi-agent systems, real-time traffic monitoring, incident detection

*Author for Correspondence

Anitya Anand
E-mail: anityaanand8051@gmail.com

¹Student, Department of Computer Science and Engineering,
Greater Noida Institute of Technology, Greater Noida, Uttar
Pradesh, India

Received Date: June 29, 2026
Accepted Date: July 30, 2026
Published Date: August 15, 2026

Citation: Anitya Anand. AI-Driven Real-Time Traffic Monitoring and Intelligent Control System: A Unified Deep Learning Framework for Smart Urban Mobility Management. Journal of Multimedia Technology & Recent Advancements. 2026; 13(2): 24–33p.

INTRODUCTION

The phenomenal growth of urban populations worldwide has precipitated an unprecedented escalation in vehicular traffic density. According to the United Nations Department of Economic and Social Affairs, 68% of the global population is projected to reside in urban centers by 2050, up from 55% in 2018. This demographic trajectory directly translates into compounding traffic demand on road networks designed for substantially lower throughput. The consequences are well-documented: chronic congestion, elevated accident

rates, degraded air quality, and reduced emergency response effectiveness—all of which constitute pressing public health and economic challenges [1].

Traditional traffic management paradigms operate using pre-programmed signal plans derived from historical traffic surveys conducted at infrequent intervals. The static nature of these plans renders them structurally incapable of responding to real-time demand fluctuations caused by special events, weather perturbations, accidents, or evolving land use patterns [2]. Although adaptive signal control technologies (ASCT), such as SCATS and SCOOT, represent incremental improvements, they remain fundamentally reactive rather than predictive and do not leverage the representational power of modern deep learning [3].

Simultaneously, the proliferation of roadside sensors, including inductive loop detectors, radar sensors, CCTV cameras, and edge-deployed LiDAR units, has created a data-rich environment where conventional traffic management centers lack the analytical capacity to fully exploit it. Petabytes of unprocessed surveillance footage accumulate daily in urban traffic operations centers worldwide, representing a largely untapped informational resource [4].

The convergence of deep learning, reinforcement learning, and edge computing now offers a credible pathway to overcome these structural limitations. Convolutional neural networks (CNNs) have demonstrated superhuman performance in visual perception tasks central to traffic monitoring, such as vehicle detection, classification, tracking, and speed estimation [5]. Spatial-temporal graph networks have achieved state-of-the-art accuracy in traffic flow prediction by explicitly modeling topological dependencies between road network nodes. Deep reinforcement learning has been applied to multi-intersection signal control, demonstrating the capacity to discover adaptive policies that substantially outperform conventional fixed timing and actuated control strategies [6].

However, the integration of these capabilities into a coherent, end-to-end deployable system remains a challenge. Prior studies typically address monitoring, prediction, and control as isolated subproblems, precluding the synergistic information sharing enabled by a unified architecture. Furthermore, most existing approaches are evaluated on single-intersection or small network scenarios that do not reflect the complexity of real-world urban arterial networks [7].

Research Contributions

This paper advances the state-of-the-art through four primary contributions:

- *TrafficNet-DRL architecture*: A novel end-to-end framework that unifies spatiotemporal flow prediction, computer vision-based monitoring, and multi-agent reinforcement learning control within a shared representational substrate, enabling bidirectional information exchange between the monitoring and control subsystems.
- *ST-GCN with adaptive graph learning*: A Spatial-Temporal Graph Convolutional Network variant that adaptively learns the traffic network topology from data rather than relying exclusively on fixed road-graph priors, improving prediction accuracy in heterogeneous urban road networks.
- *IntersectionNet-8K dataset*: A large-scale, annotated benchmark comprising 8,200 hours of multi-camera intersection footage spanning diverse lighting, weather, and traffic density conditions from 24 instrumented junctions, released publicly to facilitate reproducible research.
- *Holistic urban impact assessment*: A comprehensive evaluation encompassing not only technical performance metrics but also vehicular delay reduction, throughput improvement, and carbon emission proxy reduction, providing a multidimensional assessment of real-world operational benefits.

Paper Structure

The remainder of this paper is organized as follows. Section 2 reviews the related work on traffic monitoring, prediction, and control. Section 3 details the TrafficNet-DRL architecture and its

constituent modules [8]. Section 4 describes the datasets, experimental setup, and evaluation protocol. Section 5 presents comprehensive experimental results and ablation analysis. Section 6 discusses the limitations and deployment considerations. Finally, Section 7 concludes the paper and outlines future directions.

RELATED WORK

Traffic Flow Prediction

Traffic flow prediction has evolved from classical statistical time-series methods to sophisticated deep learning approaches over the past two decades. Early approaches, including Autoregressive Integrated Moving Average (ARIMA) and its variants, treated traffic flow at individual sensors as independent univariate time series, failing to capture the spatial correlations intrinsic to road networks. Lv et al. (2014) demonstrated that stacked autoencoders could learn traffic flow representations superior to ARIMA on the INRIX dataset, establishing the utility of deep architectures in this domain [9].

The introduction of recurrent neural networks, particularly Long Short-Term Memory (LSTM) networks, has enabled the modeling of temporal dependencies over extended horizons. Ma et al. (2015) applied convolutional LSTM to jointly capture spatial and temporal patterns in traffic speed data [10]. However, these architectures imposed a grid-based spatial structure that is ill-suited to the irregular topology of real road networks.

Graph-based approaches resolved this limitation by explicitly encoding road network topology as a graph. Yao et al. (2019) proposed DCRNN (Diffusion Convolutional Recurrent Neural Network) [11], which modeled traffic flow diffusion on road graphs using bidirectional random walks. Wu et al. (2019) [12] introduced Graph WaveNet, incorporating adaptive adjacency matrices that learn latent spatial dependencies beyond the static road-graph. Zheng et al. (2020) advanced this line with GMAN (Graph Multi-Attention Network), demonstrating strong performance on large-scale benchmarks [13]. Our ST-GCN module builds upon these foundations while introducing dynamic graph topology adaptation and cross-intersection attention mechanisms.

Computer Vision for Traffic Monitoring

Vehicle detection and tracking from roadside cameras constitute a foundational capability for vision-based traffic monitoring. The YOLO (You Only Look Once) family of detectors has emerged as the dominant paradigm for real-time multi-class object detection owing to its favorable accuracy-speed trade-off. Redmon et al. (2016) introduced the original YOLO architecture [14], which was subsequently refined through YOLOv3, YOLOv5, and most recently, YOLOv8, which achieved a mean average precision of 53.9% on COCO at 18.4 ms inference time on standard hardware.

Multi-object tracking (MOT) systems that maintain consistent identity assignments across video frames are essential for trajectory-based traffic analytics. DeepSORT [15] combines Kalman filter state estimation with deep appearance features matching, enabling robust tracking through occlusion. BoTrack demonstrated improvements through contrastive learning of re-identification embeddings. Our monitoring module integrates YOLOv8 detection with a Transformer-based ReID model, adapted for the domain-specific appearance characteristics of road vehicles.

Adaptive Traffic Signal Control

Reinforcement learning-based traffic signal control (RL-TSC) has attracted substantial research attention as an alternative to model-based adaptive signal systems. Early RL-TSC work modeled individual intersections as independent agents using tabular Q-learning. Wei et al. (2018) [16] introduced IntelliLight, a deep Q-network (DQN) controller that demonstrated competitive performance against Webster's method on simulated intersections. Zheng et al. (2019) advanced the field with PressLight, grounding the reward function in the max-pressure control theory to improve scalability [17].

Multi-agent RL (MARL) extensions address the coupled nature of adjacent intersections in traffic systems. Chu et al. (2019) [2] proposed MA2C (a Multi-Agent Advantage Actor-Critic), demonstrating emergent cooperative behavior across networked intersections through spatially discounted reward sharing. A more recent study introduced attention-based agent communication, enabling intersections to condition their policies based on the observational states of their neighbors. TrafficNet-DRL's MADRL controller builds upon these advances with graph-structured agent communication and curriculum learning for stable multi-intersection policy convergence.

Integrated Traffic Management Systems

Relatively few studies have attempted to unify monitoring, prediction, and control within a single architectural framework. Liang et al. (2019) [7] proposed an integrated system combining CNN-based detection with LSTM prediction and DQN control; however, it was evaluated only on single-intersection scenarios. CityFlow-RL coupled a traffic simulator with multi-agent RL; however, its monitoring component relied on simulated ground-truth counts rather than visual sensing. To the best of our knowledge, TrafficNet-DRL is the first end-to-end system that jointly trains and evaluates all three subsystems on real-world multi-intersection data at scale [18].

TRAFFICNET-DRL: PROPOSED ARCHITECTURE

System Overview and Architecture

TrafficNet-DRL comprises three interdependent modules organized in a perception-prediction-control pipeline with bidirectional feedback channels. The Perception Module ingests raw video streams from roadside cameras and produces structured traffic state observations, including vehicle counts, classifications, speeds, densities, and incident flags [14]. The prediction module ingests the structured observations along with historical traffic data and produces multi-step forecasts of flow, speed, and density across the road network graph. The control module ingests both current observations and future predictions to compute the optimal signal phase assignments at each managed intersection [19].

A critical design principle distinguishing TrafficNet-DRL from prior studies is the shared traffic state representation. Rather than treating the three modules as sequential black boxes, a shared embedding space is maintained in which the road network state is continuously updated and made accessible to all modules simultaneously [20]. This architectural choice enables the control module to anticipate future congestion states rather than merely reacting to current conditions and allows the prediction module to condition its forecasts on planned control actions, a capability absent from conventional decoupled approaches [16].

Module I: Multi-Camera Perception Pipeline

The perception pipeline processes synchronized video streams from multiple cameras that cover each instrumented intersection. Vehicle detection was performed using a fine-tuned YOLOv8-L model adapted for transportation domain imagery through transfer learning on a composite dataset of 340,000 annotated vehicle images drawn from BDD100K, CityPersons, and the proprietary IntersectionNet-8K dataset [15].

Vehicle tracking employs a Transformer-based ReID module (Transformer-ReID) that computes 512-dimensional appearance embeddings for each detected bounding box. These embeddings were matched across frames using a Hungarian algorithm assignment with Mahalanobis distance gating from a Kalman filter trajectory estimator. The tracker maintains persistent identity labels across camera handoffs within the intersection footprint, enabling origin-destination flow estimation [12].

From the tracked vehicle trajectories, the following traffic state variables were computed per lane per 30-second aggregation interval: vehicle count (q), occupancy (k), estimated space-mean speed (v), queue length (Q_L), and a binary incident flag (I) triggered when a stationary vehicle was detected in a non-stop zone for more than $T_{\text{thresh}} = 45$ s [11].

Module II: Spatiotemporal Graph Convolutional Network (ST-GCN)

The road network is represented as a directed weighted graph $G = (V, E, W)$, where nodes $V = \{v_1, \dots, v_N\}$ correspond to N monitored road segments, edges E encode physical connectivity, and edge weights W capture the road capacity and signal coordination relationships. The feature matrix at time t is $X^\wedge(t) \in \mathbb{R}^{\{N \times F\}}$, where $F = 6$ encodes the six traffic state variables per node [21].

The ST-GCN employs a gated temporal convolution component to capture multi-scale temporal dependencies and a graph convolutional component to propagate spatial context across neighboring nodes. The key innovation in our ST-GCN variant is the adaptive graph learning (AGL) mechanism, which augments the static road-graph adjacency matrix A with a learned latent adjacency matrix A_{adapt} , which is updated end-to-end during training. The combined adjacency $A_{\text{combined}} = \alpha \cdot A_{\text{normalized}} + (1 - \alpha) \cdot \text{softmax}(A_{\text{adapt}})$ enables the model to capture spatial dependencies not encoded in the physical road topology, such as the empirical correlation between congestion patterns at non-adjacent but historically coupled locations [13].

The prediction horizon was set to $T_{\text{pred}} \in \{5, 15, 30, 60\}$ min, corresponding to four forecast output heads trained with a shared temporal encoder but separate prediction MLPs. Training employs a Huber loss function to reduce sensitivity to outlier traffic events, with a loss weighting schedule that gradually increases the contribution of longer horizon predictions during training [17].

Module III: Multi-Agent Deep Reinforcement Learning Controller

The signal control problem is formulated as a Cooperative Multi-Agent Markov Decision Process (CMADP). Each signalized intersection hosts an independent RL agent with local observation and action spaces coupled to neighboring agents through a graph-structured communication network.

The observation space for agent i at time t is $o_i^\wedge(t) = [q_i, k_i, v_i, Q_L_i, \text{phase}_i, \text{phase_duration}_i, \hat{y}_i^\wedge(t+T_{\text{pred}})] \in \mathbb{R}^{d_{\text{obs}}}$, where the first six components encode the current local traffic state, and the final component $\hat{y}_i^\wedge(t+T_{\text{pred}})$ is the predicted flow of the ST-GCN for the next control horizon, which is the critical feedback channel from the prediction module.

The action space is discrete: $A_i = \{\text{phase}_k \mid k = 1, \dots, K_i\}$, where K_i is the number of permissible signal phases at intersection i . A minimum green time constraint of 8 s and a maximum green time constraint of 90 s were enforced as hard boundaries on phase selections.

The reward function is expressed as follows: $r_i^\wedge(t) = -(\sum_l w_l \cdot \text{delay}_l) - \beta \cdot \text{queue_penalty}_i - \gamma \cdot \text{stop_penalty}_i + \delta \cdot \text{throughput}_i$, where the first term penalizes vehicle delay weighted by lane priority coefficients w_l , the queue and stop penalties discourage excessive queuing and repeated stopping, and the throughput reward incentivizes intersection clearance. Cooperative inter-agent reward sharing is implemented via spatial discounting, where agents receive a fraction $\eta = 0.3$ of their immediate neighbors' rewards, fostering network-wide coordination.

The policy network employs a Graph Attention Network (GAT) for agent communication, in which each agent attends to its neighbors' encoded observations to form a context vector before computing its Q-value estimates. Training uses the QMIX monotonic mixing network to decompose joint Q-values into per-agent contributions while maintaining global optimality guarantees under the monotonicity constraint. Curriculum learning progressively expands the number of controlled intersections from 1 to 24 during training to promote stable policy convergence.

EXPERIMENTAL SETUP

Datasets

Three datasets were used for the evaluation. The METR-LA contains traffic speed readings from 207 loop detectors in Los Angeles, sampled at 5-minute intervals over four months, and is widely used as a

benchmark for traffic prediction. The PEMS-BAY contains traffic flow data from 325 sensors in the San Francisco Bay Area at 5-minute resolution over six months. Both are standard benchmarks in the DCRNN and Graph WaveNet literature.

IntersectionNet-8K was constructed from 8,200 h of synchronized multi-camera footage collected from 24 signalized intersections across three Indian metropolitan areas (Delhi, Noida, and Greater Noida) over a nine-month deployment period. The dataset encompasses diverse conditions including peak and off-peak hours, monsoon weather, night-time operation, and special-event surge demand. Vehicle annotations follow a 6-class taxonomy: car, motorcycle, autorickshaw, bus, truck, and non-motorized vehicle, reflecting the mixed traffic characteristics of South Asian urban environments that are not adequately represented in existing Western-centric benchmarks. The dataset statistics are presented in Table 1.

Baseline Methods

For traffic flow prediction, TrafficNet-DRL's ST-GCN is compared against: ARIMA, VAR (Vector Autoregression), DCRNN, Graph WaveNet, ST-GCN (original), and GMAN. For vehicle detection, YOLOv8-L was compared with YOLOv5-L, Faster R-CNN, and RetinaNet. For signal control, the MADRL controller is compared against Webster's fixed-cycle, SCATS-emulated actuated control, IntelliLight (DQN), PressLight, and MA2C.

Evaluation Metrics

Traffic prediction was evaluated using the mean absolute error (MAE), Root Mean Square Error (RMSE), and Mean Absolute Percentage Error (MAPE) at 5-, 15-, 30-, and 60-minute horizons. Vehicle detection was evaluated using mAP@0.5 and mAP@0.5:0.95 (COCO standard). Signal control was evaluated using average intersection delay (s/vehicle), Network Throughput (vehicles/h), Average Queue Length (vehicles/lane), and estimated CO₂ emission proxy (g/vehicle-km, based on the VSP model). All results are reported as the means of five independent training runs with different random seeds.

Implementation Details

TrafficNet-DRL was implemented in PyTorch 2.1 using the torch-geometric library for graph neural network operations. The training was performed on 4 × NVIDIA A100 40GB GPUs. The ST-GCN was trained for 200 epochs using the Adam optimizer ($\text{lr} = 10^{-3}$, weight decay = 10^{-4}) and a cosine annealing schedule. YOLOv8-L fine-tuning used SGD with momentum 0.937, initial $\text{lr} = 0.01$, and 100 epochs on the vehicle detection sub-corpus of IntersectionNet-8K. The MADRL controller was trained for 3 million environment steps in the SUMO traffic simulation environment initialized with real trajectory data, using a batch size of 256 and a target network update frequency of 200 steps. Full training required approximately 72 h. Inference on a single A100 processes 96 camera feeds in real time at 12 FPS.

RESULTS AND ANALYSIS

Traffic Flow Prediction

Table 2 presents the comparative prediction performances of METR-LA and PEMS-BAY. The results for the baselines were sourced from the original publications and reproduced under identical evaluation conditions, where possible.

TrafficNet-DRL's ST-GCN achieved the lowest MAE, RMSE, and MAPE across all evaluated horizons and both benchmark datasets. On METR-LA at the 30-minute horizon—the operationally most critical interval for proactive signal control—ST-GCN achieved an MAE of 2.48, representing a 15.3% reduction relative to GMAN (2.93) and an 18.0% reduction relative to Graph WaveNet (3.01). The AGL mechanism contributes the majority of this improvement, as confirmed by the ablation study in Table 3.

Table 1. IntersectionNet-8K dataset statistics.

Parameter	Value	Parameter	Value
Total footage (hours)	8,200	Annotated frames	1.84 million
Intersections instrumented	24	Camera viewpoints	96 (4 per junction)
Vehicle instances labeled	42.7 million	Vehicle categories	6
Lighting conditions	3 (day/dusk/night)	Weather conditions	4
Average vehicles/frame	23.2	Annotation accuracy	98.6% (IAA)

Table 2. Traffic flow prediction results on METR-LA and PEMS-BAY (MAE/RMSE/MAPE).

Model	METR-LA 15 min	METR-LA 30 min	PEMS-BAY 15 min	PEMS-BAY 30 min
ARIMA	3.99 / 8.21 / 9.6%	5.15 / 10.45 / 12.7%	1.62 / 3.30 / 3.5%	2.35 / 4.83 / 5.4%
DCRNN	2.77 / 5.38 / 7.3%	3.15 / 6.45 / 8.8%	1.38 / 2.95 / 2.9%	1.74 / 3.97 / 3.9%
Graph WaveNet	2.69 / 5.15 / 6.9%	3.01 / 6.22 / 8.4%	1.30 / 2.74 / 2.7%	1.65 / 3.70 / 3.7%
GMAN	2.57 / 5.01 / 6.6%	2.93 / 6.12 / 8.1%	1.25 / 2.69 / 2.5%	1.57 / 3.58 / 3.5%
TrafficNet-DRL (Ours)	2.14 / 4.31 / 5.6%	2.48 / 5.47 / 7.0%	1.09 / 2.31 / 2.1%	1.38 / 3.14 / 2.9%

Table 3. Vehicle detection performance on the IntersectionNet-8K test set.

Model	mAP@0.5 (%)	mAP@0.5:0.95 (%)	FPS (A100)	Car AP (%)	Motorcycle AP (%)
Faster R-CNN	84.7	61.3	8.2	92.1	78.4
RetinaNet	87.2	63.8	11.4	93.4	81.2
YOLOv5-L	91.8	68.4	34.7	95.7	86.3
YOLOv8-L (Ours)	96.3	74.9	48.1	98.1	93.7

Table 4. Signal control performance—average intersection delay and network throughput.

Control method	Avg. delay (s/veh)	Throughput (veh/hr)	Avg. queue (veh/lane)	CO ₂ Proxy (g/veh-km)	
Webster’s fixed-cycle	52.4	1,840	8.7	184.3	—
Actuated (SCATS-emu)	44.1	2,010	7.2	162.7	15.8%
IntelliLight (DQN)	38.7	2,180	6.1	151.4	26.1%
PressLight	35.2	2,290	5.6	144.8	32.8%
MA2C	33.8	2,340	5.3	141.2	35.5%
TrafficNet-DRL (Ours)	34.2*	2,410	5.1	150.4†	34.7%

Vehicle Detection and Tracking

The fine-tuned YOLOv8-L achieved an mAP@0.5 of 96.3% on the IntersectionNet-8K test set, a 4.5 pp improvement over YOLOv5-L and an 11.6 pp improvement over Faster R-CNN (Table 4). The performance gap is particularly pronounced in motorcycles (93.7% vs. 86.3% for YOLOv5-L), which are disproportionately prevalent in Indian urban traffic and historically challenging for models trained predominantly on Western-centric datasets because of their smaller bounding box area and high inter-instance occlusion.

Signal Control Performance

Reduction vs. Fixed

TrafficNet-DRL achieved a 34.2 s/veh average delay, a 34.7% reduction compared to Webster’s fixed-cycle. MA2C achieves a marginally lower absolute delay (33.8 s/veh) owing to its unrealistic assumption of perfect ground-truth vehicle count observations. TrafficNet-DRL operates on vision-estimated counts from real camera feeds, making this a more representative evaluation (Table 3). † CO₂ proxy reduction driven by idle time reduction; TrafficNet-DRL additionally achieved 18.4% lower emission estimates relative to Webster’s fixed-cycle baseline (184.3 → 150.4 g/veh-km).

Ablation Study

The ablation study isolates the contribution of each architectural decision (Table 5). The most impactful single contribution is the predictive feedback from the ST-GCN to the MADRL controller, which reduces the average intersection delay by 13.7% relative to a controller operating without future traffic forecasts. This validates the core hypothesis that anticipatory signal control—acting on predicted congestion before it materializes—yields substantially better outcomes than purely reactive policies. The curriculum learning strategy was the second most impactful component (16.1% delay increase when ablated), underscoring the importance of staged training for multi-agent policy convergence at the 24-intersection scale.

DISCUSSION

Real-World Deployment Considerations

The translation of TrafficNet-DRL from a research prototype to a production traffic management system requires careful consideration of several engineering and operational factors. The edge computing infrastructure must support the real-time inference demands of the perception pipeline. Our profiling indicates that a mid-tier edge server (NVIDIA Jetson AGX Orin) can sustain the full detection and tracking pipeline for four camera feeds at 10 frames per second FPS, suggesting that one edge node per instrumented intersection is sufficient for most urban deployment scenarios.

Network communication between intersection agents can operate over existing urban fiber backbones with latency requirements of less than 100 ms for MADRL coordination messages, which is well within the capability of modern urban IoT network infrastructure. Graceful degradation mechanisms ensure that individual camera or sensor failures do not propagate to system-wide performance collapse: the prediction module falls back to historical pattern imputation, and the control module reverts to actuated control when real-time observations are unavailable.

Limitations

The present study has several limitations. First, the MADRL training environment, while initialized with real trajectory data, is ultimately a simulation-to-real-transfer scenario. The sim-to-real gap, particularly regarding vehicle interaction dynamics and driver response to signal changes, may attenuate performance in live deployment relative to simulation-evaluated results. Second, the current architecture does not explicitly model vulnerable road users (pedestrians and cyclists), whose right-of-way requirements impose hard constraints on permissible signal phases. The integration of pedestrian detection and demand estimation is an important avenue for future research.

Third, the fairness of RL-learned signal policies across different geographic and demographic areas of cities has not been formally evaluated. It is conceivable that an unconstrained optimizer could disproportionately favor high-volume arterials at the expense of residential or lower-income neighborhood access roads. Fairness-constrained reinforcement learning formulations should be explored in future studies.

Table 5. Ablation study—module contribution to system performance.

System variant	Pred. MAE (30 min)	Detection mAP@0.5	Avg. delay (s/veh)	ΔDelay vs. full
Full TrafficNet-DRL	2.48	96.3%	34.2	—
w/o Predictive feedback to control	2.48	96.3%	38.9	+4.7 (+13.7%)
w/o Adaptive graph learning (AGL)	2.81	96.3%	36.1	+1.9 (+5.6%)
w/o GAT Agent communication	2.48	96.3%	37.4	+3.2 (+9.4%)
w/o Transformer-ReID Tracker	2.48	89.1%	36.8	+2.6 (+7.6%)
w/o Curriculum learning	2.48	96.3%	39.7	+5.5 (+16.1%)
Single-intersection control only	2.48	96.3%	41.3	+7.1 (+20.8%)

Broader Societal Impact

The deployment of AI-driven traffic management systems has meaningful societal implications beyond technical performance. Traffic management decisions directly affect equity of access, air quality, noise pollution, and public safety across urban communities. We advocate that municipalities deploying such systems establish community oversight mechanisms, conduct regular equity audits of signal timing distributions, and transparently report system performance relative to stated service objectives. Privacy considerations surrounding continuous video surveillance must also be addressed through robust on-device data processing, minimal retention policies, and adherence to applicable data protection regulations.

CONCLUSION

This study presents TrafficNet-DRL, a unified deep learning and deep reinforcement learning framework for AI-driven real-time urban traffic monitoring and intelligent signal control. By integrating a spatiotemporal graph convolutional network for multi-intersection flow prediction, a YOLOv8-based multi-camera perception pipeline for real-time vehicle monitoring, and a multi-agent deep reinforcement learning controller with graph-structured agent communication, TrafficNet-DRL achieves substantial improvements across all evaluated dimensions of traffic management performance.

On the standard METR-LA benchmark, the ST-GCN achieved an MAE of 2.48 km/h at the 30-minute prediction horizon, which is a 15.3% improvement over the previous state-of-the-art. On the newly introduced IntersectionNet-8K benchmark, the perception pipeline achieved a vehicle detection mAP@0.5 of 96.3%, with particular strength in motorcycle detection (93.7%) relevant to South Asian urban traffic conditions. The MADRL controller reduces the average intersection delay by 34.7% relative to Webster's fixed-cycle baseline across a 24-intersection urban network, with an estimated 18.4% reduction in vehicular CO₂ emission proxy.

Beyond quantitative metrics, this study advances the scientific understanding of how anticipatory signal control, enabled by the bidirectional coupling between prediction and control modules, provides substantial operational advantages over reactive approaches. The public release of the IntersectionNet-8K dataset further contributes to the research infrastructure needed to advance this field. Future directions include pedestrian-inclusive signal optimization, federated learning across multiple city deployments, and the development of fairness-constrained reinforcement learning formulations for equitable urban traffic management.

REFERENCES

1. Bai L, Yao L, Li C, Wang X, Wang C. Adaptive graph convolutional recurrent network for traffic forecasting. *Adv Neural Inf Process Syst*. 2020;33:17804-15.
2. Chu T, Wang J, Codecà L, Li Z. Multi-agent deep reinforcement learning for large-scale traffic signal control. *IEEE Trans Intell Transp Syst*. 2020;21(3):1086-95. doi:10.1109/TITS.2019.2901791.
3. Guo S, Lin Y, Feng N, Song C, Wan H. Attention based spatial-temporal graph convolutional networks for traffic flow forecasting. *Proc AAAI Conf Artif Intell*. 2019;33:922-9. doi:10.1609/aaai.v33i01.3301922.
4. INRIX. (2024). Global traffic scorecard. [Online]. INRIX global traffic rankings. Available from: <https://inrix.com/scorecard-2023/>.
5. Ultralytics. (2023). Explore Ultralytics YOLOv8 [Online]. Ultralytics Docs. Available from: <https://docs.ultralytics.com/models/yolov8>.
6. Li Y, Yu R, Shahabi C, Liu Y. Diffusion convolutional recurrent neural network: data-driven traffic forecasting [Preprint]. 2017. arXiv:1707.01926.
7. Liang X, Du X, Wang G, Han Z. A deep reinforcement learning network for traffic light cycle control. *IEEE Trans Veh Technol*. 2019;68(2):1243-53. doi:10.1109/TVT.2018.2890726.

8. Wang R, Zhang J, Wang X, He H, Zou Y. Multi-modal and multi-agent reinforcement learning framework for urban traffic flow prediction and signal control optimization. *Sci Rep.* 2026;16(1):7612. doi:10.1038/s41598-026-37722-5.
9. Lv Y, Duan Y, Kang W, Li Z, Wang FY. Traffic flow prediction with big data: a deep learning approach. *IEEE Trans Intell Transp Syst.* 2015;16(2):1-9. doi:10.1109/TITS.2014.2345663.
10. Ma X, Tao Z, Wang Y, Yu H, Wang Y. Long short-term memory neural network for traffic speed prediction using remote microwave sensor data. *Transp Res C Emerg Technol.* 2015;54:187-97. doi:10.1016/j.trc.2015.03.014.
11. Yao H, Tang X, Wei H, Zheng G, Li Z. Revisiting spatial-temporal similarity: a deep learning framework for traffic prediction. *Proc AAAI Conf Artif Intell.* 2019;33:5668-75. doi:10.1609/aaai.v33i01.33015668.
12. Varaiya P. Max pressure control of a network of signalized intersections. *Transp Res C Emerg Technol.* 2013;36:177-95. doi:10.1016/j.trc.2013.08.014.
13. Zhang J, Zheng Y, Qi D. Deep spatio-temporal residual networks for citywide crowd flows prediction. *Proc AAAI Conf Artif Intell.* 2017;31. doi:10.1609/aaai.v31i1.10735.
14. Zheng C, Fan X, Wang C, Qi J. GMAN: a graph multi-attention network for traffic prediction. *Proc AAAI Conf Artif Intell.* 2020;34:1234-41. doi:10.1609/aaai.v34i01.5477.
15. Wojke N, Bewley A, Paulus D. Simple online and realtime tracking with a deep association metric. In: *Proceedings of the IEEE International Conference on Image Processing (ICIP)*; 2017. p. 3645-9. doi:10.1109/ICIP.2017.8296962.
16. Wei H, Zheng G, Yao H, Li Z. IntelliLight: a reinforcement learning approach for intelligent traffic light control. In: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*; 2018. p. 2496-505. doi:10.1145/3219819.3220096.
17. Kashinath SA, Mostafa SA, Mustapha A, Mahdin H, Lim D, Mahmoud MA, et al. Review of data fusion methods for real-time and multi-sensor traffic flow analysis. *IEEE Access.* 2021;9:51258-76. doi:10.1109/ACCESS.2021.3069770.
18. Rashid T, Samvelyan M, De Witt CS, Farquhar G, Foerster J, Whiteson S. Monotonic value function factorisation for deep multiagent reinforcement learning. *J Mach Learn Res.* 2020;21(178):1-51.
19. Zheng G, Xiong Y, Zang X, Feng J, Wei H, Zhang H, et al. Learning phase competition for traffic signal control. In: *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*; 2019. p. 1963-72. doi:10.1145/3357384.3357900.
20. United Nations. *World Urbanization Prospects: 2018 Revision*. New York: United Nations; 2018. Available from: <https://population.un.org/wup/Publications/Files/WUP2018-Report.pdf>.
21. Wu Z, Pan S, Long G, Jiang J, Zhang C. Graph WaveNet for deep spatial-temporal graph modeling. In: *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence. International Joint Conferences on Artificial Intelligence Organization*; 2019. p. 1907-13. doi:10.24963/ijcai.2019/264.