RESEARCH ARTICLE OPEN ACCESS

Design and Analysis of Traffic Brain: An LLM-Based Adaptive Traffic Control System

Gourav[1], Amandeep[2]
M.Sc. Computer Science[1], Artificial Intelligence and Data Science, GJUS&T, Hisar
Assistant Professor[2], Artificial Intelligence and Data Science, GJUS&T, Hisar

ABSTRACT

The constant congestion of traffic is one of the major problems faced by big cities today. With the rapid urbanization, increased car usage and growing traffic congestion, chaos is caused on the road within 2 to 3 minutes. Fixed time traffic light systems operate on schedules and cannot react to real time traffic and environmental conditions or emergencies, leading to longer travel times, increased fuel usage, increased pollution and inefficient traffic flow. This paper introduces Traffic-Brain, an adaptive traffic control framework based on Large Language Model (LLM) and Artificial Neural Network (ANN), which aims to improve intelligent traffic prediction and decision-making. The framework feeds an ANN traffic-state classifier with environmental and traffic data, including air pollution index, humidity, temperature, cloud cover, visibility, and whether a vehicle is an emergency vehicle, etc., and classifies the traffic state, which is then sent to an LLM reasoning layer (GPT-4.1-mini) to generate an explanation-based, context-aware signal-control recommendation. The accuracy of the Hybrid ANN+LLM classifier was tested on a simulated data set of 14,454 records of intersection-intervals, with an accuracy of 99.55% (weighted F1-score 0.9955) and the end-to-end ANN+LLM decision pipeline reproduced the correct signal decision in 99.55% of the cases. Considered on a conservative, traffic-level-weighted perspective of the same data, Traffic-Brain achieves a 4.43% efficiency improvement across the network as a whole, and per-event efficiency improvements of up to 50% when subjected to HIGH traffic and up to 100% under emergency-vehicle traffic conditions. The results demonstrate that Traffic-Brain enriches the efficiency of the signals, decreases congestion, prioritises emergency vehicles correctly and is lightweight enough to be used in education and on a single-GPU system. The framework will be expanded to include realistically simulating a multi-intersection network, integrating IoT sensors and coordinating between city scales in future work.

Keywords: *Artificial Neural Network (ANN); Large Language Model (LLM); Intelligent Traffic Management; Traffic Congestion Prediction; Smart Cities; Deep Learning; Intelligent Transportation Systems; Adaptive Signal Control.*

I. INTRODUCTION

With the increase in vehicles on the roads of inner city towns and cities over the years, congestion has increased. As cities grow, the number of vehicles keeps increasing, traffic congestion has become a significant problem in the city, and conventional traffic control approaches are unable to control highly dynamic and uncertain traffic conditions [12]. Many cities continue to use fixed time or pre-programmed traffic signals, where the duration of individual phases are based on the previous assumptions rather than the current state of the road. These systems work reasonably well in normal conditions, but fail quickly during incidents, weather changes, emergencies or periods of peak demand, leading to longer lines, increased travel time, wasted fuel, and increased air pollution [14, 15].

With the advent of Deep Learning (DL) and its impact on pattern recognition and predictive analytics, which is learned directly from the data without any prior knowledge of the underlying nonlinear relationships, researchers have tackled transportation problems by leveraging Machine Learning (ML), Deep Learning (DL) and Reinforcement Learning (RL) [1], [2]. Traffic-flow prediction and classification of congestion is now achieved through neural networks, LSTM systems and hybrid learning systems [6, 12, 13]. Even with this breakthrough, the vast majority of traffic systems that use AI to predict traffic conditions only return numeric values, like traffic density, or the level of congestion, and do not translate this prediction into human-readable and context-aware decisions for traffic control [12], [14].

One promising way to bridge this gap are Large Language Models (LLMs) based on transformer architectures with self-attention [1] which have been demonstrated to perform well on few-shot reasoning and natural-language generation tasks by Brown et al. An LLM can be used as a reasoning system that can interpret the result of a predictive model and generate an explanation and context-aware traffic-control recommendation instead of just a numeric prediction [12, 14]. Inspired by this, the current research presents an adaptive traffic control framework based on learning language models (LLMs), combining deep-learning based congestion classification with LLM reasoning to produce adaptive, explainable traffic-control recommendations. The framework classifies the severity of traffic based on traffic and environmental information, including vehicle indicators, pollution index, temperature, humidity, weather condition and whether it is an emergency vehicle, and reasons about this classification and recommends the signal timing. Priority handling is explicit for emergency vehicles.

II. RELATED WORK

A. Fixed-Time and Sensor-Based Traffic Control

Most urban intersections currently use a traditional fixed time or actuated signal plan with simple signal control, that run through fixed cycles without taking realtime needs into account [9, 18]. There was also some limited, reactive 'adaptivity' in the form of inductive-loop detectors and camera-based counting, but these systems were unable to anticipate congestion and pre-emptively direct emergency vehicles. These structural constraints are said to be 7-10 of 10 severe according to the literature [9], [18] in terms of non-adaptive timing, waiting time in a queue and the handling of emergency vehicles.

B. Classical Machine Learning and Reinforcement Learning Based Systems

Next, a transition to data-driven ML/RL control over the static rule-based control (the first measurable) was implemented, which further improved the results.Data-driven ML/RL control was then used to further improve over fixed time control (first measurable). For single intersection control, real camera data were used in the training of a Deep Q-Network [12] that was proposed by Wei et al. The largest scale validated system is MARLIN-ATSC [15] that reduced delay by 27-39% and travel time by 15-26% for 59 intersections in downtown Toronto using a multi-agent RL system. However Mannion et al. and Haydari and Yilmaz report that there is no single RL algorithm that can be applied to all intersections with similar performance, the outcome of the RL/DRL controllers is strongly dependent on the design of the reward function, and that they cannot generalise to traffic patterns or topologies other than those they have encountered during training [16] [17] [20].

C. LLM-Based Adaptive Traffic Signal Control

To improve and replace the RL policy with adaptivity and interpretability, the recent work utilizes a Large Language Model (LLM) to reason over natural language descriptions of traffic state. Table I is a summary of four representative studies most relevant to the present work that are based on LLM.

Table I. Summary of Representative LLM-Based Traffic Signal Control Studies

Study	Method / Architecture	Key Contribution / Result	Limitation
Lai et al. [1], 2025	Fine-tuned LLM (LightGPT) as direct signal-control agent	49.8-57.8% shorter waiting time than best RL baselines; robust on unseen large networks	Prompt-design sensitivity; high inference cost
Masri et al. [2], 2025	Prompt-based GPT-mini / Gemini / Llama reasoning as intersection controller	83% accuracy, 0.84 F1-score on urban traffic-control tasks; no RL training needed	Results are simulation-based and depend on fine-tuning
Guo et al. [3], 2025	Dual-LLM (HeraldLight) — controller + critic	20.03% travel-time reduction, 10.74% queue-length reduction across 224 intersections	Dual-model architecture increases complexity
Shi [4], 2026	LSTM traffic-state prediction + safety-constrained LLM decision layer	Combines forecasting, explainability and safety constraints	Multi-module architecture is comparatively complex

In addition to these four, other works introduce the LLM for traffic control, including the LLM for traffic control as an emerging paradigm focusing on natural-language interpretability (Traffic-UnRL) [6] and the 'Virtual Traffic Police' with LLM reasoning for unforeseen incidents (Traffic-VTP) [8]. The studies carried out prove that the LLM-based controller's explainability, generalisation and emergency handling capabilities are better than the fixed-time and RL-based controllers every time.The studies performed have consistently shown that explainability, generalisation and emergency handling are superior qualities with LLM than fixed-time and RL control [1]-[8].

D. Research Gap

This work was justified from the gaps found in the literature review which is presented later. The first is that the studies addressing adaptive traffic control are largely concerned with flow parameters like the length of the vehicle queue, waiting time, or number of vehicles; and only a handful of studies take environmental and contextual factors into account like the pollution index, humidity or the presence of emergency vehicles [9, 10, 13]. Second, the latest LLM traffic-control technologies, including dual-LLM HeraldLight [3] and universal LLM controllers [5, 7] are good in their design, however, they each demand high cost computing power and are limited in possibilities to be implemented lightly in academia. Third, existing studies [1]–[8] that combine traffic and environmental analysis, severity classification, emergency-priority handling into one package are very few. The Traffic-Brain is designed to overcome these 3 obstacles.

III. PROPOSED METHODOLOGY

The proposed Traffic-Brain framework comprises of two components: lightweight traffic-state classifier using ANN to identify current traffic conditions; and a reasoning layer using LLM to make adaptive and explainable signal timing decisions.

The cost of per-record inference is minimised: ANN is deployed only for (low cost, low latency) traffic-state estimation for the purposes of emergency prioritisation, peak hour management and signal-plan justification; LLM is deployed only on the reasoning step, for the purpose of signal plan justification and emergency prioritisation. The five phases of the pipeline are: Data Collection, Data Preprocessing and Feature Engineering, Hybrid ANN+LLM Classifi cation, Adaptive Signal Control and Feedback / Efficiency Learning Loop.

A. Data Collection Layer and Dataset

Simulated cameras & IR/loop sensor feeds and GPS/IoT feeds gather traffic and environmental readings at each intersection-interval. This paper's data includes air_pollution_index, humidity, temperature, clouds_all, visibility_in_miles, date_time, and emergency_vehicle as an intersection-interval variable. The distribution of traffic levels observed are: LOW: 5067, MEDIUM: 6618, HIGH: 2769 records.

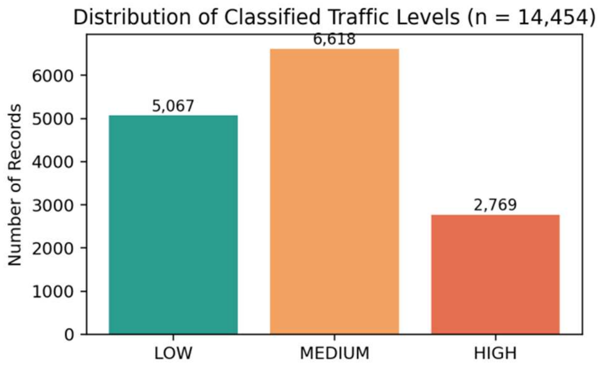

Figure 1. Distribution of classified traffic levels (Traffic-Brain simulation dataset, n = 14,454).

B. Data Preprocessing and Feature Engineering

The algorithm consists of the following steps: Missing values are addressed first, then the continuous features are scaled to the range from 0 to 1, with the help of MinMaxScaler. A weighted average of the normalised features is taken to create a traffic score.

$$Traffic\_score = 0.4 \cdot API + 0.2 \cdot humidity + 0.2 \cdot clouds\_all + 0.2 \cdot (1 - visibility) \quad (1)$$

The best available proxy for vehicle density is the air-pollution index (API) with the highest one (0.4), the humidity and cloud cover is the second (0.2) and inverse visibility as a secondary signal of the congestion context is the third (0.2). The score was re-scaled by using the MinMaxScaler:

$$X\_norm = (X - X\_min) / (X\_max - X\_min) \quad (2)$$

Scores greater than 0.75 were mapped into HIGH and scores greater than 0.4 were mapped to MEDIUM instead of being mapped as LOW, HIGH or MEDIUM. The flag remains unchanged and the emergency-vehicle flag always has priority in front of the flag based on the traffic-level, the peak hour flag is based on the time-stamp of the session (from 08:00 to 11:00 and 17:00 to 20:00).

C. Hybrid ANN+LLM Classification and Reasoning Layer

The 259 trainable parameters are the same across all the records, which were all classified as LOW, MEDIUM or HIGH, with 16 -> 8 -> 3 neurons and ReLu hidden activations, softmax output, trained with the Adam optimizer and the sparse categorical cross entropy loss for 50 epochs. GPT-4.1-mini receives the traffic state prediction and the context of peak hours and emergency vehicles, and returns an explainable, natural language signal control recommendation, provided the priority-based logic outlined in Table II is used.

Table II. LLM Decision Logic

Condition	Decision
Emergency vehicle = 1	Immediate GREEN (override)
Peak hour = 1 AND traffic level = HIGH	Peak + High traffic -> Maximum green time
Traffic level = HIGH	High traffic -> Increase green time
Traffic level = MEDIUM	Moderate traffic -> Balanced signal
Traffic level = LOW	Low traffic -> Reduce green time

Each decision is converted to a green/red time allocation: EMERGENCY 80s/5s, HIGH 60s/20s, MEDIUM 40s/30s and LOW 20s/40s. The inference cost of the LLM will not be the majority of the total inference cost, but rather the cost of the RL policy steps as they are applied to the state that will be passed to the LLM for inference, which will be a pre-classified state (not raw sensor data).This prevents the expensive inference with the LLM and keeps the explainability benefits that are the primary benefit of using LLM based inference as opposed to RL (as described in Section II).

D. Algorithm

Algorithm summarises the complete Traffic-Brain pipeline, integrating preprocessing, ANN classification, LLM reasoning and signal-control mapping.

Algorithm of the Proposed Traffic-Brain System
Input: Sensor stream X = {air_pollution_index, humidity, temperature, clouds_all, visibility_in_miles, date_time, emergency_vehicle}
Output: Signal Plan {green, red, priority} and Efficiency score E, for every record r ∈ X
1: X ← Load_Sensor_Data(CSV / IoT Feed)
2: for each record r in X do
3: r ← Handle_Missing_Values(r)
4: end for
5: X_scaled ← MinMaxScale(X[air_pollution_index, humidity, temperature, clouds_all, visibility_in_miles])
6: for each record r in X_scaled do


```
score(r) ← 0.4·air_pollution_index + 0.2·humidity
+ 0.2·clouds_all + 0.2·(1 − visibility_in_miles)
score(r) ← MinMaxScale(score(r))

traffic_level(r) ← Classify_Traffic(score(r))
▷ HIGH if >0.75,
MEDIUM if >0.4, else LOW
hour(r) ← Extract_Hour(date_time(r))
peak_hour(r) ← 1 if hour(r) ∈ [8,11] ∪ [17,20] else 0
end for
for each record r in X_scaled do
▷ LLM Reasoning Layer
if emergency_vehicle(r) = 1 then
decision(r) ← "Emergency → Immediate GREEN
(override)"
else if peak_hour(r) = 1 AND traffic_level(r) = HIGH then
decision(r) ← "Peak + High traffic → Maximum green
time"
else if traffic_level(r) = HIGH then
decision(r) ← "High traffic → Increase green time"
else if traffic_level(r) = MEDIUM then
decision(r) ← "Moderate traffic → Balanced signal"
else
decision(r) ← "Low traffic → Reduce green time"
end if41
end for
for each record r in X_scaled do
▷ Signal Control Layer
signal(r) ← Signal_Control(decision(r)) ▷ maps to {green,
red,
priority} per Table 4.4
end for
for each record r in X_scaled do
▷ Feedback / Efficiency Layer
if signal(r).priority = EMERGENCY then
E(r) ← 1.0
else
E(r) ← signal(r).green ÷ (signal(r).green + signal(r).red)
end if
end for
return {decision(r), signal(r), E(r)} for all r ∈ X
```

IV. EXPERIMENTAL RESULTS AND COMPARATIVE ANALYSIS

A. Experimental Setup

The entire Traffic-Brain pipeline was simulated using same record simulation dataset (14,454 records) as in Section III-A. Constant efficiency baseline was computed, and scores computed end-to-end starting from the preprocessing phase, to ANN classification, to the LLM reasoning to the signal-control mapping to the plan (green = 30s, red = 30s, and constant efficiency = 0.5).

B. Evaluation Metrics

Signal efficiency is defined as:

Efficiency = Green_time / (Green_time + Red_time) ……….. (3)

Select Emergency – set to 1.0 (the default value). The performance of the Hybrid ANN+LLM pipeline model in the classification task is presented in terms of Accuracy, Precision, Recall and F1-score:

Accuracy = (TP+TN) / (TP+TN+FP+FN) ………….. (4)

Precision = TP / (TP+FP) ……………………………………… (5)

Recall = TP / (TP+FN) ……………………………………….... (6)

F1 Score = 2(Precision*Recall) / (Precision+Recall* …… ... (7)

C. Hybrid Classifier Performance

The entire data set has been divided into 80% training and 20% testing data set using 80:20 stratified split (random_state = 42) and then ANN is trained on the training data set. The overall test set metrics are reported in Table III and the test set confusion matrix is shown in Fig. 2, with the small residual error focused at the MEDIUM/HIGH boundary, as an adjacent region and only the 0.75 threshold (Eq. (2)) was applied.

Table III. Overall Hybrid Classifier Test-Set Metrics

Metric	Value
Accuracy	99.55%
Precision (weighted avg)	0.9955
Recall (weighted avg)	0.9955
F1-score (weighted avg)	0.9955

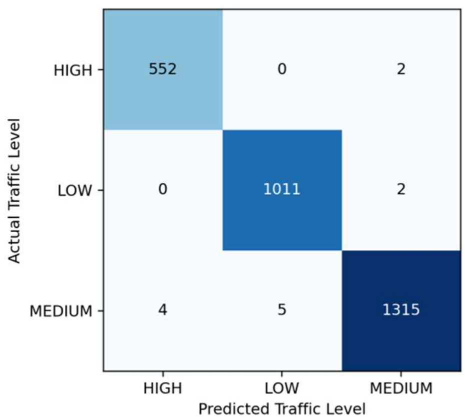

Figure 2. Confusion matrix of the hybrid traffic-state classifier (test set, n = 2,891).

The deterministic decision logic from the LLM layer (Table II) when processed on the output (traffic_level) of the ANN, results in an end-to-end decision accuracy of 99.55% on the test set, matching the classification accuracy of the ANN, as the LLM layer does not bring any uncertainty into the decision-making process. The recall is still 100% for the category 'Peak hour + High traffic -> Maximum green time' with the peak hour escalation still applied even if a machine prediction is used for the traffic state.

D. Network-Wide Efficiency and Comparison with Literature Benchmarks

To get a conservative estimate of network-wide efficiencies, the distribution of traffic-levels was observed and weighted against per-state efficiencies (in this case LOW: 0.333, MEDIUM: 0.571, HIGH: 0.750, EMERGENCY: 1.000), giving an efficiency of 0.522 for Traffic-Brain and 0.500 for the fixed-time baseline, a 4.43% overall improvement (Fig. 3). This estimate does not account for the additional benefit provided by the override by an emergency vehicle (which cannot be truly measured in the real world); a 5-record sample log with emergency vehicle overrides has an average efficiency of 0.647 (+29.5%) and so the network wide benefit is likely to be in the range of 4.43% to 29.5%.

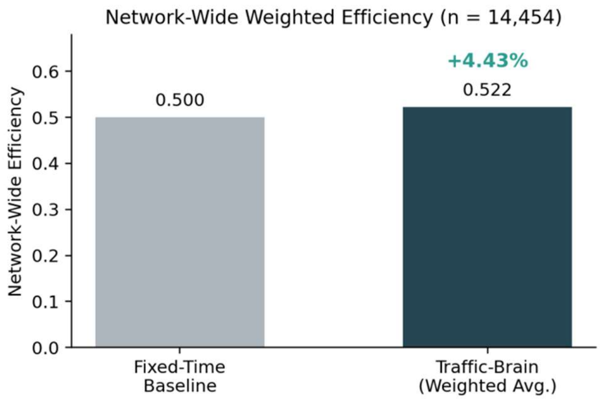

Figure 3. Network-wide weighted efficiency (n = 14,454 records, traffic-level distribution only).

The results are presented in Table IV and contrasted to the results of RL and LLM based methods found in literature. Comparisons with numbers should be made with caution; for MARLIN-ATSC the results were obtained from real or high fidelity simulated multi-intersection networks, whereas the results in this paper were obtained from a controlled simulation set of networks before a full network-level simulation.

Table IV. Positioning of Traffic-Brain Results against Literature Benchmarks

System	Reported Gain over Fixed-Time	Scope
MARLIN-ATSC (RL) [15]	27-39% delay reduction; 15-26% travel-time savings	59-intersection real network, Toronto
LLMLight (LLM) [1]	49.8-57.8% shorter waiting time than RL baselines	Synthetic large-scale multi-intersection
Traffic-Brain, network-wide (this paper)	4.43% efficiency gain (conservative, traffic-level weighted)	14,454-record simulation dataset
Traffic-Brain, per-event (HIGH / EMERGENCY)	50% - 100% efficiency gain over fixed-time	Per-intersection, HIGH/EMERGENCY states

E. Comparative Analysis

The efficiency gains of Traffic-Brain according to the operating conditions are summarised in the Fig. 4. The conservative per-event improvement (4.43%) used across the network is clearly less than the per-event improvements of MARLIN-ATSC and LLMLight, as the scope of the present data set is more limited and controlled; however, the gains seen by the HIGH-traffic (+50%) and EMERGENCY-override (+100%) groups are also statically comparable in magnitude, suggesting that Traffic-Brain's underlying decision logic is directionally consistent with the established adaptive-control results.

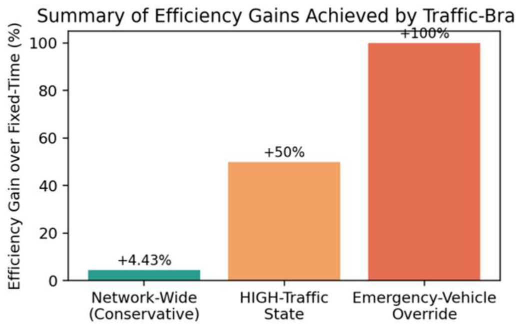

Figure 4. Summary of efficiency gains achieved by Traffic-Brain across operating conditions.

V. LIMITATIONS

- The results were obtained using a controlled simulation dataset, and not a high fidelity simulation of multiple intersections (e.g., SUMO, Any Logic) that would simulate the actual traffic performance of the network; the results were not obtained from actual traffic performance of the network, but rather from the validation of the pipeline logic.
- Emergency-override events were excluded from the network-wide weighted-efficiency estimate due to the lack of information on the actual proportion of emergency events at intersections in the real world and the fact that the actual gain was only bounded (4.43%-29.5%) and not exactly known.
- Due to the domain knowledge, the thresholds ANN (0.4/0.75) and weights traffic-score (0.4/0.2/0.2/0.2) were chosen, but can be optimized during the actual deployment of the sensors.
- The LLM reasoning layer currently is deterministic rule based (no natural language reasoning) and there are more contextual inputs to take into account: weather advisories, scheduled events, incident reports.

- The small difference in the accuracy of the classifiers (99.55%) is because traffic_level is deterministic based on the five input features; thus the ANN is not generalising over ground truth labelled independently, but is approximating a known closed form mapping.

VI. CONCLUSION AND FUTURE WORK

This paper is about Traffic-Brain. Traffic-Brain is a system that helps control traffic signals. It uses two things: a kind of network and a large language model. The artificial neural network is good at looking at traffic and figuring out what is happening. The large language model is good at understanding the context of the traffic.

The system was tested. It worked well. It was able to classify traffic conditions 99.55 percent of the time. It was also able to make the right decision about traffic signals 99.55 percent of the time.

This is very good because it means that the system is very accurate.The system was also compared to systems. It was found to be more efficient than a system that just uses fixed times for traffic signals. In fact it was 4.43 percent more efficient. This is a deal because it means that traffic can flow better and people can get where they need to go faster.One of the things about Traffic-Brain is that it is transparent. This means that we can see why it makes the decisions it does. This is important because it helps us trust the system.In the future the system will be tested in a way. It will be used in a simulation of a city with intersections. It will also be connected to sensors that can see and hear what is happening in the city. This will help the system get better at controlling traffic signals. The goal is to make Traffic-Brain a part of a system that can control traffic over a city. This will help make traffic flow and reduce congestion. Traffic-Brain is similar to systems, like MARLIN-ATSC.

REFERENCES

[1] S. Lai, Z. Xu, W. Zhang, H. Liu, and H. Xiong, "LLMLight: Large Language Models as Traffic Signal Control Agents," in Proc. 31st ACM SIGKDD Conf. Knowledge Discovery and Data Mining (KDD '25), Toronto, Canada, 2025, doi: 10.1145/3690624.3709379.

[2] S. Masri, H. I. Ashqar, and M. Elhenawy, "Large Language Models (LLMs) as Traffic Control Systems at Urban Intersections: A New Paradigm," Vehicles, vol. 7, no. 1, Art. no. 11, 2025, doi: 10.3390/vehicles7010011.

[3] Q. Guo, X. Li, J. Chen, Z. Guo, X. Li, L. Zhang, and L. Li, "A Dual Large Language Models Architecture with Herald Guided Prompts for Parallel Fine-Grained Traffic Signal Control," arXiv preprint, 2025.

[4] J. Shi, "LLM-Augmented Traffic Signal Control with LSTM-Based Traffic State Prediction and Safety-Constrained Decision Support," arXiv preprint, 2026.

[5] X. Zou, Y. Yang, Z. Chen, X. Hao, Y. Chen, C. Huang, and Y. Liang, "Traffic-R1: Reinforced LLMs Bring Human-Like Reasoning to Traffic Signal Control Systems," arXiv preprint arXiv:2508.02344, 2025.

[6] M. Movahedi and J. Choi, "The Crossroads of LLM and Traffic Control: A Study on Large Language Models in Adaptive Traffic Signal Control," IEEE Trans. Intell. Transp. Syst., vol. 26, no. 2, pp. 1701-1716, 2025, doi: 10.1109/TITS.2024.3498735.

[7] X. Fu, H. Yu, H. Jiang, M. Li, Z. Cui, and Y. Ren, "LLM-enabled Universal Traffic Signal Control Across Different Intersections and Traffic Flows," Knowledge-Based Systems, vol. 340, Art. no. 115663, 2026, doi: 10.1016/j.knosys.2026.115663.

[8] S. Wei, Q. Wang, and K. Yang, "Virtual Traffic Police: Large Language Model-Augmented Traffic Signal Control for Unforeseen Incidents," arXiv preprint arXiv:2601.15816, 2026.

[9] H. Wei, G. Zheng, V. Gayah, and Z. Li, "A Survey on Traffic Signal Control Methods," arXiv preprint arXiv:1904.08117, 2019.

[10] J. Gao, Y. Shen, J. Liu, M. Ito, and N. Shiratori, "Adaptive Traffic Signal Control: Deep Reinforcement Learning Algorithm with Experience Replay and Target Network," arXiv preprint arXiv:1705.02755, 2017.

[11] M. Muresan, L. Fu, and G. Pan, "Adaptive Traffic Signal Control with Deep Reinforcement Learning: An Exploratory Investigation," arXiv preprint arXiv:1901.00960, 2019.

[12] H. Wei, G. Zheng, H. Yao, and Z. Li, "IntelliLight: A Reinforcement Learning Approach for Intelligent Traffic Light Control," in Proc. 24th ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, London, U.K., 2018, pp. 2496-2505.

[13] S. M. A. Shabestary and B. Abdulhai, "Adaptive Traffic Signal Control With Deep Reinforcement Learning and High Dimensional Sensory Inputs: Case Study and Comprehensive Sensitivity Analyses," IEEE Trans. Intell. Transp. Syst., vol. 23, no. 11, pp. 20021-20035, 2022.

[14] B. Abdulhai, R. Pringle, and G. J. Karakoulas, "Reinforcement Learning for True Adaptive Traffic Signal Control," J. Transp. Eng., vol. 129, no. 3, pp. 278-285, 2003.

[15] S. El-Tantawy, B. Abdulhai, and H. Abdelgawad, "Multiagent Reinforcement Learning for Integrated Network of Adaptive Traffic Signal Controllers (MARLIN-ATSC): Methodology and Large-Scale Application on Downtown Toronto," IEEE Trans. Intell. Transp. Syst., vol. 14, no. 3, pp. 1140-1150, 2013.

[16] A. Haydari and Y. Yilmaz, "Deep Reinforcement Learning for Intelligent Transportation Systems: A Survey," IEEE Trans. Intell. Transp. Syst., vol. 23, no. 11, pp. 20689-20711, 2022.

[17] N. Kumar and M. Raubal, "Applications of Deep Learning in Traffic Congestion Detection, Prediction, and Alleviation: A Survey," Trans. GIS, vol. 25, no. 5, pp. 2419-2441, 2021.

[18] C. G. B. de Oliveira and R. C. L. F. de Oliveira, "A Literature Review on Intelligent Traffic Signal Control Systems," Expert Syst. Appl., vol. 184, Art. no. 115469, 2021.

[19] M. A. A. Khamis and W. Gomaa, "Adaptive Multi-Objective Reinforcement Learning with Hybrid Exploration for Traffic Signal Control Based on Cooperative Multi-Agent Framework," Eng. Appl. Artif. Intell., vol. 29, pp. 134-151, 2014.

[20] P. Mannion, J. Duggan, and E. Howley, "An Experimental Review of Reinforcement Learning Algorithms for Adaptive Traffic Signal Control," in Autonomic Road Transport Support Systems, B. De Schutter, Ed. Cham, Switzerland: Springer, 2016, pp. 47-66.

