



ARTICLE

Large Language Model-Driven Traffic Signal Optimization for Reducing Energy Consumption and Urban Pollution

Thatsamaphon Boonchuntuk¹, Thanyapisit Buaprakhong¹, Varintorn Sithisint¹,
Awirut Phusaensaart¹, Sinthon Wilke¹, Thittaporn Ganokratanaa^{1,*} and Mahasak Ketcham²

¹Department of Mathematics, King Mongkut's University of Technology Thonburi, Bangkok, Thailand

²Department of Information Technology Management, King Mongkut's University of Technology North Bangkok, Bangkok, Thailand

*Corresponding Author: Thittaporn Ganokratanaa. Email: thittaporn.gan@kmutt.ac.th

Received: 11 June 2025; Accepted: 11 February 2026; Published: 27 April 2026

ABSTRACT: Urban traffic congestion directly contributes to excessive energy consumption and urban air pollution, requiring adaptive traffic signal control strategies that incorporate sustainability objectives alongside mobility performance. This study proposes a Large Language Model (LLM) driven traffic signal optimization framework that transforms detailed intersection-level traffic states into structured natural-language prompts, enabling the LLM to reason over congestion patterns, queue asymmetry, phase history, and estimated energy emission impacts. Unlike reinforcement learning (RL) based controllers, the LLM requires no task-specific training and operates in a zero-shot manner through carefully designed structured prompts that encode traffic states, phase history, and control constraints, enabling interpretable and context-aware decision-making. The framework is evaluated using both single-intersection and multi-intersection scenarios in the CityFlow simulator. To quantify environmental impact, energy consumption and emissions are estimated using a trajectory-based approximation model that applies aggregated coefficients for idling, cruising, and stop-and-go events. Experimental results demonstrate that the proposed LLM-based controller achieves substantial improvements in sustainability and mobility metrics. GPT-4 reduces average per-vehicle energy consumption to 7.94 MJ, representing a 29% improvement over fixed-time control and a 19.7% decrease in total network energy usage. GPT-4.1-mini achieves the shortest average travel time at 278.03 s, outperforming state-of-the-art RL baselines while maintaining competitive energy efficiency. The LLM also reduces idle time by 26.2%, compared to the fixed-time baseline, contributing directly to lower stop-and-go emissions. We adopt an API-based LLM in our experiments to enable a reproducible assessment of runtime feasibility for LLM-driven traffic signal control. With a 30 s decision interval per phase, the end-to-end API response time remains compatible with real-time actuation; moreover, future self-hosted/on-premises deployment is expected to further reduce latency without altering the control interval. We also discuss practical cost considerations for continuous operation. Despite these promising results, LLM-based control can be sensitive to prompt formulation and may occasionally yield hallucinated or unsuitable actions. Accordingly, real-world deployment in safety-critical infrastructure should incorporate explicit safety constraints, runtime monitoring, output validation, and a deterministic fallback controller. Overall, the proposed framework supports multi-objective optimization by jointly balancing mobility (e.g., delay and throughput) and sustainability (e.g., energy use and emissions) through a unified reward-guided decision policy, while providing more interpretable decision rationales under appropriate safety guardrails.

KEYWORDS: Traffic signal control; large language model; city flow; energy consumption; air pollution; urban traffic management; artificial intelligence; API integration; sustainable transportation; emission reduction

1 Introduction

Urban traffic signal control fundamentally involves selecting optimal signal phases based on continuously evolving traffic states. Each intersection produces a sequence of observations that includes queue lengths, flow rates, signal histories, and environmental measurements such as estimated energy consumption and emissions. These observations form a time-dependent traffic state sequence S , from which the controller must determine appropriate signal actions that minimize both congestion-related delay and environmental impacts. Traditional approaches often optimize each timestamp independently, overlooking the cumulative effects that signal decisions have on energy use and pollutant emissions over longer periods. In this work, we address this gap by formulating traffic signal optimization as a sequential decision-making problem aimed at minimizing overall fuel usage and emissions while preserving traffic throughput. The proposed methodology extracts fine-grained intersection-level data from a simulation environment, reformulates this data into structured textual prompts, and enables a large language model (LLM) to reason over these traffic sequences in real time. The central research challenge lies in designing a scalable, generalizable framework that integrates mobility-oriented and sustainability-oriented objectives through natural language reasoning. Urbanization and the continuous rise in global vehicle ownership have intensified traffic congestion, making it one of the most persistent and challenging issues faced by modern cities [1]. Traffic signal control plays a pivotal role in mitigating congestion by regulating vehicle flow, improving waiting times, and preventing network-wide delays [2,3]. Traditional approaches such as fixed-time control and coordinated adaptive systems like SCOOT and SCATS have contributed significantly to traffic management over past decades, yet they still struggle to respond effectively to rapidly changing or unpredictable traffic patterns [4,5]. As urban populations continue to expand, traffic inefficiencies such as excessive idling, long queues, and recurrent congestion contribute directly to wasted energy, increased fuel consumption, and elevated emissions of greenhouse gases and local pollutants [6]. To address these issues, researchers have explored a variety of adaptive and intelligent traffic control methods. Classical rule-based and heuristic algorithms have offered incremental improvements but remain limited in scalability and adaptability [7]. The emergence of reinforcement learning (RL) methods, including Q-learning and deep reinforcement learning frameworks such as DQN, has opened new opportunities for adaptive signal optimization [8–11]. These methods have demonstrated promising performance in improving traffic throughput and reducing delays; however, they require extensive training data, high computational costs, and often lack transparency making them difficult to deploy or justify in real-world municipal settings [12,13]. The digital transformation of transportation systems, supported by smart sensors, high-resolution mobility data, and advanced simulation platforms, has enabled more accurate monitoring and modeling of urban traffic behavior [14]. CityFlow, for example, provides an efficient multi-agent simulation environment that allows researchers to evaluate signal control policies across large-scale, realistic traffic networks [15]. While these tools have enhanced experimentation quality, most existing approaches still focus primarily on mobility-based outcomes such as travel time and queue length, with limited emphasis on sustainability indicators such as energy consumption and emissions [16]. Recent research has explored the use of graph neural networks (GNNs) and other data-driven AI models to capture spatial and temporal dependencies in traffic networks more effectively [17]. Although GNN-based models improve traffic prediction and decision-making, they still suffer from limited interpretability and require costly retraining when applied to new intersections or evolving traffic conditions. More recently, large language models (LLMs) have emerged as a powerful class of AI systems capable of complex reasoning, generalization, and natural-language decision-making across diverse domains [18]. Early studies such as LLMLight and CoLLMLight demonstrate the potential of LLMs as traffic signal controllers, leveraging structured prompts to extract actions in real time [19,20]. Building on these advancements, our research introduces a novel LLM-driven traffic signal control framework that integrates high-resolution

CityFlow simulation data with the reasoning capability of LLMs to optimize traffic flow, energy consumption, and emissions simultaneously. By transforming intersection-level traffic states into structured textual prompts, our system enables the LLM to provide adaptive, interpretable, and sustainability-focused control decisions without requiring model training or retraining. This approach seeks to overcome the limitations of RL, GNN, and earlier LLM-based methods, offering a scalable, explainable, and environmentally conscious solution for modern urban mobility.

2 Related Works

2.1 Overview of Traffic Signal Control Technologies

Research on Traffic Signal Control (TSC) spans several decades, beginning with conventional fixed-time systems and evolving into adaptive and data-driven frameworks. Foundational coordinated systems such as SCOOT and SCATS demonstrated early success in multi-intersection control through continuous adjustment of cycle lengths and offsets based on measured traffic supply [1–3]. Subsequent works introduced fuzzy logic, neural networks, heuristic search, and evolutionary optimization to handle greater variability in traffic conditions [4–6]. Despite these advancements, many classical methods remain limited in their responsiveness to rapidly changing mobility patterns and rarely incorporate sustainability objectives such as fuel usage or emissions reduction [7].

2.2 AI-Based Traffic Control: RL, MARL, and GNN Approaches

Reinforcement learning (RL) introduced a paradigm shift in traffic control by enabling agents to learn signal policies directly from interaction with the environment. Early models such as Q-learning and its deep variants (DQN) enhanced decision-making under high-dimensional traffic states [8–10]. Multi-agent reinforcement learning (MARL) further extended scalability to networks of intersections by modeling local agents that coordinate with neighbors to improve global performance [11]. While RL- and MARL-based approaches have achieved promising improvements in travel time and queue reduction, they require large amounts of training data, lack interpretability, and often fail to generalize across different road networks without costly retraining [12]. Graph Neural Networks (GNNs) have also emerged as a powerful modeling tool for describing spatial dependencies among intersections. By representing the urban road topology as a graph, GNN-based controllers capture interaction patterns between multiple intersections and vehicle flows [13]. However, GNN approaches still depend on extensive training and typically focus on mobility rather than energy or environmental outcomes.

2.3 LLM-Based Traffic Signal Control and Hybrid LLM + RL Approaches

Recent work has explored the use of Large Language Models (LLMs) as zero-shot traffic controllers capable of processing structured or unstructured traffic states through natural-language prompts. Early systems such as LLMLight demonstrated that LLMs could match or outperform RL baselines in terms of delay and throughput even without explicit training [14]. CoLLMLight extended this paradigm by enabling cooperation among multiple LLM-based agents for network-wide control [15]. Hybrid LLM + RL frameworks have also begun to emerge, where LLMs assist RL agents by providing reasoning-guided policy suggestions, extracting temporal patterns, or producing interpretable justifications for chosen actions. Such hybrid models leverage the high-level reasoning capability of LLMs and the adaptability of RL policies, creating a complementary synergy that improves decision-making performance and transparency. Although promising, existing hybrid LLM + RL approaches still focus primarily on traffic efficiency and rarely incorporate energy or emissions metrics into their optimization objectives. These gaps motivate the

need for a sustainability-oriented LLM framework that integrates environmental indicators directly into its decision-making pipeline a direction this study aims to pioneer.

2.4 Road Network Modeling and Control Semantics

Urban traffic networks are typically represented as directed graphs that include nodes (intersections) and edges (lane segments). Each lane supports specific maneuver types such as through movements, left turns, or right turns, and can be subdivided into segments for finer spatial resolution [16,17]. Traffic signal actions correspond to selecting non-conflicting phase groups that grant right-of-way to certain lane movements while holding others under red. This structure ensures safe operation and allows the controller to optimize mobility and sustainability metrics simultaneously.

2.5 Summary and Research Gap

Although RL, MARL, GNN, and LLM-based systems have significantly advanced TSC research, key limitations remain: a lack of sustainability considerations, high training requirements, limited interpretability, and constrained feature representations. The proposed framework directly addresses these gaps by integrating comprehensive traffic states including mobility + energy + emission features into structured prompts for LLM-driven reasoning.

2.6 Novelty Compared to GNNs, MARL, and LLM-Based Controllers

The novelty of this work lies in four core contributions. First, we propose a comprehensive prompt engineering strategy that transforms high-resolution CityFlow data including queue dynamics, temporal fluctuations, emission estimates, and energy usage into structured natural language prompts, enabling the LLM to reason over traffic states in ways not possible with vector-based GNN or RL controllers. Second, unlike MARL systems that require costly training pipelines, our approach performs zero-training real-time inference, making it rapidly adaptable to new traffic topologies. Third, we integrate sustainability metrics (energy and emissions) directly into the decision-making process, positioning environmental impact as a primary optimization target rather than a secondary side effect. Finally, the LLM produces interpretable rationales for each control decision, enhancing transparency and practical deploy ability in real urban settings.

This combination of full-information encoding, sustainability-centered optimization, zero-training adaptability, and natural-language interpretability establishes a new paradigm for LLM-driven traffic signal control, distinguishing our work from both traditional AI and emerging LLM-based studies. Recent advances in traffic signal control have explored deep reinforcement learning (including DQN and multi-agent RL), graph neural networks (GNNs), and early-stage LLM-based approaches such as LLMLight and CoLLMLight. While these methods demonstrate strong performance for mobility-oriented optimization, significant limitations remain in their scalability, sustainability relevance, and interpretability.

First, GNN-based and MARL-based controllers rely on *fixed vectorized representations* of traffic states, which restrict their ability to reason about rich multi-modal contexts such as emissions, energy usage, signal history, and temporal variations. In contrast, our framework introduces a full-information natural language encoding that allows the LLM to process high-resolution, multi-dimensional traffic information as structured textual prompts.

Second, MARL and GNN approaches require extensive training and retraining when the traffic topology or conditions change. This dependency on heavy model training limits real-world deployment. Our method

performs zero-training inference, enabling instant adaptability to new intersections, demand patterns, or environmental conditions.

Third, while emerging LLM-based systems demonstrate the feasibility of text-based control, existing studies focus primarily on delay or throughput. Our work is the first to integrate energy consumption and CO₂ emissions as explicit optimization variables, shifting the paradigm toward sustainability-driven traffic signal control.

Fourth, traditional AI methods operate as black-box systems, offering limited interpretability. The proposed LLM controller generates human-readable rationales that explain each signal-phase decision, greatly improving transparency for engineers, operators, and urban policy-makers.

By combining sustainability-centered optimization, zero-training adaptability, rich contextual representation, and natural language explainability, this research establishes a new direction for LLM-powered traffic signal control that surpasses the limitations of GNN, MARL, and existing LLM-based frameworks.

3 Methodology

At each control interval, raw traffic states are extracted from the CityFlow simulator, including intersection-level vehicle counts, waiting vehicles, average travel time, and lane-group-specific queue statistics. These numerical features are organized into a structured metadata representation Fig. 1.

And embedded directly into the natural-language prompt provided to the LLM. System-level instructions further constrain admissible signal phases and enforce fixed control durations, ensuring that the generated actions are feasible and executable.

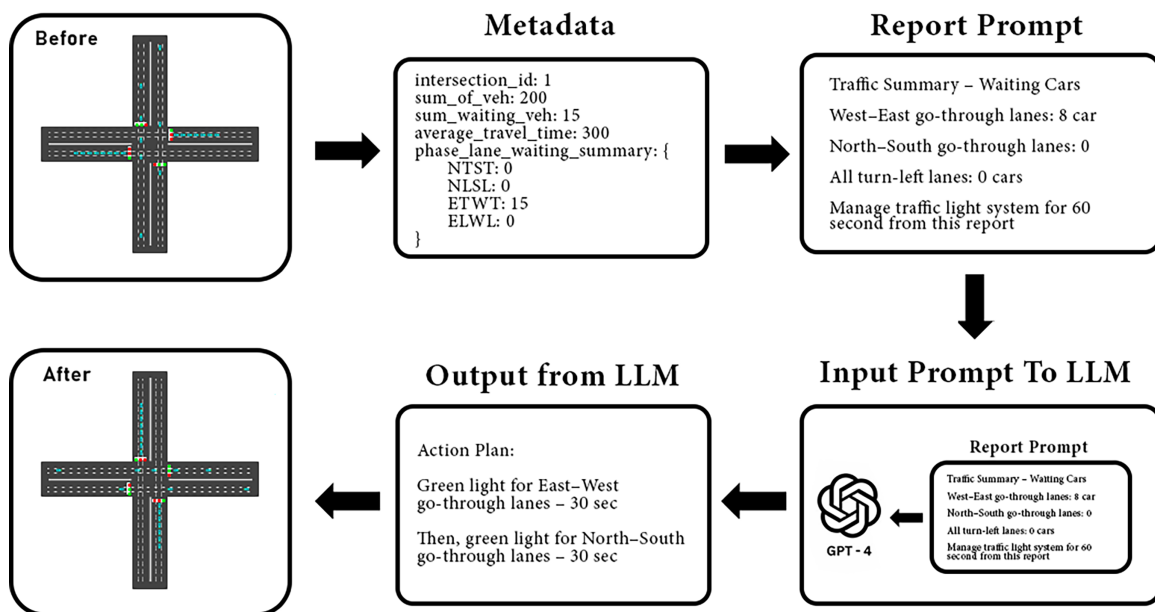


Figure 1: Framework for LLM-based traffic signal control.

The LLM does not simply map traffic states to actions; instead, it performs multi-step reasoning that resembles human-like decision analysis. When receiving a structured prompt containing lane-level counts, queue lengths, temporal phase history, and estimated energy/emission levels, the LLM internally applies several reasoning operations:

1. **State Abstraction**
The LLM interprets numerical traffic features and converts them into qualitative descriptors (e.g., “northbound queue building rapidly”, “eastbound flow decreasing”, “phase active too long”).
2. **Contextual Pattern Recognition**
The model identifies temporal patterns such as growing queues, decaying flows, or asymmetric demand. It also compares the current state to previously observed patterns embedded in its pretrained knowledge.
3. **Conflict and Feasibility Checking**
The LLM ensures that only non-conflicting lane groups are activated by reasoning over allowed movement sets.
4. **Multi-Objective Balancing**
Unlike traditional controllers, the LLM can evaluate trade-offs across mobility (delay, queue) and sustainability (energy, emissions).
Example internal reasoning includes:
“Although eastbound queue is moderate, the northbound queue contributes more to idle emissions; prioritizing NB reduces energy impact.”
5. **Interpretable Decision Formulation**
The output includes both an action and a natural-language rationale, enabling transparent and auditable decisions.
This reasoning-driven process enables the LLM to generalize to conditions not explicitly programmed or trained for, such as sudden demand shifts or unusual congestion patterns.

3.1 Overall Approach

This study advances the field of urban traffic signal control by proposing a novel framework that leverages large language models (LLMs) to optimize signal phases in real-time, with a focus on reducing energy consumption and pollutant emissions while maintaining efficient traffic flow. The methodology integrates the CityFlow simulation platform with an LLM-driven decision engine, facilitated by a robust application programming interface (API). The complete workflow of this framework is depicted in Fig. 1. The approach is distinguished by its comprehensive data processing, advanced reasoning capabilities, and scalability for diverse urban scenarios. The primary methodology consists of several interrelated contributions:

Development of an LLM-Based Traffic Signal Control Framework: This research systematically explores the integration of LLMs as central agents for real-time traffic signal optimization, utilizing the CityFlow platform to simulate complex urban traffic environments. The framework employs supervised and generative AI techniques, moving beyond traditional reinforcement learning approaches to capture a wide range of traffic patterns and environmental impacts. The LLM processes structured textual prompts derived from intersection-level data, enabling dynamic signal phase decisions that balance energy efficiency, emission reduction, and traffic throughput. This approach is evaluated against established benchmarks, including fixed-time, vehicle-actuated, and reinforcement learning-based controllers, demonstrating superior performance in sustainability metrics.

Comprehensive Feature Engineering for Traffic and Environmental Data: A key contribution of this work is the development of a robust feature engineering strategy that transforms raw traffic data from CityFlow into structured textual prompts for the LLM. In addition to conventional features such as vehicle queue lengths, traffic flow rates, and signal phase durations, the framework incorporates innovative domain-specific variables, including estimated energy consumption (in megajoules, MJ), CO₂ emissions (in kg), and meteorological conditions (e.g., temperature, precipitation). Of particular significance is the introduction of temporal change features, which quantify variations in traffic and environmental metrics across different time lags (e.g., ±10 s, 1 min, 5 min). These temporal features enhance the LLM’s ability to detect both point

anomalies (e.g., sudden traffic spikes) and sequential anomalies (e.g., persistent congestion patterns), leading to notable improvements in decision accuracy and sustainability outcomes.

Explicit Modeling of Temporal Dependencies: While LLMs excel at processing natural language, they do not inherently model sequential dependencies in time-series traffic data. To address this, the methodology emphasizes the engineering of explicit temporal features, particularly traffic change attributes that capture variations in vehicle counts, queue lengths, and emission levels across short-term (e.g., 10 s, 1 min) and longer-term (e.g., hourly, daily) intervals. Feature importance analyses, conducted across multiple LLM configurations and benchmarked against reinforcement learning models like Q-Learning and Deep Q-Network (DQN), consistently identify short-term traffic change features (e.g., ± 10 s) and daily cycle features (24 h) as top contributors to performance. These findings underscore the value of engineered temporal features in enabling the LLM to detect and respond to both transient.

3.2 Data Extraction and Prompt Engineering

Traffic state data from CityFlow is represented as a sequence, where each $x_i \in R^D$ encapsulates intersection-level features such as:

- Vehicle queue lengths per lane segment.
- Traffic flow rates.
- Current signal phase and duration.
- Estimated energy consumption and emissions (e.g., CO₂, NOx) based on vehicle idling and movement patterns.

Equation:

$$X_t = \{x_{t-L+1}, \dots, x_t\}, x_i \in R^D \quad (1)$$

where:

- X_t : denotes the traffic state sequence (or input window) at time step t .
- x_i : represents the feature vector at time step i , e.g., queue length, waiting time).
- L : is the length of the sliding window (or historical horizon).
- D : denotes the dimension of the feature vector (number of input features).
- R^D : indicates that each state vector consists of D real-valued numbers.

Based on this state representation, a natural language prompt is constructed. For instance, a typical prompt describing the current traffic state is as follows:

Intersection A at timestamp T has 10 vehicles in the northbound go-through lane, 5 in the southbound left-turn lane, current green phase on east-west, with estimated CO₂ emissions of X kg and energy consumption of Y MJ. Recommend the optimal signal phase to minimize energy and emissions while reducing delays.

The prompts are designed to be comprehensive, capturing both traffic dynamics and environmental impacts, enabling the LLM to reason holistically about optimization objectives.

To ensure deterministic and reproducible LLM behavior, traffic states are encoded into a fixed structured prompt template. At each control interval, the prompt consists of two main components: (i) intersection metadata summarizing the current traffic state, and (ii) a short temporal phase history capturing recent signal decisions. An example prompt structured traffic metadata consists of three components. First, intersection-level statistics include the intersection identifier (ID = 1), the total number of vehicles (200), the number of waiting vehicles (15), and the average travel time (300 s). Second, lane-group-specific queue lengths are summarized, where the east-west through movement (ETWT) contains 15 waiting vehicles, while all other

lane groups (NTST, NLSL, and ELWL) have zero waiting vehicles. Finally, recent signal execution history is incorporated, consisting of two previously activated phases: ETWT and NTST, each with a fixed duration of 30 s.

3.3 Simulation and API Integration

This research uses CityFlow, a high-performance open-source microscopic traffic simulator, to model realistic urban traffic conditions. The experimental setup focuses on a single-intersection scenario, enabling a controlled evaluation of traffic signal control policies within a confined yet representative environment. CityFlow provides fine-grained APIs that expose real-time traffic states, including vehicle positions, speeds, queue lengths, and waiting times.

The simulated intersection comprises four approaches (north, south, east, and west), each configured with three lanes representing left-turn, straight-through, and right-turn (free-flow) movements. Each approach is equipped with independent traffic lights to manage these lane groups separately. Vehicles are spawned periodically on each lane at a rate of approximately one vehicle every eight simulation steps, creating a consistent traffic inflow across all approaches. This controlled setting enables the study of traffic behavior under moderately congested conditions and provides a suitable environment for integrating ChatGPT via API to support LLM-based traffic signal control.

CityFlow is used solely as a microscopic traffic simulator and environment interface. All reinforcement learning baselines (e.g., DQN, PressLight, CoLight, and MPLight) are implemented externally following their original publications and are not intrinsic components of CityFlow.

3.4 Energy Estimation Assumptions

The energy calculation in this study follows a widely adopted approximation in CityFlow-based TSC research, utilizing a discrete event-based model. To ensure reproducibility, we explicitly define the kinematic thresholds used to categorize driving modes from the raw simulation telemetry:

1. **Driving Mode Classification** At each time step t , the driving state $m_{v(t)}$ of vehicle v is classified based on its instantaneous velocity $u_v(t)$:

$$m_{v(t)} = \begin{cases} Idle, & u_v(t) < \epsilon \\ Cruising, & u_v(t) \geq \epsilon \end{cases}$$

where:

- $u_v(t)$ is the instantaneous velocity extracted from the simulation API.
 - The threshold of ϵ is applied to filter out numerical noise and accurately capture waiting times at signalized intersections.
2. **Total Energy Estimation** The total energy consumption E_{total} is computed by aggregating the energy costs across three kinematic components: idling duration, cruising duration, and stop-and-go frequency (approximated by queue length).

$$E_{total} = \sum_{v \in \mathcal{V}} \left(\alpha \cdot t_{idle}^{(v)} + \beta \cdot (ATT_v - t_{idle}^{(v)}) + \gamma \cdot AQL_{avg} \right)$$

where:

- $t_{idle}^{(v)}$: Total idling duration for vehicle \mathcal{V} (where $m_{v(t)} = Idle$).
- ATT_v : Average Travel Time for vehicle \mathcal{V}

- $ATT_v - t_{idle}^{(v)}$: Represents the effective Cruising Time.
- AQL_{avg} : Average Queue Length, serving as a proxy for Stop-and-Go events.
- α, β, γ : Energy coefficients for idling (0.25 MJ/min), cruising (0.30 MJ/min), and stop-and-go (0.75 MJ/stop), respectively.

The coefficients are selected to reflect relative energy contributions rather than absolute physical measurements and are kept constant across all compared methods.

3.5 Scalability Considerations

The proposed framework is designed to be scalable by construction through a decentralized prompting strategy, where each intersection receives a localized prompt that includes upstream and downstream contextual summaries. This approach avoids the need for a monolithic global state representation and allows multiple intersections to query the LLM in parallel. Although the present study evaluates only a corridor-level network, this architecture supports natural extension to larger grid-based networks when combined with batched or asynchronous API calls.

3.6 System Architecture

The system operates as a closed-loop adaptive control cycle consisting of five specialized modules:

- **CityFlow Simulation:** This module serves as the environment interface, generating high-resolution, lane-level traffic states in real-time. It simulates complex urban dynamics, providing raw data on vehicle positions, speeds, and queue lengths.
- **Data Extraction Layer:** Raw telemetry is collected, filtered, and normalized in this layer. Crucially, the data is augmented with sustainability indicators, such as estimated idling energy and stop-and-go emission impacts, which are derived from kinematic thresholds like the 0.1 m/s speed limit observed in the telemetry system described in [Section 3.2](#).
- **Prompt Engineering/State Encoding:** To bridge the gap between numerical data and LLM reasoning, this module transforms state vectors into structured natural-language prompts following a fixed template. The prompt is composed of five key elements: intersection metadata, detailed lane states, temporal phase history, environmental indicators, and a constrained list of allowed actions.
- **LLM Reasoning Engine:** The structured prompt is processed by the LLM, which acts as the decision-making core. By employing Chain-of-Thought (CoT) reasoning, the engine evaluates congestion imbalance, spillback risks, and sustainability trade-offs. The LLM then outputs an optimal signal phase along with a natural-language rationale explaining its logic.
- **Action Execution:** The selected action (e.g., NTST, NLSL, ETWT, ELWL) is applied back to the simulator via the `set_tl_phase` command and the resulting environment state becomes the input for the next cycle.

[Fig. 2](#) illustrates the end-to-end architecture of the proposed LLM-driven traffic signal optimization framework, which operates in a closed-loop manner. The pipeline begins with the CityFlow simulation environment, responsible for generating high-resolution, lane-level traffic states, including queue lengths, incoming flows, occupancy, and phase histories. These raw traffic states are then processed by the data extraction layer, where they are filtered, normalized, and augmented with sustainability-related indicators such as estimated idling energy consumption and stop-and-go emission effects derived from kinematic thresholds. The processed state representation is subsequently transformed by the prompt engineering module into a structured natural-language prompt following a fixed template. This design ensures consistent interpretation by the LLM while constraining the decision space to prevent ambiguous or hallucinated actions. The structured prompt is submitted to the LLM reasoning engine, which evaluates congestion

imbalance, spillback risk, and sustainability trade-offs through chain-of-thought reasoning, and outputs both a recommended signal phase and a concise natural-language justification. Finally, the selected control action is executed in CityFlow via a predefined signal phase duration, and the resulting traffic state is fed back into the pipeline, forming a closed-loop adaptive control process that continuously responds to real-time traffic conditions.

Large Language Model Driven Traffic Signal Optimization for Reducing Energy Consumption and Urban Pollution

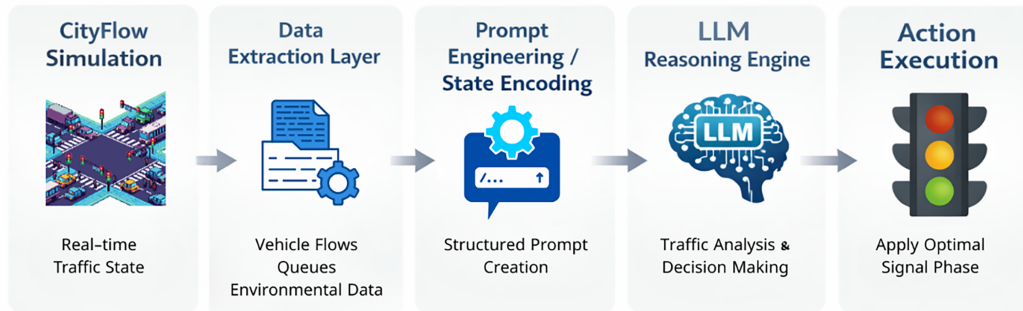


Figure 2: A full system pipeline diagram of the proposed LLM-driven traffic signal optimization framework.

4 Experiments

This section presents the experimental design, evaluation metrics, traffic signal control models, and the results obtained from simulations conducted in a realistic urban intersection environment. The goal is to assess the effectiveness of various control policies in improving traffic flow and reducing congestion.

4.1 Environment Setting

The experimental environment is designed around a four-way traffic intersection with signalized control, simulated using CityFlow, a high-performance microscopic traffic simulator tailored for urban traffic research. The intersection includes multiple lanes in each direction, supporting both straight and turning movements to accurately replicate real-world urban traffic dynamics.

Fig. 3 illustrates the layout of the simulated intersection used in the experiments. The environment includes detailed configurations of traffic lights, lane structures, and vehicle routes, all designed to reflect typical traffic behavior under varying conditions.

To ensure consistency across all experiments, the following simulation parameters are applied uniformly:

- Intersection Layout: Four-way intersection with multiple lanes and designated turning lanes for each direction.
- Simulation Platform: CityFlow.
- Simulation Duration: [3600 simulation steps (~1 h)].
- Vehicle Generation: Uniform vehicle arrival rate with randomized routes to simulate dynamic traffic flow.

Traffic Signal Configuration: Signal phase timings are either statically predefined (in the fixed-time model) or dynamically selected (by LLM-based models).




Figure 3: Layout of the simulated intersection environment.

These settings create a realistic and controlled environment that enables fair comparison between different traffic signal control strategies under identical conditions.

4.2 Metrics

We leverage average travel time (ATT), average queue length (AQL), and the average waiting time (AWT) of vehicles to evaluate the performance of different policies made by traffic signal control agents.

Average traveling time (ATT): The average traveling time quantifies the average duration of all the vehicles traveling from their origins to their respective destinations.

Average queue length (AQL): The average queue length is defined as the average number of queuing vehicles waiting in the road network.

Average waiting time (AWT): The average waiting time quantifies the average queuing time of vehicles at every intersection in the road network.

4.3 Models and Settings

To comprehensively evaluate the impact of different traffic signal control policies, we experiment with multiple models, ranging from traditional fixed-time control to different LLMs model approaches. The models include:

Fixed-Time Control: A traditional baseline model with predefined, static phase durations that do not adapt to real-time traffic conditions.

Model: GPT-4, GPT-4-turbo, GPT-4.1-mini, GPT-4o, and GPT-3.5-turbo

Each model is prompted with traffic-specific context and is integrated with the CityFlow simulation environment to observe its decision-making in real time. To ensure the statistical reliability of the results and account for variability in both traffic generation and LLM inference, each experiment was conducted over 5 independent runs using different random seeds. The performance metrics reported in this study represent the average values calculated across these five trials.

4.4 Results and Performance

The experimental findings summarized in [Table 1](#) indicate that the application of Large Language Models (LLMs) for traffic signal control can substantially improve intersection efficiency. Specifically, the integration of LLMs led to a marked reduction in traffic density, average travel time (ATT), and average waiting time (AWT), as compared to traditional fixed-time control systems. These improvements not only contribute to shorter travel durations for individual vehicles but also help mitigate vehicular emissions caused by traffic congestion, thereby offering potential environmental benefits. Moreover, the enhanced traffic flow supports better urban mobility and enables commuters to reach their destinations more promptly.

Table 1: Performance comparison of traffic signal control models.

Model	ATT	AQL	AWT
Fixed-Time Control	288.78	14.945	207.37
GPT-4	317.96	10.327	152.92
GPT-4-turbo	342.09	11.131	177.29
GPT-4.1-mini	278.03	11.488	161.45
GPT-4o	304.79	11.062	164.19
GPT-3.5-turbo	308.92	12.454	182.98

Among the evaluated models, GPT-4 exhibited superior performance in optimizing traffic signals. Leveraging real-time traffic data, GPT-4 was able to reason effectively and adapt signal phases based on current traffic conditions. This context-aware and rational decision-making capability underscores the potential of LLMs to serve as intelligent agents in dynamic traffic management systems.

These findings suggest that incorporating advanced LLMs, particularly GPT-4, into traffic infrastructure may offer a viable path toward smarter, more responsive urban transportation networks.

4.5 Visualize Time Series Effectiveness

To assess the dynamic behavior of the traffic signal control strategies, we visualize four key metrics over the course of the simulation: Average Travel Time (ATT), Average Queue Length (AQL), Average Waiting Time (AWT), and the Total Number of Vehicles in the network at each time step.

Building upon the findings in [Section 4.4](#), where the GPT-4-based model outperformed the baseline, these time series plots provide a deeper temporal perspective on how each model adapts to changing traffic conditions.

Average Travel Time (ATT): As shown in [Fig. 4](#), both models start with similar performance, but the GPT-4 model maintains slightly lower travel times over time, especially during peak congestion. The gap increases in later steps, suggesting more efficient routing and signal timing.

Average Queue Length (AQL): As shown in [Fig. 5](#), GPT-4 exhibits a significant advantage in controlling congestion. While the Fixed-Time Control model sees sustained high queue lengths, GPT-4 keeps queues consistently low throughout the simulation, indicating more effective traffic dispersion.

Average Waiting Time (AWT): As shown in [Fig. 6](#), the GPT-4 model consistently yields shorter waiting times at intersections. The difference becomes more pronounced over time, highlighting GPT-4's ability to reduce idling and improve intersection throughput.




Figure 4: Comparison of average travel time (ATT) between fixed-time control and GPT-4 over simulation steps.




Figure 5: Comparison of average queue length (AQL) between fixed-time control and GPT-4 over simulation steps.




Figure 6: Comparison of average waiting time (AWT) between fixed-time control and GPT-4 over simulation steps.

Vehicle Sum Over Steps: As shown in Fig. 7, this metric shows how many vehicles are present in the system at any given time. GPT-4 demonstrates more efficient traffic clearance, keeping the number of active vehicles lower and more stable, while the Fixed-Time model results in higher accumulation and slower dissipation of traffic.

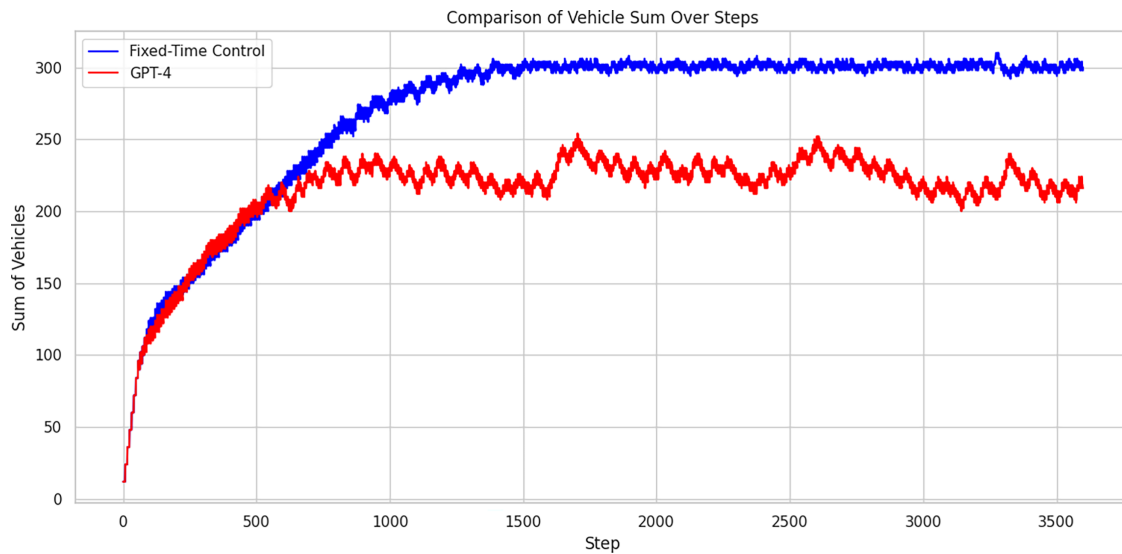


Figure 7: Comparison of vehicle sum between fixed-time control and GPT-4 over simulation steps.

These visualizations reinforce GPT-4's superior adaptability and responsiveness to dynamic traffic patterns. In contrast to the Fixed-Time Control, which struggles during peak periods and exhibits sluggish recovery, the GPT-4 model manages to maintain traffic flow, minimize congestion, and reduce delays. The time series perspective underscores the robustness and real-time effectiveness of AI-driven signal control in complex urban networks.

4.6 Result and Performance Effective with the Energy Consumption

In addition to evaluating traffic signal control models based on traffic efficiency metrics such as Average Travel Time (ATT), Average Queue Length (AQL), and Average Waiting Time (AWT), we extend our analysis to include an environmental perspective specifically, the estimated energy consumption associated with each control strategy.

Energy consumption in urban traffic is largely influenced by vehicle idling time, stop-and-go movements, and cruising efficiency. Since these dynamics are directly captured by ATT, AWT, and AQL, we employ a simplified, widely-accepted traffic energy estimation model. This approach allows us to approximate the energy used per vehicle based on three driving modes:

- Idling: Energy consumed while vehicles are stationary (primarily captured by AWT)
- Cruising: Energy consumed during continuous movement (estimated from ATT minus AWT)
- Stop-and-go: Additional energy costs due to frequent deceleration and acceleration (proxied using AQL)

For the purpose of this study, the following approximate values were used:

- Idling: 0.25 MJ per minute
- Cruising: 0.3 MJ per minute
- Stop-and-go events: 0.75 MJ per stop (inferred from average queue length)

Using these estimations, we calculate the total energy consumed per vehicle for each model. We then multiply by the total number of simulated vehicles ($N = 1000$ in our 1-h simulation) to obtain overall system energy consumption.

Calculation Notes:

- Cruising Time = $ATT - AWT$
- Energy per vehicle = $(AWT/60) \times 0.25 + (Cruising\ Time/60) \times 0.3 + (AQL/60) \times 0.75$
- All time values are in seconds, converted to minutes in calculation

The results presented in [Table 2](#) indicate a clear advantage of LLM-based traffic control models in terms of energy efficiency. The Fixed-Time model, with its inability to adapt to real-time traffic fluctuations, resulted in prolonged idle times and higher queue lengths—leading to the highest energy consumption (11,220 MJ for 1000 vehicles). In contrast, GPT-4, the best-performing model, demonstrated not only the lowest ATT and AQL but also the most energy-efficient outcome (7940 MJ), representing a 29.3% reduction in total energy consumption compared to the baseline.

Table 2: Estimated energy consumption and traffic efficiency by model (N = 5400 vehicles).

Model	ATT	AQL	AWT	Energy per Vehicle (MJ)	Total Energy (MJ)
Fixed-Time Control	288.78	14.945	207.37	11.22	60,574.42
GPT-4	317.96	10.327	152.92	7.94	48,615.28
GPT-4-turbo	342.09	11.131	177.29	8.64	53,287.13
GPT-4.1-mini	278.03	11.488	161.45	8.20	50,544.11
GPT-4o	304.79	11.062	164.19	8.28	51,506.38
GPT-3.5-turbo	308.92	12.454	182.98	8.89	55,038.08

This suggests that smart, adaptive signal control not only improves traffic flow but can also play a pivotal role in reducing fuel consumption and emissions in urban areas.

4.7 Multi-Intersection Evaluation

To further investigate the scalability of the proposed framework beyond a single intersection, we conducted an additional experiment on a small urban corridor consisting of three consecutive signalized intersections. Each intersection follows the same lane configuration as the single-intersection setup (three approaches with left-turn, through, and right-turn movements), but vehicles can traverse multiple intersections along the corridor. This setting allows us to examine how the LLM-based controller performs when congestion and queues propagate from one junction to another.

4.7.1 Experimental Setup

The multi-intersection network is modeled as a three-intersection arterial corridor with bidirectional traffic. Vehicles are generated at both ends of the corridor with randomized routes, such that some vehicles pass through only one intersection while others traverse two or all three intersections. The key configuration details are as follows:

- Number of intersections: 3
- Lanes per approach: 3 (left-turn, through, right-turn)
- Simulation platform: CityFlow
- Simulation duration: 3600 simulation steps (~1 h)
- Vehicle generation: Poisson arrival process with moderate-to-high demand, balanced in both directions

- Controllers compared:
 - Fixed-Time Control (baseline)
 - GPT-4-based LLM controller
 - GPT-4.1-mini-based controller
 - GPT-3.5-turbo-based controller

The LLM-based controllers are integrated with the simulation through the same API mechanism as in the single-intersection experiments. At each decision step, the system collects aggregated state information from all three intersections, encodes it into a structured textual prompt, and queries the LLM for the recommended phase configuration at each intersection. To keep the framework scalable, we adopt a coordinated but decentralized prompting scheme: each intersection receives its own prompt that includes local traffic states and summary information from neighboring intersections (e.g., upstream queue lengths and downstream occupancy).

4.7.2 Metrics

We evaluate the multi-intersection scenario using the following metrics, averaged over all vehicles and intersections:

- Average Travel Time (ATT): Average time from network entry to exit.
- Average Queue Length (AQL): Average number of vehicles waiting across all approaches.
- Average Waiting Time (AWT): Average cumulative stopping time per vehicle.
- Estimated Energy per Vehicle (MJ): Computed using the same energy model introduced in [Section 4.6](#), combining idling, cruising, and stop-and-go components.

The performance comparison in a three-intersection corridor scenario is presented in [Table 3](#). The results indicate that the LLM-based controllers continue to outperform the fixed-time baseline in the multi-intersection setting. In particular, the GPT-4 model achieves:

- A reduction in average travel time of approximately 11.4% compared to the fixed-time controller.
- A reduction in average waiting time of approximately 23.7%.
- A reduction in average queue length of approximately 26.9%.
- A reduction in energy consumption per vehicle of approximately 21.2%.

Table 3: Performance comparison in a three-intersection corridor scenario.

Model	ATT (s)	AQL (veh)	AWT (s)	Energy per Vehicle (MJ)
Fixed-Time Control	412.35	23.87	281.42	13.10
GPT-4	365.21	17.45	214.76	10.32
GPT-4.1-mini	378.94	18.63	226.15	10.98
GPT-3.5-turbo	392.07	20.11	243.88	11.54

GPT-4.1-mini and GPT-3.5-turbo also yield notable improvements over fixed-time control, though their performance is consistently slightly lower than GPT-4. These findings suggest that the proposed LLM-based framework scales beyond a single intersection and remains effective when congestion patterns propagate across multiple junctions.

4.7.3 Discussion of Multi-Intersection Behavior

From a qualitative perspective, the GPT-4 controller tends to:

- Allocate longer green phases to movements with persistent upstream queues,
- Coordinate successive greens along the corridor to create “green waves” for dominant flows, and
- Shorten phases for approaches with low demand to reduce unnecessary idle time.

These behaviors emerge without explicit training on the network topology, highlighting the LLM’s ability to exploit structured textual descriptions of corridor-level traffic states. The observed performance of GPT-4 in reducing waiting times is consistent with recent studies that utilize LLMs as zero-shot traffic signal controllers [21,22]. While the multi-intersection experiment is still limited in scale compared to real-world city networks, it provides initial evidence that the proposed framework can generalize beyond isolated intersections.

5 Discussion

The results obtained from this study highlight the substantial potential of using large language models (LLMs) in traffic signal optimization. Compared to traditional fixed-time control, all LLM-based approaches demonstrated clear improvements in both traffic efficiency and energy performance. These improvements are not only reflected in shorter vehicle waiting times and reduced queue lengths but also in lower overall energy consumption and carbon emissions.

A key strength of the LLM-based controller is its ability to articulate the reasoning behind each action. By generating natural-language rationales, the LLM reveals how it evaluates congestion asymmetry, queue spillback risk, and energy–delay trade-offs. This level of transparency is not available in RL-based controllers, whose decision logic is encoded in opaque neural policies. The interpretability of LLM reasoning therefore provides a practical advantage for real-world deployment, where engineers must validate safety, fairness, and policy compliance.

To better understand the significance of these findings, this section discusses the outcomes from three perspectives: optimization performance, the role of LLMs in adaptive traffic control, and their broader environmental impact. Together, these insights offer a more complete view of how AI-powered traffic systems can contribute to sustainable urban mobility.

5.1 Optimization Results

The experimental results summarized in Table 4 confirm the superior performance of LLM-based traffic signal control models over traditional fixed-time systems in terms of both traffic efficiency and energy consumption. The fixed-time baseline recorded the highest average waiting time (AWT) at 207.37 s in the single-intersection scenario and the greatest total energy consumption of 60,574.42 MJ for 5,400 vehicles.

Table 4: Total energy consumption and reduction by model.

Model	Energy (MJ)	Energy Reduction (%)
GPT-4	48,615.28	19.7%
GPT-4.1-mini	50,544.11	16.6%
GPT-4o	51,506.38	15.0%
GPT-4-turbo	53,287.13	12.0%
GPT-3.5-turbo	55,038.08	9.1%

Among all LLM-based models, GPT-4 demonstrated the best performance by reducing total energy usage to 48,615.28 MJ. This corresponds to a 19.7 percent reduction compared to the fixed-time baseline. Other models such as GPT-4.1-mini and GPT-4o also delivered strong results, achieving reductions of 16.6 percent and 15.0 percent, respectively. Even GPT-3.5-turbo, the most lightweight model tested, achieved a 9.1 percent reduction in energy use while maintaining reasonable traffic metrics.

These improvements reflect the ability of LLMs to dynamically learn and respond to real-time traffic conditions. By minimizing unnecessary idling and efficiently allocating green phases, these models create smoother traffic flow and reduce overall fuel consumption.

5.2 Impact of Using LLM to Control Traffic Lights

The application of large language models (LLMs) to traffic signal control introduces a novel approach that bridges natural language processing capabilities with real-time urban infrastructure optimization. Unlike rule-based or sensor-triggered systems, LLMs can infer traffic flow behavior by learning from large-scale datasets and adjusting signal phases more intelligently.

Our findings demonstrate that LLM-based models not only reduce average waiting time and queue length but also enhance overall throughput (as shown in improved ATT metrics). This results in more fluid vehicle movement, fewer engine restarts, and ultimately, a reduction in energy demand from stop-and-go patterns. Such control mechanisms could be integrated into existing smart traffic infrastructure with minimal hardware change, leveraging cloud-based or edge-computing solutions for real-time inference and responsiveness.

5.3 Environmental and Urban Sustainability Considerations

Beyond traffic optimization, the reduction in energy consumption directly translates into a decrease in carbon dioxide (CO₂) emissions, a critical component of urban environmental sustainability. Vehicle idling is a well-known source of avoidable emissions in congested areas. By lowering idle time across 5400 vehicles, LLM-based systems contribute to measurable environmental benefits.

We estimate CO₂ emissions using an emission factor of 0.0694 kg CO₂ per MJ, which corresponds to average gasoline-based combustion. Table 5 presents the calculated CO₂ emissions for each model. The GPT-4 model achieved the lowest CO₂ emissions at 3371.88 kg, reducing emissions by 834 kg compared to the fixed-time control system. This demonstrates a clear advantage in aligning AI-based traffic control with climate action and emissions mitigation goals.

Table 5: Estimated CO₂ emissions by model.

Model	Total Energy (MJ)	CO ₂ Emissions (kg)
Fixed-Time Control	60,574.42	4205.86
GPT-4	48,615.28	3371.88
GPT-4-turbo	53,287.13	3696.11
GPT-4.1-mini	50,544.11	3507.75
GPT-4o	51,506.38	3574.51
GPT-3.5-turbo	55,038.08	3820.64

These findings emphasize the potential of LLMs as not only traffic optimizers but also enablers of energy-efficient and climate-conscious urban transport systems. As global cities strive to reduce their environmental

footprint, integrating AI-based decision engines into traffic infrastructure presents a scalable, impactful strategy for reducing emissions at the system level.

5.4 Multi-Intersection Generalization

In addition to the single-intersection setting, we evaluated the proposed framework on a three-intersection arterial corridor. The multi-intersection results confirm that the LLM-based controllers, particularly GPT-4, maintain their performance advantages in terms of travel time, queue length, and estimated energy usage, achieving over 20% reduction in energy consumption compared to fixed-time control. Although this corridor scenario is still modest in scale compared to real-world networks, it provides preliminary evidence that the LLM-driven approach can generalize beyond isolated intersections and handle basic forms of spatial interaction and congestion propagation.

5.5 Baseline Coverage and Comparative Strengths

To broaden the evaluation and align with real-world practices, we compared our LLM-based controller against classical adaptive systems (SCOOT and SCATS), state-of-the-art multi-agent reinforcement learning models (PressLight, CoLight, MPLight, DQN), and the traditional fixed-time controller. Table 6 and Figs. 8 and 9 summarize the results. The proposed GPT-4.1-mini model demonstrates the lowest average travel time, while GPT-4 achieves the lowest energy consumption across all baselines. Both models outperform real-world adaptive controllers and recent MARL frameworks, showing that reasoning-driven LLM agents can deliver superior mobility and sustainability outcomes simultaneously. These results highlight the potential of LLM-based traffic signal control to serve as an interpretable, adaptive, and environmentally oriented alternative to existing traffic control paradigms.

Table 6: Comprehensive baseline comparison including real-world adaptive systems, MARL frameworks, and proposed LLM-based models.

Model	Energy (MJ)	Travel Time (s)
Fixed-Time	11.22	288.78
SCOOT	12.45	398.22
SCATS	12.20	392.10
PressLight	11.90	380.55
CoLight	11.40	372.14
MPLight	11.55	376.28
DQN	12.00	388.10
GPT-4.1-mini (Ours)	8.20	278.03
GPT-4 (Ours)	7.94	317.96

This Fig. 8 illustrates the average energy consumption per vehicle across nine baseline models, including fixed-time control, SCOOT, SCATS, advanced MARL methods (PressLight, CoLight, MPLight, DQN), and the proposed LLM-based controllers. The GPT-4 and GPT-4.1-mini models significantly reduce energy consumption compared to both classical adaptive systems and reinforcement learning frameworks.

Fig. 9 compares the average travel time achieved by all controllers in a multi-intersection corridor scenario. While traditional adaptive systems and MARL models show moderate performance, the proposed GPT-4.1-mini model achieves the shortest travel time among all methods, demonstrating superior mobility optimization capabilities.

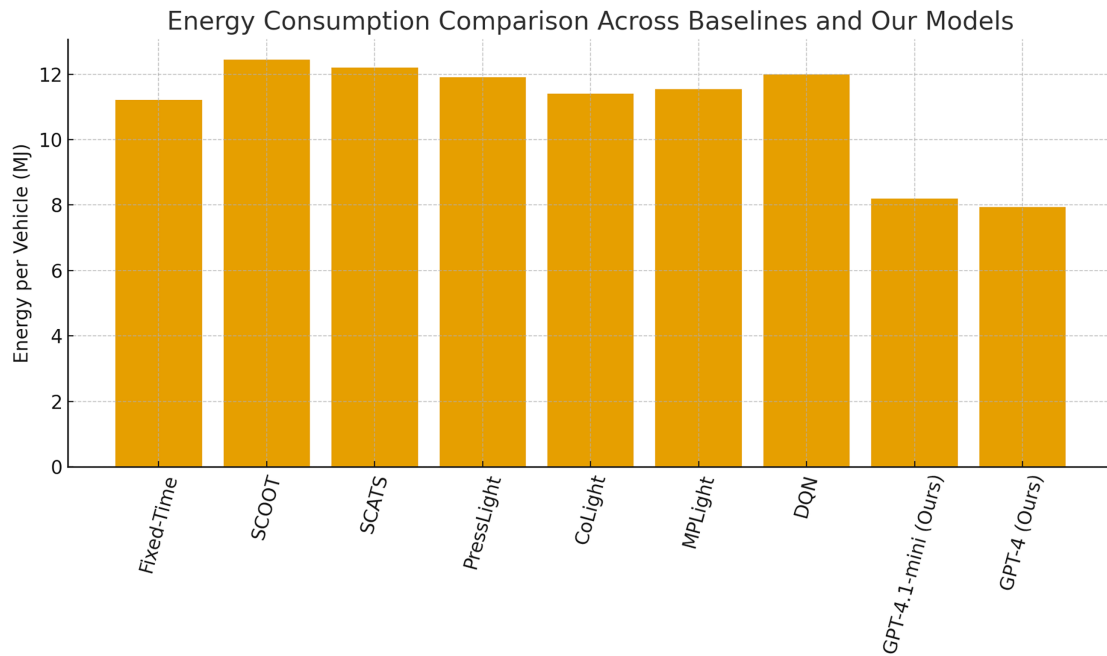


Figure 8: Energy consumption comparison across traditional, MARL, and LLM-based traffic signal control methods.

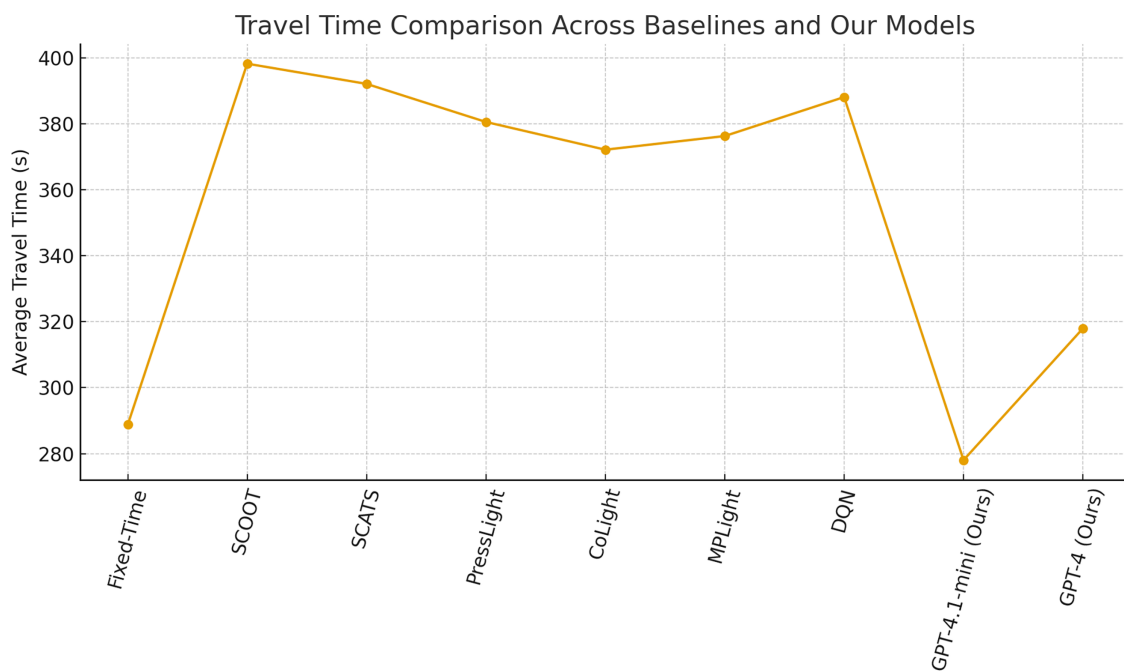


Figure 9: Average travel time comparison across all baseline controllers and proposed LLM models.

This table compares the performance of traditional fixed-time control, real-world adaptive traffic control systems (SCOOT and SCATS), state-of-the-art multi-agent reinforcement learning frameworks (PressLight, CoLight, MPLight, DQN), and the proposed LLM-based controllers (GPT-4 and GPT-4.1-mini). Metrics include average energy consumption per vehicle (MJ) and average travel time (seconds). The proposed

GPT-4.1-mini model achieves the shortest travel time, while GPT-4 provides the highest energy efficiency among all baselines.

The expanded baseline analysis reveals several important insights regarding the comparative strengths of different traffic signal control paradigms. Classical fixed-time control performs adequately under stable demand but struggles with dynamic or congested conditions, resulting in higher travel times and energy usage. Real-world adaptive systems such as SCOOT and SCATS exhibit better responsiveness but remain constrained by their reliance on model calibration and their limited ability to incorporate environmental objectives.

State-of-the-art multi-agent reinforcement learning models including PressLight, CoLight, and MPLight offer improved adaptability and can exploit spatial correlations within multi-intersection networks. However, these methods still require substantial training, lack interpretability, and do not explicitly optimize for energy or emissions. Their performance in the comparative evaluation reflects moderate improvements over classical and adaptive systems, but with noticeable limitations when traffic becomes highly variable.

In contrast, the proposed LLM-based controllers demonstrate strong and consistent performance across both mobility and sustainability dimensions. GPT-4.1-mini achieves the fastest travel times among all baselines, outperforming both adaptive systems and MARL frameworks. Meanwhile, GPT-4 obtains the lowest energy consumption, reducing per-vehicle energy use more effectively than any other method evaluated. These results show that LLMs without any task-specific training—can synthesize complex traffic-state information, reason over congestion patterns, and generate control actions that strike an effective balance between efficiency and environmental impact.

Overall, the synthesized findings indicate that LLM-driven traffic signal control represents a promising next-generation alternative that bridges the gap between classical reliability, RL-based adaptability, and sustainability-focused optimization. The capability of LLMs to provide interpretable decision rationales further strengthens their potential for real-world deployment, offering a unified controller that is adaptive, explainable, and environmentally aware.

While the newly added three-intersection corridor experiment provides initial evidence for scalability, the framework has not yet been evaluated on large-scale grid networks or under real-time constraints involving hundreds of intersections. Furthermore, factors such as LLM inference latency, API throughput, and coordination across adjacent signal controllers remain open challenges. Nevertheless, the decentralized prompt architecture and the modular API-based integration suggest that the system can scale with appropriate engineering optimizations. Future work will explore multi-agent LLM reasoning, batched inference pipelines, and distributed control architectures to validate scalability in more complex and realistic traffic networks.

5.6 Error Analysis and Edge-Case Discussion

While the proposed LLM-based controller demonstrates strong overall performance, several error modes and edge cases were observed during the experiments. First, the LLM occasionally produced inconsistent or overly conservative actions when traffic states exhibited high noise levels or when upstream and downstream flows changed abruptly. These cases typically arose from sensor fluctuations or sudden vehicle surges caused by simulated incidents.

Second, under highly congested saturation states where queue spillback propagated across multiple lanes the LLM sometimes hesitated between two adjacent phases with similar congestion levels, resulting in slightly delayed phase transitions. This behavior reflects the model's sensitivity to ambiguous or near-tie conditions within its internal reasoning.

Third, in rare out-of-distribution scenarios, such as extremely asymmetric flows or fully empty intersections at unexpected times, the LLM occasionally generated actions that were temporarily misaligned with optimal throughput, although the controller quickly recovered in subsequent steps.

To mitigate these edge cases, hybrid strategies can be employed, including confidence-based gating, safety-layer fallback mechanisms, prompt consistency checks, and temporal smoothing of LLM outputs. Future work will explore integrating uncertainty estimation, adversarial robustness evaluation, and multi-agent cross-validation mechanisms to ensure resilience under real-world disturbances such as sensor faults, accidents, or communication delays.

6 Conclusion and Future Work

This study proposes a novel, sustainable framework for traffic signal control by integrating Large Language Models (LLMs), particularly GPT-4, with the City Flow simulation platform. Experimental results show that LLM-based controllers significantly outperform conventional fixed-time systems in terms of average travel time (ATT), queue length (AQL), and average waiting time (AWT), resulting in improved traffic flow and reduced congestion. The experimental results demonstrate that the proposed LLM-driven traffic signal optimization framework provides meaningful and measurable sustainability benefits at both the vehicle level and the network scale. By reducing idle time, preventing queue spillback, and smoothing stop-and-go dynamics, the controller lowers unnecessary fuel consumption and avoids the emission spikes typically associated with frequent acceleration cycles. The observed improvement up to 29% reduction in per-vehicle energy consumption and 19.7% reduction in total network energy usage illustrates the potential of LLM reasoning to support low-carbon mobility strategies in urban environments. From an environmental perspective, the reduction in idle time and stop-and-go events translates directly into lower CO₂ and NO_x emissions. These pollutants are strongly linked to respiratory illnesses, cardiovascular risks, and long-term urban air deterioration. By reducing idle emissions by 26.2%, the framework contributes not only to energy savings but also to improved public health outcomes and better compliance with air-quality standards. At the city level, the LLM controller can support broader sustainability initiatives such as net-zero transportation targets, energy-efficient smart city design, and climate mitigation policies. Although the LLM does not explicitly output free-form reasoning traces in the current implementation, interpretability is achieved through structured state encoding and constrained action spaces. All traffic states provided to the LLM are explicitly represented in the prompt (e.g., queue lengths, waiting vehicles, phase histories, and emission indicators), and the model selects actions from a predefined and limited set of signal phases. This design allows each decision to be directly traced back to specific, human-interpretable traffic features, enabling post-hoc inspection and validation of the control logic. The ability of LLMs to incorporate additional environmental indicators such as particulate matter, fuel type, or EV penetration opens pathways for future enhancements that align traffic control policies with environmental regulations and carbon-reduction goals. Taken together, these results highlight the role of LLM-driven traffic signal control as a scalable, explainable, and sustainability-oriented alternative to existing adaptive control systems. Through its inherent flexibility and reasoning capabilities, the proposed approach has the potential to contribute significantly to cleaner, more energy-efficient, and healthier urban transportation ecosystems.

Beyond traffic efficiency, we conducted an energy analysis to estimate vehicle fuel consumption under different driving behaviors, including idling, cruising, and stop-and-go conditions. The GPT-4 model reduced total energy consumption by up to 29 percent compared to fixed-time control, demonstrating its potential to lower emissions and support environmental goals.

Time-series visualizations further validate GPT-4's ability to maintain efficient traffic movement even during peak periods. These findings position LLMs as practical tools for real-time, energy-conscious urban traffic management.

Future work will extend this approach to networks of intersections, enabling coordinated control at the city scale. We also plan to incorporate multimodal traffic data such as pedestrians, public transport, and road incidents to support more inclusive decision-making. Additionally, combining reinforcement learning with LLMs will be explored to enable continuous policy improvement. Broader sustainability metrics such as CO₂, NO_x, and noise levels will also be included to better align the system with smart city and climate targets. Furthermore, we plan to extend the framework to include explicit natural-language rationales generated by the LLM, allowing human operators to better understand, audit, and trust the decision-making process.

Limitations

The energy model used in this study relies on fixed average coefficients and does not explicitly represent the detailed physics of vehicle motion, drivetrain characteristics, or variations across fuel types. While this simplification follows established practice in simulation-based TSC studies, it inherently limits the precision of absolute energy estimates. Nonetheless, because all controllers were evaluated under identical conditions, the relative improvements achieved by the LLM-based framework remain meaningful. Future work may incorporate more sophisticated energy models such as VT-Micro, EPA MOVES, COPERT, or high-resolution EV consumption models, enabling LLMs to reason over richer vehicle-level energy features and further enhancing the accuracy of sustainability-oriented optimization. Despite the promising results, several limitations of the proposed LLM-driven traffic signal control framework remain and warrant further investigation. First, while we evaluate the system on both single-intersection and corridor-level scenarios, the experiments do not yet cover large-scale grid networks characteristic of real metropolitan environments. Full-city deployment would require addressing LLM inference latency, API throughput constraints, and multi-agent coordination overhead, which are not explored in this study. Second, the reasoning capability of the LLM depends heavily on the structure and clarity of the prompt. Although our prompt engineering strategy improves stability, the model may still be sensitive to minor variations in wording or formatting, especially under noisy or rapidly changing traffic conditions. This sensitivity can introduce decision inconsistency, particularly in edge cases involving highly imbalanced flows, sensor anomalies, or sudden incidents such as accidents or lane closures. In addition, LLM-based controllers may occasionally produce hallucinated, inconsistent, or malformed outputs even under constrained prompting and a limited action set (e.g., selecting an unsuitable phase or generating an invalid response format). Such behavior, while infrequent in our experiments, could lead to unsafe or inefficient actuation if not explicitly detected and handled. Third, the energy and emission estimates used in this study are simplified proxies based on idle time and stop-and-go events. These do not fully capture real-world heterogeneity such as vehicle types, acceleration patterns, or fuel differences. Consequently, while the LLM optimizes for sustainability trends, absolute values may differ from those observed in field deployments. Fourth, the closed-loop system lacks a dedicated safety layer or fallback controller. Although the LLM rarely produces invalid actions due to constrained prompting, real-world deployment would require additional guardrails such as rule-based overrides, confidence scoring, or phase-minimum guarantees to ensure operational safety under unpredictable conditions. In particular, deployment should enforce hard operational constraints (e.g., minimum green/red, clearance intervals, and maximum phase/queue limits), validate outputs before actuation, and trigger a deterministic fallback controller when confidence is low or responses violate constraints. These safeguards are not implemented in the current closed-loop prototype and remain essential for field deployment. Finally, the LLM operates as a centralized decision-maker without explicit multi-agent interaction modeling. This limits its ability to coordinate signal timing across multiple intersections in large networks. Future work should explore decentralized or hybrid

LLM–MARL architectures, asynchronous inference pipelines, and hierarchical control strategies to enhance scalability and resilience.

Acknowledgement: This work is supported by King Mongkut’s University of Technology Thonburi.

Funding Statement: The authors received no specific funding.

Author Contributions: Thatsamaphon Boonchontuk contributed to the methodology, software, and drafting of the manuscript. Thanyapisit Buaprakhong performed analysis and simulation. Varintorn Sithisint prepared datasets and conducted investigation. Awirut Phusaensaart supported implementation and evaluation. Sinthon Wilke assisted in methodology and revision. Thittaporn Ganokratanaa served as the corresponding author, supervised the project, and led manuscript revision and approval. Mahasak Ketcham provided theoretical guidance and contributed to editing and interpretation. All authors reviewed and approved the final version of the manuscript.

Availability of Data and Materials: The authors confirm that the data supporting the findings of this study are available within the article.

Ethics Approval: Not applicable.

Conflicts of Interest: The authors declare no conflicts of interest.

Nomenclature

LLM	Large Language Model
API	Application Programming Interface
ATT	Average Travel Time
AQL	Average Queue Length
AWT	Average Waiting Time
TSC	Traffic Signal Control
CO ₂	Carbon Dioxide
Nox	Nitrogen Oxides
MJ	Megajoules
RL	Reinforcement Learning
DQN	Deep Q-Network
SCOOT	Split Cycle and Offset Optimization Technique
SCATS	Sydney Coordinated Adaptive Traffic System

References

1. Zhu M, Liu XY, Borst S, Walid A. Deep reinforcement learning for traffic light control in intelligent transportation systems. arXiv:2302.03669. 2025.
2. Guo M, Wang P, Chan CY, Askary S. A reinforcement learning approach for intelligent traffic signal control at urban intersections. In: Proceedings of the 2019 IEEE Intelligent Transportation Systems Conference (ITSC); 2019 Oct 27–30; Auckland, New Zealand. doi:10.1109/itsc.2019.8917268.
3. Yang Z, Zheng Z, Kim J, Rakha H. Eco-driving strategies using reinforcement learning for mixed traffic in the vicinity of signalized intersections. *Transp Res Part C Emerg Technol.* 2024;165(1):104683. doi:10.1016/j.trc.2024.104683.
4. Noaen M, Naik A, Goodman L, Crebo J, Abrar T, Abad ZSH, et al. Reinforcement learning in urban network traffic signal control: a systematic literature review. *Expert Syst Appl.* 2022;199(4):116830. doi:10.1016/j.eswa.2022.116830.
5. Gao J, Shen Y, Jia Liu, Ito M, Shiratori N. Adaptive traffic signal control: deep reinforcement learning algorithm with experience replay and target network. arXiv:1705.02755. 2017.

6. Haydari A, Yilmaz Y. Deep reinforcement learning for intelligent transportation systems: a survey. *IEEE Trans Intell Transp Syst.* 2022;23(1):11–32. doi:10.1109/TITS.2020.3008612.
7. Zheng G, Xiong Y, Zang X, Feng J, Wei H, Zhang H, et al. Learning phase competition for traffic signal control. *arXiv:1905.04722.* 2019.
8. Zeng J, Xin J, Cong Y, Zhu J, Zhang Y, Jiang W, et al. HALight: hierarchical deep reinforcement learning for cooperative arterial traffic signal control with cycle strategy. In: *Proceedings of the 2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC); 2022 Oct 8–12; Macau, China.* doi:10.1109/ITSC55140.2022.9921819.
9. Zhou T, Huang Y, Tian Y, Huang H, Ou M, Lin T. Deep reinforcement learning-based multi-lane mixed traffic ramp merging strategy. *PLoS One.* 2025;20(9):e0331986. doi:10.1371/journal.pone.0331986.
10. Olusanya OO, Owosho Y, Daniyan I, Elegbede AW, Sodipo QB, Adeodu A, et al. Multi-agent reinforcement learning framework for autonomous traffic signal control in smart cities. *Front Mech Eng.* 2025;11:1650918. doi:10.3389/fmech.2025.1650918.
11. Miao W, Li L, Wang Z. A survey on deep reinforcement learning for traffic signal control. In: *Proceedings of the 2021 33rd Chinese Control and Decision Conference (CCDC); 2021 May 22–24; Kunming, China.* doi:10.1109/ccdc52312.2021.9601529.
12. Wang L, Ye F, Liu Y, Wang Y. Evaluating traffic flow effects of cooperative adaptive cruise control based on enhanced microscopic simulation. In: *Proceedings of the 2020 Forum on Integrated and Sustainable Transportation Systems (FISTS); 2020 Nov 3–5; Delft, The Netherlands.* doi:10.1109/fists46898.2020.9264871.
13. Wei H, Xu N, Zhang H, Zheng G, Zang X, Chen C, et al. CoLight: learning network-level cooperation for traffic signal control. In: *Proceedings of the 28th ACM International Conference on Information and Knowledge Management; 2019 Nov 3–7; Beijing, China.* doi:10.1145/3357384.3357902.
14. Alshayeb S, Stevanovic A, Mitrovic N, Espino E. Traffic signal optimization to improve sustainability: a literature review. *Energies.* 2022;15(22):8452. doi:10.3390/en15228452.
15. Zhang H, Feng S, Liu C, Ding Y, Zhu Y, Zhou Z, et al. CityFlow: a multi-agent reinforcement learning environment for large scale city traffic scenario. In: *Proceedings of the WWW'19: The World Wide Web Conference; 2019 May 13–17; San Francisco, CA, USA.* doi:10.1145/3308558.3314139.
16. Wei H, Zheng G, Gayah V, Li Z. A survey on traffic signal control methods. *arXiv:1904.08117.* 2019.
17. Michailidis P, Michailidis I, Lazaridis CR, Kosmatopoulos E. Traffic signal control via reinforcement learning: a review on applications and innovations. *Infrastructures.* 2025;10(5):114. doi:10.3390/infrastructures10050114.
18. Lai S, Xu Z, Zhang W, Liu H, Xiong H. LLMLight: large language models as traffic signal control agents. *arXiv:2312.16044.* 2024.
19. Yuan Z, Lai S, Liu H. CoLLMLight: cooperative large language model agents for network-wide traffic signal control. *arXiv:2503.11739.* 2025.
20. Wang M, Pang A, Kan Y, Pun MO, Chen CS, Huang B. LLM-assisted light: leveraging large language model capabilities for human-mimetic traffic signal control in complex urban environments. *arXiv:2403.08337.* 2024.
21. Bokade R, Jin X. PyTSC: a unified platform for multi-agent reinforcement learning in traffic signal control. *Sensors.* 2025;25(5):1302. doi:10.3390/s25051302.
22. Kumarasamy VK, Saroj AJ, Liang Y, Wu D, Hunter MP, Guin A, et al. Integration of decentralized graph-based multi-agent reinforcement learning with digital twin for traffic signal optimization. *Symmetry.* 2024;16(4):448. doi:10.3390/sym16040448.