vehicles

MDPI

Article

Large Language Models (LLMs) as Traffic Control Systems at Urban Intersections: A New Paradigm

Sari Masri [1], Huthaifa I. Ashqar [2,3,*] and Mohammed Elhenawy [4]

[1] Natural, Engineering and Technology Sciences Department, Arab American University, 13 Zababdeh, Jenin P.O. Box 240, Palestine; s.masri3@student.aaup.edu

[2] AI and Data Science Department, Arab American University, 13 Zababdeh, Jenin P.O. Box 240, Palestine

[3] Artificial Intelligence Program, Fu Foundation School of Engineering and Applied Science, Columbia University, 500 W 120th St., New York, NY 10027, USA

[4] Centre for Accident Research & Road Safety—Queensland (CARRS-Q) and Centre for Data Science, Queensland University of Technology, 130 Victoria Park Rd., Kelvin Grove, QLD 4059, Australia; mohammed.elhenawy@qut.edu.au

* Correspondence: huthaifa.ashqar@aaup.edu

Abstract: This study introduces a novel approach for traffic control systems by using Large Language Models (LLMs) as traffic controllers. The study utilizes their logical reasoning, scene understanding, and decision-making capabilities to optimize throughput and provide feedback based on traffic conditions in real time. LLMs centralize traditionally disconnected traffic control processes and can integrate traffic data from diverse sources to provide context-aware decisions. LLMs can also deliver tailored outputs using various means such as wireless signals and visuals to drivers, infrastructures, and autonomous vehicles. To evaluate LLMs' ability as traffic controllers, this study proposed a four-stage methodology. The methodology includes data creation and environment initialization, prompt engineering, conflict identification, and fine-tuning. We simulated multi-lane four-leg intersection scenarios and generated detailed datasets to enable conflict detection using LLMs and Python simulation as a ground truth. We used chain-of-thought prompts to lead LLMs in understanding the context, detecting conflicts, resolving them using traffic rules, and delivering context-sensitive traffic management solutions. We evaluated the performance of GPT-4o-mini, Gemini, and Llama as traffic controllers. Results showed that the fine-tuned GPT-mini achieved 83% accuracy and an F1-score of 0.84. The GPT-4o-mini model exhibited a promising performance in generating actionable traffic management insights, with high ROUGE-L scores across conflict identification of 0.95, decision making of 0.91, priority assignment of 0.94, and waiting time optimization of 0.92. This methodology confirmed LLMs' benefits as a traffic controller in real-world applications. We demonstrated that LLMs can offer precise recommendations to drivers in real time including yielding, slowing, or stopping based on vehicle dynamics. This study demonstrates LLMs' transformative potential for traffic control, enhancing efficiency and safety at intersections.

Keywords: urban intersection; traffic control systems; Large Language Models (LLMs); logical reasoning

Academic Editor: Elżbieta Macioszek

Received: 18 November 2024
Revised: 24 January 2025
Accepted: 25 January 2025
Published: 27 January 2025

Citation: Masri, S.; Ashqar, H.I.; Elhenawy, M. Large Language Models (LLMs) as Traffic Control Systems at Urban Intersections: A New Paradigm. *Vehicles* **2025**, *7*, 11. https://doi.org/10.3390/vehicles7010011

1. Introduction

Traffic control systems are essential for maintaining safety and efficiency in urban transportation networks, where traffic of vehicles, pedestrians, and public transportation vehicles mixes at intersections and other conflict points. By managing the flow of traffic, these systems reduce congestion [1], prevent crashes [2], and minimize delays [3]. Effective and timely traffic control is especially crucial in urban areas, where densely populated environments and high mobility demand can quickly lead to congestion, pollution, and safety-critical events if not managed well [4]. Moreover, traffic volume characteristics vary on different days of the week, different weeks of the month, as well as different months of the year, which makes it hard to manage using a static schedule [5]. Traffic signals, signage, and intelligent transportation systems (ITSs) are the backbone of this control system, which guides traffic flow based on established algorithms, sensor inputs, and historical data to optimize intersection usage, and increase its efficiency [6,7].

However, modern urban traffic control systems face significant challenges due to the increasing complexity and dynamism of city traffic patterns [8,9]. Traditional systems often struggle to adapt to real-time changes in traffic conditions, such as sudden surges in vehicle numbers, diverse modes of transportation, and unpredictable pedestrian and driver behaviors, particularly in mixed traffic scenarios with pedestrians, cyclists, and vehicles. For instance, Yarlagadda and Pawar [10] explore driver heterogeneity under diverse conditions, while Bella and Silvestri [11] investigate driver–bicyclist interactions. Further studies [12,13] highlight the impact of unpredictable behaviors on intersection safety. Additionally, optimizing multiple objectives simultaneously such as safety, operational efficiency, and environmental impact further challenges many conventional traffic control systems. For example, ref. [14] examines the complexities of balancing safety and efficiency, whereas [15,16] underscore the limitations of traditional methods in addressing these competing objectives. Thus, there is a persistent need for more advanced, flexible, and adaptive traffic control solutions that can respond dynamically to real-time data and manage traffic flows through a holistic method. LLMs have shown promise in various domains. For instance, refs. [17,18] focus on LLMs' logical reasoning, while refs. [19–22] emphasize their ability to process diverse data inputs for context-aware decision making.

In the context of traffic management, LLMs can handle data efficiently and create context-appropriate responses and predictions. Previous studies indicated the effectiveness of LLMs in traffic signal management, including systems labeled as LLMLight [23] and TrafficGPT [24], which both demonstrate how LLMs can optimize traffic signal timing under varying conditions. Nevertheless, investigating the use of LLMs as traffic controllers is still missing in the literature [4]. The contributions of the paper are as follows:

1. Proposing a novel paradigm where LLMs act as dynamic traffic controllers, leveraging their logical reasoning, scene understanding, and decision-making capabilities.
2. Presenting an LLM-based method through the concept of 4D traffic control system (i.e., Detect, Decide, Disseminate, Deploy) that centralizes traffic control processes traditionally managed by disconnected components.
3. Outlining a four-stage process for integrating LLMs into traffic control, including data creation and initialization, prompt generation using the chain-of-thought (CoT) method, conflict identification and resolution, and fine-tuning and performance analysis.
4. Evaluating the performance of LLMs using metrics such as accuracy, precision, recall, F1-score, and Recall-Oriented Understudy for Gisting Evaluation (ROUGE-L). It identifies the Generative Pre-trained Transformer (GPT-4o-mini) model as the most effective in conflict detection and traffic management tasks.

2. LLMs as Traffic Control Systems: A New Paradigm

This study introduces a new paradigm for traffic control systems using LLMs as traffic controllers that can serve as practical and efficient alternatives to the current signalized intersection systems. LLMs have shown potential in several domains. For example, ref. [21] discusses their logical reasoning strengths, while refs. [25–29] illustrate decision-making advantages in complex environments. This study argues that LLMs offer a transformative approach to traffic control, where their inherent flexibility and adaptability can dynamically respond to changing traffic conditions, unlike traditional fixed or adaptive control systems. Notably, refs. [9,30] examine the limitations of static signal timing, whereas [31] explores adaptive but still constrained approaches. By deploying LLMs in traffic management, intersections could benefit from more adaptive, context-aware decision making that addresses real-time traffic flows and diversely mixed traffic types (vehicles, pedestrians, cyclists, and public transportation) and provides tailored-module feedback to drivers with minimal delay [32]. Replacing (or enhancing) existing systems with LLM-based controllers could lead to several advantages, including optimized traffic throughput, enhanced response to fluctuating traffic demands, and a reduction in congestion and waiting times [33]. This LLM-driven control could significantly reduce environmental impact by optimizing fuel consumption and minimizing idling times, which might contribute to eco-friendly and sustainable urban mobility solutions.

LLMs may also provide a unique capability to handle multi-objective optimization tasks in traffic management. LLMs can concurrently address and prioritize multiple objectives. For example, ref. [34] focuses on balancing speed with environmental considerations, while [35] highlights the role of data-driven optimization in reducing emissions. This flexibility enables LLMs to balance factors like speed, timing, and flow with environmental considerations. Using pretrained knowledge and real-time data processing, LLMs can adjust signal timing or alternatively provide direct instructions to the drivers to adjust their speed to decrease emissions and fuel consumption while preserving safety and reducing operational delays. Moreover, the adaptive nature of LLMs allows them to evolve and refine their decision making based on historical and real-time data, enabling them to learn and adapt traffic patterns over time. For instance, ref. [9] evaluates real-time data streams, and ref. [30] emphasizes historical data integration for continuous improvement.

That said, the adoption of LLMs as traffic controller introduces a new paradigm that expands and enhances traditional definitions and features. Traffic control systems are conventionally defined as managing transportation system users to ensure conflict-free movement. For example, ref. [8] highlights complexities in adaptive signal systems, while [36] analyzes the broader scope of traffic management devices. Traditionally, traffic control at intersections involves various devices including signals, signs, and pavement markings, which operate manually or by rule-based algorithms [37]. LLMs, however, allow for a more integrated and centered approach that can combine control components. They have the capacity to interpret and process vast amounts of heterogeneous data inputs. For example, refs. [38,39] highlight multimodal sensor fusion techniques, while [40] explores real-time analytics. This can adapt to diverse intersection designs, including four-leg intersections, roundabouts, interchanges, and emerging urban intersection models. Moreover, LLMs can communicate instructions through voice, signals, code, visuals, or wirelessly using vehicle-to-infrastructure (V2I) and vehicle-to-vehicle (V2V) channels. For instance, ref. [18] details V2I protocol designs, while refs. [41,42] discuss advancements in low-latency V2V communication, which can enhance existing control frameworks.

This study presents LLM as a traffic controller through the concept of *4D traffic control system* (i.e., Detect, Decide, Disseminate, Deploy), as shown in Figure 1 [8,37,43]. LLMs could centralize and integrate the disconnected process of detecting and responding to various traffic conditions, which characterizes the traditional traffic controller systems, as shown in Figure 1a. We argue that LLMs can leverage data from a multitude of detection technologies at the same time, such as loop detectors, Global Positioning System (GPS), video imaging, and dedicated short-range communications (DSRC), processing both individual and group trajectories to gain a comprehensive understanding of traffic flows in urban intersections [34]. This detection informs their decision-making process, which allows LLMs to logically reason "where", "when", "who", and "how" to control traffic based on real-time and predictive analysis. Moreover, the integration of multiple objectives in the "Decide" phase, from safety to operational efficiency and energy savings, ensures that LLMs align with specific key goals in traffic management.

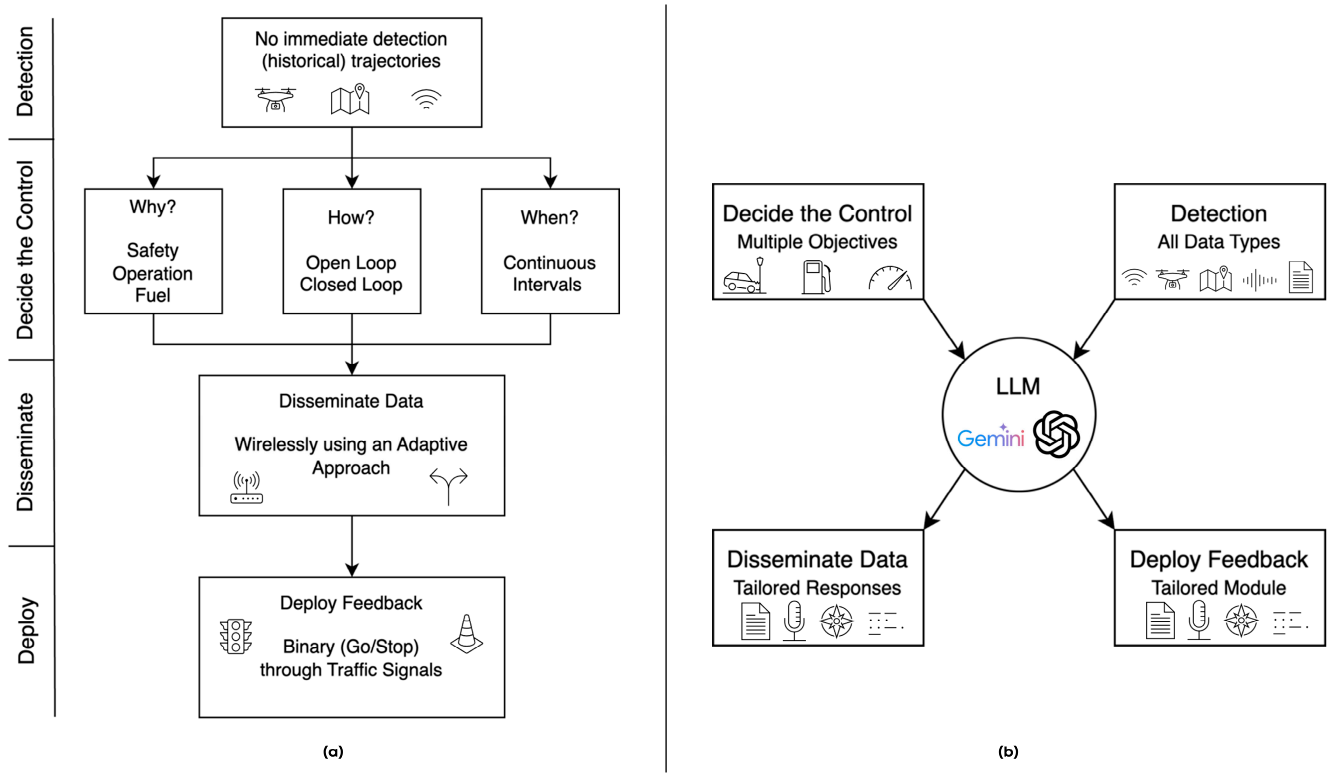

Figure 1. A comparison between (**a**) the traditional traffic control systems and (**b**) the proposed new paradigm of using LLMs as a traffic controller. The traditional system uses historical data and structured intervals for decision making and control dissemination through standard traffic signals. In contrast, the LLM-based system leverages diverse real-time data sources and advanced decision-making capabilities, allowing for adaptive, multi-objective control dissemination and tailored feedback deployment.

In the "Disseminate" and "Deploy" phases, LLMs have the potential to show more versatility. LLMs can interact with human, machine, and other road infrastructures through multiple communication channels and through different means including voice, visual signals, wireless signals, code, or DSRC. LLMs can also provide timely and effective control outputs that can be binary (stop/go) or non-binary (adjustments to speed or path through logical and actionable instructions to drivers). Moreover, the control outputs generated by LLMs can be modulated based on user type—whether human drivers, pedestrians, or autonomous vehicles—enabling tailored responses that increase compliance and reduce risks associated with complex traffic conditions [26,27,38].

The suitability of LLMs for traffic control lies in their unique capabilities, such as semantic understanding, multimodal processing, and adaptive decision making [2,3]. Unlike traditional rule-based or optimization algorithms, LLMs like GPT-4o-mini can interpret complex, context-rich scenarios involving vehicle movement, pedestrian interactions, and environmental constraints. Their ability to integrate visual and textual data enables them to process multimodal inputs like traffic camera feeds and movement logs effectively. Moreover, prompt engineering allows for dynamic adaptation to nuanced traffic conditions, outperforming static models. The scalability and near-real-time processing of LLMs make them ideal for responding to time-sensitive situations, while transfer learning facilitates fine-tuning for domain-specific tasks. In this study, GPT-4o-mini provided high-quality explanations and actionable recommendations. These capabilities position LLMs as transformative tools for real-time, scalable, and context-aware traffic management systems, addressing limitations of conventional methods.

3. Background

LLMs have the potential to enable more flexible, adaptable, explainable traffic control systems as well as able to provide actionable feedback to the drivers, traffic engineers, and policymakers. Previous studies presented the integration of road traffic management solutions using LLMs and Internet of Things (IoT) technologies [31,33,44]. Different solutions were introduced such as routing mechanisms, intelligent transportation light solutions, or network traffic management strategies, thereby classifying them [31,33,44]. A survey on traffic management with machine and deep learning has been conducted to exemplify the advantages and disadvantages of such techniques [45]. Other surveys highlighted future research directions and gave insight into how machine learning and deep learning can help solve traffic management problems [46].

Leveraging LLMs in recent work for traffic prediction involved incorporating sequence and graph embedding layers to obtain features compatible with the input formats of LLMs, followed by utilizing efficient fine-tuning techniques [23]. Experiments showed that these frameworks achieved impressive historical sample size and few-shot prediction performance [23]. Innovative combinations of multiple LLMs and traffic foundation models can be an attractive approach for using LLMs with the proper ability to perceive, analyze, and manipulate transportation data [24]. Moreover, other studies proposed a multi-task decision-making system for autonomous driving [47] that utilizes reinforcement learning and LLM sequence modeling to perform effectively in complicated settings [48] such as uncontrolled intersections [49]. Furthermore, it has been shown that using LLMs with input from several sensors, among them cameras and LiDAR, enhances theme perception of transport scenes in a sensible concession of an all-round transportation survey [25]. In the context of advanced connected and automated driving, including LLMs also helps harness traffic management effectiveness [50].

Moreover, the capacity of LLMs to develop a safety case with minimum human effort in traffic management shows and ensures compliance and safety [51]. Upgrading capabilities of prognosis of transportation functions based on LLMs presents a new transformative approach to advance transportation systems [52].

4. Methodology

The proposed framework for introducing LLMs as traffic control system and evaluating its performance in real time is shown in Figure 2. It broadly comprises four key stages, including data creation and initialization, prompt generation, conflict identification, and fine-tuning and model analysis.

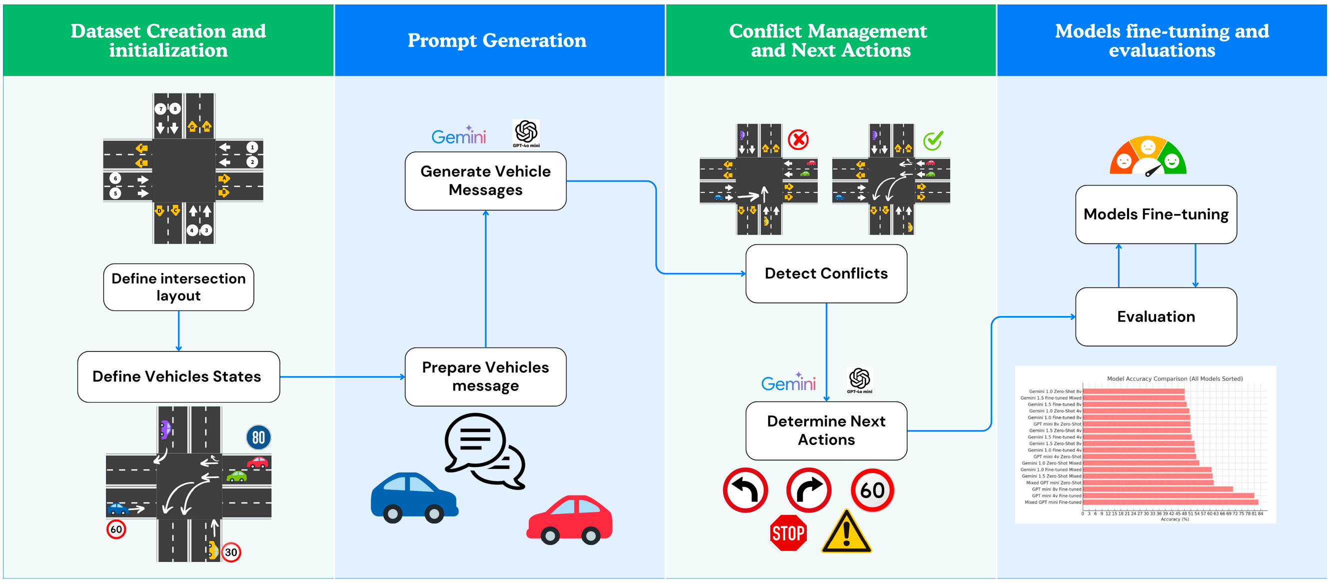

Figure 2. Flowchart of the proposed framework.

4.1. *Dataset Creation and Initialization*

The framework starts with dataset creation and initializations of the environment, where it explains the intersection layout to the LLM, including lanes, entering and exiting points, and initial parameters for each vehicle such as speed, initial lane, direction, and destination through the intersection. A traffic dataset was developed to cover intersection scenarios in real-world urban settings. This was achieved by creating a system based on rational thought, calculation methods, and traffic laws. The following section describes the steps involved in producing and generating the actual datasets for the scenarios. The process can be found in Figure 3, which outlines the workflow for scenario preparation and data collection.

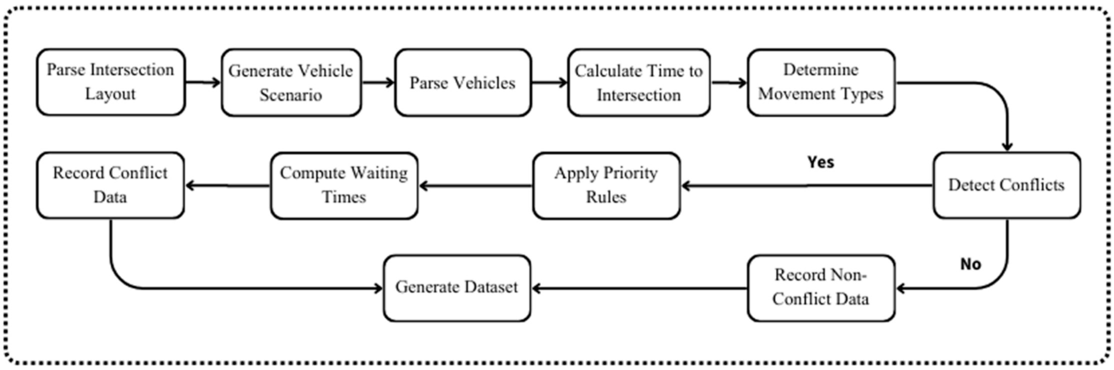

Figure 3. Workflow for dataset creation system.

To transform the generated scenario data into a format suitable for the LLM, we employ JSON (JavaScript Object Notation) as an intermediate structure. Each vehicle's attributes—such as ID, speed, direction, and lane—are stored as key–value pairs. These JSON records are then parsed by our prompt generation script, which uses chain-of-thought prompting to guide the LLM in conflict detection and resolution. Specifically, each JSON field is converted into a compact textual description (e.g., 'Vehicle V101 in lane 2, speed 40 km/h'), which the LLM interprets to identify conflicts and propose decisions.

First, the intersection layout is parsed to obtain information concerning individual lanes and possible destinations from any given lane. Vehicle movement analysis is made more accessible and facilitated by storing the layout in a structured JSON format. Parsing helps in understanding how the intersection is arranged, which is essential when it comes to determining where vehicles will pass and where there may be danger of accidents occurring. We assumed the initial testing to be carried out on a multi-lane four-legged urban intersection, where vehicles can travel in different directions, as shown in Figure 4. Each of its legs has lanes marked for specific purposes, including moving forward only and taking left- or rightward turns. This allows traffic from all directions across the 'cross' section, with left and right turns facilitated by specific lanes to avoid any conflicts with traffic. As shown in Figure 4, each approach leg is marked with the corresponding lane arrows to accurately reflect valid turning movements.

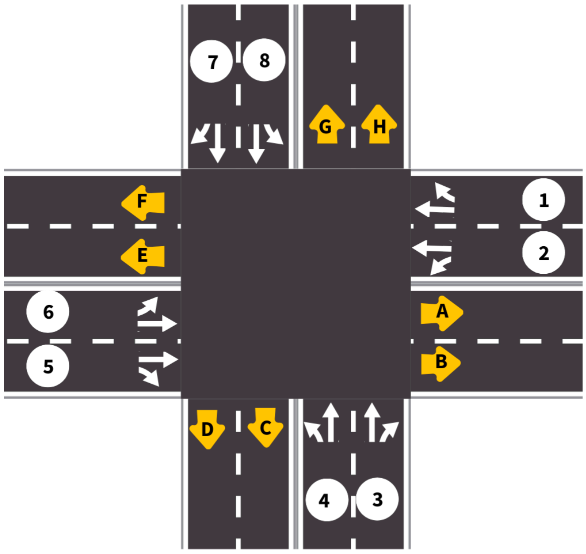

Figure 4. Multi-lane four-leg intersection showing accurate lane markings for each permissible maneuver (left turn, right turn, straight).

To handle shared lanes, each vehicle's intended maneuver (straight, left turn, right turn) is assigned when the vehicle is still between 80 and 100 m away from the intersection. The system continuously calculates each vehicle's time-to-intersection (TTI). If two vehicles in the same lane have conflicting maneuvers, the conflict detection module flags a potential issue. This advanced detection ensures adequate space and time to prompt drivers or vehicles to yield, slow down, or switch lanes (where applicable) before reaching the stop line.

For all simulations, we used a discrete time-step environment coded in Python, where each time step corresponds to 0.5 s. Vehicle positions and speeds are updated at each step, with collision checks handled via lane-based trajectory overlaps. Signal timing adjustments or driver instructions (yield, slow, and stop) are executed instantaneously at the start of each time step. This environment ensures a consistent and controlled platform for generating and testing diverse intersection scenarios.

Table 1 shows the different attributes that were used to define the environment to the LLM. Each instance of the vehicle class represents each vehicle approaching the intersection. The methods in the vehicle class are used for validating input, calculating time, and determining movement. Input validation is responsible for checking the correctness of all vehicle data attributes, including non-negative speed, distance greater than zero, and valid directions. The time calculation method is meant to determine how much longer it will take this vehicle to reach the intersection, given its current rate of travel as well as how far away it is from there. Consequently, movement determination stands here to classify the vehicle movement into certain types in advance, such as at intersections, lanes, and destinations.

Table 1. Vehicle attributes defined for the LLM.

Attribute	Description
Vehicle_ID	Unique identifier for the vehicle
Lane	The lane number the vehicle is in
Speed	Vehicle's speed in km/h
Distance_to_Intersection	Distance from the vehicle to the intersection (in meters)
Direction	Direction of approach (north, east, south, and west)
Destination (Egress)	Intended exit lane number from the intersection

A vehicle is parsed by reading the JSON scenario data of the vehicle and forming an object for each vehicle. To avoid duplicating Vehicle IDs, the system confirms them while ensuring that all requested information is correct. This turns raw scenario data into organized objects, which can be subject to further analysis.

4.2. Prompt Design

To properly navigate the frame and utilize LLM features, we created prompts utilizing the chain-of-thoughts method. These make it possible for the LLM to undertake a situational analysis, identify probable conflicts and give advice. The prompts were structured to have details on a traffic situation, vehicle conditions and potential conflicts, hence enabling the LLM to make decisions concerning the case. Table 2 illustrates a sample prompt that was applied. This prompt allows the LLM to have background information about traffic and give specific location-based guidance on managing the intersection best.

Table 2. An example of a used prompt.

You are an Urban Intersection Traffic Conflict Detector, responsible for monitoring a four-way intersection with traffic coming from the north, east, south, and west. Each direction has two lanes guiding vehicles to different destinations: - North: Lane 1 directs vehicles to F and H, Lane 2 directs vehicles to E, D, and C. - East: Lane 3 leads to H and B, Lane 4 leads to G, E, and F. - South: Lane 5 directs vehicles to B and D, Lane 6 directs vehicles to A, G, and H. - West: Lane 7 directs vehicles to D and F, Lane 8 directs vehicles to B, C, and A. Analyze the traffic data from all directions and lanes, and determine if there is a potential conflict between vehicles at the intersection. Respond only with 'Yes' or 'No' for conflict detection. Output: - If there **is a conflict**, provide a report with the following structure: **Conflict Status**: State whether a conflict is detected (e.g., "Conflict detected."). **Conflicts Overview**: Mention the number of conflicts and any vehicles involved (e.g., "Number of conflicts: 1. Involved vehicles: Vehicle V1234 and Vehicle V5678."). **Actions & Decisions**: Summarize any key decisions or actions taken (e.g., "Decisions: Vehicle V5678 must yield to Vehicle V1234."). **Priority Assignment**: List the vehicles and their assigned priorities (e.g., "Vehicle V1234: Priority 1, Vehicle V5678: Priority 2."). **Vehicle Waiting Times**: Provide waiting times for each vehicle (e.g., "Vehicle V1234: 5 s, Vehicle V5678: 10 s."). **The output format must exactly follow this structure in case of conflict:**

4.3. Conflict Detection

The creation of vehicle scenarios requires careful consideration of conflict detection, which involves several successive steps. First, the system analyzes whether a particular type of movement may cause intersecting paths between vehicles. For example, vehicles moving in opposite directions without turning may not cross each other's paths at all. Then, it assesses whether any vehicles are likely to arrive at the intersection within a close time interval, potentially creating conflicts. To handle these identified conflicts, the system applies traffic priority rules to determine if a vehicle should yield. Factors such as arrival times and relative positions are considered when developing these rules based on movement types—left, straight, or right. The system also calculates waiting times at the junction for lower-priority vehicles, as illustrated in Figure 4, where conflicts and priority decisions, including waiting durations before movement, are captured.

The process begins with system initialization, where intersection layout data—essential for understanding possible movements—are read and organized. Then, a random vehicle scenario with various types of vehicles is generated, and a vehicle object is created for each vehicle. Predefined traffic rules identify conflicts and specify which vehicles should yield, but the LLM goes beyond simple rule application. After evaluating initial conditions, the LLM reviews each detected conflict in real time, providing contextual explanations. It considers factors such as vehicle positions, speeds, and predicted arrival times to recommend specific actions. By offering context-aware driver guidance on yielding decisions and speed adjustments, the LLM enhances traditional rule-based systems, making the interaction more adaptable and human-centered. In cases with no conflicts, scenario data are retained for completeness, while a textual description is generated for human understanding. The procedure continues until all steps are complete, with waiting durations determined for yielding vehicles and conflicts, decisions, and wait times documented.

The system also utilizes standard road traffic priority rules to establish the right of way between conflicting vehicles. Vehicles intending to go straight generally have priority over those planning to turn. For instance, if a vehicle wants to make a left turn, it must wait until all vehicles turning right have cleared the intersection. If two vehicles arrive at the same time with no clear priority, the right-hand rule is applied, where the vehicle on the right has precedence. These rules are essential for resolving conflicts in alignment with common traffic regulations.

As an example, consider a scenario where four vehicles are approaching an intersection in Figure 5. Vehicle V217 is coming from the north in lane 2 at 40 km/h, 80 m from the intersection, heading toward destination E. Vehicle V218 is approaching from the east in lane 4 at 40 km/h, also 80 m away, heading toward destination G. Vehicle V219 is approaching from the south in lane 6 at 40 km/h, 80 m away, heading toward destination A. Lastly, vehicle V220 is approaching from the west in lane 8 at 40 km/h, 80 m away, heading toward destination C. In this scenario, the system creates vehicle objects for V217, V218, V219, and V220, computing each vehicle's time to reach the intersection. Based on the lanes and destinations, it determines that all vehicles are likely making left turns, leading to potential conflicts as they approach simultaneously. Using traffic rules, the system decides which vehicles must yield. For instance, vehicle V217 may yield to vehicle V218 if the latter is coming from the right. Waiting times are calculated for each yielding vehicle, and conflict data, including decisions and waiting durations, are recorded and described for human interpretation.

However, in order to make scenarios more interpretable, a system is used to transform them into human-readable textual descriptions. It takes scenario data and pulls out key details concerning every single vehicle. For each vehicle, explanations specify the lane, destination point, speed, distance from the intersection or crossed path, and approach direction. Such descriptions are then aggregated into a coherent whole, which outlines what went down in an entire scene. This is illustrated in Table 3.

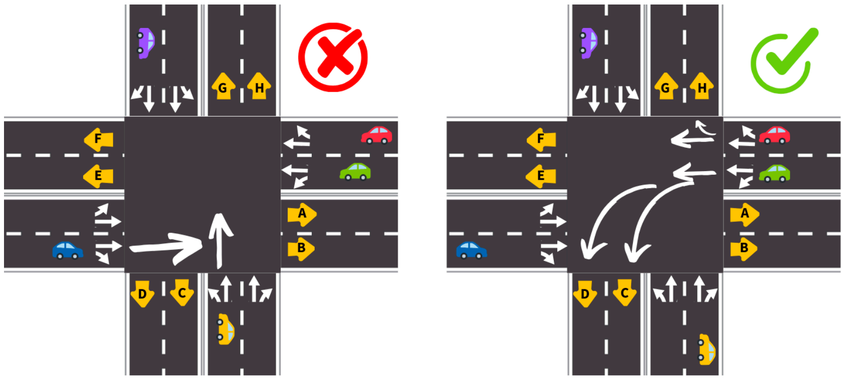

Figure 5. Comparison of conflict (**left**) and non-conflict (**right**) of vehicles.

Table 3. An example of parsing vehicle data.

JSON Input	**Input**

{
"vehicles_scenario": [
{
"vehicle_id": "V1151",
"lane": "2",
"speed": 39.995323464891484,
"distance_to_intersection": 388.95660041889687,
"direction": "north",
"destination": "C"
},
{
"vehicle_id": "V5173",
"lane": "8",
"speed": 68.0915930855088,
"distance_to_intersection": 150.82949998592466,
"direction": "west",
"destination": "B"
},
{
"vehicle_id": "V8617",
"lane": "1",
"speed": 43.411746756299856,
"distance_to_intersection": 180.7639436593828,
"direction": "north",
"destination": "F"
},
{
"vehicle_id": "V2618",
"lane": "4",
"speed": 63.24744202519462,
"distance_to_intersection": 366.3390574707282,
"direction": "east",
"destination": "F"
}
]
}
	Vehicle V7155 is in lane 2, moving north at a speed of 30.86 km/h, and is 88.54 m away from the intersection, heading towards D. Vehicle V6439 is in lane 3, moving east at a speed of 53.37 km/h, and is 107.50 m away from the intersection, heading towards B. Vehicle V5182 is in lane 7, moving west at a speed of 47.69 km/h, and is 94.67 m away from the intersection, heading towards D. Vehicle V2432 is in lane 1, moving north at a speed of 46.17 km/h, and is 74.59 m away from the intersection, heading towards H.

Table 3. *Cont.*

JSON Output
{
“is_conflict”: “yes”,
“number_of_conflicts”: 4,
“places_of_conflicts”: [
“intersection”,
“intersection”,
“intersection”,
“intersection”
],
“conflict_vehicles”: [
{
“vehicle1_id”: “V7155”,
“vehicle2_id”: “V6439”
},
{
“vehicle1_id”: “V7155”,
“vehicle2_id”: “V5182”
},
{
“vehicle1_id”: “V6439”,
“vehicle2_id”: “V2432”
},
{
“vehicle1_id”: “V5182”,
“vehicle2_id”: “V2432”
}
],
“decisions”: [
“Potential conflict: Vehicle V7155 must yield to Vehicle V6439”,
“Potential conflict: Vehicle V7155 must yield to Vehicle V5182”,
“Potential conflict: Vehicle V6439 must yield to Vehicle V2432”,
“Potential conflict: Vehicle V5182 must yield to Vehicle V2432”
],
“priority_order”: {
“V2432”: 1,
“V5182”: 2,
“V6439”: 3,
“V7155”: 4
},
“waiting_times”: {
“V2432”: 0,
“V5182”: 1,
“V6439”: 3,
“V7155”: 2
}
}
	Output
	Conflict Status: Conflict detected.
	Conflicts Overview: Number of conflicts: 4. Involved vehicles: Vehicle V7155 and Vehicle V6439, Vehicle V7155 and Vehicle V5182, Vehicle V6439 and Vehicle V2432, Vehicle V5182 and Vehicle V2432.
	Actions & Decisions: Decisions: Potential conflict: Vehicle V7155 must yield to Vehicle V6439, Potential conflict: Vehicle V7155 must yield to Vehicle V5182, Potential conflict: Vehicle V6439 must yield to Vehicle V2432, Potential conflict: Vehicle V5182 must yield to Vehicle V2432
	Priority Assignment: Vehicle V2432: Priority 1, Vehicle V5182: Priority 2, Vehicle V6439: Priority 3, Vehicle V7155: Priority 4.
	Vehicle Waiting Times:
	- Vehicle V2432: 0 s
	- Vehicle V5182: 1 s
	- Vehicle V6439: 3 s
	- Vehicle V7155: 2 s

4.4. Model Selection and Fine-Tuning

We tested our novel paradigm (i.e., using LLMs as a traffic controller) on three different models, including GPT-mini, Gemini, and Llama. They were chosen for their ability to solve natural language problems and their potential to understand traffic situations in real time under changing environments. We tested the models using two learning methods, including zero-shot and fine-tuning. This allowed us to analyze their effectiveness in different types of tasks.

The GPT-model mini has been scaled down for it to process faster than the bigger ones and was fine-tuned on three datasets, the first one being a four-vehicle dataset, while the second contained eight groups of vehicle data, such as mixed data with examples ranging between two and eight cars in each scenario. This mixed dataset was created based on the assumption that performance would improve over a wide range of traffic complexities.

The Gemini models were also tested. Like GPT-mini, they are transformer-based, but their training methods and design differ. This study used two versions of the Gemini model: Gemini 1.0, evaluated in both fine-tuned and zero-shot conditions, and Gemini 1.5, a more advanced version with improved capabilities for scenario understanding and conflict resolution. Additionally, we used two versions of the Llama model, Llama-3.1-8B and Llama-3.1-70B.

For fine-tuned models, the training and validation datasets were used to customize the models to the specific task of conflict detection in urban intersections. On the other hand, zero-shot models were not fine-tuned for the specific task and were evaluated directly on the testing dataset. This approach tested the models' ability to generalize to new tasks without specialized training, which offers insights into their base knowledge in urban traffic control systems.

We examined the performance of these models across both fine-tuned and zero-shot settings. We aimed to determine how effectively LLMs can adapt to real-world traffic management tasks. Thus, each LLM model was analyzed using classification metrics such as accuracy, precision, recall, and F1-score, alongside their ability to generate detailed analyses of traffic scenarios, measured through ROUGE-L scores. We tested the outputs of LLMs using a simulation model that was built on Python.

5. Analysis and Results

5.1. Logical Reasoning Results

In this section, we present the results of using LLMs to identify conflicts in intersection scenarios. Specifically, an evaluation was conducted on different versions of GPT-4o-mini, Gemini, and Llama under both fine-tuned and zero-shot conditions. It must be noted that the datasets were designed to consist of varying numbers of cars to evaluate how well these models generalize within different traffic conditions.

Depending on the number of vehicles involved, datasets were divided. The four-vehicle dataset involves situations with just four vehicles on the road intersection, while the eight-vehicle dataset is all about having eight cars in similar spots. We also used another mixed-vehicle dataset, in which different numbers of vehicle-based scenarios ranging between two and eight were fed to the LLM. They are organized into three groups of 10,000 scenarios in total, including 7000 that were catered for training purposes, 2000 available for testing, and 1000 reserved for validation purposes. Zero-shot evaluation was only performed using test datasets. Test datasets were used for assessment of fine-tuned models, and training and validation were used for fine-tuning as well. The performance of the LLM in detecting conflicts at the defined intersection is summarized in Table 4.

The GPT-mini fine-tuned model achieved the highest accuracy of 83% on the mixed-vehicle scenario, indicating its superior ability to generalize across scenarios with varying numbers of vehicles. Fine-tuning on a diverse dataset likely contributed to its robust performance. The high recall of 0.85 suggests that the model effectively identifies actual conflict scenarios, which is critical in applications where overlooking a conflict can have severe consequences. The GPT-4o-mini four-vehicle fine-tuned model also performed well, with an accuracy of 81%. This demonstrates that even when trained on scenarios with a fixed number of vehicles, the model can achieve high performance, although slightly lower than the mixed dataset counterpart. The reduction in accuracy compared to the mixed model may be due to overfitting to scenarios with only four vehicles, limiting its ability to generalize to more varied situations. The GPT-4o-mini eight-vehicle fine-tuned model's accuracy dropped to 71%, indicating that scenarios with a higher number of vehicles pose a greater challenge.

Table 4. Performance of the LLM in detecting conflicts. Results for fine-tuned models are highlighted. The highest performance of each scenario across all models are in bold.

Model	Learning Method	Scenario	Accuracy	Precision	Recall	F1-Score
GPT-4o-mini	fine-tuning	mixed-vehicle	**83.0**	**0.83**	**0.85**	**0.84**
	fine-tuning	4-vehicle	**81.0**	**0.80**	**0.83**	**0.82**
	fine-tuning	8-vehicle	**71.0**	**0.70**	**0.74**	**0.72**
	zero-shot	mixed-vehicle	61.9	0.59	0.62	0.60
	zero-shot	4-vehicle	53.6	0.53	0.54	0.53
	zero-shot	8-vehicle	50.8	0.50	0.51	0.42
Gemini 1.5	zero-shot	8-vehicle	52.8	0.55	0.77	0.62
	fine-tuning	4-vehicle	51.5	0.51	0.48	0.50
	zero-shot	4-vehicle	50.9	0.51	0.54	0.52
	fine-tuning	8-vehicle	49.1	0.49	0.50	0.49
	fine-tuning	mixed-vehicle	48.2	0.49	0.72	0.58
	zero-shot	mixed-vehicle	61.4	0.62	0.60	0.61
Gemini 1.0	fine-tuning	mixed-vehicle	60.9	0.61	0.62	0.61
	zero-shot	mixed-vehicle	55.1	0.57	0.42	0.49
	fine-tuning	4-vehicle	52.9	0.55	0.73	0.61
	fine-tuning	8-vehicle	50.8	0.60	0.05	0.09
	zero-shot	4-vehicle	50.4	0.50	0.58	0.54
	zero-shot	8-vehicle	48.2	0.47	0.48	0.45
Llama-3.1-8B-Instruct	zero-shot	mixed-vehicle	50.4	0.52	0.50	0.37
	fine-tuning	4-vehicle	49.6	0.49	0.50	0.43
Llama-3.1-70B-Instruct	fine-tuning	4-vehicle	51.5	0.51	0.51	0.51

The complexity introduced by additional vehicles may require more sophisticated modeling or larger datasets to capture the nuances of interactions in such scenarios. Zero-shot models generally underperformed compared to their fine-tuned counterparts. The Mixed GPT-mini Zero-Shot model achieved an accuracy of 61.9%, highlighting the importance of fine-tuning for specialized tasks. Without task-specific training, models struggle to capture the intricacies of conflict detection in intersection scenarios. The Gemini models consistently exhibited lower performance across both fine-tuned and zero-shot settings. The highest accuracy among them was 61.4% for the Gemini 1.5 Zero-Shot Mixed model. This suggests that the Gemini architecture may not be as well suited for this classification task or that the fine-tuning process was less effective compared to GPT-mini models. The Meta-Llama models, particularly the Meta-Llama-3.1-70B fine-tuned on four vehicles, showed moderate improvements with fine-tuning, achieving a balanced F1-score of 0.51 and an accuracy of 51.5%. This indicates some potential for scalability with larger parameter sizes; however, they still trail behind the GPT-mini models, likely due to a lack of optimization for complex traffic conflict detection. The comparison underscores the effectiveness of fine-tuning on diverse datasets and highlights the suitability of the GPT-mini architecture for nuanced traffic scenarios, as shown in Figure 6.

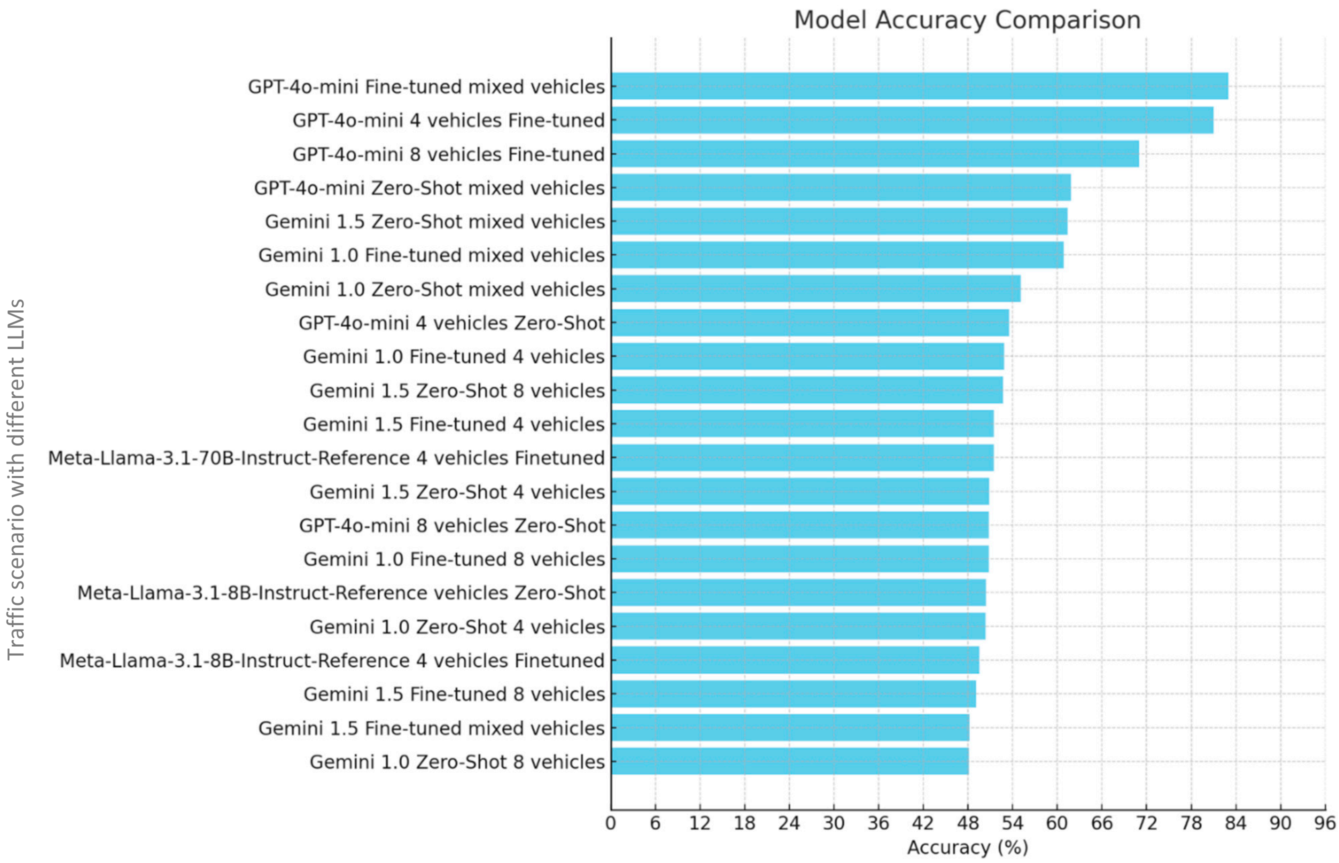

Figure 6. Accuracy comparison across various LLMs, learning methods, and vehicle scenarios.

To understand the performance of the fine-tuned GPT-mini models, confusion matrices were analyzed for the best-performing models. Figure 7 presents the confusion matrices for the GPT-4o-mini models in different scenarios. The mixed-vehicle fine-tuned GPT-mini model exhibits the best performance among the three, with a higher number of true positives (820) and true negatives (849), and lower false negatives (180) and false positives (151). This indicates that the model is proficient at correctly identifying both conflict and non-conflict scenarios. Comparing the fine-tuned GPT-4o-mini model for four vehicles to the mixed model, we observe slightly lower true positives (797) and true negatives (830), and higher false negatives (203) and false positives (170). This suggests that the model trained exclusively on four-vehicle scenarios may not generalize as effectively as the mixed model, leading to a higher rate of misclassification. The eight-vehicle fine-tuned GPT-mini model shows a decrease in performance compared to the mixed and four-vehicle models, with a higher number of false negatives (324) and false positives (255). This reinforces the earlier observation that scenarios with more vehicles introduce complexity that challenges the model's ability to accurately classify conflict scenarios.

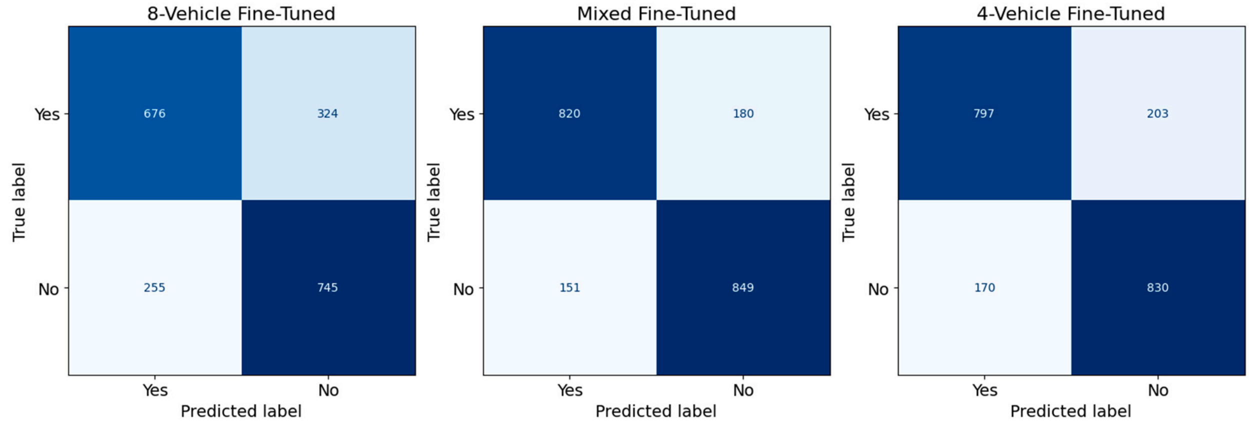

Figure 7. Confusion matrix for fine-tuned GPT-4 Mini across different vehicle scenarios.

5.2. Decision-Making and Feedback Results

In this section, we tested the decision making and feedback that the fine-tuned GPT-mini model provided to the drivers in the four-vehicle and mixed-vehicle scenarios. An evaluation was conducted using ROUGE-L score based on four measurements, including identifying conflicts, decisions, priority assignments, and waiting times. The truth values were found using Python simulation and were validated by traffic engineering experts. Firstly, identifying conflicts involves detecting potential points where vehicles' paths may intersect or where their arrival times may lead to a collision risk at an intersection. Secondly, decisions refer to determining actions for each vehicle in a conflict scenario, such as whether a vehicle should yield, proceed, or adjust its speed to avoid collision. Third, priority assignments involve assigning right-of-way to vehicles based on traffic rules, such as giving priority to vehicles going straight over those making turns or applying the right-hand rule. Finally, waiting times measure the duration a vehicle must remain at an intersection before it can proceed, often due to lower priority or the need to yield to conflicting traffic.

Figure 8 presents the ROUGE-L scores for various scenarios. In the four-vehicle scenarios, the fine-tuned GPT-mini model achieved high ROUGE-L scores across all components, with the conflicts overview scoring 0.95. This indicates that the model can accurately and coherently summarize conflicts in simpler scenarios. The decisions component scored 0.91, demonstrating the model's ability to interpret and explain the decisions made based on traffic rules. The high scores in priority assignment (0.94) and vehicle waiting times (0.92) further showcase the model's proficiency in assigning priorities and calculating waiting times for vehicles. For mixed-vehicle scenarios, the ROUGE scores are slightly lower but still strong. The conflicts overview scored 0.84, and the decisions component scored 0.81. These results suggest that as the complexity of the scenarios increases with varying numbers of vehicles, the model faces more challenges in summarizing conflicts accurately. However, the model maintains respectable performance, indicating its ability to handle complex scenarios effectively. The priority assignment and vehicle waiting times components scored 0.87 and 0.86, respectively, reflecting the model's competence in applying traffic rules and calculating necessary delays even in more complex situations. The consistently high scores in the decisions component across both scenario types underscore the model's strong grasp of traffic rules and its ability to apply them in diverse situations. The slight decline in scores for mixed-vehicle scenarios highlights the increased complexity but also shows that the model adapts reasonably well. These results demonstrate that the fine-tuned GPT-mini model excels not only in binary classification of conflict scenarios but also in generating detailed, coherent analyses. The model's ability to maintain high performance across key components, even in complex scenarios, is promising for practical applications where detailed explanations are necessary.

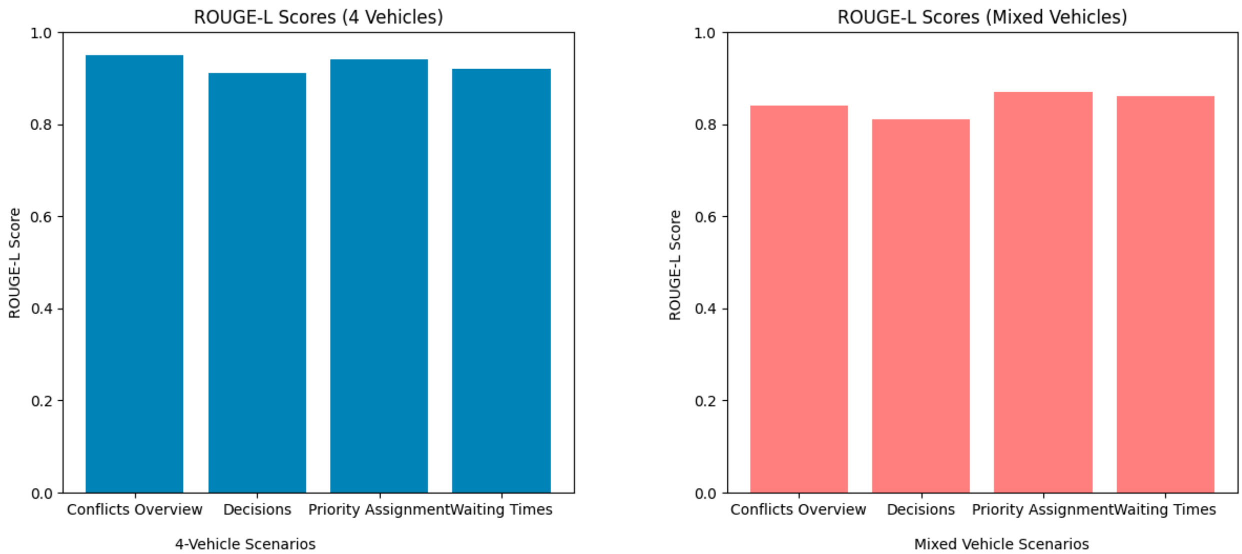

Figure 8. ROUGE score based on identifying conflicts, decisions, priority assignments, and waiting times.

The high ROUGE scores in critical components indicate that the model can effectively support traffic management tasks by providing accurate conflict detection and actionable decisions. In real-world applications, such capabilities can enhance traffic flow efficiency and safety at intersections by enabling proactive measures to prevent accidents and reduce congestion.

The fine-tuned GPT model's outputs in Table 5 showcase its capability to identify potential conflicts and provide precise recommendations for vehicles approaching an intersection. In scenarios where conflicts are detected, such as between vehicles V4625 and V1909 or vehicles V7019 and V5264, the model suggests actionable measures like yielding or slowing down based on the vehicles' distances, speeds, and intended paths, prioritizing safety and minimizing collision risks. For instance, it recommends that vehicle V1909 yield to vehicle V4625 due to the latter's advantage in distance and suggests that vehicle V7019 yield to vehicle V5264 to avoid intersection congestion. In contrast, the model advises all vehicles to continue without adjustment for conflict-free situations, accurately discerning safe scenarios. This approach extends to more complex cases, such as between northbound and westbound vehicles V9867 and V6070, where the model proposes a slight yield from V9867 to ensure safe passage, demonstrating its nuanced understanding of cooperative traffic flow. Combining conflict detection with tailored, real-time guidance, the model displays a structured and responsive framework for intersection management, supporting automated safety and efficiency in multi-vehicle environments.

Table 5. Examples of outputs of LLMs' decision making and feedback to the drivers.

Scenario Input	Fine-Tuned GPT-Mini Output
Vehicle V7406 is in lane 7, moving west at a speed of 27.59 km/h, and is 182.12 m away from the intersection, heading towards D. Vehicle V4625 is in lane 3, moving east at a speed of 49.97 km/h, and is 189.12 m away from the intersection, heading towards B. Vehicle V1909 is in lane 8, moving west at a speed of 75.90 km/h, and is 263.27 m away from the intersection, heading towards C. Vehicle V2327 is in lane 8, moving west at a speed of 26.36 km/h, and is 445.70 m away from the intersection, heading towards A.	**Conflict Status**: Yes **Conflict Analysis**: A conflict is detected between Vehicle V4625 and Vehicle V1909. Both vehicles are approaching the same intersection from opposite directions, and they are likely to reach the intersection at the same time due to their respective speeds and distances. **Recommendations**: To prevent a potential collision, it is recommended that Vehicle V1909 yield to Vehicle V4625, as V4625 has a slight advantage in distance. **Next Actions for Each Vehicle**: - Vehicle V4625: Continue straight towards the intersection - Vehicle V1909: Yield and prepare to adjust speed or trajectory if necessary
Vehicle V1975 is in lane 2, moving north at a speed of 61.28 km/h, and is 431.49 m away from the intersection, heading towards D. Vehicle V7019 is in lane 4, moving east at a speed of 50.89 km/h, and is 195.12 m away from the intersection, heading towards E. Vehicle V5264 is in lane 6, moving south at a speed of 73.17 km/h, and is 207.91 m away from the intersection, heading towards G. Vehicle V8370 is in lane 3, moving east at a speed of 25.37 km/h, and is 393.16 m away from the intersection, heading towards H.	**Conflict Status**: Yes **Conflict Analysis**: A potential conflict is detected between vehicles V7019 and V5264. Vehicle V7019 is approaching the intersection, while Vehicle V5264 is already entering the intersection; both are on a collision course. **Recommendations**: It is recommended that Vehicle V7019 yield to Vehicle V5264 to avoid conflict. **Next Actions for Each Vehicle**: - Vehicle V1975: No action - Vehicle V7019: Slow down and yield - Vehicle V5264: Continue at current speed - Vehicle V8370: No action
Vehicle V5009 is in lane 6, moving south at a speed of 45.37 km/h, and is 217.81 m away from the intersection, heading towards H. Vehicle V3191 is in lane 6, moving south at a speed of 68.86 km/h, and is 298.54 m away from the intersection, heading towards H. Vehicle V6111 is in lane 4, moving east at a speed of 78.52 km/h, and is 106.46 m away from the intersection, heading towards E. Vehicle V7721 is in lane 7, moving west at a speed of 44.52 km/h, and is 309.51 m away from the intersection, heading towards F.	**Conflict Status**: No **Conflict Analysis**: No conflicts detected among the vehicles. **Recommendations**: None **Next Actions for Each Vehicle**: - V5009: Continue moving south - V3191: Continue moving south - V6111: Continue moving east - V7721: Continue moving west
Vehicle V9867 is in lane 2, moving north at a speed of 65.73 km/h, and is 270.95 m away from the intersection, heading towards C. Vehicle V6094 is in lane 6, moving south at a speed of 61.93 km/h, and is 344.02 m away from the intersection, heading towards G. Vehicle V6070 is in lane 7, moving west at a speed of 59.35 km/h, and is 296.58 m away from the intersection, heading towards F. Vehicle V2692 is in lane 2, moving north at a speed of 46.17 km/h, and is 118.97 m away from the intersection, heading towards E.	**Conflict Status**: Yes **Conflict Analysis**: There is a conflict between Vehicle V9867 (northbound) and Vehicle V6070 (westbound). V9867 is approaching the intersection from the north at a speed of 65.73 km/h and is positioned 270.95 m away, while V6070 is approaching from the west at a speed of 59.35 km/h and has a distance of 296.58 m to the intersection. The distance to the intersection for both vehicles is similar, which can lead to an intersection conflict. **Recommendations**: The drivers of both vehicles should communicate well and gauge each other's positions. It may be best for Vehicle V9867 to yield slightly to allow for safer navigation through the intersection. **Next Actions for Each Vehicle**: - Vehicle V9867: Yield slightly, if necessary, to allow V6070 to proceed safely through the intersection. - Vehicle V6070: Proceed without delay, being aware of V9867's proximity. - Vehicle V6094: No action needed concerning the conflict. - Vehicle V2692: No action needed concerning the conflict.
Vehicle V6758 is in lane 4, moving east at a speed of 61.16 km/h, and is 413.33 m away from the intersection, heading towards F. Vehicle V5284 is in lane 2, moving north at a speed of 59.44 km/h, and is 248.61 m away from the intersection, heading towards D. Vehicle V6460 is in lane 5, moving south at a speed of 73.17 km/h, and is 363.19 m away from the intersection, heading towards B. Vehicle V9110 is in lane 7, moving west at a speed of 20.11 km/h, and is 65.18 m away from the intersection, heading towards D.	**Conflict Status**: Yes **Conflict Analysis**: A conflict exists between Vehicle V5284 and Vehicle V6460. They approach the intersection from different directions and will reach it at approximately the same moment, creating a potential collision risk. **Recommendations**: To resolve this conflict, Vehicle V6460 should yield to Vehicle V5284, allowing the smaller vehicle to navigate past the intersection first. **Next Actions for Each Vehicle**: - Vehicle V5284: Continue to the intersection as planned - Vehicle V6460: Slow down and prepare to yield to Vehicle V5284 - Vehicle V6758: No action - Vehicle V9110: No action

6. Conclusions

This study presents a new paradigm of traffic control systems, proposing LLMs as dynamic traffic controllers. Leveraging their logical reasoning, scene understanding, and decision-making abilities, LLMs can adaptively respond to real-time traffic conditions and optimize throughput while considering safety, reducing congestion and environmental impact. Through the 4D traffic control system framework (i.e., Detect, Decide, Disseminate, and Deploy), LLMs centralize processes traditionally managed by disconnected components. LLMs integrate diverse data sources (e.g., loop detectors, GPS, video imaging) to make context-aware decisions and provide tailored control outputs and feedback using diverse modules via various communication channels (e.g., voice, wireless signals, code, and visuals) to drivers, infrastructures, and autonomous vehicles. With their ability to handle multi-objective optimization functions, LLMs can transform urban mobility. We argued in this study that as LLMs offer adaptability in communication, interfacing with humans and machines, this will potentially ensure compliance and safety in complex traffic scenarios.

We presented a methodology to integrate and evaluate LLMs as a controller into real-time traffic control systems, which comprises four key stages, including data creation and initialization, prompt generation, conflict identification, and fine-tuning with model analysis. Initially, we proposed a system that generated datasets reflecting real-world multi-lane four-leg urban intersection scenarios. We also parsed detailed intersection layouts, lane configurations, and vehicle attributes such as speed, direction, and destination of each vehicle. These datasets are organized into vehicle objects, enabling scenario simulation and conflict detection to find the truth values. Prompts were designed using the chain-of-thought method to guide the LLM in analyzing intersection conditions, detect potential conflicts, and suggest priority-based traffic decisions. Conflict resolution contains traditional traffic rules such as the right-hand rule and yielding priorities. These conflict resolutions were augmented by LLMs' context-aware analyses of vehicle positions, speeds, and arrival times. This multi-stage approach not only enhances conflict detection but also offers adaptive, context-sensitive traffic management solutions that align with real-world regulations. We tested three different LLM models including GPT-mini, Gemini, and Llama. We evaluated the results using accuracy, precision, recall, F1-score, and ROUGE-L.

Results demonstrated that LLMs (specifically, a fine-tuned GPT-4o-mini model) have a significantly high ability to identify conflicts and support decision making in traffic intersection scenarios. It is also worth noting that the Gemini and Llama models showed potential for improvement with fine-tuning; however, their performance remained lower than that of the GPT-4o-mini models. This might be due to architectural limitations or less effective optimization for this specific task. Nonetheless, the fine-tuned GPT-4o-mini achieved a high accuracy of about 83% and an F1-score of 0.84 in mixed-vehicle scenarios, which reflects its ability to generalize across diverse traffic conditions. This performance is further supported by high recall and precision rates of about 0.85 and 0.83, respectively. This shows that LLMs can be reliable in detecting actual conflicts, a critical aspect for real-world intersection control applications.

Results also showed the fine-tuned GPT-4o-mini model excelled in generating detailed, actionable insights for decision making and traffic management. With the help of traffic experts, the LLM model achieved high ROUGE-L scores across the four key components, including conflict identification of about 0.95, decision making of about 0.91, priority assignment of about 0.94, and waiting times of about 0.92. These results showed that LLMs can be practical and applicable in providing coherent and precise recommendations such as yielding, slowing down, and stopping based on vehicles' speeds, distances, and paths and considering the other vehicles approaching the intersection.

The findings of this study demonstrated LLMs' capabilities and potential to transform traffic control systems. It can enhance traffic flow efficiency and safety at intersections. Its ability to handle both simple and complex scenarios makes it a promising tool for real-world traffic management systems in real time and for providing feedback to drivers, infrastructures, and autonomous vehicles. However, future work will focus on further improving performance in high-density traffic scenarios and integrating the model into real-time traffic management applications. In future work, we will compare the LLM-generated traffic flow and conflict predictions to real-world intersection data using GEH statistics, which was named after Geoffrey E. Havers. The GEH measure is widely used in traffic engineering to evaluate how closely model outputs match observed counts. By calculating the GEH for multiple intersections, we can quantitatively assess and refine the alignment of the LLM-based traffic control model with actual traffic conditions.

Despite these promising findings, several limitations warrant caution. The performance of LLM-based traffic control heavily relies on the availability of high-quality data for fine-tuning, which may be challenging to obtain in some regions. Additionally, computational and memory requirements for large-scale LLMs can be substantial, potentially limiting real-time deployment in resource-constrained environments. Moreover, unexpected traffic behaviors or non-traditional road layouts may require model retraining or specialized prompt engineering to ensure robust performance. Addressing these drawbacks will be crucial for widespread adoption of this paradigm.

Author Contributions: Conceptualization, S.M., H.I.A. and M.E.; methodology, S.M., H.I.A. and M.E.; software, S.M.; formal analysis, S.M.; investigation, S.M. and H.I.A.; resources, S.M.; data curation, S.M.; writing—original draft preparation, S.M.; writing—review and editing, H.I.A. and M.E.; visualization, S.M.; supervision, H.I.A. and M.E.; project administration, H.I.A. and M.E. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The data and code that support the findings of this study are openly available in Zenodo at https://doi.org/10.5281/zenodo.14171745, reference number 14171745.

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

References

1. Bucsuházy, K.; Matuchová, E.; Zůvala, R.; Moravcová, P.; Kostíková, M.; Mikulec, R. Human factors contributing to the road traffic accident occurrence. *Transp. Res. Procedia* **2020**, *45*, 555–561. [CrossRef]
2. Movahedi, M.; Choi, J. The Crossroads of LLM and Traffic Control: A Study on Large Language Models in Adaptive Traffic Signal Control. *IEEE Trans. Intell. Transp. Syst.* **2024**. [CrossRef]
3. Mahmud, D.; Hajmohamed, H.; Almentheri, S.; Alqaydi, S.; Aldhaheri, L.; Khalil, R.A.; Saeed, N. Integrating LLMs with ITS: Recent Advances, Potentials, Challenges, and Future Directions. *arXiv* **2025**, arXiv:2501.04437. [CrossRef]
4. Rakha, H.; Amer, A.; El-Shawarby, I. Modeling Driver Behavior within a Signalized Intersection Approach Decision–Dilemma Zone. *Transp. Res. Rec. J. Transp. Res. Board* **2008**, *2069*, 16–25. [CrossRef]
5. Macioszek, E.; Kurek, A. Road traffic distribution on public holidays and workdays on selected road transport network elements. *Transp. Probl.* **2021**, *16*, 127–138. [CrossRef]
6. Ashqer, M.I.; Ashqar, H.I.; Elhenawy, M.; Almannaa, M.; Aljamal, M.A.; Rakha, H.A.; Bikdash, M. Evaluating a signalized intersection performance using unmanned aerial Data. *Transp. Lett.* **2024**, *16*, 452–460. [CrossRef]
7. Distefano, N. Physiological and driving behaviour changes associated to different road intersections. *Eur. Transp. Eur.* **2020**, *77*, 4–16. [CrossRef]
8. Stevanovic, A.; Dobrota, N.; Mitrovic, N. *NCHRP 20-07/Task 414: Benefits of Adaptive Traffic Control Deployments—A Review of Evaluation Studies*; NCHRP: Washington, DC, USA, 2019.
9. Lai, S.; Xu, Z.; Zhang, W.; Liu, H.; Xiong, H. Large language models as traffic signal control agents: Capacity and opportunity. *arXiv* **2023**, arXiv:2312.16044.

10. Khanfar, N.O.; Elhenawy, M.; Ashqar, H.I.; Hussain, Q.; Alhajyaseen, W.K.M. Driving behavior classification at signalized intersections using vehicle kinematics: Application of unsupervised machine learning. *Int. J. Inj. Control Saf. Promot.* **2022**, *30*, 34–44. [CrossRef]
11. Yarlagadda, J.; Pawar, D.S. Heterogeneity in the Driver Behavior: An Exploratory Study Using Real-Time Driving Data. *J. Adv. Transp.* **2022**, *2022*, 4509071. [CrossRef]
12. Bella, F.; Silvestri, M. Interaction driver–bicyclist on rural roads: Effects of cross-sections and road geometric elements. *Accid. Anal. Prev.* **2017**, *102*, 191–201. [CrossRef] [PubMed]
13. Abu Hamad, F.; Hasiba, R.; Shahwan, D.; Ashqar, H.I. Driver Behavior at Roundabouts in Mixed Traffic: A Case Study Using Machine Learning. *J. Transp. Eng. Part A Syst.* **2024**, *150*, 05024004. [CrossRef]
14. Hadi, M.U.; Qureshi, R.; Shah, A.; Irfan, M.; Zafar, A.; Shaikh, M.B.; Akhtar, N.; Wu, J.; Mirjalili, S. A Survey on Large Language Models: Applications, Challenges, Limitations, and Practical Usage. *TechRxiv* **2023**. [CrossRef]
15. Mohammadnazar, A.; Patwary, A.L.; Moradloo, N.; Arvin, R.; Khattak, A.J. Incorporating driving volatility measures in safety performance functions: Improving safety at signalized intersections. *Accid. Anal. Prev.* **2022**, *178*, 106872. [CrossRef]
16. Ashqar, H.I.; Obaid, M.; Jaber, A.; Ashqar, R.; Khanfar, N.O.; Elhenawy, M. Incorporating driving behavior into vehicle fuel consumption prediction: Methodology development and testing. *Discov. Sustain.* **2024**, *5*, 344. [CrossRef]
17. Sha, H.; Mu, Y.; Jiang, Y.; Chen, L.; Xu, C.; Luo, P.; Li, S.E.; Tomizuka, M.; Zhan, W.; Ding, M. Languagempc: Large language models as decision makers for autonomous driving. *arXiv* **2023**, arXiv:2310.03026.
18. Tong, K.; Solmaz, S. ConnectGPT: Connect Large Language Models with Connected and Automated Vehicles. In Proceedings of the 2024 IEEE Intelligent Vehicles Symposium (IV), Jeju-si, Republic of Korea, 2–5 June 2024; IEEE: New York, NY, USA, 2024; pp. 581–588.
19. Zhang, Z.; Sun, Y.; Wang, Z.; Nie, Y.; Ma, X.; Sun, P.; Li, R. Large Language Models for Mobility in Transportation Systems: A Survey on Forecasting Tasks. *arXiv* **2024**, arXiv:2405.02357.
20. Hadi, M.U.; Qureshi, R.; Shah, A.; Irfan, M.; Zafar, A.; Shaikh, M.B.; Akhtar, N.; Wu, J.; Mirjalili, S. Large language models: A comprehensive survey of its applications, challenges, limitations, and future prospects. *TechRxiv* **2023**. [CrossRef]
21. Ashqar, H.I.; Jaber, A.; Alhadidi, T.I.; Elhenawy, M. Advancing Object Detection in Transportation with Multimodal Large Language Models (MLLMs): A Comprehensive Review and Empirical Testing. *arXiv* **2024**, arXiv:2409.18286.
22. Zheng, O.; Abdel-Aty, M.; Wang, D.; Wang, Z.; Ding, S. Chatgpt is on the horizon: Could a large language model be all we need for intelligent transportation? *arXiv* **2023**, arXiv:2303.05382.
23. Ren, Y.; Chen, Y.; Liu, S.; Wang, B.; Yu, H.; Cui, Z. TPLLM: A Traffic Prediction Framework Based on Pretrained Large Language Models. *arXiv* **2024**, arXiv:2403.02221.
24. Zhang, S.; Fu, D.; Liang, W.; Zhang, Z.; Yu, B.; Cai, P.; Yao, B. TrafficGPT: Viewing, processing and interacting with traffic foundation models. *Transp. Policy* **2024**, *150*, 95–105. [CrossRef]
25. Jain, S.; Thapa, S.; Chen, K.-T.; Abbott, A.L.; Sarkar, A. Semantic Understanding of Traffic Scenes with Large Vision Language Models. In Proceedings of the 2024 IEEE Intelligent Vehicles Symposium (IV), Jeju-si, Republic of Korea, 2–5 June 2024; IEEE: New York, NY, USA, 2024; pp. 1580–1587. [CrossRef]
26. Cui, C.; Ma, Y.; Cao, X.; Ye, W.; Wang, Z. Receive, reason, and react: Drive as you say, with large language models in autonomous vehicles. *IEEE Intell. Transp. Syst. Mag.* **2024**, *16*, 81–94. [CrossRef]
27. Ramanishka, V.; Chen, Y.-T.; Misu, T.; Saenko, K. Toward driving scene understanding: A dataset for learning driver behavior and causal reasoning. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 18–23 June 2018; pp. 7699–7707.
28. Elhenawy, M.; Abutahoun, A.; Alhadidi, T.I.; Jaber, A.; Ashqar, H.I.; Jaradat, S.; Abdelhay, A.; Glaser, S.; Rakotonirainy, A. Visual Reasoning and Multi-Agent Approach in Multimodal Large Language Models (MLLMs): Solving TSP and mTSP Combinatorial Challenges. *Mach. Learn. Knowl. Extr.* **2024**, *6*, 1894–1920. [CrossRef]
29. Jaradat, S.; Nayak, R.; Paz, A.; Ashqar, H.I.; Elhenawy, M. Multitask Learning for Crash Analysis: A Fine-Tuned LLM Framework Using Twitter Data. *Smart Cities* **2024**, *7*, 2422–2465. [CrossRef]
30. Tang, Y.; Dai, X.; Zhao, C.; Cheng, Q.; Lv, Y. Large Language Model-Driven Urban Traffic Signal Control. In Proceedings of the 2024 Australian & New Zealand Control Conference (ANZCC), Gold Coast, Australia, 1–2 February 2024; pp. 67–71. [CrossRef]
31. Zhang, D.; Zheng, H.; Yue, W.; Wang, X. Advancing ITS Applications with LLMs: A Survey on Traffic Management, Transportation Safety, and Autonomous Driving. In *International Joint Conference on Rough Sets*; Springer: Cham, Switzerland, 2024; pp. 295–309.
32. Pang, A.; Wang, M.; Pun, M.-O.; Chen, C.S.; Xiong, X. iLLM-TSC: Integration reinforcement learning and large language model for traffic signal control policy improvement. *arXiv* **2024**, arXiv:2407.06025.
33. Wang, M.; Pang, A.; Kan, Y.; Pun, M.-O.; Chen, C.S.; Huang, B. LLM-assisted light: Leveraging large language model capabilities for human-mimetic traffic signal control in complex urban environments. *arXiv* **2024**, arXiv:2403.08337.
34. Li, L.; Jiang, R.; He, Z.; Chen, X.M.; Zhou, X. Trajectory data-based traffic flow studies: A revisit. *Transp. Res. Part C Emerg. Technol.* **2020**, *114*, 225–240. [CrossRef]

35. Xu, Z.; Zhang, Y.; Xie, E.; Zhao, Z.; Guo, Y.; Wong, K.Y.K.; Li, Z.; Zhao, H. Drivegpt4: Interpretable end-to-end autonomous driving via large language model. *arXiv* **2023**, arXiv:2310.01412. [CrossRef]
36. Gordon, R.L.; Reiss, R.A.; Haenel, H.; Case, E.; French, R.L.; Mohaddes, A.; Wolcott, R. *Traffic Control Systems Handbook*; United States Federal Highway Administration Office of Technology Applications: Washington, DC, USA, 1996.
37. Day, C.M.; Langdon, S.; Stevanovic, A.; Tanaka, A.; Lee, K.; Smaglik, E.J.; Overn, L.; Agarwal, N.; Richardson, L.; Philips, S. Traffic signal systems research: Past, present, and future trends. *Centen. Pap.* **2019**. Available online: http://onlinepubs.trb.org/onlinepubs/centennial/papers/AHB25-Final.pdf (accessed on 1 November 2023).
38. Abu Tami, M.; Ashqar, H.I.; Elhenawy, M.; Glaser, S.; Rakotonirainy, A. Using Multimodal Large Language Models (MLLMs) for Automated Detection of Traffic Safety-Critical Events. *Vehicles* **2024**, *6*, 1571–1590. [CrossRef]
39. Zhou, X.; Liu, M.; Zagar, B.L.; Yurtsever, E.; Knoll, A.C. Vision language models in autonomous driving and intelligent transportation systems. *arXiv* **2023**, arXiv:2310.14414.
40. Ashqar, H.I.; Alhadidi, T.I.; Elhenawy, M.; Khanfar, N.O. Leveraging Multimodal Large Language Models (MLLMs) for Enhanced Object Detection and Scene Understanding in Thermal Images for Autonomous Driving Systems. *Automation* **2024**, *5*, 508–526. [CrossRef]
41. Yan, Z.; Zhou, H.; Tabassum, H.; Liu, X. Hybrid LLM-DDQN based Joint Optimization of V2I Communication and Autonomous Driving. *arXiv* **2024**, arXiv:2410.08854.
42. Desai, B.; Patil, K. Secure and Scalable Multi-Modal Vehicle Systems: A Cloud-Based Framework for Real-Time LLM-Driven Interactions. *Innov. Comput. Sci. J.* **2023**, *9*, 1–11.
43. Rakha, H.; Zohdy, I.; Du, J.; Park, B.B.; Lee, J.; El-Metwally, M. *Traffic Signal Control Enhancements Under Vehicle Infrastructure Integration Systems*; Mid-Atlantic Universities Transportation Center: Charlottesville, VA, USA, 2011.
44. Tashtoush, Y.M.; Darweesh, D.A.; Hussari, G.; Darwish, O.A.; Darwish, Y.; Issa, L.B.; Ashqar, H.I. Agile Approaches for Cybersecurity Systems, IoT and Intelligent Transportation. *IEEE Access* **2022**, *10*, 1360–1375. [CrossRef]
45. Ouallane, A.A.; Bahnasse, A.; Bakali, A.; Talea, M. Overview of Road Traffic Management Solutions based on IoT and AI. *Procedia Comput. Sci.* **2022**, *198*, 518–523. [CrossRef]
46. Almukhalfi, H.; Noor, A.; Noor, T.H. Traffic management approaches using machine learning and deep learning techniques: A survey. *Eng. Appl. Artif. Intell.* **2024**, *133*, 108147. [CrossRef]
47. Raiaan, M.A.K.; Mukta, M.S.H.; Fatema, K.; Fahad, N.M.; Sakib, S.; Mim, M.M.J.; Ahmad, J.; Ali, M.E.; Azam, S. A review on large Language Models: Architectures, applications, taxonomies, open issues and challenges. *IEEE Access* **2024**, *12*, 26839–26874. [CrossRef]
48. Mani, S.K.; Zhou, Y.; Hsieh, K.; Segarra, S.; Eberl, T.; Azulai, E.; Frizler, I.; Chandra, R.; Kandula, S. Enhancing Network Management Using Code Generated by Large Language Models. In Proceedings of the 22nd ACM Workshop on Hot Topics in Networks, Cambridge, MA, USA, 28–29 November 2023; ACM: New York, NY, USA, 2023; pp. 196–204. [CrossRef]
49. Liu, J.; Hang, P.; Qi, X.; Wang, J.; Sun, J. MTD-GPT: A Multi-Task Decision-Making GPT Model for Autonomous Driving at Unsignalized Intersections. In Proceedings of the 2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC), Bilbao, Spain, 24–28 September 2023; IEEE: New York, NY, USA, 2023; pp. 5154–5161. [CrossRef]
50. Elhenawy, M.; Ashqar, H.I.; Rakotonirainy, A.; Alhadidi, T.I.; Jaber, A.; Tami, M.A. Vision-Language Models for Autonomous Driving: CLIP-Based Dynamic Scene Understanding. *arXiv* **2025**, arXiv:2501.05566.
51. Sivakumar, M.; Belle, A.B.; Shan, J.; Shahandashti, K.K. Prompting GPT –4 to support automatic safety case generation. *Expert Syst. Appl.* **2024**, *255*, 124653. [CrossRef]
52. Jiang, R.; Yin, D.; Wang, Z.; Wang, Y.; Deng, J.; Liu, H.; Cai, Z.; Deng, J.; Song, X.; Shibasaki, R. DL-Traff: Survey and Benchmark of Deep Learning Models for Urban Traffic Prediction. In Proceedings of the 30th ACM International Conference on Information & Knowledge Management, Gold Coast, Australia, 1–5 November 2021; ACM: New York, NY, USA, 2021; pp. 4515–4525. [CrossRef]