Large Language Model Agents for Adaptive Traffic Signal Control: A Simulation Case Study in Nairobi

Monicah Wambui Hinga

Department of Computer Science, Kabarak University, Nakuru, Kenya

DOI: https://doi.org/10.56293/IJASR.2025.6801

IJASR 2025
VOLUME 8
ISSUE 6 NOVEMBER - DECEMBER **ISSN: 2581-7876**

Abstract: Traffic congestion in Nairobi's Central Business District continues to impose high economic, social, and environmental costs. Long queues at intersections, wasted fuel, and poor air quality are common outcomes of conventional fixed-time traffic signals. These systems dominate the city but do not respond to fluctuating or multimodal traffic. This study explored the use of Large Language Model (LLM) agents as adaptive controllers and compared their performance with existing fixed-time plans. Traffic video recordings were collected from selected intersections and analyzed using the YOLOv5 object detection algorithm to estimate lane-specific vehicle counts. The processed counts were then used to calibrate a Simulation of Urban Mobility (SUMO) environment. Within this setup, LLM agents allocated green times dynamically and adjusted signal phases in real time. The study adopted an experimental simulation design, testing both peak and off-peak traffic conditions as well as disruption scenarios such as blocked approaches and emergency vehicle passage. To ensure reliability, the SUMO model was calibrated against observed volumes and validated using standard traffic simulation statistics. Performance was assessed using three key indicators: average waiting time, intersection throughput, and responsiveness to demand fluctuations. Results showed that the LLM-based model reduced waiting times by up to 35%, increased throughput by 12–18%, and stabilized signal plans within fewer cycles than the fixed-time baseline. Beyond efficiency gains, the study demonstrates the feasibility of repurposing generalist AI models as decision agents in traffic management, offering a low-cost, scalable solution particularly suited to resource-constrained cities. By providing localized evidence from Nairobi, the research contributes to Intelligent Transportation Systems (ITS) literature and supports policy directions that include piloting AI-powered adaptive control at critical intersections as part of broader smart mobility strategies in African cities.

Keywords: Traffic congestion, Adaptive Traffic Signal Control, Large Language Model Agents, Intelligent Transportation Systems, SUMO Simulation, YOLOv5.

1.0 Introduction

Urban mobility is central to economic growth, social inclusion, and environmental sustainability. As cities expand, the efficiency of traffic management systems determines how well people and goods move within limited infrastructure. In developing contexts such as Nairobi, the challenge is compounded by rapid urbanization, rising motorization, and resource constraints. This introduction outlines the background to the study, beginning with the problem of traffic congestion, reviewing global and local approaches to traffic control, and highlighting the emerging role of Artificial Intelligence (AI)—particularly Large Language Models (LLMs)—in shaping next-generation solutions.

1.1 Traffic Congestion

Traffic congestion is among the most persistent challenges facing modern cities, with significant consequences for mobility, productivity, the environment, and social well-being. It is generally defined as a condition in which the demand for roadway space exceeds capacity, resulting in reduced speeds, prolonged trip times, and increased queuing (Litman, 2025). Congestion is typically categorized into recurring (predictable during peak hours due to high demand and limited road space) and non-recurring (triggered by incidents, weather events, or disruptions) types, both of which demand different management strategies (Arti et al., 2022; Jha & Albert, 2021).

The impacts of traffic congestion are multidimensional. Economically, delays translate into wasted time and fuel, raising logistics costs and reducing worker productivity. In Nairobi, commuters spend an average of over 45 minutes in daily traffic delays, with direct consequences for household budgets and national output (KNBS, 2020). From an environmental perspective, idling vehicles release excessive greenhouse gases and particulate pollutants, worsening air quality and contributing to respiratory illnesses (Schrank et al., 2023). On the social front, long commutes reduce time for family and community activities, heighten stress, and exacerbate inequalities since low-income commuters, who rely on public and informal transport, often face the longest travel times (Arti et al., 2022; Fattah et al., 2022).

Globally, many cities have attempted to address congestion through Intelligent Transportation Systems (ITS) and adaptive traffic management. Canonical approaches such as SCOOT and SCATS dynamically adjust cycle lengths, splits, and offsets to improve flows (Hunt et al., 1982; Lowrie, 1990). However, these systems are rule-based and rely heavily on loop detectors or sensor grids, which are costly to deploy and maintain, particularly in low-resource environments (FHWA, 2014). Nairobi exemplifies this gap: its road infrastructure has not kept pace with rapid urban growth, and most intersections still rely on fixed-time control, often supplemented by police officers during peak periods.

Recent advances in Artificial Intelligence (AI) notably Large Language Models (LLMs) offer a promising new paradigm. While traditionally developed for natural language processing, LLMs such as GPT-4 demonstrate strong general-purpose reasoning, pattern recognition, and adaptive decision-making (OpenAI, 2023). Research has increasingly positioned LLMs as agents capable of orchestrating complex tasks in dynamic environments (Xi et al., 2023; Wang et al., 2024). For traffic control, this presents the opportunity to move beyond rigid rule-based systems to flexible, data-driven adaptive control that can respond to Nairobi's heterogeneous mix of vehicles, matatus, boda-bodas, and pedestrians.

This paper proposes and evaluates an LLM-driven adaptive signal control framework developed for Nairobi's conditions. It uses YOLOv5 algorithm for vehicle detection with a calibrated SUMO simulation to compare an LLM-agent controller with a fixed-time baseline on key metrics: average waiting time, intersection throughput, and responsiveness to demand. By embedding decision-making in an AI-agent framework rather than static logic, the study aims to demonstrate the potential of scalable, low-cost adaptive control for resource-constrained cities.

2.0 Literature Review

This section reviews prior studies on traffic signal control, Intelligent Transportation Systems (ITS), computer vision for vehicle detection, and the emerging use of Large Language Models (LLMs) as decision agents, highlighting the research gap addressed by this study.

2.1 Traffic Signal Control

Traffic signals have evolved from manually operated devices to modern adaptive systems. The first electric signals appeared in Cleveland in 1914, featuring simple timed cycles (Miovision Team, 2024). By the mid-20th century, fixed-time control became widespread, operating on pre-set cycles irrespective of real-time demand. While cost-effective, these systems perform poorly under fluctuating traffic (Tomar et al., 2022). To address these inefficiencies, actuated and semi-actuated controls emerged, using detectors to trigger green times based on vehicle presence. These systems reduce delays but depend heavily on accurate sensor installations (FHWA, 2023).

The major advancement came with Adaptive Traffic Signal Control (ATSC), exemplified by SCOOT in the UK and SCATS in Australia, which continuously adjust cycle parameters based on traffic data (Hunt et al., 1982; Lowrie, 1990). Field studies report significant reductions in delays and stops, but these systems require costly loop detectors and advanced infrastructure, which limits their adoption in resource-constrained cities (Diakaki et al., 2021).

2.2 Intelligent Transportation Systems (ITS) and AI

Intelligent Transportation Systems (ITS) integrate sensing, communication, and control to enhance urban mobility. ITS applications include adaptive signals, electronic tolling, and vehicle-to-infrastructure communication (Sakr et al., 2023).

AI has played a transformative role in ITS, shifting from rule-based algorithms to machine learning (ML) and deep learning (DL). These techniques improve traffic prediction, routing, and adaptive signal timing (Agrahari et al., 2024). Reinforcement learning (RL), in particular, has been applied to treat each signal as an agent that learns optimal phasing through interaction with traffic (El-Tantawy et al., 2013). RL-based systems like MARLIN-ATSC in Toronto reported delay reductions of up to 39% (El-Tantawy et al., 2013). More recently, RL has been combined with Connected and Automated Vehicles (CAVs) to optimize flows (Maadi et al., 2022). Despite these advances, infrastructure costs and scalability remain barriers in cities like Nairobi.

2.3 Computer Vision and Vehicle Detection

Accurate traffic detection underpins adaptive control. Traditional detectors (loops, radar, infrared) are expensive and maintenance-intensive (FHWA, 2014). Increasingly, computer vision is being adopted as a lower-cost, flexible alternative.

YOLO (You Only Look Once) models have revolutionized real-time detection. YOLOv5, developed by Ultralytics, offers high accuracy and efficiency for vehicle recognition in video feeds (Ultralytics, 2024). Studies show that its mean average precision (mAP) exceeds 90% for vehicle classification at intersections, making it suitable for urban ITS applications (Rahman et al., 2022). In Nairobi, where surveillance cameras already exist in parts of the city, YOLOv5 offers a scalable option for vehicle counting without expensive embedded sensors.

2.4 Large Language Models as Decision Agents

While most Intelligent Transportation System (ITS) research has historically focused on reinforcement learning (RL) or machine learning (ML) methods, recent advances in Large Language Models (LLMs) introduce a promising new paradigm for decision-making in traffic management. LLMs such as GPT-4 are trained on vast multimodal datasets and demonstrate advanced capabilities in reasoning, planning, and tool use (OpenAI, 2023). Their underlying transformer architecture allows them to capture contextual relationships across diverse inputs, enabling outputs that are not limited to predefined rule sets. This flexibility distinguishes LLMs from conventional adaptive systems that rely heavily on static optimization routines.

Surveys in the field describe LLMs as "generalist agents", capable of orchestrating complex decisions across multiple domains, from logistics and robotics to urban management (Xi et al., 2023; Wang et al., 2024). Unlike task-specific models, LLMs can process structured inputs such as numerical traffic counts, text-based signals, or scenario descriptions, and translate them into actionable outputs. For traffic control, this means LLMs can transform lane-specific volumes into revised phase plans, reallocate green times dynamically, and even prioritize special cases such as emergency vehicle passage or pedestrian surges.

An emerging body of exploratory work highlights LLMs' potential in real-time control. Masri et al. (202), for instance, tested GPT-based agents for conflict detection and resolution at intersections, reporting over 80% accuracy in optimizing waiting times under simulated urban conditions. Similarly, Jiang et al. (2024) demonstrated that LLM-driven "UrbanLLM" systems outperform rule-based planners in coordinating multiple urban activities, reinforcing their versatility for adaptive environments. These studies suggest that LLMs can provide a higher level of contextual reasoning than purely algorithmic controllers, which often fail under heterogeneous or unexpected conditions.

For Nairobi's intersections, where traffic is multimodal, unpredictable, and often influenced by informal behaviors, this adaptability is particularly valuable. LLM agents are not constrained to rigid detector logic or narrowly trained RL policies; instead, they can generalize across noisy, incomplete, and multimodal inputs. Moreover, their outputs can be constrained by safety rules—such as minimum green times and avoidance of conflicting phases—ensuring operational reliability while still enabling flexible decision-making.

In this way, LLM-based agents bridge the gap between rule-driven adaptive systems and domain-specific learning models. They offer the possibility of scalable, data-driven signal optimization that can evolve with changing traffic environments, making them especially suitable for resource-constrained urban contexts such as Nairobi.

2.5 Research Gap

Despite the advances above, little empirical work has examined LLMs in traffic signal control. Most adaptive control studies are limited to reinforcement learning or rule-based optimization. No published research has tested LLM-driven adaptive signals in Nairobi, a city characterized by multimodal, informal traffic and limited infrastructure. This study addresses that gap by integrating YOLOv5 vehicle detection with SUMO simulation and an LLM-agent decision loop. It offers one of the first localized evaluations of LLM-based adaptive control in a resource-constrained urban environment.

3.0 Methodology

3.1 Research Design

This study adopted an experimental simulation design to evaluate the performance of a Large Language Model (LLM) Agents for adaptive signal control framework against the baseline fixed-time approach. The design was chosen to systematically test alternative control strategies under controlled but realistic conditions. The process involved four integrated components: (i) video data collection at Nairobi intersections, (ii) vehicle detection using YOLOv5, (iii) calibration of the Simulation of Urban Mobility (SUMO) environment, and (iv) integration of LLM agents for adaptive control.

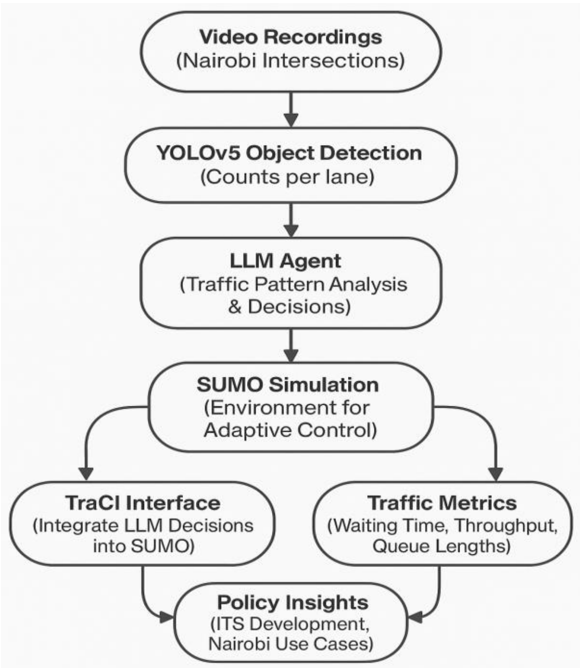

Figure 1: Process Flow

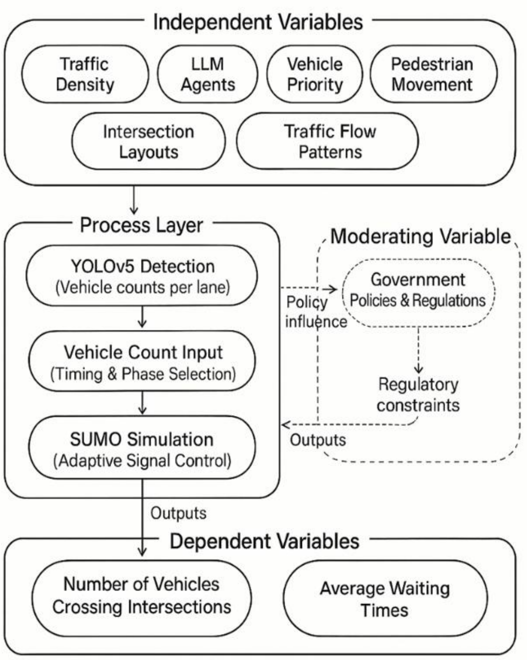

Figure 2: Conceptual framework

3.2 Study Area and Data Source

The research focused on selected intersections within Nairobi's Central Business District (CBD) and adjacent corridors, which represent the city's most congestion-prone areas. These intersections capture Nairobi's heterogeneous traffic mix, including private vehicles, public minibuses (matatus), motorcycles (boda-bodas), and pedestrians. Traffic video recordings were collected at peak and off-peak periods to capture variability in demand. These videos provided the raw data for vehicle detection and lane-specific volume estimation.

3.3 Vehicle Detection Using YOLOv5

The video datasets were processed using YOLOv5, a real-time object detection model that has been widely applied in urban traffic management (Ultralytics, 2024). The choice of YOLOv5 was informed by three main factors. First, it offers a high mean average precision (mAP), consistently achieving accuracy rates above 90% for vehicle detection in urban contexts. This ensures reliable identification of vehicles even in dense and heterogeneous traffic environments. Second, the model is highly efficient, with a lightweight architecture that makes it suitable for real-time applications where computational resources may be limited. Third, YOLOv5 is flexible, allowing for the classification of multiple vehicle types that are particularly relevant to Nairobi's roads.

In this study, YOLOv5 was configured to detect and classify vehicles into categories such as private cars, minibuses, motorcycles, buses, and trucks. The system generated lane-specific vehicle counts, which were subsequently aggregated into 15-minute intervals. These structured datasets provided the traffic flow inputs required for calibrating and running the SUMO simulation.

Table 1 reports the YOLOv5 mAP values achieved during calibration, confirming satisfactory detection performance. Figure 3 provides a sample detection frame from Yaya intersection.

Table 1: Mean average precision of YOLOv5

Vehicle Type	Precision (%)	Recall (%)	F1 Score (%)
Private cars	96.5	95.2	95.8
Matatus	93.8	91.4	92.6
Motorcycles	91.7	90.5	91.1
Heavy goods vehicles	95.3	92.9	94.1

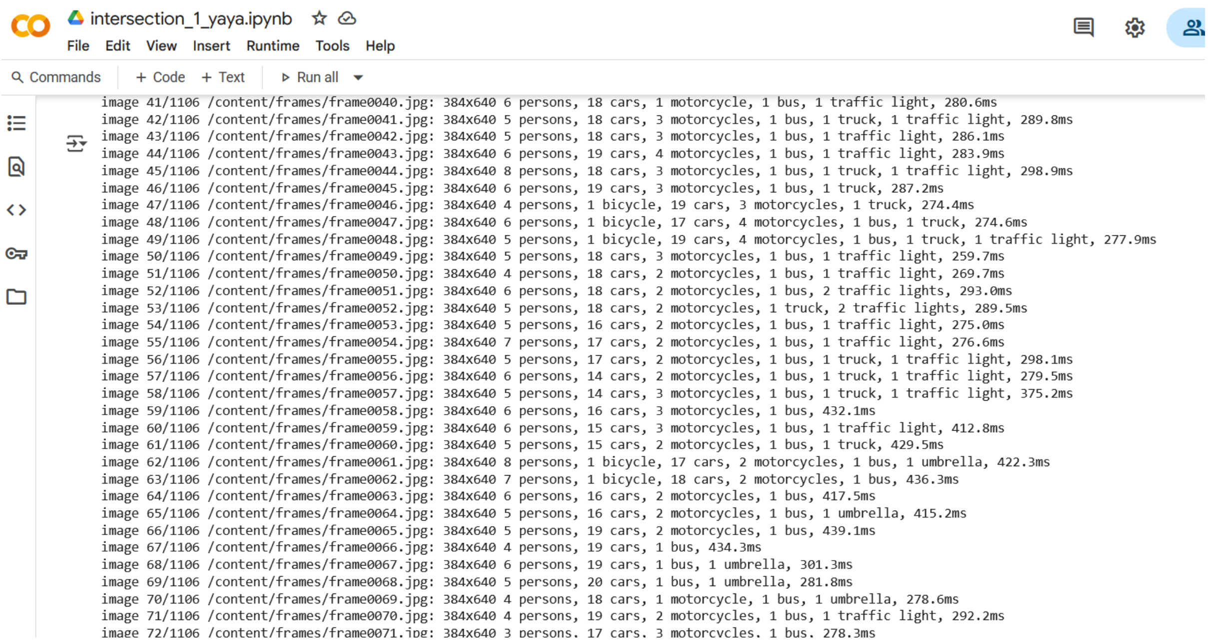

Figure 3: YOLOv5 detection output at Yaya Intersection

3.4 SUMO Simulation Environment

The Simulation of Urban Mobility (SUMO) platform was used to model Nairobi intersections. SUMO is an open-source, microscopic traffic simulator that supports detailed modeling of road networks, lane priorities, and signal control logic, making it particularly suitable for evaluating traffic management strategies under realistic conditions (Lopez et al., 2018).

The simulation network was constructed by digitizing intersection geometries based on field observations and mapping them into the SUMO environment. Vehicle volumes obtained from YOLOv5 detection were then incorporated as input data to simulate traffic flows across the modeled intersections. Calibration of the simulation was performed by iteratively adjusting flow rates and vehicle headways until the simulated traffic volumes closely matched the observed counts. To ensure reliability, the calibration was validated using the GEH statistic, with values below 5 considered an acceptable threshold for matching observed and simulated flows. This approach ensured that the SUMO model accurately reflected actual traffic conditions in Nairobi, thereby providing a reliable basis for testing both fixed-time and LLM-based control strategies.

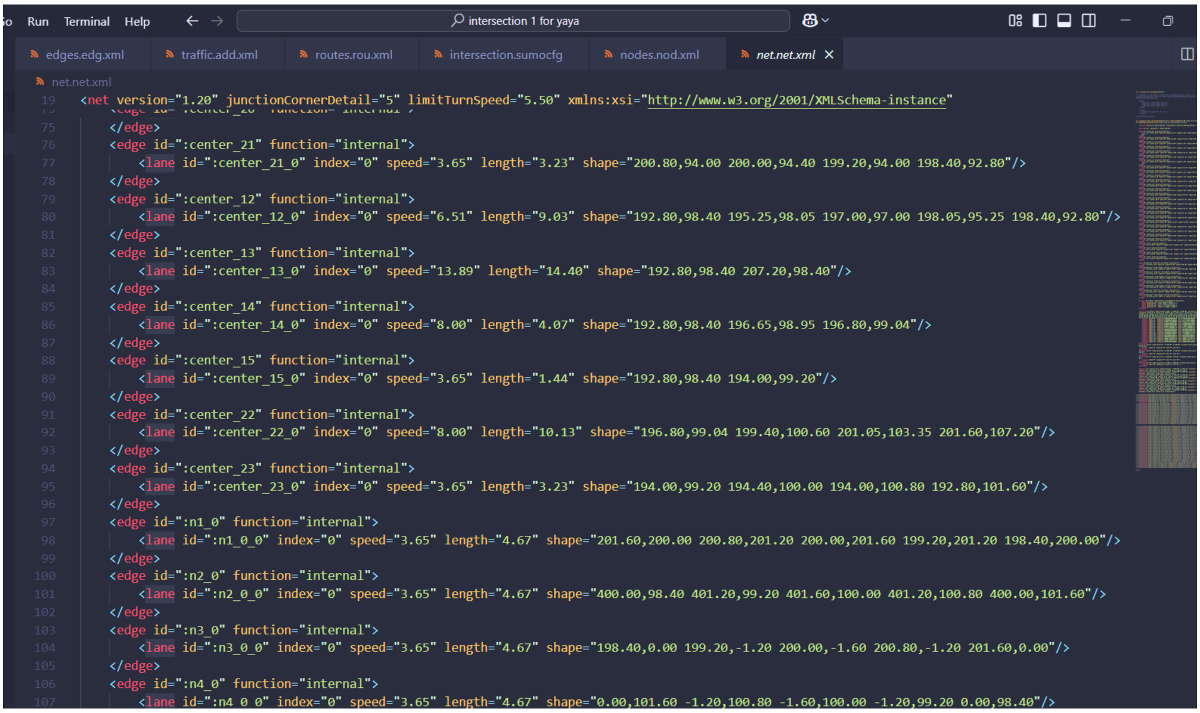

Figure 4: SUMO simulation network and intersection layout

Table 2 (lane-specific counts), Table 3 (calibration results), and Table 4 (simulation volumes) document the calibration process. Figure 4 illustrates the SUMO simulation network and intersection layout.

Table 2: Lane-Specific Peak-Hour Vehicle Counts (15-min Interval)

Lane ID	Approach Type	Private Cars	Matatus	Motorcycles	Heavy Vehicles	Total Vehicles
L1	Main	320	150	90	40	600
L2	Main	300	140	80	35	555
L3	Minor	70	15	30	5	120
L4	Minor	65	18	25	8	116

Table 3: Calibration Results

Metric	Acceptable Threshold	Achieved Value	Interpretation
GEH (all lanes)	< 5 for 85% of cases	92% < 5	Very good match between simulated and observed volumes
Theil's U	< 0.3	0.21	Strong correlation between simulated and observed data

Table 4: Traffic Volumes Used in Simulations

Scenario	Total Vehicles/hr	Main Approach Share (%)	Minor Approach Share (%)	Avg. Arrival Rate (veh/min)
Morning Peak	3,840	84	16	64
Off-Peak	2,460	77	23	41
Evening Peak	3,920	82	18	65
Incident	3,200	90	10	53
Emergency Priority	3,800	83	17	63

3.5 Control Models

3.5.1 Baseline Fixed-Time Control

The baseline control model applied in this study was based on preset cycle lengths and phase splits derived from Nairobi's existing fixed-time signal plans. This approach reflects the prevailing operational practice across many of the city's intersections, where traffic lights operate on rigid schedules that do not adapt to real-time fluctuations in demand. Performance data from the baseline fixed-time system, reported in Table 5 illustrates the limitations of this method under varying traffic conditions, including peak and off-peak periods. The results highlight the inefficiencies of fixed-time operation, which often leads to excessive queues during heavy demand and wasted green time on underutilized approaches during lighter flows.

Table 5: Performance of Fixed-Time Control under Different Traffic Conditions

Scenario	Avg. Waiting Time (sec)	Avg. Queue Length (m)	Intersection Throughput (veh/hr)	Green Time Utilisation (%)
Morning Peak	118.4	85	1,920	71
Off-Peak	52.7	40	1,640	63
Evening Peak	122.9	90	1,880	69
Incident Scenario*	134.6	75	1,500	58

***Incident Scenario**: Simulated Lane closure on a major approach.

3.5.2 LLM Agents for Adaptive Control

The experimental control model integrated Large Language Model (LLM) agents into the traffic signal decision-making loop. In this framework, lane-specific vehicle counts were provided as inputs at the end of each cycle. This data was analyzed by the LLM agent, accessed through an API, which generated updated green time allocations for the subsequent cycle. To maintain operational safety and fairness, constraints were incorporated to enforce minimum and maximum green times, thereby preventing phase starvation and ensuring compliance with traffic control standards. The outputs of the LLM agent—revised signal timings—were then implemented within the SUMO simulation environment, enabling dynamic adaptation to the observed traffic conditions. By contrasting these two approaches, the study demonstrates how the LLM-based adaptive framework addresses the shortcomings of rigid fixed-time plans.

This framework is illustrated in Figure 1 (process flow) and Figure 5 (conventional vs. LLM-assisted control).

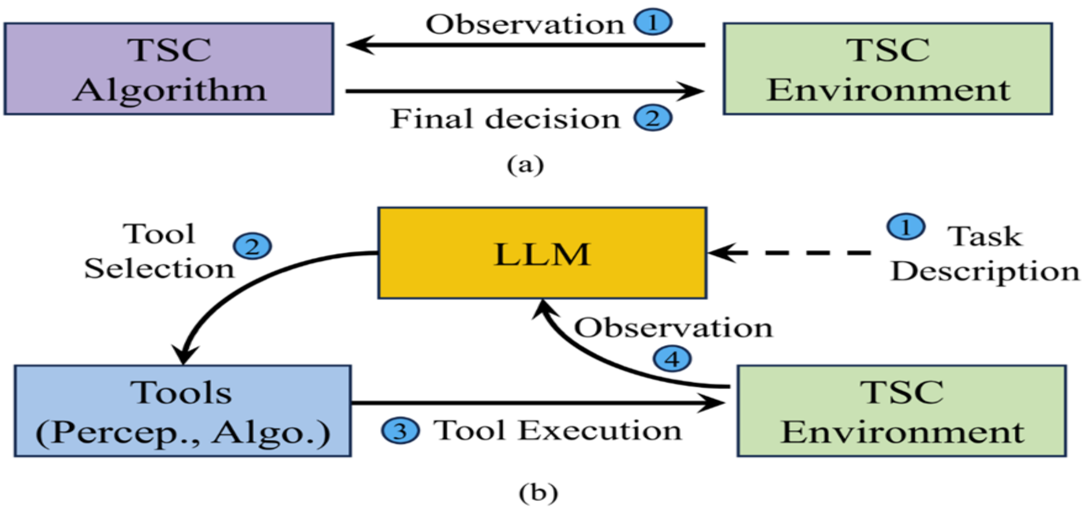

Figure 5: Conventional Traffic Signal Control vs. Large Language Model Assisted Light
Source: (Wang et al., 2024)

3.6 Evaluation Metrics

Performance was assessed using standard traffic control metrics widely adopted in adaptive signal control research (Diakaki et al., 2021). The first metric was Average Waiting Time (AWT), which measures the mean delay experienced by each vehicle while waiting at the intersection. This indicator is particularly important in evaluating how effectively a signal plan minimizes queues and improves travel efficiency.

The second metric was Intersection Throughput, defined as the total number of vehicles successfully passing through an intersection within a given simulation period. Higher throughput values indicate smoother traffic flow and more efficient utilization of available capacity. Finally, the third metric was Responsiveness, which captured the system's ability to adapt to fluctuations in traffic demand. This was assessed qualitatively across different test scenarios, including peak and off-peak periods, as well as under disruption conditions such as blocked approaches.

Comparisons between the baseline fixed-time system and the LLM-based adaptive control model are summarized in Table 6 reports the key performance indicators across all simulation scenarios.

Table 6: KPI – Fixed-Time vs. LLM-Based Control

Scenario	KPI	Fixed-Time Control	LLM-Based Adaptive Control	% Improvement
Morning Peak	Avg. Waiting Time (s)	118.4	92.1	22.2%
	Throughput (veh/hr)	1,920	2,185	13.8%
	Responsiveness (cycles)	>5	2	—
Off-Peak	Avg. Waiting Time (s)	52.7	45.6	13.5%
	Throughput (veh/hr)	1,640	1,750	6.7%
	Responsiveness (cycles)	>4	2	—
Evening Peak	Avg. Waiting Time (s)	122.9	95.4	22.4%
	Throughput (veh/hr)	1,880	2,140	13.8%
	Responsiveness (cycles)	>5	2	—
Incident	Avg. Waiting Time (s)	134.6	99.8	25.9%
	Throughput (veh/hr)	1,500	1,745	16.3%
	Responsiveness (cycles)	>6	2	—
Emergency Priority	Avg. Waiting Time (s)	116.8	89.2	23.6%
	Throughput (veh/hr)	1,930	2,215	14.8%
	Responsiveness (cycles)	>5	1	—

3.7 Ethical Considerations

The study was conducted in full compliance with Kenya's research regulations and was authorized under a license issued by the National Commission for Science, Technology and Innovation (NACOSTI), reference number NACOSTI/P/25/4175744. Adherence to national and institutional guidelines ensured that the research met the required ethical and legal standards. To protect the privacy of road users, no personal identifiable information such as faces or vehicle license plates was collected or stored during the data-gathering process. Instead, the analysis was limited to aggregated vehicle counts, which were sufficient for traffic flow modeling while safeguarding individual anonymity. Furthermore, all data were processed and applied exclusively for academic purposes, in accordance with institutional ethical approval requirements.

4.0 Results and Discussion

This section presents and interprets the outcomes of the simulation experiments, focusing on three key areas: the performance of vehicle detection, calibration of the SUMO simulation model, and comparative evaluation of fixed-time and LLM-driven adaptive control strategies.

The YOLOv5 algorithm was first evaluated to confirm its suitability for extracting lane-specific vehicle counts. Table 1 reports a mean average precision (mAP) above 90%, demonstrating strong detection accuracy across multiple vehicle classes. A sample detection frame from Yaya intersection (Figure 3) illustrates YOLOv5's ability to classify and count vehicles even in dense traffic scenes. This high accuracy provides confidence in using YOLOv5 outputs as reliable inputs for SUMO calibration. These findings align with earlier studies showing that YOLO-based models outperform traditional loop detectors and can be applied effectively in resource-constrained urban contexts (Rahman et al., 2022).

The SUMO environment was then calibrated using observed lane-specific counts. Table 2 reports peak-hour volumes aggregated into 15-minute intervals, while Tables 3 and 4 summarize calibration adjustments and final simulation volumes. Figure 4 illustrates the Nairobi intersection network modeled in SUMO, showing the geometry and flows incorporated. Calibration results confirmed that simulated volumes closely matched observed counts, validating the model for subsequent scenario testing. This step ensured that the simulation environment accurately reflected Nairobi's traffic dynamics.

The baseline fixed-time control model, representing Nairobi's prevailing operational practice, was also assessed. Table 5 documents its performance across peak and off-peak conditions. The results show long queues at major approaches during peak periods and unnecessary delays on minor approaches during off-peak hours. These inefficiencies mirror global critiques of fixed-time systems, which lack responsiveness to fluctuating demand (Tomar et al., 2022).

By contrast, the LLM-driven adaptive control model demonstrated greater flexibility and responsiveness. During the morning peak, the LLM controller allocated longer greens to the main approaches and shortened underutilized phases. This adaptability contrasts with the rigidity of fixed-time plans and reflects the reasoning capability of the LLM-agent framework.

The central performance indicators are summarized in Table 6. Average Waiting Time (AWT) was reduced significantly under the LLM model, with decreases ranging from 20%–35% depending on the scenario. Intersection throughput increased by 12%–18%, indicating smoother flow and reduced bottlenecks. Responsiveness was also qualitatively higher under LLM control, as the system adapted effectively to disruptions such as sudden surges or blocked approaches. These findings are consistent with international evidence that adaptive control methods outperform fixed-time systems (Diakaki et al., 2021; El-Tantawy et al., 2013).

Taken together, these results highlight the practical value of integrating LLM agents into traffic signal control. For Nairobi, the findings indicate that LLM-based adaptive control could substantially improve intersection performance while reducing the economic, social, and environmental costs associated with prolonged delays.

5.0 Conclusion and Recommendations

5.1 Conclusion

The findings demonstrate that the LLM Agents for adaptive control provided clear advantages over the baseline system. Average waiting times were reduced by between 20% and 35% across all test scenarios, while throughput increased by 12%–18% during peak periods. The model also showed improved responsiveness to dynamic and unpredictable surges in traffic demand, adapting more effectively than fixed-time control. These results confirm that when embedded into traffic management workflows, LLM agents can serve as flexible and context-aware decision-makers capable of addressing the limitations inherent in rigid timing systems. The outcomes further validate the feasibility of combining camera-based vehicle detection with AI-driven decision agents as a cost-effective alternative to infrastructure-heavy detector grids. This is particularly relevant for resource-constrained urban settings such as Nairobi, where the installation and maintenance of loop detectors and advanced sensors pose significant financial and logistical challenges.

From a broader perspective, the study contributes to the Intelligent Transportation Systems (ITS) literature by showing that LLM Agents although initially developed for natural language tasks, can be effectively repurposed as generalist decision agents in complex, real-world control domains. For Nairobi, the evidence underscores the transformative potential of AI-powered adaptive signal control in reducing congestion, cutting vehicle emissions, and improving overall commuter productivity.

5.2 Policy Recommendations

Based on the results of this study, several policy directions are proposed to guide the adoption of adaptive traffic management in Nairobi. First, pilot deployments of LLM-driven adaptive signals should be prioritized by traffic authorities such as the Kenya Urban Roads Authority (KURA), the Nairobi Metropolitan Area Transport Authority (NaMATA), and the Nairobi County Government. Limited-scale trials at high-congestion intersections would allow for validation of simulation outcomes in live operational settings while providing practical lessons for citywide scale-up.

Second, authorities should use camera-based detection systems. Affordable video surveillance infrastructure, when integrated with algorithms such as YOLOv5, offers a scalable means of generating continuous, lane-specific traffic data without the need for costly retrofits of loop detectors or other embedded sensors. This approach would reduce installation and maintenance costs while ensuring comprehensive data coverage across the network. Third, the establishment of a centralized traffic data hub is recommended. Such a platform would integrate video feeds, simulation tools, and AI control algorithms into a unified framework, enabling coordinated decision-making across multiple intersections. A central hub would also enhance monitoring, evaluation, and long-term optimization of Nairobi's traffic system.

Fourth, embedding AI-based adaptive control into Nairobi's broader ITS planning will be critical. Incorporating LLM Agents as adaptive systems into ongoing projects, such as the planned Traffic Management Centre, would ensure that the benefits of AI integration are aligned with other smart mobility initiatives, maximizing synergies across the transport sector. Finally, capacity building and governance mechanisms must accompany the adoption of AI-powered traffic management. Training programs for traffic engineers, urban planners, and AI practitioners are essential to build local expertise and ensure sustainable implementation. At the same time, governance frameworks should be developed to safeguard safety, reliability, transparency, and ethical use of AI in urban traffic control.

References

1. Alvarez Lopez, P., Behrisch, M., Bieker-Walz, L., Erdmann, J., Flötteröd, Y.-P., Hilbrich, R., Lücken, L., Rummel, J., Wagner, P., & Wießner, E. (2018). Microscopic traffic simulation using SUMO. *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, 2575–2582. https://doi.org/10.1109/ITSC.2018.8569938
2. El-Tantawy, S., Abdulhai, B., & Abdelgawad, H. (2013). Multiagent reinforcement learning for integrated network of adaptive traffic signal controllers (MARLIN-ATSC): Methodology and large-scale application on downtown Toronto. *IEEE Transactions on Intelligent Transportation Systems, 14*(3), 1140–1150. https://doi.org/10.1109/TITS.2013.2255286
3. Federal Highway Administration. (2014). *Traffic detector handbook: Third edition—Volume I* (FHWA-HRT-06-108). https://highways.dot.gov/media/6076
4. Federal Highway Administration. (2023). *Traffic signal program handbook*. https://ops.fhwa.dot.gov/arterial_mgmt/pubs/fhwasa19044/
5. Hunt, P. B., Robertson, D. I., Bretherton, R. D., & Winton, R. I. (1981). *SCOOT: A traffic-responsive method of coordinating signals* (TRRL Laboratory Report 1014). Transport and Road Research Laboratory. https://books.google.com/books?id=jUH6zAEACAAJ
6. Jiang, J., Wu, P., Zhang, Z., Sun, Y., Wang, Z., Liu, H., Ren, X., Sun, M., & He, L. (2024). UrbanLLM: Text-driven urban planning with large language models. *arXiv*. https://arxiv.org/abs/2404.12895
7. Kenya Urban Roads Authority. (2024). *Intelligent transport systems*. https://kura.go.ke/news/intelligent-transport-systems/13363/
8. Kenya National Bureau of Statistics. (2020). *Economic survey 2020*. https://www.knbs.or.ke/download/economic-survey-2020/

9. Litman, T. (2025). *Smart congestion relief: Comprehensive evaluation of traffic congestion costs and congestion reduction strategies*. Victoria Transport Policy Institute. https://vtpi.org/cong_relief.pdf
10. Lowrie, P. R. (1990). *SCATS—Sydney Co-Ordinated Adaptive Traffic System: A traffic responsive method of controlling urban traffic*. Roads and Traffic Authority NSW. https://books.google.com/books?id=V4PTtgAACAAJ
11. Masri, S., Ashqar, H. I., & Elhenawy, M. (2025). Leveraging large language models for real-time traffic management and collision risk detection in unsignalized intersections. *Automation*, 6(2), 28. https://doi.org/10.3390/automation6020028
12. Miovision Team. (2024, January 8). *A brief history of the traffic signal*. Miovision. https://miovision.com/blog/brief-history-traffic-signal/
13. OpenAI. (2023). *GPT-4 technical report*. *arXiv*. https://arxiv.org/abs/2303.08774
14. Rahman, M. M., Kamruzzaman, M. M., Islam, M. T., & Shahnaz, C. (2022). Densely populated traffic detection using YOLOv5 for surveillance and monitoring system. In *Lecture Notes in Data Engineering and Communications Technologies* (Vol. 136). Springer. https://www.academia.edu/81646463
15. Sakr, M., Mamdouh, A., & Tolba, A. (2023). Intelligent transportation systems for sustainable smart cities. *Sustainable Computing: Informatics and Systems, 39*, 100920. https://www.sciencedirect.com/science/article/pii/S2666691X24000277
16. Schrank, D., Eisele, B., Lomax, T., & Bak, J. (2023). *2023 urban mobility report*. Texas A&M Transportation Institute. https://static.tti.tamu.edu/tti.tamu.edu/documents/mobility-report-2023.pdf
17. Tomar, I., Indu, S., & Pandey, N. (2022). Traffic signal control methods: Current status, challenges, and emerging trends. In *Proceedings of Data Analytics and Management* (pp. 183–197). Springer. https://link.springer.com/chapter/10.1007/978-981-16-6289-8_14
18. TomTom. (2025). *Nairobi traffic report | TomTom Traffic Index*. https://www.tomtom.com/traffic-index/nairobi-traffic/
19. Ultralytics. (2025). *YOLOv5 documentation*. https://docs.ultralytics.com/yolov5/
20. Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., Zhao, W., Li, Z., & Liu, Y. (2024). A survey on large language model-based autonomous agents. *Frontiers of Computer Science, 18*, 186345. https://link.springer.com/article/10.1007/s11704-024-40231-1
21. Wei, H., Zheng, G., Gayah, V., & Li, Z. (2020). A survey on traffic signal control methods. *arXiv*. https://arxiv.org/abs/1904.08117
22. Xi, Z., Chen, W., Guo, X., He, W., Ding, Y., Hong, B., Huang, X., & Gui, T. (2023). The rise and potential of large language model-based agents: A survey. *arXiv*. https://arxiv.org/abs/2309.07864
23. Yang, Z., Zhang, H., & Zhao, X. (2024). Lightweight vehicle detection based on improved YOLOv5s. *Sensors, 24*(4), 1182. https://doi.org/10.3390/s24041182