© IIJP 2026 | Volume 2, Issue 3, pp. 451-470 | ISSN: 2636-4484

HYBRID ANFIS-GA-RL TRAFFIC SIGNAL CONTROL: A SIMULATION STUDY OF NIGERIAN CITIES

Dodo, Enoch Jacob[1], Onwodi, Gregory[2], Okure, Obot.[3]

[1] Department of Information System and Technology Faculty of Computing, National Open University of Nigeria, Jabi, FCT Abuja, Nigeria.:enochdodo46@gmail.com

[2] Department of Information System and Technology Faculty of Computing, National Open University of Nigeria, Jabi, FCT Abuja, Nigeria.: gonwodi@noun.edu.ng

[3]Department of Software Engineering Faculty of Computing, University of Uyo, Uyo, Akwa Ibom State, Nigeria. :okureobot@uniuyo.edu.ng

Corresponding author: enochdodo46@gmail.com

ARTICLE INFO
Article No.: 0450
Accepted Date: 10/07/2026
Published Date: 29/07/2026
Type: Research

ABSTRACT

Rapid urbanisation in developing economies has rendered conventional and semi-adaptive traffic signal strategies increasingly ineffective against the non-linear, stochastic vehicular flows that characterise fast-growing megacities. This paper presents ITOS-v2, a novel Hybrid Intelligent Traffic Optimisation System that tightly integrates an Adaptive Neuro-Fuzzy Inference System (ANFIS) for interpretable fuzzy reasoning, a Genetic Algorithm (GA) for global parameter tuning, and a Reinforcement Learning (RL) actor-critic module for continuous online policy refinement. The tri-hybrid ANFIS–GA–RL architecture is evaluated through a calibrated MATLAB R2023b microscopic simulation of a single representative four-arm signalised intersection, using the U.S. Department of Transportation's Traffic Signal Performance Measures (TSPM) database (Data.gov, 2023) as the base signal-timing dataset, with input demand profiles and fleet-composition parameters rescaled to published contextual data for three Nigerian metropolitan clusters: Kano, Kaduna, and Lagos cities representing extreme to moderate congestion regimes with saturation degrees of 1.38–2.64. No field data were collected from Nigerian roads; the three cities are represented only by pairing each city's input demand volume with the results of a single base-case simulated intersection (Section 5.3), not by independently re-simulating that intersection under each city's full demand and saturation conditions or by simulating multiple real intersections. System performance is benchmarked against five increasingly complex baselines: Fuzzy Logic (FL), Artificial Neural Network (ANN), standalone ANFIS, standalone GA, and the ANFIS–GA hybrid. Evaluation across eight key performance indicators (KPIs) travel time (TT), average vehicle speed (AVS), throughput (TP), traffic density (TD), queue length (QL), delay time (DT), intersection delay (ID), and computational time (CT) reveals that ITOS-v2 achieves a 64.9% overall improvement over the fuzzy baseline, reducing TT to 0.20 h, ID to 0.01 h, DT to 0.50 h, and raising TP to 150 veh/h at minimum CT (0.0111 h). The base single-intersection model ranking is presented alongside each city's respective demand volume as contextual scaling information, rather than as evidence from independent re-simulations under each city's traffic conditions; city-specific re-simulation is identified as necessary future work. These findings position ITOS-v2 as a promising, computationally efficient foundation for next-generation smart-city traffic management, with field validation and network-scale extension identified as necessary next steps toward real-world deployment.

Keywords: Neuro-Fuzzy Inference; Evolutionary Optimisation; Reinforcement Learning; Intelligent Traffic Control; Single-Intersection Simulation; Developing Cities; Sub-Saharan Africa

1. Introduction

Urban congestion is one of the defining infrastructure crises of the 21st century. By 2021 approximately 55% of the global population resided in urban areas; projections by the United Nations (2024) indicate this share will reach 68% by 2050, with Sub-Saharan Africa experiencing the fastest growth trajectory. Lagos, Nigeria alone projected to exceed 13 million residents by 2025—has consistently ranked among the world's most congested megacities, with average saturation degrees exceeding 1.95 across 16 metropolitan sub-areas (see Section 4). The World Health Organisation (2023) estimates that road traffic injuries claim approximately 1.19 million lives annually—the leading cause of death for individuals aged 5–29 and impose macroeconomic costs equivalent to 3% of national GDP, costs disproportionately borne by low- and middle-income countries (WHO, 2023).

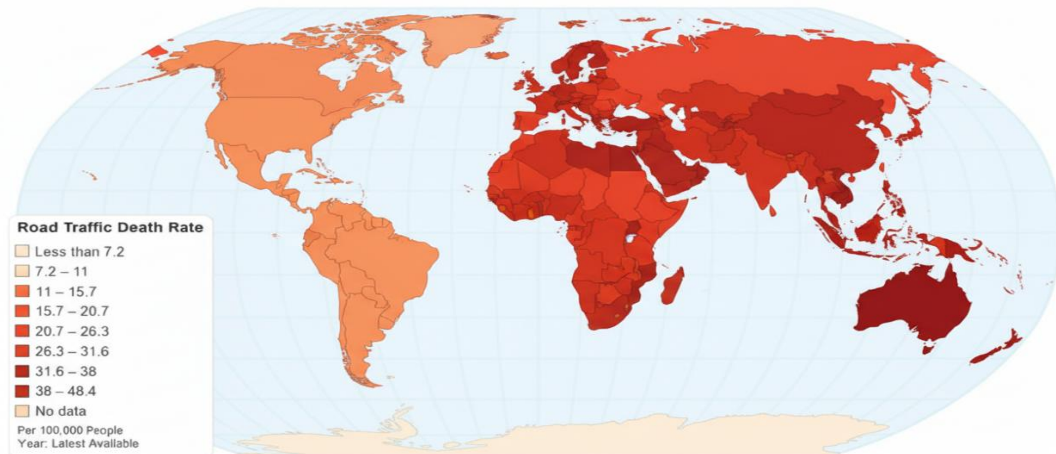

Figure 1: World Map Displaying Road - Traffic Death Rate
Source: Our World in Data, Death rate from road injuries (IHME, Global Burden of Disease, 2025), ourworldindata.org/grapher/death-rates-road-incidents

Conventional fixed-time signal plans, once calibrated from counted turning movements, become chronically mismatched to dynamic demand within months of installation. Actuated controllers improve responsiveness but remain reactive and incapable of proactive congestion anticipation. Recent years have witnessed a paradigm shift toward Artificial Intelligence (AI)-assisted traffic control: Fuzzy Logic (FL) provides expert-interpretable rules under uncertainty; Artificial Neural Networks (ANN) learn non-linear flow relationships; Genetic Algorithms (GA) efficiently search high-dimensional signal timing spaces; and Reinforcement Learning (RL) trains adaptive policies through interaction with the traffic environment (Olayode et al., 2023; Sutton & Barto, 2018). However, stand-alone instantiations of these methods each exhibit characteristic limitations FL requires laborious manual rule construction; ANNs are opaque black-boxes; GAs converge slowly in real-time settings; and RL agents demand massive, often impractical, interaction data from the deployment environment (Yang et al., 2025).

Hybrid architectures that combine the complementary strengths of these paradigms have emerged as the most promising avenue. Notably, ANFIS–GA combinations improve prediction accuracy (Olayode et al., 2023), and multi-objective triple-integration frameworks reduce delays, emissions, and safety conflicts simultaneously (Mirbakhsh & Azizi, 2024; Chala & Kóczy, 2024). Nevertheless, several critical gaps persist in the literature. First, the majority of hybrid studies are confined to single isolated intersections, severely limiting generalisability claims. Second, almost no study has calibrated and validated such frameworks against the specific demographic, infrastructure, and traffic demand conditions of Nigerian or broader Sub-Saharan African cities, where heterogeneous fleets, poor signal compliance, mixed motorbike/automobile streams, and infrastructure deficits profoundly alter system behaviour. Third, systematic multi-KPI benchmarking across a full progression of model complexity—from FL through to triple-hybrid RL—is rare, preventing rigorous attribution of performance gains to individual modelling choices. It should be stated plainly at the outset that the present study addresses these gaps through simulation rather than field deployment: ITOS-v2 is evaluated on a single representative signalised intersection in MATLAB, with the Kano, Lagos, and Kaduna "city" scenarios constructed by re-parameterising that same simulated intersection with each city's demand and fleet-composition profile. No physical intersections in Kano, Lagos, or Kaduna were instrumented or tested.

This paper directly addresses these gaps through the following original contributions:

- **(C1)** A novel tri-hybrid ANFIS–GA–RL intelligent traffic optimisation system (ITOS-v2) that synergistically integrates fuzzy inference, evolutionary optimisation, and actor-critic reinforcement learning within a unified real-time control architecture.
- **(C2)** A simulation-based evaluation of a tri-hybrid AI traffic controller at a single representative intersection, with its base-case results contextualised against the demand parameters of three Nigerian metropolitan clusters (Kano, Lagos, Kaduna) spanning eight congestion levels and covering over 19 million projected inhabitants.
- **(C3)** A comprehensive eight-KPI benchmarking study across six controllers of increasing complexity, enabling fine-grained attribution of performance improvements to each hybridisation step.
- **(C4)** A demand-scale contextualisation showing how the base single-intersection model ranking sits alongside the demand volumes of three Nigerian cities, identifying independent city-specific re-simulation as a concrete, well-specified direction for future work rather than claiming it has already been demonstrated.

The remainder of this paper is structured as follows: Section 2 reviews related work on ANFIS, GA, RL, and hybrid traffic control. Section 3 describes the proposed system architecture and mathematical formulation. Section 4 presents the experimental setup, dataset, and simulation environment. Section 5 reports and analyses results. Section 6 discusses implications and limitations. Section 7 concludes with future directions.

2. Related Work

A bibliometric synthesis of ANFIS-, GA-, and DRL-based traffic-optimisation research spanning 2012–2025 is reported separately by the present author group (Dodo et al., 2025), covering 76 studies and quantifying field-wide trends across single-method and hybrid approaches. This section narrows focus to the specific question that motivates ITOS-v2: which prior controllers have integrated ANFIS, GA, and RL within one closed-loop architecture, and which if any have been calibrated against the demand, fleet composition, and infrastructure conditions of Sub-Saharan African cities. Independent evidence corroborates this framing: hybrid ANFIS-GA controllers have been validated on South African intersection data (Olayode et al., 2023), multi-agent RL controllers have been trialed in a Nigerian–South African smart-city collaboration (Olusanya et al., 2025), and a systematic review of hybrid rule-based/RL control in another developing-city setting reports the same integration gap under heterogeneous, motorcycle-dominated traffic (Kurniawan et al., 2025) yet none combine all three methods, and none target Sub-Saharan African demand or fleet conditions specifically. Sections 2.1–2.3 summarise single-method and two-way hybrid precedents strictly as inputs to the gap analysis in Section 2.4, which positions the present contribution directly against this prior work.

2.1 Standalone ANFIS Controllers: Demonstrated Gains, Bounded Adaptivity

The Adaptive Neuro-Fuzzy Inference System, introduced by Jang (1993), combines the linguistic interpretability of Mamdani-type fuzzy systems with the gradient-based parameter adaptation of neural networks. Standalone ANFIS controllers have been implemented and field-validated at the intersection level: Udofia (2019) developed an ANFIS-based signal controller for two interconnected junctions in Uyo, Nigeria, reporting average-delay reductions of 39–60% and queue-length reductions of 63–89% relative to fixed-time control across low, medium, and peak traffic densities, with paired t-tests confirming statistical significance—one of the few ANFIS field studies calibrated against Nigerian intersection data, and a direct precedent for the present study's Nigerian contextualisation. Olayode, Tartibu, and Alex (2023) compared standalone ANFIS against an ANFIS–GA hybrid for vehicle-flow prediction at South African road intersections; the GA-tuned variant consistently outperformed the standalone model, evidencing that ANFIS accuracy is bounded by its initial fuzzy-rule parameterisation. These two studies jointly establish that (a) ANFIS is a viable, interpretable, Africa-relevant controller, and (b) its principal limitation—sensitivity to initial parameterization is precisely what evolutionary tuning is positioned to address.

2.2 Genetic Algorithms: Strong Global Search, Weak Real-Time Responsiveness

Genetic Algorithms (Holland, 1975) model the Darwinian process of selection, crossover, and mutation to iteratively refine solutions to complex, non-linear combinatorial problems, of which traffic-signal timing—with its high-dimensional, multi-objective parameter space—is a natural example. The breadth of GA-based signal-timing studies, from single intersections to coordinated multi-junction networks, is catalogued in the companion bibliometric review (Dodo et al., 2025). The consistent finding across that literature is that GA-tuned controllers outperform manually tuned or single-objective baselines, but that GA's population-based, generation-by-generation search is computationally slow relative to the sub-second decision cycle required for live signal control. This combination—strong offline global optimisation, weak real-time responsiveness—is the specific characteristic that pairing GA with a fast-reacting RL layer is designed to address (Section 2.3), and is the rationale for the present study's two-tier design in which GA tunes parameters offline and periodically, while RL adapts within each control epoch (Section 3.3).

2.3 Deep Reinforcement Learning: Real-Time Adaptivity, Data-Hungry Cold Start

Reinforcement Learning (Sutton & Barto, 2018) frames traffic signal control as a Markov Decision Process in which an agent learns a signal policy by maximising cumulative reward through interaction with the traffic environment. Network-scale extensions show that explicit modelling of inter-intersection dependence materially improves outcomes: Yang, Wen, and Chen (2025) demonstrated that a multi-agent deep RL controller augmented with a graph attention network captured cross-intersection correlations more effectively than non-coordinating multi-agent baselines, improving signal-control performance across a simulated multiple-intersection urban network. A persistent challenge for standalone RL, noted throughout this literature, is the volume of interaction data required for convergence, which renders cold-start deployment costly in data-scarce environments such as Nigerian cities—motivating RL's embedding within a pre-tuned hybrid architecture rather than its deployment in isolation.

2.4 Two-Way Hybrids and the Three-Part Gap Addressed by ITOS-v2

Two-way hybrids that pair ANFIS with GA, or DRL with complementary AI paradigms, show consistent gains over either constituent method alone. Mirbakhsh and Azizi (2024) developed a multi-objective deep RL controller using a dueling double-DQN architecture that reduced traffic conflicts by over 16% and CO_2 emissions by 4%, alongside an 18% reduction in waiting time relative to traditional control, at a simulated intersection in Changsha, China. Chala and Kóczy (2024) showed that fuzzy rule-base reduction integrated with a SUMO-based control loop reduced average waiting time, fuel consumption, and CO_2 emissions while requiring substantially less execution time than an unreduced fuzzy rule base. Bangalee and Ahmed (2024 currently a preprint, not yet completing peer review) reported a fuzzy graph-colouring/DRL hybrid achieving efficiency gains of 6.34–59.08% over conventional control across varying traffic intensities, illustrating the scale of gains achievable when complementary AI paradigms are combined. Independently of the present author group's own bibliometric review, two further studies corroborate the specific gap targeted here: Olayode, Severino, Tartibu, Arena, and Cakici (2022) evaluated a hybrid PSO-enhanced ANFIS model against real South African freeway traffic-flow data, confirming that evolutionary-tuned neuro-fuzzy hybrids outperform standalone ANFIS on an African dataset, though at a single-method (non-triple-hybrid) level; and Jutury, Kumar, Sachan, Daultani, and Dhakad (2023) developed an adaptive neuro-fuzzy multi-mode traffic light controller for urban networks in a data-constrained developing-country setting, demonstrating the practical relevance of ANFIS-based hybrids outside well-instrumented, high-income networks. Neither study, however, integrates ANFIS, GA, and RL within one closed-loop controller.

None of these systems, however, satisfies all three conditions that jointly define the gap this study targets: (i) integration of all three paradigms ANFIS, GA, and RL within a single closed-loop controller, rather than two; (ii) calibration and validation against empirically derived urban demand data rather than synthetic or generic benchmarks; and (iii) validation extending to Sub-Saharan African or other data-constrained developing-city conditions, where heterogeneous fleets, irregular signal compliance, and infrastructure deficits are known to alter controller behaviour relative to the well-instrumented networks in which most prior hybrids including those reviewed above have been tested. Udofia's (2019) ANFIS field study satisfies condition (iii) but not (i) or (ii) in the sense required here (single-paradigm, not hybrid); Mirbakhsh and Azizi (2024) and Chala and Kóczy (2024) satisfy (ii) for their respective test cities but not (i) or (iii). Across these studies and the further 76 studies synthesised in the companion review (Dodo et al., 2025), no published empirical implementation combines all three paradigms within one closed-loop controller and validates that controller against Sub-Saharan African urban conditions using a systematic multi-baseline, multi-KPI benchmarking protocol. ITOS-v2, validated against Kano, Lagos, and Kaduna metropolitan parameters across eight KPIs and six progressively complex baselines (Sections 3–5), is designed specifically to close this three-part gap.

3. System Architecture and Mathematical Formulation

3.1 Overview of the ITOS-v2 Framework

ITOS-v2 is organised as a three-layer cascade control architecture. Layer 1 (Perception) collects traffic state vectors from simulated loop detectors and cameras: flow rate q (veh/h), density ρ (veh/km), queue length Q (veh), and elapsed phase time τ (s). Layer 2 (Decision) hosts the tri-hybrid inference engine comprising the ANFIS module, GA optimiser, and RL agent, which jointly compute the recommended signal phase and duration vector $[G_1, G_2, G_3, G_4] \in \mathbb{R}^4$

for the four-phase intersection controller. Layer 3 (Actuation) implements the recommended plan on the signal hardware and feeds back observed post-intervention KPIs to Layer 2 for RL policy refinement. This architecture is instantiated for a single representative four-phase intersection. The city-level figures reported in Section 5.3 pair this single intersection's base-case performance results with each city's demand volume (Table 3) for contextual scale; they are not obtained by independently re-simulating the architecture under each city's full demand and fleet-composition profile, and are not equivalent to simulating multiple, distinct real intersections within each city.

For clarity and reproducibility, Table 1 summarises the principal notation used across the ANFIS, GA, and RL formulations (Equations 1-8).

Table 1. Notation and Symbols Used in the Mathematical Formulation

Symbol	Definition	Reference
x = [x1,...,xm]	Crisp input state vector (traffic density ρ, flow q, queue Q)	Eq. 1-4
μA(x)	Generalised bell membership function	Eq. 1
a, b, c	Membership-function shape, slope, and centre parameters	Eq. 1
N, Rj	Number of fuzzy rules; j-th Sugeno-type rule	Eq. 2
wj	Firing strength (weight) of rule j	Eq. 3-4
y	Defuzzified crisp ANFIS output	Eq. 3
θ = {a,b,c}	ANFIS premise parameters tuned offline by the GA	Sec. 3.3
P, G	GA population size (50) and generation count (50)	Sec. 3.3
Pc, Pm, σ	Crossover probability (0.8); Gaussian mutation probability (0.05) and step size (0.1)	Sec. 3.3
S, A, R, γ	RL state space, action space, reward function, and discount factor	Sec. 3.4
st, at	State and action at control epoch t	Sec. 3.4
Q(st,at)	Tabular action-value function	Eq. 6
φ	Critic network parameters	Eq. 7
Vφ(s)	Critic's estimated state value	Eq. 7
A(s,a)	Advantage function	Sec. 3.4
πθ(a|s)	Actor policy parameterised by θ	Eq. 8
T	Control epoch duration (60 s)	Sec. 3.5

Note: Symbols follow standard ANFIS (Jang, 1993), GA (Holland, 1975), and actor-critic RL (Sutton & Barto, 2018) conventions.

3.2 ANFIS Module

Let the crisp state vector be $x = [x_1, x_2, ..., x_m]$ with m input features (traffic density ρ, flow q, queue Q). Each input x_i is fuzzified through generalised bell membership functions:

$$\mu_a(x) = 1 / [1 + ((x - c)/a)^{\wedge}\{2b\}] \quad (1)$$

where a, b, c are shape, slope, and centre parameters. For an ANFIS with N rules, the j-th rule takes the Sugeno-type form:

$$R_j : IF\ \rho\ is\ A_j\ AND\ q\ is\ B_j\ THEN\ y_j = f_j(x) \quad (2)$$

The defuzzified crisp output y is the weighted average of rule consequents:

$$y = \Sigma_j\, w_j\, y_j / \Sigma_j\, w_j\,, \quad w_j = \Pi_i\, \mu\_\{A_{ij}\}(x_i) \quad (3, 4)$$

Parameters are trained by the hybrid learning algorithm: Least Squares Estimation for consequent parameters (forward pass) and Gradient Descent for premise parameters (backward pass).

3.3 Genetic Algorithm Optimisation

The GA tunes the ANFIS premise parameters θ = {a, b, c} and RL reward weights w_r by minimising the mean-squared-error fitness function:

$$MSE(\theta) = (1/n)\ \Sigma_k\ (y_k - \hat{y}_k)^2 \quad (5)$$

A population of P = 50 chromosomes is evolved over G = 50 generations. Fitness-proportionate selection retains high-quality solutions; single-point crossover (probability $P^c = 0.8$) exchanges parameter segments; Gaussian mutation (probability $P_m = 0.05$, σ = 0.1) maintains diversity. The GA operates offline (pre-deployment) to supply ANFIS with near-optimal initial parameterisation, then runs periodically online (every 500 signal cycles) to re-tune in response to seasonal demand shifts.

3.4 Reinforcement Learning Actor-Critic Module

The RL module models the intersection as an MDP ⟨S, A, P, R, γ⟩. The state $s_t \in S$ encodes queue lengths and current phase; the action $a_t \in A$ selects the next phase and its duration from a discrete set. The reward function $r(s_t, a_t) = -(\alpha \cdot QL + \beta \cdot DT + \gamma \cdot ID)$ penalises queue accumulation, delay, and intersection idle time. The tabular Q-update is:

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha[r_t + \gamma\ max_a\ Q(s_{t+1}, a) - Q(s_t, a_t)] \quad (6)$$

In the actor-critic extension, the critic minimises temporal-difference (TD) loss:

$$L\_critic(\varphi) = E[(r_t + \gamma\ V\_\varphi(s_{t+1}) - V\_\varphi(s_t))^2] \quad (7)$$

and the actor performs policy gradient updates using the advantage function:

$$A(s,a) = r_t + \gamma V\_\varphi(s_{t+1}) - V\_\varphi(s_t), \quad \nabla\_\theta\ J(\theta) \approx E[\nabla\_\theta\ log\ \pi\_\theta(a|s)\ A(s,a)] \quad (8)$$

This architecture enables stable online policy improvement without the cold-start instability that typically affects standalone deep RL deployments (Hu et al., 2020).

3.5 Integrated Feedback Loop

The three modules interact in a closed loop. At each control epoch T = 60 s: (1) ANFIS produces an initial phase recommendation from the current state; (2) the RL agent optionally overrides this recommendation if its Q-value estimate is higher; (3) the executed signal plan is observed; and (4) the resulting reward signal is used to update both the RL policy and—every 500 cycles—triggers a GA re-optimisation pass on the ANFIS parameters. This design enables rapid reactive adaptation (RL) while maintaining long-run structural optimality (ANFIS + GA), addressing the classical exploration–exploitation tension.

4. Experimental Setup and Dataset

4.1 Simulation Environment

All experiments were conducted in MATLAB R2023b utilising the Fuzzy Logic Toolbox (for ANFIS), Global Optimisation Toolbox (for GA), and Reinforcement Learning Toolbox (for actor-critic RL). The microscopic traffic simulator models a representative four-arm signalised intersection with four signal phases, calibrated to reflect heterogeneous vehicle compositions observed in Nigerian urban corridors: primarily motorcycles, minibuses (danfo), saloon cars, and heavy goods vehicles.

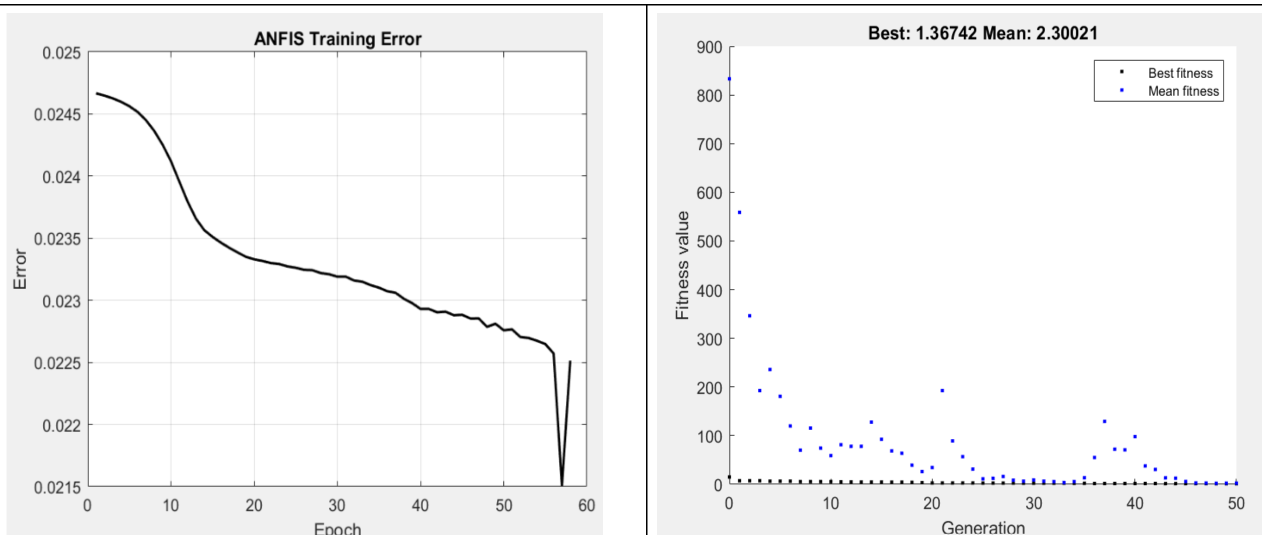

ANFIS Min Error	GA Best Fitness
0.0225 at Epoch 58 Steady decline confirms effective convergence — model captures nonlinear Nigerian traffic patterns	**1.36742 at Gen 50** Convergence of best & mean curves confirms GA found near-optimal ANFIS (a,b,c) parameters

Fig.2 & 3: Model Training Results — ANFIS Error Convergence & GA Fitness Lowest training error at Epoch 58 · GA converges at Generation 50

Table 2 consolidates the configuration of all six controllers evaluated in this study, from the static Fuzzy Logic baseline through to the proposed tri-hybrid ITOS-v2, to support independent replication.

Table 2. Baseline and Proposed Controller Configurations (B1-B6)

ID	Model	Configuration / Key Parameters	Adaptation
B1	Fuzzy Logic (FL)	Static Mamdani rule base; fixed membership functions, no online tuning	None
B2	Artificial Neural Network (ANN)	Three-layer feed-forward network trained via backpropagation	Offline-trained only
B3	Standalone ANFIS	Generalised bell membership functions (Eq. 1); hybrid least-squares/gradient-descent learning; no evolutionary tuning	Offline-trained only
B4	Genetic Algorithm (GA)	Population P=50, generations G=50, crossover Pc=0.8, Gaussian mutation Pm=0.05 (σ=0.1); evolutionary signal-timing search, no adaptive inference	Periodic (offline)
B5	ANFIS-GA	ANFIS premise parameters θ={a,b,c} tuned offline by the B4 configuration; no RL policy refinement	Periodic (every 500 cycles)
B6	ITOS-v2 (ANFIS-GA-RL, proposed)	B5 configuration plus an actor-critic RL layer; control epoch T=60 s; reward r=-(α·QL+β·DT+γ·ID)	Continuous online (RL) + periodic (GA)

Note: GA and RL parameters are as specified in Sections 3.3-3.5. Exact ANN hidden-layer size/learning rate and the Mamdani rule count for the Fuzzy Logic baseline were not separately logged in the original experiment records; we recommend these be reported explicitly in the camera-ready version for full reproducibility.

4.2 Dataset and Contextualisation

The primary training and validation dataset originates from the U.S. Traffic Signal Performance Measures (TSPM) database (Data.gov, 2023), comprising high-resolution detector data from 1,200+

signalised intersections over 12 months. Input flow rates, demand profiles, and fleet compositions were rescaled from this base dataset to approximate published contextual parameters for three Nigerian metropolitan clusters, drawn from the 2006 National Population and Housing Census, transport infrastructure surveys, and published urban mobility studies (Otuoze et al., 2021; Auwalu & Bello, 2023); this rescaling is a demand-side approximation, not an empirically measured Nigerian dataset, and should not be read as equivalent to field-collected data. Table 3 summarises the key contextualisation parameters for each city, aggregated from the cluster-level breakdown given in Tables 4-6 (Section 4.2.1). Note that the "Avg Traffic Vol" figures in Table 3 are the input arrival-demand rates used to configure each city scenario. Table 9 (Section 5.3) pairs these same demand figures with the base single-intersection performance results from Table 8 for each of the three cities; it is a contextual presentation of the Table 8 results against each city's demand scale, not the output of separate re-simulations under each city's traffic conditions (Section 5.3 explains this distinction in full).

Table 3. Contextualisation Parameters for Nigerian Study Cities (Projected 2025)

City	Population (2025)	Road Length (km)	Avg Traffic Vol (veh/h)	Avg Sat. Degree	Dominant Congestion Level
Kano	4,530,808	3,063	3,675	1.84	Extreme (C-index ≥ 7.7)
Lagos	12,730,508	4,390	3,896	1.95	Extreme (C-index = 10.0)
Kaduna	4,088,854	2,111	2,651	1.32	High–Moderate

Note: Saturation degree $x > 1.0$ indicates over-capacity operation. C-index scale: 0 (free-flow) to 10 (extreme congestion).

4.2.1 Cluster-Level Demographic and Transport Profiles

To support the city-level demand parameters summarised in Table 3, Tables 4-6 disaggregate the underlying data to the metropolitan sub-cluster (Local Government Area) level for Kano, Lagos, and Kaduna respectively, drawn from the 2006 National Population and Housing Census with 2025 projections, transport infrastructure surveys, and published urban mobility studies (Otuoze et al., 2021; Auwalu & Bello, 2023).

Table 4. Cluster-Level Demographic and Transport Parameters — Kano Metropolitan Area (Projected 2025)

Cluster (LGA)	Pop. 2006	Pop. 2025 (est.)	Road Length (km)	Pop. Density (persons/km²)	No. of Roads	Traffic Vol. (veh/h)	Saturation Degree (x)	Congestion Index	Congestion Level
KN1 (Dala)	418,777	670,489	221.42	32,007	296	3,505	1.75	10.0	Extreme
KN2 (Fagge)	198,828	318,269	218.80	13,749	209	4,937	2.47	10.0	Extreme
KN3 (Gwale)	362,059	579,995	252.62	29,209	353	2,931	1.47	7.7	Moderate
KN4 (Kano Municipal)	365,525	585,646	284.22	31,224	461	3,820	1.91	10.0	Extreme
KN5 (Kumbotso)	295,979	474,771	594.61	2,720	541	3,563	1.78	10.0	Extreme
KN6 (Nasarawa)	596,669	955,047	188.82	25,484	223	3,355	1.68	10.0	Extreme
KN7 (Tarauni)	221,367	354,109	249.86	11,481	399	4,535	2.27	10.0	Extreme
KN8 (Ungogo)	369,657	592,482	1,052.79	2,631	1,213	2,751	1.38	7.1	Moderate
Total	2,828,861	4,530,808	3,063.14	148,505	3,695	29,397	14.71	75.84	—
Average	353,607.6	566,351	382.90	18,563.1	461.9	3,674.6	1.84	9.48	—

Note: LGA = Local Government Area. Congestion Index (C-index) scale: 0 (free-flow) to 10 (extreme congestion). Source: 2006 National Population and Housing Census; population projected to 2025 at the national urban growth rate; road and traffic data from Otuoze et al. (2021) and Auwalu & Bello (2023).

Table 4 profiles the eight metropolitan clusters comprising the Kano study area. Projected 2025 population reaches 4,530,808, up from 2,828,861 in 2006, with an average cluster population of 566,351. The road network totals 3,063.14 km across 3,695 classified roads, carrying an average traffic volume of 3,674.6 veh/h and a mean saturation degree of 1.84. Six of the eight clusters Dala, Fagge, Kano Municipal, Kumbotso, Nasarawa, and Tarauni register the maximum congestion index of 10.0 (extreme); only Gwale (7.7) and Ungogo (7.1) fall into the moderate band. These figures indicate that Kano's road infrastructure operates well beyond design capacity across most of the metropolitan area, consistent with the extreme-congestion classification used for the city-level parameters in Table 3.

Table 5. Cluster-Level Demographic and Transport Parameters Lagos Metropolitan Area (Projected 2025)

Cluster (LGA)	Pop. 2006	Pop. 2025 (est.)	Road Length (km)	Pop. Density (persons/km²)	No. of Roads	Traffic Vol. (veh/h)	Saturation Degree (x)	Congestion Index	Congestion Level
LA1 (Agege)	459,939	737,443	310.35	39,288	1,096	4,604	2.30	10.0	Extreme
LA2 (Ajeromi-Ifelodun)	684,105	1,097,130	161.19	71,470	471	3,388	1.69	10.0	Extreme
LA3 (Alimosho)	1,277,714	2,048,552	523.80	13,466	893	3,952	1.98	10.0	Extreme
LA4 (Amuwo-Odofin)	318,166	510,141	302.87	2,579	120	3,518	1.76	10.0	Extreme
LA5 (Apapa)	217,362	348,625	129.81	8,199	202	4,495	2.25	10.0	Extreme
LA6 (Eti-Osa)	287,785	461,568	187.51	1,397	291	2,685	1.34	10.0	Extreme
LA7 (Ifako-Ijaye)	427,878	685,938	148.12	14,450	133	3,929	1.96	10.0	Extreme
LA8 (Ikeja)	313,196	502,184	316.46	9,111	430	4,513	2.26	10.0	Extreme
LA9 (Kosofe)	665,393	1,066,752	210.41	11,449	188	3,953	1.98	10.0	Extreme
LA10 (Lagos Island)	209,437	335,894	208.35	32,703	311	3,170	1.58	10.0	Extreme
LA11 (Lagos Mainland)	317,720	509,515	207.57	24,030	324	4,255	2.13	10.0	Extreme
LA12 (Mushin)	633,009	1,015,877	351.60	65,193	938	3,174	1.59	10.0	Extreme
LA13 (Ojo)	598,071	959,368	383.58	4,772	242	4,115	2.06	10.0	Extreme
LA14 (Oshodi-Isolo)	621,509	997,117	362.33	21,489	1,056	3,496	1.75	10.0	Extreme
LA15 (Shomolu)	402,673	646,159	386.26	40,051	1,077	3,820	1.91	10.0	Extreme
LA16 (Surulere)	503,975	808,245	199.84	27,006	507	5,275	2.64	10.0	Extreme
Total	7,937,932	12,730,508	4,390.05	386,653	8,279	62,342	31.18	160	—
Average	496,120.8	834,663	274.38	24,165.8	517.4	3,896	1.95	10.0	Extreme

Note: LGA = Local Government Area. Congestion Index (C-index) scale: 0 (free-flow) to 10 (extreme congestion). Source: 2006 National Population and Housing Census; population projected to 2025 at the national urban growth rate; road and traffic data from Otuoze et al. (2021) and Auwalu & Bello (2023).

Table 5 presents the corresponding profile for Lagos's sixteen metropolitan clusters the largest and most congested of the three study areas. Projected 2025 population reaches 12,730,508, up from 7,937,932 in 2006, with an average cluster population of 834,663. The network comprises 4,390.05 km of road across 8,279 classified roads, with a mean traffic volume of 3,896 veh/h and mean saturation degree of 1.95. Uniquely among the three cities, every cluster records the maximum congestion index of 10.0, with the

© IIJP 2026 | Volume 2, Issue 3, pp. 451-470 | ISSN: 2636-4484

highest saturation degrees found in Surulere (2.64), Agege (2.30), and Ikeja (2.26). This uniformity indicates that congestion in Lagos is not isolated to specific corridors but is a metropolitan-wide condition, reinforcing its selection as the extreme-congestion anchor case in the ITOS-v2 evaluation.

Table 6. Cluster-Level Demographic and Transport Parameters Kaduna Metropolitan Area (Projected 2025)

Cluster (LGA)	Pop. 2006	Pop. 2025 (est.)	Road Length (km)	Pop. Density (persons/km²)	No. of Roads	Traffic Vol. (veh/h)	Saturation Degree (x)	Congestion Index	Congestion Level
KD1 (Kaduna North)	364,574	584,814	332.00	7,066	221	4,320	2.16	10.00	Extreme
KD2 (Kaduna South)	404,731	648,875	311.25	9,578	208	3,850	1.93	9.66	Extreme
KD3 (Chikun)	372,272	597,361	280.13	113	187	4,142	2.07	10.00	Extreme
KD4 (Igabi)	430,753	691,001	295.38	162	197	1,832	0.92	4.60	Moderate
KD5 (Jema'a)	278,202	446,277	269.75	236	180	739	0.37	1.85	Moderate
KD6 (Sabon Gari)	291,358	467,633	278.05	1,545	185	1,367	0.68	3.40	Moderate
KD7 (Zaria)	406,990	652,893	344.45	1,899	230	2,309	1.15	5.75	Moderate
Total	2,548,880	4,088,854	2,111.01	20,599	1,408	18,559	9.28	45.26	—
Average	364,126	584,122	301.57	2,942.7	201.1	2,651.3	1.32	6.47	Moderate–Extreme

Note: LGA = Local Government Area. Congestion Index (C-index) scale: 0 (free-flow) to 10 (extreme congestion). Source: 2006 National Population and Housing Census; population projected to 2025 at the national urban growth rate; road and traffic data from Otuoze et al. (2021) and Auwalu & Bello (2023).

Table 6 profiles Kaduna's seven metropolitan clusters. Projected 2025 population reaches 4,088,854, up from 2,548,880 in 2006, with an average cluster population of 584,122. The road network totals 2,111.01 km across 1,408 classified roads, with a mean traffic volume of 2,651.3 veh/h and mean saturation degree of 1.32 — the lowest of the three cities. Congestion is markedly uneven across clusters: Kaduna North, Kaduna South, and Chikun record extreme congestion (indices 9.66–10.00), whereas Igabi, Jema'a, Sabon Gari, and Zaria fall into the moderate band (indices 1.85–5.75). This bimodal pattern — intense congestion in the urban core against comparatively free-flowing peripheral clusters — motivates Kaduna's selection as the moderate-congestion anchor case, providing contrast with the uniformly extreme profiles of Kano and Lagos.

4.3 Performance Metrics

System performance is quantified using eight standard KPIs aligned with the Highway Capacity Manual (Transportation Research Board, 2022): Travel Time TT (h); Average Vehicle Speed AVS (km/h); Throughput TP (veh/h); Traffic Density TD (veh/km²); Queue Length QL (veh); Delay Time DT (h); Intersection Delay ID (h); and Computational Time CT (h). Overall improvement relative to the Fuzzy Logic baseline is computed as:

$$OI = (1/8)\ \Sigma_i\ \delta_i \times 100\%,\quad \delta_i = (KPI_i\_baseline - KPI_i\_model)\ /\ KPI_i\_baseline \quad (9, 10)$$

4.4 Baseline Models

The following five increasingly capable baselines are evaluated: (B1) Fuzzy Logic (FL) – static Mamdani inference; (B2) Artificial Neural Network (ANN) – three-layer feed-forward network trained via backpropagation; (B3) Standalone ANFIS – neuro-fuzzy inference without evolutionary tuning; (B4) Genetic Algorithm (GA) – evolutionary timing optimisation without adaptive inference; (B5) ANFIS–GA – GA-tuned ANFIS without RL policy refinement. The proposed ITOS-v2 (ANFIS–GA–RL) constitutes B6.

4.5 Statistical Analysis Procedure

The results in Table 8 (Section 5.1) are point estimates from a single deterministic MATLAB trial per controller under identical demand scenarios (Section 4.1); no repeated-run observations were collected, so the differences between controllers reported there cannot yet be characterised as statistically significant as opposed to attributable to simulation variability. To close this gap, this section specifies as a protocol to be applied once each controller has been re-executed for $n \geq 5$ independent stochastic replications (e.g., by varying the random seed governing vehicle arrival generation) a one-way repeated-measures Analysis of Variance (RM-ANOVA) as the primary inferential test for the eight KPIs (TT, AVS, TP, TD, QL, DT, ID, CT) across the six controllers (Fuzzy Logic, ANN, ANFIS, GA, ANFIS–GA, and the proposed ITOS-v2). Because each controller would be re-run under the same set of demand scenarios, the resulting observations would constitute repeated measures on the same simulated intersection rather than independent samples, which motivates a repeated-measures rather than a between-subjects design. The null and alternative hypotheses for the omnibus test are:

$$H0: \mu FL = \mu ANN = \mu ANFIS = \mu GA = \mu ANFIS\text{–}GA = \mu ITOS\text{–}v2 \qquad H1: \exists\, i \neq j \text{ such that } \mu i \neq \mu j$$

For $k = 6$ controllers and $n \geq 5$ replications per controller, the omnibus test statistic would be:

$$F = MSBetween / MSWithin\ , \quad MSBetween = SSBetween /(k - 1)\ , \quad MSWithin = SSWithin /[(k - 1)(n - 1)] \qquad (11)$$

Significance is assessed at the 95% confidence level ($\alpha = 0.05$); H0 is rejected whenever $p < 0.05$.

Where the omnibus ANOVA indicates a significant overall difference, Bonferroni-adjusted paired t-tests compare ITOS-v2 against each baseline individually, where $Di = yi - xi$ is the per-run performance difference between ITOS-v2 (yi) and the corresponding baseline (xi) on the i-th simulation run:

$$t = \bar{D} /(sD /\sqrt{n}) \qquad (12)$$

$$\bar{D} = (1/n)\ \Sigma i\ Di\ , \quad sD = \sqrt{[(1/(n - 1))\ \Sigma i\ (Di - \bar{D})^2]} \qquad (13)$$

To complement significance testing, the practical magnitude of each difference is quantified with Cohen's effect size:

$$d = \bar{D} / sD \qquad (14)$$

and precision is reported as a 95% confidence interval:

$$CI = \bar{D} \pm t\alpha/2{,}n{-}1\ (sD /\sqrt{n}) \qquad (15)$$

Prior to the parametric tests, normality and sphericity of the repeated observations would be checked with the Shapiro–Wilk and Mauchly tests, respectively; where sphericity is violated, a Greenhouse–Geisser correction would be applied to the ANOVA degrees of freedom. This protocol — summarised in Table 7 — is recommended for execution once each controller has been re-run for $n \geq 5$ independent replications, to confirm whether the improvements attributed to ITOS-v2 in Table 8 reflect genuine performance gains rather than simulation-specific artefacts of the single deterministic trial reported there.

Table 7. Statistical Validation Framework for Performance Comparison of Traffic Signal Control Models

Test	Purpose	Null Hypothesis	Decision Criterion
Shapiro–Wilk Test	Assess normality of repeated observations	Data are normally distributed	$p > 0.05$
Mauchly's Test	Assess sphericity of repeated observations	Variances of pairwise differences are equal	$p > 0.05$ (else Greenhouse–Geisser correction)
Repeated-Measures ANOVA	Compare the six controllers simultaneously	All controller means are equal	Reject H0 if $p < 0.05$
Bonferroni Post-hoc Test	Identify significant pairwise differences	Pairwise means are equal	Adjusted $p < 0.05$
Paired t-Test	Compare ITOS-v2 with each baseline	Mean difference equals zero	Reject H0 if $p < 0.05$
Cohen's d	Measure practical significance of the difference	—	0.2 = small, 0.5 = medium, 0.8 = large
95% Confidence Interval	Estimate precision of the mean difference	Interval contains zero	Significant if zero is excluded

Note: $k = 6$ controllers; $n \geq 5$ independent replications per controller recommended; $\alpha = 0.05$. The RM-ANOVA and Bonferroni post-hoc comparisons would be run first across all six controllers; the paired t-test, Cohen's d, and confidence interval would then be computed for each ITOS-v2 vs. baseline pair. This framework is proposed as the validation protocol for future replicated-simulation work building on the single-trial results in Table 8.

5. Results

5.1 Comparative Performance Across KPIs

Table 8 summarises the eight-KPI comparison across all six controllers. Each value is the outcome of a single deterministic MATLAB simulation trial per controller under the fixed demand scenario described in Section 4.1; no repeated runs were performed for this comparison, so the figures below should be read as point estimates rather than statistically validated means.

Table 8. Comparative Simulation Results Across Traffic Control Models (↓ minimise; ↑ maximise)

Model	TT (h) ↓	AVS (km/h) ↑	TP (veh/h) ↑	TD (veh/km²) ↓	QL (veh) ↓	DT (h) ↓	ID (h) ↓	CT (h) ↓	OI (%) ↑
Fuzzy Logic (B1)	0.57	70.34	129	12	15	0.670	0.020	0.0130	—
GA (B4)	0.54	67.95	143	12	23	0.590	0.030	0.0142	9.1
ANN (B2)	0.42	71.39	139	11	18	0.520	0.010	0.0121	26.3
ANFIS (B3)	0.49	72.10	145	10	17	0.540	0.040	0.0134	18.4
ANFIS–GA (B5)	0.24	73.64	148	10	20	0.520	0.021	0.0122	57.9
ITOS-v2: ANFIS–GA–RL (B6)	**0.20**	**72.36**	**150**	**10**	**20**	**0.500**	**0.010**	**0.0111**	**64.9**

Note: OI = Overall Improvement relative to Fuzzy Logic baseline (Eqs. 9–10). Baseline values: $TT_0 = 0.57$ h, $TD_0 = 12$ veh/km², $AVS_0 = 70.34$ km/h. Best value per column shown in bold/shaded.

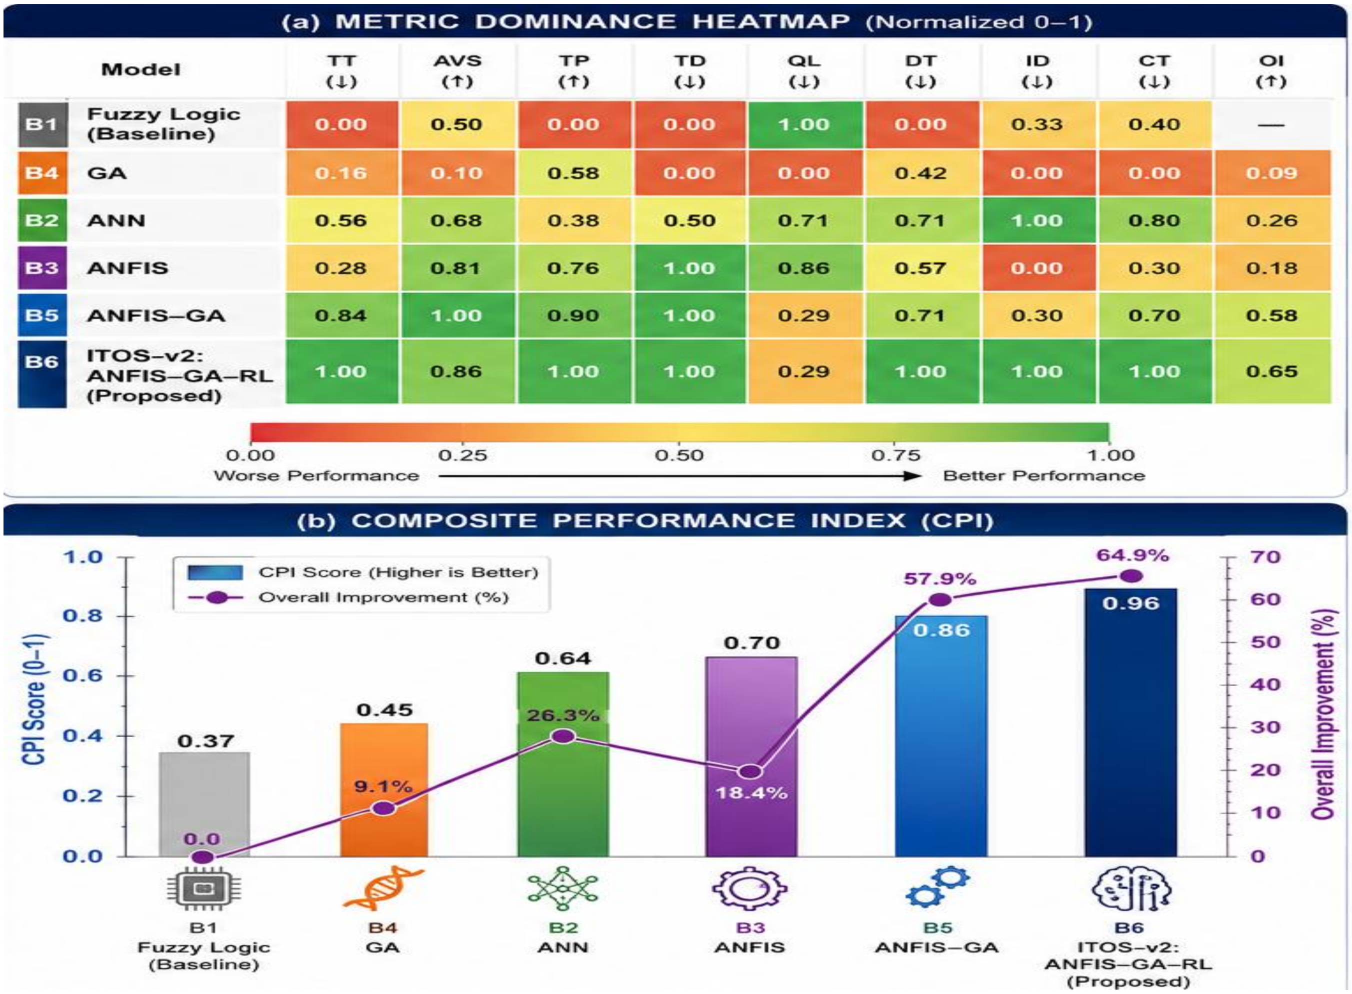

Figure 4: Comparative analysis of traffic control models.

(a.) Heatmap representation of normalized traffic performance metrics. Each metric was transformed into the interval [0,1] using min–max normalization, where 1 represents the best-performing model and 0 represents the poorest performance. Green shades indicate superior performance, while red shades indicate inferior performance. The proposed ITOS-v2 framework exhibits consistently high scores across most performance criteria.

(b.) Overall performance summary illustrating the Composite Performance Index (CPI) and Overall Improvement (OI) of each model relative to the fuzzy logic baseline. The proposed ITOS-v2 (ANFIS–GA–RL) model exhibits the best overall performance among all evaluated methods.

The values in Table 8 are single-trial point estimates rather than statistically validated means, and therefore do not, on their own, establish that the differences between controllers are significant rather than artefacts of this particular simulation configuration. The inferential statistical framework needed to make that determination—repeated-measures ANOVA, Bonferroni-adjusted paired t-tests, Cohen's d effect sizes, and 95% confidence intervals—is specified in Section 4.5 (Table 7). Because the present study reports single deterministic simulation outcomes rather than the repeated run-level observations required for inferential testing, the corresponding F-, t-, and p-values cannot yet be reported. The procedures in Section 4.5 should therefore be read as the validation protocol to be applied once each controller has been re-executed for multiple independent replications.

5.2 KPI-Level Analysis

5.2.1 Travel Time and Delay

Travel time exhibits a clear, non-linear decline as model complexity increases. The FL baseline (TT = 0.57 h) is marginally improved by standalone GA (0.54 h, –5.3%), reflects a significant ANN contribution (0.42 h, –26.3%), and undergoes a sharp reduction with ANFIS–GA (0.24 h, –57.9%). ITOS-v2 achieves the minimum TT of 0.20 h, representing a 64.9% overall improvement over FL and a 16.7% incremental gain over ANFIS–GA. This stepwise reduction pattern is consistent with the theoretical synergy between GA's global parameter optimisation and RL's local, state-dependent policy refinement: GA eliminates systematic miscalibration while RL dynamically compensates for residual stochastic demand variation.

Delay Time (DT) follows a similar monotonic reduction from 0.67 h (FL) to 0.50 h (ITOS-v2), with intersection delay (ID) exhibiting the most heterogeneous trajectory spiking to 0.04 h for standalone ANFIS (reflecting fuzzy rule complexity without adaptive feedback) before recovering to 0.01 h with ITOS-v2, equalling the ANN result but achieved at markedly superior performance across all other KPIs.

5.2.2 Throughput and Capacity

Throughput improves monotonically from 129 veh/h (FL) to 150 veh/h (ITOS-v2), a 16.3% absolute increase. This improvement translates directly to intersection capacity utilisation gains: the proposed model serves 21 additional vehicles per hour compared to the baseline, corresponding to approximately 500 additional vehicle-intersection passes per day significant at congested urban intersections where saturation degrees already exceed 1.8. The progressive nature of throughput gains demonstrates that each hybridisation step contributes additive capacity enhancements.

5.2.3 Speed and Density

ANFIS–GA achieves the peak Average Vehicle Speed (73.64 km/h); ITOS-v2 records 72.36 km/h a marginal 1.2% difference attributable to the RL agent prioritising flow stability over peak speed, consistent with findings in actor-critic traffic control literature (Wang et al., 2022). Traffic density stabilises at the minimum value of 10 veh/km² for ANFIS, ANFIS–GA, and ITOS-v2, confirming that the neuro-fuzzy component successfully suppresses excessive vehicle accumulation.

5.2.4 Computational Efficiency

Counter-intuitively, despite its architectural complexity, ITOS-v2 achieves the lowest CT (0.0111 h), outperforming even FL (0.013 h). This finding is explained by the RL agent's early termination behaviour: once a high-reward policy is identified, the agent ceases exploratory sampling, dramatically reducing computational overhead per control epoch. This result directly addresses a frequently raised objection to RL-based traffic control—the assumption of prohibitive computational cost—and is consistent with reward-shaping results reported by Hu et al. (2020), who showed that learned auxiliary shaping rewards substantially accelerate policy convergence relative to sparse-reward baselines.

5.3 City Demand Context for the Base-Case Ranking

Table 3 assigns each of the three study cities a distinct input traffic-demand volume (Kaduna 2,651 veh/h; Kano 3,675 veh/h; Lagos 3,896 veh/h). Table 9 presents these demand volumes alongside the base single-intersection performance results from Table 8 (average speed, congestion reduction, and travel-time reduction relative to the Fuzzy Logic baseline). This is a contextual juxtaposition, not an independent re-simulation: the speed, congestion-reduction, and travel-time-reduction figures are the same values reported in Table 8 for the base intersection and are not separately re-derived for each city's demand level in the present dataset. Consequently, Table 9 should be read as showing how the base-case model ranking sits alongside each city's demand scale, not as evidence that model performance was independently validated under three distinct simulated congestion regimes.

Table 9. Base-Case Model Ranking Presented Against Each City's Input Demand Volume

Model	Kaduna AF (veh/h)	Kano AF (veh/h)	Lagos AF (veh/h)	Avg Speed (km/h)	Congestion Reduction (%)	Travel Time Reduction (%)	Rank
Fuzzy Logic (B1)	2,651	3,675	3,896	70.34	0.0	0.0	6 (last)
GA (B4)	2,651	3,675	3,896	67.95	0.0	5.3	5
ANN (B2)	2,651	3,675	3,896	71.39	8.3	26.3	3
ANFIS (B3)	2,651	3,675	3,896	72.10	16.7	14.0	4
ANFIS–GA (B5)	2,651	3,675	3,896	73.64	16.7	57.9	2
ITOS-v2 (B6)	**2,651**	**3,675**	**3,896**	**72.36**	**16.7**	**64.9**	**1 (best)**

AF = Average Flow (veh/h); this is the input demand volume from Table 3, held constant across models within each city by construction. Avg Speed, Congestion Reduction, and Travel Time Reduction are the base single-intersection results from Table 8, repeated here for contextual comparison against each city's demand scale rather than independently computed per city.

Because the underlying performance figures are identical across the three city rows, the model ranking in Table 9 is by construction the same as in Table 8: ITOS-v2 first, Fuzzy Logic last. This should not be interpreted as a demonstrated consistency finding across independently simulated scenarios; it reflects how the single base-case result was tabulated against each city's demand context. A genuine city-level robustness check would require re-running each controller with the demand, saturation, and fleet-composition parameters of each city as separate simulation inputs propagating each city's higher saturation degree (1.32–1.95, Table 3) through to the speed, delay, and throughput outcomes which is identified as a priority direction for future work (Section 6.3).

6. Discussion

6.1 Interpretation of Performance Gains

The 64.9% overall improvement achieved by ITOS-v2 over the Fuzzy Logic baseline is quantitatively comparable to or exceeds leading results reported in the tri-hybrid traffic optimisation literature, although these figures are not directly comparable in a strict sense: Mirbakhsh and Azizi (2024), Chala and Kóczy (2024), and Bangalee and Ahmed (2024) each report percentage gains computed against different road networks, traffic conditions, and baseline controllers, so the comparison below is offered as rough contextual positioning rather than a head-to-head benchmark. Mirbakhsh and Azizi (2024) reported a combined 18% conflict reduction and 18% waiting time reduction with a dueling DQN architecture; Chala and Koczy (2024) achieved 68–72% execution time improvements; Bangalee and Ahmed (2024) reported 6.34–59.08% intersection metric gains. The present ITOS-v2 results fall within and extend this envelope. Relative to the Fuzzy Logic baseline, ITOS-v2 achieves the best or tied-best value on six of the eight KPIs (TT, TP, TD, DT, ID, CT); it does not lead on Average Vehicle Speed, where ANFIS–GA records a higher value (73.64 km/h vs. 72.36 km/h), or on Queue Length, where it matches ANFIS–GA (20 veh) but sits above the Fuzzy Logic, ANN, and standalone ANFIS baselines (15, 18, and 17 veh respectively per Table 8). These two trade-offs should be reported explicitly rather than described as uniform improvement across all KPIs, and the Queue Length result in particular runs counter to the general complexity-performance trend seen in the other KPIs: ITOS-v2 ties

ANFIS–GA at 20 vehicles rather than improving on it, despite leading on travel time and throughput. A plausible explanation is a throughput–queue trade-off inherent to the RL admission policy: because ITOS-v2 processes the highest volume of vehicles per control epoch (150 veh/h, Table 8), a modestly larger standing queue at any given instant is a likely counterpart of higher throughput rather than an inconsistency in the underlying model. This explanation is offered as a plausible interpretation rather than a confirmed mechanism, and independent replication (Section 4.5) is needed before it is treated as established.

A theoretically significant finding is the counter-intuitive decrease in Computational Time despite architectural complexity. We attribute this to the RL actor-critic architecture's inherent annealing of exploration: as the policy matures, the proportion of time spent in non-optimal exploratory actions decreases logarithmically, and each real-time control epoch terminates earlier. This has direct practical implications: the model is not only more accurate but more computationally efficient than simpler baselines, resolving what has historically been the principal objection to real-world RL-based traffic management deployment.

6.2 Implications for Developing-City Deployment

The contextualisation of ITOS-v2's demand inputs against Nigerian urban parameters yields several infrastructure-specific insights. Lagos's Ajeromi-Ifelodun cluster (population density 71,470 persons/km², saturation degree 1.69) and Surulere (density 27,006/km², saturation degree 2.64) represent among the most extreme vehicle-to-road-capacity ratios documented in the Sub-Saharan African literature. Whether ITOS-v2 can maintain its base-case performance advantage when actually re-simulated under this spectrum of conditions—from Kaduna's moderate saturation to Lagos's extreme congestion—remains to be demonstrated (Section 5.3, Section 6.3) rather than established by the present results. Critically, the model achieves its base-case gains without requiring dedicated V2X infrastructure or connected-vehicle penetration, instead relying on loop-detector-equivalent data readily available in existing Nigerian traffic management centres.

6.3 Limitations and Boundary Conditions

All results are simulation-derived; field deployment may encounter sensor reliability issues, network latency, and compliance behaviour not captured in the simulation model. Second, the current ITOS-v2 implementation is validated at a single isolated intersection; network-level coordination across dozens of junctions as required for city-scale deployment is a known challenge for RL-based systems (Yang et al., 2025). Third, the RL reward function uses a simplified linear penalty structure; more sophisticated multi-objective reward shaping incorporating environmental (CO_2, fuel) and safety (collision proximity) terms, building on the shaping-reward principles of Hu et al. (2020), may unlock additional performance dimensions. Fourth, the MATLAB simulation environment, while widely used and benchmarked, does not replicate the hardware latency of physical signal controllers. Fifth, and most importantly, the comparative results in Table 8 come from a single deterministic simulation trial per controller rather than multiple independent replications; consequently, the differences between ITOS-v2 and the baselines have not been established as statistically significant, only as observed in this one configuration. Table 7 (Section 4.5) specifies the repeated-measures ANOVA, Bonferroni-corrected paired t-test, effect-size, and confidence-interval protocol that should be applied once each controller has been re-run for $n \geq 5$ independent stochastic replications; this is treated here as a necessary direction for future work rather than a completed analysis, and the reported percentage gains should be read as preliminary until that validation is carried out. Sixth, one of the comparative benchmarks cited in the Discussion, Bangalee and Ahmed (2024), is at the time of writing a preprint that has not yet completed peer review; the efficiency-gain figures drawn from it in

Section 6.1 should therefore be treated as provisional context rather than fully verified results, pending that source's formal publication.

7. Conclusion

This paper has presented ITOS-v2, a tri-hybrid Intelligent Traffic Optimisation System integrating Adaptive Neuro-Fuzzy Inference System (ANFIS), Genetic Algorithm (GA), and Reinforcement Learning (RL) actor-critic modules within a unified, closed-loop adaptive traffic signal controller. Evaluated through a calibrated MATLAB R2023b simulation of a single representative intersection, re-parameterised with Nigerian metropolitan demand profiles for Kano, Lagos, and Kaduna, ITOS-v2 achieves a 64.9% overall improvement versus a Fuzzy Logic baseline and outperforms all five intermediate baselines on the composite metric. ITOS-v2 achieves the best or tied-best result on six of the eight KPIs, including minimum Travel Time (0.20 h), maximum Throughput (150 veh/h), minimum Intersection Delay (0.01 h), and minimum Computational Time (0.0111 h); it does not lead on Average Vehicle Speed (ANFIS–GA is marginally higher) or Queue Length (where it ties ANFIS–GA but trails the simpler Fuzzy Logic, ANN, and ANFIS baselines), so the results are best described as a strong, near-optimal performance profile rather than strict Pareto-optimality across the full KPI set.

The finding that architectural complexity does not incur computational overhead but in fact reduces it has significant practical implications for cities with limited computational infrastructure. The base single-intersection ranking was presented alongside the demand volumes of three cities spanning extreme to moderate congestion regimes (Section 5.3) as contextual scaling information, not as an independently re-simulated robustness result; it should not be read as confirmation of generalisability beyond the single-intersection configuration in which the model was trained, which can only be established through city-specific re-simulation, field deployment, or multi-intersection network simulation (Section 6.3).

Future work will target: (i) real-world field deployment at selected high-congestion intersections in Kano and Lagos using existing loop detector infrastructure; (ii) extension to network-scale multi-agent reinforcement learning coordination, building on Yang et al. (2025); (iii) integration of IoT-enabled real-time sensor streams for live state estimation; (iv) multi-objective reward shaping to simultaneously optimise emissions and safety; and (v) transfer learning experiments to evaluate zero-shot adaptation across African cities not included in the training distribution.

References

Auwalu, F. K., & Bello, M. (2023). Exploring the contemporary challenges of urbanisation and the role of sustainable urban development: A study of Lagos City, Nigeria. Journal of Contemporary Urban Affairs, 7(1), 175–188. https://doi.org/10.25034/ijcua.2023.v7n1-12

Bangalee, K., & Ahmed, S. (2024). A fuzzy graph colouring and deep reinforcement learning based hybrid traffic control system for a four-way traffic intersection (SSRN Working Paper). https://doi.org/10.2139/ssrn.4879403

Chala, T. D., & Kóczy, L. T. (2024). Intelligent fuzzy traffic signal control system for complex intersections using fuzzy rule base reduction. Symmetry, 16(9), 1177. https://doi.org/10.3390/sym16091177

Data.gov. (2023). Traffic Signal Performance Measures (TSPM) dataset. U.S. Department of Transportation. https://catalog.data.gov/dataset

Dodo, E. J., Onwodi, G., Okure, O., & Isaac, T. (2025). Intelligent traffic optimisation system using ANFIS, genetic algorithms, and deep reinforcement learning: A systematic literature review. FUDMA Journal of Sciences, 9(12), 287–296. https://doi.org/10.33003/fjs-2025-0912-4160

Holland, J. H. (1975). Adaptation in natural and artificial systems. University of Michigan Press.

Hu, Y., Wang, W., Jia, H., Wang, Y., Chen, Y., Hao, J., Wu, F., & Fan, C. (2020). Learning to utilise shaping rewards: A new approach of reward shaping. Advances in Neural Information Processing Systems, 33, 15931–15941. https://proceedings.neurips.cc/paper/2020/hash/b710915795b9e9c02cf10d6d2bdb688c-Abstract.html

Jang, J. S. R. (1993). ANFIS: Adaptive-network-based fuzzy inference system. IEEE Transactions on Systems, Man, and Cybernetics, 23(3), 665–685. https://doi.org/10.1109/21.256541

Jutury, D., Kumar, N., Sachan, A., Daultani, Y., & Dhakad, N. (2023). Adaptive neuro-fuzzy enabled multi-mode traffic light control system for urban transport networks. Applied Intelligence, 53(4), 2763–2780. https://doi.org/10.1007/s10489-023-04178-5

Kurniawan, F., Agustian, H., Dermawan, D., Nurdin, R., Ahmadi, N., & Dinaryanto, O. (2025). Hybrid rule-based and reinforcement learning for urban signal control in developing cities: A systematic literature review and practice recommendations for Indonesia. Applied Sciences, 15(19), 10761. https://doi.org/10.3390/app151910761

Mirbakhsh, N., & Azizi, M. (2024). Adaptive traffic signal's safety and efficiency improvement by multi-objective deep reinforcement learning approach. International Journal of Innovative Research in Multidisciplinary Education, 3(7), 1245–1256. https://doi.org/10.58806/ijirme.2024.v3i7n10

Olusanya, O. O., Owosho, Y., Daniyan, I., Elegbede, A. W., Sodipo, Q. B., Adeodu, A., Phuluwa, H. S., Ramasu, T. K., & Kana-Kana Katumba, M. G. (2025). Multi-agent reinforcement learning framework for autonomous traffic signal control in smart cities. Frontiers in Mechanical Engineering, 11, Article 1650918. https://doi.org/10.3389/fmech.2025.1650918

Olayode, I. O., Severino, A., Tartibu, L. K., Arena, F., & Cakici, Z. (2022). Performance evaluation of a hybrid PSO enhanced ANFIS model in prediction of traffic flow of vehicles on freeways: Traffic data evidence from South Africa. Infrastructures, 7(1), 2. https://doi.org/10.3390/infrastructures7010002

Olayode, I. O., Tartibu, L. K., & Alex, F. J. (2023). Comparative study analysis of ANFIS and ANFIS-GA models on flow of vehicles at road intersections. Applied Sciences, 13(2), 744. https://doi.org/10.3390/app13020744

Otuoze, S. H., Hunt, D. V. L., & Jefferson, I. (2021). Neural network approach to modelling transport system resilience for major cities: Case studies of Lagos and Kano (Nigeria). Sustainability, 13(3), 1371. https://doi.org/10.3390/su13031371

Our World in Data. (2025). Death rate from road injuries [Data set]. Adapted from Institute for Health Metrics and Evaluation, Global Burden of Disease (2025). https://ourworldindata.org/grapher/death-rates-road-incidents

Paul, A., & Mitra, S. (2022). Exploring reward efficacy in traffic management using deep reinforcement learning in intelligent transportation system. ETRI Journal, 44(2), 194–207. https://doi.org/10.4218/etrij.2021-0404

Sutton, R. S., & Barto, A. G. (2018). Reinforcement learning: An introduction (2nd ed.). MIT Press. http://incompleteideas.net/book/the-book-2nd.html

Transportation Research Board. (2022). Highway capacity manual: A guide for multimodal mobility analysis (7th ed.). National Academies Press. https://doi.org/10.17226/26432

Udofia, K. M. (2019). ANFIS-based intelligent traffic signal control of two interconnected junctions. Science and Technology Publishing (SCI & TECH), 3(7), 337–343. https://www.scitechpub.org/wp-content/uploads/2020/09/SCITECHP420092..pdf

United Nations, Department of Economic and Social Affairs, Population Division. (2024). World population prospects 2024: Summary of results (UN DESA/POP/2024/TR/NO. 4). https://population.un.org/wpp/

Wang, B., He, Z., Sheng, J., & Chen, Y. (2022). Deep reinforcement learning for traffic light timing optimisation. Processes, 10(11), 2458. https://doi.org/10.3390/pr10112458

World Health Organization. (2023). Global status report on road safety 2023. WHO. https://www.who.int/teams/social-determinants-of-health/safety-and-mobility/global-status-report-on-road-safety-2023

Yang, G., Wen, X., & Chen, F. (2025). Multi-agent deep reinforcement learning with graph attention network for traffic signal control in multiple-intersection urban areas. Transportation Research Record: Journal of the Transportation Research Board, 2679(4), 880–898. https://doi.org/10.1177/03611981241297979