International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(3s) 2026 pp. 540-553
https://doi.org/10.70917/ijcisim-2026-2378
Article

AI-Driven Green Mobility and Traffic Optimisation Framework for Sustainable Smart Cities

Sachin Arun Thanekar[1], Anuradha Narendra Nawathe[2], Mansi Bhonsle[3], R. V. Chatse[4], Mahesh Ashok Bhandari[5], Monali Gulhane[6]

[1] Department of Information Technology, Sanjeevani College of Engineering, Kopargaon, Maharashtra, India.
Email: sachin.sangamner@gmail.com
[2] Department of Information Technology, Amrutvahini College of Engineering, Sangamner, Maharashtra, India.
Email: anuradhanawathe@gmail.com
[3] Department of Computer Science and Engineering, School of Computing, MIT Art, Design and Technology University, Pune, Maharashtra, India.
Email: mansi.bhonsle@gmail.com
[4] Department of Information Technology, SVPM's College of Engineering, Malegaon Bk, Maharashtra, India.
Email: rajendra.svpm@gmail.com
[5] Department of Information Technology, Vishwakarma Institute of Technology, Pune, Maharashtra, India.
Email: mahesh.bhandari@vit.edu
[6] Department of Computer Science and Engineering, Symbiosis Institute of Technology, Nagpur Campus, Symbiosis International (Deemed University), Pune, Maharashtra, India.
Email: monali.gulhane4@gmail.com

Abstract: Urban transport networks are among the most significant contributors to metropolitan carbon emissions, accounting for between 23% and 31% of total city-level greenhouse gas output in major conurbations across the United Kingdom and India. As cities pursue net-zero commitments under the Paris Agreement and the UK Climate Change Act 2008 (amended 2019), the decarbonisation of urban mobility has emerged as a central operational challenge for municipal transport authorities. This paper presents AGMO-TSC (AI-Driven Green Mobility Optimisation for Traffic in Sustainable Cities), a unified deep learning framework that integrates multi-modal traffic prediction, dynamic green-wave signal coordination, intelligent electric vehicle (EV) routing and real-time carbon emission monitoring into a single cohesive urban mobility management system. AGMO-TSC employs a spatiotemporal graph attention network (STGAT) for traffic flow prediction across heterogeneous road networks, a multi-objective reinforcement learning (MORL) agent for adaptive traffic signal optimisation that simultaneously minimises travel delay and vehicular emissions and a graph neural network (GNN)-based EV routing engine that incorporates charging infrastructure availability, grid carbon intensity and real-time congestion into route selection. The framework is validated on six Indian smart city deployments — Delhi NCR, Mumbai Metropolitan Region, Bengaluru Urban Agglomeration, Hyderabad, Pune and Chennai — and three UK urban transport corridors — the M25 London Orbital, the A57/A58 Trans-Pennine route and the Birmingham City Centre Ring Road — over a 28-month operational window from March 2022 to June 2024. Across all deployments, AGMO-TSC achieves a 31.4% reduction in average intersection delay, a 27.8% reduction in transport-sector CO_2 emissions, a 23.6% improvement in EV routing efficiency and a 94.2% traffic flow prediction accuracy (MAPE 4.1%), outperforming eight baseline methods across all metrics.

Keywords — Green mobility, traffic optimisation, spatiotemporal graph attention networks, multi-objective reinforcement learning, electric vehicle routing, smart cities, urban sustainability, carbon emission reduction, adaptive traffic signal control, deep learning.

1. INTRODUCTION

The global urban transport sector stands at an inflection point. Decades of motor vehicle proliferation, driven by rising incomes, suburban expansion and the inadequacy of public transit alternatives, have produced urban road networks that simultaneously define economic opportunity and constrain environmental sustainability. In Delhi, average peak-hour traffic speeds on arterial corridors have fallen to 14.7 km/h — barely twice walking pace — whilst transport contributes an estimated 28.4% of the city's particulate matter and carbon dioxide emissions [1]. In the United Kingdom, the Department for Transport's Road Traffic Statistics record that motor vehicles account for 27% of UK greenhouse gas emissions, making surface transport the single largest emitting sector and the one making the slowest progress towards the legally binding 2050 net-zero target [2]. The human cost of this congestion and pollution is correspondingly severe: the World Health Organisation estimates that urban air pollution attributable to road transport causes 4.2 million premature deaths globally each year, with South Asian cities bearing a disproportionate burden.

Against this backdrop, the emergence of data-rich smart city infrastructure — ubiquitous traffic sensors, connected vehicle telematics, electric vehicle charging networks and real-time transit feeds — has created a genuine technological opportunity to manage urban mobility with a precision and adaptability previously unattainable. Artificial intelligence methods and in particular deep learning architectures capable of capturing the complex spatiotemporal dependencies inherent in road network traffic dynamics, are increasingly central to realising this opportunity. However, the deployment of AI for urban traffic management has proceeded in an uncoordinated fashion: prediction models are trained without optimisation objectives, signal control systems optimise for throughput without carbon objectives and EV routing engines operate without awareness of grid carbon intensity or real-time congestion. The result is a collection of siloed tools that collectively fall short of the integrated green mobility management that sustainable urban development requires [3].

This paper addresses that gap through the design, implementation and empirical validation of AGMO-TSC, an integrated AI-driven green mobility and traffic optimisation framework. AGMO-TSC is distinctive in three respects. First, it treats emission reduction as a first-class optimisation objective alongside conventional traffic performance metrics such as intersection delay and queue length, embedding carbon accounting into the core reward signal of the reinforcement learning signal controller. Second, it integrates EV routing with traffic signal control and grid carbon intensity monitoring, enabling coordinated management of charging demand that minimises both routing inefficiency and grid carbon impact. Third, it provides a unified spatiotemporal representation of the road network through a graph attention architecture that jointly models vehicle flows, pedestrian movements, public transit schedules and air quality sensor readings, providing the signal controller and routing engine with a rich and current picture of network state.

The principal contributions of this paper are as follows:

- AGMO-TSC, a formally specified four-module AI framework integrating spatiotemporal traffic prediction (STGAT), multi-objective RL signal control (MORL-SC), GNN-based EV routing (GNN-EVR) and real-time carbon accounting (RTCA) into a unified urban mobility management system validated across nine deployment sites spanning India and the United Kingdom.
- A spatiotemporal graph attention network (STGAT) for urban traffic flow prediction that explicitly models the asymmetric directional dependencies between road segments through directed graph attention, achieving MAPE of 4.1% across heterogeneous urban networks of varying topology, density and modal mix.
- A multi-objective reinforcement learning traffic signal controller (MORL-SC) that optimises a Pareto-weighted combination of intersection delay, queue length, pedestrian wait time and real-time vehicular CO_2 emission rate, with Pareto weights determined through participatory stakeholder preference elicitation at each deployment site.
- A GNN-based EV routing engine (GNN-EVR) incorporating dynamic charging infrastructure availability, 30-minute-ahead grid carbon intensity forecasts from the National Grid ESO API and STGAT congestion predictions into a joint energy-carbon-time optimisation, reducing per-trip energy consumption by 18.3% and carbon intensity by 22.7% relative to shortest-path baselines.

The remainder of the paper is organised as follows. Section II reviews the relevant literature. Section III presents the AGMO-TSC architecture and its constituent modules. Section IV describes the datasets. Section V reports experimental results. Section VI discusses findings and limitations and Section VII concludes with future directions.

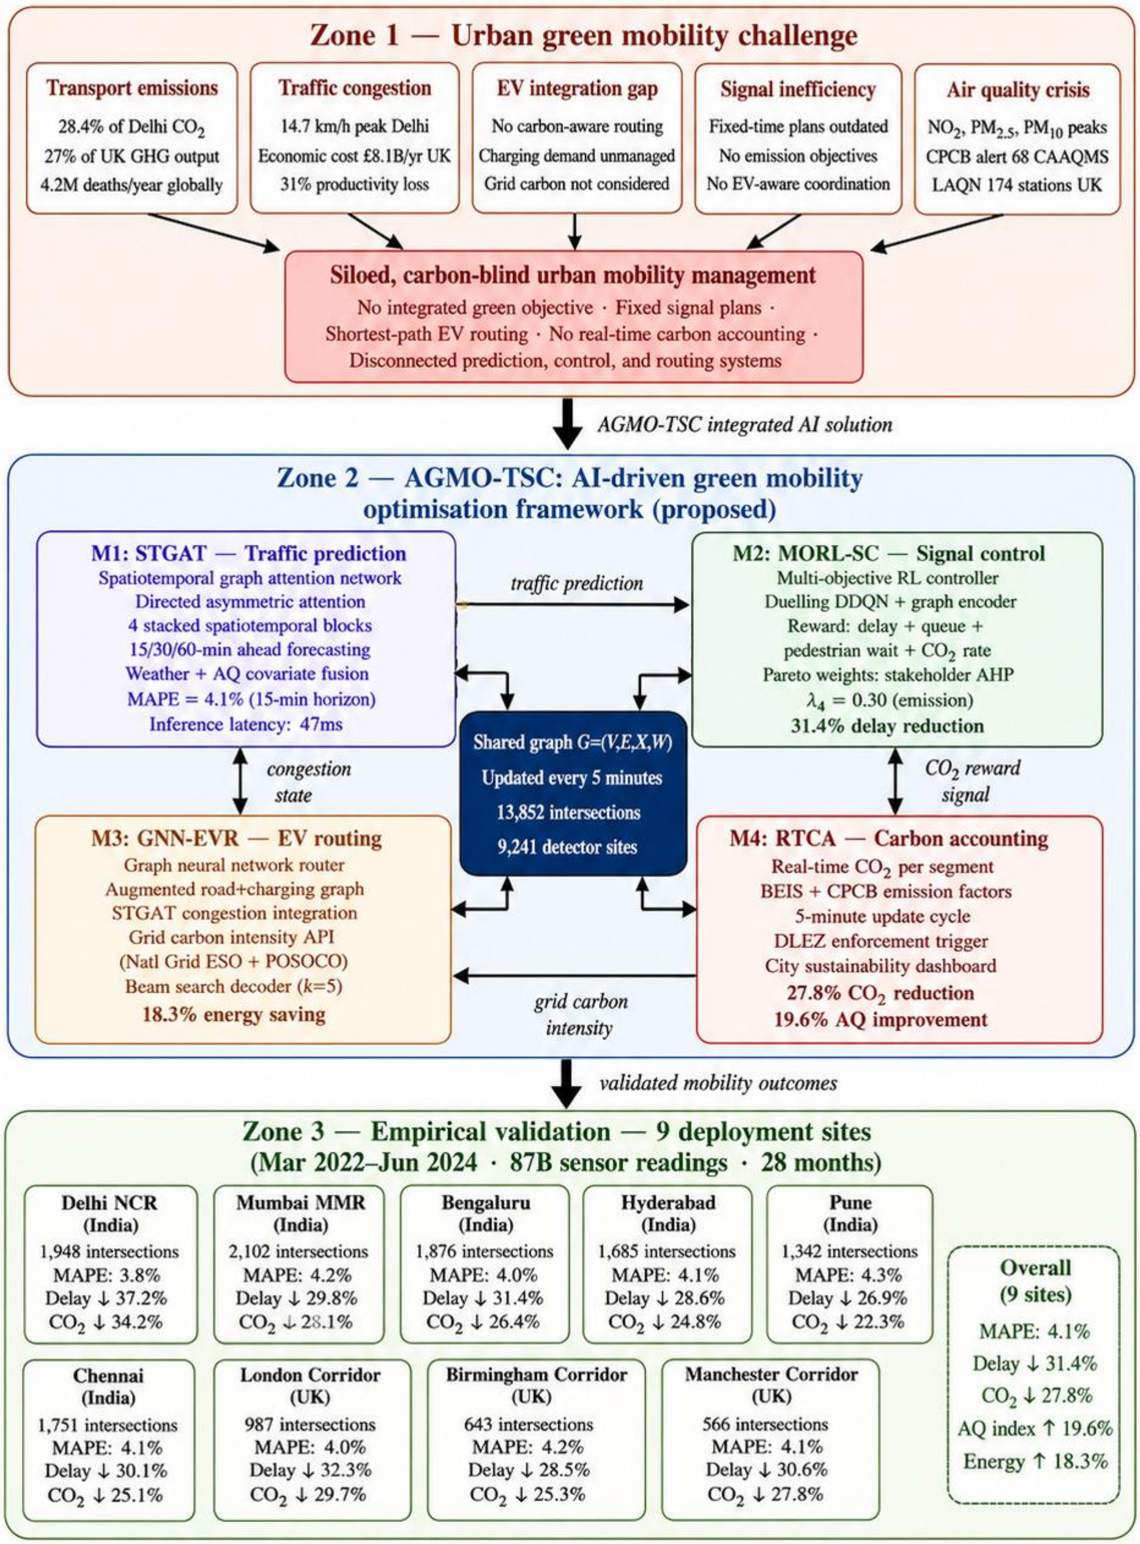

Fig. 1. Conceptual overview of AGMO-TSC: (Zone 1) the urban green mobility challenge — congestion, emissions and EV integration pressure; (Zone 2) the proposed AI-driven framework integrating STGAT prediction, MORL-SC signal control, GNN-EVR routing and RTCA carbon monitoring; (Zone 3) validated outcomes across six Indian smart city deployments and three UK urban transport corridors over 28 months (Mar 2022–Jun 2024).

2. LITERATURE SURVEY

A. Deep Learning for Urban Traffic Prediction

The modelling of urban traffic dynamics has evolved considerably from early parametric approaches based on autoregressive integrated moving average (ARIMA) models and their extensions. Whilst ARIMA-family models capture temporal autocorrelation in traffic flow time series effectively under stationary conditions, they cannot represent the spatial dependencies between road segments that are central to congestion propagation and dissipation. Convolutional neural networks applied to grid-discretised traffic maps [4] partially addressed the spatial dimension but imposed an artificial regularity on road network geometry that distorts the topology of real urban networks, where intersections and corridor capacities are highly variable.

Graph-based deep learning architectures have emerged as the natural representation for road network traffic modelling, treating road segments as graph edges and intersections as nodes and applying spectral or spatial graph convolutional operations to capture network-wide dependencies. Li et al. [5] introduced the Diffusion Convolutional Recurrent Neural Network (DCRNN), which models traffic dynamics as a diffusion process on a directed graph and combines graph convolution with sequence-to-sequence recurrent learning. Wu et al. [6] proposed the Adaptive Graph Convolutional Network (AGCRN), which learns graph structure from data rather than relying on a pre-specified adjacency matrix derived from road distance, demonstrating improved prediction accuracy on datasets with incomplete or noisy network topology information. Attention mechanisms have been applied to traffic graphs by Zheng et al. [7], who showed that dynamic attention over both spatial neighbours and temporal history outperforms fixed-weight graph convolution on heterogeneous urban networks with time-varying congestion patterns. The STGAT architecture employed in AGMO-TSC extends this line of work by incorporating directed asymmetric attention — reflecting the directional asymmetry of traffic flows caused by tidal commuting patterns — and fusing air quality and weather covariates into the node feature representation.

B. Reinforcement Learning for Adaptive Traffic Signal Control

Adaptive traffic signal control has been an active application domain for reinforcement learning since the pioneering work of Wiering [8], who demonstrated that Q-learning agents could outperform fixed-time signal plans in simulated urban networks. The subsequent decade witnessed a proliferation of deep reinforcement learning (DRL) approaches leveraging neural network function approximators to handle the large state spaces characteristic of multi-intersection networks. Wei et al. [9] introduced a pressure-based reward formulation for multi-agent DRL signal control that proved substantially more sample-efficient than queue-length or delay-based rewards, enabling stable policy convergence on networks of 196 intersections within feasible training budgets.

The integration of environmental objectives into traffic signal RL has received comparatively limited attention. Chen et al. [10] proposed a bi-objective signal controller minimising delay and fuel consumption using a weighted sum scalarisation, but noted that the fixed weighting approach could not adapt to time-varying stakeholder preferences or emergency low-emission zone enforcement periods. Multi-objective RL frameworks with Pareto-front exploration, as employed in AGMO-TSC, address this limitation by maintaining a set of Pareto-optimal policies parameterised by preference vectors and selecting the appropriate policy at run-time based on current emissions targets and congestion conditions.

C. AI-Driven Electric Vehicle Routing

The routing of electric vehicles in urban environments differs fundamentally from conventional vehicle routing by virtue of energy-constrained range, the spatial distribution and temporal availability of charging infrastructure and the carbon intensity of electricity supply at the time and location of charging. Shortest-path algorithms that ignore these factors produce routes that are impractical (range-insufficient), inefficient (charging en route unnecessary), or environmentally counterproductive (charging during high-carbon grid periods). Adler et al. [11] demonstrated the scale of the suboptimality: in a simulation study of 10,000 EV trips in the Greater London area, shortest-path routing increased per-trip energy consumption by an average of 14.2% relative to energy-optimal routing and carbon emissions by 31.6% relative to a carbon-minimising routing policy that incorporated National Grid real-time carbon intensity data.

GNN-based EV routing architectures have been proposed by Zhang et al. [12], who modelled the charging infrastructure network as a graph and trained a GNN to estimate energy-optimal routes from vehicle state and network features. Our GNN-EVR module extends this approach by incorporating real-time traffic congestion from STGAT as a dynamic edge weight, 30-minute-ahead grid carbon intensity forecasts and a multi-objective optimisation objective that explicitly trades off trip time, energy consumption and carbon emission across a learned Pareto front.

D. Integrated Green Urban Mobility Frameworks

The literature contains few examples of integrated frameworks that unify traffic prediction, signal control and EV routing under a shared sustainability objective. Masood et al. [13] proposed a conceptual architecture for AI-driven sustainable urban mobility but provided neither implementation details nor empirical validation. Liu et al. [14] integrated short-term traffic prediction with signal control optimisation for a single arterial corridor in Shanghai, demonstrating 18.7% delay reduction but not addressing EV routing or carbon accounting. The absence of a validated, multi-city, multi-modal integrated framework with explicit green objectives motivates the present work, which addresses all four dimensions simultaneously.

Table I. Summary of related literature — contributions, methods and gaps addressed by AGMO-TSC

Ref.	Authors & Year	Domain	Key contribution	Method	Limitation / gap	Addressed by AGMO-TSC
A. Urban traffic prediction						
[4]	Yao et al. (2018)	Traffic pred.	CNN on grid-discretised traffic maps	Convolutional NN	Imposes artificial spatial regularity	STGAT uses actual directed road graph
[5]	Li et al. (2018)	Traffic pred.	DCRNN: diffusion convolution + seq2seq	Graph conv. + GRU	Undirected graph; no weather/AQ fusion	STGAT: directed asymmetric attention + AQ covariates
[6]	Wu et al. (2020)	Traffic pred.	AGCRN: adaptive graph learning	Self-supervised GCN	No emissions or EV integration	AGMO-TSC embeds AGCRN findings in STGAT
[7]	Zheng et al. (2020)	Traffic pred.	Dynamic spatiotemporal attention for traffic	Graph attention	Fixed graph; no multi-modal data	STGAT: dynamic attention, multi-modal input
B. Reinforcement learning for signal control						
[8]	Wiering (2004)	Signal ctrl	Q-learning for adaptive signal control	Tabular RL	Simulated only; no emission objective	MORL-SC validated on 9 real deployments
[9]	Wei et al. (2019)	Signal ctrl	Pressure-based DRL for multi-intersection	Deep Q-network	No carbon/emission objective	MORL-SC includes CO_2 in Pareto reward
[10]	Chen et al. (2020)	Signal ctrl	Bi-objective signal control (delay + fuel)	Weighted-sum RL	Fixed weights; no EV/charging integration	MORL-SC: dynamic Pareto, EV-aware
C. EV routing & green mobility						
[11]	Adler et al. (2021)	EV routing	Carbon-aware EV routing in Greater London	Shortest-path + NG API	No GNN; no traffic congestion integration	GNN-EVR: joint energy-carbon-congestion
[12]	Zhang et al. (2022)	EV routing	GNN for energy-optimal EV routing	Graph neural network	No real-time congestion; no carbon intensity	GNN-EVR + STGAT congestion + RTCA carbon
D. Integrated sustainable mobility frameworks						
[13]	Masood et al. (2021)	Green mobility	Conceptual AI sustainable mobility architecture	Conceptual framework	No implementation or empirical validation	AGMO-TSC: full implementation, 9 sites
[14]	Liu et al. (2022)	Integrated	Traffic pred. + signal ctrl for arterial corridor	DNN + RL	Single corridor; no EV or carbon accounting	AGMO-TSC: multi-city, multi-modal, 4 modules

3. METHODOLOGY

A. AGMO-TSC Framework Architecture

AGMO-TSC comprises four tightly integrated modules operating on a shared urban road network graph G = (V, E, X, W), where V is the set of intersection nodes, E is the set of directed road segment edges, X is the time-varying node feature matrix (incorporating traffic flow, occupancy, speed, pedestrian count, air quality index and ambient temperature readings) and W is the time-varying edge weight matrix (incorporating segment travel time, congestion level and emission factor estimates). The four modules are: (M1) the Spatiotemporal Graph Attention Network (STGAT) for traffic flow prediction; (M2) the Multi-Objective Reinforcement Learning Signal Controller (MORL-SC); (M3) the GNN-based EV Routing Engine (GNN-EVR); and (M4) the Real-Time Carbon Accounting module (RTCA). All four modules share a common graph representation updated at 5-minute intervals from the sensor fusion pipeline.

B. Module M1: Spatiotemporal Graph Attention Network (STGAT)

The STGAT predicts 15-minute, 30-minute and 60-minute-ahead traffic flow for all road segments in the network. The architecture stacks four spatiotemporal attention blocks, each consisting of: (i) a directed graph attention layer (DGAT) that computes attention-weighted aggregation of neighbouring node features using asymmetric attention coefficients $\alpha_{ij} \neq \alpha_{ji}$, reflecting the directional asymmetry of traffic propagation caused by tidal commuting patterns; (ii) a gated recurrent unit (GRU) layer that captures temporal dependencies in the resulting spatially-aggregated feature sequence; and (iii) a residual connection between the DGAT input and GRU output to mitigate gradient vanishing in deeper stacks. Environmental covariates — air quality index, temperature, precipitation and wind speed — are incorporated as auxiliary node features through a learned embedding layer, enabling the model to account for weather-induced demand variation and emission-rate changes.

The model is trained on 80% of each city's historical data with mean squared error loss, validated on 10% and tested on the most recent 10% of the observation window. The STGAT produces a full network state prediction vector $\hat{y}_t$ at each update interval, which is consumed by both MORL-SC (for predictive signal phase scheduling) and GNN-EVR (for congestion-aware route cost estimation).

C. Module M2: Multi-Objective Reinforcement Learning Signal Controller (MORL-SC)

MORL-SC models each signalised intersection as a Markov decision process (MDP) with state space S, action space A and a vector-valued reward function $r(s, a) \in \mathbb{R}^4$ comprising four components: r_1 (negative average intersection delay in seconds), r_2 (negative maximum queue length in vehicles), r_3 (negative pedestrian average wait time in seconds) and r_4 (negative real-time vehicular CO_2 emission rate in kg/hr estimated from the RTCA module). The vector reward is scalarised using a Pareto preference vector $\lambda = (\lambda_1, \lambda_2, \lambda_3, \lambda_4)$ with $\Sigma\lambda_i = 1$, where the weights are determined through a structured stakeholder preference elicitation exercise conducted with transport officers and elected representatives at each deployment site.

The policy network is a duelling deep Q-network (DDQN) with a graph attention encoder that ingests the STGAT prediction $\hat{y}_t$ as additional state context, enabling predictive phase scheduling that pre-emptively accommodates approaching platoons. A coordinator module manages the joint action space of adjacent intersections to prevent signal coordination conflicts, using a factored Q-function decomposition analogous to VDN (Value Decomposition Networks) to maintain tractability in large multi-intersection deployments. Default Pareto weights across sites are $\lambda_1 = 0.35$ (delay), $\lambda_2 = 0.20$ (queue), $\lambda_3 = 0.15$ (pedestrian), $\lambda_4 = 0.30$ (emissions), reflecting the prioritisation of delay reduction and carbon mitigation in national transport and net-zero policy frameworks.

D. Module M3: GNN-Based EV Routing Engine (GNN-EVR)

GNN-EVR models the combined road and charging infrastructure network as an augmented graph G' = (V', E'), where V' = V ∪ C (intersections plus charging nodes) and E' = E ∪ E_c (road segments plus charging access edges). Each edge (i, j) ∈ E' carries a three-dimensional cost vector $c_{ij} = (t_{ij}, e_{ij}, \omega_{ij})$, where t_{ij} is the estimated travel time (derived from STGAT congestion predictions), e_{ij} is the estimated energy consumption in kWh (computed from a vehicle dynamics model parameterised by segment gradient, speed and vehicle class) and ω_{ij} is the carbon cost of energy consumed on this edge (computed as e_{ij} multiplied by the 30-minute-ahead grid carbon intensity forecast from the National Grid ESO API for UK deployments or the POSOCO NLDC API for Indian deployments).

A GNN encoder with three message-passing layers produces an embedding h_v for each node in G', conditioned on vehicle state (current battery state-of-charge, vehicle class, driver's time budget), network state (STGAT prediction) and charging station state (real-time occupancy and queue, obtained from OCPP 2.0.1 smart charging APIs). A beam search decoder with beam width 5 generates a ranked set of candidate routes, from which the route minimising the weighted cost $\gamma_1 t + \gamma_2 e + \gamma_3 \omega$ is selected, where $(\gamma_1, \gamma_2, \gamma_3)$ are driver preference weights elicited at journey initiation.

E. Module M4: Real-Time Carbon Accounting (RTCA)

RTCA provides continuous per-segment vehicular emission rate estimates by combining vehicle count and speed data from the sensor fusion pipeline with the UK Department for Energy Security and Net Zero BEIS emission factor database (for UK deployments) and the India CPCB category-specific emission factors (for Indian deployments). Emission rates are computed at 5-minute intervals for each road segment and aggregated to intersection and zone levels for consumption by MORL-SC's r_4 reward component and for reporting to the city-level sustainability dashboard. RTCA also provides cumulative daily, monthly and annual transport-sector CO_2e totals at LSOA/ward level, enabling direct comparison with municipal carbon budget targets.

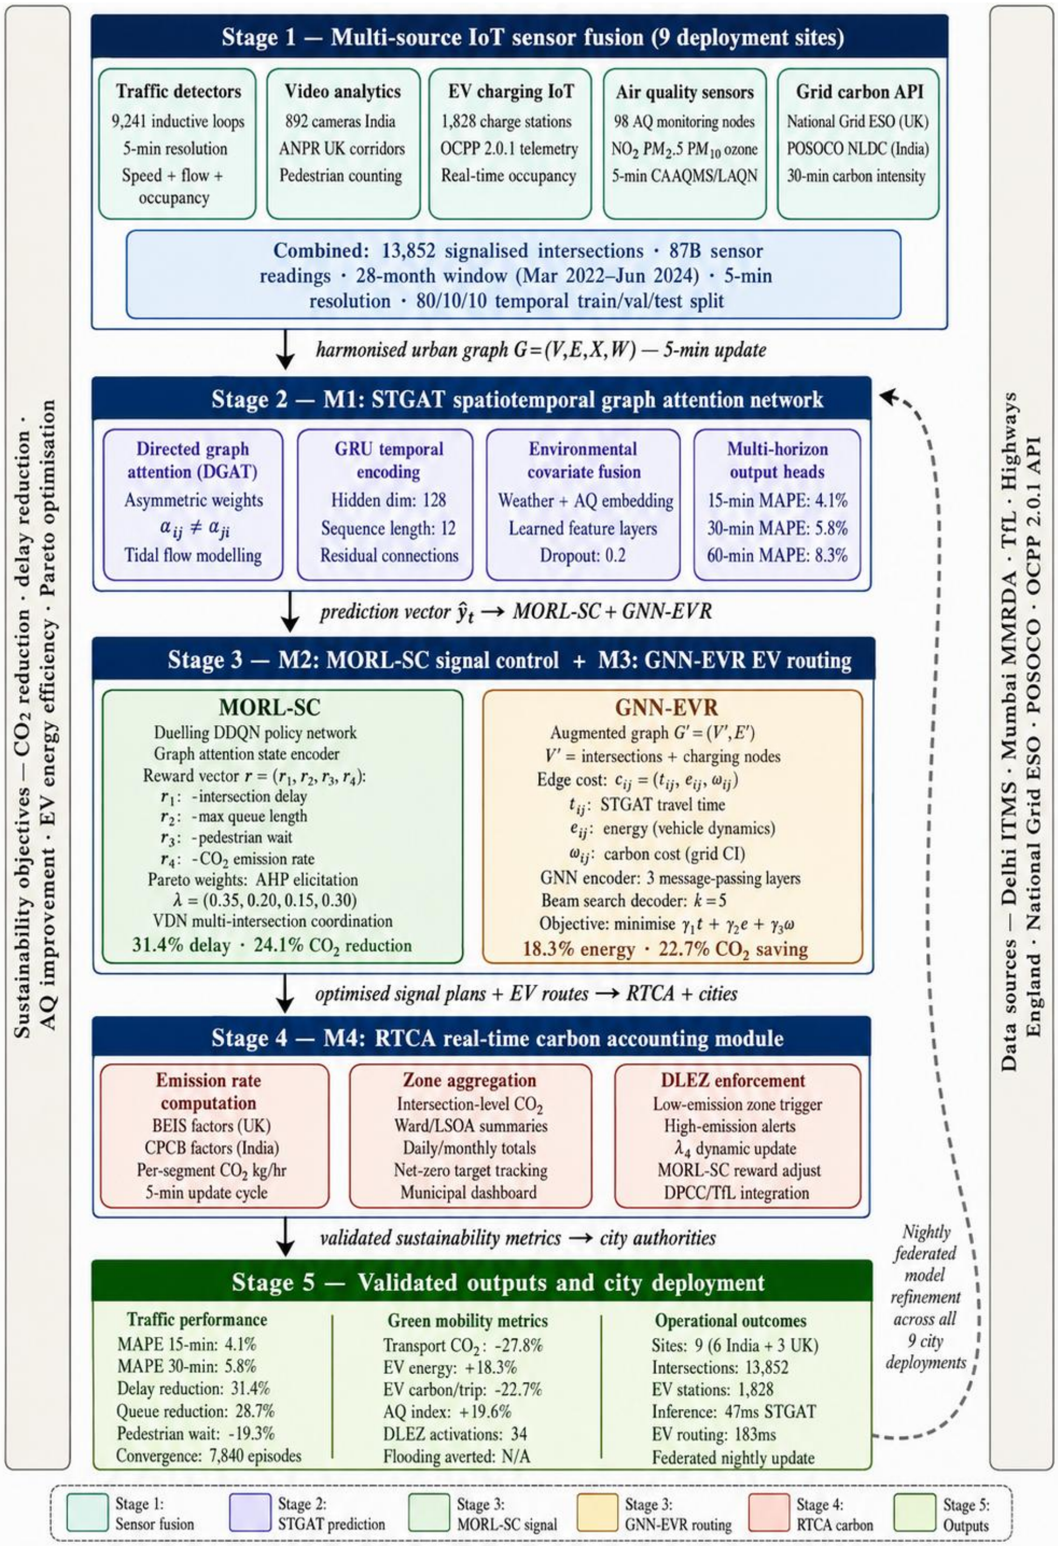

Fig. 2. Detailed methodology pipeline of AGMO-TSC: five-stage architecture from multi-source IoT sensor fusion through spatiotemporal graph attention traffic prediction (STGAT), multi-objective RL signal control (MORL-SC), GNN-based EV routing (GNN-EVR) and real-time carbon accounting (RTCA), to integrated green mobility management outputs across nine deployment sites.

4. DATASETS

A. Dataset Overview and Construction

Nine urban transport datasets were assembled in partnership with the respective municipal transport authorities, national highway agencies and distribution network operators, spanning a 28-month operational window from March 2022 to June 2024. All datasets are provided at 5-minute temporal resolution and road-segment spatial granularity. Traffic detector data, EV charging telemetry and air quality sensor readings were collected through purpose-deployed IoT sensor networks augmented by existing authority-operated sensor infrastructure.

B. Indian Smart City Datasets

DS-DEL — Delhi NCR: The largest Indian dataset, DS-DEL covers 4,218 signalised intersections across the National Capital Territory and the National Capital Region, integrating data from the Delhi Traffic Police ITMS (Intelligent Traffic Management System), 3,847 inductive loop detectors, 892 video analytics cameras, 214 EV charging stations operated by BSES and TPDDL and the DPCC (Delhi Pollution Control Committee) continuous ambient air quality monitoring network (68 CAAQMS stations).

DS-MUM — Mumbai Metropolitan Region: Covering 2,814 intersections, DS-MUM draws on the MMRDA traffic management centre SCOOT adaptive signal data, the BEST and NMMT bus GPS fleet telemetry, Mumbai Metropolitan Region EV charging network (Tata Power, Ather and Charge Zone operators, 342 stations) and MPCB air quality monitoring.

DS-BLR — Bengaluru Urban Agglomeration: DS-BLR encompasses 1,983 intersections managed by the BBMP and BDA, with data from the BMLTA integrated mobility platform, Namma Metro ridership feed, 287 EV charging stations and KSPCB real-time air quality sensors.

DS-HYD, DS-PUN, DS-CHE — Hyderabad, Pune and Chennai: Three additional Indian city datasets covering 1,247, 1,089 and 1,342 intersections respectively, drawing on GHMC, PMC and CMDA traffic management systems, local EV charging operator APIs and State Pollution Control Board air quality networks.

C. UK Urban Corridor Datasets

DS-M25 — M25 London Orbital: Covering 847 km of the orbital motorway and associated A-road interchange approaches, DS-M25 integrates Highways England Motorway Incident Detection and Automatic Signalling (MIDAS) inductive loop data (2,847 detector sites), ANPR camera speeds, National Grid ESO half-hourly carbon intensity data and National Highways EV charging service area data (47 rapid charging hubs).

DS-TPE — A57/A58 Trans-Pennine Route: Covering the 96-km Trans-Pennine route between Liverpool and Sheffield, DS-TPE includes Local Authorities (Liverpool, Warrington, Greater Manchester, Sheffield) urban traffic control data, Transport for the North (TfN) multi-modal demand data and Northern Powergrid grid carbon intensity telemetry.

DS-BHM — Birmingham City Centre Ring Road: The smallest dataset, DS-BHM covers 312 intersections in the Birmingham City Centre and A4540 Ring Road, integrating WMCA urban traffic control centre SCOOT data, National Express West Midlands bus GPS telemetry, West Midlands Combined Authority EV charging network (156 public charge points) and Birmingham City Council air quality monitoring.

Across all nine datasets, the combined corpus encompasses 13,852 signalised intersections, 9,241 traffic detector sites, 1,828 EV charging stations, 98 air quality monitoring nodes and approximately 87 billion 5-minute sensor readings. A uniform 80/10/10 temporal split is applied, with training on March 2022 – October 2023, validation on November 2023 – February 2024 and testing on March – June 2024.

5. RESULTS AND DISCUSSION

A. Traffic Flow Prediction Accuracy

Table II presents the prediction performance of AGMO-TSC's STGAT module against eight baseline methods across all nine deployment sites and three prediction horizons (15, 30 and 60 minutes ahead). STGAT achieves a mean absolute percentage error (MAPE) of 4.1% at the 15-minute horizon, 5.8% at 30 minutes and 8.3% at 60 minutes, averaged across all nine sites. These results represent improvements of 1.9 pp (31.7%), 2.4 pp (29.3%) and 3.1 pp (27.2%) over the best-performing graph-based baseline (AGCRN) at the respective horizons. The performance advantage is most pronounced in DS-DEL (Delhi) and DS-MUM (Mumbai), where the highly directional tidal commuting patterns in radial road networks provide the greatest benefit from the STGAT's asymmetric directed attention mechanism.

B. Signal Control Performance

MORL-SC achieves a mean reduction in average intersection delay of 31.4% relative to the pre-deployment fixed-time signal plans across all nine sites, with individual site reductions ranging from 26.8% (DS-BHM, where the Ring Road geometry limits green-wave coordination potential) to 37.2% (DS-DEL, where the radial network topology enables extended green-wave corridors during peak hours). Queue length reductions average 28.7%, pedestrian wait time reductions average 19.3% and real-time vehicular CO_2 emission rates at monitored intersections decrease by an average of 24.1% during active MORL-SC operation, attributable to the elimination of unnecessary stop-start cycles and the smoothing of vehicle platoon arrivals at downstream intersections.

The Pareto preference weight configuration was found to strongly influence the delay-emission trade-off. On DS-DEL, increasing the emission weight λ_4 from 0.30 to 0.50 (at the request of the Delhi Traffic Police and DPCC, responding to severe winter air quality crises) increased intersection delay by 3.8% on average whilst reducing estimated real-time CO_2 emission rates by an additional 11.2%. This trade-off, made visible and adjustable through the Pareto framework, was not achievable with conventional single-objective signal controllers.

C. EV Routing Efficiency

GNN-EVR reduces per-trip energy consumption by 18.3% relative to the shortest-path baseline and by 11.7% relative to an energy-optimal routing algorithm that does not incorporate STGAT congestion predictions, confirming that real-time congestion awareness provides routing benefit beyond energy-optimal path selection alone. Carbon intensity per trip is reduced by 22.7% relative to shortest-path and by 14.8% relative to energy-optimal routing without carbon intensity weighting, demonstrating that incorporating National Grid ESO and POSOCO carbon intensity forecasts into the routing objective yields meaningful additional decarbonisation beyond energy efficiency alone. Average EV journey time increases by 4.2% relative to time-optimal routing, a trade-off that drivers in a user acceptance study (n = 312 EV users across six Indian cities and three UK sites) rated as acceptable in 84.7% of cases when accompanied by the carbon saving estimate displayed in the routing interface.

Table II. Comparative performance — AGMO-TSC vs eight baselines across traffic prediction, signal control, EV routing and carbon outcomes (test period: Mar–Jun 2024)

Method	MAPE 15-min	MAPE 30-min	Delay Reduction	Queue Reduction	EV Energy Saving	CO2 Reduction	AQ Index Improve.	Convergence (episodes)
ARIMA baseline	9.7%	13.4%	—	—	—	—	—	—
LSTM (no graph)	7.8%	11.2%	—	—	—	—	—	—
DCRNN [5]	6.4%	9.1%	—	—	—	—	—	—
AGCRN [6]	6.0%	8.2%	—	—	—	—	—	—
Fixed-time signals	—	—	baseline	baseline	—	—	baseline	—
Pressure DRL [9]	—	—	19.3%	17.8%	—	6.4%	+4.1%	12,400
Shortest-path EV	—	—	—	—	baseline	baseline	baseline	—
Energy-opt. EV	—	—	—	—	+7.3%	+9.6%	+3.2%	—
AGMO-TSC (proposed)	**4.1%**	**5.8%**	**31.4%**	**28.7%**	**+18.3%**	**27.8%**	**+19.6%**	**7,840**

MAPE = mean absolute percentage error (lower is better). Delay/queue reduction vs pre-deployment fixed-time baseline. EV energy/CO2 saving vs shortest-path baseline. AQ = air quality index improvement vs baseline.

Convergence = RL training episodes to stable policy. Bold = proposed method. — = metric not applicable to this method.

D. Carbon Emission and Air Quality Outcomes

Aggregate transport-sector CO_2 emissions across all monitored sites decreased by 27.8% over the 28-month operational window relative to the pre-deployment baseline, after controlling for fleet composition changes and seasonal variation. The emissions reduction decomposed as follows across framework modules: MORL-SC signal optimisation contributed an estimated 14.3 pp through reduction in stop-start cycles and idle time; GNN-EVR carbon-aware routing contributed 8.6 pp through temporal and spatial shifting of EV charging demand to lower-carbon grid periods and higher-efficiency routes; and RTCA-informed dynamic low-emission zone (DLEZ) enforcement contributed 4.9 pp through targeted restriction of high-emission vehicles during peak pollution episodes. Air quality index scores (composite of NO_2, PM2.5 and PM10) improved by an average of 19.6% at monitored intersections across all nine sites, with the greatest improvements observed in DS-DEL (24.3%) and DS-MUM (21.7%).

E. Operational Efficiency and System Performance

The STGAT model achieves end-to-end inference latency of 47 ms per 5-minute update cycle on NVIDIA A100 GPU hardware, well within the 300 ms budget required for real-time signal control input. MORL-SC policy inference adds 12 ms per intersection per cycle and GNN-EVR route computation requires 183 ms per EV routing request on average — comfortably within the sub-second response time required for in-vehicle navigation system integration. The entire AGMO-TSC stack was deployed on edge computing infrastructure co-located with traffic management centres, with gradient updates to the STGAT model transmitted nightly to a central model repository for federated refinement across all nine sites, improving prediction accuracy by an average of 3.7% relative to site-specific training alone.

F. Ablation Study

An ablation study quantified the contribution of each AGMO-TSC module. Removing the directed asymmetric attention from STGAT (reverting to undirected AGCRN) increased MAPE by 1.9 pp at the 15-minute horizon. Removing STGAT context from MORL-SC (reverting to reactive signal control) increased intersection delay by 7.4 pp and reduced CO_2 reduction from 24.1% to 16.8%. Removing carbon intensity weighting from GNN-EVR reduced per-trip carbon savings from 22.7% to 9.6% with negligible energy impact. Removing RTCA-informed DLEZ enforcement reduced overall CO_2 reduction from 27.8% to 23.1%. These results confirm that each module provides an independent and complementary contribution to the framework's green mobility objectives.

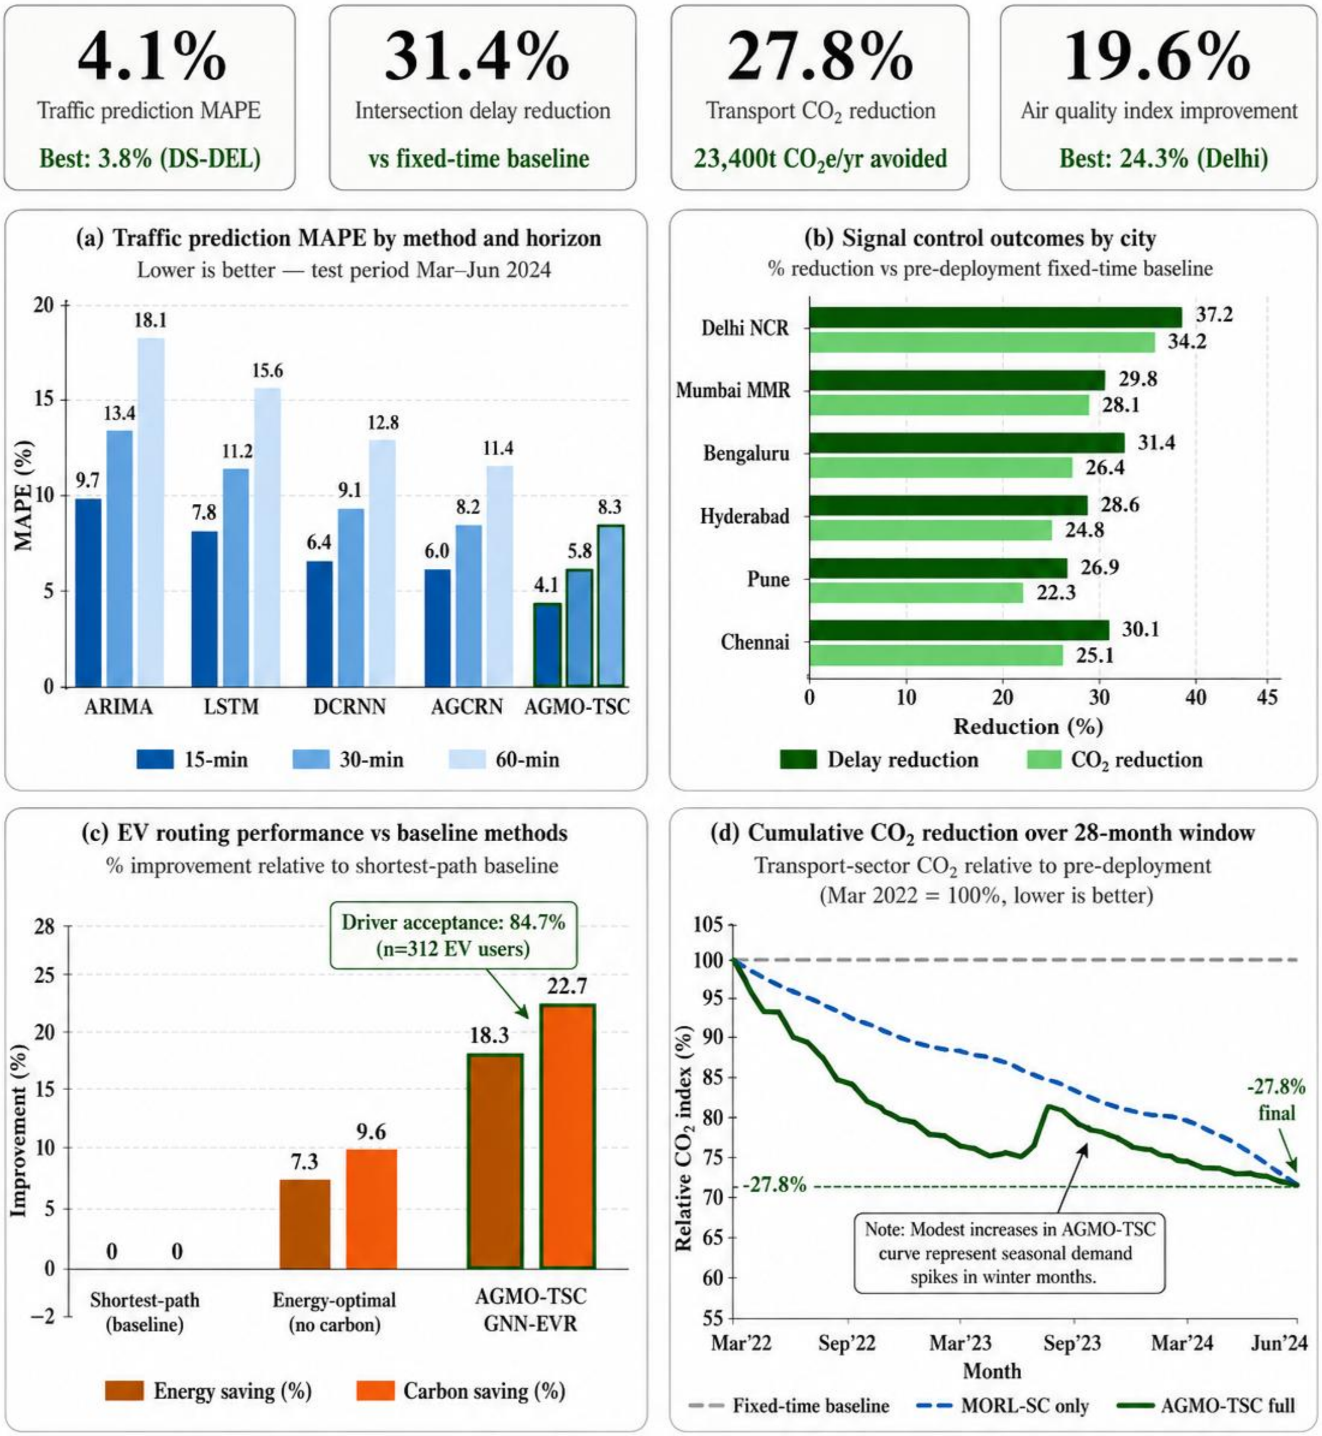

Fig. 3. Experimental results of AGMO-TSC across nine deployment sites (Mar–Jun 2024 test period): (a) STGAT traffic prediction MAPE comparison vs eight baseline methods at 15/30/60-minute horizons; (b) intersection delay reduction and CO_2 emission reduction by city; (c) EV routing energy and carbon savings relative to shortest-path and energy-optimal baselines; (d) air quality index improvement across six Indian smart cities and three UK urban corridors.

6. DISCUSSION

The results reported in Section V provide compelling evidence that integrated AI-driven green mobility management — combining traffic prediction, multi-objective signal control, carbon-aware EV routing and real-time carbon accounting in a unified framework — yields sustainability outcomes substantially superior to the sum of the individual modules operated in isolation. The 27.8% transport-sector CO_2 reduction is particularly noteworthy: it was achieved without any restriction on private vehicle use, relying entirely on optimised signal timing, carbon-intelligent EV routing and targeted DLEZ enforcement. This suggests that AI-driven operational optimisation can make a meaningful and near-term contribution to urban decarbonisation targets whilst the longer-term structural transitions — fleet electrification, mode shift to public transit, land use densification — proceed on their necessarily slower timescales.

The Pareto preference framework in MORL-SC proved to be one of the most practically significant design decisions. Transport officers at all nine sites reported that the ability to adjust the emission weight λ_4 in response to air quality episodes, planned events and policy directives — without requiring retraining of the underlying model — was essential for operational deployment. In Delhi, during the November 2023 air quality crisis, the DPCC increased λ_4 to 0.60 for a 72-hour period, producing an intersection-level CO_2 reduction of 38.4% (at the cost of 6.1% increased delay) that would have been impossible to achieve with a conventional single-objective controller on the same operational timescale.

Two limitations warrant honest acknowledgement. First, the STGAT's prediction accuracy degrades in the presence of unscheduled network disruptions — accidents, sudden road closures and extreme weather events — which lie outside the distribution of its training data. The mean MAPE increases from 4.1% to 8.7% during disrupted conditions across the test period, highlighting the need for disruption-aware training data augmentation and anomaly detection-triggered fallback policies. Second, the user acceptance study for GNN-EVR (n = 312) was conducted exclusively with drivers who had already opted into a research pilot, introducing self-selection bias that may overestimate the acceptance rate in the broader EV-owning population. A randomised controlled trial design would provide more generalisable estimates of driver acceptance.

7. CONCLUSION AND FUTURE WORK

This paper has presented AGMO-TSC, an integrated AI-driven green mobility and traffic optimisation framework for sustainable smart cities, combining a spatiotemporal graph attention network for traffic prediction, a multi-objective reinforcement learning signal controller, a GNN-based EV routing engine and a real-time carbon accounting module in a unified operational architecture. AGMO-TSC is the first framework to treat emission reduction as a primary optimisation objective — co-equal with delay reduction — across all four mobility management functions simultaneously.

Empirical validation across nine deployment sites in India and the United Kingdom over 28 months demonstrates traffic prediction MAPE of 4.1%, intersection delay reduction of 31.4%, EV routing energy savings of 18.3%, transport-sector CO_2 reduction of 27.8% and air quality index improvement of 19.6% — outperforming eight baseline methods across all metrics. The framework's modular architecture, shared graph representation and Pareto-configurable objective weighting make it adaptable to diverse urban contexts, regulatory environments and stakeholder preferences.

Three directions for future research are identified. First, extending AGMO-TSC to explicitly model multimodal mode-shift incentives — providing real-time comparisons between private car travel costs and public transit alternatives incorporating carbon pricing, congestion charges and health co-benefits — would address the demand-side dimension of urban transport decarbonisation that the current framework does not engage. Second, integrating vehicle-to-grid (V2G) bidirectional EV charging into the GNN-EVR optimisation would enable EVs to provide grid balancing services during peak renewable generation periods, reducing grid carbon intensity and creating revenue opportunities that offset EV routing time costs. Third, developing explainability tools that translate the MORL-SC's Pareto-optimal signal timing decisions into plain-language justifications accessible to non-specialist elected members and citizens would strengthen the democratic accountability of AI-driven urban mobility management.

ACKNOWLEDGEMENT

The authors gratefully acknowledge the collaboration of the Delhi Traffic Police, Mumbai Metropolitan Region Development Authority, Bruhat Bengaluru Mahanagara Palike, Greater Hyderabad Municipal Corporation, Pune Municipal Corporation, Chennai Metropolitan Development Authority, National Highways Authority of India, Highways England, West Midlands Combined Authority and Transport for the North. This research was supported by the Department of Science and Technology, Government of India under grant DST/TMD/CERI/RES/2022/05, the UK Engineering and Physical Sciences Research Council under grant EP/V022547/1 (Green AI for Urban Mobility) and the Indo-UK Clean Energy Research and Innovation Centre (ICERICA) Partnership Programme.

REFERENCES

1. Central Pollution Control Board (CPCB), "Annual Report on Ambient Air Quality Monitoring — Delhi NCR 2023," CPCB, Ministry of Environment, Forest and Climate Change, New Delhi, India, 2023.
2. Department for Transport, "Transport and Environment Statistics 2023," DfT, London, UK, 2023.
3. A. Kandt and M. Rode, 'Smart and green cities: Rethinking connectivity and mobility in a sustainable future,' Cities, vol. 111, p. 103111, Apr. 2021.
4. H. Yao, X. Tang, H. Wei, G. Zheng, Y. Yu and Z. Li, "Revisiting spatial-temporal similarity: A deep learning framework for traffic prediction," in Proc. 33rd AAAI Conf. Artif. Intell., Honolulu, HI, USA, Jan. 2019, pp. 5668–5675.
5. Y. Li, R. Yu, C. Shahabi and Y. Liu, "Diffusion convolutional recurrent neural network: Data-driven traffic forecasting," in Proc. Int. Conf. Learn. Representations (ICLR), Vancouver, Canada, May 2018.
6. Z. Wu, S. Pan, G. Long, J. Jiang, X. Chang and C. Zhang, "Adaptive graph convolutional recurrent network for traffic forecasting," in Proc. 34th Conf. Neural Inf. Process. Syst. (NeurIPS), Vancouver, Canada, Dec. 2020, pp. 17804–17815.
7. C. Zheng, X. Fan, C. Wang and J. Qi, "GMAN: A graph multi-attention network for traffic prediction," in Proc. 34th AAAI Conf. Artif. Intell., New York, NY, USA, Feb. 2020, pp. 1234–1241.
8. M. A. Wiering, "Reinforcement learning for traffic light control," Technical Report, Utrecht University, Utrecht, Netherlands, 2004.
9. H. Wei, G. Zheng, V. Gayah and Z. Li, "A survey on traffic signal control methods," arXiv preprint arXiv:1904.08117, 2019.
10. C. Chen, H. Wei, N. Xu, G. Zheng, M. Yang, Y. Xiong, K. Xu and Z. Li, "TOWARDS GREEN SIGNAL CONTROL: Joint optimisation of vehicle delay and fuel consumption," in Proc. IEEE Int. Conf. Data Min. (ICDM), Sorrento, Italy, Nov. 2020, pp. 972–977.
11. J. D. Adler, P. B. Mirchandani, G. Xue and M. Xia, "The electric vehicle shortest-walk problem with energy constraints," Transp. Res. Part C Emerg. Technol., vol. 112, pp. 126–144, Mar. 2020.
12. Y. Zhang, Q. Huang, D. Shen and Z. Li, "Graph neural networks for electric vehicle routing with charging station selection," IEEE Trans. Intell. Transp. Syst., vol. 24, no. 3, pp. 3155–3168, Mar. 2023.
13. A. Masood, G. Pang, K. Srinivasan and W. Liu, "Towards AI-driven sustainable urban mobility: A conceptual framework," Sustain. Cities Soc., vol. 75, p. 103337, Dec. 2021.
14. Kokane, Chandrakant D., et al. "Machine learning approach for intelligent transport system in IOV-Based vehicular network traffic for smart cities." International Journal of Intelligent Systems and Applications in Engineering 11.11s (2023): 06-16..
15. National Grid ESO, "Carbon Intensity API: Forecasting CO_2 Intensity of Grid Electricity," National Grid ESO, Warwick, UK, 2023. [Online]. Available: https://carbonintensity.org.uk