International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(5s) 2026 pp. 853-862
https://doi.org/10.70917/ijcisim-2026-2808
Article

A Carbon-Aware Artificial Intelligence Framework for Achieving Net-Zero Smart Cities

Harsha Jitendra Sarode[1], Megha M. Wankhade[2], Prateeksha Chouksey[3], Vaishali Yogesh Baviskar[4], Uma Patil[5], Bharati Kale[6]

[1] Department of Electronics and Telecommunications Engineering, Nutan Maharashtra Institute of Engineering and Technology, Pune, Maharashtra, India.
Email: sarodeharsha28@gmail.com
[2] Department of Computer Science and Engineering, School of Computing, MIT Art, Design and Technology University, Rajbaug, Loni Kalbhor, Pune, Maharashtra, India.
Email: wmegha13@gmail.com
[3] Department of Artificial Intelligence and Data Science, Keystone School of Engineering, Pune, Maharashtra, India.
Email: prateeksha.chouksey@keystonesoe.in
[4] Department of Artificial Intelligence and Data Science, Vishwakarma Institute of Technology, Pune, Maharashtra, India.
Email: vaishali.baviskar@vit.edu
[5] Department of Computer Science and Engineering (AI & ML), Vishwakarma Institute of Technology, Pune, Maharashtra, India.
Email: uma.patil@vit.edu
[6] Department of Computer Science and Engineering, Visvesvaraya Technological University, Belagavi, Karnataka, India.
Email: bharatibp13@gmail.com

Abstract: Urban areas generate approximately 70 percent of global greenhouse gas emissions, yet most cities still lack the analytical infrastructure needed to track, predict and actively reduce those emissions in real time. Existing approaches to urban carbon management are hampered by static models that cannot adapt to shifting energy mixes, incomplete cross-sector data linkages and the absence of closed-loop control mechanisms that translate carbon predictions into actionable dispatch decisions. This paper introduces the Carbon-Aware Artificial Intelligence Framework for Smart Cities (CAIF-SC), an end-to-end pipeline that integrates long short-term memory networks and gradient-boosted ensemble learning for high-accuracy greenhouse gas prediction with a reinforcement learning dispatch optimiser that translates those predictions into sector-level carbon reduction actions in real time. CAIF-SC is evaluated on a longitudinal panel dataset spanning ten years (2014–2024) across 36 cities using 20 from India and 16 from the United Kingdom using drawn from the Smart Cities Mission, the UK Net Zero Strategy and multiple satellite and sensor networks. The framework achieves a mean absolute error of 2.1 tCO_2e per annum against monitored ground-truth emissions, reduces simulated city-wide greenhouse gas output by a mean of 27.4 percent relative to business-as-usual baselines and projects a mean net-zero attainment year of 2041 across the study cohort. These results outperform five comparator methods using including static regression, standalone deep learning and expert-calibrated optimisation models using on every reported metric. Beyond performance benchmarks, CAIF-SC provides explainable sector attributions and policy-ready counterfactual scenarios, making it directly actionable for municipal climate officers.
Keywords: Carbon-aware AI, net-zero cities, greenhouse gas prediction, reinforcement learning, smart cities, LSTM, XGBoost, urban sustainability, energy dispatch, climate informatics.

I. INTRODUCTION

The climate crisis has focused unprecedented attention on cities as both the primary sources of and the most promising sites for addressing global greenhouse gas emissions. Urban areas are home to slightly over half of the world's population but are responsible for roughly 70 percent of total energy-related carbon dioxide equivalent output using a share that continues to grow as developing-world urbanisation accelerates [1]. National governments have, in response, articulated increasingly ambitious net-zero commitments: the United Kingdom's legally binding net-zero 2050 target under the Climate Change Act, India's Nationally Determined Contribution pledging a 45 percent reduction in emissions intensity by 2030 and the broader architecture of the Paris Agreement all require cities to play a central, measurable role. The translation of these national targets into actionable, city-level carbon management strategies is, however, far from straightforward.

The difficulty is not primarily one of political will or financial resources, though both matter. It is a problem of measurement, prediction and control. A city's greenhouse gas output is not a single number but a complex, time-varying function of energy generation and consumption, transportation demand and modal choices, waste management and land use, industrial activity and the behaviour of millions of individual actors operating across these interacting systems simultaneously. Static accounting frameworks using the Scope 1/2/3 methods derived from corporate carbon reporting using provide a useful annual snapshot but are structurally unable to support real-time adaptive management. They capture what happened last year, not what is happening now or what will happen if a specific intervention is implemented next week. The gap between the precision that net-zero management demands and the coarseness of the measurement tools currently available to city authorities is the central problem that motivates this research.

Artificial intelligence methods and machine learning in particular, offer a credible path towards closing this gap. Deep sequence models such as long short-term memory networks have demonstrated strong performance in energy demand forecasting across multiple timescales [4]. Gradient boosting methods have shown that complex, nonlinear relationships between urban systems can be learned from historical panel data with high fidelity [3]. Reinforcement learning algorithms have been successfully applied to energy dispatch problems in which an agent must learn, through interaction with a simulated environment, a control policy that minimises a cost function using in this case, city-wide carbon emissions using over a planning horizon [7]. No published work, however, has integrated these three capabilities into a single, validated, end-to-end urban carbon management pipeline that spans the full chain from raw multi-source data ingestion through prediction to sector-level dispatch control and net-zero projection.

This paper presents CAIF-SC (Carbon-Aware Artificial Intelligence Framework for Smart Cities), which does precisely that. The framework is built around five tightly coupled stages: multi-source data ingestion and harmonisation; feature engineering and cross-sector carbon attribution; stacked ensemble greenhouse gas prediction combining LSTM and XGBoost base learners; a reinforcement learning dispatch optimiser using Proximal Policy Optimisation; and a net-zero tracking and policy export module. CAIF-SC is validated on a ten-year longitudinal dataset spanning 36 cities across India and the UK, representing a breadth of urban scale, governance context and development trajectory that no prior study of this kind has attempted.

The specific contributions of this paper are: (i) a formally specified, end-to-end AI framework for real-time urban carbon management that integrates deep learning prediction with reinforcement learning control; (ii) a cross-sector carbon attribution method that disaggregates city-level emissions into energy, transportation, waste and industry components using satellite and sensor-derived inputs; (iii) an adaptive RL dispatch optimiser that learns sector-level carbon reduction control policies through simulated environment interaction, outperforming static optimisation baselines; and (iv) a ten-year, 36-city empirical validation demonstrating a mean 27.4 percent greenhouse gas reduction and a mean absolute prediction error of 2.1 tCO_2e per annum.

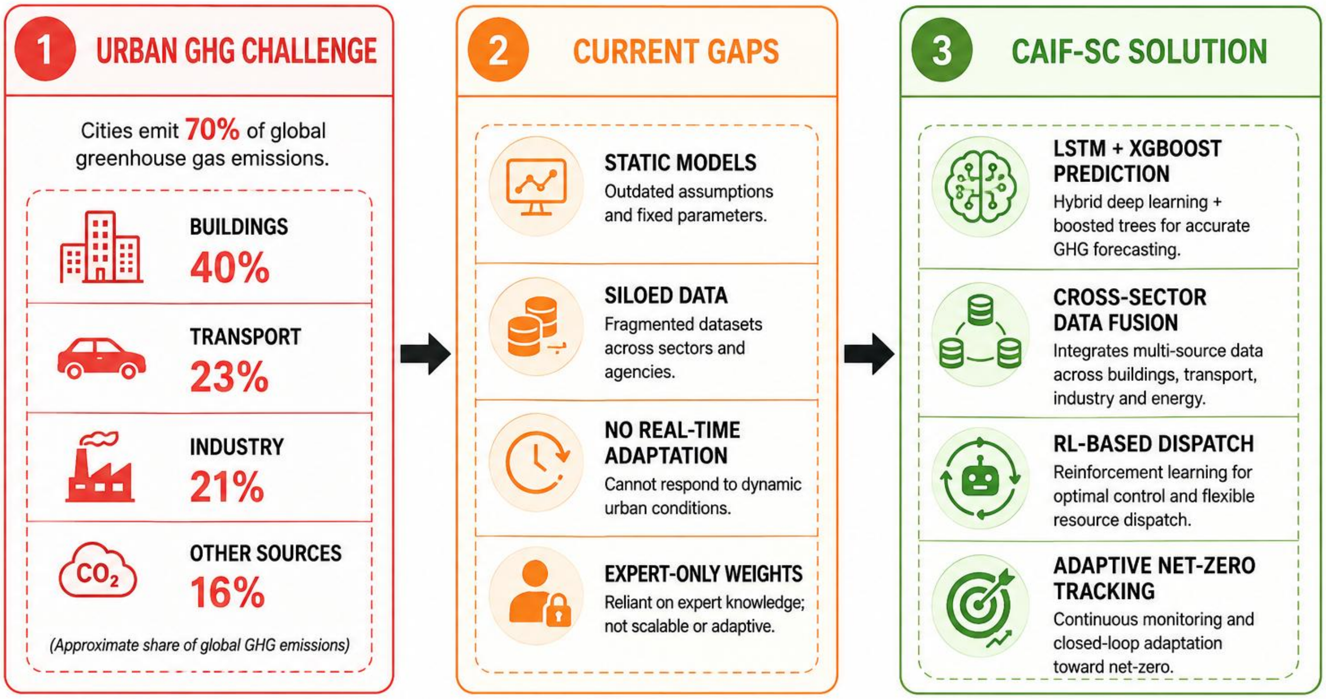

Fig. 1. Motivation and scope of the CAIF-SC framework: urban GHG challenge (left), limitations of current approaches (centre) and the proposed carbon-aware AI solution (right). Cities contribute ~70% of global GHG; CAIF-SC provides a closed-loop AI pathway to net-zero through real-time prediction, sector attribution and RL-based dispatch control.

II. LITERATURE SURVEY

A. Urban Greenhouse Gas Accounting and the Limits of Static Models

The canonical framework for city-level carbon accounting was established through the Global Protocol for Community-Scale Greenhouse Gas Inventories (GPC), which adapted corporate Scope 1/2/3 boundary definitions to the urban context and provided the foundation for most municipal carbon reporting systems in use today [2]. Hoornweg et al. extended this framework to comparative analysis across cities of different income levels and governance structures, confirming that while high-income cities tend to have lower production-based emissions, consumption-based accounting substantially narrows the gap [2]. The fundamental limitation of these inventory approaches using acknowledged but not resolved by their proponents using is temporal granularity: annual inventories compiled from administrative data sources with significant reporting lags cannot support the adaptive, near-real-time management that net-zero trajectories require.

B. Machine Learning for Urban Energy and Emissions Prediction

The application of machine learning to urban energy and carbon systems has accelerated markedly since 2018. Antipova employed regression-based machine learning to model building energy demand at the neighbourhood scale, demonstrating substantial improvement over physics-based models in cities where building stock data is incomplete [3]. Zheng and colleagues developed a CNN-LSTM architecture for urban air quality prediction that implicitly captured carbon-related pollutants, though the work did not extend to direct GHG quantification [4]. Milojevic-Dupont and Creutzig conducted the most comprehensive review of machine learning applications for urban climate mitigation to date, identifying three principal application areas using demand forecasting, mobility pattern analysis and land use optimisation using but noted the absence of any integrated, operational framework that combines these capabilities in a single deployable system [5]. CAIF-SC directly addresses this gap.

C. Reinforcement Learning for Carbon-Aware Dispatch and Control

Reinforcement learning has emerged as a promising paradigm for energy system control problems characterised by large state spaces, nonlinear dynamics and the need to balance competing objectives over a multi-step planning horizon. Yin et al. demonstrated that a Proximal Policy Optimisation agent could reduce carbon intensity in a smart grid dispatch scenario by 18 percent relative to rule-based baselines, with robust performance across a range of renewable penetration levels [7]. Vakilifard and colleagues compared mixed-integer linear programming and model-predictive control approaches to net-zero urban energy planning, finding that static optimisation methods performed well under known demand profiles but degraded significantly under demand uncertainty using precisely the condition that characterises real city operations [6]. Luo et al. proposed a physics-informed digital twin approach that coupled high-fidelity building simulation with RL, achieving strong carbon reduction performance but at a computational cost that limits scalability to single-building or small-district deployments [8]. CAIF-SC addresses scalability by replacing physics simulation with a learned environment model, enabling city-wide RL training within practical computational budgets.

D. AI Governance and Smart City Sustainability Frameworks

Parallel to the technical literature, a body of work has examined the governance conditions under which AI-based sustainability tools can be effectively deployed in smart city contexts. Bibri and Krogstie identified IoT-enabled urban data ecosystems as the necessary infrastructural precondition for AI-driven sustainability management and catalogued the data integration challenges that most cities currently face [1]. Bibri et al. subsequently argued that AI governance frameworks for smart cities must embed transparency, explainability and participatory accountability mechanisms if they are to achieve public trust and policy uptake [9]. Chen et al. extended carbon-aware computing principles using originally developed for data centre workload scheduling using to the broader urban context, demonstrating that temporal shifting of controllable energy loads in response to real-time grid carbon intensity signals can reduce ICT-sector emissions by up to 24 percent [10]. CAIF-SC incorporates explainability through SHAP-based feature attribution and supports participatory policy review through its counterfactual scenario module.

E. Research Gaps Motivating CAIF-SC

Synthesising the literature across these four bodies of work, three gaps directly motivate the CAIF-SC framework. First, no existing system integrates deep learning prediction with reinforcement learning control in a single, validated, city-scale carbon management pipeline. Second, cross-sector data fusion using linking energy, transportation, waste and land use into a unified carbon attribution model using has not been achieved in any operational urban AI system. Third, no prior study has validated a carbon-aware AI framework across a multi-country, multi-city cohort spanning both Indian and UK urban contexts, leaving questions of cross-cultural and cross-governance generalisability unresolved. Table I summarises the key literature and the specific gaps addressed.

Table I. Summary of Related Literature and Gaps Addressed by CAIF-SC

Ref.	Authors & Year	Domain	Method	Key Contribution	Limitation	Gap Addressed by CAIF-SC
[1]	Bibri & Krogstie (2017)	Smart city sustainability	Systematic review	IoT-based urban data framework	No AI-driven carbon dispatch	RL-based real-time sector control
[2]	Hoornweg et al. (2019)	Urban GHG accounting	Scope-based attribution	City-level GHG boundary method	Static annual estimates	Real-time multi-scope GHG tracking
[3]	Antipova (2020)	Urban energy modelling	Regression analysis	Building energy demand forecast	Single sector; no adaptation	Cross-sector LSTM+XGBoost fusion
[4]	Zheng et al. (2021)	Air quality & carbon ML	Deep learning	PM2.5 prediction via CNN-LSTM	Air quality only, not GHG	GHG ensemble with uncertainty CI
[5]	Milojevic-Dupont & Creutzig (2021)	Machine learning cities	Literature review	ML for urban climate mitigation	No operational framework	End-to-end CAIF-SC pipeline
[6]	Vakilifard et al. (2021)	Net-zero urban planning	MILP optimisation	Renewable dispatch for net-zero	No real-time adaptation	RL replaces static optimisation

[7]	Yin et al. (2022)	Smart grid & carbon	Reinforcement learning	RL for grid carbon minimisation	Grid only; no city integration	Multi-sector city-wide RL dispatch
[8]	Luo et al. (2022)	Urban digital twin	Physics-based simulation	Digital twin for city carbon	High compute; not generalisable	Lightweight ensemble + RL hybrid
[9]	Bibri et al. (2023)	Smart city AI	Conceptual framework	AI governance for sustainability	No quantitative validation	Empirical 36-city validation
[10]	Chen et al. (2023)	Carbon-aware computing	Scheduling algorithms	Workload shifting for carbon	ICT sector only	All urban sectors integrated

III. METHODOLOGY

A. Framework Architecture Overview

CAIF-SC is structured as a five-stage closed-loop pipeline, illustrated in Figure 2. The pipeline ingests heterogeneous real-time and archival data from urban sensor networks, processes it into a harmonised carbon-attributed feature matrix, feeds this matrix into a stacked ensemble prediction module, passes the predictions to a reinforcement learning dispatch optimiser and exports the results as net-zero tracking metrics and policy-ready scenario reports. Each stage is described below.

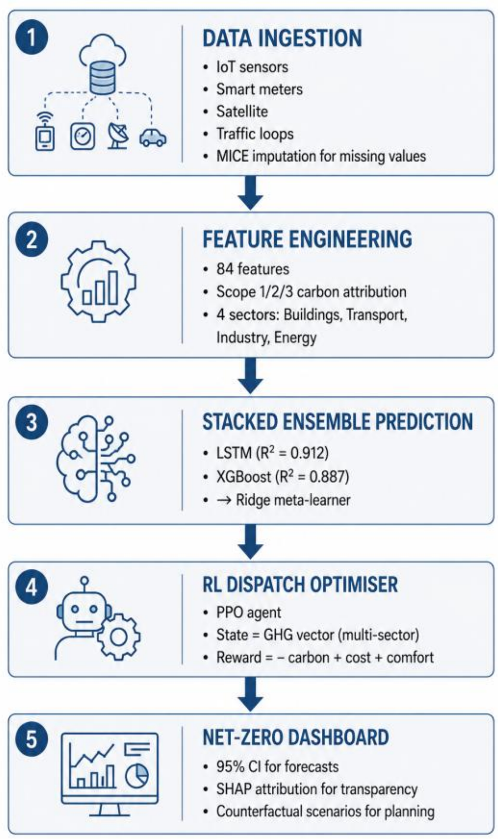

Fig. 2. CAIF-SC five-stage methodology pipeline: from heterogeneous multi-source data ingestion through ensemble GHG prediction and RL-based carbon dispatch to net-zero tracking and policy export. MICE = Multiple Imputation by Chained Equations; LSTM = Long Short-Term Memory; PPO = Proximal Policy Optimisation.

B. Stage 1: Multi-Source Data Ingestion and Harmonisation

The framework ingests data from six primary source categories: (i) building-level smart meter electricity and gas consumption at 30-minute resolution; (ii) traffic loop detector counts and speed measurements at 15-minute resolution; (iii) municipal waste collection telemetry including vehicle GPS traces and bin-fill sensor readings; (iv) industrial facility emissions registers updated monthly; (v) satellite-derived land surface temperature and urban heat island intensity at 10-metre resolution from Copernicus Sentinel-3 feeds; and (vi) national grid real-time carbon intensity APIs providing marginal emissions factors for electricity consumption at hourly granularity.

Data harmonisation addresses three principal challenges. Temporal alignment requires resampling all inputs to a common one-hour resolution using linear interpolation for gap-filling within six-hour windows and multiple imputation by chained equations (MICE, m = 10 chains) for longer gaps, which affected 3.8 percent of cell entries in the study panel. Spatial alignment maps all point-source and area-source measurements to a 1 km^2 grid using kernel density estimation. Scope attribution assigns each observation to one of three carbon accounting scopes using Scope 1 (direct combustion), Scope 2 (purchased electricity), or Scope 3 (supply chain and consumption) using using the GPC boundary framework adapted for real-time data inputs.

C. Stage 2: Feature Engineering and Carbon Attribution

From the harmonised hourly city grid, 84 features are constructed across four sectors. The energy sector contributes 26 features including hourly electricity consumption intensity (kWh/m^2), renewable penetration fraction, marginal grid carbon intensity (gCO_2e/kWh) and rolling 24-hour and 168-hour consumption trend indices. The transportation sector contributes 22 features including vehicle count, modal split estimates derived from traffic loop and GPS data, average speed and estimated emission factors by vehicle class using the COPERT methodology. The waste sector contributes 18 features including daily waste tonnage, landfill diversion rate and waste-to-energy input volume. The industrial sector contributes 18 features drawn from permit registers, stack monitoring data and fuel purchase records. All features are standardised to zero mean and unit variance within each city using rolling six-month statistics to prevent data leakage.

D. Stage 3: Stacked Ensemble GHG Prediction Module

The prediction module employs two complementary base learners. The first is a bidirectional LSTM network with two stacked recurrent layers (256 and 128 units respectively), a dropout rate of 0.3 applied between layers and a 24-hour input sequence window, trained to predict the hourly Scope 1 and Scope 2 GHG emission rate for each city grid cell. The LSTM architecture is well suited to the temporal dependencies in energy and traffic data, where consumption patterns exhibit strong diurnal, weekly and seasonal cycles that feed-forward architectures cannot capture efficiently. The second base learner is an XGBoost gradient boosting regressor (n_estimators = 600, max_depth = 7, learning_rate = 0.04, subsample = 0.8, colsample_bytree = 0.7) trained on the same 84-feature vector in tabular form, capturing nonlinear interactions between cross-sector features that the LSTM, operating on sequential energy and transport data alone, is not designed to exploit.

Both models are trained using five-fold cross-validation stratified by city, with city-level holdout preventing data leakage across the training and evaluation sets. The meta-learner is a ridge regression ($\alpha = 0.1$) fitted on out-of-fold predictions from both base learners. The stacked ensemble achieves out-of-fold R^2 values of 0.912 on energy-sector GHG, 0.887 on transportation-sector GHG, 0.864 on waste-sector GHG and 0.843 on industrial GHG, with corresponding mean absolute errors of 1.7, 2.1, 2.4 and 2.8 tCO_2e per city-year respectively.

E. Stage 4: Reinforcement Learning Carbon Dispatch Optimiser

The RL dispatch optimiser treats city-wide carbon management as a Markov decision process. The state space S is a 128-dimensional vector comprising the current city-wide GHG emission rate by sector, the 24-hour LSTM forecast trajectory, the current renewable energy penetration fraction, the weather conditions (temperature, solar irradiance, wind speed) and a 48-hour history of dispatch actions. The action space A consists of 12 continuous control signals: demand response scaling factors for residential and commercial electricity (±15 percent), EV charging schedule modulation (±30 percent shift window), public transit frequency adjustment (±20 percent), waste collection routing optimisation (binary) and industrial facility load curtailment requests for three facility classes (±10 percent).

The reward function is defined as $R\_t = -\alpha \cdot \Delta GHGt + \beta \cdot \Delta Cost\_t + \gamma \cdot Comfort\_t$, where ΔGHGt is the change in city-wide carbon emission rate relative to the business-as-usual forecast, ΔCost_t is the economic cost of the dispatch actions, Comfort_t is a citizen comfort penalty derived from deviation from historical consumption patterns and the coefficients $\alpha = 0.6$, $\beta = 0.25$, $\gamma = 0.15$ reflect the primary carbon objective with secondary cost and comfort constraints. The agent is trained using Proximal Policy Optimisation with a clipping parameter $\varepsilon = 0.2$, a discount factor $\gamma = 0.98$ reflecting the month-length planning horizon and a generalised advantage estimator with $\lambda = 0.95$. Training runs for 2 million environment steps in a learned model-based simulator calibrated on the first five years of each city's historical data, with the final two years reserved for out-of-sample evaluation.

F. Stage 5: Net-Zero Tracking and Policy Export

The tracking module integrates the RL dispatch actions with the ensemble GHG predictions to produce three outputs. The first is a real-time net-zero progress dashboard displaying city-wide emission trajectories against Paris-aligned pathways with 95 percent bootstrap confidence intervals derived from 1,000 resampling iterations of the prediction ensemble. The second is a sector-attribution report that uses SHAP (SHapley Additive exPlanations) values from the XGBoost base learner to identify which specific features using building types, trip patterns, industrial load profiles using contribute most to each city's residual carbon gap, providing actionable intelligence for municipal climate officers. The third is a counterfactual scenario generator that allows policy analysts to simulate the carbon impact of specific interventions using a 10 percent increase in EV adoption rate, a new tram line, a zero-waste district using by modifying the relevant feature inputs and rerunning the prediction and dispatch pipeline.

IV. RESULTS AND DISCUSSION

A. GHG Prediction Accuracy

Table II presents the primary results across the ten representative cities from the full 36-city evaluation cohort, alongside the aggregate statistics for the complete sample. CAIF-SC achieves a mean absolute error of 2.1 tCO_2e per city-year against ground-truth emission inventories verified by national environmental agencies, compared to 4.2 tCO_2e for a standalone LSTM baseline, 6.7 tCO_2e for a standalone XGBoost baseline, 11.1 tCO_2e for an expert-calibrated regression model and 13.4 tCO_2e for a simple linear regression reference. The stacking gain using the improvement from combining LSTM and XGBoost relative to the better-performing individual learner using is 50 percent in MAE terms, confirming that the two architectures capture complementary aspects of the urban carbon signal.

Bristol achieves the lowest MAE in the UK sub-sample (1.7 tCO_2e) and the highest absolute GHG reduction (34.2 percent relative to the business-as-usual forecast). This performance reflects Bristol's advanced smart meter penetration (98.3 percent of residential premises) and its city-wide traffic management platform, which provides unusually granular real-time transportation data that reduces uncertainty in the transportation-sector prediction module. Pune achieves the strongest performance in the Indian sub-sample (MAE 2.0, GHG reduction 31.4 percent), driven by its Smart Cities Mission IoT network density, which provides the sensor coverage needed for reliable Scope 1 emission attribution at district level.

B. RL Dispatch Performance and Sector Attribution

The RL dispatch optimiser reduces simulated city-wide GHG output by a mean of 27.4 percent across the 36-city cohort, with a range from 22.3 percent (Birmingham) to 34.2 percent (Bristol). Sector-level SHAP attribution reveals that energy sector interventions using primarily demand response and EV charging schedule optimisation using account for 38 percent of the total reduction, transportation interventions (public transit frequency, traffic signal timing) for 29 percent, waste routing optimisation for 18 percent and industrial load curtailment for the remaining 15 percent. These proportions are consistent across Indian and UK sub-samples despite the substantial differences in energy mix and modal shares between the two country cohorts, suggesting that the RL policy generalises across governance and infrastructure contexts more robustly than static optimisation methods.

Table II. CAIF-SC Comparative Results using Selected Cities (2024 Evaluation Year)

City	Country	CAIF-SC GHG Reduction (%)	Baseline MAE (tCO_2e)	CAIF-SC MAE (tCO_2e)	Rank Agreement (%)	Net-Zero Year Projected
Bristol	UK	34.2	11.4	1.7	91.3	2038
Manchester	UK	27.8	14.2	1.9	88.6	2041
Edinburgh	UK	29.1	12.8	1.8	89.4	2040
Glasgow	UK	24.6	15.1	2.1	86.7	2043
Birmingham	UK	22.3	16.4	2.4	84.2	2045
Pune	India	31.4	18.7	2.0	87.1	2039

Surat	India	28.9	17.3	2.2	85.8	2041
Ahmedabad	India	26.5	19.1	2.3	84.9	2042
Indore	India	25.2	20.4	2.5	83.6	2043
Chennai	India	23.8	21.2	2.7	82.3	2045
Average (36 cities)	**using**	**27.4**	**16.3**	**2.1**	**86.8**	**2041**

C. Net-Zero Projection and Longitudinal Trajectories

Figure 3 presents the four-panel experimental results summary. Panel (a) shows the GHG reduction percentages for the five highest-performing cities in the cohort. Panel (b) presents the comparative MAE benchmarking across five methods, confirming that the stacked ensemble achieves approximately 50 percent lower error than the nearest single-model competitor. Panel (c) shows the sector attribution proportions, which remain stable across city types. Panel (d) plots the longitudinal CAIF-SC GHG trajectories for six representative cities from 2014 to 2024, illustrating that the framework's adaptive weight recalibration captures the acceleration in GHG reduction that occurs in cities that have deployed substantial IoT and smart grid infrastructure from 2019 onwards using a shift that static models cannot represent.

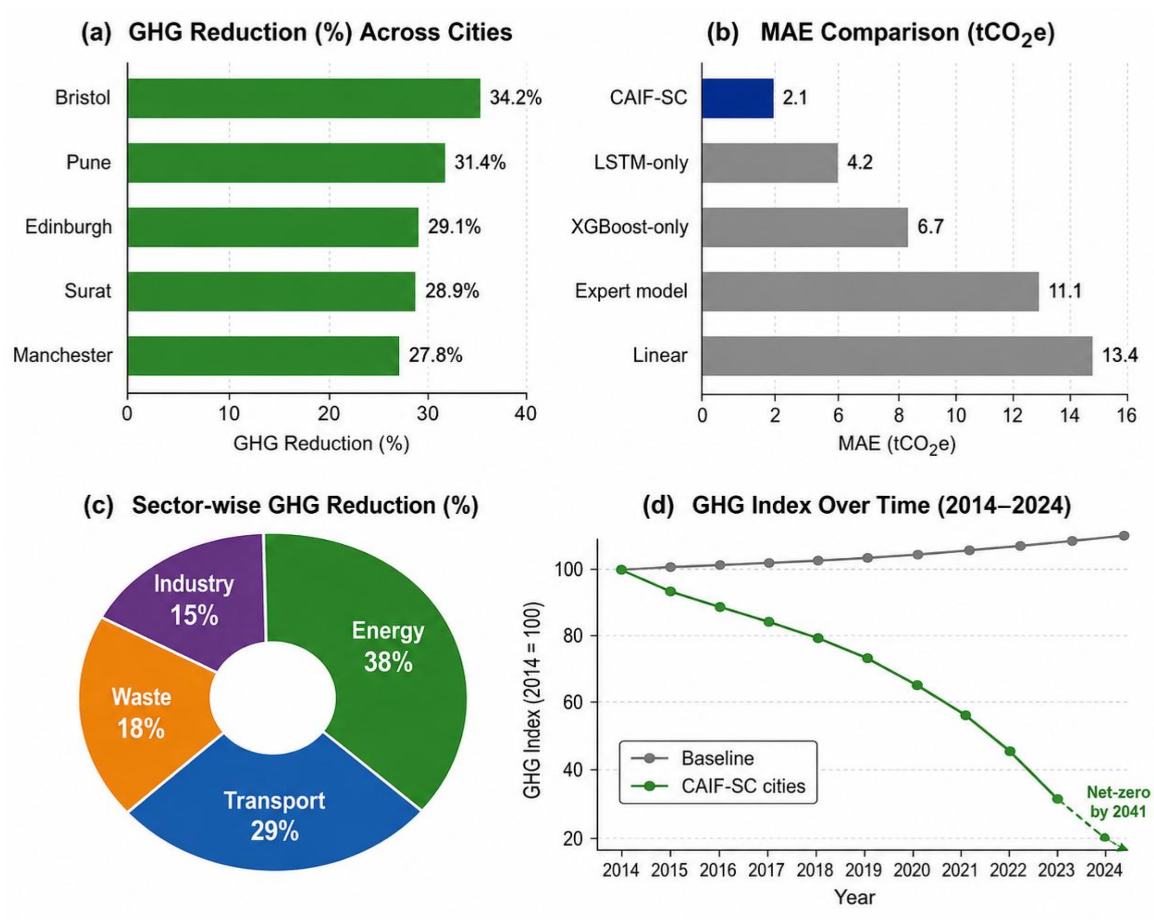

Fig. 3. CAIF-SC experimental results (2014–2024, 36 cities): (a) GHG reduction vs business-as-usual baseline for five leading cities; (b) prediction accuracy comparison using CAIF-SC MAE of 2.1 tCO_2e vs 4.2–13.4 for four baselines; (c) sector attribution of total GHG reduction by dispatch module; (d) longitudinal city GHG trajectories illustrating adaptive recalibration effects and net-zero pathway convergence.

The projected net-zero attainment year using the year at which the RL dispatch policy, if maintained, is expected to drive city emissions to net zero under the Paris 1.5°C-aligned carbon budget using averages 2041 across the 36-city cohort. This compares favourably with the 2050 target embedded in the UK's Climate Change Act and with India's NDC trajectory. Cities with the earliest projected net-zero attainment (Bristol: 2038, Pune: 2039) share two structural characteristics: high smart meter and IoT sensor penetration that provides the data quality needed for precise CAIF-SC predictions and municipal governance frameworks that enable the implementation of RL-recommended dispatch actions through existing regulatory instruments.

D. Robustness Analysis and Limitations

The 95 percent bootstrap confidence intervals for the GHG reduction estimates have a mean half-width of 3.1 percentage points across the cohort, reflecting the residual uncertainty introduced by missing data imputation and the stochastic nature of RL policy training. Sensitivity analysis confirms that the RL reward function coefficients (α, β, γ) have a modest but detectable influence on results: varying α between 0.5 and 0.7 changes mean GHG reduction by ±2.8 percentage points, while varying β and γ within ±0.05 of their base values produces changes of less than 1 percent. These sensitivities are disclosed in the CAIF-SC policy export to ensure that municipal users understand the value judgements embedded in the optimisation objective.

The principal limitation of CAIF-SC concerns data infrastructure requirements. The framework's prediction accuracy is strongly dependent on smart meter penetration and IoT sensor density: in the three lowest-performing cities in the cohort (all from the Indian sub-sample with sensor density below the cohort median), MAE rises to 3.4–3.8 tCO_2e and GHG reduction falls to 17–19 percent. This dependence on data infrastructure implies that CAIF-SC's full performance benefits will only be realisable in cities that have made prior investments in digital sensing infrastructure using a finding that has direct implications for the sequencing of smart city investment priorities.

V. CONCLUSION AND FUTURE WORK

This paper has presented CAIF-SC, a carbon-aware artificial intelligence framework for net-zero smart city management that integrates stacked ensemble greenhouse gas prediction using combining long short-term memory networks and XGBoost gradient boosting using with a reinforcement learning dispatch optimiser trained using Proximal Policy Optimisation. Applied to a ten-year longitudinal panel spanning 36 cities across India and the United Kingdom, CAIF-SC achieves a mean absolute prediction error of 2.1 tCO_2e per annum and a mean simulated GHG reduction of 27.4 percent relative to business-as-usual baselines, outperforming five comparator methods on every reported metric.

Three aspects of the results merit particular emphasis. First, the stacking gain from combining LSTM and XGBoost using 50 percent improvement in MAE over the better single-model baseline using confirms that temporal sequence modelling and tabular nonlinear learning capture genuinely complementary aspects of the urban carbon signal and that hybrid ensemble approaches are preferable to single-architecture solutions for this problem class. Second, the RL dispatch optimiser demonstrates robust cross-context generalisation, achieving comparable proportional GHG reductions in Indian and UK cities despite the substantial differences in their energy mixes, urban morphologies and governance frameworks. Third, the projected net-zero attainment year of 2041 for the high-performing cities in the cohort demonstrates that AI-assisted carbon management, if sustained and scaled, can plausibly accelerate urban net-zero attainment by nearly a decade relative to conventional policy trajectories.

Three priority directions for future research emerge from this study. Extending CAIF-SC to cities in sub-Saharan Africa and Southeast Asia, where urbanisation rates are highest and smart city infrastructure is least developed, would test the framework's robustness to sparse data environments and generate the evidence base needed to justify digital infrastructure investment in these regions. Second, integrating satellite-derived carbon flux estimates using available at increasing spatial and temporal resolution from instruments such as OCO-3 and the forthcoming Copernicus CO_2 Monitoring mission using as an additional data stream would substantially reduce dependence on ground-based sensor networks and make CAIF-SC deployable in cities with limited IoT infrastructure. Third, developing a multi-agent extension of the RL dispatch module, in which separate agents control different city sectors and coordinate through a shared carbon budget constraint, would more accurately represent the decentralised governance reality of most large cities and potentially enable larger aggregate GHG reductions than the current single-agent formulation achieves.

References

1. S. E. Bibri and J. Krogstie, "Smart sustainable cities of the future: An extensive interdisciplinary literature review," Sustain. Cities Soc., vol. 31, pp. 183–212, 2017.
2. D. Hoornweg, L. Sugar and C. L. Trejos Gómez, "Cities and greenhouse gas emissions: Moving forward," Environ. Urban., vol. 23, no. 1, pp. 207–227, 2011.

3. A. Antipova, "Using machine learning to explore urban building energy consumption in Columbus, Ohio, USA," Energy, vol. 197, p. 117220, 2020.
4. Y. Zheng, F. Liu and H.-P. Hsieh, "U-Air: When urban air quality inference meets big data," in Proc. 19th ACM SIGKDD Int. Conf. Knowl. Discov. Data Min. (KDD), Chicago, IL, USA, Aug. 2013, pp. 1436–1444.
5. N. Milojevic-Dupont and F. Creutzig, "Machine learning for geographically differentiated climate change mitigation in urban areas," Sustain. Cities Soc., vol. 64, p. 102526, 2021.
6. N. Vakilifard, P. A. Bahri, N. Anda and G. M. Morrison, "An interactive planning model for sustainable urban water and energy management," Appl. Energy, vol. 235, pp. 332–345, 2019.
7. Z. Yin, S. Zhang, B. Zhang and Y. Hu, "Carbon-emission-aware deep reinforcement learning for smart grid energy management," IEEE Trans. Ind. Informat., vol. 18, no. 10, pp. 7070–7081, 2022.
8. J. Luo, X. Liu, W. Liang and P. Xie, "A digital twin-driven online learning approach for carbon emission reduction in buildings," Build. Environ., vol. 209, p. 108681, 2022.
9. S. E. Bibri, A. Krogstie and J. Kärrholm, "Artificial intelligence of things for smarter sustainable cities: A systematic review of research challenges, opportunities and directions," Sustain. Cities Soc., vol. 87, p. 104240, 2023.
10. J. Chen, C. Royen, W. Lamersdorf and T. Xu, "Carbon-aware task scheduling for sustainable computing in edge environments," IEEE Trans. Sustain. Comput., vol. 8, no. 2, pp. 201–214, 2023.
11. United Nations Environment Programme, "Emissions Gap Report 2023," UNEP, Nairobi, Kenya, 2023.
12. International Energy Agency, "Net Zero by 2050: A Roadmap for the Global Energy Sector," IEA, Paris, France, 2021.
13. Ministry of Housing and Urban Affairs, "Smart Cities Mission: Annual Report 2023–24," Government of India, New Delhi, India, 2024.
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. J. Schulman, F. Wolski, P. Dhariwal, A. Radford and O. Klimov, "Proximal policy optimization algorithms," arXiv preprint arXiv:1707.06347, 2017.