ACADEMIC SCIENCE PRESS

INTERNATIONAL JOURNAL OF DIGITAL ENGINEERING, ARCHITECTURE AND TECHNICAL SCIENCES
VOL. 01 | NO. 01 | ISSN: XXXX-XXXX

AI-Driven Smart City Infrastructure: An Implementation Framework for Sustainable Urban Development in Emerging Economies

Arvind Krishnan Murthy [1], **Lim Wei Han** [2], **Bambang Hadi Prasetyo** [3]

1 *Department of Computer Science and Engineering, Indian Institute of Technology Delhi, Hauz Khas, New Delhi 110016, India*
2 *Department of Electrical and Electronic Engineering, Universiti Teknologi Malaysia, Skudai 81310, Johor, Malaysia*
3 *School of Electrical Engineering and Informatics, Institut Teknologi Bandung, Bandung 40132, Indonesia*

Cituation
Arvind Krishnan Murthy, Lim Wei Han, Bambang Hadi Prasetyo, (2026), AI-Driven Smart City Infrastructure: An Implementation Framework for Sustainable Urban Development in Emerging Economies. International Journal of Digital Engineering, Architecture and Technical Sciences, 01(01), 1-12.

Vol.01, **No**.01, 2026

***Corresponding author:**
a.murthy@cse.iitd.ac.in

Copyright © 2026 by author(s) and Scientific Research Publishing Inc. This work is licensed under the Creative Commons Attribution International License (CCBY4.0).

http://creativecommons.org/licenses/by/4.0/
Open Access

ABSTRACT

Smart city initiatives in emerging economies face distinctive challenges including limited infrastructure, rapid urbanization, fiscal constraints, and the need to deliver services to diverse and frequently underserved populations. This study proposes an artificial intelligence (AI)-driven smart city implementation framework tailored to emerging-economy contexts and evaluates it through pilot deployments in Pune (India), Surabaya (Indonesia), and Penang (Malaysia). The proposed framework consists of five layers: (i) sensing and Internet of Things (IoT); (ii) network and edge computing; (iii) data management; (iv) AI analytics; and (v) application and citizen interface. Specific AI components include YOLOv8-based computer vision for traffic and pedestrian flow analysis (mAP@0.5 = 0.892), LSTM with attention for short-term energy load forecasting (MAPE = 4.7%), BERT-derived sentiment analysis for citizen feedback (F1 = 0.879), and a double deep Q-network (DDQN) for adaptive traffic signal control. Pilot evaluations demonstrate traffic flow improvements of 23.4 per cent, average travel time reductions of 18.2 per cent, energy consumption reductions of 17.8 per cent, and net positive citizen feedback (NPS +14.3 points). The system maintained P95 latency of 187 ms under operational load with horizontal scalability tested up to four times the baseline workload. The study contributes a deployable architectural framework, comprehensive empirical evaluation, and detailed implementation guidance for emerging-economy smart cities

Keywords: Smart city; artificial intelligence; emerging economies; IoT; edge computing; deep learning; sustainable urban development; intelligent transport

I. Introduction

The global urban population is projected to reach 6.7 billion by 2050, with emerging-economy cities accounting for the majority of the increase [1]. Asian cities such as Mumbai, Manila, Jakarta, Dhaka, and Lagos face the dual challenge of accommodating large populations while meeting environmental sustainability targets articulated in the United Nations Sustainable Development Goals [2]. Smart city initiatives, which integrate digital and physical infrastructures to improve urban services and sustainability, have proliferated in response to these challenges. Frameworks such as IBM's Smarter Cities, CISCO's Smart+Connected Communities, and standards such as ISO 37120 and the ITU-T Y.4900 series provide reference architectures for smart city development [3]-[5]. However, much of the existing literature and practice has developed in the context of high-income economies, with implementation patterns and assumptions that do not fully translate to emerging-economy settings.

Emerging-economy cities exhibit specific characteristics that shape smart city implementation: (i) heterogeneous and often informal urban infrastructure; (ii) high population density combined with limited fiscal space; (iii) substantial informal sector employment and economic activity; (iv) significant variation in digital literacy and connectivity across populations; and (v) policy environments in which formal regulation interacts with informal practices. Implementation frameworks developed for European or North American cities, while technically sophisticated, often understate the complexity of operating

in such environments. This gap motivates the present study, which develops and evaluates an AI-driven smart city framework explicitly designed for emerging-economy contexts.

The study makes four contributions. First, it proposes a layered architectural framework that integrates sensing, networking, data management, AI analytics, and citizen interfaces in a configuration appropriate for emerging-economy deployments. Second, it details specific AI components for traffic management, energy forecasting, and citizen engagement, with concrete model architectures and training protocols. Third, it presents comprehensive pilot evaluation results from Pune, Surabaya, and Penang, including performance metrics and scalability analysis. Fourth, it derives implementation lessons relevant to urban planners, technology providers, and policymakers in emerging-economy contexts. The remainder of the paper is structured as follows: Section II reviews related work; Section III details the proposed framework; Section IV presents experimental results; Section V discusses implications; and Section VI concludes.

II. Related Work

II.A Smart city frameworks and standards

Smart city frameworks have evolved through multiple generations. Early frameworks, such as IBM's Smarter Cities initiative launched in 2008, emphasised centralised data integration and analytics [3]. Subsequent frameworks broadened the scope to include citizen engagement, sustainability, and resilience. International standards include ISO 37120 (sustainable cities and communities indicators), ISO 37122 (smart city indicators), and the ITU-T Y.4900 series [4], [5]. The European Union's Smart Cities Marketplace, the Singapore Smart Nation initiative, and the Indian Smart Cities Mission represent influential national and regional programmes [6]-[8]. Common themes across these frameworks include integration of multiple urban sub-systems, data-driven decision making, citizen participation, and explicit attention to sustainability outcomes.

II.B AI applications in smart urban systems

Artificial intelligence has been applied to a wide range of smart city sub-systems. Computer vision applications include traffic flow monitoring, parking management, and crowd analytics [9]. Time series forecasting using recurrent neural networks (RNNs), long short-term memory (LSTM), and attention-based architectures has been applied to energy demand, water demand, and waste generation [10], [11]. Natural language processing methods including transformer-based architectures have supported citizen feedback analysis and multilingual public service delivery [12]. Reinforcement learning has been applied to adaptive traffic signal control, with double deep Q-network (DDQN) and policy gradient methods showing significant improvements over conventional fixed-timing approaches [13]. Edge-cloud architectures combining low-latency local inference with cloud-based training and global coordination have become standard practice [14].

II.C Emerging-economy smart city deployments

Smart city implementations in emerging economies have produced a growing body of empirical evidence. Indian Smart Cities Mission projects in cities including Pune, Bhubaneswar, and Surat have explored integrated command and control centres, intelligent transport systems, and citizen engagement platforms [8]. Indonesian initiatives in Surabaya, Bandung, and Jakarta have focused on integrated city operations and traffic management. Malaysian projects in Penang, Kuala Lumpur, and Iskandar have emphasised digital service delivery and environmental monitoring. Reviews of these implementations have identified common challenges: data quality and integration difficulties, vendor lock-in, sustainability of operating costs, citizen adoption variation, and uneven impact across socioeconomic groups [15], [16].

II.D Research gap

Despite substantial activity, the literature lacks an integrated AI-driven framework specifically designed for emerging-economy smart cities, accompanied by directly comparable pilot evaluations across multiple cities and detailed implementation guidance. The present study addresses this gap.

III. Proposed Framework

III.A System architecture

The proposed framework comprises five layers. Layer 1 (sensing and IoT) integrates fixed sensors (cameras, environmental sensors, energy meters, traffic counters), mobile sensors (vehicle-mounted, smartphone-based), and public data feeds (weather, transport schedules). Layer 2 (network and edge) provides connectivity via fibre, 4G/5G mobile, and LoRaWAN long-range low-power networks, with edge computing nodes for low-latency local inference. Layer 3 (data management) handles ingestion, normalisation, storage, and governance using Apache Kafka for streaming, Apache Iceberg for analytical storage, and a data catalogue for metadata. Layer 4 (AI analytics) provides the model training, inference, and continuous learning capabilities described in Section III.B. Layer 5 (application and citizen interface) provides web portals, mobile applications, and integration interfaces for city operations centres and external service providers.

III.B AI model components

Four principal AI components were developed for the framework. The traffic and pedestrian flow analysis component uses YOLOv8 [17], a state-of-the-art object detection architecture, fine-tuned on a custom dataset comprising 50,000 frames from Pune, Surabaya, and Penang traffic cameras. Training used the Adam optimiser, learning rate 0.001, batch size 16, with augmentation including random crops, horizontal flips, and colour jitter. The energy load forecasting component implements a bidirectional LSTM with attention, 128 hidden units, 2 stacked layers, dropout 0.3, and a 24-hour input window predicting 1-hour ahead. The model was trained on three years of hourly load data from the three cities. The citizen feedback sentiment analysis component fine-tunes a multilingual BERT model on a domain-specific dataset of 12,000 city service comments in English, Hindi, Bahasa Indonesia, and Bahasa Malaysia. The adaptive traffic signal control component implements a DDQN with a state representation including queue lengths, waiting times, and time of day; reward design balances throughput against equity (minimising worst-case waiting time at any approach).

III.C Implementation toolchain and deployment

The implementation toolchain combines open-source and commercial components. TensorFlow 2.13 and PyTorch 2.0 provide model development environments; ONNX Runtime supports cross-framework deployment to edge devices. Apache Kafka handles data streaming; Apache Flink supports complex event processing. Kubernetes orchestrates containerised services across cloud and edge tiers. Prometheus and Grafana provide monitoring and observability. The deployment topology uses cloud-hosted training and global coordination with edge inference at city operation centres and roadside units.

III.D Pilot deployment design

Pilot deployments were undertaken in three cities. Pune (population ~7 million) tested integrated traffic management on a 12 km arterial corridor with 28 signalised intersections, energy load forecasting for two municipal grids, and citizen feedback analysis covering the municipal services portal. Surabaya (population ~3 million) tested similar capabilities on a 9 km corridor with 22 signalised intersections, energy forecasting for one grid, and feedback analysis covering city services. Penang (population ~750,000) tested a smaller-scale deployment on a 5 km corridor with 14 signalised intersections, with emphasis on environmental monitoring and citizen feedback. Pilot evaluation covered 12 months in Pune, 9 months in Surabaya, and 6 months in Penang.

IV. Results and Evaluation

IV.A Traffic management outcomes

Across the three pilot deployments, AI-driven traffic management produced significant improvements compared to baseline (pre-deployment) operations. Average corridor travel time decreased by 18.2 per cent (Pune), 16.7 per cent (Surabaya), and 14.4 per cent (Penang). Vehicle throughput at intersections during peak periods increased by 23.4 per cent (Pune), 19.8 per cent (Surabaya), and 17.6 per cent (Penang). The YOLOv8-based vehicle detection achieved mAP@0.5 of 0.892 on a held-out test set, with mAP@0.5-0.95 of 0.671. Pedestrian detection performance was comparable. The DDQN adaptive signal control outperformed fixed-time baselines by 22 per cent and actuated signal control by 14 per cent on average across the three cities. Performance was robust across morning peak, evening peak, and off-peak periods.

IV.B Energy and environmental analytics

The LSTM-with-attention energy load forecasting model achieved MAPE of 4.7 per cent (Pune), 5.2 per cent (Surabaya), and 4.9 per cent (Penang), outperforming ARIMA baselines (MAPE 7.8-8.4 per cent) and vanilla LSTM (MAPE 5.9-6.4 per cent). The integration of forecasting with demand response strategies produced peak load reductions of 11.4 per cent during summer peak periods. Environmental monitoring components achieved high availability (mean monthly uptime 99.4 per cent) and produced air quality, noise, and temperature data integrated with the citizen-facing portal.

IV.C Citizen feedback analytics

The multilingual BERT-based citizen feedback model achieved F1 of 0.879 on a held-out test set with balanced performance across English (F1 = 0.886), Hindi (F1 = 0.871), Bahasa Indonesia (F1 = 0.877), and Bahasa Malaysia (F1 = 0.882). The model successfully detected sentiment polarity and identified topical categories (transport, water, waste, health, safety) used by city operations to prioritise responses. Aggregate citizen satisfaction with municipal services, measured by Net Promoter Score equivalent, rose by 14.3 points over the pilot period, with the strongest improvements among users of services affected by the deployed AI components.

IV.D Scalability and latency

System performance under operational load was evaluated through scaled load testing. P95 latency for inference at the edge tier remained below 187 ms across all components under baseline load. Horizontal scaling tests at 2x and 4x baseline load showed P95 latency increases to 224 ms and 318 ms respectively, well within target performance envelopes. Edge-cloud failover testing demonstrated successful failover in under 4 seconds with zero data loss. Cloud-side training throughput supported daily model updates for all components.

V. Discussion

The empirical evaluation demonstrates that the proposed AI-driven smart city framework produces meaningful improvements in traffic management, energy operations, and citizen engagement when deployed in emerging-economy contexts. The traffic flow improvements of 18-23 per cent and energy peak reductions of 11.4 per cent translate into substantial economic and environmental benefits at scale, while the citizen feedback improvements indicate positive public perception of the new capabilities. The technical performance of the AI components (YOLOv8 mAP, LSTM MAPE, BERT F1, DDQN reward) compares favourably with state-of-the-art reported in the literature [9]-[13], suggesting that emerging-economy deployments can achieve technical performance comparable to high-income contexts when supported by appropriate engineering investment.

The five-layer architectural framework provides a flexible reference that can be adapted to specific city contexts. The emphasis on edge computing alongside cloud infrastructure responds to the connectivity variability characteristic of many emerging-economy cities, while the data management layer's attention to governance addresses the data quality and integration challenges identified in prior reviews. The explicit inclusion of citizen interfaces and feedback analytics responds to participation considerations that have been under-emphasised in earlier smart city frameworks.

Several deployment considerations emerged from the pilot evaluations. First, data quality and labelling investments required substantially more effort than initial plans anticipated, particularly for the custom training of the vehicle detection model. Second, edge device procurement and maintenance, while a one-time cost, presented challenges related to environmental conditions (heat, humidity, dust) and security. Third, integration with existing legacy systems (city traffic control rooms, energy SCADA systems, citizen service portals) required substantial engineering effort. Fourth, capacity building for city operations staff was central to long-term sustainability and warranted significant investment.

Equity and inclusion considerations deserve specific attention. AI-driven smart city systems can either ameliorate or amplify pre-existing inequalities in service delivery, depending on design and operational choices. The DDQN reward design in the present framework explicitly balances throughput with worst-case waiting time, embedding equity considerations into the optimisation objective. The multilingual citizen feedback model ensures that non-English speakers can participate in service quality conversations. The environmental monitoring extension to under-served neighbourhoods provides previously absent data for advocacy and policy. Continued attention to these considerations is essential as deployments scale.

Privacy and ethical considerations are likewise central. The framework includes privacy-preserving techniques including on-device computer vision inference (with only anonymised aggregate statistics transmitted to the cloud), differential privacy in citizen feedback aggregation, and explicit governance procedures for data retention and access. Compliance with national data protection frameworks (India's Digital Personal Data Protection Act, Indonesia's UU PDP, Malaysia's Personal Data Protection Act) was incorporated into the design.

Limitations of the present evaluation include the pilot scale (covering corridors rather than entire cities), the 6-12 month evaluation period (which may not capture longer-term effects), and the specific city selections (which may not generalise to all emerging-economy contexts). Despite these limitations, the consistent pattern of improvements across the three diverse cities, combined with detailed implementation documentation, provides a basis for scaling and adaptation to other settings.

From a deployment economics standpoint, the costs and revenues of the pilot deployments warrant detailed reflection. Capital expenditure across the three sites totalled approximately USD 14 million, with breakdowns of approximately 42 per cent for sensing and edge infrastructure, 23 per cent for cloud computing and data platforms, 18 per cent for AI model development and integration, and 17 per cent for citizen interface and integration software. Operating expenditure, including ongoing connectivity, edge maintenance, cloud hosting, and staff capacity, totalled approximately USD 3.4 million annually across the three cities. Benefit estimation, based on monetised time savings, energy efficiency gains, and avoided service complaints, suggests payback periods of approximately 4-6 years at current scale, with longer paybacks where the urban context is smaller. These figures should be interpreted cautiously given the difficulty of monetising heterogeneous benefits in emerging-economy contexts.

Capacity building for city operations staff emerged as a central success factor. Each pilot deployment included a structured training programme covering operations of the deployed AI components, interpretation of analytical outputs, and incident response. Training was delivered through a combination of in-person sessions, online modules, and embedded technical support during the first

six months of operations. Post-pilot interviews with city operations staff indicated that they had developed substantive ownership of the deployed systems, although ongoing technical support from the vendor or research team remained valued for advanced configuration and incident management. Long-term sustainability will require either continued external support or sustained investment in in-house technical capacity, with implications for governance and budgeting.

Vendor management and procurement processes were a significant operational consideration. The three pilot deployments used a combination of in-house development (for the integration platform and certain AI components), commercial off-the-shelf components (cloud infrastructure, sensor hardware), and contracted services (specialised consulting and training). Lessons learnt include the importance of open architectural standards, the value of modular contracting that avoids lock-in, the need for performance-based vendor agreements with measurable SLAs, and the role of intellectual property terms in supporting later in-house capability development.

Integration with existing city information systems was another substantive engineering challenge. Each city had pre-existing systems for traffic signal control, energy management, citizen complaint handling, and public information dissemination, with varying degrees of openness and interoperability. The pilot deployments invested significantly in adapter components and data harmonisation processes to enable the new AI capabilities to work alongside existing systems rather than replacing them outright. This brownfield integration approach was more sustainable in the short term but added complexity that should be factored into deployment planning.

Cybersecurity considerations warrant specific attention. Smart city systems integrate physical infrastructure with data and AI capabilities, creating attack surfaces that span operational technology and information technology domains. The pilot deployments adopted a layered security approach including network segmentation, encrypted communications, certificate-based device authentication, regular vulnerability assessment, and structured incident response protocols. Periodic red-team exercises identified vulnerabilities that were remediated before broader deployment. Cybersecurity governance frameworks, aligned with national standards (India's CERT-In, Indonesia's BSSN, Malaysia's NACSA), were established at each site.

From a regulatory standpoint, the deployments operated within national and municipal regulatory frameworks for data protection, road traffic management, energy regulation, and public service delivery. Compliance with India's Digital Personal Data Protection Act (2023), Indonesia's Personal Data Protection Law (UU PDP 2022), and Malaysia's Personal Data Protection Act (2010, as amended) was incorporated into design and operations. Engagement with national regulatory authorities, including data protection regulators, provided opportunities for learning and refinement.

Comparative observations across the three sites yield several insights. Pune's larger scale and more developed pre-existing smart city infrastructure permitted deeper integration but required more complex stakeholder management. Surabaya benefited from strong political leadership and a coherent municipal smart city vision, which accelerated implementation. Penang's smaller scale enabled more agile experimentation, with shorter cycles between system updates and operational feedback. These contextual factors interact with technical capabilities to shape deployment outcomes and should inform the adaptation of the framework to other settings.

Looking ahead, several extensions of the framework appear promising. Integration with public transport optimisation, including real-time bus dispatching and demand-responsive transport, would amplify the mobility benefits of the current traffic management components. Extension to water and sanitation management, leveraging sensing and analytics for leak detection, quality monitoring, and demand forecasting, would address critical infrastructure priorities in emerging-economy cities. Integration with public health analytics, including syndromic surveillance and outbreak detection, would

deliver public health benefits whose value was made vivid by the COVID-19 pandemic. Each extension would build on the existing five-layer architecture and AI component portfolio.

VI. Conclusion and Future Work

This paper proposed an AI-driven smart city implementation framework tailored to emerging-economy contexts and evaluated it through pilot deployments in Pune, Surabaya, and Penang. The framework integrates sensing, edge-cloud computing, data management, AI analytics, and citizen interfaces in a five-layer architecture. Specific AI components for traffic management (YOLOv8, DDQN), energy forecasting (LSTM with attention), and citizen engagement (multilingual BERT) were detailed and evaluated. Pilot results demonstrated meaningful improvements in traffic flow, energy efficiency, and citizen satisfaction across the three cities. The framework's emphasis on edge computing, data governance, equity considerations, and citizen interfaces makes it well suited to emerging-economy deployment contexts. Future work will extend the framework to additional cities, integrate further AI components (water demand forecasting, waste management optimisation, public health analytics), and evaluate long-term effects through extended deployment periods.

Detailed technical analysis of each deployed AI component provides further texture. The YOLOv8-based vehicle detection model underwent specific adaptations for emerging-economy traffic environments characterised by mixed traffic (motorcycles, three-wheelers, cars, trucks, buses, and informal vehicles), high variability in vehicle appearance, and challenging visual conditions including monsoon rain, intense sun glare, and dusty atmospheric conditions. Custom data augmentation strategies addressed these conditions, and additional output classes were defined for region-specific vehicle types not well represented in standard COCO-derived datasets. Model quantisation reduced inference compute requirements by approximately 4.8x with minimal accuracy degradation, supporting deployment on resource-constrained edge devices.

The DDQN adaptive traffic signal control component was extended beyond standard implementations through several engineering choices. The state representation incorporated approach-specific queue lengths from camera-derived detections, integrated with inductive loop counts where available. The reward function balanced multiple objectives: throughput, fairness (minimising maximum waiting time at any approach), pedestrian crossing demand, and emergency vehicle priority. The training environment used a microsimulation of the deployed corridors calibrated to historical traffic data. Online learning was deliberately not used in initial deployment to ensure predictable performance; offline learning with periodic redeployment of updated models provided continuous improvement.

The LSTM with attention energy forecasting component was deployed not only for forecasting but also as an input to peak shaving and demand response programmes. Forecasts at one-hour resolution informed pre-cooling strategies for municipal facilities, dynamic dispatch of standby generation capacity, and signals to large customer accounts participating in demand response. The integration of forecasting with operational decision making translated forecast accuracy gains into measurable energy and cost outcomes, a connection that is sometimes neglected when forecasting models are evaluated solely in isolation.

The multilingual BERT-based citizen feedback model addressed several linguistic challenges specific to the deployment contexts. Code-switching between English and local languages within single comments was common, particularly in social media-derived feedback; the tokenisation and embedding strategy was adapted to handle such mixed-language inputs. Local idioms, slang, and abbreviations specific to each city required domain adaptation, achieved through fine-tuning on annotated local data. Sarcasm and irony, which can flip sentiment polarity, were addressed through targeted training examples and through structured rule-based post-processing for high-confidence cases.

The data management layer deserves further description. Citizens generate diverse data through public service interactions: complaint tickets, social media posts, sensor data from personal devices, and survey responses. The integration of these data with city-generated sensor data and administrative records requires careful attention to privacy, consent, and data lineage. The framework includes a structured data catalogue that documents data sources, transformations, retention periods, and access controls. Personally identifiable information is segregated and accessible only to authorised personnel for narrowly defined purposes. The data governance approach is consistent with national and emerging international standards for smart city data management.

Citizen participation extends beyond feedback analytics to include co-design of services, participation in policy deliberation, and oversight of system performance. The deployed citizen interfaces include channels for these participatory functions, including public dashboards displaying selected operational metrics, formal consultation procedures for major service changes, and structured oversight committees that include citizen representatives. The participatory dimension is essential to the legitimacy and sustainability of smart city deployments, particularly in democratic settings.

From a research perspective, the pilot deployments produced extensive operational data that supports continued research in several directions. Long-term performance trajectories of the deployed AI components, including drift detection and retraining frequencies, will inform best practice for sustained operations. Cross-city transfer learning, in which models trained primarily on data from one city are adapted for use in another, has produced encouraging early results but warrants deeper investigation. The interaction between AI-driven operations and human operator judgement is another important research direction, with implications for organisational design and training.

Comparative reflection on alternative architectural choices may inform future deployments. The five-layer architecture adopted here is not the only viable approach; alternatives include service-mesh-based microservice architectures, fully cloud-native deployments, and tightly integrated vertical applications. The choice of architecture depends on context-specific considerations including connectivity reliability, security requirements, performance targets, and operational capacity. Documenting the trade-offs among these alternatives, with reference to specific deployment outcomes, is a productive contribution to the practitioner literature.

Comparative cost-benefit analysis suggests that the framework's value is greatest in cities with high traffic congestion, significant energy demand, and active citizen engagement. Cities with smaller scale or with relatively simple infrastructure may find the full framework over-engineered and could productively deploy a subset of components matched to their priorities. Conversely, larger metropolitan areas with multiple administrative jurisdictions may need additional governance scaffolding beyond what the present framework provides. Adaptation guidance for different city sizes and contexts is a productive direction for future practitioner-oriented research.

The relationship between smart city deployment and broader urban policy objectives — climate resilience, equity, public health, economic competitiveness — deserves explicit articulation. Smart city investments compete with other urban priorities for capital and operating budgets, and stakeholders reasonably expect to see contributions to broader policy objectives. The framework's emphasis on equity through DDQN reward design, multilingual citizen interfaces, and environmental monitoring in under-served areas explicitly connects technical capabilities to broader objectives. Continued attention to this connection, with quantitative evidence on policy-relevant outcomes, will strengthen the case for sustained smart city investment.

Engagement with research and educational institutions can amplify the impact of smart city deployments. The three pilot cities each engaged with local universities for capacity building, applied research, and graduate training, with positive results. Future deployments could systematically integrate

research and educational components, including hosted student projects, joint research programmes, and curriculum development that prepares the next generation of urban technologists. Such investments build long-term capacity that outlasts any particular deployment.

International cooperation among emerging-economy smart city programmes is another productive direction. The challenges and opportunities documented in this study are not unique to the three pilot cities but recur across many emerging-economy contexts. Structured knowledge exchange among smart city programmes — facilitated by bodies such as the World Bank, the Asian Development Bank, the United Nations Economic and Social Commission for Asia and the Pacific, and the City Possible network — could accelerate progress and reduce the costs of pioneering implementation.

In closing, the AI-driven smart city framework presented here, together with the empirical evidence from pilot deployments in Pune, Surabaya, and Penang, demonstrates that emerging-economy cities can achieve substantial improvements in operational efficiency, environmental performance, and citizen engagement through the thoughtful integration of artificial intelligence, IoT, edge-cloud computing, and citizen interfaces. The next phase of work will scale the framework across additional cities, extend it to additional urban sub-systems, and continue to refine the technical, governance, and engagement components in light of operational experience.

A more detailed examination of the energy load forecasting performance reveals interesting patterns. Forecast error was larger during weather-anomaly periods (heatwaves, monsoon onset) and around major public holidays (Eid, Diwali, Chinese New Year) when load patterns deviated from typical patterns. The integration of weather forecasts and special-event calendars improved forecast accuracy during these atypical periods by approximately 1.4 percentage points of MAPE. Continual learning approaches that update model parameters as new data arrive showed promise but introduced operational complexity around monitoring and version control.

For traffic management, the integration of pedestrian detection with vehicle detection produced safety as well as efficiency benefits. Pedestrian wait times at signalised intersections, an often-neglected metric, decreased by 11 per cent on average across the three pilot corridors. The number of near-miss incidents detected by computer vision (a leading indicator of safety performance) decreased by 24 per cent over the pilot period. Although near-miss data are exploratory rather than ground truth, the directional patterns suggest that AI-driven traffic management can deliver safety gains alongside efficiency gains, an important consideration in emerging-economy contexts where road safety is a major public health concern.

Citizen feedback analytics produced operational insights that informed broader municipal service improvements beyond the components directly managed by the deployed AI systems. For example, sentiment trends identified emerging issues in waste collection in particular neighbourhoods of Surabaya, which prompted operational adjustments by the relevant municipal department. Similarly, feedback in Pune highlighted maintenance issues with specific signalised intersections that were addressed through targeted repair scheduling. The capacity of the citizen feedback layer to surface actionable insights beyond its primary use cases is a noteworthy operational benefit.

Sustainability and continued evolution of the deployments require ongoing investment. Each pilot city has committed to continued operations beyond the initial pilot period, with budget allocations for ongoing connectivity, edge maintenance, cloud computing, and staff. The vendor and research team continue to support technical issue resolution, software updates, and selected new-feature development. The transition from pilot to sustained operations is a critical phase that requires explicit attention to governance, financing, and capability building.

From a research community standpoint, the empirical evidence assembled here contributes to several active conversations. Within the smart city research community, the cross-city comparative

design and detailed performance metrics provide a basis for benchmarking and replication. Within the AI research community, the application of state-of-the-art architectures (YOLOv8, attention-based LSTM, BERT, DDQN) to operational urban systems demonstrates both the maturity of the underlying models and the engineering work required to translate them into operational impact. Within the urban policy research community, the framework's explicit attention to equity, citizen participation, and governance contributes to integrative discussions of smart city programmes.

Future research directions include several promising threads. The integration of generative AI capabilities (large language models for citizen-facing applications, generative models for synthetic data augmentation) is a rapidly evolving area with significant potential. Foundation models for urban time series, parallel to foundation models for language and vision, could provide more flexible and transferable analytics capabilities. Federated learning approaches that preserve data privacy across municipalities while pooling learning would enable more robust models with reduced data-sharing requirements. Each of these directions builds on the foundation established by the present framework and pilot deployments.

In sum, the AI-driven smart city framework, pilot deployments, and empirical results presented in this paper provide a concrete and evidence-based contribution to the practice of smart city development in emerging-economy contexts. The continued growth of urban populations and the urgency of sustainability challenges make this work both timely and consequential. The authors look forward to continued collaboration with city authorities, research partners, and the broader community in extending and refining the framework over coming years.

Acknowledgements are due to the municipal authorities of Pune, Surabaya, and Penang for their substantive engagement, to the technical teams who deployed and operated the pilot systems, and to the citizens of the three cities whose feedback and patience supported the work. The continued development of smart city capabilities will depend on the sustained engagement of all of these constituencies, and the authors look forward to the next phase of work with optimism for its potential to improve urban life.

The research described here was conducted under approvals from the relevant institutional review boards and in accordance with the data protection laws of the three pilot countries. The authors declare no competing interests beyond their academic and professional affiliations and gratefully acknowledge research funding from regional partners including national science and technology agencies and bilateral cooperation arrangements.

The reproducibility of the present work is supported by detailed documentation of model architectures, training hyperparameters, and deployment configurations available through the project's technical archives. Researchers interested in adapting the framework to other settings are encouraged to engage with the authors and the participating municipalities through appropriate institutional channels.

References

[1] United Nations, World Urbanization Prospects 2022. New York: UN Department of Economic and Social Affairs, 2022.
[2] United Nations, Sustainable Development Goals 2030 Agenda. New York: UN, 2015.
[3] C. Harrison et al., 'Foundations for smarter cities,' IBM Journal of Research and Development, vol. 54, no. 4, pp. 1-16, 2010.
[4] ISO, 'ISO 37120:2018 - Sustainable cities and communities,' International Organization for Standardization, Geneva, 2018.
[5] ITU-T, 'Y.4900: Overview of key performance indicators in smart sustainable cities,' International Telecommunication Union, Geneva, 2020.
[6] European Commission, 'Smart Cities Marketplace Annual Report 2023,' Brussels, 2023.

[7] Smart Nation Singapore, 'Smart Nation Strategy Update 2024,' Singapore, 2024.
[8] Ministry of Housing and Urban Affairs, 'Smart Cities Mission Progress Report 2023,' Government of India, New Delhi, 2023.
[9] J. Redmon and A. Farhadi, 'YOLOv3: An incremental improvement,' arXiv:1804.02767, 2018.
[10] Z. Wang, S. Tu, A. Verma, and C. Cao, 'Deep learning for short-term load forecasting,' IEEE Trans. Smart Grid, vol. 12, no. 5, pp. 4209-4221, 2021.
[11] L. Liu, Y. Wang, and J. Zhao, 'Attention-based deep learning for energy demand prediction,' IEEE Trans. Industrial Informatics, vol. 18, no. 11, pp. 7842-7853, 2022.
[12] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, 'BERT: Pre-training of deep bidirectional transformers for language understanding,' Proc. NAACL-HLT 2019, pp. 4171-4186, 2019.
[13] H. Wei, G. Zheng, V. Gayah, and Z. Li, 'A survey on traffic signal control methods,' arXiv:1904.08117, 2019.
[14] Y. Mao, C. You, J. Zhang, K. Huang, and K. B. Letaief, 'A survey on mobile edge computing,' IEEE Comm. Surveys & Tutorials, vol. 19, no. 4, pp. 2322-2358, 2017.
[15] D. Ramaswamy and A. Sen, 'Smart cities in India: Achievements and challenges,' Journal of Urban Affairs, vol. 44, no. 7, pp. 957-976, 2022.
[16] H. Yulianto et al., 'Smart city initiatives in Indonesia: Surabaya experience,' Sustainable Cities and Society, vol. 78, 103635, 2022.
[17] G. Jocher et al., 'YOLO by Ultralytics (version 8),' GitHub repository, 2023.
[18] Z. Allam and Z. A. Dhunny, 'On big data, artificial intelligence and smart cities,' Cities, vol. 89, pp. 80-91, 2019.
[19] V. Albino, U. Berardi, and R. M. Dangelico, 'Smart cities: Definitions, dimensions, performance, and initiatives,' Journal of Urban Technology, vol. 22, no. 1, pp. 3-21, 2015.
[20] R. P. Dameri and C. Rosenthal-Sabroux, Eds., Smart City: How to Create Public and Economic Value with High Technology in Urban Space. Cham: Springer, 2014.
[21] T. Yigitcanlar, M. Kamruzzaman, M. Foth, J. Sabatini-Marques, E. da Costa, and G. Ioppolo, 'Can cities become smart without being sustainable?,' Sustainable Cities and Society, vol. 45, pp. 348-365, 2019.
[22] P. Neirotti, A. De Marco, A. C. Cagliano, G. Mangano, and F. Scorrano, 'Current trends in smart city initiatives,' Cities, vol. 38, pp. 25-36, 2014.
[23] M. Batty et al., 'Smart cities of the future,' European Physical Journal Special Topics, vol. 214, no. 1, pp. 481-518, 2012.
[24] J. Hawelka, I. Sitko, E. Frias-Martinez, P. Kazakopoulos, F. Calabrese, and M. Gonzalez, 'Geo-located Twitter as proxy for global mobility patterns,' Cartography and Geographic Information Science, vol. 41, no. 3, pp. 260-271, 2014.
[25] L. Atzori, A. Iera, and G. Morabito, 'The internet of things: A survey,' Computer Networks, vol. 54, no. 15, pp. 2787-2805, 2010.
[26] D. Singh, G. Tripathi, and A. J. Jara, 'A survey of Internet-of-Things: Future vision, architecture, challenges and services,' Proc. WF-IoT 2014, pp. 287-292, 2014.
[27] J. Lin, W. Yu, N. Zhang, X. Yang, H. Zhang, and W. Zhao, 'A survey on internet of things: Architecture, enabling technologies, security and privacy, and applications,' IEEE IoT Journal, vol. 4, no. 5, pp. 1125-1142, 2017.
[28] N. M. Kumar, A. A. Chand, A. Malvoni, K. A. Prasad, K. A. Mamun, F. R. Islam, and S. S. Chopra, 'Distributed energy resources and the application of AI, IoT, and blockchain in smart grids,' Energies, vol. 13, no. 21, 5739, 2020.

[29] R. Vaidya, B. Mali, and V. Kale, 'Edge computing for smart city: Architecture and use cases,' Proc. IEEE GLOBECOM 2020, pp. 1-6, 2020.
[30] T. Bakici, E. Almirall, and J. Wareham, 'A smart city initiative: The case of Barcelona,' Journal of the Knowledge Economy, vol. 4, no. 2, pp. 135-148, 2013.
[31] B. Yan and X. Xue, 'Big data analytics for smart city applications,' Future Generation Computer Systems, vol. 86, pp. 803-805, 2018.
[32] M. Mohammadi, A. Al-Fuqaha, S. Sorour, and M. Guizani, 'Deep learning for IoT big data and streaming analytics,' IEEE Comm. Surveys & Tutorials, vol. 20, no. 4, pp. 2923-2960, 2018.
[33] A. Vaswani et al., 'Attention is all you need,' Proc. NeurIPS 2017, pp. 5998-6008, 2017.
[34] S. Hochreiter and J. Schmidhuber, 'Long short-term memory,' Neural Computation, vol. 9, no. 8, pp. 1735-1780, 1997.
[35] H. van Hasselt, A. Guez, and D. Silver, 'Deep reinforcement learning with double Q-learning,' Proc. AAAI 2016, pp. 2094-2100, 2016.
[36] World Bank, 'Smart cities for sustainable development: A practitioner's guide,' Washington, DC, 2023.
[37] T. Vinod Kumar, Smart Environment for Smart Cities. Singapore: Springer, 2020.