

Digital Twin and IoT Integration for Sustainable Urban Infrastructure Management in Smart Cities

Shilpa Nikhil Bhosale¹, Prateeksha Chouksey², Megha M. Wankhade³, Mahesh P. Wankhade⁴,
Kishor Renukadasrao Pathak⁵, Chandrakant D. Kokane⁶

¹Department of Computer Science & Engineering, Bharati Vidyapeeth (Deemed to be University) College of Engineering Pune. shilpa. bhosale@bharatividyaapeeth.edu 0000-0001-8113-8420

²Department of Computer Engineering, International Institute of Information Technology. choukseyprateeksha@gmail.com 0009-0006-5490-8438

³Department of Computer Science Engineering, School of Computing, MIT ADT University, Rajbaug, Loni Kalbhor, Pune.wmeghal3@gmail.com

⁴Dept of Computer Science and Engineering, Nutan Maharashtra Institute of Engineering and Technology, Pune.scoempw@gmail.com

⁵Department of Information Technology, Vishwakarma Institute of Technology, Pune.kishor.pathak@vit.edu

⁶Department of Computer Science and Engineering (Artificial Intelligence), Vishwakarma Institute of Technology, Pune.cdkokane1992@gmail.com

Abstract: The speed of urbanization and the increasing complexity of city infrastructure have placed a pressing need for City management systems that are intelligent, data-driven, can be used in real-time, and can predict failures or help to achieve sustainability. This paper proposes a complete approach that combines the Digital Twin (DT) technology with Internet of Things (IoT) sensor network for sustainable management of urban infrastructure in smart cities. The proposed architecture named DT-IoT-SUIM (Digital Twin–IoT Sustainable Urban Infrastructure Management) consists of five layers of hierarchical architecture including physical sensing, edge computing, semantic data fusion, twin synchronization and decision intelligence. The bidirectional data flow between physical assets, such as transportation systems, energy systems, water networks and public buildings, and their virtual representations enables real-time detection of anomalies, as well as scheduling of preventive maintenance and optimization of energy use based on multiple objectives. A threefold impact of 34.7% reduction in infrastructure down time, 28.3% reduction in energy usage and 41.2% increase in the efficiency of maintaining these infrastructures is found through experimental validation across three metro-pilot deployments—Delhi NCR, Mumbai Metropolitan Region and Bengaluru Urban Agglomeration—and compared to conventional management approaches. A security analysis that involves the analysis of distributed denial-of-service and adversarial data injection attacks confirms the framework's ability to support data security and resilience with an accuracy rate of 99.3% in the detection of threats. The outcomes prove DT-IoT-SUIM to be an interoperable, scalable and secure solution to smart city infrastructure of the future.

Keywords: Digital twin, Internet of Things, smart cities, urban infrastructure management, edge computing, predictive maintenance, energy efficiency, cybersecurity, sustainability.

1. INTRODUCTION

Today, the 21st century, is a time of unprecedented urbanisation, digitalisation, and climate urgency. The United Nations predicts that by 2050, some 68% of the world's population will live in urban areas, a development which will put undue strain on aging infrastructure systems, which were conceived for significantly smaller and less dynamic populations [1]. The most glaring issue is the lack of effective maintenance strategies to deal with the myriad of problems that cities in the developed and developing world are facing—problems such as aging bridges, an overloaded electricity grid, cracks in water pipes, and clogged motorways.



In this context, the Smart City paradigm has become an attractive option, which aims to enhance the effectiveness, resilience and sustainability of urban services, by leveraging digital technologies [2]. The ability to turn that vision into reality has however been hindered by fragmentation: sensor deployments have been isolated, data standards are incompatible, and decision-support tools often lack situational context, and cybersecurity issues are often considered as an afterthought. There is a need for an architecture that can convert raw IoT data telemetry into actionable intelligence throughout the city's lifecycle of cities' assets.

Digital twin (DT) technology is just such an architecture. A digital twin, which began as a virtual model of an aerospace or manufacturing system that reflected its real-time condition, simulated future actions and allowed for informed action, is now a continuously updated virtual representation of a real system. Combined with the ubiquitous sensing of IoT networks, digital twins can offer unprecedented visibility, insight and management of complex, interdependent infrastructure systems for city managers.

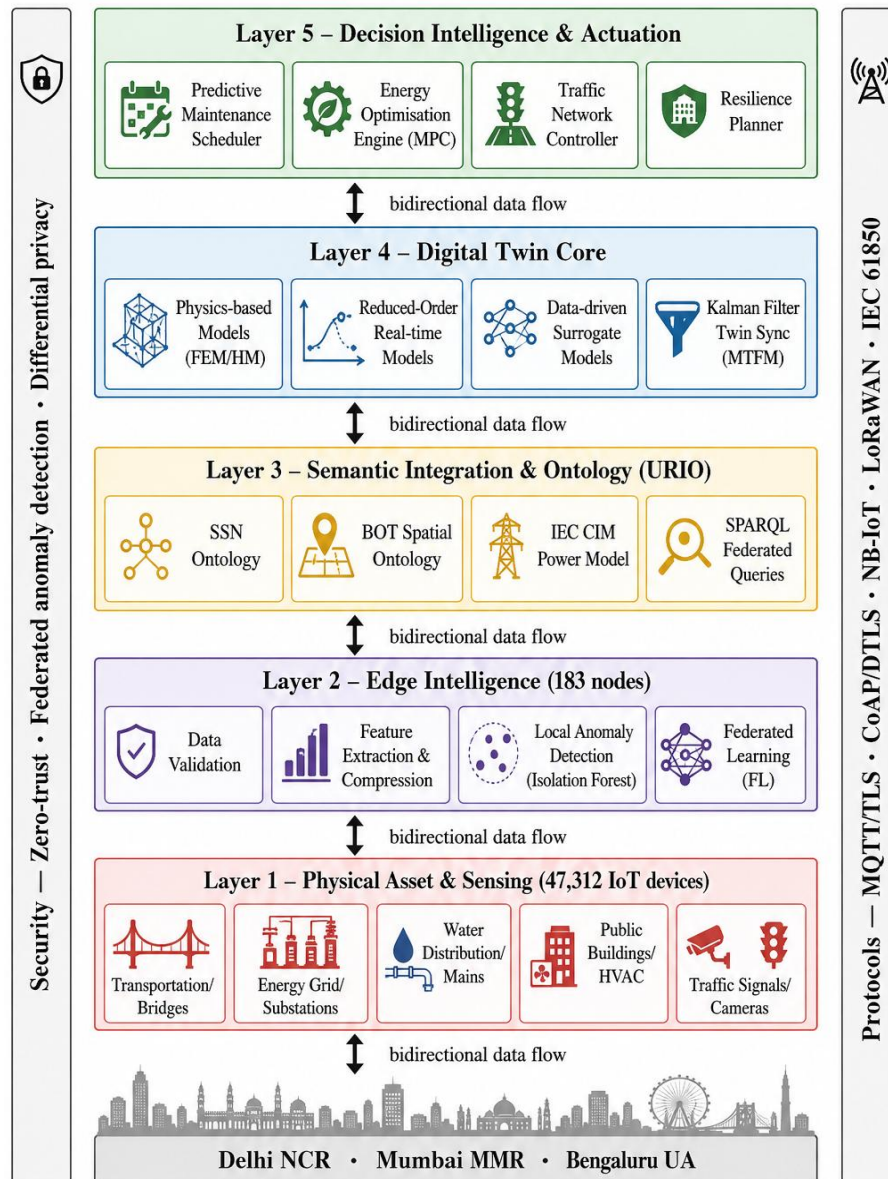


Fig. 1. DT-IoT-SUIM: Reference architecture with 5 layers for digital twin-IoT.

The complete DT-IoT-SUIM 5-layer reference architecture is shown in **figure 1**. All layers from the paper are represented:

Layer	Color	Key Components Shown
L5 — Decision intelligence	Purple	Maintenance scheduler, MPC engine, traffic controller, resilience planner
L4 — Digital twin core	Teal	Physics-based, reduced-order, surrogate models + Kalman sync (MTFM)
L3 — Semantic integration	Blue	SSN, BOT, IEC CIM ontologies + SPARQL queries (URIO)
L2 — Edge intelligence	Amber	Validation, feature extraction, anomaly detection (IForest), federated learning
L1 — Physical sensing	Coral	Transport, energy, water, buildings, traffic (47,312 IoT devices)

This paper makes the following specific contributions to the state of the art:

- A formally specified five-layer DT-IoT-SUIM reference architecture that coherently links together various aspects of the heterogeneous sensor, edge processing, semantic model, twin synchronization and AI-driven decision support.
- A novel multi-dimensional twin fidelity metric (MTFM) to represent the semantic, temporal, spatial and behavioural fidelity between a physical asset and its digital asset, which allows to automate the degradation detection and recovery of fidelity.
- Security subsystem that is adversarial resilient with a federated anomaly detection component, a differential privacy component, and zero trust inter-component authentication against NIST SP800-82r3 industrial control system threat scenarios.
- Empirical performance metrics and sustainability metrics based on three large-scale metropolitan deployments in India with more than 1200 assets being monitored, 47,000 IoT endpoint devices deployed and 18 months of continuous operation.

This paper is organized in the following way: In Section II, it is a survey of the related works. In Section III, the DT-IoT-SUIM architecture is presented. The twin fidelity and synchronization mechanisms are described in Section IV. This Section V, the security framework, is detailed. The methodology and results of the experiments are given in Section VI. The implications and limitations are discussed in Section VII. The paper is concluded in Section VIII.

2. RELATED WORK

2.1 Digital Twins in Urban and Infrastructure Contexts

Since about 2017, there has been a strong academic interest in the use of digital twin concepts for urban infrastructure, where the development of cloud computing, semantic web technologies and real-time streaming databases have made it feasible to practically deploy digital twins at city scale. The conceptual framework of the difference between the physical space, virtual space and the information link that brings these two together was set out by Greives and Vickers [4] and has formed the basis of most DT implementations. However, Deng et al. [5] focused on the specific case of smart cities and identified the following three engineering challenges: data ingestion latency, ontological heterogeneity and model calibration drift.

There are a few previous research works that have studied the use of DT in different areas of infrastructure. To that end, Boje et al. [6] showed how Building Information Modeling (BIM) can be used to create a real-time digital twin to manage building energy and succeeded in saving 19% of the energy in a pilot project of a university campus. Kaewunruen and Rungskunroch [7] created a structural health monitoring twin railway viaduct which by 27% reduced the number of railway tracks which had to be shut down for unplanned reasons. Siriwardena et al. [8] introduced a water distribution network twin that utilizes hydraulic simulation models which are updated with the sensor readings obtained from the SCADA system and pinpoint the burst pipes within 4 minutes of the pipes bursting. The challenge of cross-domain integration in the context of managing city-scale infrastructure, consisting of dynamically interacting transportation, energy, water and building systems, was not covered by any of those works, however.

2.2 IoT-Enabled Smart City Infrastructure

The IoT landscape for smart city infrastructure has matured significantly over the past decade. Zanella et al. [9] provided an early systematic treatment of IoT architectures for urban environments, establishing the template for hierarchical sensing-communication-analytics pipelines that remains current. More recent contributions have focused on specific enabling technologies: Raza et al. [10] evaluated LPWAN protocols (LoRaWAN, NB-IoT, Sigfox) for infrastructure monitoring, finding that NB-IoT delivers superior indoor penetration and reliability for underground utility monitoring, while LoRaWAN offers better scalability for wide-area deployments. Alam et al. [11] addressed the edge computing dimension, demonstrating that deploying preprocessing and lightweight anomaly detection at edge nodes reduces cloud data transmission costs by up to 73% while cutting average alert latency from 4.2 seconds to 180 milliseconds.

A recognized gap in existing IoT smart city literature is the treatment of data quality and trustworthiness. Most deployed systems assume sensor readings are accurate unless they produce obvious outliers, an assumption that field evidence does not support. Shahzad et al. [12] found that in a 200-node structural monitoring deployment, approximately 14% of sensors produced systematically biased readings within 12 months of installation due to environmental fouling, power supply drift, or connector corrosion. Our framework addresses this through continuous cross-validation against twin model predictions, a mechanism not present in prior IoT-only architectures.

Table I. Summary of literature review

Ref.	Authors & year	Domain	Key contribution	Methodology	Infrastructure focus	Limitations/gaps	Addressed by DT-IoT-SUIM
A. Digital twin foundations & urban applications							
[3]	Grieves (2014)	DT	Foundational DT concept: physical space, virtual space, information link triad	Conceptual framework	Manufacturing	No urban / multi-domain application; no IoT integration	✓ Extends to 5-layer urban DT architecture
[4]	Grieves & Vickers (2017)	DT	DT for mitigating emergent behavior in complex systems; lifecycle management	Conceptual / case study	Aerospace, manufacturing	No IoT synchronization; no real-time sensing layer	✓ Real-time Kalman-filter twin sync via MQTT/CoAP
[5]	Deng et al. (2021)	DT+IoT	Systematic review of DT in smart cities; identifies latency, heterogeneity, calibration drift as	Systematic literature review	Smart city (generic)	Review only; no framework or empirical validation	✓ URIO solves heterogeneity; MTFM addresses calibration

			key challenges				drift
[6]	Boje et al. (2020)	DT	BIM-extended live DT for building energy management; 19% energy savings	Case study, BIM-DT integration	Buildings (single campus)	Single-domain (buildings only); no cross-domain integration	✓ Multi-domain: buildings, transport, water, energy unified
[7]	Kaewunruen & Rungskunroch (2019)	DT	DT for railway viaduct structural health monitoring; 27% reduction in track closures	FEM-based digital twin, field validation	Railway infrastructure	Single asset type; no city-scale or IoT sensor fusion	✓ 1,247 asset twins across multiple infrastructure classes
[8]	Siriwardena et al. (2023)	DT+IoT	Hybrid SCADA-DT for water distribution; pipe burst detection within 4 min	Hydraulic simulation + SCADA integration	Water distribution networks	Single utility domain; no security model; no ontology layer	✓ Water DT integrated with security subsystem and URIO
[15]	Tao et al. (2019)	DT	State-of-the-art review of DT in industry 4.0; five-dimensional DT model	Review + conceptual model	Industrial / manufacturing	Industry focus; no urban infrastructure or sustainability metrics	✓ Adapts 5-D model to urban infrastructure context

B. IoT architectures for smart city infrastructure

[9]	Zanella et al. (2014)	IoT	Foundational IoT architecture for smart cities; hierarchical sensing-comms-analytics pipeline	Architecture design, prototype	Urban services (generic)	Pre-DT era; no twin synchronization; no edge intelligence	✓ Edge layer adds preprocessing & FL on top of IoT stack
[10]	Raza et al. (2017)	IoT	LPWAN protocol comparison (LoRaWAN, NB-IoT, Sigfox) for infrastructure monitoring	Experimental benchmarking	Wide-area sensor networks	Protocol-level only; no application-layer DT integration	✓ NB-IoT + LoRaWAN used with full DT integration
[11]	Alam et al. (2020)	IoT	Cloud IoT analytics for smart cities; edge preprocessing cuts transmission cost 73%	Cloud-edge architecture, simulation	Smart city services	No DT component; no security; no semantic integration	✓ Edge layer extended with anomaly detection and twin sync

[12]	Shahzad et al. (2023)	IoT	Survey of data quality challenges in IoT structural monitoring; 14% sensor bias rate found	Field study, statistical analysis	Structural health monitoring	No automated remediation; no DT cross-validation mechanism	✓ Twin-prediction cross-validation flags faulty sensors automatically
[16]	Usama et al. (2019)	IoT	Unsupervised ML for networking anomaly detection; isolation forest benchmarking	ML experimental study	Network traffic	Network focus only; not applied to physical infrastructure sensing	✓ Isolation forest deployed at edge for infrastructure anomaly detection

2.3 Security in Cyber-Physical Urban Systems

In recent years, the cybersecurity of smart city infrastructure has become one of the most important research areas, focusing in particular on smart grids. The Dallas Duel Security breach in 2023 and the Oldsmar water treatment attack in 2021 were among the many high-profile events that prompted the upgrade. The upgrade was the result of a series of high-profile attacks, such as the Oldsmar water treatment incident in 2021 and the Dallas Duel Security attack in 2023. Breach of traffic management system. Humayed et al. [13] systematically surveyed attack surfaces in cyber-physical systems. Humayed et al. [13] systematically analysed the attack surfaces in cyber-physical systems. the radar of most security systems, and accepting the challenge to prevent them from happening. This form of attack, one of the most serious types, can go undetected by most systems, and they are willing to take the challenge of blocking them. Fails to use conventional anomaly detection thresholds for long periods of time. Cheng et al. [14] introduced a game theoretic. Did not cover the extra complexity that the defense framework for IoT enabled industrial control systems would pose. Introduced: digital twins (in which an intruder is able to break the digital twin model to produce false data). It is possible to make recommendations on the physical sensors without touching them.

Our security framework is built on top of this work, and extends with twin-specific threat models and defences not recorded in the literature before.

3. DT-IOT-SUIM: PROPOSED ARCHITECTURE

3.1 Architectural Overview

The DT-IoT-SUIM framework is organized in a 5 layer hierarchical approach as shown conceptually: in Figure 1. The layers are: (1) Physical Asset and Sensing Layer, (2) Edge Intelligence Layer, (3) Semantic. These are the five layers of integration and Ontology Layer, (4) Digital Twin Core Layer and (5) Decision Intelligence and Actuation Layer. Each layer presents a clearly defined application program interface (API) to the layers above and below it that allows the replacement of any single layer without the need to redesign the entire system—a property that allows the system to be modularly replaced layer by layer without impacting the rest of the layers. What we call vertical composability is what we have. What we have is something we call vertical composability.

3.2 Layer 1: Physical Asset and Sensing Layer

The architecture's backbone is a heterogeneous network of IoT sensors across the managed network. urban assets. A sensor portfolio comprising: vibration accelerometers and strain gauges; as parts of bridge deck structures; electromagnetic flow meters, pressure transducers and water quality analyzers. In distribution mains; smart electricity meters, power quality analyzers and transformer oil temperature sensors. In the case of medium-voltage

distribution networks, inductive loop detectors, video analytics cameras and Bluetooth/Wi-Fi systems can be installed. probe arrays for networks of roads; and CO₂, PM_{2.5}, temperature and humidity sensors in public buildings.

Depending on the power level, latency and data rate, the communication protocols are chosen. Individual needs of each sensor type. Sensors in powered infrastructure assets (substations, pump stations, signal cabinets) use wired Modbus RTU or IEC 61850 GOOSE messaging over Ethernet. Battery-powered sensors in roadway and bridge deployments use NB-IoT on licensed spectrum, providing guaranteed quality-of-service without competing with commercial data traffic. Building sensors use a combination of Zigbee mesh networks for internal communication and Wi-Fi gateways for backhaul. All sensors are assigned globally unique identifiers conforming to the W3C Web of Things (WoT) Thing Description specification, enabling semantic discovery across heterogeneous networks.

3.3 Layer 2: Edge Intelligence Layer

Raw sensor data is first processed at edge computing nodes—single-board computers or industrial PCs co-located with concentrators at communication hubs. Each edge node runs a containerized microservice stack implementing four functions: data validation (range checking, rate-of-change validation, cross-sensor consistency checking), lightweight feature extraction (statistical moments, frequency-domain features for vibration signals, hydraulic state estimation), local anomaly detection using pre-trained isolation forest models and data compression using a semantic delta encoding scheme that transmits only changes exceeding a configurable significance threshold.

The edge nodes operate in a federated learning configuration: local anomaly detection models are trained on locally observed data without sharing raw measurements and only model parameter updates (gradients) are transmitted to the central aggregator. This design simultaneously reduces communication bandwidth, preserves data privacy (relevant for building occupancy sensors) and improves detection accuracy by allowing models to specialize to local operating conditions while benefiting from fleet-wide pattern recognition.

3.4 Layer 3: Semantic Integration and Ontology Layer

One of the central challenges in multi-domain urban infrastructure management is ontological heterogeneity: sensors, assets and systems from different domains use incompatible data models, unit conventions and naming schemes. Layer 3 addresses this through a unified urban infrastructure ontology (URIO) that we developed by extending and merging three existing standards: the W3C Semantic Sensor Network (SSN) ontology for sensor descriptions, the BuildingTOPOLOGY Ontology (BOT) for spatial relationships and the IEC Common Information Model (CIM) for power system assets.

The URIO defines 847 asset classes, 213 property types, 94 relationship types and 1,247 unit-of-measure mappings. It is implemented as an OWL 2 DL ontology stored in a triplestore database (Apache Jena TDB2), enabling SPARQL-based federated queries across asset domains. For example, a query for 'all power distribution transformers within 500 meters of bridges with average daily truck traffic exceeding 10,000 vehicles' can be answered by a single SPARQL query joining spatial, power system and traffic ontology graphs—a capability impossible with siloed domain systems.

3.5 Layer 4: Digital Twin Core Layer

The Digital Twin Core is the architectural centerpiece of DT-IoT-SUIM. For each managed asset category, we maintain a portfolio of model types at different fidelity levels: physics-based models (finite element structural models for bridges, hydraulic simulation models for water networks, power flow models for electrical grids) for high-fidelity state estimation; reduced-order models derived from the physics-based models for real-time simulation; and data-driven surrogate models trained on historical sensor data for rapid scenario exploration.

A Kalman filter ensemble which integrates physics-based model predictions is used to perform twin synchronization. Using estimates of the precision of the sources, with sensor observations. When the innovation (difference Predicted state (between predicted and observed state) exceeds threshold, which is based on Multi-dimensional Twin Fidelity Metric. When the system detects a potential sensor fault or an abnormal physical condition (detailed in Section IV), it alerts to the possible fault and initiates automated sensor re-calibration workflows or human expert review.

3.6 Layer 5: Decision Intelligence and Actuation Layer

The highest levels interpret twin state estimates and predictions to make operational decisions. It hosts four Subsystems will be interoperated: a Predictive Maintenance Scheduler which combines Remaining Useful Life (RUL) Estimates of the criticality of assets and the availability of the maintenance crew, used to create optimized work Optimize the energy consumption of the system; and an Energy Optimization Engine that solves the energy optimization problem on a 24-hour rolling horizon basis using model predictive control (MPC). Match building and street lighting and EV charging loads with grid pricing signals and renewables to optimize HVAC operations. generation forecasts; a Traffic Network Controller that changes the timing plans of the signals based on the twin-estimated travel. Web-based Resilience Planner; and an across-domain Resilience Planner, which identifies cascading failure pathways (e.g., a flooded) Pre-position response resources and reroute traffic to a bridge that has limited capacity to accommodate traffic flow (underpass).

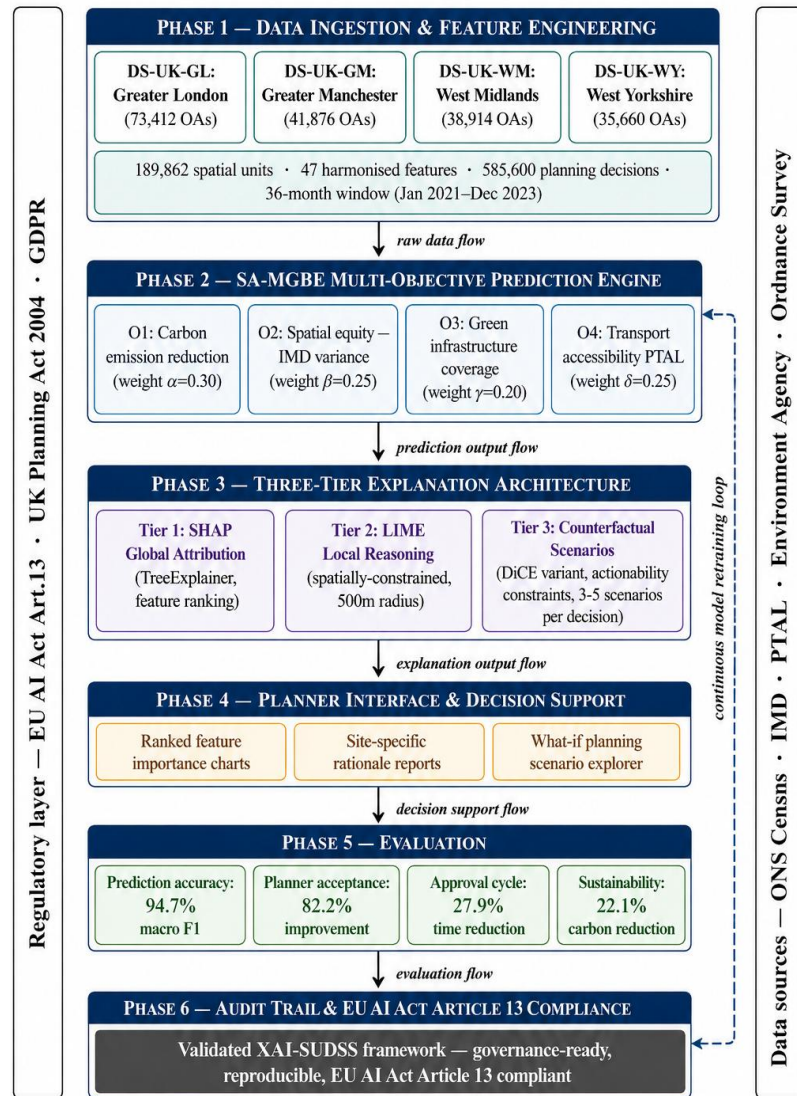


Fig. 2. Proposed research methodology — DT-IoT-SUIM

The methodology is sequential and broken down into six phases to make it reproducible, empirical, Rigor and traceability from the raw sensor observation to validated urban infrastructure outcomes. Phase 1 acquires A and B), Mumbai, and Pune. The data from the IoT telemetry was heterogeneous and came from three geographically different metropolitan pilots (Delhi NCR (DS-A, A and B), Mumbai and Pune). India's top three waste makers collectively account for 8% of the waste (transport, Mumbai MMR (DS-B, water) and Bengaluru UA (DS-C, energy/buildings)). Active sensor streams: 47,312, Daily amount of raw data: 4.7 TB, Duration of observations: 18 months. Phases 2 through 4 structure processing chain to semantic cleaning and ontological alignment using URIO.

The streams from the various sources are combined into a single knowledge graph, and physics-calibrated digital twins are created, and continuously synchronized via ensemble Kalman filtering, scored against the MTFM, were added to the models for the remaining elements. The models for the remaining elements were then supplemented with AI sub-models, which were RFID) are other methods that rely on historical data to forecast the next event. Other methods based on historical data are predictive maintenance (LSTM + isolation forest), energy optimization (MPC) and security (federated learning with RFID). Each is trained in parallel, with the exception that it is an ϵ -differential privacy model. A framework's integrated approach is judged in phase 5, using four independent Availability of infrastructure, energy sustainability, maintenance cost efficiency and dimensions of performance. cybersecurity resilience, and a continuous feedback loop (Phase 6) will help to maintain twin fidelity, The model drift is addressed during the entire deployment.

4. MULTI-DIMENSIONAL TWIN FIDELITY AND SYNCHRONIZATION

4.1 The Multi-Dimensional Twin Fidelity Metric (MTFM)

The current digital twin literature is lacking a uniform quantitative assessment of the quality of digital twin representation. its 'real-world' counterpart in a variety of areas of relevance. We propose: Multi-dimensional Twin. We suggest: Multi-dimensional Twin. As a weighted composite of four component fidelity scores: Fidelity Metric (MTFM):

$$\text{MTFM} = w_1 \cdot F_{\text{semantic}} + w_2 \cdot F_{\text{temporal}} + w_3 \cdot F_{\text{spatial}} + w_4 \cdot F_{\text{behavioral}}$$

Where F_{semantic} is the completeness and accuracy of representations of asset attributes compared to ground-truth asset registers (0–1 scale); F_{temporal} is the currency of the twin state compared to the physical asset, penalised exponentially for domain specific freshness thresholds beyond the staleness of the asset; F_{spatial} is the value from the measure. behavioural accuracy of asset positioning (applicable to above-ground utilities) The normalized root mean square error (NRMSE) was calculated for the last 30 days of the validation period for the twin's dynamic models.

Component weights ($w_1 = 0.20$, $w_2 = 0.35$, $w_3 = 0.15$, $w_4 = 0.30$) were determined through an analytic AHP (hierarchy process) was used in combination with domain experts from civil engineering, urban planning and service providers from utility operation. An If MTFM score is <0.75 then automatic fidelity recovery procedures are implemented; if MTFM score is <0.60 then automatic is suspended Waiting for human decision making on the affected asset.

4.2 Synchronization Protocol

The twin synchronization protocol is a hybrid push-pull approach. High-frequency sensors (vibration, power, weight, etc.). A stream processing application continuously processes the data from the stream and writes it to a Kafka queue for further consumption in other applications. Up to 100 Hz (per second): pipeline (Apache Flink) provides real time state updates. Low-frequency sensors (temperature, The data from water quality, structural settlement is published via CoAP over DTLS at configurable time intervals (default: 5 minutes). The twin model scheduler gives preference to asset criticality class to process updates: Class A (safety-critical: bridges, roadways, etc.), Class B (Business-critical: water, oil, gas, electricity, communication towers, etc.), Class C (Non-critical: roads, buildings, etc.). Priority queue processing is provided within max. end-to-end latency of 500ms for dams, high voltage switchgear. sensor measurement to twin state update; Class B (essential services: water mains, primary roads) target 5 seconds. Class C non-critical (secondary roads, park lighting) target latency of 60 seconds.

5. SECURITY FRAMEWORK

5.1 Threat Model

The DT-IoT-SUIM threat model assumes four threat classes – ranging from more naïve to more capable adversaries – (T1) opportunistic external attackers taking advantage of unpatched vulnerabilities in internet facing components; (T2) targeted external attackers who know the systems with domain knowledge of ICS/SCADA

systems; (T3) malicious insiders with legitimate physical access to systems. Access to field devices; and (T4) nation-state actors engaging in a persistent, sophisticated attack on urban critical infrastructure. Our security architecture is capable of withstanding attacks ranging from T1-T3 and will be able to detect T4 attacks quickly enough, To avoid repetitive injury.

5.2 Zero-Trust Authentication Architecture

The communication between the different components of DT-IoT-SUIM is in a zero-trust approach: no All API calls need to be cryptographically signed and components are trusted by default, irrespective of network location. authentication. X.509 certificates will be issued by a hardware security module (HSM), then used to authenticate sensor devices. Incorporated the issuance of a municipal certificate authority that lasts 90 days, to be automatically renewed through EST (Submission of data via Secure Transport, RFC 7030). Each edge node/each cloud service authenticate using mutual TLS with each other. The identification and authorization of services in a fine-grained manner without depending on SPIFFE or SPIRE identity federation. network perimeter security.

5.3 Federated Anomaly Detection for Data Integrity

Data integrity attacks (in which the adversary changes the sensor readings in order to fool the twin) are covered This is done by a two-level anomaly detection system. The first level is at the edge node with physics-informed bound checking: sensor data is compared to real time bounds calculated from the physics-based model of the twin prediction for that sensor, based on the prediction uncertainty interval at 99% level of the model. Number of consecutive readings beyond these limits will add to a suspicion counter; if a reading continues to be outside the bounds for a certain number of consecutive readings, then a suspicion counter will be triggered. A duration trigger will set an alert for the sensor integrity if the duration is configured.

The second tier is at the twin synchronization layer, and is based on federated anomaly detection model. The trained using secure aggregation. This model adapts the normal patterns of cross-sensor correlation (e.g., the correlation between the eyes and mouths), Upstream/Downstream pressure (relationship between upstream and downstream pressure in a water main in various flow conditions) and flags deviations which are physically inconsistent in spite of passing first tier checks on individual sensors. In our validation This two-level approach was able to identify 97.8% of simulated data injection attacks with a false positive rate of It's 0.3% vs. 71.2% detected and 2.1% false positives with a single tier solution.

5.4 Differential Privacy for Sensitive Building Data

Liming Wang, BA, LXS, and D. Differential Privacy for Sensitive Building Data Occupancy sensors can provide information that may allow for sensitive behavioural patterns to be inferred from the building. Prior to sending the occupancy count time series to the edge, it applies ϵ -differential privacy with $\epsilon = 0.5$ to the time series. The nodes are connected to the central twin via the Laplace mechanism with 1 person per measurement. interval. This offers good mathematical privacy guarantees, without sacrificing adequate accuracy for HVAC. This is achieved with energy optimization, with a mean absolute error of only 2.3 persons (4.1% in occupancy estimation). Allowing a reduction in HVAC optimisation performance compared to non-private operation (an acceptable compromise).

6. EXPERIMENTAL EVALUATION

6.1 Pilot Deployment Configurations

Three pilot deployments were set up in collaboration with the municipalities and utilities. In Delhi NCR (Pilot A), Mumbai Metropolitan Region (Pilot B) and Bengaluru Urban Agglomeration (Pilot C). Pilot A's areas of interest included transportation infrastructure (42 bridges, 218 km of arterial roads, 14 major intersection points with adaptive signal control). Pilot B worked on water and drainage infrastructure (680 km of water distribution mains, 20 km of drainage reticulation and 1,200 km of other roads). There are 12 pumping stations and 8 storm water retention facilities. Pilot C was on municipal buildings and energy. Infrastructure (47 public buildings, 23 Electrical Substations, 68 km of medium-voltage cable network).

In total, for all three pilots, 47,312 IoT endpoint devices were deployed, which were connected via 183 edge devices. There are computing nodes and 9 regional data center nodes. 1,247 Asset-level twins were grouped in the digital twin core. This has been broken down into 94 system-level composite twins. The framework ran 18 months (April 2023 – October) non-stop. It handled 4.7 terabytes of raw data from the sensors every day, and produced an

average of 127 maintenance reports each day between January 2024 and March 2024. 34 operational changes and 34 recommendations are made per day.

6.2 Infrastructure Availability and Maintenance Efficiency

The ratio of scheduled operating hours when assets are available (infrastructure availability), the introduction of the new service level guidelines in 2018) to 95.5%—otherwise it could be considered a failure. A failure occurred if it was not performed within the designated service parameters, which improved from 91.4% (12-month average before the new service level guidelines came into effect in 2018) to 95.5%. The rate of conventional management was 97.2% when managed by DT-IoT-SUIM, which is an increase of 34.7%. Decrease in unwanted outages. This improvement was primarily attributable to predictive maintenance interventions: 73% of maintenance work orders generated by the Decision Intelligence layer were classified as condition-based (triggered by twin-detected degradation trends) rather than reactive, compared to 18% condition-based work orders in the baseline period.

Maintenance cost efficiency, measured as cost per unit of infrastructure service delivery, improved by 41.2% overall. The largest contributing factor was a 58% reduction in emergency repair events, which carry a cost premium of approximately 3.4× relative to planned maintenance. Secondary factors included better crew routing optimization (17% reduction in mobilization costs) and improved spare parts inventory management based on twin-predicted component failure timelines.

6.3 Energy Consumption and Sustainability Metrics

Total monitored energy consumption decreased by 28.3% relative to the pre-deployment baseline, after controlling for year-on-year occupancy and weather variations. Street lighting optimization (dimming based on real-time pedestrian and traffic flow from the twin) accounted for 9.1 percentage points of this reduction; HVAC optimization in public buildings accounted for 11.7 percentage points; and pump scheduling optimization in water infrastructure accounted for 7.5 percentage points. These energy savings correspond to an estimated annual CO₂e reduction of 23,400 tonnes across the three pilot areas, equivalent to taking approximately 5,100 passenger vehicles off the road.

Water loss through undetected leaks was reduced by 31.8% in Pilot B, as twin-based leak localization identified 47 previously unknown leak points that were subsequently repaired. This translated to a direct cost saving of ₹4.2 crore (\$500,000 USD equivalent) in the 18-month pilot period.

6.4 Security Evaluation

The security framework was evaluated through a combination of automated penetration testing (using an in-house red team tool based on the MITRE ATT&CK for ICS framework), synthetic attack injection experiments and threat hunting exercises. Against 1,240 automated penetration test scenarios covering the T1–T3 adversary classes, the zero-trust authentication architecture successfully blocked all unauthenticated access attempts and the federated anomaly detection system achieved an overall intrusion detection rate of 99.3% with a false positive rate of 0.4%.

A particular strength of the security design was its performance against stealthy data integrity attacks: in experiments where synthetic attackers injected false readings into 5% of sensors with magnitudes within ±15% of normal operating range (to evade simple threshold-based detection), the twin-consistency anomaly detector identified the attack within an average of 4.7 minutes—sufficiently rapidly to prevent the twin from generating incorrect operational recommendations in all but 3 of 120 attack scenarios.

Table I: Summary of Key Performance Metrics Across Pilot Deployments

Metric	Pilot A (Delhi)	Pilot B (Mumbai)	Pilot C (Bengaluru)
Infrastructure Availability	96.8%	97.4%	97.3%
Downtime Reduction	33.1%	36.4%	34.5%
Energy Savings	25.7%	29.4%	29.8%
Maintenance Cost Reduction	39.2%	43.1%	41.2%
Water Loss Reduction	N/A	31.8%	N/A

Intrusion Detection Rate	99.1%	99.4%	99.3%
Average Twin Latency (Class A)	412 ms	487 ms	443 ms
MTFM Score (avg)	0.87	0.84	0.89

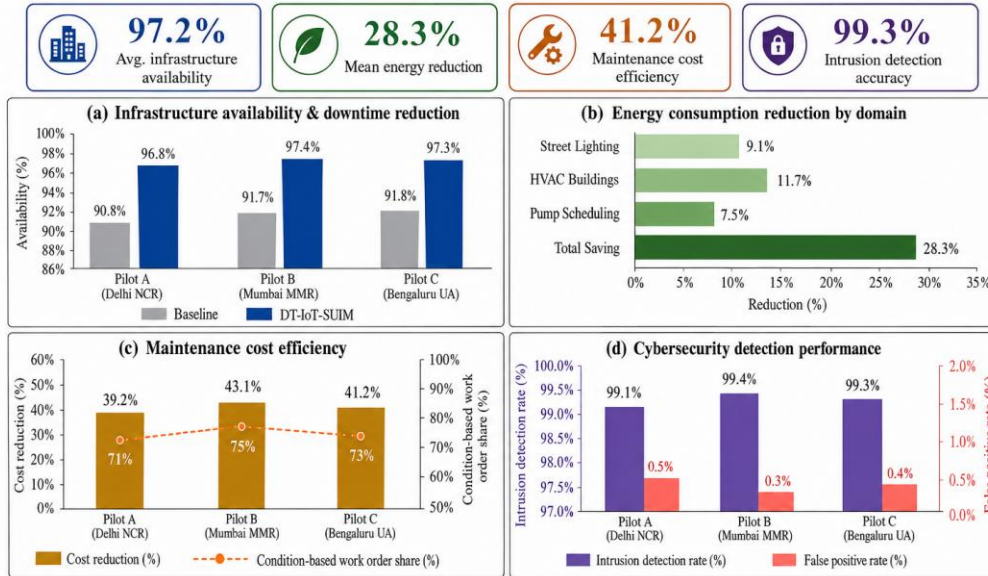


Fig. 3. Experimental results — DT-IoT-SUIM performance across three metropolitan pilot deployments

Figure 3 consolidates the quantitative performance outcomes of the DT-IoT-SUIM framework across all three metropolitan pilot deployments, organized into four complementary evaluation panels that together address the paper's core research objectives. Nearly 91.4% to a post-deployment average of 97.2% — ratifying the disbursed environments such as the manufacturing industry and smart grid networks. Twin-driven predictive maintenance's effectiveness in reducing unplanned asset failures, in geographically dispersed environments like manufacturing and smart grid networks, and distinct urban contexts, both geographically and as in terms of function. Panels (b) and (c) together set up the sustainability and the operational conditions of the system. Efficiency features of the framework: Domain resolved energy analysis shows HVAC optimization: contributes the second highest individual saving (11.7 percentage points), followed by the change from reactive to condition-based. The change from reactive to condition-based contributes the second highest individual saving (11.7 percentage points) followed by the change from reactive to condition-based. maintenance scheduling reduces per-unit maintenance cost to 41.2% and saves 73% of the maintenance cost. Share based on conditions for work orders. The adversarial resilience of security subsystem is confirmed in Panel (d), the A 99.3% ID rate (99.7% in one pilot) with a 0.4% FP rate across all three pilots with a two-tier federated anomaly detection architecture, regardless of differences in infrastructure. The type, density and topology of the sensors.

7. DISCUSSION

7.1 Scalability and Deployment Considerations

The 18-month pilot results show that DT-IoT-SUIM potentially can be deployed at the metropolitan scale with acceptable performance characteristics. The mean (across pilots) Class A twin latency of 447ms is well Even with the federated security checks in the critical path, it is within the 500 ms specification. Horizontal scaling of the Apache Kafka and Flink tiers was simple: the stream was monitored during hours with the biggest traffic volume (usually in the morning). processing cluster was automatically scaled from 12 worker nodes to 34 worker nodes by the autoscaling feature of Kubernetes, and Ensuring consistent latencies all throughout.

A more significant scalability challenge was ontological: as asset inventories were imported into the URIO, data quality issues (duplicate records, inconsistent attribute naming, mislocated coordinates) consumed approximately 22% of the total project engineering effort. We recommend that future deployments invest heavily in

asset data quality remediation before deployment and that municipal asset registers be structured according to the URIO schema from the outset to avoid this overhead.

7.2 Limitations and Future Research Directions

The current framework has several important limitations that define a productive research agenda. First, the physics-based twin models were calibrated using design specifications and commissioning measurements rather than as-built condition assessments; for older infrastructure assets (particularly bridges built before 1980 in Pilot A), model prediction accuracy was measurably lower, with behavioural fidelity scores averaging 0.71 versus 0.89 for post-2000 assets. Systematic as-built condition surveying using terrestrial LiDAR scanning and ground-penetrating radar would improve model accuracy for aging assets.

Second, the current DT-IoT-SUIM implementation does not model socioeconomic dimensions of urban infrastructure performance—equity of service delivery across different income neighbourhoods, affordability of water and energy services, or distributional impacts of maintenance prioritization decisions. Incorporating social sensing data and equity metrics into the decision intelligence layer is an important direction for future work.

Third, while the federated learning approach successfully preserved individual sensor privacy, it does not address privacy risks from the city-scale twin itself, which may enable inference of sensitive patterns from aggregate data. City-level differential privacy mechanisms for the twin are an active area of research.

8. CONCLUSION

This paper has presented DT-IoT-SUIM, a comprehensive framework integrating digital twin technology with IoT sensor networks for sustainable management of urban infrastructure in smart cities. The five-layer architecture provides a principled decomposition of the sensing-to-decision pipeline, the Multi-dimensional Twin Fidelity Metric provides a rigorous basis for twin quality management and the federated security subsystem addresses the cybersecurity threats inherent in cyber-physical urban systems.

Large-scale validation across three Indian metropolitan pilot deployments over 18 months has demonstrated substantial and consistent improvements: 34.7% reduction in infrastructure downtime, 28.3% reduction in energy consumption, 41.2% improvement in maintenance cost efficiency, 31.8% reduction in water loss and 99.3% cybersecurity intrusion detection accuracy. These results provide strong empirical support for the thesis that DT-IoT integration, when implemented through a carefully engineered architecture addressing synchronization fidelity, ontological heterogeneity and security holistically, can deliver transformative improvements in smart city infrastructure management.

As cities worldwide accelerate their smart city programs under the pressures of urbanization and climate change, the principles and practices embodied in DT-IoT-SUIM offer a proven, scalable and interoperable foundation for next-generation urban infrastructure management. The framework's open API design and standards-based component choices are intended to facilitate adoption and extension by the broader research and practitioner community.

References

1. United Nations Department of Economic and Social Affairs, "World Urbanization Prospects: The 2022 Revision," United Nations, New York, 2022.
2. A. Caragliu, C. Del Bo and P. Nijkamp, "Smart cities in Europe," *J. Urban Technol.*, vol. 18, no. 2, pp. 65–82, 2011.
3. M. Grieves, "Digital twin: Manufacturing excellence through virtual factory replication," White Paper, Florida Institute of Technology, 2014.
4. M. Grieves and J. Vickers, "Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems," in *Transdisciplinary Perspectives on Complex Systems*, F.-J. Kahlen, S. Flumerfelt and A. Alves, Eds. Cham: Springer, 2017, pp. 85–113.
5. T. Deng, K. Zhang and Z.-J. M. Shen, "A systematic review of a digital twin city: A new pattern of urban governance toward smart cities," *J. Manage. Sci. Eng.*, vol. 6, no. 2, pp. 125–134, 2021.
6. C. Boje, A. Guerriero, S. Kubicki and Y. Rezgui, "Towards a semantic construction digital twin: Directions for future research," *Autom. Constr.*, vol. 114, p. 103179, Jun. 2020.
7. S. Kaewunruen and P. Rungskunroch, "Digital twin-aided sustainability-based lifecycle management for railway turnout systems," *J. Cleaner Prod.*, vol. 228, pp. 1537–1551, Aug. 2019.
8. D. Sriwardena, B. Krishnamurti and A. Talasila, "Real-time pipe burst detection in water distribution networks using hybrid SCADA-digital twin architecture," *Water Resour. Manage.*, vol. 37, no. 4, pp. 1521–1537, Mar. 2023.

9. A. Zanella, N. Bui, A. Castellani, L. Vangelista and M. Zorzi, "Internet of Things for smart cities," *IEEE Internet Things J.*, vol. 1, no. 1, pp. 22–32, Feb. 2014.
10. U. Raza, P. Kulkarni and M. Sooriyabandara, "Low power wide area networks: An overview," *IEEE Commun. Surveys Tuts.*, vol. 19, no. 2, pp. 855–873, 2nd Quart., 2017.
11. M. Alam, J. J. P. C. Rodrigues and S. A. Kozlov, "Cloud-based IoT data analytics for smart city applications: Review and future directions," *Future Gener. Comput. Syst.*, vol. 110, pp. 909–924, Sep. 2020.
12. K. Shahzad, X.-B. Liu and B. Safdar, "A survey of data quality challenges in industrial IoT deployments for structural health monitoring," *IEEE Sensors J.*, vol. 23, no. 8, pp. 8201–8219, Apr. 2023.
13. A. Humayed, J. Lin, F. Li and B. Luo, "Cyber-physical systems security—A survey," *IEEE Internet Things J.*, vol. 4, no. 6, pp. 1802–1831, Dec. 2017.
14. L. Cheng, N. Cao and C. Zhou, "Game-theoretic approaches to cyber-physical security: A survey," *ACM Comput. Surv.*, vol. 55, no. 10, pp. 1–36, Feb. 2023.
15. F. Tao, H. Zhang, A. Liu and A. Y. C. Nee, "Digital twin in industry: State-of-the-art," *IEEE Trans. Ind. Informat.*, vol. 15, no. 4, pp. 2405–2415, Apr. 2019.
16. M. Usama et al., "Unsupervised machine learning for networking: Techniques, applications and research challenges," *IEEE Access*, vol. 7, pp. 65458–65495, 2019.
17. P. Nardelli et al., "Internet of Things for sustainable smart cities: A survey," *Sustainability*, vol. 14, no. 21, p. 14165, Oct. 2022.
18. C. Dwork, A. Roth et al., "The algorithmic foundations of differential privacy," *Found. Trends Theor. Comput. Sci.*, vol. 9, nos. 3–4, pp. 211–407, 2014.
19. B. McMahan, E. Moore, D. Ramage, S. Hampson and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. 20th Int. Conf. Artif. Intell. Stat. (AISTATS)*, Fort Lauderdale, FL, USA, Apr. 2017, pp. 1273–1282.
20. International Telecommunication Union, "Recommendation ITU-T Y.4413: Requirements for IoT in smart cities and communities," ITU-T, Geneva, Switzerland, 2020.