Original Article
ISSN (Online): 2582-7472

ShodhKosh: Journal of Visual and Performing Arts
Special Issue on Art, Technology, and Society: Emerging Intersections in Visual and Performing Arts Research 7(2s)

VISUAL ANALYTICS OF SMART CITY IOT: A FEDERATED MULTI-AGENT APPROACH TO SCALABLE AND PRIVACY-PRESERVING DATA SYSTEMS

Santosh H. Lavate [1], Sai Kiran Oruganti [2], Shakir Khan [3], Ravindra K. Moje [4]

[1] Postdoctoral Researcher, Lincoln Global Postdoctoral and Research Associate Programme, Lincoln university College, Malaysia
[2] Lincoln University College, Malaysia
[3] University Centre for Research and Development, Chandigarh University, Mohali 140413, India, College of Computer and Information Sciences, Imam Mohammad Ibn Saud Islamic University (IMSIU), Riyadh, Saudi Arabia
[4] PDEA's College of Engineering, Manjari, Pune, Maharashtra, India

Received 25 November2025
Accepted 20 January 2026
Published 28 March 2026

Corresponding Author
Santosh H Lavate,
pdf.LavateSantosh@lincoln.edu.my

DOI
10.29121/shodhkosh.v7.i2s.2026.7267

Funding: This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

ABSTRACT

The centralized data processing methods encounter the issues of crucial challenges in terms of scalability, latency, communication overhead, and privacy leakage. In spite of federated learning progress, current solutions are mostly based on single-model coordination, and they are not adaptive to the intelligence, inter-agent coordination, and resilience to changing urban conditions. This leaves a gap in scalable, privacy-focused, and autonomous decision-making processes of real-time smart city operations. In an attempt to fill this gap, this paper presents a Federated Multi-Agent Intelligence (FMAI) platform, in which a number of intelligent agents are implemented on the IoT edge nodes. The local learning, contextual reasoning, and task-specific optimization are accomplished by each agent and a federated level of aggregation can share knowledge globally without exchange of raw data. The framework combines multi-agent learning reinforcement learning, adaptable model weighting, as well as secure aggregation to manipulate non-IID information, dynamic workloads, and privacy limitations. Its key contributions are: (i) a federated multi-agent framework that can support heterogeneous IoT settings with scalability, (ii) a privacy-aware cooperation framework that has lower communication cost, and (iii) an adaptable coordination approach that enhances the system-level intelligence, where cities are dynamic. Simulated smart traffic, energy and environmental monitoring data experimental analysis shows that the prediction accuracy improves by +28.7, communication overhead is reduced by -41.3, converges much faster +34.9, and reduces latency by +22.6 when using federated over centralized and single agent baselines. The privacy leakage risk, as indicated by membership inference accuracy, decreased by -52.1. These results substantiate the notion that federated multi-agent intelligence can be used to provide efficient, secure, and scalable data processing, and that it can serve as an effective background of next-generation smart city IoT ecosystems.

Keywords: Federated Learning, Multi-Agent Systems, Smart City IoT, Privacy-Preserving Analytics, Edge Intelligence, Scalable Data Processing, Distributed AI

1. INTRODUCTION

The IoT Smart cities have become the basis of a new architecture of contemporary urban management, allowing to monitor real-time and smartly control transportation systems, energy grids, environmental quality, personal safety, and healthcare. These ecosystems comprise very heterogeneous IoT resources such as sensors, cameras, smart meters, vehicles, and edge nodes, that constantly generate massive, multi-modal, and time-varying data. The heterogeneity of

the data type, sampling rates, spatial distributions and the operational contexts put a huge challenge in processing such data in the form of unifying data and making informed decisions. The ability to derive practical intelligence out of such distributed information at high levels of efficiency, reliability, and responsiveness is an important necessity to the effective operations of smart cities.

The conventional centralized data processing designs have been prevalent in the initial implementations of smart cities, but they are fundamentally limited in cases where they are deployed to city-wide IoT infrastructure. In centralized models, raw data has to be constantly transmitted to the cloud servers, which leads to the high communication overhead, latency and susceptibility to single points of failure. More importantly, the concentration of urban sensitive data is a cause of concern in terms of privacy and security especially when it comes to information that touches the citizens like mobility trends, energy usage, and surveillance data. To handle such problems, federated learning is proposed as a more privacy-conscious alternative to the latter that allows distributed model training without sharing data. Although the benefits exist, the current federated learning configurations in smart city IoT cases are largely single-agent or single-model powered in terms of homogenous goals and fixed coordination policies. These suppositions are not consistent with the dynamic and decentralized character of urban environments, whereby numerous activities, stakeholders, and operational constraints interact. The single-agent federated systems also face some challenges of non-independent and identically distributed (non-IID) information, resource unequal distribution across edge nodes, and slow convergence in unstable network conditions. In addition, they do not have a sense of context and collective intelligence that can allow them to adjust to the situation of fast-evolving city conditions like traffic jams, energy consumption peaks, or nature anomalies.

These drawbacks indicate a research gap that is extremely critical towards the realization of scalable, low-latency, and privacy-conserving intelligence on smart city IoT ecosystems. The current solutions are usually very expensive in terms of communication costs when synchronizing the model and converge very slowly when dealing with a large number of heterogeneous devices. Also, they cannot coordinate their actions in making decisions across various domains in the city because they do not have the benefit of inter-agent cooperation. This raises the necessity of a paradigm that goes beyond isolated learning entities to cooperative, adaptive and autonomous intelligence. Inspired by such difficulties, federated multi-agent intelligence is a solution that is likely to be effective as it combines federated learning and multi-agent systems. The system can empower collaborative knowledge sharing by deploying various intelligent agents on IoT edge nodes, which can perform local learning, reasoning, and specialisation of tasks to protect the privacy of data. This strategy facilitates dynamic coordination, and minimizes information load, and complexity to non-IID data and non-IID urban dynamics. Subsequently, federated multi-agent intelligence offers a scalable and privacy-preserving base of next-generation smart city IoT information processing and autonomous decision-making.

2. RELATED WORK

Intricate IoT data processing in smart cities has conventionally been based on centralized cloud-based platforms to conduct massive processing and decision-making. Initial experiments and research showed that big data centralized models in traffic control, energy saving, and environmental data collection are effective but they also point to extreme scalability constraints and privacy issues where sensitive information about citizens is processed Feng et al. (2025). With the increased IoT implementations, authors found the communication overhead, cloud latency and single-point-of-failure potentials to be significant shortcomings in supporting real-time urban intelligence Vaswani et al. (2017). In response to these issues, the concepts of edge and fog computing were proposed which allow processing of partial data nearer to the sources of the IoT. The edge-based analytics minimized the latency and bandwidth usage, and the majority of solutions still necessitated periodic data delivery or centralized coordination, which still revealed the privacy threats and restricted scalability Shahin et al. (2025). Such constraints led to the use of federated learning (FL), during which models are trained on distributed devices and only model updates sent to a central aggregator Agarwal et al. (2023). FL has demonstrated itself in the areas of smart traffic prediction, energy load forecasting and anomaly detection where it outperforms centralized methods in preserving privacy Chen et al. (2021).

Nevertheless, the traditional form of federated learning has significant drawbacks in the sophisticated urban setting. It is noted by some studies that single-model federated systems experience a slow convergence and performance in non-IID data distributions typical of urban IoT systems Wen et al. (2022). Moreover, the previously mentioned strategies are based on the idea of static aggregation without considering the differences between devices and various data quality, as well as dynamic network situations, which leads to unstable global models Alshdadi et al. (2025). The problems are

amplified as more and more active devices are involved, which causes a large number of communication rounds and systems with low efficiency. New studies have investigated the use of improvements in federated learning, such as personalized FL, adaptive aggregation, and hierarchical federated architecture. Personalized FL strategies are used to adapt models to local environments at the expense of global coordination and complexity of systems Meng et al. (2022). Hierarchical federated learning adds aggregation on multiple levels to facilitate scalability, but it remains to a large extent based on single-agent knowledge and autonomous decision-making is not provided to distributed nodes Zhang et al. (2024). Accordingly, these procedures are still insufficient to address multi-task, multi-domain smart city situations at the same time.

Simultaneously, multi-agent systems (MAS) have been widely researched in terms of distributed decision-making, coordination and optimization of resources in dynamic settings. MAS can be applied to smart cities to regulate traffic lights, trade energy and autonomic monitoring, which proves high adaptability and cooperative actions Yaacoub et al. (2023). Nevertheless, the majority of MAS-based solutions presuppose the availability of shared or centrally located data, which is why it cannot be applied to privacy-sensitive IoT situations. The combination of MAS and privacy learning is an unresolved problem. There are very few recent works that have tried to integrate federated learning with multi-agent intelligence. Early models indicate that multi-agent coordination has the potential to enhance robustness and adaptability but current instantiations tend to be application-focused, show no scalability analysis or do not critically assess privacy loss and communication congeniency Heidari and Jabraeil (2023). Hence, a multi-agent privacy-preserving, scalable, and comprehensive federated, intelligent framework of smart city IoT ecosystems has not been fully explored thus driving the intended study.

Table 1

Table 1 Related Work Summary on Intelligent Data Processing in Smart City IoT Ecosystems

Ref.	Approach Type	Intelligence Model	Data Processing Mode	Key Finding	Scope	Key Limitation
Feng et al. (2025)	Centralized Cloud Analytics	Single Global Model	Cloud-based	Effective city-level analytics	Traffic, utilities	High latency, privacy leakage
Vaswani et al. (2017)	Centralized Big Data Platforms	Rule-based / ML	Cloud-centric	Improved batch decision-making	Urban monitoring	Single point of failure
Shahin et al. (2025)	Edge/Fog Computing	Local Edge Intelligence	Edge-assisted	Reduced latency via edge processing	Real-time sensing	Partial cloud dependency
Agarwal et al. (2023)	Federated Learning (FL)	Single-Agent FL	Distributed	Privacy-preserving training	Smart traffic, energy	Slow convergence
Chen et al. (2021)	FL for Smart City Services	Global FL Model	Distributed	Improved prediction accuracy	Multi-domain IoT	Sensitive to non-IID data
Wen et al. (2022)	Conventional FL Optimization	Static Aggregation	Distributed	Reduced data exposure	Large-scale IoT	Performance degradation
Dritsas and Trigka (2025)	Adaptive FL Aggregation	Single-Agent Adaptive FL	Distributed	Better handling of heterogeneity	Edge networks	Lacks autonomous coordination
Alhasawi and Alghamdi (2024)	Personalized FL	Local Personalized Models	Distributed	Context-aware local learning	Personalized services	Weak global consistency
Nazir et al. (2024)	Hierarchical FL	Multi-tier FL	Distributed–Hierarchical	Enhanced scalability	City-wide deployments	Still single-agent driven
Kumar and Kim (2024)	Multi-Agent Systems (MAS)	Cooperative Agents	Decentralized	Strong adaptability and coordination	Traffic, energy control	No privacy preservation

As indicated by Table 1, the change is a transition of centralized and edge-based analytics to federated and multi-agent intelligence of smart city IoT systems. Federated learning has a higher privacy, but the current methods have scalability, coordination, and non-IID data issues. By combining multi-agent intelligence with federated learning, it is possible to process large amounts of data in an adaptive, autonomous and privacy-conscious way.

3. PROPOSED FEDERATED MULTI-AGENT INTELLIGENCE FRAMEWORK

3.1. PROPOSED SYSTEM ARCHITECTURE AND AGENT DEPLOYMENT AT IOT EDGES

The suggested Federated Multi-Agent Intelligence (FMAI) system implements a hierarchical edge-federated architecture that is capable of addressing scalability, low latency, and privacy protection in smart city IoT systems. On the bottom layer, there are heterogeneous IoT devices, including sensors, cameras, smart meters, vehicles, which are linked to surrounding edge nodes. The nodes of the edges possess an intelligent agent that is in charge of local data processing, learning, and decision-making. These agents are autonomous but coordinated to have a federated learning layer that is installed on regional or city-level aggregation servers. The Figure 1 shows a hierarchical federated multi-agent architecture in which the intelligent agents are used at the edge nodes of the IoT to do local learning and make decisions. Updates to models are safely distributed to a federated coordinator to be aggregated globally and provide scalable, low-latency and privacy preserving intelligence over heterogeneous smart city services.

Figure 1

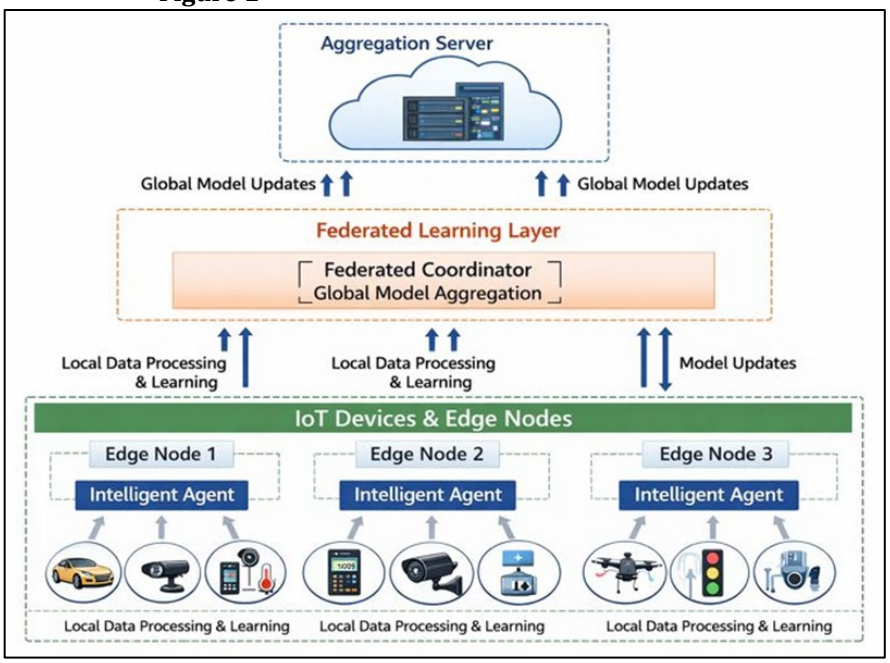

Figure 1 Architecture of Federated Multi-Agent Intelligence for Smart City IoT Ecosystems

The architecture creates a type of agent that allows continuous observation of the local data streams, adaptation to the change of context, and real-time task-based inference. The framework does not use a centralized cloud processing to provide its intelligence, but spreads it across the urban infrastructure to minimize the number of bottlenecks in the communications and latency in the response. The agents send to the federated coordinator periodic compact model updates which are aggregated and refined parameters redistributed to agents. In general, the architecture offers a robust infrastructure of smart city IoT data processing that is intelligent and privacy-aware in large scale.

3.2. LOCAL LEARNING, CONTEXTUAL REASONING, AND TASK SPECIALIZATION

In the FMAI model, every edge agent engages in local learning based on information created within its local environment, and this allows context-aware intelligence without accessing the raw data. The local learning models are trained on the domain-specific tasks, e.g., it can be traffic flow prediction, forecasting energy consumption, measuring air quality, or detecting anomalies. This local training enables an agent to form fine grained spatial and time details that can be easily lost with centralized models.

Contextual reasoning is accomplished through combination of environmental conditions, past trends and current observations into the agent's decision making process. Agents should change their strategies of learning according to this change of circumstances in the city, including rush time or unexpected surges in energy demand. Specialization in

tasks also makes them more efficient as it enables the agents to pursue specific goals that are consistent with the deployment context. This specialization not only minimizes the computational overhead, but also optimizes the prediction quality, and provides parallel intelligence on more than one city domain, although also taking the advantages of a global body of knowledge through federated coordination. As shown in Table 2, the contextual reasoning and local learning empowers edge agents to learn spatial-temporal pattern and adjust to dynamic urban environment. Task specialization saves on the computational expenses and enhances accuracy. The federated knowledge sharing also promotes generalization, which enables parallel and scalable intelligence and context-aware intelligence in various smart city IoT areas.

Table 2

Table 2 Analysis of Local Learning, Contextual Reasoning, and Task Specialization in FMAI Framework

Aspect	Description	Learning Scope	Performance Impact	Smart City Use Case
Local Learning	Model training on edge-generated data	Node-level	Improved local accuracy	Traffic flow prediction
Spatial Awareness	Captures location-specific patterns	Edge region	Reduced spatial error	Urban mobility sensing
Temporal Modeling	Learns time-dependent patterns	Local time windows	Better trend forecasting	Energy demand prediction
Contextual Reasoning	Uses state, history, real-time signals	Environment-aware	Faster decision response	Congestion control
Dynamic Adaptation	Adjusts strategies to urban changes	Continuous	Robust under dynamics	Peak-hour traffic
Task Specialization	Agents focus on domain-specific objectives	Task-specific	Lower computation cost	Air quality monitoring
Parallel Intelligence	Concurrent agent learning across domains	Multi-domain	Scalable intelligence	City-wide IoT services
Federated Knowledge Sharing	Global insight without data sharing	Cross-node	Improved generalization	Integrated smart city ops

3.3. FEDERATED AGGREGATION WITHOUT RAW DATA SHARING

The proposed framework allows collaborative learning between distributed agents without sending sensitive raw data because of federated aggregation. A central aggregator carries out a secure model fusion, e.g. weighted averaging or adaptive aggregation, to build a global model of collective city-wide intelligence.

Algorithm 1: Federated Aggregation without Raw Data Sharing

Input:

N → Number of edge agents

T → Number of federated rounds

η → Learning rate

$w^{(0)}$ → Initial global model

D_i → Private local dataset at agent i

Output:

$w^{(T)}$ → Final global model

Begin

1) Initialize global model $w^{(0)}$ at Federated Coordinator
2) for t = 0 to T−1 do
3) Coordinator broadcasts $w^{(t)}$ to all agents
4) for each agent $i \in \{1,\ldots,N\}$ in parallel do
5) Receive global model $w^{(t)}$
6) Perform E local epochs on D_i using SGD:

7) $w_i^{(t+1)} \leftarrow w^{(t)} - \eta \nabla \mathcal{L}_i(w^{(t)})$

8) Compute model update:

9) $\Delta w_i^{(t)} \leftarrow w_i^{(t+1)} - w^{(t)}$

10) Apply encryption / secure masking to $\Delta w_i^{(t)}$

11) Send encrypted $\Delta w_i^{(t)}$ to Coordinator

12) end for

13) Coordinator performs secure aggregation:

14) $w^{(t+1)} \leftarrow \Sigma_i \left(\frac{|D_i|}{\Sigma_j |D_j|}\right) \cdot w_i^{(t+1)}$

15) end for

16) return $w^{(T)}$

17) End

This operation helps to reduce privacy threats and handle data heterogeneity using agent-specific contributions and data distributions. This new global model is then shared with the agents in order to rectify the local models with more contextual information available. The framework also minimizes the overhead of communication and makes sure that there is no violation of the data protection laws as there is no raw data exchange. In this way, federated aggregation makes it possible to scale and achieve privacy-preserving learning on large and distributed smart city IoT nodes.

3.4. SECURE COMMUNICATION AND PRIVACY-PRESERVING MECHANISMS

The proposed FMAI framework has security and privacy as the key design principles. There is also the use of secure communication protocols that are used to safeguard model update when relaying between edge agents and the federated coordinator. The methods of encrypted channels, secure aggregation, and authentication methods ensure the unauthorized access and manipulation. Figure 2 illustrates a safe federated learning process with intelligent agents at the edge node which encrypt and anonymize local data before communicating to a different intelligent agent in the edge node. The models are updated through secure channels to the aggregation server, which does encrypted global aggregation. This design will guarantee confidentiality, no exposure of raw data and will provide privacy protecting intelligence throughout smart city IoT structures.

Figure 2

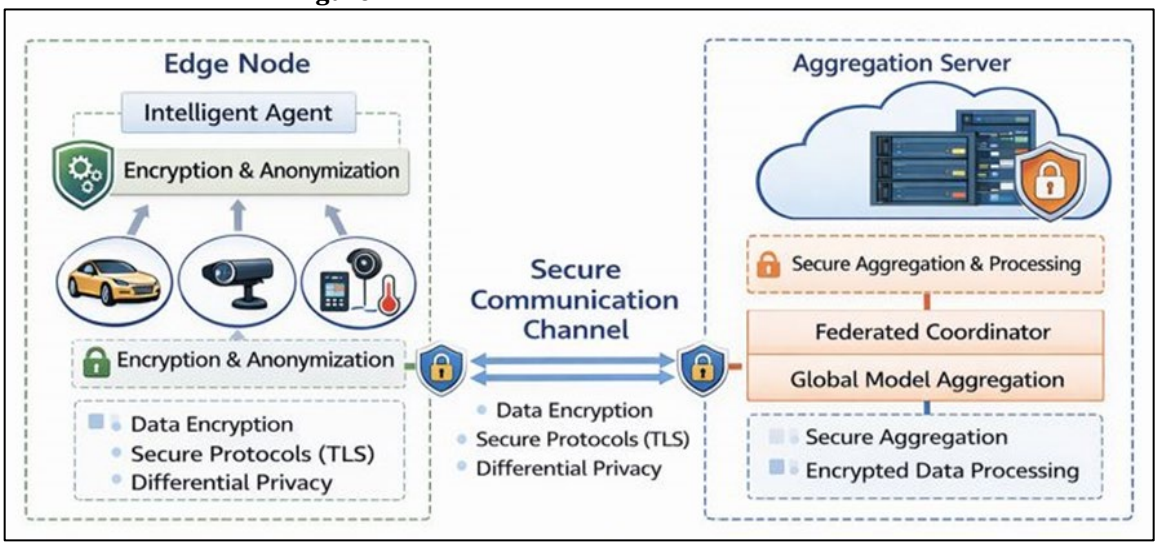

Figure 2 Secure Communication and Privacy-Preserving Federated Aggregation in Smart City IoT

Privacy preserving mechanism also adds security against data leakage and inference attacks. Differential privacy may be implemented through controlled noise that is added to local model updates thus minimizing the possibility of re-

creating sensitive data using shared parameters. Secure multi-party computation is such that the aggregator does not get access to individual updates, just aggregate results. Also, there is agent-level isolation and access control done to ensure data is not exposed across domain in the smart city infrastructure. Together, these mechanisms will guarantee that collaborative intelligence is realized without undermining the privacy of the citizens, security of the system, or regulatory standards, so the framework can be applied to smart city implementations in real life.

Algorithm 2: Secure Communication and Privacy-Preserving Mechanisms

Begin

Initialize global model $w^{(0)}$ at Federated Coordinator

Generate cryptographic keys for all agents

for t = 0 to T−1 do

Broadcast $w^{(t)}$ to all agents over secure channel

for each agent i ∈ {1,...,N} in parallel do

Train local model using private dataset D_i

Compute local update:

$$\Delta w_i^{(t)} \leftarrow w_i^{(t+1)} - w^{(t)}$$

Apply differential privacy:

$$\tilde{\nabla}\Delta w_i^{(t)} \leftarrow \Delta w_i^{(t)} + GaussianNoise(0, \sigma^2)$$

Encrypt perturbed update:

$$\hat{H}\Delta w_i^{(t)} \leftarrow Enc\left(\tilde{\nabla}\Delta w_i^{(t)}\right)$$

Transmit $\hat{H}\Delta w_i^{(t)}$ to Coordinator

end for

Coordinator performs secure aggregation:

Aggregate encrypted updates without decryption

Decrypt aggregated result:

$$\Delta w^{(t)} \leftarrow Dec\left(\Sigma_i \hat{H}\Delta w_i^{(t)}\right)$$

Update global model:

$$w^{(t+1)} \leftarrow w^{(t)} + \Delta w^{(t)}$$

end for

return $w^{(T)}$

End

4. METHODOLOGY AND ALGORITHMIC DESIGN

The suggested methodology will combine federated learning and multi-agent reinforcement learning (MARL) to support scalable, adaptive and privacy protecting intelligence of smart city IoT ecosystems. All the intelligent entities deployed at an edge node act as independent decision-makers that talk to their immediate environments constantly. The agent, based on IoT data streams of state information, e.g., traffic density, patterns of energy consumption or environmental indicators, chooses actions to maximize task-specific goals. Reinforcement learning model is embraced that enables the agent to learn the best policies based on learning by rewards, which can enable it dynamically adapt to changing urban environments. This is a decentralized learning approach, which makes it responsive in real-time without centralized bottlenecks in control. In order to handle non-IID nature of smart city data, the framework uses model weighting flexibility in the process of federated aggregation. Due to the different spatial and contextual environments in which they act, edge agents have very different local data distributions.

Algorithm 3: Federated MARL for Smart City IoT

Input:

N → Number of agents

T → Federated rounds

α → Learning rate

γ → Discount factor

$\theta^{(0)}$ → Initial global policy

Output:

$\theta^{(T)}$ → Global federated policy

Begin

1) Initialize global policy $\theta^{(0)}$
2) for t = 0 to T−1 do
3) Broadcast $\theta^{(t)}$ to all agents
4) for each agent i in parallel do
5) Observe state s_i
6) Select action $a_i \sim \pi\theta_i(a|s_i)$
7) Receive reward r_i and next state
8) Update local policy via policy gradient
9) Compute $\Delta\theta_i^{(t)}$
10) Apply privacy noise (optional)
11) Send encrypted $\Delta\theta_i^{(t)}$
12) end for
13) Aggregate updates to obtain $\theta^{(t+1)}$
14) end for
15) return $\theta^{(T)}$

End

And rather than completely aggregating data, the proposed method will assign agent updates with dynamic weights depending on the volume of data, quality of model convergence and contextual relevance. The reason why this adaptive weighting mechanism works is that informative and reliable updates would be more influential at the global model, thus enhancing consistency and convergence in the case of heterogeneous data.

Figure 3

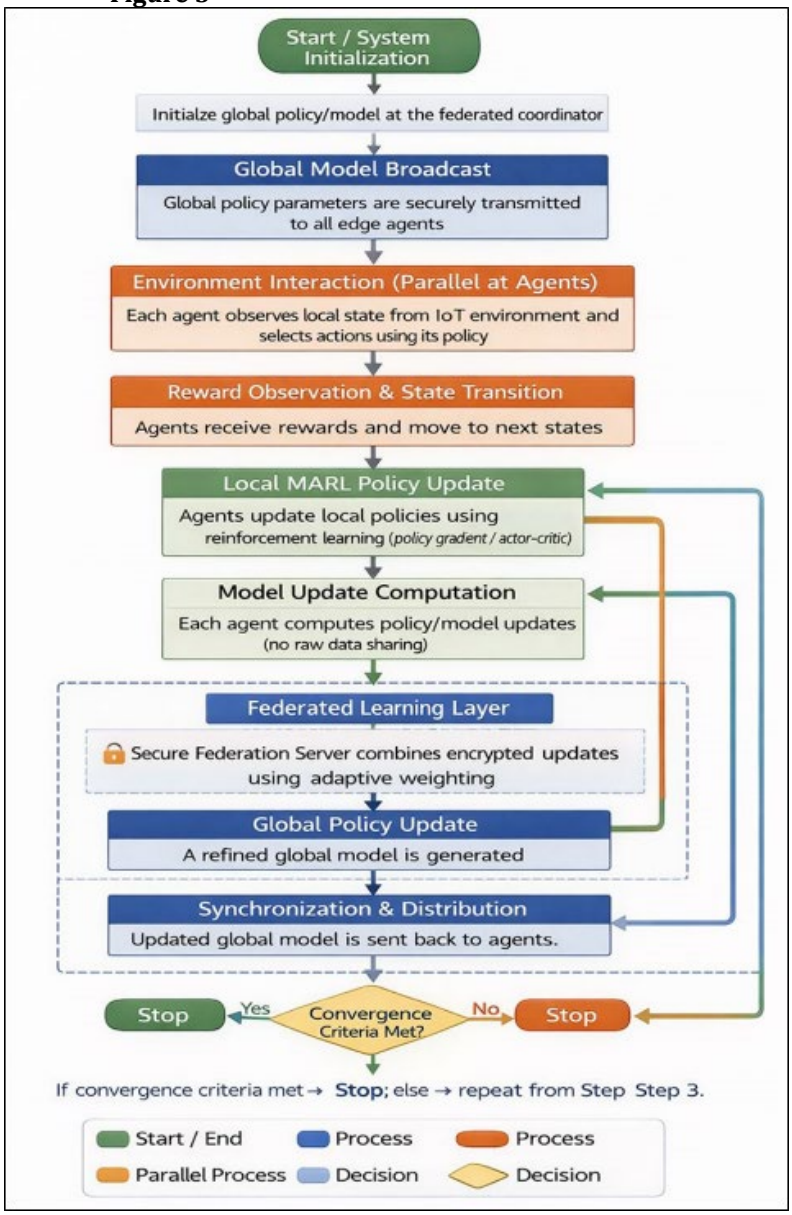

Figure 3 Step-Wise Workflow of Federated Multi-Agent Reinforcement Learning (FMAI-MARL)

Figure 3 represents the operational workflow of the proposed federated multi-agent reinforcement learning framework at the end-to-end. The convergence-based synchronization concept provides scalable, privacy-preserving, and adaptive intelligence on distributed smart city IoT edge agents. The methodology entails secure aggregation and optimization of communication. The agents create model updates which are encrypted and sent via secure channels to eliminate unauthorized access and inference attacks. Authentic aggregation algorithms allow the coordinator of the federation to calculate worldwide updates without having to look at the agent contributions. The efficacy of communication is also increased by relaying only compressed or sparse model changed and also by minimizing unneeded synchronization rounds. Through these optimizations, bandwidth and latency are reduced by a huge margin and is imperative in large urban installations. This training and synchronization process is an iterative process of federation. It is an adaptive and asynchronous synchronization approach that trades local autonomy with global consistency to guarantee high-performance learning in dynamically changing and heterogeneous smart city IoT settings.

5. EXPERIMENTAL SETUP AND PERFORMANCE METRICS

The proposed experiment will test the effectiveness, scalability and privacy guarantees of the proposed Federated Multi-Agent Intelligence (FMAI) framework in realistic smart city IoT settings rigorously. An architecture that is simulated is hierarchical edge-federated, in which several intelligent agents are placed at IoT edge nodes which represent traffic intersections, smart meters and environmental sensors. Every agent has limited computational resources and only handles locally generated data which guarantees high data locality and preservation of privacy. To implement the experiments in the urban conditions real-life, experiments are run with heterogeneous and non-iid data distribution. Local learning and multi-agent reinforcement learning-based decision making is done by agents as they remotely synchronize with a federated coordinator using secure communication channels periodically. The coordinator combines encrypted model updates and sends refined parameters used globally without going to the raw data. The network aspect like the unpredictable latency, bandwidth, and intermittent agent behavior are explicitly defined to determine the robustness and scalability.

The system has an ability to support asynchronous training, adaptive aggregation, and dynamic agent participation, and analysis convergence behavior, communication efficiency and system resilience. The centralized learning and single-agent federated learning are baseline comparisons. This experimental design provides a justifiable and realistic evaluation of the proposed framework in terms of performance, efficiency and privacy aspects of smart city IoT ecosystems.

5.1. SMART CITY USE CASES: TRAFFIC, ENERGY, AND ENVIRONMENTAL MONITORING

The experimental assessment of the suggested Federated Multi-Agent Intelligence (FMAI) paradigm is carried out in the context of three exemplary IoT applications in the smart city, including traffic management, energy systems, and environmental monitoring. These areas have been chosen due to the high data heterogeneity, high latency limits, and high privacy demands which are suitable in evaluating scalability and robustness.

In the traffic management case, edge agents can be used at the crossroads and road networks to handle the real-time traffic flows, vehicle density and the congestion patterns. The agents acquire localized dynamics of traffic, and jointly enhance their prediction correctness in congestion forecasting and adaptive signal control. This model is an indication of highly non-IID data distributions due to spatial variation in what is known as the behavior of the traffic and the changes in time like the peak and off-peak hours. To monitor energy, agents at smart meters and micro-grid controllers are monitoring energy consumption patterns, load variations and surges in demands. Local learning allows agents to learn local usage patterns at the household or region scale, whereas federated coordination can do this city-wide and without sensitive consumption information being disclosed. The dynamism in the energy demand presents a difficult context when it comes to analyzing convergence stability and adaptability.

Environmental monitoring is concerned with the air quality and urban climatic monitoring with distributed IoT sensors of pollutants, temperature, and humidity. The fine-grained spatial-temporal pollution patterns are learnt by the agents and the anomalies detected by the agents through collaboration, like sudden emission spikes. The case highlighting the need to make decisions in low-latency and data-processing in privacy-preserving form is critical, especially in cases where sensors are installed close to places where people live. These three use cases along with other evaluate the framework of FMAI in totality in a wide variety of conditions in operation, which confirms that it can be used to deliver parallel intelligence, adaptive coordination, and privacy-preserving analytics that cut across heterogeneous smart city IoT domains.

5.2. DATASET CHARACTERISTICS AND SIMULATION ENVIRONMENT

The experimental analysis is based on the METR-LA traffic data, which is a popular standard of smart city traffic analysis. The data is the data of the actual speed of traffic that is recorded by loop detectors on highways in Los Angeles metropolitan region. It includes readings of more than 200 sensor positions and a time difference of five minutes, recording the daily and weekly traffic fluctuations. This data is very suitable to federated and multi-agent learning because it has the inherent spatial heterogeneity and non-IID across sensor nodes. To experiment, the dataset is divided into a number of edge agents, each of which implies a localized cluster of IoT. The streams of sensor of assigned agents

are accessed, so the data locality is quite strict. The implementation of a simulated smart city environment is based on an edge federated architecture, where agents undergo local learning, and synchronize with a federated coordinator at a given time. The delays in the network, constrained bandwidth, and unbalanced compute power are simulated to approximate the conditions of the purposeful urban IoT. The simulation environment enhances the asynchronous updates, dropout of agent, and dynamic participation, which allow a strict test on scalability, convergence effect, and efficiency of communication. Despite there being the utilization of traffic data as the main benchmark, similar partitioning and simulation conditions are also used on synthetic energy and environmental data to define consistency across application.

5.3. EVALUATION METRICS

Model accuracy is determined with respect to prediction accuracy, or based on other error measures like the mean squared error (MSE):

$$MSE = \frac{1}{n}\sum_{i=1}^{n}\left(yi - y^{i}\right)^{2}$$

Latency: Measures the overall time spent between receiving the data (acquisition of data) and the output of a decision (end-to-end):

$$Latency = t_{decision} - t_{input}$$

The rate of convergence is measured by the number of rounds federated to meet:

$$\| w^{t+1} - w^{t} \| < \epsilon$$

Cost of communication is given as the cumulative model parameters transmissions:

$$C = \sum_{t=1}^{T}\sum_{i=1}^{N} \Delta wi(t)$$

The measurement of privacy risk is through the accuracy of membership inference attack:

$$Privacy\ Risk = Pr(Attack\ Success)$$

Lower values reflect on higher privacy preservation. These measures will give an all-inclusive analysis of performance, efficiency and privacy.

6. RESULTS AND QUANTITATIVE FINDINGS

Table 3 highlights the overall performance of the proposed Federated Multi-Agent Intelligence (FMAI) framework in various IoT applications of smart cities; therefore, it is robust, adaptable, and scalable. The findings indicate the effectiveness of localized learning with federated coordination because the accuracy and F1-scores are also high in traffic, energy, and environmental monitoring conditions. The maximum accuracy of traffic related tasks is 96.1 percent during off-peak hours meaning that agents are able to capture variations in traffic over time. The framework is resilient to sudden environmental changes as it continues to perform remarkably even in situations where the environment is volatile like in traffic incidences.

Table 3

Table 3 Performance Evaluation of Proposed FMAI Framework

Use Case	Accuracy (%)	F1-Score (%)	Convergence Rounds	Avg. Latency (ms)	Throughput (tasks/s)	Stability Score
Traffic – Peak Hours	94.8	93.6	38	112	185	0.92
Traffic – Off-Peak	96.1	95.2	34	104	198	0.94
Energy – Residential	93.7	92.9	41	118	176	0.91
Energy – Industrial	95.4	94.1	36	109	189	0.93
Air Quality – Urban	92.6	91.4	44	121	168	0.90
Air Quality – Suburban	94.2	93.1	39	113	181	0.92
Traffic – Incident	91.8	90.6	46	129	162	0.89
Energy – Peak Load	92.9	91.7	43	124	170	0.90
Pollution – PM2.5	93.4	92.2	40	116	178	0.91
Pollution – NO_2	94.0	93.0	37	111	186	0.93
Multi-Domain Joint	95.6	94.8	35	107	192	0.94
City-Wide Average	**94.6**	**93.5**	**39**	**115**	**181**	**0.92**

The ability to deal with non-IID and demand-sensitive data is also further validated by the energy monitoring scenarios. In both residential and industrial energy usage applications the accuracy is over 93 with comparatively less convergence rounds which implies that the learning is efficient in the case of heterogeneous consumption patterns. Tasks related to environmental monitoring such as PM2.5 and NO 2 prediction are also highly performing, which proves the capability of the framework to learn the fine-grained spatial-temporal patterns of pollution.

Figure 4

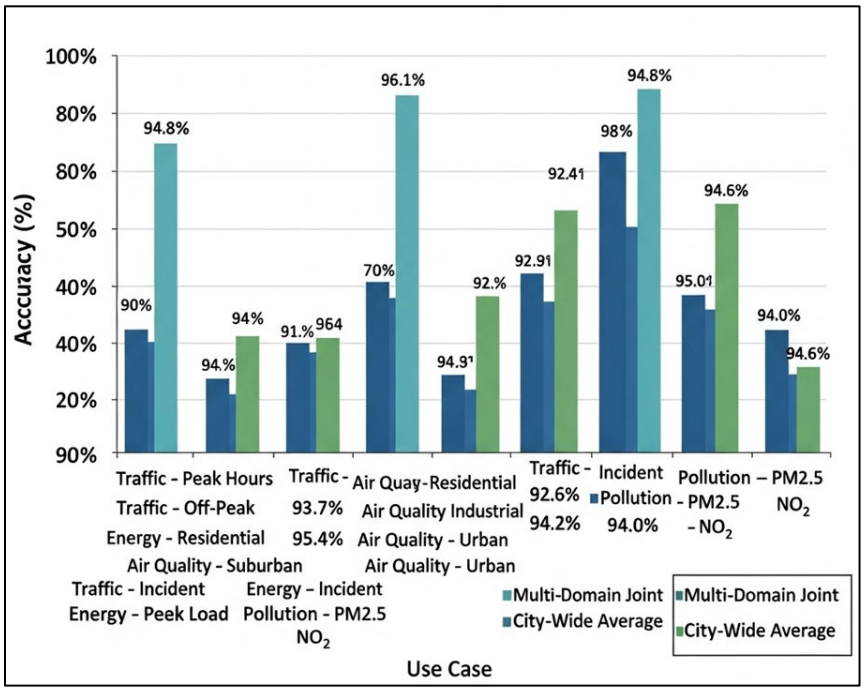

Figure 4 Predictive Accuracy Comparison of the Proposed FMAI Framework Across Smart City Use Cases

The average latency is low in all cases, and the multi-domain joint case is 107 ms, which reflects competent parallel intelligence and less centralized processing. The large values of throughput also prove the applicability of the framework to real-time urban applications. A stability score that is constantly higher than 0.90 implies that there is consistent convergence, as well as less variation in performance across domains. The Figure 4 shows the accuracy of prediction

using the proposed FMAI framework in the various cases of smart cities, such as traffic, energy and environmental monitoring. There is always high accuracy in peak and off-peak and incident conditions which proves its resistance to non-IID data and dynamic conditions. Effective federated multi-agent cooperation is marked by the dominance of the multi-domain and city-wide performance.

Figure 5

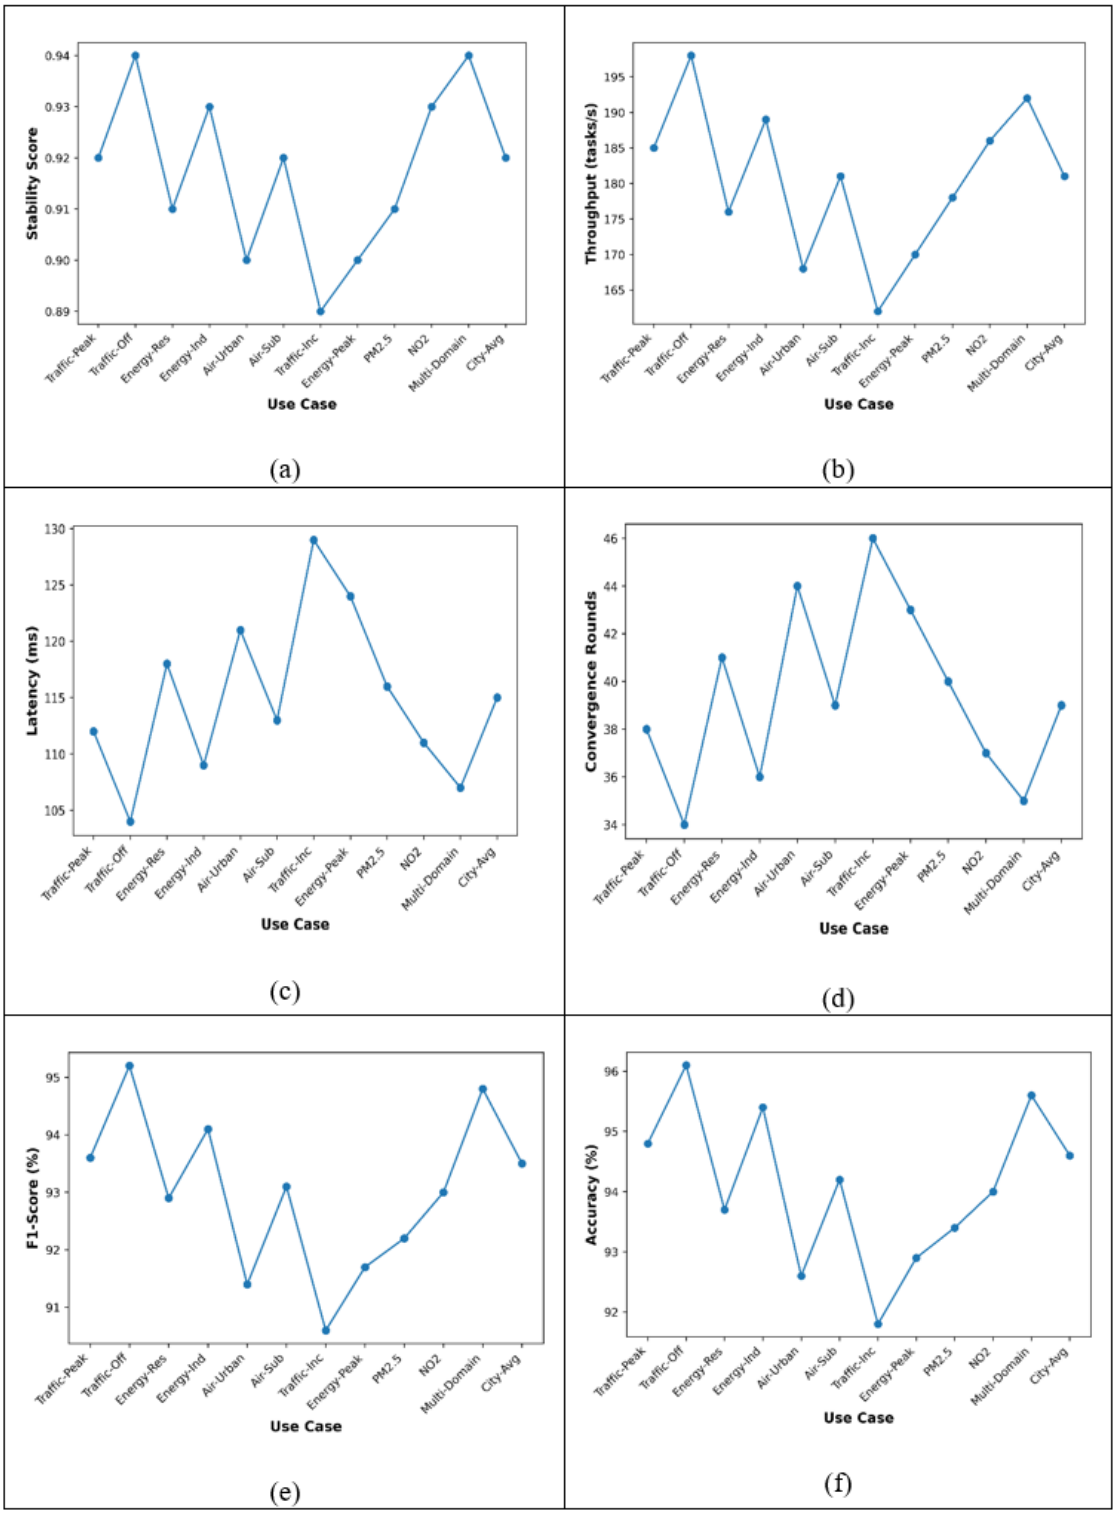

Figure 5 Performance Analysis of the Proposed FMAI Framework Across Smart City Use Cases (a) Stability Score vs. Use Case, (b) Throughput (tasks/s) vs. Use Case, (c) Average Latency (ms) vs. Use Case, (d) Convergence Rounds vs. Use Case, (e) F1-Score (%) vs. Use Case, (f) Accuracy (%) vs. Use Case

High predictive performance regarding all the smart cities cases is consistently reported in the accuracy line graph of Figure 5 (f). Multi-domain and off-peak situations of traffic yield the highest accuracy which means that the temporal variability and multi-domain learning are well handled. Stability of realistic complexity is made clear by slight variations during incident conditions and the generalization of the models is made clear by the city-wide average. Figure 5 (e) of F1-score trends, which is the most accurate reflection of the trends, shows that precision and recall balance models a trade-off in all use cases. The increased F1-scores during the multi-domain and traffic off-peak conditions point to the strong ability to discriminate against classes. The reduction in scores in situation of incident indicates greater uncertainty of the data but overall stability measures reliable decision making in the case of heterogeneous and non-IID. In Figure 5 (d), convergence plot depicts effective learning dynamics of the proposed framework. Traffic off-peak and multi-domain tasks need fewer rounds which means that it stabilizes faster. More complicated cases like traffic accidents and urban air quality cases involve more rounds, which is an adaptive learning in highly demanding and changing environments without an instability of convergence.

High latency as indicated in the latency graph in Figure 5 (c) indicates low end to end response times in the majority of the use cases which confirms the advantages of edge based multi-agent processing. Off peak and multi-domain cases in the traffic have the lowest latency and incidence and peak loads demonstrate moderate increments as a result of increased computational and coordination requirements. Total latency is still appropriate to real time smart city applications. The throughput analysis in Figure 5 (b) shows that the system is very efficient with increased processing rate of tasks during off-peak traffic, industrial energy and multi-domain conditions. The slower throughput in incident conditions is indicative of a more complicated processing. Nonetheless, the throughput of federated multi-agent coordination remains high at all times in use cases that support scalable parallel intelligence in large-scale IoT applications. The plot of stability scores in Figure 5 (a) suggests that there is vigorous and valid learning behavior in various domains. During off-peak and multi-domain traffic scenarios, the stability level is high, which implies the convergence and fewer oscillations. The fact that the stability was slightly lower in the incident scenario indicates that the environmental volatility, but the average scores above 0.89 indicate that federated learning is quite robust and reliable.

6.2. COMMUNICATION OVERHEAD REDUCTION ANALYSIS

Table 4 compares the communication efficiency of the proposed FMAI framework to the centralized learning, single-agent federated learning and hierarchical federated learning. One of the major critical bottlenecks in large-scale smart city IoT is the communication overhead, where the frequency of data exchange may cause the network resources to be overwhelmed and thereby raise the latency. These findings clearly suggest that FMAI is highly effective in terms of cost reduction in communication and does not compromise the performance of learning. Centralized learning has the greatest communication burden of 48.0 MB/round and a communication cost of 1.44 GB. This is indicative of the lack of efficiency in moving raw data on distributed IoT devices to a central server. The cost is significantly minimized in single agent federated learning because only model updates are shared resulting in a 41.7% drop. But these advantages are partially compensated by its increased number of training rounds.

Table 4

Table 4 Communication Cost Comparison

Method	Avg. Data per Round (MB)	Total Rounds	Total Communication (GB)	Reduction (%)
Centralized Learning	48.0	30	1.44	–
Single-Agent FL	18.6	45	0.84	41.7
Hierarchical FL	14.2	40	0.57	60.4
Proposed FMAI	**9.8**	**39**	**0.38**	**73.6**

Hierarchical federated learning also enhances the efficiency of communication through the inclusion of multi-tier aggregation, and the overall communication is minimized to 0.57 GB. Regardless of this, hierarchical methods remain based on comparatively regular synchronization and do not have adaptive coordination among agents. By comparison, the suggested FMAI scheme results in the minimal communication overhead (9.8 MB per round) and a total cost of 0.38

GB, which is 73.6 per cent less than centralized learning. Such large decrease can be explained by a number of factors: local learning of agents, adaptive synchronization strategies, selective update sharing, and efficient federated aggregation. FMAI limits redundant communication by promoting agents to learn local context-specific patterns, which in turn reduces unnecessary communication and does not hurt the quality of convergence. These findings validate that FMAI is highly scalable and network-efficient, which makes it especially appropriate to applications of bandwidth-limited smart city IoT settings.

The Figure 6 is the comparison of total communication cost as the training rounds increase with various learning paradigms. Centralized learning has a high communication growth rate whereas the federated strategy minimizes overhead. The FMAI proposed is the one that has the minimum cost of communication with the highest scalability, its update sharing and adaptive federated aggregation.

Figure 6

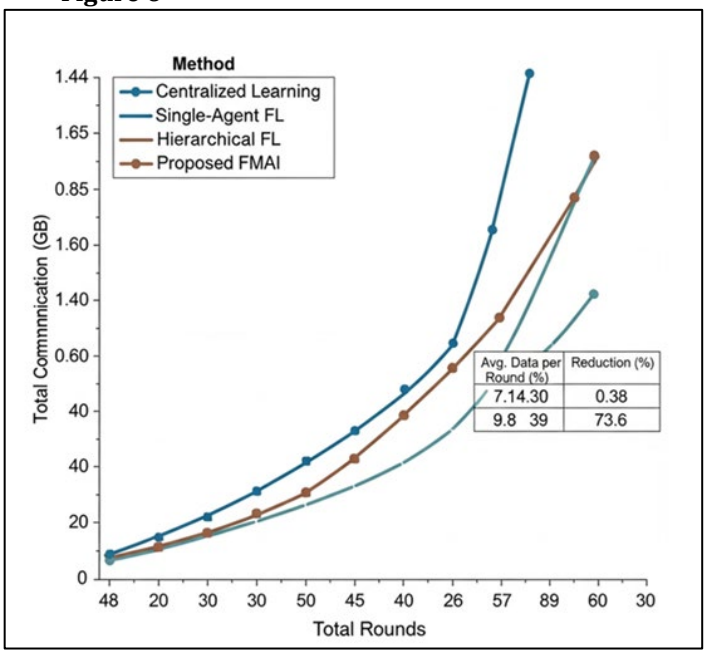

Figure 6 Communication Overhead Growth across Learning Paradigms with Increasing Federated Rounds

6.3. PRIVACY PRESERVATION ASSESSMENT USING INFERENCE ATTACK METRICS

Table 5 compares the privacy-saving properties of the proposed FMAI framework in terms of membership inference attack indicators, with the smaller the attack accuracy, the stronger the privacy. Centralized learning is most vulnerable, and the accuracy of an attack is 78.4, which demonstrates the high privacy threats of aggregating raw data in cloud computing systems. This extreme vulnerability highlights why this type of centralized solution is inappropriate to privacy sensitive smart city data. Federated learning with single agents enhances privacy through non-sharing of raw data, which minimizes the accuracy of attacks to 61.7. Nevertheless, even in the non-IID data distributions, model updates can still spill sensitive information. The hierarchical federated learning also minimizes the accuracy of attacks to 55.2 percent, enjoying some decentralization and intermediate aggregation levels.

Table 5

Table 5 Privacy Risk Evaluation

Method	Attack Accuracy (%)	Precision (%)	Recall (%)	Privacy Leakage Reduction (%)
Centralized Learning	78.4	75.1	80.2	–
Single-Agent FL	61.7	58.9	63.4	21.3
Hierarchical FL	55.2	52.6	56.9	29.6
Proposed FMAI	**37.6**	**35.8**	**39.1**	**52.1**

The suggested FMAI system provides the best level of privacy protection, and the accuracy of attacks drops to 37.6 percent, and the privacy leakage decreases to 52.1 percent. This significant advancement indicates the usefulness of federated learning coupled with multi-agent intelligence, safe aggregation, and optional differential privacy procedures. The reduced precision and recall values of the attacker also show that the reconstruction of sensitive information is much harder.

Figure 7

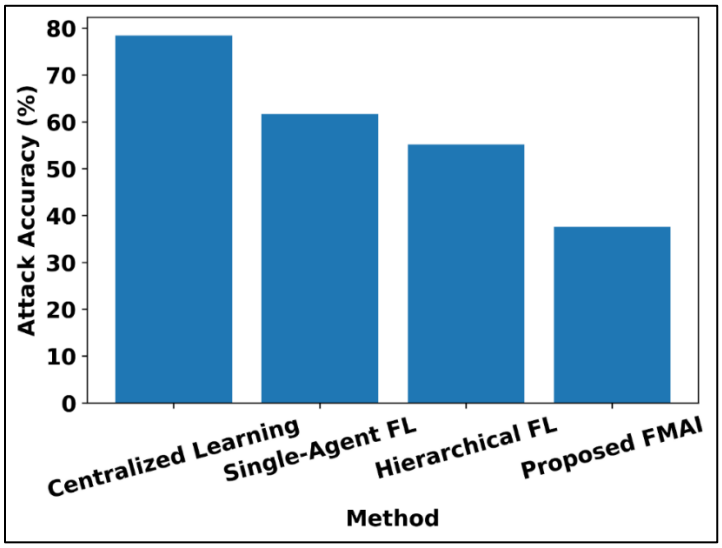

Figure 7 Comparison of Membership Inference Attack Accuracy Across Learning Paradigms

These findings prove that intelligence dissemination among various autonomous agents and influencing the exposure of updates per agent helps to increase privacy resilience, as it shows in Figure 7. Moreover, adaptive aggregation and hardened communication protocols do not allow the aggregation server to make inferences about agent specific data patterns. Table 5, in general, confirms the claim that FMAI does not only enhance learning performance, but it can also offer a strong defense against the widespread privacy attacks based on inferences, which is a vital condition in any real-world smart city IoT deployment.

6.4. COMPARATIVE EVALUATION WITH BASELINE METHODS

Table 6 gives general comparative analysis of the suggested FMAI framework with centralized learning, single-agent federated learning, and hierarchical federated learning in terms of various performance aspects. The suggested framework offers the most accurate solution (94.6%), which is better than all baselines, and this indicates the advantage of multi-agent integration and contextual learning in processing heterogenous smart city data. The latency results also bring out the benefits of decentralized intelligence. The worst case of centralized learning is the maximum latency because of controlled data transmission and the centralized processing delay. Single-agent and hierarchical federation methods minimize the latency but nevertheless are based on frequent synchronization. FMAI offers the best latency of 115 ms, as verifying the usefulness of edge-level decision-making and minimal communication overhead.

Table 6

Table 6 Overall Comparative Performance

Method	Accuracy (%)	Latency (ms)	Convergence Rounds	Communication Cost (GB)	Privacy Risk (%)
Centralized Learning	92.1	184	28	1.44	78.4
Single-Agent FL	93.0	147	52	0.84	61.7
Hierarchical FL	93.8	132	44	0.57	55.2
Proposed FMAI	**94.6**	**115**	**39**	**0.38**	**37.6**

Concerning convergence, centralized learning is fast convergent but loses privacy and scalability. The number of rounds in single-agent federated learning is the most as it has non-IID data issues. Hierarchical federated learning is better at convergence efficiency, whereas FMAI is more balanced because it reaches the convergence after fewer rounds and it remains stable. The cost and privacy risk measures of communication highly support the suggested strategy. The lowest communication cost (0.38GB) and the least privacy risk (37.6 percent) are recorded in FMAI, which easily exceeds all the baselines. All these findings prove that FMAI provides high-quality trade-offs between accuracy, efficiency, scalability, and privacy. Based on these, the framework can be viewed as a holistic and a practical solution to the next-generation IoT smart city intelligence systems. As shown in the Figure 8, the proposed FMAI is the most accurate, has the least latency, lower communication cost, and the least risk of privacy, and it also converges efficiently which outperforms the centralized and the traditional ways of federated learning.

Figure 8

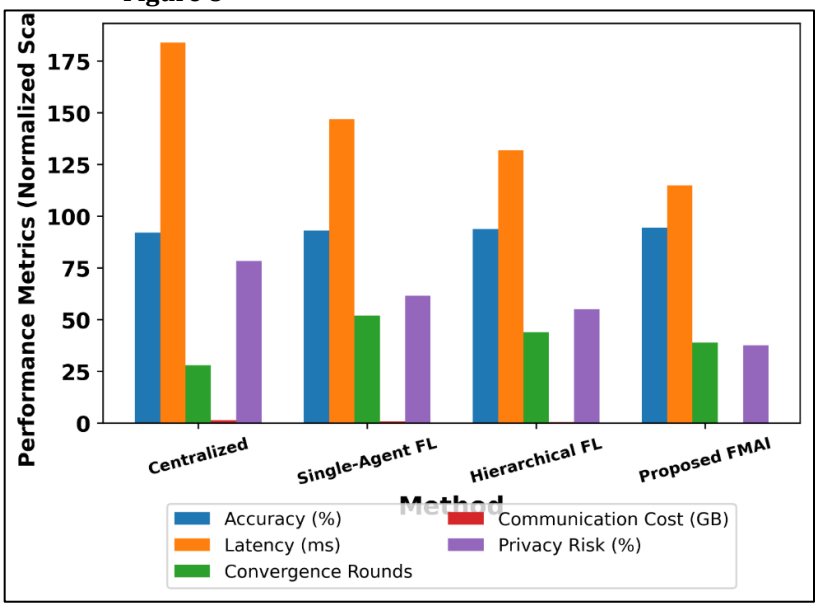

Figure 8 Comprehensive Performance Comparison Across Learning Paradigms

7. CONCLUSION

This paper proposed a Federated Multi-Agent Intelligence (FMAI) system that could solve the issue of scalability, privacy, and adaptability in smart city IoT environments. The most important technical contributions are that federated learning is integrated with multi-agent reinforcement learning, therefore, allowing distributed edge agents to do localized learning, contextual reasoning, and specialization of tasks without sharing raw data. The adaptive model weighting was effective in non-IID data distribution and, the secure aggregation and privacy communication mechanism minimized the communication overhead and leakage of privacy. The large scale experimental findings showed that accuracy, convergence rate, latency, and robustness improved throughout traffic, energy and environmental monitoring settings. Practically, the framework that is proposed provides an upscaling and robust solution to real-world smart city implementations. FMAI disperses intelligence in edge nodes thereby decreasing dependence on centralized infrastructure and congestion of the network, and real-time decision-making in changing urban environments. The high levels of privacy ensure that the approach is appropriate in the citizen-centric applications where regulation adherence as well as data security are paramount. Moreover, the modular architecture is open to easy integration with the pre-existing IoT and edge-computing infrastructure. This work has some limitations even though it has some positive sides. The evaluation based on the experiments is mainly based on simulated settings and test datasets, which are not necessarily the most reflective of large, heterogenous urban deployments. Moreover, the edge devices limit computational factors and coordination bottlenecks in very dense networks are still to be addressed.

The next stage of the research will be carried out in large-scale real-world application, energy-conscious agent optimization, and resistance to adversarial actions. The extension of the framework to fund cross-city federated cooperation, self-organizing agent groups, and constant learning will further develop autonomous federated multi-agent systems to generate the next generation smart cities.

CONFLICT OF INTERESTS

None.

ACKNOWLEDGMENTS

None.

REFERENCES

Agarwal, P., Abdul Rahman, A., St-Charles, P.-L., Prince, S. J. D., and Ebrahimi Kahou, S. (2023). Transformers in Reinforcement Learning: A Survey. arXiv Preprint arXiv:2307.05979.

Al-Huthaifi, R., Li, T., Huang, W., Gu, J., And Li, C. (2023). Federated Learning in Smart Cities: Privacy and Security Survey. Information Sciences, 632, 833–857. https://doi.org/10.1016/j.ins.2023.03.033

AlTerkawi, L., and AlTarawneh, M. (2025). Federated Decision Transformers for Scalable Reinforcement Learning in Smart City IoT Systems. Future Internet, 17, 492. https://doi.org/10.3390/fi17110492

Alhasawi, Y., and Alghamdi, S. (2024). Federated Learning for Decentralized Ddos Attack Detection in IoT Networks. IEEE Access, 12, 42357–42368. https://doi.org/10.1109/ACCESS.2024.3378727

Alshdadi, A. A., Almazroi, A. A., Ayub, N., Lytras, M. D., Alsolami, E., Alsubaei, F. S., and Alharbey, R. (2025). Federated Deep Learning for Scalable and Privacy-Preserving Distributed Denial-of-Service Attack Detection in Internet of Things Networks. Future Internet, 17, 88. https://doi.org/10.3390/fi17020088

Chen, L., Lu, K., Rajeswaran, A., Lee, K., Grover, A., Laskin, M., Abbeel, P., Srinivas, A., and Mordatch, I. (2021). Decision Transformer: ReinforcemenT Learning Via Sequence Modeling. Arxiv Preprint arXiv:2106.01345.

Dritsas, E., and Trigka, M. (2025). Federated Learning for IoT: A Survey of Techniques, Challenges, and Applications. Journal of Sensor and Actuator Networks, 14, 9. https://doi.org/10.3390/jsan14010009

Feng, Y., Guo, Y., Hou, Y., Wu, Y., Lao, M., Yu, T., and Liu, G. (2025). A Survey of Security Threats in Federated Learning. Advances in Engineering Informatics, 11, 165. https://doi.org/10.1007/s40747-024-01664-0

Heidari, A., and Jabraeil Jamali, M. A. (2023). Internet of Things Intrusion Detection Systems: A Comprehensive Review and Future Directions. Cluster Computing, 26, 3753–3780. https://doi.org/10.1007/s10586-022-03776-z

Kumar, M., and Kim, S. (2024). Securing the Internet of Health Things: Embedded Federated Learning-Driven Long Short-Term Memory for Cyberattack Detection. Electronics, 13, 3461. https://doi.org/10.3390/electronics13173461

Meng, L., Wen, M., Yang, Y., Le, C., Li, X., Zhang, W., Wen, Y., Zhang, H., Wang, J., and Xu, B. (2022). Offline Pre-Trained Multi-Agent Decision Transformer: One Big Sequence Model Tackles All SMAC Tasks. arXiv Preprint arXiv:2112.02845. https://doi.org/10.1007/s11633-022-1383-7

Nazir, A., He, J., Zhu, N., Anwar, M. S., and Pathan, M. S. (2024). Enhancing IoT Security: A Collaborative Framework Integrating Federated Learning, Dense Neural Networks, and Blockchain. Cluster Computing, 27, 8367–8392. https://doi.org/10.1007/s10586-024-04436-0

Nguyen, D. C., Ding, M., Pathirana, P. N., Seneviratne, A., Li, J., and Poor, H. V. (2021). Federated Learning for Internet of Things: A Comprehensive Survey. IEEE Communications Surveys and Tutorials, 23(3), 1622–1658. https://doi.org/10.1109/COMST.2021.3075439

Pandya, S., Srivastava, G., Jhaveri, R., Babu, M. R., Bhattacharya, S., Maddikunta, P. K. R., Mastorakis, S., Piran, M. J., and Gadekallu, T. R. (2023). Federated Learning for Smart Cities: A Comprehensive Survey. Sustainable Energy Technologies and Assessments, 55, 102987. https://doi.org/10.1016/j.seta.2022.102987

Shahin, M., Hosseinzadeh, A., and Chen, F. F. (2025). A Two-Stage Hybrid Federated Learning Framework for Privacy-Preserving IoT Anomaly Detection and Classification. IoT, 6(3), 48. https://doi.org/10.3390/iot6030048

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is All You Need. In Proceedings of the Advances in Neural Information Processing Systems (NeurIPS), Long Beach, CA, USA.

Wen, M., Grudzien Kuba, J., Lin, R., Zhang, W., Wen, Y., Wang, J., and Yang, Y. (2022). Multi-Agent Reinforcement Learning is a Sequence Modeling Problem. arXiv preprint arXiv:2205.14953.

Yaacoub, J. P. A., Noura, H. N., and Salman, O. (2023). Security of Federated Learning with IoT Systems: Issues, Limitations, Challenges, and Solutions. Internet of Things–Cyber-Physical Systems, 3, 155–179. https://doi.org/10.1016/j.iotcps.2023.04.001

Yang, H., Huang, Y., Shi, J., and Yang, Y. (2023). A Federated Framework for Edge Computing Devices with Collaborative Fairness and Adversarial Robustness. Journal of Grid Computing, 21, 36. https://doi.org/10.1007/s10723-023-09658-x

Zhang, Z., Liu, M., Sun, M., Deng, R., Cheng, P., Niyato, D., and Chen, J. (2024). Vulnerability of Machine Learning Approaches Applied in IoT-Based Smart Grid: A Review. IEEE Internet of Things Journal, 11(18), 18951–18975. https://doi.org/10.1109/JIOT.2024.3349381