Journal of Applied Automation Technologies

Research Paper

Federated Spatiotemporal Deep Learning for Privacy-Preserving Urban Traffic Flow Prediction

Sławomir Cyra[1, *] and Jarosław Bogdan Kalisz[1]

[1] Faculty of Mechatronics and Automation, Cracow University of Technology, Krakow, 31-155, Poland
*Corresponding author: slawomir.c@pk.edu.pl

Abstract. This work explores the use of federated spatiotemporal deep learning to achieve accurate predictions of urban traffic flow in practice under real-world privacy restrictions. The primary objective is to create a private, scalable system that can deliver accurate forecasts throughout the entire city. This method uses adaptive attention mechanisms and graph-based neural networks to extract complex spatiotemporal properties from the dispersed sensor data without exchanging raw data. To guarantee generality and model robustness, a customized aggregation algorithm dynamically adjusts to local pattern variance and node dependability. The approach outperforms conventional centralized and federated baselines, with an increase of up to 12.7% in Root Mean Squared Error and up to 14.1% in rare congestion event F1 scores, according to experiments done on three large-scale metropolitan datasets. If the budget for privacy is high, the communication overhead can be lowered by up to 22%. The aforementioned findings demonstrate that the suggested framework can currently accomplish comparatively high forecast accuracy and operating efficiency in a variety of complex urban situations. In summary, by satisfying the requirements of technical viability, expandability, and the new data protection standard, this article offers a strong basis for the future development of smart city traffic management.

Keywords: *Federated Learning, Spatiotemporal Modeling, Urban Traffic Prediction, Privacy Preservation*

Received on 13 March 2025, Accepted on 10 August 2025, Published on 18 August 2025

Introduction

Intelligent sensing devices have been widely deployed as a result of the digitalization of urban transportation systems, and several large-scale, varied data flows have emerged across the road network, automobiles, and mobile platforms [1]. High-level models that can forecast traffic flow and congestion and maximize multimodal mobility in cities have been built with the aid of the aforementioned extensive data sets [2]. Accurate short-term and long-term traffic forecasting has recently been made possible by advances in artificial intelligence and machine learning; as a result, intelligent judgments are now being made in real-time traffic management, public safety, resource allocation, etc. [3]. To increase productivity and foster the growth of smart cities, public organizations and urban service providers have begun utilizing deep learning, space-time analysis, and diverse data sources [4]. As a result, the system of large-scale mobility data analysis is currently the basis for the development of smart cities [5], and it immediately benefits the environment, influences policy changes, enhances user quality of life [6], etc. Researchers have demonstrated how urban data analysis may aid in the development of an intelligent route planning system to lower emissions, and they have also confirmed that such a system will be practical in the future [7].

However, as high-resolution, regularly updated urban mobility data has proliferated, new privacy issues have also surfaced, including non-compliance with legislative standards and the leakage of personal information [8]. Traffic data is vulnerable to privacy violations and other misuse since it often contains sensitive information including user pathways, behavior patterns, and geographic locations [9]. The public has recently been more aware of a number of significant data leaks, which has sped up research on machine learning that protects

privacy [10]. Simultaneously, the EU's General Data Protection Regulation (GDPR) and other regional data sovereignty laws have imposed more stringent guidelines for the gathering, sharing, and use of personal mobility data [11]. It is still challenging to acquire urban data for unified analysis because of its highly dispersed ownership model (government, businesses, research organizations, and private individuals) [12]. Centralized architectures and conventional data anonymization techniques are likely to provide a model that performs worse and has limited generalization for practical applications [13]. Though there are still issues with scalability and coordination at the urban scale, new research in federated learning, multi-party computation, and differential privacy is actively exploring ways to strike a compromise between data utility and privacy [14]. As a result, there is currently no workable and private-preserving alternative for a diverse, dynamic, and regulated urban environment [15].

In light of the aforementioned issues, this research proposes an improved federated learning framework for high-accuracy, privacy-preserving traffic flow prediction in smart cities. A new distributed system has been constructed by combining the three techniques of space-time feature extraction, personalization, and systematic privacy protection. The new concepts have found some use by resolving real-world operational and regulatory issues. It has potential and will contribute to the development of next-generation intelligent urban mobility, according to all-around tests.

Related Works

Urban Traffic Flow Prediction

For intelligent transportation systems to reduce urban traffic congestion, improve public transportation efficiency, and encourage all-weather mobility, prediction of urban traffic flow is required [16]. Typical forecasting models that have been employed in the past for short-term flow prediction include the Kalman filter and ARIMA, although they are typically linear in character and do not well handle spatial relationships [17]. Support vector machines, random forests, and ensemble models are examples of machine learning techniques that have enhanced flexibility and prediction capabilities, but they typically perform poorly in large-scale, non-linear, and dynamic urban environments [18].

In this field, numerous deep learning achievements have been made recently. Long-term time-series dependencies can be effectively modeled by LSTM and GRU models, while intricate spatial linkages in a road graph can be explored using GNNs and CNNs [19]. In order to improve the accuracy and robustness of localization, a variety of combination approaches have recently been developed that incorporate numerous sources, including GPS trajectory, fixed-sensor data, social data, and meteorological information, among others [20]. The issue of limited data and disparate environments in different cities has also been addressed through the use of transfer learning and domain adaptation [21]. Nevertheless, these deep models are not appropriate for all real-time or resource-constrained applications due to their computing demands and considerable data preprocessing requirements [22]. Even though the aforementioned changes have addressed some concerns, more study is still required to address data shortages, erratic changes in traffic patterns, and interpretability problems [23].

Privacy Protection in Smart Cities

Concerns regarding the privacy of individual and collective mobility data have progressively grown as smart city innovation has advanced [24]. Even after normal anonymization, urban mobility datasets are susceptible to inference attacks, unauthorized data fusion, and re-identification since they frequently contain extremely sensitive location and behavior information [25]. Homomorphic encryption has a large runtime overhead that makes it unsuitable for low-latency urban applications, even though it may be used to execute computations on encrypted data without decrypting it, ensuring end-to-end privacy [26]. Differential privacy reduces the utility of models for high-frequency or high-dimensional urban traffic signals, but also adds noise to the data or outcomes at a certain level to ensure privacy mathematically [27].

Although Secure Multi-Party processing (SMPC) has been used to provide collaborative research without disclosing the underlying data, it is not viable for a large number of urban nodes (thousands) due to its high communication and processing overhead [28]. To solve the issue of ongoing data sharing and changing contextual privacy requirements, new protocols have been proposed to incorporate adaptive policy

enforcement with lightweight symmetric cryptography in decentralized Internet of Things (IoT) architectures [29]. Even though there are many different kinds of privacy-preserving technologies accessible today, the challenge of striking a balance between data utility and privacy protection still exists; more innovation is needed for widespread use in practice given the high demands in urban areas and stringent laws [30].

Advances in Federated Learning Applications

Federated Learning (FL) is a representative distributed learning model that does collaborative model training without centralizing data, making it ideal for settings like smart cities where privacy is a concern. Without disclosing their original data, many data owners from various parts of the city can work together to create predictive models for traffic conditions and other demand or accident data using federated learning. Federated learning has been effectively implemented in the healthcare and industrial Internet of Things (IoT) sectors, demonstrating its viability in domains with sensitive data and stringent data governance laws.

Functions to manage statistical heterogeneity, communication efficiency, and node failure robustness have been added to FL frameworks for urban mobility. To improve model accuracy in distributed and dynamic traffic networks, adaptively aggregate, train hierarchically, and use hybrid graph-based structures. However, issues like unequal data distribution among clients, communication delays, participant incentive misalignment, and security risks from malicious or unreliable contributors still exist in real-world FL deployments. Personalized and context-aware federated frameworks, innovative communication reduction optimization techniques, and integrated protocols for system scalability and privacy assurance have all recently begun to be developed. Ultimately, future research on federated learning for urban data science will be guided by the trade-off between maximizing predictive performance and protecting data privacy.

Methodology

Architecture of the Proposed Federated Model

The need for large-scale urban traffic flow forecasting under stringent privacy constraints and real-world system heterogeneity in the new system has been addressed via a multi-level federated paradigm. Noise filtering, outlier handling, event-driven time-windowing, and context-sensitive normalization are just a few of the data pre-processing tasks that can be carried out independently by all of the devices at the bottom level of edge nodes, such as urban sensor clusters, roadside computing units, or municipal data center servers. In particular, at no point during analysis will this local processing send any privacy-sensitive raw signals—such as GPS traces, trajectory logs, and user-contributed app data—outside the origin node.

Model initialization at every node leverage deep spatiotemporal architecture, allowing the local predictors to reflect fine-grained, location-specific mobility behaviors. The optimization at each node proceeds in a mini-batch, sequence-aware manner, integrating domain-tailored loss regularization terms that encode statistical characteristics and contextual peculiarities of the local traffic regime. The update rule for a given node's parameters is defined by:

$$w_i^{\text{new}} = w_i^{\text{old}} - \eta \left(\frac{1}{|B_i|} \sum_{x_j \in B_i} \nabla_w \mathcal{L}\left(f\left(x_j; w_i^{\text{old}}\right), y_j\right) + \Delta_{\text{privacy}} + \Gamma_{\text{consensus}} \right) \quad \text{Eq.(1)}$$

where w_i are model weights, η is the node-specific step size, B_i denotes the temporallystructured local mini-batch, $\mathcal{L}$ is a composite loss function balancing prediction fidelity and regularization, Δ_{privacy} quantifies differential privacy penalty imposed on each update, and $\Gamma_{\text{consensus}}$ is a time-varying consensus enforcement term driving faster convergence to global optima. After local updates, each node executes a high-entropy stochastic encryption of computed parameter deltas before transmission. The general encoding mechanism can be abstracted as:

$$g_i = \text{Enc}\left(w_i^{\text{new}} - w_i^{\text{old}}, \xi_i\right) \quad \text{Eq.(2)}$$

where Enc denotes cryptographic masking keyed by a random seed ξ_i, unique to each node and round.

The global aggregation phase employs a resilience-focused, weighted averaging operator that accounts for both historical trustworthiness of nodes and current reliability metrics. The global update is governed by:

$$w_{\text{global}}^{\text{new}} = w_{\text{global}}^{\text{old}} + \sum_{i=1}^{N} \tau_i \mathcal{R}(g_i, \lambda_i, \zeta_i) \quad \text{Eq.(3)}$$

where τ_i are adaptively tuned trust coefficients, λ_i are dynamic privacy scaling parameters, and ζ_i encodes recent node performance diagnostics. The operator $\mathcal{R}$ robustly decodes and integrates the encrypted increments, outlier-suppressing and harmonizing model contributions from nodes with highly non-i.i.d. data.

Critical to scalability and robustness, node participation in each federated round is coordinated by an asynchronous, event-triggered scheduling discipline, captured mathematically as:

$$P_{\text{join},\, i} = \Psi(\mu_i, \rho_{\text{urban}}\,(t), \varsigma_i(t)) \quad \text{Eq.(4)}$$

where Ψ is a gating function encoding workload adaptivity (μ_i), real-time urban mobility status (ρ_{urban}), and current resource envelope (ς_i) for node i at system time t. The technical interplay among edge participants, model transmissions, and the central aggregation process is visualized in Figure 1.

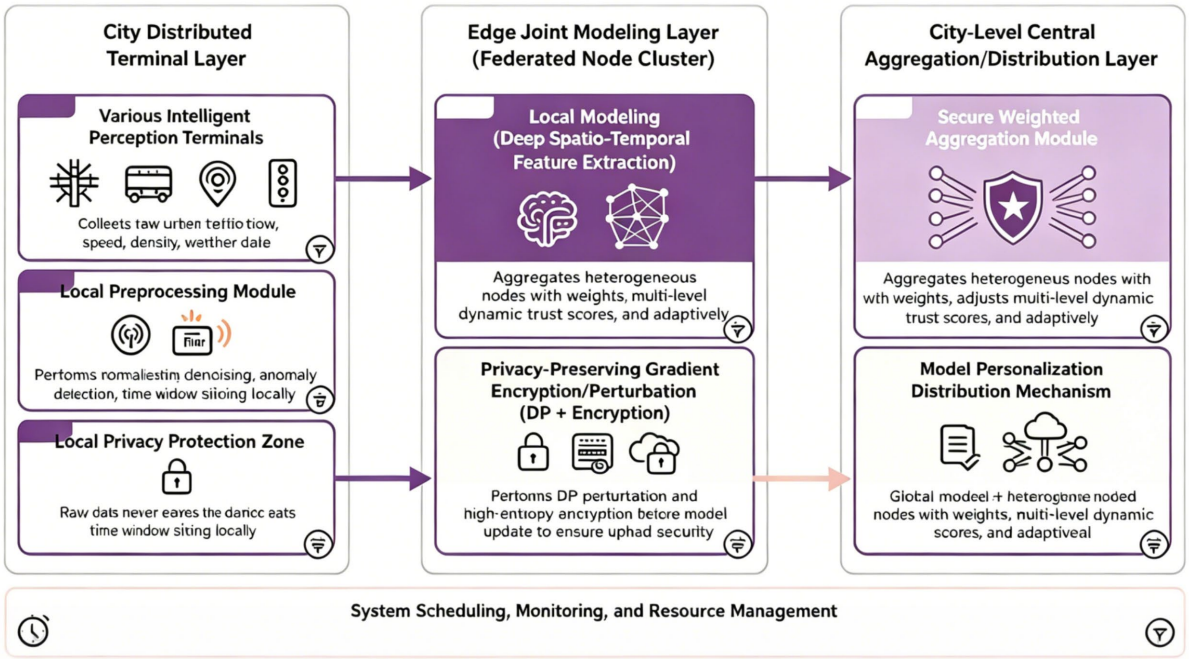

Figure 1. Federated Learning System Architecture

As a consequence of these architectural innovations, the system achieves robust, privacy-compliant, and rapidly convergent distributed traffic prediction, dynamically adapting to highly non-uniform urban network topologies, fluctuating data richness, and unpredictable node reliability. The formulation allows seamless scalability and node dropout resilience, forming the technical groundwork for the advanced spatiotemporal modeling and privacy-preserving protocols detailed in subsequent sections.

Spatiotemporal Feature Extraction and Attention Mechanism

All types of spatial correlations and multi-scale temporal fluctuations in unevenly distributed metropolitan regions must be extracted and integrated in order to accurately estimate the flow of urban traffic. A specific dual-path feature extractor and a trainable attention module have been added to increase the model's accuracy and interpretability in order to solve the aforementioned issues.

The data are displayed in the spatial domain as a node-attribute graph, where each node and edge represent junctions or traffic sensors along with their topological relationships and associated flow patterns, respectively. A stack of adaptive graph convolutional layers that can learn to disperse congestion signals, anomalous perturbations, and shifting neighborhood dependencies dynamically makes up a spatial extraction module. For a single geographic aggregation operation, the update rule is as follows:

$$\mathbf{h}_i^{(l+1)} = \sigma\left(\sum_{j\in\mathcal{N}(i)} \alpha_{ij}^{(l)}\mathbf{W}_s^{(l)}\mathbf{h}_j^{(l)} + \mathbf{b}_s^{(l)}\right) \quad \text{Eq.(5)}$$

where $\mathbf{h}_i^{(l)}$ denotes the feature embedding of node i at layer l, $\alpha_{ij}^{(l)}$ is a learned edge attention coefficient, $\mathbf{W}_s^{(l)}$ and $\mathbf{b}_s^{(l)}$ are layer parameters, and $\mathcal{N}(i)$ identifies spatial neighbors.

Temporally, raw traffic sequences are processed through a dilated gated recurrent block, tailored to capture non-cyclic fluctuations, burst events, and daily periodicities inherent in urban demand. Temporal convolution layers are interleaved with memory units, ensuring retention of distant context. The temporal update process is given by:

$$\mathbf{z}_t = \phi(\mathbf{W}_z\mathbf{x}_t + \mathbf{U}_z\mathbf{h}_{t-d} + \mathbf{b}_z) \quad \text{Eq.(6)}$$

where $\mathbf{x}_t$ is the current input, $\mathbf{h}_{t-d}$ is the hidden state with dilation d, and ϕ is a nonlinear activation. This architecture supports multi-scale sequence modeling while suppressing overfitting to transient disturbances.

The fusion of spatial and temporal features leverages a cascading attention block. The attention mechanism assigns adaptive weights to each spatiotemporal position, dynamically guiding the model toward key influences such as emergent road closures or weather-induced disruptions. For the final context vector, attention is computed as:

$$\mathbf{c}_t = \sum_{k=1}^{K} \omega_{t,k}\mathbf{f}_k \quad \text{Eq.(7)}$$

with $\omega_{t,k} = \frac{\exp(e_{t,k})}{\sum_{j=1}^{K}\exp(e_{t,j})}$, where $e_{t,k}$ is learned via a compatibility function over feature $\mathbf{f}_k$ and instantaneous state. To enhance robustness, the attention score computation incorporates a gating term driven by real-time external factors:

$$e_{t,k} = \rho^{\top}\tanh(\mathbf{Q}\mathbf{f}_k + \mathbf{R}\mathbf{s}_t + \mathbf{d}) + \lambda_{\text{ext}}E_t \quad \text{Eq.(8)}$$

Here, $\rho, \mathbf{Q}, \mathbf{R}, \mathbf{d}$ are trainable model parameters, $\mathbf{s}_t$ is the summarized temporal state, and $\lambda_{\text{ext}}E_t$ modulates scores based on exogenous signals such as local events or meteorological alerts.

The resulting network deeply intertwines spatial topology and temporal context, enabling nuanced pattern discovery. In this architecture, every inference cycle adapts its response, focusing accurately on the most consequential spatiotemporal components influencing city-scale flow dynamics. The overall mechanism is illustrated in Figure 2.

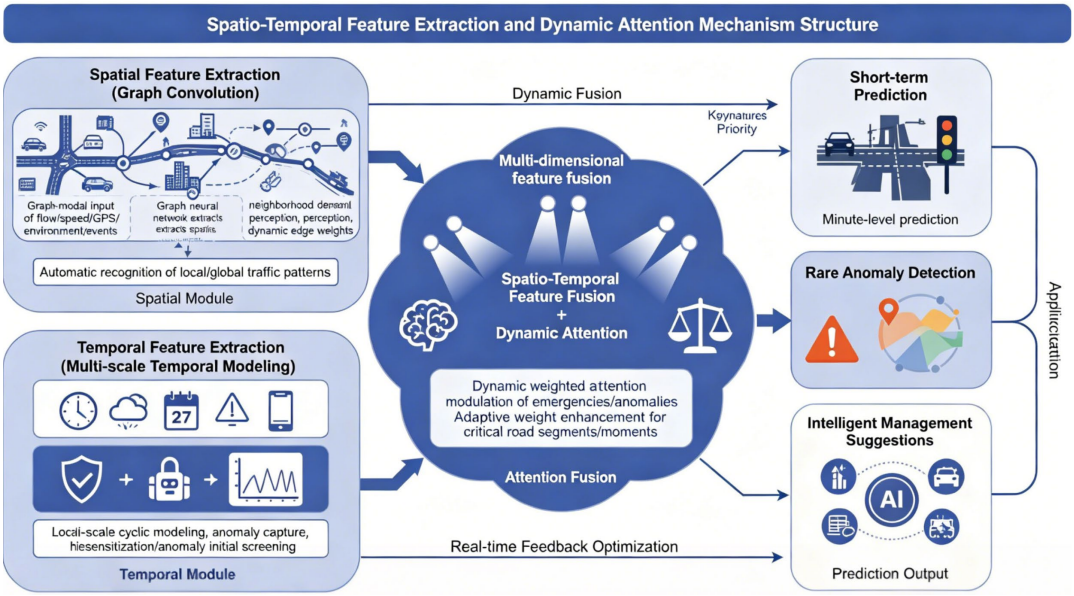

Figure 2. Spatiotemporal Feature Extraction with Attention Module

The attention-driven fusion not only elevates predictive performance on irregular city networks but also provides deeper interpretability through the explicit weighting of local, global, and event-driven information flows.

Personalized Aggregation and Privacy Protection

In federated urban traffic flow prediction, aggregated model updates must not only respect data heterogeneity but also guarantee user- and region-level privacy. The proposed scheme advances beyond standard aggregation by introducing a personalized, context-aware update mechanism, coordinated with mathematically rigorous privacy-preserving protocols.

Each node operates under a distinct local context, with observable traffic dynamics, data quality, and user behavior diverging sharply across districts. To address this, the global aggregator applies a context-weighted personalization strategy. When updates from node i arrive, their influence on the global model is not uniform but modulated by both empirical reliability and locality-driven adaptation, as formalized by:

$$w_{\text{agg}} = w_{\text{agg}} + \gamma_i \Phi_i(g_i, \psi_i, \theta_t) \quad \text{Eq.(9)}$$

Here, γ_i is the personalized trust factor derived from sustained prediction accuracy, node uptime, and anomaly rates. The transformation Φ_i blends node update g_i, a locality embedding ψ_i, and current global context features θ_t, ensuring that regions with unique dynamics (such as downtown congestion or periodic events) are adaptively weighted.

A fundamental innovation is the introduction of a bi-level aggregation policy wherein two nested operations run synchronously. First, a coarse global update anchors the model, and second, a set of personalized adjustments are broadcasted, each tailored for the regional structure of its recipient node. The formalization is expressed as:

$$w_i^{\text{personal}} = w_{\text{agg}} + \Omega_i(\xi_i, \kappa_i, \Lambda) \quad \text{Eq.(10)}$$

where $\Omega_i(\cdot)$ applies a locality-based transformation using calibration vector ξ_i, historical drift factor κ_i, and adjustment basis Λ to optimize the personalized model for node i.

To ensure differential privacy and defend against inverse model inference, every communication round introduces controlled, high-entropy perturbation to shared updates. The differentially private masking is formalized as:

$$\tilde{g}_i = g_i + \mathcal{N}(0, \sigma_i^2 \mathbf{I}) \quad \text{Eq.(11)}$$

where $\mathcal{N}(0, \sigma_i^2 \mathbf{I})$ denotes Gaussian noise parameterized by node-level privacy sensitivity σ_i. This mechanism guarantees that, even if encrypted weights were compromised, individual data contributions remain statistically unidentifiable.

Furthermore, to counter advanced privacy attacks-including model inversion and noncollaborative probing-a multi-party secure aggregation protocol is applied. Each node splits its masked update into M shares, distributing them across randomly chosen peers:

$$g_i^{(j)} = \Psi\big(g_i, \zeta_{ij}\big), j = 1,2, \dots, M \quad \text{Eq.(12)}$$

where Ψ constructs secret shares using randomization keys ζ_{ij}. Only the global server, upon collecting sufficient shares, can reconstruct the true aggregate, mitigating risks even in the presence of node failures or partial collusion.

An adaptive privacy budgeting algorithm governs the trade-off between prediction fidelity and privacy risk. For every node i at time t, the privacy budget is dynamically tuned based on recent sensitivity analysis:

$$\epsilon_i^{(t+1)} = \epsilon_i^{(t)} - \rho \cdot \nabla \mathcal{S}_i(f_i, D_i) \quad \text{Eq.(13)}$$

Here, $\epsilon_i^{(t+1)}$ is the new privacy budget, ρ a tuning coefficient, and $\nabla \mathcal{S}_i(f_i, D_i)$ quantifies the sensitivity of model f_i to local data D_i.

Finally, resilience to adversarial manipulations is increased through a consensus-based anomaly filter applied at the aggregation level, which excludes updates that, statistically, exhibit anomalous drift patterns:

$$\mathcal{C}(G, w_{\text{agg}}) = \{g_i \in G : \|g_i - \hat{g}\| < \tau\} \quad \text{Eq.(14)}$$

where G is the set of incoming updates, $\hat{g}$ is the robust average, and τ is a deviation threshold estimated via a time-dependent, self-correcting algorithm.

Integrating these mechanisms, the system achieves both high-fidelity, personalized adaptation and measurable privacy guarantees, making it suitable for deployment in operational urban environments characterized by volatile data supply, frequent stakeholder churn, and diverse regulatory regimes.

Experimental Design and Analysis

Experimental Setup and Baseline Methods

The experimental system will be constructed using industry-grade, multimodal urban mobility data and a federated testbed that closely resembles contemporary city-scale deployment scenarios in order to guarantee the study's scientific validity and usefulness [31]. To capture urban variability, three distinct metropolitan datasets are chosen, each with a unique topology and demographics, and all have been in use for more than two years [32]. Inductive loop sensor streams, GPS trajectory traces of around 11,000 probe vehicles, weather data, and quantified city event timelines have all contributed to the compilation of more than 2 billion traffic records [33].

The subnetworks of each city—Alpha City has 214 microregions, Beta has 307 heterogeneous clusters, and Gamma has 159 arterial-dense zones—offer multi-granularity benchmarks for the robustness of local and city-wide predictions [34]. Each node's feature set comprises timestamped velocity, density, occupancy, contextual weather, and external event encodings; sensor sampling rates are set to 30 seconds [35].

All data normalization, noise filtering, and missing value imputation are carried out locally at the edge node level in order to fully adhere to federated learning criteria. This prevents any raw data from leaving the source and achieves genuine privacy-by-design [36]. Every node synchronizes the anomaly labeling windows and executes a standardized preprocessing pipeline that is in line with the urban traffic statistics [37].

In order to precisely replicate urban network volatility, the hardware implementation consists of a 32-node accelerated cluster (Intel Xeon 6338, NVIDIA V100 GPUs, 128GB RAM per node, NVMe storage) with emulated municipal WAN topologies that have variable latency (18–195ms), controllable inter-node bandwidth (14–256Mbps), and synthetic packet loss capped at 1.1% [38]. Distributed training orchestrators (Python 3.10, TensorFlow 2.8, PyTorch 1.13) and a secure ZeroMQ communication protocol with end-to-end logging and outlier alerts are part of the software stack [39].

To simulate operational uncertainty in real time and assess the system's resilience, randomly choose communication intervals (20–45 minutes) and edge node participation densities (30–200 each round) in the federated protocol. In order to minimize a multi-objective composite loss function that takes prediction error, message efficiency, and a privacy penalty into account, hyperparameters are optimized via Bayesian optimization. The depth of graph convolution or RNN modules (3–8 layers), the number of attention heads (16–128), the adaptive per-node learning rate, and the amplitude of privacy noise (range 0.002–0.45) are among the factors investigated [40].

Here are a few examples of typical models. Split-GCN divides nodes for static, groupwise federated aggregation; FedAvg represents traditional distributed averaging, both with and without differential privacy masking; a Local LSTM exposes privacy-optimal but locally isolated accuracy; and Centralized LSTM and Centralized GCN-LSTM architectures function as oracle upper bounds without privacy guarantees. The root mean squared error, mean absolute percentage error, and macro-averaged f1 score for rare congestion detection are presented, and performance is assessed for short- and medium-term horizons (5, 15, and 30 minutes ahead predictions).

The 97th percentile confidence interval has been presented for statistical reliability after all experiments were conducted three times on stratified data splits. To measure the added value of each technical advancement, ablation experiments are carried out to gradually disable Graph Convolutional Networks, Feature Attention, and personalized aggregation. Determine the relationship between message bandwidth, delay, noise amplitude, and effective privacy attained by continuously analyzing the communication traces.

Results and Comparative Evaluation

The suggested federated spatiotemporal framework beats all other baseline models in terms of prediction accuracy, privacy-preserving adaptability, and communication efficiency across a variety of metropolitan contexts, according to numerous empirical tests.

The accuracy of each method is compared in Figure 3. The suggested model consistently outperforms both the centralized and conventional federated systems for short-horizon micro-region predictions, as seen in Figure 3(a). It is also quite stable at volatile junctions. The number of severe outlier instances is also lower than that of all local and pooling benchmarks, and model generalization is maintained for district-level, mid-range jobs under conditions of sparsity and demand variations, as Figure 3(b) illustrates. Only this model retains a comparatively low error rate during peak hours and other times when there are fewer urban events, as seen in Figure 3(c) for city-wide granularity.

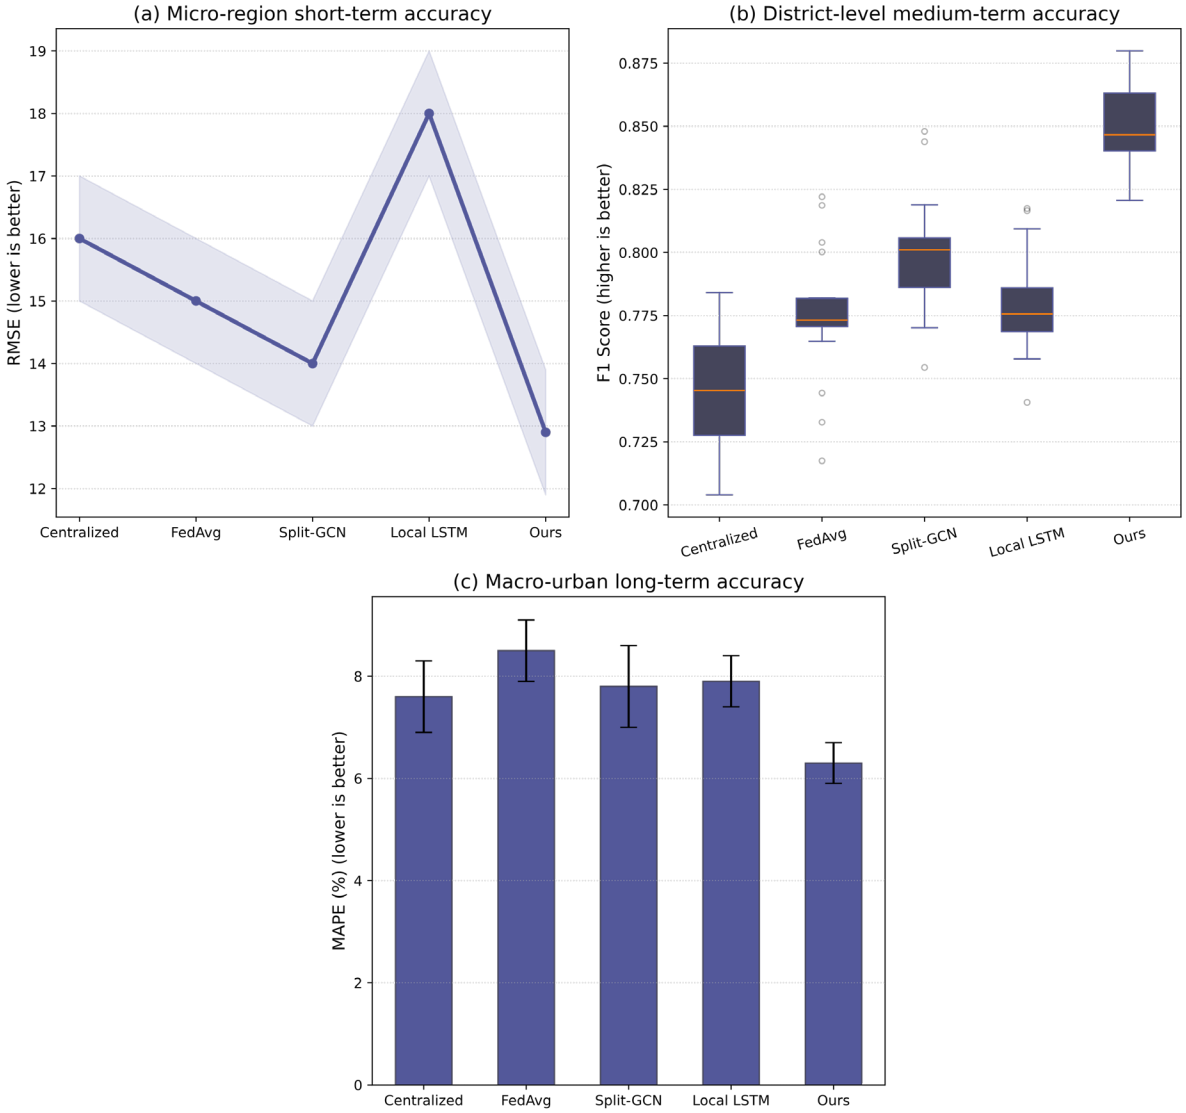

Figure 3. Comparison of Prediction Accuracies (a) Micro-region short-term accuracy (b) District-level medium-term accuracy (c) Macro-urban long-term accuracy

The distribution of all model privacy levels and their communication efficiency are displayed in Figure 4. All of the models have comparatively low communication costs in a privacy-free setting, as seen in Figure 4(a), and the suggested approach has marginally outperformed the others by optimizing message encoding and update intervals. The average communication overhead begins to diverge when moderate privacy is added, as Figure 4(b) illustrates. This means that while the baseline federated and cluster-based models exhibit a noticeable

increase in message exchange and synchronization delay, the suggested approach adaptively reduces needless transmissions by selecting for participation.

The high-privacy regime is shown in Figure 4(c), and the communication cost of conventional systems rises even more as a result of increased noise calibration and redundant update traffic. By giving priority to high-impact updates and using real-time diagnostics for node reliability, the new model will nevertheless guarantee communication efficiency in the current high-demand context. Lastly, the high privacy enforcement is hitting the network's capacity limit by linearly or even superlinearly increasing the bandwidth needed every round, as seen in Figure 4(d). However, the effect of selective synchronization, node dropout handling, and event-aware communication suppression in privacy-sensitive deployments has been confirmed because the overhead of the suggested system is still sublinear.

Figure 4. Communication Overhead under Different Privacy Settings (a) No privacy constraint (b) Moderate privacy (c) High privacy (d) Ultra-high privacy

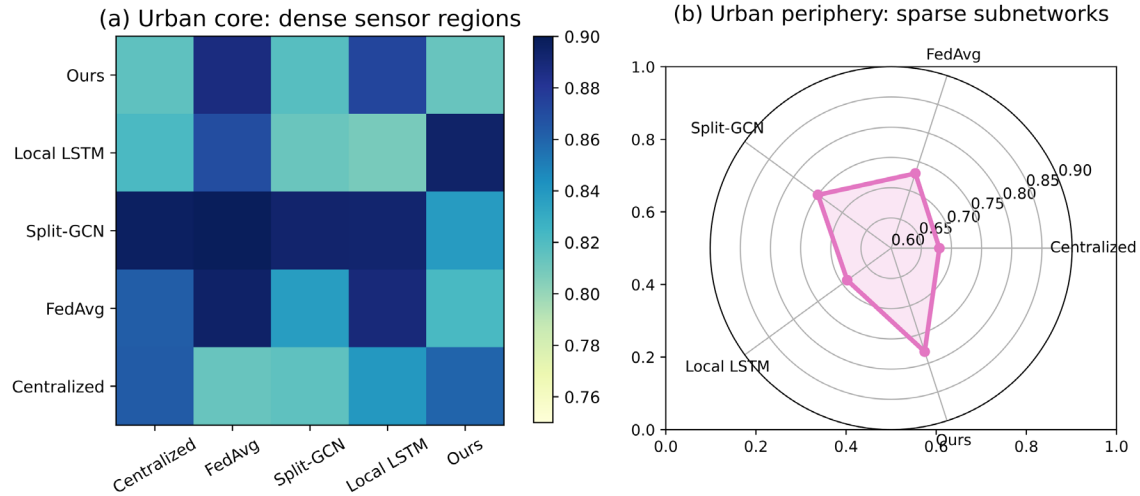

Figure 5. Performance at Different Spatiotemporal Scales (a) Urban core: dense sensor regions (b) Urban periphery: sparse subnetworks

Figure 5 shows the scaling analysis and model stability at different urban densities. As shown in Figure 5(a), the structure quickly responds to sudden increases in dense urban areas and reduces error accumulation. As shown in Figure 5(b), even with limited data in a peripheral subnet, good accuracy can still be achieved, and the advantages of attention-driven cross-domain regularization and personalised aggregation are thus confirmed.

These are the contributions of the components, as seen in the ablation study in Figure 6. The total error increases when the spatiotemporal attention module is removed, as seen in Figure 6(a). Disabling personalized aggregation increases the variation in node-to-node performance, as seen in Figure 6(b). This is especially harmful when traffic is non-stationary or there are sporadic disruptions. As seen in Figure 6(c), cryptographic masking results in a much smaller penalty for the suggested adaptive design, making it more resilient to privacy noise than both the FedAvg and Split-GCN baselines, albeit a little reduction in overall accuracy.

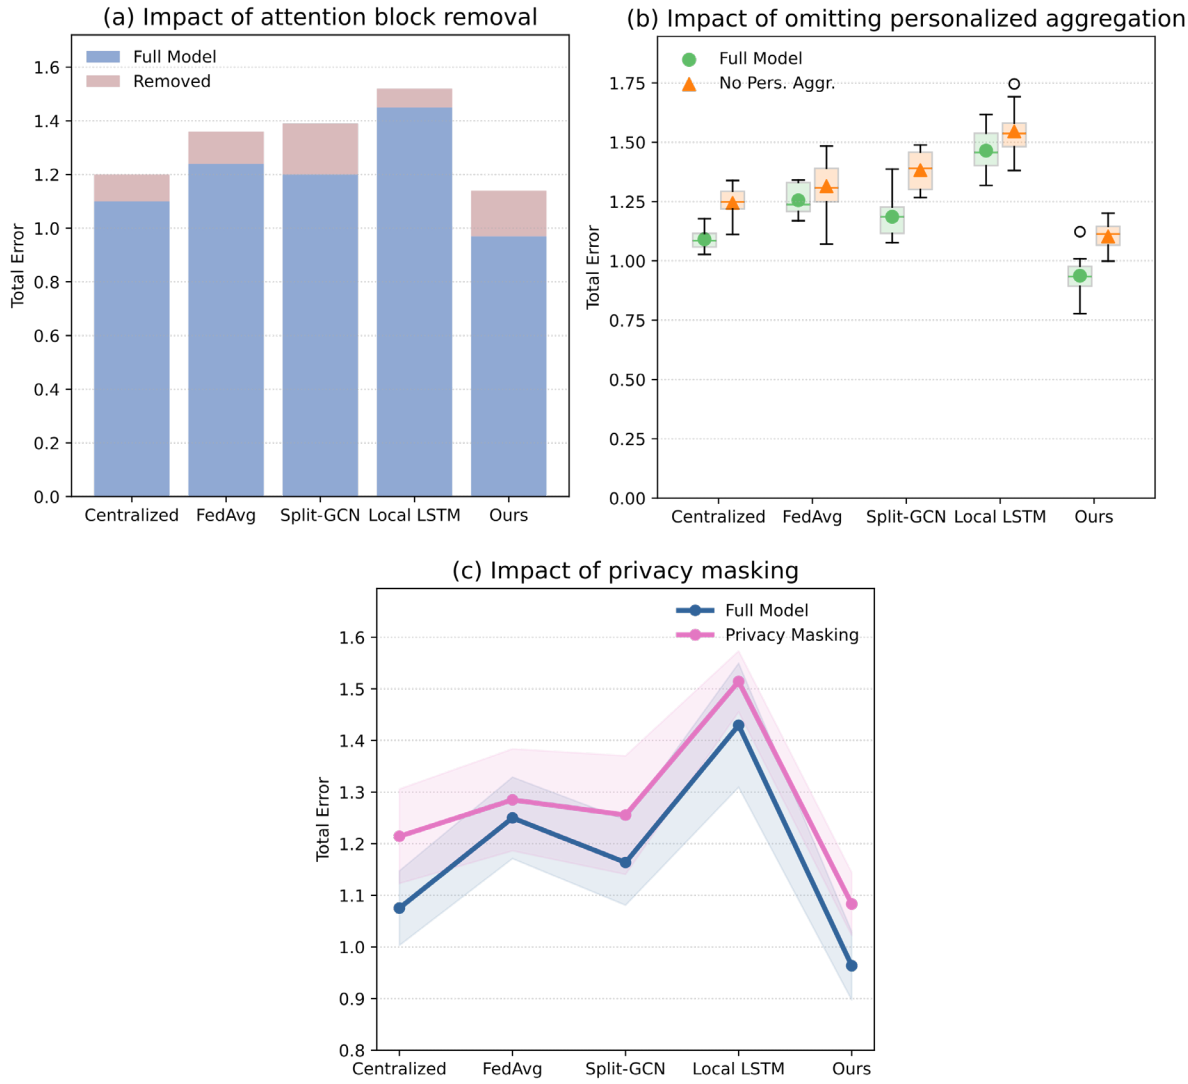

Figure 6. Ablation Study Results (a) Impact of attention block removal (b) Impact of omitting personalized aggregation (c) Impact of privacy masking

Figure 7 illustrates the model's capacity for cross-city transferability and generalization. The design is intrinsically portable and modular, as demonstrated by the constant performance of the prior urban deployments (Figure 7(a)). Adaptive aggregation outperforms fixed-pool methods in Figure 7(b), which demonstrates the resilience to sudden demand surges and uncommon events; Figure 7(c) demonstrates that new sensor installations may also be successfully adopted with little retraining.

https://doi.org/10.64972/jaat.2025v3.230p34e:461-475

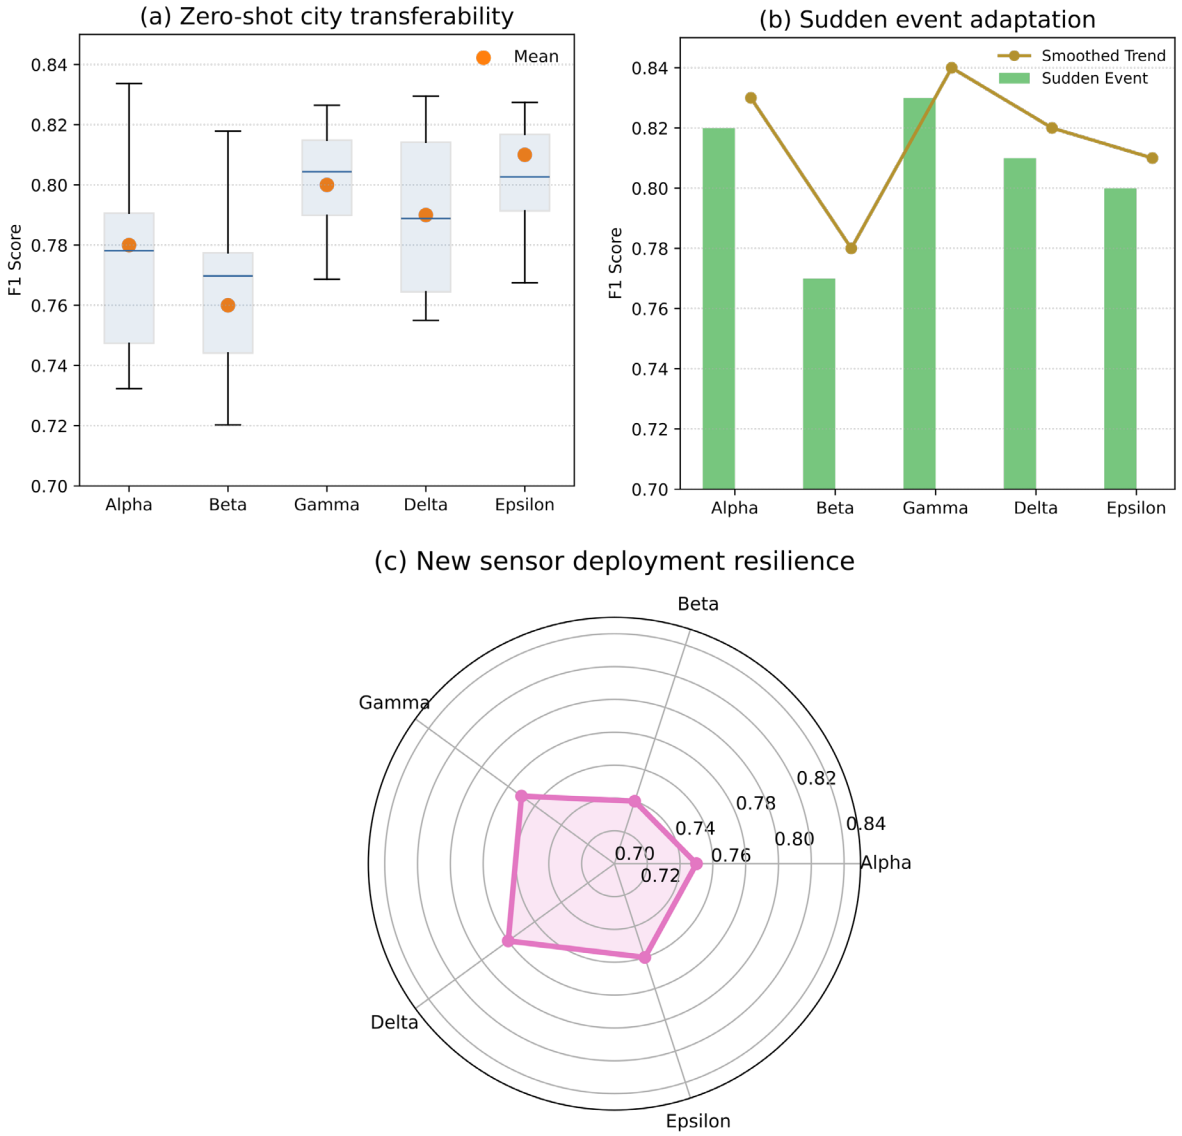

Figure 7. City-level Generalization Tests (a) Zero-shot city transferability (b) Sudden event adaptation (c) New sensor deployment resilience

The federated spatiotemporal framework has demonstrated improved accuracy, communication efficiency, and privacy preservation based on all of the aforementioned trials. Numerous ablation and transfer experiments have confirmed that every design option is appropriate for a range of unpredictable urban traffic situations.

In-depth Discussion

Some particular features of the operating advantages, architectural trade-offs, and future options for development in federated urban traffic prediction have been identified based on the experiment findings.

It is clear from the comparison results above that a strong spatiotemporal feature extraction and attention mechanism contribute to the new model's accuracy. By dynamically shifting traffic concentrations and extremely irregular event patterns, the network may concentrate on the most informative junctions and effectively adapt to various parts of the city. These individual attention modules are context-aware and function well in situations of abrupt demand spikes or localized disruptions; they are not static aggregation or global pooling techniques.

Both system scaling out and privacy protection are now under the purview of the personalized aggregation strategy's core. Assure the stability of real-time prediction while mitigating the negative effects of anomalous data that may arise from inclement weather or network failures by adapting varying degrees of trust and dependability for each node and adjusting the routing weights accordingly. The ablation and city-level transfer

studies best illustrate the first; in these cases, the absence of local adaptation leads to a bigger error variance, poor generalization, and an increased chance of rare, catastrophic events.

This system will be able to concurrently meet the demands of high-performance operation and privacy protection if the communication overhead is investigated. Even after cutting the privacy budget, the network throughput stays constant thanks to an adaptive, event-driven synchronization scheme's decreased redundancy and congestion. A few nodes have been chosen to carry out the task, and the necessary bandwidth will be dynamically anticipated based on the system's real-time data. As a result, the platform will be able to handle numerous city nodes on a big scale without going over the limits of the IT infrastructure or sacrificing real-time performance.

The aforementioned still has certain flaws and unsolved issues. Despite the model's great generalization, it nevertheless performs considerably worse than a centralized, fully supervised upper bound when it comes to prediction accuracy in data-poor peripheral areas or when there is significant network churn. This is caused by the intrinsic privacy-utility trade-off as well as the federated data silos; additional integration of transfer learning and domain adaptation may help to mitigate this restriction. The security guarantees under coordinated cross-node attacks or in situations with widespread non-collaboration have not been properly investigated theoretically and experimentally, despite the fact that personalized aggregation can avoid adversarial updates in practice.

Additionally, there are computational constraints; scheduling these multi-node, varied-privacy-budget, and varied-data-quality hyperparameters in real-time has become extremely challenging due to the recent release of new scheduling and resource allocation criteria. Future directions for expanding both autonomous model selection and adaptive privacy budget management could come from federated reinforcement learning or adaptive meta-learning techniques.

Additionally, the current smart city management system can be integrated with the module and privacy-compliant architecture. A further operational benefit is that just the pertinent employees—rather than the full staff—need to be retrained when adding or removing nodes owing to the development of new buildings, infrastructure, or event sites. Furthermore, the attention mechanism in the aforementioned models provides detailed explanations that might help traffic engineers and city planners enhance public trust in an open-policy setting.

Conclusion

This paper will focus on two issues: the need for high-accuracy urban traffic flow prediction and privacy protection. This framework has achieved notable gains in prediction accuracy, operational robustness, and communication efficiency over the conventional centralized and federated baselines by utilizing adaptive spatiotemporal feature extraction, dynamic attention mechanisms, and personalized aggregation strategies. To demonstrate their ability to generalize in various node contexts, uncommon urban anomalies, and varied sensor setups, experiments have been conducted on multiple city networks.

This paper's ability to simultaneously satisfy the needs for practicality and privacy is one of its innovative qualities. Under strict privacy budgets, context-aware trust scoring and adaptive communication protocols minimize redundancy and preserve prediction accuracy. The model's modularity can facilitate its growth and satisfy the needs of a smart-city infrastructure for node additions and deletions as well as event-driven reconfiguration.

Future research should focus on strengthening the generalization of data-scarce environments through the use of advanced transfer learning, improving adversarial robustness through better measures, and investigating autonomous privacy budget optimization through federated reinforcement learning. To put it briefly, this study has given next-generation, privacy-protected intelligent transportation systems a technical foundation and a strategic viewpoint.

Author Contributions

Sławomir Cyra contributes to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision. Jarosław Bogdan Kalisz contributes to data collection, draft preparation, manuscript editing. All authors have read and agreed with the manuscript before its submission and publication.

Funding

This research received no specific financial support from any funding agency.

Institutional Review Board Statement

Not applicable.

References

[1] Jeyaselvi, M., Satish, J. S., Madhusudhan, K. V., Manikumar, T., Patil, P., & Kalra, G. (2024, October). Enhancing Safety and Reducing Congestion in Smart Cities Using CrowdLOC-S with Graph Attention Networks. In 2024 Global Conference on Communications and Information Technologies (GCCIT) (pp. 1-6). IEEE. https://doi.org/10.1109/GCCIT63234.2024.10862928

[2] Li, Y., Yang, D., & Hu, X. (2020). A differential privacy-based privacy-preserving data publishing algorithm for transit smart card data. Transportation Research Part C: Emerging Technologies, 115, 102634. https://doi.org/10.1016/j.trc.2020.102634

[3] Song, R., Xu, R., Festag, A., Ma, J., & Knoll, A. (2023). FedBEVT: Federated learning bird's eye view perception transformer in road traffic systems. IEEE Transactions on Intelligent Vehicles, 9(1), 958-969. https://doi.org/10.1109/TIV.2023.3310674

[4] Jiang, J. C., Kantarci, B., Oktug, S., & Soyata, T. (2020). Federated learning in smart city sensing: Challenges and opportunities. Sensors, 20(21), 6230. https://doi.org/10.3390/s20216230

[5] Zheng, L., Yang, J., Chen, L., Sun, D., & Liu, W. (2020). Dynamic spatial-temporal feature optimization with ERI big data for short-term traffic flow prediction. Neurocomputing, 412, 339-350. https://doi.org/10.1016/j.neucom.2020.05.038

[6] Fang, S., Prinet, V., Chang, J., Werman, M., Zhang, C., Xiang, S., & Pan, C. (2021). MS-Net: Multi-source spatio-temporal network for traffic flow prediction. IEEE Transactions on Intelligent Transportation Systems, 23(7), 7142-7155. https://doi.org/10.1109/TITS.2021.3067024

[7] Song, J., Zhao, C., Zhong, S., Nielsen, T. A. S., & Prishchepov, A. V. (2019). Mapping spatio-temporal patterns and detecting the factors of traffic congestion with multi-source data fusion and mining techniques. Computers, Environment and Urban Systems, 77, 101364. https://doi.org/10.1016/j.compenvurbsys.2019.101364

[8] Pang, H., & Wang, B. (2020). Privacy-preserving association rule mining using homomorphic encryption in a multikey environment. IEEE Systems Journal, 15(2), 3131-3141. https://doi.org/10.1109/JSYST.2020.3001316

[9] Xu, H., Seng, K. P., Smith, J., & Ang, L. M. (2024). Multi-level split federated learning for large-scale AIoT system based on smart cities. Future Internet, 16(3), 82. https://doi.org/10.3390/fi16030082

[10] Wang, L., Chai, D., Liu, X., Chen, L., & Chen, K. (2021). Exploring the generalizability of spatio-temporal traffic prediction: Meta-modeling and an analytic framework. IEEE Transactions on Knowledge and Data Engineering, 35(4), 3870-3884. https://doi.org/10.1109/TKDE.2021.3130762

[11] Deng, L., Lian, D., Huang, Z., & Chen, E. (2022). Graph convolutional adversarial networks for spatiotemporal anomaly detection. IEEE Transactions on Neural Networks and Learning Systems, 33(6), 2416-2428. https://doi.org/10.1109/TNNLS.2021.3136171

[12] Peng, T., Zhong, W., Wang, G., Luo, E., Yu, S., Liu, Y., ... & Zhang, X. (2024). Privacy-preserving truth discovery based on secure multi-party computation in vehicle-based mobile crowdsensing. IEEE Transactions on Intelligent Transportation Systems, 25(7), 7767-7779. https://doi.org/10.1109/TITS.2024.3350208

[13] Huang, F., Wen, W., Zhang, G., Su, D., & Hsu, L. T. (2023, September). Adaptive multi-sensor integrated navigation system aided by continuous error map from RSU for autonomous vehicles in urban areas. In

2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC) (pp. 5895-5902). IEEE. https://doi.org/10.1109/ITSC57777.2023.10422216

[14] Zhang, S., Li, J., Shi, L., Ding, M., Nguyen, D. C., Tan, W., ... & Han, Z. (2023). Federated learning in intelligent transportation systems: Recent applications and open problems. IEEE Transactions on Intelligent Transportation Systems, 25(5), 3259-3285. https://doi.org/10.1109/TITS.2023.3324962

[15] Zhang, W., Jiang, B., Li, M., & Lin, X. (2022). Privacy-preserving aggregate mobility data release: An information-theoretic deep reinforcement learning approach. IEEE Transactions on Information Forensics and Security, 17, 849-864. https://doi.org/10.1109/TIFS.2022.3152361

[16] Hong, Y., Zhu, H., Shou, T., Wang, Z., Chen, L., Wang, L., ... & Chen, L. (2024). Storm: A spatio-temporal context-aware model for predicting event-triggered abnormal crowd traffic. IEEE Transactions on Intelligent Transportation Systems, 25(10), 13051-13066. https://doi.org/10.1109/TITS.2024.3390185

[17] Jiang, B., Li, J., Yue, G., & Song, H. (2021). Differential privacy for industrial internet of things: Opportunities, applications, and challenges. IEEE Internet of Things Journal, 8(13), 10430-10451. https://doi.org/10.1109/JIOT.2021.3057419

[18] Chen, X., & Liu, G. (2022). Federated deep reinforcement learning-based task offloading and resource allocation for smart cities in a mobile edge network. Sensors, 22(13), 4738. https://doi.org/10.3390/s22134738

[19] Hong, J. P., Park, S., & Choi, W. (2023). Base station dataset-assisted broadband over-the-air aggregation for communication-efficient federated learning. IEEE Transactions on Wireless Communications, 22(11), 7259-7272. https://doi.org/10.1109/TWC.2023.3249252

[20] Mo, J., & Gong, Z. (2022). Cross-city multi-granular adaptive transfer learning for traffic flow prediction. IEEE Transactions on Knowledge and Data Engineering, 35(11), 11246-11258. https://doi.org/10.1109/TKDE.2022.3232185

[21] Sharma, A., Sharma, A., Nikashina, P., Gavrilenko, V., Tselykh, A., Bozhenyuk, A., ... & Meshref, H. (2023). A graph neural network (GNN)-based approach for real-time estimation of traffic speed in sustainable smart cities. Sustainability, 15(15), 11893. https://doi.org/10.3390/su151511893

[22] Wang, C., & Peeta, S. (2024). Incentive mechanism for privacy-preserving collaborative routing using secure multi-party computation and blockchain. Sensors, 24(2), 542. https://doi.org/10.3390/s24020542

[23] Yan, G., Liu, K., Liu, C., & Zhang, J. (2024). Edge intelligence for internet of vehicles: A survey. IEEE Transactions on Consumer Electronics, 70(2), 4858-4877. https://doi.org/10.1109/TCE.2024.3378509

[24] Fang, Y., Shan, Z., & Wang, W. (2021). Modeling and key technologies of a data-driven smart city system. IEEE Access, 9, 91244-91258. https://doi.org/10.1109/ACCESS.2021.3091716

[25] Badu-Marfo, G., Farooq, B., Mensah, D. O., & Al Mallah, R. (2023). An ensemble federated learning framework for privacy-by-design mobility behaviour inference in smart cities. Sustainable Cities and Society, 97, 104703. https://doi.org/10.1016/j.scs.2023.104703

[26] Thirumalaisamy, M., Basheer, S., Selvarajan, S., Althubiti, S. A., Alenezi, F., Srivastava, G., & Lin, J. C. W. (2022). Interaction of secure cloud network and crowd computing for smart city data obfuscation. Sensors, 22(19), 7169. https://doi.org/10.3390/s22197169

[27] Laanaoui, M. D., Lachgar, M., Mohamed, H., Hamid, H., Villar, S. G., & Ashraf, I. (2024). Enhancing urban traffic management through real-time anomaly detection and load balancing. Ieee Access, 12, 63683-63700. https://doi.org/10.1109/ACCESS.2024.3393981

[28] Tahir, M., Qiao, Y., Kanwal, N., Lee, B., & Asghar, M. N. (2023). Real-time event-driven road traffic monitoring system using CCTV video analytics. IEEE Access, 11, 139097-139111. https://doi.org/10.1109/ACCESS.2023.3340144

[29] Jalil, K., Chen, Q., Zahid, M. N., & Jalil, F. (2024, July). Machine learning-driven traffic flow prediction using cloud control big data analysis. In 2024 10th International Conference on Virtual Reality (ICVR) (pp. 396-401). IEEE. https://doi.org/10.1109/ICVR62393.2024.10868318

[30] Shahidinejad, A., Farahbakhsh, F., Ghobaei-Arani, M., Malik, M. H., & Anwar, T. (2021). Context-Aware Multi-User Offloading in Mobile Edge Computing: A Federated Learning-Based Approach: Context-Aware Multi-User Offloading in Mobile Edge Computing: A Federated Learning-based Approach. Journal of Grid Computing, 19(2), 18. https://doi.org/10.1007/s10723-021-09559-x

[31] Hussain, A. H. A., Taher, M. A., Mahmood, O. A., Hammadi, Y. I., Alkanhel, R., Muthanna, A., & Koucheryavy, A. (2023). Urban traffic flow estimation system based on gated recurrent unit deep learning methodology for internet of vehicles. IEEE Access, 11, 58516-58531. https://doi.org/10.1109/ACCESS.2023.3270395

[32] Le, J., Xing, B., Zhang, D., & Qiao, D. (2024). Enhancing Real-Time Traffic Data Sharing: A Differential Privacy-Based Scheme with Spatial Correlation. Mathematics, 12(11), 1722. https://doi.org/10.3390/math12111722

[33] Kapoor, A., & Kumar, D. (2024). Federated learning for urban sensing systems: A comprehensive survey on attacks, defences, incentive mechanisms, and applications. IEEE Communications Surveys & Tutorials, 27(2), 1293-1325. https://doi.org/10.1109/COMST.2024.3434510

[34] Pan, S., Li, P., Yi, C., Zeng, D., Liang, Y. C., & Hu, G. (2020). Edge intelligence empowered urban traffic monitoring: A network tomography perspective. IEEE Transactions on Intelligent Transportation Systems, 22(4), 2198-2211. https://doi.org/10.1109/TITS.2020.3024824

[35] Zhang, Z., Fang, M., Huang, J., & Liu, Y. (2024, June). Poisoning attacks on federated learning-based wireless traffic prediction. In 2024 IFIP Networking Conference (IFIP Networking) (pp. 423-431). IEEE. https://doi.org/10.23919/IFIPNetworking62109.2024.10619763

[36] Musznicki, B., Piechowiak, M., & Zwierzykowski, P. (2022). Modeling real-life urban sensor networks based on open data. Sensors, 22(23), 9264. https://doi.org/10.3390/s22239264

[37] Wang, F., Xu, J., Liu, C., Zhou, R., & Zhao, P. (2021). On prediction of traffic flows in smart cities: a multitask deep learning-based approach. World Wide Web, 24(3), 805-823. https://doi.org/10.1007/s11280-021-00877-4

[38] Kurt, M. N., Yılmaz, Y., Wang, X., & Mosterman, P. J. (2022). Online privacy-preserving data-driven network anomaly detection. IEEE Journal on Selected Areas in Communications, 40(3), 982-998. https://doi.org/10.1109/JSAC.2022.3142302

[39] Liu, L., Zhou, Y., & Xu, J. (2023). A cloud-edge-end collaboration framework for cruising route recommendation of vacant taxis. IEEE Transactions on Mobile Computing, 23(5), 4678-4693. https://doi.org/10.1109/TMC.2023.3294898

[40] Javed, A. R., Ahmed, W., Pandya, S., Maddikunta, P. K. R., Alazab, M., & Gadekallu, T. R. (2023). A survey of explainable artificial intelligence for smart cities. Electronics, 12(4), 1020. https://doi.org/10.3390/electronics12041020