FedTP: A Federated Learning Framework for Traffic Prediction

Baobao Chai[1], Zhongyuan Yu[2,*], Tianqing He[3], Yang Cao[3,*], Qingze He[4], Jianyuan Li[5]
{bbchai_915@sdust.edu.cn, 1099166860@qq.com, sunny.he@std.uestc.edu.cn, caoyoung@uestc.edu.cn, qingzehe@outlook.com, lijianyuan@sd.chinamobile.com}
[1] Department of Computer Science and Engineering, Shandong University of Science and Technology, 266590 Qingdao, China
[2] College of Computer Science and Technology, China University of Petroleum, 266580 Qingdao, China
[3] School of Computer Science and Engineering, University of Electronic Science and Technology of China, 611731 Chengdu, China
[4] College of Engineering and Computer Science, University of California, Davis, CA 95616, United States
[5] Government and Enterprise Solutions Department, China Mobile Shandong Co., Ltd., 250102 Jinan, China
* Corresponding authors

Abstract. Accurate traffic forecasting is essential for intelligent transportation systems, yet most existing approaches rely on centralized training paradigms that raise privacy concerns and often neglect fairness across heterogeneous nodes. Federated learning (FL) provides a privacy-preserving alternative by enabling collaborative training without raw data sharing. However, standard FL algorithms, such as FedAvg and FedProx, primarily optimize global accuracy while exacerbating inter-client disparities, leading to biased service quality in critical regions. To address these challenges, we propose FedTP, a federated learning framework that integrates gradient conflict elimination into the aggregation process. By projecting conflicting client gradients onto a conflict-free subspace, FedTP harmonizes local updates, thereby improving fairness across clients without compromising overall predictive accuracy. Extensive experiments on real-world traffic datasets demonstrate that FedTP achieves competitive accuracy compared to centralized baselines while significantly reducing performance disparities among clients.

Keywords: Federated learning, traffic prediction, fairness, gradient conflict

1 Introduction

In recent years, urban traffic forecasting [1] has become an indispensable task for intelligent transportation systems[2, 3], enabling accurate decision-making for congestion mitigation, route planning [4], and smart city management. Traditional approaches often rely on centralized training paradigms, where data collected from multiple nodes (e.g., sensors, road segments, or edge devices) is aggregated into a central server for model optimization. While these centralized schemes

IIKI 2025, December 18-21, Chengdu, People's Republic of China

DOI 10.4108/eai.18-12-2025.2365258

Fig. 1. Geographical distribution of sensors in the METR-LA dataset.

have demonstrated strong predictive performance, they suffer from two fundamental drawbacks. First, they raise significant privacy concerns, since raw traffic data—potentially containing sensitive spatiotemporal information—must be shared with the central server. Second, they often overlook fairness across heterogeneous nodes: the global model may favor regions with abundant or cleaner data while yielding degraded performance on nodes with limited or noisy data. Such disparities undermine the reliability and equity of intelligent transportation services. Federated learning (FL) has recently emerged as a promising paradigm to address the privacy issue by enabling collaborative training without sharing raw data [5]. In FL, clients (e.g., traffic monitoring nodes) perform local training on their private data, and only model parameters or gradients are shared with the server for aggregation. However, standard FL algorithms such as FedAvg primarily optimize global accuracy, which can exacerbate inter-client disparities in performance. In traffic forecasting, this imbalance is particularly problematic, as underperforming nodes may correspond to critical road segments, leading to systemic inefficiencies and biased service quality. Ensuring fairness across all nodes is therefore essential for practical deployment of FL in traffic prediction.

To more explicitly illustrate the influence of federated learning and fairness on traffic prediction, we take the actual spatial distribution of sensor nodes in the METR dataset as an example, as shown in Fig. 1. In practice, sensors are deployed across different regions, and the data they generate are owned by distinct institutions, which makes centralized training infeasible due to privacy concerns. Furthermore, traffic characteristics vary significantly across regions: sensors located near urban peripheries often capture limited traffic dynamics compared with those on major roads. Consequently, during federated collaborative training, the edge nodes tend to be dominated by the data collected from central road networks, leading to degraded predictive performance at the urban periphery.

To this end, we propose a novel federated learning framework, FedTP, which integrates gradient conflict mitigation into the training process. Inspired by recent advances in fairness-aware federated optimization, our method detects conflicts between gradients of different clients and applies a projection-based adjustment to eliminate such conflicts. This mechanism encourages consistent updates across clients and mitigates the dominance of clients with more favorable data distributions. As a result, FedTP not only preserves data privacy but also improves fairness by balancing predictive performance across all participating nodes.

The main contributions of this work can be summarized as follows.

- We highlight the limitations of centralized traffic prediction approaches and standard FL in terms of privacy and fairness, motivating the need for a new framework.
- We design a federated traffic prediction framework that incorporates gradient conflict elimination to harmonize local updates, ensuring that no client's performance is severely compromised.
- We conduct extensive experiments on real-world traffic datasets, demonstrating that FedTP achieves competitive accuracy compared to centralized baselines while significantly improving fairness across clients.

2 Related Work

In this section, we introduce traffic forecasting and fair federated learning.

2.1 Traditional Traffic Prediction

Traffic forecasting has been extensively studied in centralized settings using spatio-temporal graph neural networks. Models like DCRNN [6], STGCN [7], Graph WaveNet [8], GMAN, and ASTGCN [9] effectively capture temporal dynamics and spatial dependencies among sensors to produce high accuracy predictions. However, these methods require aggregating raw data to a central server, raising privacy concerns and ignoring performance heterogeneity across different sensor nodes.

To address privacy and distribution heterogeneity, federated learning (FL) emerges as an alternative. Algorithms like FedAvg and FedProx enable collaborative training without data sharing, improving robustness in non-IID settings. In traffic prediction, works such as FedAGCN [10] demonstrate application of FL with graph neural networks for networked traffic forecasting. Nevertheless, these FL methods typically optimize global accuracy and lack explicit mechanisms to ensure fairness across nodes.

2.2 Fairness-Aware Federated Learning

Fairness-oriented FL methods aim to mitigate inter-client performance disparity. For instance, q-FFL introduces a re-weighted loss objective that penalizes clients with high loss to enforce more uniform performance [11]. Agnostic FL optimizes for worst-case mixtures of client data distributions. Hierarchical fairness methods and weighted regularization ensure better balance among clients

[12]. These works improve fairness in heterogeneous data regimes but seldom integrate gradient-level conflict resolution tailored to spatio-temporal forecasting[13, 14, 15].

Despite the rapid progress in traffic forecasting and federated learning, prior studies exhibit several limitations. Centralized forecasting models, while achieving high predictive accuracy, inherently assume the availability of all data at a central server, thereby neglecting privacy concerns and overlooking fairness among heterogeneous nodes. On the other hand, existing FL approaches such as FedAvg, FedProx [16], and their extensions have primarily focused on coping with non-IID data distributions and improving average accuracy, but they seldom provide explicit mechanisms to ensure fairness across different participants. Fairness-aware FL methods, including q-FFL and agnostic FL, introduce re-weighting strategies or worst-case optimization to reduce performance disparities; however, these approaches do not address the underlying conflicts among gradients during aggregation, which may result in suboptimal fairness and degraded node-level performance in spatio-temporal traffic prediction tasks.

3 Preliminaries

This section provides an overview of the primary technologies used in this paper, including federated learning and gradient conflict.

3.1 Federated Learning

In federated learning, the ultimate goal is to calculate the optimization problem in Eq. 1, where ℓ denotes the loss function. x_i and y_i represent input and the corresponding true value associated with client i, respectively. $f(x_i;\theta)$ represents the predicted value, and N represents the total number of clients involved in the overall training process.

$$\min_{\theta} \frac{1}{N} \sum_{i=1}^{N} \ell_i \left(f\left(x_i;\theta\right), y_i \right) \tag{1}$$

The above equation is actually designed to minimize client loss. However, in fair federated learning, it is imperative not only to ensure the minimization of the overall global loss but also to guarantee that the losses incurred by each individual client are minimized, as depicted in the following Eq. 2.

$$\min_{\theta} \left(\ell_1(\theta), \ell_2(\theta), ..., \ell_N(\theta) \right) \tag{2}$$

It can be characterized as a multi-objective optimization problem that necessitates the simultaneous optimization of the training loss across all clients [17]. θ_1 is considered fairer than θ_2 if the standard deviation of $\ell(\theta_1)$ is smaller than that of $\ell(\theta_2)$.

3.2 Gradient conflict

In Non-IID scenarios, the presence of heterogeneous data may result in divergent model parameters. A conflict emerges between two clients when their direction divergence exceeds 90°. We

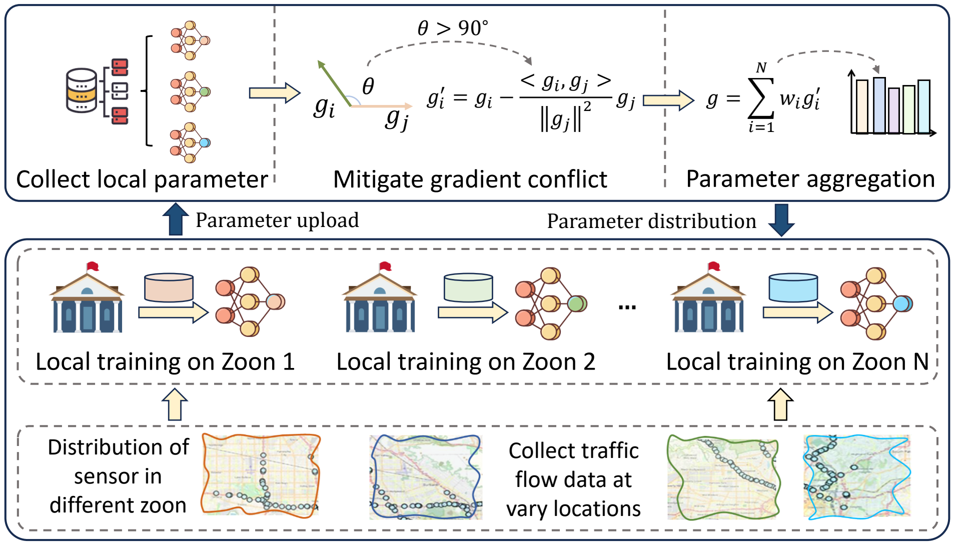

Fig. 2. Framwork of the proposed FedTP scheme.

formalize the concept of gradient direction conflict based on the definition provided in [18]. Additionally, we introduce a formal definition of gradient magnitude conflict.

Definition 1 (Gradient Conflict)**.** *The gradients of two clients are deemed to be in direction conflict iff* $g_i \cdot g_j < 0$*, which represents that the angle between* g_i *and* g_j *exceeds* $90°$.

4 FedTP Scheme

In this section, we will describe in detail the proposed fair federal traffic forecast scheme FedTP.

4.1 Advantages of Fair Federated Learning for Traffic Prediction

Federated learning (FL) enables distributed model training without centralizing raw traffic data, which is crucial for preserving privacy across different network operators and regions. In the context of traffic prediction, fairness plays a vital role because traffic patterns are highly heterogeneous: urban areas typically generate dense and volatile traffic flows, while rural or suburban regions may contribute relatively sparse data. Traditional FL methods such as FedAvg and FedProx often bias the global model towards clients with larger or more stable datasets, leading to imbalanced prediction accuracy.

Fair federated learning addresses this limitation by explicitly incorporating fairness objectives into the optimization process. By balancing global accuracy with inter-client equity, it ensures that clients with limited or noisy traffic data also benefit from collaborative learning. This not only improves the robustness of traffic forecasting models across diverse environments but also supports equitable resource allocation in real-world network management scenarios.

4.2 Framework of the Proposed Scheme

Next, we present the overall framework of FedTP, as shown in Fig. 2. Specifically, data collection nodes in different regions acquire traffic information from their local sensors, and the data are stored independently to avoid mutual interference. Each collection node then performs local training on its regional data to obtain an independent local model. After training, the nodes transmit their model parameters g_i to a central server, where parameter processing and aggregation are conducted. To ensure fairness, we employ a gradient projection method to align gradients and mitigate directional conflicts among clients. The processed parameters are subsequently aggregated using the FedAvg scheme, and the resulting global model g is redistributed to all collection nodes for continued local training. This process is iterated until a predefined number of training rounds or convergence threshold is achieved.

4.3 Conflict Mitigation via Gradient Projection

One of the key challenges in fairness-aware FL is the existence of conflicting client gradients. When the gradients of different clients point in opposing directions, naïve aggregation can hinder convergence and exacerbate fairness disparities. To mitigate this, we introduce a gradient projection mechanism that resolves conflicts before aggregation.

Specifically, given two client gradients g_i and g_j, a conflict is detected when

$$\langle g_i, g_j \rangle < 0, \tag{3}$$

where $\langle g_i, g_j \rangle$ denotes the inner product. In such cases, g_i is projected onto the normal cone of g_j, ensuring that the updated gradient does not negatively interfere. The projected gradient is computed as

$$g_i' = g_i - \frac{\langle g_i, g_j \rangle}{\|g_j\|^2} g_j, \tag{4}$$

where $\|\cdot\|$ is the Euclidean norm.

This projection guarantees that $\langle g_i', g_j \rangle = 0$, effectively removing the conflict while retaining as much useful information as possible from g_i. After conflict resolution, the server performs aggregation over the adjusted gradients:

$$g = \sum_{i=1}^{N} w_i g_i', \tag{5}$$

where N is the total number of participating clients and w_i denotes the aggregation weight.

Through this gradient projection strategy, the federated optimization process avoids destructive interference among clients, leading to more stable convergence and improved fairness in traffic prediction tasks.

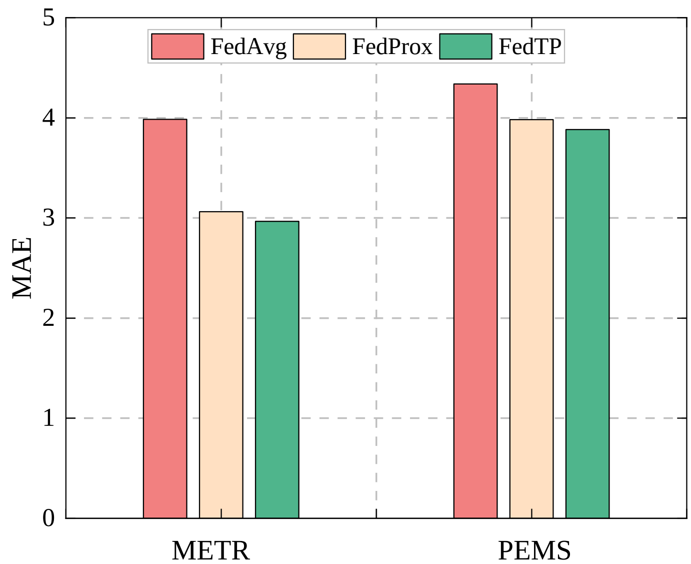

Fig. 3. Performance comparison on two datasets.

5 Performance Evaluation

In this section, we evaluate the performance of the FedTP by performing a series of experiments on two datasets.

5.1 Dataset Introduction

The experiments are conducted on two widely used real-world traffic datasets, METR-LA and PEMS-BAY [19]. The METR-LA dataset was collected from 207 loop detectors located on the highways of Los Angeles County, spanning the period from March to June 2012. The dataset contains time-series traffic measurements, including average traffic speed recorded every 5 minutes. The PEMS-BAY dataset was collected by the California Transportation Agencies (Caltrans) Performance Measurement System (PeMS) in the Bay Area, consisting of 325 sensors with traffic speed readings aggregated every 5 minutes from January to May 2017. Both datasets capture diverse spatial and temporal traffic dynamics, and have been widely adopted as benchmarks for evaluating traffic prediction models.

5.2 Parameter Setting

All experiments were carried out on a laptop with an Intel (R) Core (TM) i9-14900HX CPU @ 2.20 GHz and 32.0 GB of memory. The codes involved in the experiments are all written using PyTorch. We compare FairTP with other classic algorithms , including FedAvg [5] and FedProx

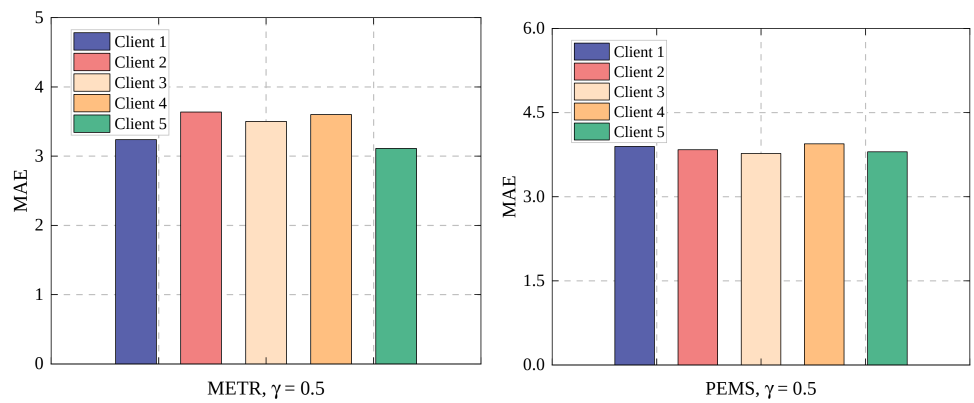

Fig. 4. Performance across five clients on two datasets.

[16] in terms of fairness performance. The dataset is partitioned into five regions, and a Dirichlet distribution [20] is employed to simulate data heterogeneity across clients. We use γ to indicate the degree of data heterogeneity. The smaller the γ, the greater the degree of data heterogeneity. For the predictive model, we adopt the classical GRU (Gated Recurrent Unit) [21] architecture and utilize SGD (Stochastic Gradient Descent) as the default optimizer. The number of global communication rounds is set to 50, with each round consisting of a single local iteration. The learning rate is fixed at 0.001.

5.3 Fairness Performance

To evaluate the predictive performance of the proposed FairTP framework, we conducted experiments on two benchmark datasets. To highlight its fairness, we first report the average prediction accuracy across all clients, as illustrated in Fig. 3. Compared with FedAvg and FedProx, FairTP consistently achieves superior fairness. While FedProx benefits from its regularization term and thus provides a slight improvement over FedAvg, it remains inferior to FairTP in balancing client performance. Furthermore, Fig. 4 presents the prediction results of individual clients. It can be observed that the prediction accuracy across the five clients in both the METR-LA and PEMS-Bay datasets is well balanced, with only minor variations, which are acceptable under heterogeneous data partitions. Overall, FairTP effectively enhances inter-client fairness through gradient de-confliction, thereby mitigating the performance degradation typically observed at urban edge nodes with limited traffic data.

5.4 Impact of γ on Fairness Performance

We examine robustness by varying the degree of data heterogeneity, configuring γ to 0.1, 0.3, 0.5, and 0.7, where smaller γ implies greater heterogeneity across clients. Unless specified, $\gamma = 0.5$ is used as the default setting. While client performance does shift with different partitioning levels,

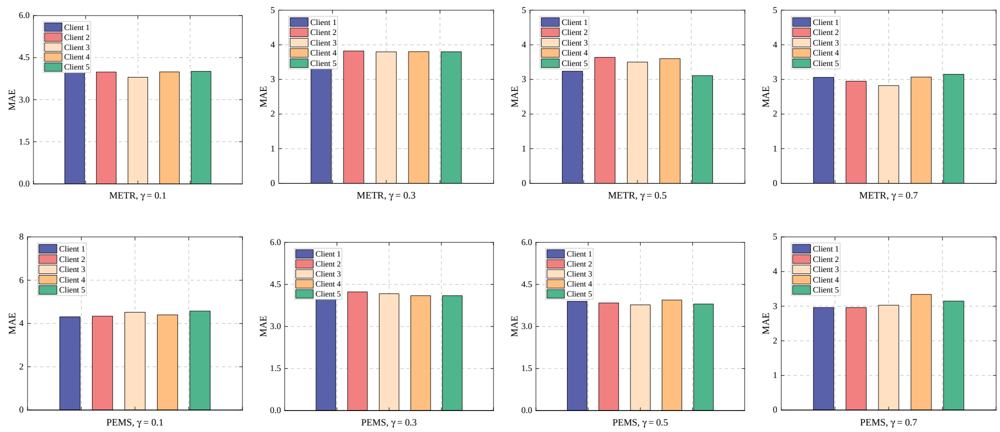

Fig. 5. Impact of γ on fairness performance across five clients on two datasets.

the performance disparities among clients are consistently modest on both datasets, demonstrating strong fairness, as shown in Fig. 5. In addition, as heterogeneity decreases, client prediction error consistently declines. In heterogeneous settings, some clients inevitably achieve lower errors than others due to their data characteristics and random sampling. Overall, even under higher heterogeneity, the differences in prediction error across clients remain modest, demonstrating the fairness of the FedTP scheme.

It is worth noting that the experiments above use a GRU for time-series prediction and do not employ popular spatiotemporal neural networks to model spatial–temporal correlations. Our goal is to evaluate federated learning and traffic prediction under heterogeneous data; all methods therefore share the same network backbone, lending credibility to the comparisons. In future work, we will investigate spatiotemporal neural architectures for traffic prediction and explore additional techniques, such as knowledge graphs and hypergraph networks to further improve federated traffic forecasting.

6 Conclusion

In this work, we proposed FedTP, a federated traffic prediction framework that jointly addresses two major limitations of traditional centralized methods, namely privacy risks and fairness degradation under heterogeneous data distributions. By adopting federated learning, FedTP enables multiple regions to collaboratively train a forecasting model without sharing raw traffic data, thereby protecting institutional privacy and reducing the security risks associated with centralized data aggregation. To further alleviate performance imbalance among clients, we introduced a gradient conflict mitigation strategy based on gradient projection, which reduces optimization inconsistency across local updates and promotes fairer model improvement for different nodes. As a result, not only high-resource nodes but also disadvantaged nodes, especially those located at the network edge, can achieve more robust and stable predictive performance. Extensive experiments on the METR-LA and PEMS-Bay datasets demonstrate that FedTP effectively preserves privacy while significantly improving fairness across heterogeneous nodes without sacrificing overall prediction accuracy. These findings confirm the effectiveness, robustness, and practical value of FedTP and suggest that federated, fairness-aware learning offers a promising direction for real-world traffic prediction in intelligent transportation systems. Future work will explore more adaptive fairness objectives and broader cross-city deployment settings to further enhance generalization and applicability.

Acknowledgments This work was supported by the Major Program of Shandong Provincial Natural Science Foundation for the Fundamental Research under Grant ZR2022ZD03, the National Natural Science Foundation of China under Grants 62272256 and 62202250, the National Science Foundation of Shandong Province under Grant ZR2021QF079, the Talent Cultivation Promotion Program of Computer Science and Technology in Qilu University of Technology (Shandong Academy of Sciences) under Grant 2023PY059, the Pilot Project for Integrated Innovation of Science, Education and Industry of Qilu University of Technology (Shandong Academy of Sciences) under Grant 2022XD001, and the Colleges and Universities 20 Terms Foundation of Jinan City under Grant 202228093.

Declaration on Generative AI

The author(s) have not employed any Generative AI tools.

References

[1] Saini K, Sharma S. Smart Road Traffic Monitoring: Unveiling the Synergy of IoT and AI for Enhanced Urban Mobility. ACM Computing Surveys. 2025;57(11):1-45.

[2] Wang R, Xu Y, Ji J. Traffic Scenario Logic: A Spatial-Temporal Logic for Modeling and Reasoning of Urban Traffic Scenarios. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 39; 2025. p. 15195-202.

[3] He S, Ji J, Lei M. Decomposed Spatio-Temporal Mamba for Long-Term Traffic Prediction. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 39; 2025. p. 11772-80.

[4] Liu H, Chen K, Li Y, Huang Z, Liu M, Ma J. UDMC: Unified Decision-Making and Control Framework for Urban Autonomous Driving With Motion Prediction of Traffic Participants. IEEE Transactions on Intelligent Transportation Systems. 2025;26(5):5856-71.

[5] McMahan B, Moore E, Ramage D, Hampson S, y Arcas BA. Communication-efficient learning of deep networks from decentralized data. In: Artificial intelligence and statistics. PMLR; 2017. p. 1273-82.

[6] Yoon C, Yim S, Yoo S, Jung C, Yeon H, Jang Y. V-DCRNN: Virtual Network-Based Diffusion Convolutional Recurrent Neural Network for Estimating Unobserved Traffic Data. IEEE Transactions on Intelligent Transportation Systems. 2025.

[7] Yu B, Yin H, Zhu Z. Spatio-temporal Graph Convolutional Networks: A Deep Learning Framework for Traffic Forecasting. In: Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI); 2018. .

[8] Wu Z, Pan S, Long G, Jiang J, Zhang C. Graph WaveNet for Deep Spatial-Temporal Graph Modeling. In: Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19. International Joint Conferences on Artificial Intelligence Organization; 2019. p. 1907-13. Available from: https://doi.org/10.24963/ijcai.2019/264.

[9] Guo S, Lin Y, Feng N, Song C, Wan H. Attention based spatial-temporal graph convolutional networks for traffic flow forecasting. In: Proceedings of the AAAI conference on artificial intelligence. vol. 33; 2019. p. 922-9.

[10] Qi T, Chen L, Li G, Li Y, Wang C. FedAGCN: A traffic flow prediction framework based on federated learning and Asynchronous Graph Convolutional Network. Applied Soft Computing. 2023;138:110175.

[11] Li T, Sanjabi M, Beirami A, Smith V. Fair resource allocation in federated learning. arXiv preprint arXiv:190510497. 2019.

[12] Mohri M, Sivek G, Suresh AT. Agnostic federated learning. In: International conference on machine learning. PMLR; 2019. p. 4615-25.

[13] Xia J, Yang Y, Shen J, Wang S, Cao J. FairTP: A Prolonged Fairness Framework for Traffic Prediction. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 39; 2025. p. 26391-9.

[14] Zhou H, Yu W, Wan S, Tong Y, Gu T, Gong C. Traffic pattern sharing for federated traffic flow prediction with personalization. In: 2024 IEEE International Conference on Data Mining (ICDM). IEEE; 2024. p. 639-48.

[15] Zhang C, Dang S, Shihada B, Alouini MS. Dual attention-based federated learning for wireless traffic prediction. In: IEEE INFOCOM 2021-IEEE conference on computer communications. IEEE; 2021. p. 1-10.

[16] Li T, Sahu AK, Zaheer M, Sanjabi M, Talwalkar A, Smith V. Federated optimization in heterogeneous networks. Proceedings of Machine learning and systems. 2020;2:429-50.

[17] Hu Z, Shaloudegi K, Zhang G, Yu Y. Federated Learning Meets Multi-Objective Optimization. IEEE Transactions on Network Science and Engineering. 2022;9(4):2039-51.

[18] Yu T, Kumar S, Gupta A, Levine S, Hausman K, Finn C. Gradient surgery for multi-task learning. Advances in neural information processing systems. 2020;33:5824-36.

[19] Li Y, Yu R, Shahabi C, Liu Y. Diffusion Convolutional Recurrent Neural Network: Data-Driven Traffic Forecasting. In: International Conference on Learning Representations (ICLR '18); 2018. .

[20] Hsu TMH, Qi H, Brown M. Measuring the effects of non-identical data distribution for federated visual classification. arXiv preprint arXiv:190906335. 2019.

[21] Cho K, Van Merriënboer B, Bahdanau D, Bengio Y. On the properties of neural machine translation: Encoder-decoder approaches. arXiv preprint arXiv:14091259. 2014.