

GraFSTNet: Graph-based Frequency SpatioTemporal Network for Cellular Traffic Prediction

Ziyi Li¹, Hui Ma^{1,*}, Fei Xing^{2,*}, Chunjong Zhang³, Ming Yan¹

¹ Xinjiang Key Laboratory of Intelligent Computing and Smart Applications, School of Software, Xinjiang University, China

² College of Geography and Remote Sensing Sciences, Xinjiang University, China

³ Department of AI Convergence Network, Ajou University, Suwon, South Korea

Emails: 2112508422@qq.com, huima@xju.edu.cn,
xingfei@xju.edu.cn, cjz@ajou.ac.kr, yanmingtop@gmail.com

Abstract—With rapid expansion of cellular networks and the proliferation of mobile devices, cellular traffic data exhibits complex temporal dynamics and spatial correlations, posing challenges to accurate traffic prediction. Previous methods often focus predominantly on temporal modeling or depend on predefined spatial topologies, which limits their ability to jointly model spatio-temporal dependencies and effectively capture periodic patterns in cellular traffic. To address these issues, we propose a cellular traffic prediction framework that integrates spatio-temporal modeling with time-frequency analysis. First, we construct a spatial modeling branch to capture inter-cell dependencies through an attention mechanism, minimizing the reliance on predefined topological structures. Second, we build a time-frequency modeling branch to enhance the representation of periodic patterns. Furthermore, we introduce an adaptive-scale LogCosh loss function, which adjusts the error penalty based on traffic magnitude, preventing large errors from dominating the training process and helping the model maintain relatively stable prediction accuracy across different traffic intensities. Experiments on three open-sourced datasets demonstrate that the proposed method achieves prediction performance superior to state-of-the-art approaches.

Index Terms—Cellular traffic prediction; Spatio-temporal modeling; Time-frequency analysis; Spatial attention; Adaptive loss function.

I. INTRODUCTION

With the advent of sixth-generation (6G) mobile communication networks and the widespread adoption of smart mobile devices, global mobile data traffic continues to surge[1]. Moreover, traffic patterns have become increasingly complex and dynamic, exhibiting pronounced temporal periodicity and burstiness, along with significant spatial dependencies arising from user mobility and geographical distribution[2]. Consequently, cellular traffic forecasting has emerged as a critical solution for capturing these complex spatiotemporal dependencies in network time series. Accurate forecasting is essential for efficient and intelligent network management, providing reliable support for resource scheduling, base station energy conservation, load balancing, and network self-optimization, thereby enhancing resource utilization and maintaining[3].

In cellular networks, traffic data exhibits strong spatiotemporal correlations, motivating the adoption of spatiotemporal modeling methods for accurate traffic forecasting[4]. For instance, numerous approaches, including graph-based models

and attention-based models, have been proposed to capture the spatial relationships by propagating and aggregating node features. However, most of these methods focus predominantly on time-domain modeling and overlook frequency-domain characteristics of traffic data, making it challenging to capture periodic variation patterns. To address this, some studies have attempted to incorporate frequency-domain or time-frequency analysis into cellular traffic forecasting. For example, Tang et al.[5] proposed a multi-scale frequency-domain attention network to explicitly extract periodic components. Nevertheless, these frequency-enhanced methods often fail to adequately model complex spatiotemporal correlations[6]. Consequently, developing a joint spatiotemporal-frequency modeling approach that simultaneously captures implicit spatial dependencies and periodic characteristics remains a significant challenge for improving forecasting performance.

Moreover, most existing cellular traffic forecasting methods employ conventional regression loss functions, such as mean absolute error (MAE) and mean squared error (MSE). However, in real-world cellular networks, the distribution of traffic data is often highly non-uniform across different cells and time periods[7]. Standard loss functions like MSE apply a unified penalty on residuals, causing samples with large magnitudes to contribute disproportionately to the overall loss and gradients. Consequently, the optimization process tends to be biased toward high-traffic regions, often at the expense of prediction accuracy in low-traffic areas.

To address the aforementioned challenges, this study proposes a cellular traffic prediction framework named Graph-based Frequency SpatioTemporal Network (GraFSTNet). The model employs two branches to learn spatial dependencies and time-frequency representations, subsequently integrating them through an attention-based fusion stage. Specifically, the spatial branch captures inter-cell dependencies through an attention mechanism rather than relying on a predefined topology. Simultaneously, the time-frequency module extracts features from the frequency domain and learns to weight information across different bands, thereby better reflecting periodic patterns in traffic sequences. Besides, a node-level attention mechanism is introduced to highlight influential regions. The spatial and time-frequency representations of each node are concatenated

to form a joint feature, which is then reweighted by a scalar importance weight and passed to the prediction head. We introduce an adaptive-scale LogCosh loss function to mitigate the impact of large residuals while maintaining a smooth error penalty, thereby improving prediction accuracy across different traffic conditions. The main contributions of this work are summarized as follows:

- We propose a spatio-temporal-frequency fusion framework for cellular traffic prediction. By employing GraphTrans and TFTransformer, we achieve a unified representation of complex spatial dependencies and time-frequency features.
- We design an adaptive-scale LogCosh loss function that dynamically scales residuals based on the magnitude of ground-truth values, thereby maintaining stable prediction performance across different traffic levels.
- We conduct extensive experiments on three public cellular traffic datasets to validate the proposed method. The results demonstrate that our approach achieves forecasting performance superior to state-of-the-art baselines.

The remainder of this paper is organized as follows. Section II reviews related work on cellular traffic prediction. Section III describes the proposed method in detail. Section IV presents the experimental setup and evaluation metrics. Section V discusses the experimental results. Finally, Section VI concludes the paper and outlines future research directions.

II. RELATED WORK

Traffic data in mobile networks exhibit strong spatiotemporal correlations and non-stationarity, which motivates the adoption of spatiotemporal modeling methods for cellular traffic prediction[8]. Representative studies treat graph nodes as graph convolutional networks (GCN), dynamic graph convolutions[9], or graph temporal networks with gated recurrent units (GRU) [10]. By coupling spatial information propagation with time-series modeling, these methods achieve improved performance in traffic forecasting[11].

Building on spatiotemporal modeling, some studies further introduce Transformer architectures to enhance global dependency modeling. For instance, Xu et al. [12] proposed the Spatial-Temporal Transformer (ST-Transformer). Liu et al. [13] introduced STGHTN, which integrates convolution and self-attention through gated hybrid mechanisms. Kong et al. [14] proposed GESTformer by incorporating graph convolution into the Transformer framework. At the same time, other research has explored graph-based modeling methods. For example, Weng et al. [11] presented DDGCRN, which introduces signal decomposition into dynamic graph convolutional networks to extract multi-scale temporal features and improve prediction accuracy. Wang et al. [15] employed graph attention network to capture complex spatiotemporal dependencies from traffic data.

Frequency-domain analysis has been introduced to capture the periodic and non-stationary characteristics of cellular traffic. Zang et al. [16] introduced wavelet decomposition to enhance the modeling of periodic fluctuations. Tian et al. [17] combined

wavelet transformation with multi-model fusion to improve robustness under non-stationary conditions. Tang et al. [5] proposed a frequency-domain MLP-based attention network to extract periodic features and fuse temporal information. Ma et al. [18] proposed MMFNet, which enhances spectral modeling through multi-scale frequency masking and feature fusion. Eldele et al. [19] introduced TSLANet, which incorporates local attention and frequency-domain feature enhancement for time-series modeling.

To jointly model spatiotemporal dependencies and frequency-domain characteristics, a representative line of work in road-traffic forecasting has explored spatiotemporal-frequency fusion. Specifically, Teng et al. [20] proposed a frequency-aware and interactively integrated spatiotemporal graph convolutional network. By introducing a frequency-domain analysis module into the conventional spatiotemporal graph convolution framework, the model separates low-frequency trends and high-frequency fluctuations and employs an interactive dynamic graph convolution mechanism to capture time-varying spatial dependencies. Building on this idea, Guo et al. [21] further incorporated spatial, temporal, and frequency dependencies into a unified framework and proposed a spatiotemporal-frequency joint attention mechanism. It is worth noting that the above studies provide useful insights into spatiotemporal-frequency fusion, but their modeling settings and data characteristics are primarily oriented toward road-traffic scenarios; directly transferring such designs to cellular traffic forecasting still requires adaptation to the data and service characteristics of cellular networks. In such spatiotemporal-frequency fusion studies, the integration of spatiotemporal and frequency-domain features is often implemented with fixed fusion strategies or approximately fixed weight allocation, making it difficult to adaptively adjust the contribution of each feature conditioned on the input; consequently, cross-domain complementary information may not be fully exploited.

III. METHOD

A. Problem Formulation

This study considers a temporal network composed of multiple nodes, where each node represents a cell and contains traffic observations over historical time steps. Specifically, a cellular network consists of N base-station nodes, and the traffic of the i -th node at time step t is denoted as $x_t^{(i)}$. The overall network state at time t is represented as $\mathbf{x}_t = [x_t^{(1)}, x_t^{(2)}, \dots, x_t^{(N)}]^\top \in \mathbb{R}^N$. By stacking the observations of T consecutive time steps in chronological order, the historical input sequence is defined as $\mathbf{X}_{t-T+1:t} = [\mathbf{x}_{t-T+1}, \mathbf{x}_{t-T+2}, \dots, \mathbf{x}_t] \in \mathbb{R}^{N \times T}$. Accordingly, this study aims to learn a parameterized mapping $f(\cdot; \Theta)$ from historical observations to predict traffic over the next H time steps based on the underlying graph structure among nodes.

$$\hat{\mathbf{X}}_{t+1:t+H} = f(\mathbf{X}_{t-T+1:t}; \Theta, \mathcal{G}), \quad (1)$$

where $\mathcal{G}$ the relational structure among nodes.




Fig. 1: Overall architecture of our proposed prediction model.

B. Model Description

1) *Overall Architecture*: This paper proposes a cellular traffic prediction model that integrates spatiotemporal correlation modeling with time–frequency features. The overall architecture of the proposed model is illustrated in Fig. 1. The framework consists of two main modules: (a) **Spatiotemporal Module**: By GraphTrans-based spatiotemporal modeling, the model captures spatial dependencies among cells. This approach does not rely on predefined topological information, instead, it learns implicit spatial correlations directly from the data using self-attention. This enables the model to capture spatial dependencies without requiring an explicit network configuration, making it more flexible and adaptable to various traffic patterns. (b) **Time–Frequency Module**: We employs the TFTransformer-based time-frequency feature extraction module fully leverages the temporal and frequency-domain characteristics of traffic data to capture periodic traffic patterns, thereby enhancing time-series forecasting capabilities and improving prediction accuracy across multiple cells.

(a) Spatiotemporal Module. In multi-cell traffic forecasting, explicit or implicit interactions often exist among cells, such as coupling between geographically adjacent cells and mutual influences from user handovers. Traditional time-series models assume sample independence, making it difficult to capture spatial dependencies. To address this, we introduce a GraphTrans module based on graph-structured self-attention for spatial feature modeling, without relying on predefined topological information. Unlike traditional graph methods that depend on explicit adjacency matrices or topological structures, GraphTrans learns latent spatial dependencies through self-attention, capturing inter-cell dependencies without fixed graph

structures. The module dynamically learns node correlations via self-attention, enabling the model to adapt to varying network configurations and data characteristics.

Specifically, the module first applies linear transformations to obtain the query ($\mathbf{Q}$), key ($\mathbf{K}$), and value ($\mathbf{V}$) representations,

$$\mathbf{Q} = \mathbf{X}_t \mathbf{W}_Q, \quad \mathbf{K} = \mathbf{X}_t \mathbf{W}_K, \quad \mathbf{V} = \mathbf{X}_t \mathbf{W}_V, \quad (2)$$

where $\mathbf{X}_t \in \mathbb{R}^{N \times D}$ denotes the latent node representation of the historical input up to time t , D is the hidden feature dimension, $\mathbf{W}_Q, \mathbf{W}_K, \mathbf{W}_V \in \mathbb{R}^{D \times d_k}$ are learnable parameter matrices, and d_k is the projection dimension.

Then, a self-attention mechanism is employed to compute the dynamic correlation matrix among nodes and aggregate node features according to the attention weights, yielding the updated spatial feature representation as follows,

$$\mathbf{H}_t = \text{Softmax} \left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d_k}} \right) \mathbf{V}. \quad (3)$$

The GraphTrans module adaptively learns latent inter-cell relationships without requiring a predefined adjacency structure, thereby enabling effective modeling of spatial dependencies.

(b) Time–Frequency Module. To enhance the model’s ability to capture periodic fluctuations and abrupt variations in traffic sequences, a parallel time–frequency modeling branch is introduced alongside the GraphTrans-based spatial modeling module. This branch performs frequency-domain decomposition on node features and maps them into multiple discrete frequency bands, yielding corresponding time–frequency representations,

$$\mathbf{H}_t^{(f)} = \mathcal{W}^{(f)}(\mathbf{X}_t), \quad (4)$$

where $\mathbf{X}_t$ denotes the historical traffic sequence up to time t , $\mathcal{W}^{(f)}(\cdot)$ represents the discrete wavelet transform operator for frequency band f , and $\mathbf{H}_t^{(f)}$ is the resulting time–frequency representation.

Subsequently, a frequency attention mechanism is employed to assign weights to different frequency bands, characterizing their relative importance for cellular traffic prediction,

$$\alpha_f = \text{Softmax}(g(\mathbf{H}_t^{(f)})), \quad (5)$$

where $g(\cdot)$ is a learnable scoring function and α_f denotes the normalized frequency-band weight. Finally, the frequency-enhanced representation is obtained by weighted fusion of features across all frequency bands,

$$\mathbf{F}_t = \sum_f \alpha_f \mathbf{H}_t^{(f)}. \quad (6)$$

This module adaptively weights features from different frequency bands through a frequency attention mechanism to obtain frequency-enhanced node representations, which are used to characterize the temporal variation patterns of traffic sequences.

After obtaining the spatial features and frequency-domain features, considering that different nodes contribute unequally to the prediction task, a node-weighted attention mechanism is further introduced to model node-level importance. Specifically, the node representations generated by the spatial modeling branch and the frequency-domain modeling branch are concatenated along the feature dimension to construct a joint node-level representation,

$$\mathbf{U}_t = \text{cat}(\mathbf{H}_t, \mathbf{F}_t) \in \mathbb{R}^{N \times (d_s + d_f)}. \quad (7)$$

Here, $\mathbf{H}_t \in \mathbb{R}^{N \times d_s}$ denotes the node-feature matrix output by the spatial modeling module at time step t , while $\mathbf{F}_t \in \mathbb{R}^{N \times d_f}$ denotes the corresponding frequency-domain feature matrix for each node. N is the number of nodes, and d_s and d_f are the feature dimensions of the spatial and frequency-domain representations, respectively. The operator $\text{cat}(\cdot, \cdot)$ denotes concatenation along the feature dimension. To quantify the importance of different nodes, the joint feature of the i -th node, $\mathbf{U}_t(i, :)$, is fed into a two-layer feed-forward network to generate a node-wise attention weight,

$$\alpha_i = \sigma(\mathbf{W}_2 \text{ReLU}(\mathbf{W}_1 \mathbf{U}_t(i, :))), \quad (8)$$

where $\alpha_i \in (0, 1)$ measures the contribution of the i -th node to the overall prediction task, and $\mathbf{W}_1$ and $\mathbf{W}_2$ are learnable parameter matrices.

The node-weighted feature representation is then updated as,

$$\tilde{\mathbf{U}}_t(i, :) = \alpha_i \cdot \mathbf{U}_t(i, :), \quad (9)$$

where $\mathbf{U}_t$ is the joint feature matrix, $\mathbf{U}_t(i, :)$ is its i -th row, and $\alpha_i \in (0, 1)$ is the node-wise scalar weight. $\tilde{\mathbf{U}}_t$ denotes the reweighted matrix used for prediction. Subsequently, the weighted node features are fed into the prediction function $f(\cdot)$ to produce the traffic prediction for the next time step,

$$\hat{\mathbf{y}}_{t+1} = f(\tilde{\mathbf{U}}_t). \quad (10)$$

2) *Loss function*: To address the unequal penalization of errors across different traffic levels, we introduce an adaptive-scale LogCosh loss, which aims to improve prediction accuracy.

In cellular traffic prediction, target values across cells and time periods can span several orders of magnitude. When MSE, MAE, or the vanilla Log-Cosh are applied directly to residuals on the original scale, samples with larger magnitudes tend to dominate the total loss, whereas errors in low-traffic regions are relatively underweighted. This results in an imbalance in gradient contributions from different traffic intensity ranges and may bias the optimization toward high-traffic scenarios, weakening the model’s fitting ability in low-traffic regions.

To mitigate this issue, we design an adaptive-scale Log-Cosh loss (AS-LogCosh), which introduces a target-magnitude–dependent scaling factor into the residual term to normalize errors under different traffic levels. It is defined as,

$$\mathcal{L}_{\text{AS-LogCosh}} = \frac{1}{n} \sum_{i=1}^n \log \left[\cosh \left(\frac{y_i - \hat{y}_i}{(|y_i| + 1)^\beta + \varepsilon} \right) \right], \quad (11)$$

where y_i and $\hat{y}_i$ denote the ground-truth and predicted values of the i -th sample, respectively, β is a scaling control coefficient, and ε is a small numerical stability constant. The scaling term $(|y_i| + 1)^\beta$ adaptively rescales the residual according to the magnitude of the ground-truth value, so that samples from different traffic intensity ranges are evaluated on a more comparable relative scale, while preserving the smooth behavior of the Log-Cosh loss between small and large residuals. In this way, AS-LogCosh aims to balance the influence of high- and low-traffic samples during training and improve the robustness and overall predictive performance across heterogeneous traffic intensity regimes.

IV. EXPERIMENTAL SETTINGS

A. Datasets

- **MobileNJ¹**: This study adopts a cellular traffic dataset collected from the urban core and surrounding areas of Nanjing. The dataset is divided into a 5×5 grid, where each node represents a group of regional base stations. It captures significant traffic fluctuations during holidays and peak hours, reflecting complex and dynamic traffic patterns in urban environments.
- **Trento²**: The Trento dataset consists of cellular network traffic measurements collected from the Trentino region in Italy. The data are sampled at an hourly resolution and cover both the city center and surrounding areas. The traffic exhibits clear periodicity and regularity, making it suitable for evaluating the generalization performance of prediction models.
- **Milano²**: The Milano dataset contains traffic records from multiple cellular cells located in the urban core area. In this study, the original data are filtered spatially and temporally according to the experimental setup.

¹<https://github.com/MultivariateTimeSeries/Nanjing-cellular-traffic>

²<http://www.openstreet.com/>

Finally, 25 representative cells are selected, and a unified temporal resolution of one hour is adopted to construct the experimental dataset for model evaluation.

B. Baselines

- **LSTM**: LSTM introduces gating mechanisms to mitigate gradient vanishing in RNNs and capture long-term temporal dependencies.
- **GRU**: GRU is a variant of RNN that reduces model complexity by simplifying gate structures and offers competitive performance for short- and medium-term forecasting.
- **Transformer**: Transformer models rely on self-attention to capture long-range dependencies without recurrent structures.
- **GCN**: Graph convolutional networks aggregate features from neighboring nodes to model spatial dependencies on graph-structured data.
- **CNN-LSTM**: CNN-LSTM combines convolutional layers for local feature extraction with LSTM for temporal modeling.
- **DDGCRN**[11]: DDGCRN integrates dynamic diffusion graph convolution with gated recurrent units to jointly model spatial and temporal dependencies.
- **TimeMixer**[22]: TimeMixer is a pure MLP-based model that performs temporal decomposition and multiscale feature mixing.
- **DeseNet**[23]: DeseNet employs dense connections between convolutional layers to enhance feature reuse across temporal scales.
- **ST-Tran**[24]: ST-Tran adopts a Transformer-based encoder-decoder architecture for joint spatiotemporal modeling.
- **OpenCity**[25]: OpenCity combines graph neural networks and Transformers with large-scale pretraining on multi-city traffic data.
- **FISTGCN**[20]: FISTGCN is a frequency-aware spatiotemporal GCN that jointly models spatial correlations and temporal dynamics for traffic forecasting.

C. Metrics

To quantitatively evaluate the performance of the proposed model and the baseline methods, this paper adopts two widely used evaluation metrics including MAE and RMSE,

$$\text{MAE} = \frac{1}{M} \sum_{i=1}^M |y_i - \hat{y}_i|. \quad (12)$$

$$\text{RMSE} = \sqrt{\frac{1}{M} \sum_{i=1}^M (y_i - \hat{y}_i)^2}. \quad (13)$$

D. Experimental Details

In the experimental setup, all datasets are split chronologically, with 80% of the samples used for training, 10% for validation, and the remaining 10% for testing. During

training, the Adam optimizer is employed for all models with an initial learning rate of 0.001. The batch size is set to 32, and the maximum number of training epochs is fixed at 300. Model parameters are iteratively updated by minimizing the prediction error. The main hyperparameters are set as follows: $d_{\text{model}} = 512$, $e_{\text{layers}} = 2$, the number of attention heads is 8, the dropout rate is 0.1, the number of heads in the time-frequency branch is 4, and the wavelet decomposition level is 2. All experiments were conducted on a Linux server equipped with an NVIDIA A30 GPU (24 GB VRAM). The models were implemented using PyTorch 2.1 and Python 3.9.

V. RESULTS AND ANALYSIS

A. Performance Comparison

To comprehensively evaluate the spatiotemporal forecasting performance of the proposed model, we compare it with 11 baseline methods on three public cellular traffic datasets (MobileNJ, Trento, and Milano). The evaluation metrics are MAE and RMSE, and we additionally report the average rank to summarize the overall performance differences among models. The comparison results are presented in Table I, where the best results are highlighted in **bold** and the second-best results are underlined.

Traditional sequence models (e.g., LSTM and GRU) achieve acceptable results on MobileNJ, but their MAE/RMSE increase notably on Trento and Milano, suggesting limited capability in capturing inter-cell spatial dependencies. Transformer-based baselines are effective in modeling temporal dependencies, while graph-based and frequency-aware methods further improve performance by incorporating spatial structure or periodic patterns; however, their gains remain limited when spatial and time-frequency information are not jointly modeled.

GraFSTNet achieves the lowest MAE and RMSE on all three datasets. This advantage mainly comes from: (i) the spatial branch modeling inter-cell dependencies via attention without relying on a predefined topology; (ii) a time-frequency branch that constructs time-frequency representations and aggregates information across frequency bands; and (iii) a fusion stage that jointly integrates and reweights spatial and time-frequency features before prediction.

B. Ablation Analysis

To evaluate the contribution of each module, we conducted an ablation study on three public cellular traffic datasets: MobileNJ, Trento, and Milano. The results are reported in Table II. As observed, the prediction error consistently decreases as components are progressively integrated, validating the efficacy of our design.

Impact of Spatial Module. Introducing the spatial modeling component (GraFSTNet w/ GT + TF) consistently reduces both RMSE and MAE compared with the base variant (GraFSTNet), indicating that modeling inter-cell interactions is beneficial for multi-cell forecasting. By using attention to derive data-driven dependencies among cells, this component alleviates the reliance on a fixed, predefined adjacency and improves the model's ability to capture cross-cell influence patterns,

TABLE I: Performance comparison of different models on the three datasets.

Models	MobileNJ		Trento		Milano		Average Rank
	MAE	RMSE	MAE	RMSE	MAE	RMSE	
LSTM	0.1386	0.2263	8.8619	13.6084	5.8694	10.1104	6.67
GRU	0.1248	0.2057	8.4501	13.5568	6.2281	10.7324	6.33
Transformer	<u>0.1148</u>	<u>0.1832</u>	7.0313	11.0876	5.2510	9.3858	3.67
TimeMixer	0.1198	0.2032	<u>5.1645</u>	<u>8.3607</u>	5.0703	9.1047	3.33
DeseNet	0.2458	0.4418	5.4945	8.5327	5.0557	8.9763	5.33
GCN	0.9667	3.0243	6.0069	9.2083	6.8262	14.2090	9.67
ST_tran	0.1996	0.3683	7.4074	11.6215	5.3988	9.4004	7.17
CNN-LSTM	0.1376	0.2180	6.6562	10.4785	<u>4.9865</u>	<u>8.9473</u>	4.50
DDGCRN	0.6105	1.1453	5.5106	8.8814	7.1688	13.1125	8.67
OpenCity	0.4085	0.8117	8.1638	12.5792	5.4398	9.2348	8.33
FISTGCN	0.2408	0.4112	6.9197	11.7857	8.7762	13.6759	8.50
GraFSTNet	0.1034	0.1682	1.4901	2.4203	1.5719	4.0469	1.00

TABLE II: Performance comparison of different GraFSTNet variants on three datasets.

Models	MobileNJ		Trento		Milano	
	MAE	RMSE	MAE	RMSE	MAE	RMSE
GraFSTNet	0.1352	0.2250	3.3825	4.1064	2.2792	4.6917
GraFSTNet w/ TF	0.1427	0.2421	1.8593	3.0548	2.3794	4.8876
GraFSTNet w/ GT + TF	0.1102	0.1743	1.8583	2.9002	1.9561	4.3566
GraFSTNet w/ GT + TF + AM	0.1063	0.1752	1.6043	2.6278	1.6231	4.2572
GraFSTNet w/ GT + TF + AM + AL	0.1034	0.1682	1.4901	2.4203	1.5719	4.0469

which translates into lower prediction errors (GraFSTNet $\rightarrow$ GraFSTNet w/ GT + TF).

Impact of Time-Frequency Module. Adding the time-frequency module (GraFSTNet w/ TF) further improves performance over the base model (GraFSTNet), suggesting that frequency-aware representations provide information beyond time-domain encoding. In particular, the module introduces band-wise time-frequency features that help the model describe recurring patterns and oscillatory variations in traffic sequences, thereby reducing errors on the three datasets (GraFSTNet $\rightarrow$ GraFSTNet w/ TF).

Impact of Adaptive Fusion. Incorporating the attention-based fusion module (AM) yields additional gains over directly using the concatenated spatial and time-frequency representations (GraFSTNet w/ GT + TF $\rightarrow$ GraFSTNet w/ GT + TF + AM). This indicates that a learnable reweighting during fusion is beneficial for integrating the two representations. Concretely, AM produces node-wise importance weights and uses them to reweight the fused node features before prediction, so that nodes contributing more to the current forecasting target have a larger influence in the final representation. As a result, the fused features become more informative for the prediction head, leading to further reductions in RMSE/MAE.

Overall, the complete model (GraFSTNet w/ GT + TF + AM) achieves the best performance across the three datasets. These results indicate that combining the spatial and time-frequency modules yields additional gains, along with the attention-based

fusion (AM) further improves performance by reweighting node features during integration.

C. Loss Function Comparison

To evaluate the effectiveness of the proposed adaptive-scale LogCosh loss, we conduct comparative experiments on three public cellular traffic datasets, namely MobileNJ, Trento, and Milano, under identical training and evaluation settings. We compare our loss with commonly used regression losses, including MSE, MAE, and the vanilla LogCosh loss. Prediction performance is assessed using RMSE and MAE.

Table III reports the quantitative results. On MobileNJ and Trento, the proposed loss achieves lower RMSE and MAE than MSE, and shows comparable or better performance relative to MAE/LogCosh. This can be attributed to the pronounced magnitude variability and bursty fluctuations of cellular traffic across cells and time. Under the quadratic penalty of MSE, samples with larger magnitudes often contribute disproportionately to the total loss, which may reduce the optimization emphasis on low-traffic regions. In contrast, the proposed loss adaptively rescales error contributions according to traffic magnitude, leading to a more balanced optimization across different traffic levels and thus reducing overall errors.

On Milano, the proposed loss achieves performance comparable to the competing losses. This is likely because the traffic distribution of Milano is relatively stable with a narrower variation range, making the benefit of adaptive rescaling less pronounced. Importantly, no obvious performance degradation

TABLE III: Performance comparison of different loss functions on three datasets.

Datasets	Loss Function	RMSE ↓	MAE ↓
MobileNJ	MSELoss	0.1752	0.1063
	MAELoss	0.1867	0.1052
	LogCosh Loss	0.2429	0.1543
	Adaptive-scale LogCosh Loss	0.1682	0.1034
Trento	MSELoss	2.6278	1.6043
	MAELoss	2.7329	1.5883
	LogCosh Loss	2.9049	1.8928
	Adaptive-scale LogCosh Loss	2.4203	1.4901
Milano	MSELoss	4.2572	1.6231
	MAELoss	4.4275	1.6069
	LogCosh Loss	4.0469	1.5719
	Adaptive-scale LogCosh Loss	4.0469	1.5719

is observed, suggesting that the proposed loss can be applied without harming accuracy when magnitude variations are mild.

D. Visualization of Experimental Results

To further provide an intuitive analysis of the predictive performance of different models, visualization results are presented for several representative cells. By comparing the ground-truth traffic curves with the corresponding model predictions, it becomes easier to observe the models' ability to capture traffic variation trends, peak positions, and fluctuation magnitudes across different time periods.




Fig. 2: Comparison between ground-truth values and predicted values on the MobileNJ dataset




Fig. 3: Comparison between ground-truth values and predicted values on the Trento dataset

To further evaluate the model performance, we compare the proposed method with the best-performing baseline on




Fig. 4: Comparison between ground-truth values and predicted values on the Milano dataset

each dataset. As shown in Figures 2–4, the predicted curves produced by our model align well with the ground-truth curves on MobileNJ, Trento, and Milano. The model tracks the overall trends and periodic fluctuations of traffic series with a close match to the ground truth. During intervals with pronounced fluctuations, the predictions do not exhibit obvious systematic bias. Overall, the proposed model achieves better predictive accuracy than the corresponding baselines across the three datasets.

E. Effect of Model Parameters




Fig. 5: Effect of the encoder depth e_{layers} on the validation sets. Subfigure (a) and (b) report RMSE and MAE, respectively, under $e_{\text{layers}} \in \{1, 2, 3, 4\}$ on Milano, Trento, and MobileNJ.

As shown in Fig. 5 and Fig. 6, we further conduct experiments on the validation sets by varying two key hyper-




Fig. 6: Effect of the representation dimension d_{model} on the validation sets. Subfigure (a) and (b) report RMSE and MAE, respectively, under $d_{\text{model}} \in \{128, 256, 512\}$ on Milano, Trento, and MobileNJ.

parameters, i.e., the encoder depth $e_{\text{layers}} \in \{1, 2, 3, 4\}$ and the representation dimension $d_{\text{model}} \in \{128, 256, 512\}$, while keeping all other settings unchanged. The results show that the overall lowest errors across the three datasets are generally achieved at $e_{\text{layers}} = 2$; increasing the depth to $e_{\text{layers}} = 3, 4$ leads to noticeably higher errors on Trento and Milano. For d_{model} , enlarging the dimension from 128 to 512 yields an overall decrease in RMSE/MAE, and $d_{\text{model}} = 512$ achieves the best or near-best results on all three datasets. Based on these observations, we set $e_{\text{layers}} = 2$ and $d_{\text{model}} = 512$ as the default configuration in subsequent experiments.

VI. CONCLUSION

In this paper, we propose a cellular traffic prediction method that integrates spatiotemporal modeling and frequency-domain feature representation to address the challenges of dynamic spatiotemporal dependencies and insufficient utilization of periodic characteristics in cellular traffic forecasting. Firstly, the proposed approach captures latent correlations among cells, while frequency-domain analysis is employed to complement periodic information in traffic sequences. In addition, an adaptive-scale loss function is incorporated to enhance the performance of the model across different traffic levels. Experimental results demonstrate that the proposed method consistently improves prediction performance across multiple datasets. In future work, we will further incorporate heterogeneous auxiliary information to better capture complex traffic dynamics. Also, we are going to explore lightweight model architectures to improve scalability for large-scale cellular networks.

REFERENCES

- [1] Y. Liu and K. Yang, "Asynchronous decentralized federated anomaly detection for 6G networks," *IEEE Transactions on Cognitive Communications and Networking*, 2025.
- [2] W. Xu, J. Liu, J. Yan, J. Yang, H. Liu, and T. Zhou, "Dynamic spatiotemporal graph wavelet network for traffic flow prediction," *IEEE Internet of Things Journal*, vol. 11, no. 5, pp. 8019–8029, 2024.
- [3] A. A. Hussien, H. Nashaat, and R. F. Abdel-Kader, "Machine learning techniques for spatiotemporal traffic prediction in 5G cellular networks," *Discover Applied Sciences*, vol. 7, no. 10, pp. 1–35, 2025.
- [4] O. G. Manzanilla-Salazar, V. Boutin, H. Mellah, C. Wetté, and B. Sansò, "Framework for vertical performance assessment in very large-scale cellular networks," *IEEE Internet of Things Journal*, vol. 9, no. 14, pp. 12 272–12 284, 2022.
- [5] C. Tang, J. Lu, W. Yang, W. Xing, C. Tang, A. Zou, and J. Guo, "A frequency-domain multilayer perceptron with multiscale attention network for cellular traffic prediction," *Computer Networks*, p. 111593, 2025.
- [6] W. Jiang, "Cellular traffic prediction with machine learning: A survey," *Expert Systems with Applications*, vol. 203, 2022.
- [7] S. Liu, M. He, Z. Wu, P. Lu, and W. Gu, "Spatial-temporal graph neural network traffic prediction based load balancing with reinforcement learning in cellular networks," *Information Fusion*, vol. 103, 2024.
- [8] H. Ma, K. Yang, and Y. Jiao, "Cellular traffic prediction via byzantine-robust asynchronous federated learning," *IEEE Transactions on Network Science and Engineering*, 2025.
- [9] H. Nan, X. Zhu, and J. Ma, "MSTL-GLTP: A global-local decomposition and prediction framework for wireless traffic," *IEEE Internet of Things Journal*, vol. 10, no. 6, pp. 5024–5034, 2023.
- [10] V. Perifanis, N. Pavlidis, R.-A. Koutsiamanis, and P. S. Efraimidis, "Federated learning for 5G base station traffic forecasting," *Computer Networks*, vol. 235, p. 109950, 2023.
- [11] W. Weng, J. Fan, H. Wu, Y. Hu, H. Tian, F. Zhu, and J. Wu, "A decomposition dynamic graph convolutional recurrent network for traffic forecasting," *Pattern Recognition*, vol. 142, 2023.
- [12] M. Xu, W. Dai, C. Liu, X. Gao, W. Lin, G.-J. Qi, and H. Xiong, "Spatial-temporal transformer networks for traffic flow forecasting," *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 8, pp. 4881–4892, 2021.
- [13] J. Liu, Y. Kang, H. Li, H. Wang, and X. Yang, "STGHTN: Spatial-temporal gated hybrid transformer network for traffic flow forecasting," *Applied Intelligence*, vol. 53, no. 10, pp. 12 472–12 488, 2023.
- [14] W. Kong, Y. Ju, S. Zhang, J. Wang, L. Huang, and H. Qu, "Graph enhanced spatial-temporal transformer for traffic flow forecasting," *Applied Soft Computing*, vol. 170, p. 112698, 2025.
- [15] J. Wang, K. Liu, H. Li, Q. Gao, X. Wang, and Y. Gong, "Multiscale temporal features-based hybrid LSTM-GAT for traffic flow prediction," *IEEE Internet of Things Journal*, vol. 12, no. 21, pp. 44 909–44 926, 2025.
- [16] Y. Zang, F. Ni, Z. Feng, S. Cui, and Z. Ding, "Wavelet transform processing for cellular traffic prediction in machine learning networks," in *IEEE China Summit and International Conference on Signal and Information Processing*, 2015, pp. 458–462.
- [17] Z. Tian, "Network traffic prediction method based on wavelet transform and multiple models fusion," *International Journal of Communication Systems*, vol. 33, no. 11, 2020.
- [18] A. Ma, D. Luo, and M. Sha, "MMFNet: Multi-scale frequency masking neural network for time series forecasting," in *Proceedings of the 41st ACM/SIGAPP Symposium on Applied Computing*, 2026, p. 8.
- [19] E. Emadeldeen, R. Mohamed, C. Zhenghua, W. Min, and L. Xiaoli, "TSLANet: Rethinking transformers for time series representation learning," in *Proceedings of the 41st International Conference on Machine Learning*, 2024, pp. 12 409–12 428.
- [20] G. Teng, H. Wu, H. Wu, J. Cao, and M. Zhao, "Frequency-aware and interactive spatial-temporal graph convolutional network for traffic flow prediction," *Applied Sciences*, vol. 15, no. 20, 2025.
- [21] Q. Guo, Q. Tan, J. Tang, and B. Shi, "Unifying spatiotemporal and frequential attention for traffic prediction," *Scientific Reports*, vol. 15, no. 1, p. 953, 2025.
- [22] S. Wang, H. Wu, X. Shi, T. Hu, H. Luo, L. Ma, J. Y. Zhang, and J. ZHOU, "TimeMixer: Decomposable multiscale mixing for time series forecasting," in *The Twelfth International Conference on Learning Representations*, 2024.
- [23] Q. Tang, S. Min, X. Shi, Q. Zhang, and Y. Liu, "DESENet: a bilateral network with detail-enhanced semantic encoder for real-time semantic segmentation," *Measurement Science and Technology*, vol. 36, no. 1, 2024.
- [24] Q. Liu, J. Li, and Z. Lu, "ST-Tran: Spatial-temporal transformer for cellular traffic prediction," *IEEE Communications Letters*, vol. 25, no. 10, pp. 3325–3329, 2021.
- [25] Z. Li, L. Xia, L. Shi, Y. Xu, D. Yin, and C. Huang, "Open spatiotemporal foundation models for traffic prediction," *ACM Transactions on Intelligent Systems and Technology*, 2025.