

Frequency Enhanced Pre-training for Cross-city Few-shot Traffic Forecasting

Zhanyu Liu, Jianrong Ding, and Guanjie Zheng(✉)

Shanghai Jiao Tong University, Shanghai, China
{zhylu00, rafaelding, gjzheng}@sjtu.edu.cn

Abstract. The field of Intelligent Transportation Systems (ITS) relies on accurate traffic forecasting to enable various downstream applications. However, developing cities often face challenges in collecting sufficient training traffic data due to limited resources and outdated infrastructure. Recognizing this obstacle, the concept of cross-city few-shot forecasting has emerged as a viable approach. While previous cross-city few-shot forecasting methods ignore the frequency similarity between cities, we have made an observation that the traffic data is more similar in the frequency domain between cities. Based on this fact, we propose a **Frequency Enhanced Pre-training Framework for Cross-city Few-shot Forecasting (FEPCross)**. FEPCross has a pre-training stage and a fine-tuning stage. In the pre-training stage, we propose a novel Cross-Domain Spatial-Temporal Encoder that incorporates the information of the time and frequency domain and trains it with self-supervised tasks encompassing reconstruction and contrastive objectives. In the fine-tuning stage, we design modules to enrich training samples and maintain a momentum-updated graph structure, thereby mitigating the risk of overfitting to the few-shot training data. Empirical evaluations performed on real-world traffic datasets validate the exceptional efficacy of FEPCross, outperforming existing approaches of diverse categories and demonstrating characteristics that foster the progress of cross-city few-shot forecasting.

1 Introduction

The field of Intelligent Transportation Systems (ITS) has recognized the significance of traffic forecasting as a critical service. By effectively employing historical traffic data to precisely forecast future traffic, it becomes feasible to enable diverse downstream applications, including traffic signal control [38] and traffic tolling [37]. In order to gather the necessary traffic data for training a traffic forecasting model, the utilization of traffic sensors [20] or vehicle devices [6] is crucial. The devices are well deployed in the developed cities and the data collection is easy. However, when it comes to developing cities, one of the main obstacles is the lack of adequate infrastructure to support the deployment of these devices. Developing cities often have limited resources and outdated technology systems, making it difficult to establish a reliable network infrastructure for data collection. Consequently, the traffic data in developing cities could be limited, and thus training a deep traffic forecasting model is hard.

✉: Corresponding Author




Fig. 1: (a) The mean cosine similarity of one-week traffic speed data within the time and frequency domains across four cities. We can observe a higher level of similarity in the frequency domain. (b) An illustrative instance showcasing a pair of data in both the time domain and frequency domain reflects that, despite containing identical data, the frequency domain exhibits a significantly higher level of cosine similarity.

To address the issue of limited data in developing cities and enhance the efficiency of traffic forecasting services, the concept of cross-city few-shot forecasting has emerged as a viable solution. This approach involves learning from cities with abundant data and transferring that knowledge to cities with limited data. In recent years, many works have focused on this problem. RegionTrans [36] and CrossTReS [15] identify the inherent regional correlation between the source city and the target city. MetaST [41] learns a global memory which is then queried by the target region. STrans-GAN [46] generates future traffic based on traffic demand with a GAN-based model. However, these methods depend on auxiliary data, such as event information, for the transfer of knowledge. This reliance on auxiliary data poses a challenge when such data is not available in the city with limited data resources. ST-MetaNet [32] and ST-GFSL [26] utilize learned meta-knowledge to generate the parameters of the spatial-temporal neural network. TransGTR [16] aims to learn the city-agnostic node features and the transferable graph structure between cities. TPB [24] aims to construct a traffic pattern bank that contains the meta-knowledge across cities.

Previous methods have predominantly concentrated on capturing temporal relationships between different cities. However, they neglect the potential knowledge that can be derived and shared from the frequency domain. In fact, as shown in Fig. 1, the traffic data between cities is more similar in the frequency domain. We calculate the mean cosine similarity of aligned one-week traffic data of different cities in the time domain and the frequency domain, which is shown in Fig. 1(a). We could observe that the traffic data exhibits notable similarity in the frequency domain across different cities. Fig. 1(b) shows an example that even the same pair of traffic data exhibits substantial similarities in the frequency domain and they are both prominent in the low-frequency range. The identification of these frequency domain patterns implies that modeling traffic data with a focus on the frequency domain can lead to a better understanding and capture of the underlying dynamics. By leveraging these patterns into the pre-training process in the data-rich source cities, it becomes possible to develop effective and accurate models for cross-city few-shot traffic forecasting.

Consequently, we propose a **Frequency Enhanced Pre-training** framework for **Cross-city** few-shot traffic forecasting, abbreviated as **FEPCross**. It contains two stages, the pre-training stage and the fine-tuning stage. **For the pre-training stage**, we design a novel frequency-enhanced pre-training framework on the data-rich source city. This framework incorporates the information from the time, amplitude, and phase domains through a Cross-Domain Spatial-Temporal Encoder. By taking masked data from these three domains as input, the framework aims to reconstruct the missing information as a self-supervised task. A contrastive loss is then added to guarantee the quality of the learned space. **For the fine-tuning stage**, we design two modules based on the pre-trained Cross-Domain Spatial-Temporal Encoder to improve the performance of the few-shot fine-tuning process. On one hand, we mask and reconstruct the few-shot training data to create more augmented training samples. On the other hand, we maintain a momentum-updated graph structure of the target city. These two modules not only improve the effectiveness but also enhance the resistance to the over-fitting of the fine-tuned model.

In summary, the main contributions are as follows.

- We delve into the cross-city few-shot traffic forecasting task and find that the frequency domain pattern exhibits a higher degree of cross-city sharability. The potential of transferring knowledge from the frequency domain has been disregarded in prior research endeavors.
- We propose a Frequency Enhanced Pre-training framework (FEPCross) specifically designed for cross-city few-shot traffic forecasting. It incorporates information on various domains and contains modules to mitigate the risk of overfitting to the few-shot training data.
- We demonstrate the effectiveness of the FEPCross framework through extensive experiments on real-world traffic datasets from multiple cities. The results demonstrate that FEPCross achieves superior performances over the state-of-the-art baselines and exhibits distinctive characteristics such as prioritizing the modeling of low-frequency components, further contributing to its exceptional performance.

2 Related Work

Traffic Forecasting The practical applications of traffic forecasting have garnered significant attention. Some work utilizes traditional statistical methods such as Kalman Filter [22], SVM [30], probabilistic model [1], or simulation [21]. Some work combines modules such as GRU and GNN to model the spatial-temporal correlation within the traffic data, such as DCRNN [20], STGCN [42], GSTNet [9], LSGCN [12], STFGNN [19], Frigate [10], HIEST [27], FDTI [23]. To better capture the varying spatial-temporal correlations between traffic nodes, some work such as AGCRN [2,7], Graph Wavenet [40], GMAN [47], D2STGNN [35], ST-WA [5], DSTAGNN [18], TrendGCN [14], MC-STL [45] utilize adaptive methods such as node embedding to reconstruct the adaptive adjacent matrix and incorporate temporal long-term relations for improved predictions. Some works utilize techniques to model the special characteristics of the traffic. STG-NCDE [4], STDEN [13] model the traffic based on ordinary differential equation




Fig. 2: Overall Framework of FEPCross.

(ODE). FOGS [33] predict the first-order gradient of traffic flow. STGBN [8] utilizes gradient boosting to enhance the model. STEP [34] adapts MAE and proposes a pipeline to pre-train a model. Nevertheless, the aforementioned methods mainly focus on predicting traffic within a single city. These methods would encounter difficulties such as distribution bias, the tendency to over-fit the single-city data, and the incompatibility of the node embedding technique in cross-city scenarios.

Cross-city Few-shot Forecasting To tackle the data scarcity problem and use the available knowledge of the data-rich cities effectively, several methods have been developed in the field of cross-city few-shot forecasting. Floral [39] learns from the multimodal data from the source cities and transfers the knowledge to the target city for the air quality classification problem. RegionTrans [36], CrossTReS [15], and MetaST [41] learn the region correlation or the memory bank between the source cities to the target city. However, these approaches utilize multimodal auxiliary information, which is not accessible in the data-scarce city. ST-MetaNet [32] and ST-GFSL [26] generate the parameters of spatial-temporal neural networks according to the learned meta-knowledge. CityTrans [31] uses domain adversarial training to transfer knowledge across cities. UniST [43] uses prompt tuning to learn the spatial-temporal shared across cities. TransGTR [16] learns city-agnostic features and transferable graphs from one city to another city. TPB [24] and MTPB [25] construct the traffic pattern bank from multiple source cities and the data of the target city could query the bank to get the meta-knowledge to reconstruct the self-expressive graph structure [17] and forecast the traffic. However, these methods ignore the spectral similarity between the time series traffic data and thus get suboptimal results.

3 Preliminary

Traffic Spatial-Temporal Graph: The traffic spatial-temporal graph is the data structure that processes the traffic data. It could be denoted as $\mathcal{G} = (\mathbf{V}, \mathbf{A}, \mathbf{X})$. Here, \mathbf{V} is the set of the nodes, and $N = |\mathbf{V}|$ is the number of nodes. $\mathbf{A} \in \mathbb{R}^{N \times N}$ is the adjacency

matrix that describes the correlation between nodes. $\mathbf{X} \in \mathbb{R}^{N \times T \times C}$ is the traffic data of N nodes, with time steps T and channel C .

Traffic Forecasting: The Traffic forecasting problem is to utilize the historical traffic to predict future traffic. Formally, given the historical traffic of T_h steps, we aim to learn a model $f(\cdot)$ to forecast the future traffic of T_f steps, which could be formulated as follows

$$[\mathbf{X}_{t-T_h+1}, \dots, \mathbf{X}_t] \xrightarrow{f(\cdot)} [\mathbf{X}_{t+1}, \dots, \mathbf{X}_{t+T_f}]. \quad (1)$$

Cross-city Few-Shot Traffic Forecasting: Given a data-rich source city \mathcal{G}^{source} and a data-scarce target city \mathcal{G}^{target} , the goal of cross-city few-shot traffic forecasting is to pre-train a model in \mathcal{G}^{source} and fine-tune the model in \mathcal{G}^{target} . The fine-tuned model is expected to conduct accurate traffic forecasting on the future data of \mathcal{G}^{target} .

4 Method

In this section, we present the FEPCross framework, which comprises two stages: pre-training and fine-tuning as shown in Fig. 2. In the pre-training stage, we aim to learn robust traffic knowledge by the Cross-Domain Spatial-Temporal Encoder from the source city. The encoder could capture the correlation across domains and spatial neighbors, and then reconstruct the masked input of three domains. To encourage the encoder to learn a representation space that is both compact and discriminative, we incorporate a contrastive loss during the pre-training process. In the fine-tuning stage, we address the challenges associated with limited data by enriching the few-shot training dataset of the target city using the pre-trained encoder. This is achieved by applying a masking and reconstruction technique. The graph structure is momentum-updated by the knowledge of the pre-trained encoder. These two novel modules are designed to avoid over-fitting and improve the performance of the fine-tuned model. Finally, a Spatial-Temporal model (STmodel) aggregates the short-term traffic knowledge and makes predictions.

4.1 Pre-Training

In this section, we will introduce the novel Cross-Domain Spatial-Temporal Encoder that captures the pattern of the traffic data across domains and spatial neighbors. To train the Cross-Domain Spatial-Temporal Encoder, we employ two essential self-supervision tasks: reconstruction loss and contrastive loss. The reconstruction loss aims to reconstruct the original masked input data from the encoded representations. This task encourages the encoder to learn meaningful and informative representations that accurately capture the input data structure.

Input: Fig. 1 in the introduction section visually demonstrates that the frequency domain exhibits a higher similarity across different cities. Leveraging this observation, we adopt a Fourier Transform on the input traffic time series to extract the amplitude and phase domains. This transformation allows us to capture the frequency domain characteristics of the data. Then, we segment the series into patches and randomly mask most of them. Formally, by taking the time domain T_i as the example, the input time series of node i , i.e., $\mathbf{X}_{i,0:T_h}^{T_i}$ is separate to P patches as $\mathbf{X}_{i,0:T_h}^{T_i} =$

$\{\mathbf{S}_{i,0}^{\text{Ti}}, \mathbf{S}_{i,1}^{\text{Ti}}, \dots, \mathbf{S}_{i,P}^{\text{Ti}}\}$ and there is a masked vector \mathbf{M} where $M_{i,k}^{\text{Ti}} = 1$ indicates $\mathbf{S}_{i,k}^{\text{Ti}}$ is masked. The masked patch series of three domains is the input of the encoder similar to previous research [29,11].

Cross-Domain Spatial-Temporal Encoder: The Cross-Domain Spatial-Temporal Encoder aims to aggregate the information of the traffic data across domains and across neighbors. Firstly, The unmasked part of the input is fed into a Transformer block. Then, a learnable masked token S_{mt} fills the place where the data is masked, i.e., $M_{i,k} = 1$.

$$\mathbf{S}_i^{\text{Ti}} = \text{TSLayer}(\{\mathbf{S}_{i,0}^{\text{Ti}}, \mathbf{S}_{mt}^{\text{Ti}}, \dots, \mathbf{S}_{i,P}^{\text{Ti}}\}) \quad (2)$$

$$\mathbf{S}_i^{\text{Am}} = \text{TSLayer}(\{\mathbf{S}_{i,0}^{\text{Am}}, \mathbf{S}_{mt}^{\text{Am}}, \dots, \mathbf{S}_{i,P}^{\text{Am}}\}) \quad (3)$$

$$\mathbf{S}_i^{\text{Ph}} = \text{TSLayer}(\{\mathbf{S}_{i,0}^{\text{Ph}}, \mathbf{S}_{mt}^{\text{Ph}}, \dots, \mathbf{S}_{i,P}^{\text{Ph}}\}) \quad (4)$$

Here, $\mathbf{S}_{0,i}^{\text{Ti}}$ indicates the patch data of node i at time 0 and \mathbf{S}^{Ti} indicates the patch series in the time domain, and $\mathbf{S}^{\text{Am}}, \mathbf{S}^{\text{Ph}}$ indicates amplitude and phase domain respectively. Next, the concatenated data of the time, amplitude, and phase domains is then fed to a Cross-Domain Aggregator to aggregate the information of each domain. We use a Transformer layer ($\text{TSLayer}(\cdot)$) as the Cross-Domain Aggregator here.

$$\mathbf{H}_i^{\text{Ti}}, \mathbf{H}_i^{\text{Am}}, \mathbf{H}_i^{\text{Ph}} = \text{TSLayer}(\text{Concat}\{\mathbf{S}_i^{\text{Ti}}, \mathbf{S}_i^{\text{Am}}, \mathbf{S}_i^{\text{Ph}}\}) \quad (5)$$

After the information of the three domains is aggregated, we expect the model to further integrate the knowledge of the neighboring nodes. To achieve this, a Graph Neural Network ($\text{GNN}(\cdot)$) encoder is incorporated into the architecture as a Cross-Space Aggregator, dedicated to aggregating information for each domain separately.

$$\mathbf{H}_i^{\text{Ti}} = \text{GNN}\{\mathbf{H}_j^{\text{Ti}} | j \in \mathcal{N}_i\} \quad (6)$$

$$\mathbf{H}_i^{\text{Am}} = \text{GNN}\{\mathbf{H}_j^{\text{Am}} | j \in \mathcal{N}_i\} \quad (7)$$

$$\mathbf{H}_i^{\text{Ph}} = \text{GNN}\{\mathbf{H}_j^{\text{Ph}} | j \in \mathcal{N}_i\} \quad (8)$$

Here, \mathcal{N}_i denotes the set of neighboring nodes of node i in the raw graph \mathbf{A} . The GNN layer is applied independently to each domain, enabling the model to capture the contextual information from neighboring nodes in the respective time, amplitude, and phase domains. Finally, we feed the data into another Cross-Domain Aggregator again.

$$\hat{\mathbf{H}}_i^{\text{Ti}}, \hat{\mathbf{H}}_i^{\text{Am}}, \hat{\mathbf{H}}_i^{\text{Ph}} = \text{TSLayer}(\text{Concat}\{\mathbf{H}_i^{\text{Ti}}, \mathbf{H}_i^{\text{Am}}, \mathbf{H}_i^{\text{Ph}}\}) \quad (9)$$

Here $\hat{\mathbf{H}}_i^{\text{Ti}}, \hat{\mathbf{H}}_i^{\text{Am}}, \hat{\mathbf{H}}_i^{\text{Ph}}$ indicates the output embedding of Cross-Domain Spatial-Temporal Encoder in the time, amplitude, phase domains respectively.

Reconstruction Module Once we obtain the output embeddings from the encoder, we can reconstruct the masked input patches. This reconstruction process involves applying linear transformations to the output embeddings.

$$\hat{\mathbf{S}}_{i,j}^{\text{Ti}} = \text{Linear}(\hat{\mathbf{H}}_{i,j}^{\text{Ti}}) \quad (10)$$

$$\hat{\mathbf{S}}_{i,j}^{\text{Am}} = \text{Linear}(\hat{\mathbf{H}}_{i,j}^{\text{Am}}) \quad (11)$$

$$\hat{\mathbf{S}}_{i,j}^{\text{Ph}} = \text{Linear}(\hat{\mathbf{H}}_{i,j}^{\text{Ph}}) \quad (12)$$

Then, the reconstruction loss is the Mean Square Error (MSE) of the masked patches.

$$\mathcal{L}_{re} = \sum_{i,j} \mathbf{M}_{ij}^{\text{Ti}} (\hat{\mathbf{S}}_{i,j}^{\text{Ti}} - \mathbf{S}_{i,j}^{\text{Ti}})^2 + \sum_{i,j} \mathbf{M}_{ij}^{\text{Am}} (\hat{\mathbf{S}}_{i,j}^{\text{Am}} - \mathbf{S}_{i,j}^{\text{Am}})^2 + \sum_{i,j} \mathbf{M}_{ij}^{\text{Ph}} (\hat{\mathbf{S}}_{i,j}^{\text{Ph}} - \mathbf{S}_{i,j}^{\text{Ph}})^2 \quad (13)$$

By optimizing the encoder to minimize the reconstruction loss, we encourage the model to learn representations that capture the essential information needed to reconstruct the masked patches accurately.

Contrastive Module To further improve the representation learning and enhance the expressiveness of the learned embedding space, a contrastive loss is introduced. We introduce a novel approach to constructing augmented data that effectively utilizes the frequency domain information. Recognizing the significant informational content present in the amplitude domain of traffic data, we randomly switch the amplitude domain within a given batch of data, thereby producing augmented samples.

$$\tilde{\mathbf{S}}_i^{\text{Ti}} = \mathbf{S}_i^{\text{Ti}}, \tilde{\mathbf{S}}_i^{\text{Ph}} = \mathbf{S}_i^{\text{Ph}} \quad (14)$$

$$\tilde{\mathbf{S}}_i^{\text{Am}} = \mathbf{S}_j^{\text{Am}} \quad i, j \in \text{SameBatch} \quad (15)$$

Here $\tilde{\mathbf{S}}_i^{\text{Ti}}, \tilde{\mathbf{S}}_i^{\text{Ph}}, \tilde{\mathbf{S}}_i^{\text{Am}}$ is the augmented data in the three domains. Then, we mask the data and feed the data into the Cross-Domain Spatial-Temporal Encoder $CDEnc(\cdot)$ to get the augmented embedding.

$$\tilde{\mathbf{H}}_i^{\text{Ti}}, \tilde{\mathbf{H}}_i^{\text{Am}}, \tilde{\mathbf{H}}_i^{\text{Ph}} = CDEnc(\tilde{\mathbf{S}}_i^{\text{Ti}}, \tilde{\mathbf{S}}_i^{\text{Am}}, \tilde{\mathbf{S}}_i^{\text{Ph}}) \quad (16)$$

Subsequently, a linear layer is introduced as a Domain Embedding Aggregator, responsible for merging the embeddings from the three domains.

$$\tilde{\mathbf{H}}_i = \text{Linear}(\text{Concat}\{\tilde{\mathbf{H}}_i^{\text{Ti}}, \tilde{\mathbf{H}}_i^{\text{Am}}, \tilde{\mathbf{H}}_i^{\text{Ph}}\}) \quad (17)$$

Similarly, the embedding of the original data $\hat{\mathbf{H}}_i$ is generated.

$$\hat{\mathbf{H}}_i = \text{Linear}(\text{Concat}\{\hat{\mathbf{H}}_i^{\text{Ti}}, \hat{\mathbf{H}}_i^{\text{Am}}, \hat{\mathbf{H}}_i^{\text{Ph}}\}) \quad (18)$$

Finally, we want the data sample and its augmented pair to have similar embeddings. By doing this, the encoder is more robust and could capture the essential dynamics of the frequency domain of the data. We employ the NT-Xent Loss [3] here to optimize the encoder.

$$\mathcal{L}_{con} = - \sum_i^N \log \frac{e^{\text{sim}(\hat{\mathbf{H}}_i, \tilde{\mathbf{H}}_i)}}{\sum_j e^{\text{sim}(\hat{\mathbf{H}}_i, \tilde{\mathbf{H}}_j)}} \quad (19)$$

Here sim is the cosine similarity and the negative samples j are randomly chosen 10% from the N nodes.

The total loss of the pre-training stage is the sum of the two losses and it aims to update the Cross-Domain Spatial-Temporal Encoder.

$$\mathcal{L} = \mathcal{L}_{re} + \alpha \mathcal{L}_{con} \quad (20)$$

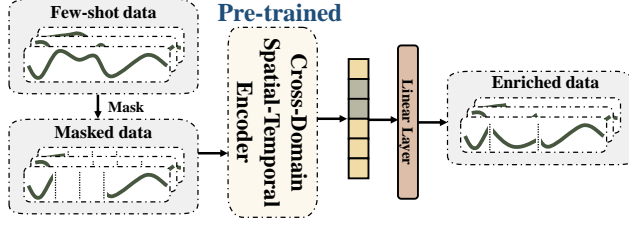


Fig. 3: Illustration of Training Data Enriching Module.

4.2 Fine-Tuning

In the fine-tuning stage, we aim to utilize the pre-trained Cross-Domain Spatial-Temporal Encoder to help forecast the traffic with few-shot training data in the target city. Firstly, we employ a masking and reconstructing technique to enrich the few-shot training data. Next, we utilize the knowledge learned by the pre-trained encoder to improve the graph structure representation of the target city. Finally, we apply a spatial-temporal short-term model to predict future traffic

Training Data Enriching: The target city may only have very little data. Directly utilizing the few-shot data could lead to over-fitting and lack of expressiveness of the fine-tuned model. Consequently, we use the pre-trained Cross-Domain Spatial-Temporal Encoder to augment the training data to create more robust training samples. As shown in Fig. 3, we first mask some part of the training data and use the masked token to replace the masked part. Then, we use the pre-trained encoder to reconstruct the masked data. The reconstructed data is combined with the unmasked data to form the augmented training set. This approach addresses the limited availability of data of the target city, mitigates the risk of over-fitting, and enhances the expressiveness of the fine-tuned model.

Momentum Graph: Due to the potential presence of diverse sources of noise and uncertainties during the generation of the initial traffic graph structure, we exploit the learned knowledge from the pre-trained encoder to enhance the graph structure of the target city. Specifically, the input data sampled from the enriched training dataset is firstly Fourier Transformed. Then, the data of the three domains is fed into the pre-trained Cross-Domain Spatial-Temporal Encoder.

$$\hat{\mathbf{H}}_i^{\text{Ti}}, \hat{\mathbf{H}}_i^{\text{Am}}, \hat{\mathbf{H}}_i^{\text{Ph}} = \text{CDEnc}(\mathbf{S}_i^{\text{Ti}}, \mathbf{S}_i^{\text{Am}}, \mathbf{S}_i^{\text{Ph}}) \quad (21)$$

Subsequently, the embedding of three domains is aggregated.

$$\hat{\mathbf{H}}_i = \hat{\mathbf{H}}_i^{\text{Ti}} + \hat{\mathbf{H}}_i^{\text{Am}} + \hat{\mathbf{H}}_i^{\text{Ph}} \quad (22)$$

The output embedding $\hat{\mathbf{H}} \in \mathbb{R}^{N \times D}$ generates the meta-graph $\hat{\mathbf{A}}_{\text{meta}}$ via inner product.

$$\hat{\mathbf{A}}_{\text{meta}} = \text{SoftMax}(\hat{\mathbf{H}}\hat{\mathbf{H}}^T) \quad (23)$$

In order to enhance the resilience of the learned graph structure and prevent it from over-fitting the limited training data, we utilize the momentum-updating technique to

Table 1: Statistical details of traffic datasets.

	PEMS-BAY	METR-LA	Chengdu	Shenzhen
# of Nodes	325	207	524	627
# of Edges	2,694	1,722	1,120	4,845
Interval	5 min	5 min	10 min	10 min
# of Time Step	52,116	34,272	17,280	17,280
Mean	61.7768	58.2749	29.0235	31.0092
Std	9.2852	13.1280	9.6620	10.9694

preserve the refined graph structure of the target city, denoted as $\hat{\mathbf{A}}$. Formally, in the k -th graph reconstruction, the graph structure is updated as follows.

$$\hat{\mathbf{A}}_{\mathbf{k}} = \tau \times \hat{\mathbf{A}}_{\text{meta}} + (1 - \tau) \times \hat{\mathbf{A}}_{\mathbf{k}-1} \quad (24)$$

$$\hat{\mathbf{A}}_0 = \mathbf{A} \quad (25)$$

Forecasting: After generating the graph structure $\hat{\mathbf{A}}_{\mathbf{k}}$, it is subsequently utilized as input to a downstream Spatial-Temporal model (STmodel), together with the short-term raw data. In this context, we use Graph Wavenet [40] as the backbone STmodel framework.

$$\hat{\mathbf{H}}^{\text{ST}} = \text{STmodel}(\mathbf{S}_{\mathbf{P}}^{\text{Ti}}, \hat{\mathbf{A}}_{\mathbf{k}}) \quad (26)$$

Finally, we concatenate the embedding of the STmodel and the pre-trained encoder to forecast future traffic.

$$\hat{\mathbf{Y}} = \text{Linear}(\hat{\mathbf{H}} \parallel \hat{\mathbf{H}}^{\text{ST}}) \quad (27)$$

The Mean Square Loss is utilized to optimize the STmodel and Linear layers.

$$\mathcal{L} = \frac{1}{NT_f C} \sum_{i=1}^N \sum_{j=1}^{T_f} \sum_{k=1}^C (\mathcal{Y}_{ijk} - \hat{\mathcal{Y}}_{ijk})^2 \quad (28)$$

5 Experiment

This section presents a thorough evaluation of our proposed framework FEPCross. Specifically, We would address the following research questions.

- How does FEPCross perform compared to other baselines in the task of few-shot traffic forecasting?
- Does each component contribute to the final result?
- Do the proposed modules effectively enhance the expressiveness of the pre-trained encoder?
- Does the Cross-Domain Spatial-Temporal Encoder successfully fuse the information of three domains?

5.1 Experiment Settings

Dataset: We evaluate our proposed framework on four real-world public datasets: *PEMS-BAY*, *METR-LA* [20], *Chengdu*, *Shenzhen* [6]. These datasets contain months of traffic speed data and the details of these data are listed in Table 1.

Few-shot Setting: Following the experimental setting of previous research [24,26,25], we choose two cities and set one as the source city and the other one as the target city. We use the full data of the source city as the pre-training data, the two-day few-shot data of the target city as the fine-tuning data, and the rest of the data of the target city as the testing data.

Details: In the pre-training stage, we use $T_h = 288$ and $P = 24$, which means one-day data is divided into 24 patches as the input. The loss weight α is set to 1. The mask ratio of Pre-training is set to 75% according to [24,34]. The learning rate of Pre-training is set to 0.0001. In the Forecasting module, we use $T_f = 12$ to forecast the future 12 steps of data. The momentum ratio τ is set to 0.1. The Adam optimizer has a learning rate of 0.001 and a weight decay of 0.01. The dimension of \mathbf{H} is set to 128. The experiment is implemented by Pytorch 1.10.0 on RTX3090. We evaluate the performances of Mean Absolute Error (MAE) and Mean Absolute Percentage Error (MAPE). Code and data are released in <https://github.com/zhyliu00/FEPCross>.

$$MAPE = \frac{1}{s} \sum_{i=1}^s \left| \frac{y_i - \hat{y}_i}{y_i} \right|, \quad MAE = \frac{1}{s} \sum_{i=1}^s |y_i - \hat{y}_i| \quad (29)$$

6 RQ1: Overall Performance

We select 13 baselines, including traditional statistic methods, deep traffic forecasting methods, time series forecasting methods, and cross-city traffic forecasting methods, to evaluate the effectiveness of FEPCross on the few-shot forecasting task. To guarantee the fairness of the comparison, the deep traffic forecasting methods and time series forecasting methods are firstly trained by Reptile [28] meta-learning framework on the data of source city and then fine-tuned on the few-shot data on target city. For each method, we keep the model architecture the same as the original paper and search the hyper-parameters in the same log space while fixing other parameters (for example, we search the fine-tune learning rate in [1e-5, 1e-4, 1e-3, 1e-2, 1e-1] for all methods).

- Traditional methods: **HA** uses the average of previous periods as predictions. **SVR** uses Support Vector Machines to do regression tasks.
- Deep traffic forecasting methods: **DCRNN** [20] uses diffusion techniques. **GWN** [40] utilizes adaptive graph and dilated causal convolution. **STFGNN** [19] uses DTW distance to construct the temporal graph. **DSTAGNN** [18] constructs spatial-temporal aware graph. **FOGS** [33] predicts the first-order difference.
- Time series forecasting methods: **DLinear** [44] uses linear layers and **FEDFormer** [48] utilized a frequency-enhanced Transformer while they both have decomposition module.
- Cross-city traffic forecasting methods: **ST-GFSL** [26] learns the metaknowledge to generate the parameter of the models. **STEP** [34] reconstructs the masked traffic to

Table 2: Overall performance of few-shot traffic forecasting on *PEMS-BAY*, *METR-LA*, *Chengdu*, and *Shenzhen*. Metr-LA→PEMS-BAY means the source city is METR-LA and the target city is PEMS-BAY. The mean and standard deviation of the results in 5 runs is shown. In each column, the best result is highlighted in bold and grey, and the second-best result is underlined. Marker * and ** indicates the mean of the results is statistically significant (* means t-test with p-value < 0.05 and ** means t-test with p-value < 0.01).

	METR-LA→PEMS-BAY						PEMS-BAY→METR-LA					
	5 min		15 min		30 min		5 min		15 min		30 min	
	MAPE(%)	MAE	MAPE(%)	MAE	MAPE(%)	MAE	MAPE(%)	MAE	MAPE(%)	MAE	MAPE(%)	MAE
HA	9.33	3.99	9.33	3.99	9.33	3.99	14.15	5.21	14.15	5.21	14.15	5.21
SVR	8.07	3.79	7.96	3.70	7.95	3.71	14.02	4.97	14.59	5.11	14.02	5.04
DCRNN	3.66±0.21	1.64±0.04	4.45±0.21	2.04±0.06	5.80±0.15	2.66±0.06	7.88±0.21	3.13±0.14	9.43±0.25	3.56±0.11	11.91±0.42	4.31±0.18
GWN	3.67±0.70	1.49±0.18	5.14±0.76	2.07±0.16	7.04±1.04	2.72±0.23	6.67±0.48	2.78±0.08	9.83±0.80	3.49±0.15	12.77±1.07	4.36±0.21
STGNN	3.21±0.50	1.49±0.11	4.05±0.17	1.93±0.06	5.68±0.18	2.66±0.13	7.57±0.46	3.40±0.32	9.78±0.42	3.93±0.20	12.50±0.65	4.52±0.25
DSTAGNN	3.36±0.49	1.67±0.32	4.80±0.80	2.34±0.54	6.29±0.66	2.85±0.37	7.11±0.57	2.82±0.14	10.18±1.23	3.59±0.47	12.56±1.11	4.42±0.37
FOGS	2.87±0.28	1.38±0.08	3.94±0.24	1.87±0.05	5.64±0.30	2.62±0.060	7.20±0.43	3.23±0.29	9.56±0.55	3.84±0.25	12.41±0.63	4.48±0.25
FEDFormer	6.01±0.55	3.08±0.37	5.87±0.51	2.89±0.24	5.65±0.46	2.69±0.24	14.35±0.92	5.02±0.40	13.12±1.00	4.98±0.35	12.24±0.34	4.79±0.40
DLinear	2.64±0.21	1.44±0.14	4.04±0.14	2.09±0.07	5.64±0.46	2.80±0.24	6.63±0.31	2.84±0.06	9.34±0.69	3.66±0.22	11.87±0.42	4.47±0.10
ST-GFSL	2.88±0.14	1.38±0.04	4.12±0.08	1.99±0.05	5.46±0.15	2.48±0.09	7.01±0.05	2.80±0.03	9.17±0.11	3.43±0.08	13.47±0.68	4.45±0.16
STEP	2.39±0.18	1.17±0.00	3.70±0.19	1.75±0.02	5.29±0.16	2.36±0.04	7.23±1.18	2.74±0.13	9.32±0.34	3.32±0.01	11.89±0.54	4.15±0.26
TransGTR	3.62±0.14	1.47±0.03	4.51±0.17	1.89±0.02	5.57±0.20	2.35±0.04	7.73±0.08	2.86±0.08	8.99±0.12	3.31±0.03	11.51±0.61	4.07±0.04
TPB	2.55±0.29	1.19±0.08	3.86±0.29	1.70±0.07	5.43±0.31	2.36±0.06	6.61±0.11	2.78±0.16	8.97±0.64	3.29±0.10	11.97±1.11	3.98±0.09
FEPCross	1.98±0.02**	1.04±0.01**	3.48±0.04**	1.71±0.01	5.04±0.07**	2.31±0.02**	5.90±0.04*	2.50±0.07*	8.52±0.26*	3.23±0.01*	11.16±0.20**	3.98±0.01
	Shenzhen→Chengdu						Chengdu→Shenzhen					
	10 min		30 min		60 min		10 min		30 min		60 min	
	MAPE(%)	MAE	MAPE(%)	MAE	MAPE(%)	MAE	MAPE(%)	MAE	MAPE(%)	MAE	MAPE(%)	MAE
HA	16.27	3.77	16.27	3.77	16.27	3.77	17.33	4.43	17.33	4.43	17.33	4.43
SVR	15.39	3.58	15.57	3.62	15.79	3.71	15.01	4.02	15.68	4.17	16.12	4.35
DCRNN	10.38±0.24	2.55±0.02	14.07±0.31	3.23±0.03	15.26±0.23	3.55±0.04	8.26±0.15	2.19±0.05	10.66±0.09	2.70±0.04	11.39±0.21	2.95±0.06
GWN	10.49±0.63	2.57±0.07	14.74±0.99	3.36±0.11	15.82±0.83	3.70±0.06	8.62±0.29	2.25±0.04	11.46±0.44	2.88±0.08	12.21±0.55	3.15±0.11
STGNN	11.79±0.86	2.90±0.11	15.78±0.98	3.59±0.09	16.95±0.85	3.95±0.07	9.96±0.54	2.55±0.08	12.60±0.81	3.09±0.10	13.21±0.74	3.35±0.09
DSTAGNN	10.37±0.39	2.55±0.04	15.31±1.15	3.46±0.31	15.47±0.88	3.62±0.10	9.27±1.00	2.48±0.32	11.94±0.80	2.94±0.15	13.17±1.18	3.41±0.43
FOGS	11.63±0.82	2.83±0.10	15.79±0.99	3.56±0.11	16.89±0.99	3.91±0.10	9.32±0.13	2.44±0.02	12.00±0.200	3.01±0.03	12.64±0.18	3.28±0.05
FEDFormer	13.99±0.76	3.26±0.17	13.53±0.86	3.06±0.17	13.49±0.62	3.17±0.19	11.88±0.87	2.87±0.29	11.55±0.70	2.76±0.31	11.65±0.35	2.78±0.23
DLinear	10.76±0.08	2.72±0.02	14.15±0.27	3.40±0.07	13.98±0.32	3.45±0.09	8.79±0.06	2.35±0.01	11.07±0.08	2.85±0.04	11.20±0.12	3.00±0.05
ST-GFSL	9.73±0.48	2.40±0.06	12.91±0.88	2.94±0.11	14.26±1.03	3.35±0.11	8.24±0.55	2.04±0.03	10.67±0.89	2.46±0.07	11.91±1.00	2.78±0.10
STEP	9.21±0.12	2.31±0.03	12.06±0.01	2.88±0.04	13.30±0.11	3.17±0.04	7.62±0.09	1.95±0.00	9.92±0.18	2.44±0.02	10.88±0.14	2.72±0.07
TransGTR	9.51±0.13	2.28±0.04	12.24±0.18	2.79±0.02	13.38±0.17	3.02±0.03	7.94±0.09	2.04±0.04	9.83±0.08	2.46±0.04	10.60±0.04	2.64±0.03
TPB	9.18±0.13	2.19±0.02	12.10±0.19	2.75±0.04	13.21±0.27	3.02±0.04	7.61±0.13	1.94±0.07	9.97±0.16	2.38±0.08	10.57±0.20	2.68±0.08
FEPCross	8.94±0.09**	2.16±0.01**	11.42±0.12**	2.63±0.02**	12.04±0.05**	2.78±0.01**	7.35±0.05**	1.89±0.03**	9.26±0.08**	2.32±0.04**	9.81±0.07**	2.46±0.03**

do per-taining on source datasets. **TransGTR** [16] build transferable graph structure across cities. **TPB** [24] constructs the traffic pattern bank from multiple source cities.

The performance evaluation results are presented in Table 2. Based on the information provided in the table, the following observations can be made. (1) FEPCross surpasses the baseline methods in both short-term and long-term traffic forecasting tasks, exhibiting a notable average improvement of 15.84% in Mean Absolute Percentage Error (MAPE) and 16.77% in Mean Absolute Error (MAE). (2) The cross-city traffic forecasting methods exhibit superior performance compared to the other three categories of baseline methods. This outcome underscores the significance of explicitly addressing the cross-city few-shot traffic forecasting problem, emphasizing the need for dedicated modeling approaches. (3) Among the cross-city traffic forecasting methods, FEPCross performs the best, which demonstrates that FEPCross captures resilient and generalizable frequency-enhanced traffic information across cities. (4) Notably, FEDFormer has high errors in short-term forecasting, which is attributed to that it is primarily tailored for long-term time series forecasting tasks.

Table 3: Ablation Study. The mean and standard deviation of the results on the target city in 5 runs are shown. In each column, the best result is highlighted in bold and grey. P-T and F-T mean the modules proposed in the pre-training stage and the fine-tuning stage respectively.

		PEMS-BAY→METR-LA						Shenzhen→Chengdu					
		5 min MAE	15 min MAE	30 min MAE	5 min MAE	15 min MAE	30 min MAE	10 min MAE	30 min MAE	60 min MAE	10 min MAE	30 min MAE	60 min MAE
P-T	Pretrain base	7.23±1.18	2.74±0.13	9.32±0.34	3.32±0.01	11.89±0.54	4.15±0.26	9.21±0.12	2.31±0.03	12.06±0.01	2.88±0.04	13.30±0.11	3.17±0.04
	Pretrain base+F	6.35±0.18	2.75±0.09	8.91±0.20	3.47±0.07	11.78±0.24	4.43±0.12	9.31±0.13	2.24±0.02	11.87±0.10	2.75±0.02	12.85±0.13	2.96±0.03
	Pretrain base+F+D	6.25±0.07	2.64±0.06	8.73±0.11	3.38±0.08	11.36±0.16	4.27±0.06	9.19±0.05	2.22±0.02	11.64±0.10	2.69±0.02	12.39±0.09	2.86±0.02
	Pretrain base+F+D+S	6.09±0.19	2.57±0.05	8.54±0.38	3.31±0.06	11.31±0.46	4.25±0.03	9.08±0.06	2.19±0.01	11.65±0.12	2.67±0.02	12.49±0.11	2.83±0.03
F-T	Finetune Base	6.13±0.15	2.58±0.03	8.80±0.38	3.34±0.06	11.49±0.48	4.24±0.13	9.44±0.04	2.33±0.03	12.22±0.14	2.92±0.09	12.67±0.14	3.01±0.09
	Finetune Base+M	6.11±0.21	2.54±0.15	8.60±0.23	3.26±0.19	11.18±0.27	4.09±0.23	8.89±0.07	2.16±0.01	11.55±0.13	2.66±0.01	12.15±0.04	2.78±0.02
	FEPCross	5.90±0.04	2.50±0.07	8.52±0.26	3.23±0.01	11.16±0.25	3.98±0.01	8.94±0.09	2.16±0.01	11.42±0.12	2.63±0.02	12.04±0.05	2.77±0.01




Fig. 4: Reconstruction Analysis. The MAE of the reconstruction on the time domain is shown. Pretrain-base and Pretrain base+F have the same performance since they both individually reconstruct the data of the time domain. FEPCross, i.e., Pretrain base+F+D+S+C, is the final version.

7 RQ2: Ablation Study

In this section, we aim to investigate the contribution of each module proposed in this paper. We test the performance of six variations of FEPCross by sequentially adding modules in pre-training and fine-tuning on METR-LA and Chengdu datasets. In the evaluation of the pre-training modules, all fine-tuning modules are added, and vice versa. (1) **Pretrain base** only uses one encoder that takes the time domain as input. (2) **Pretrain base+F(requency)** uses three individual encoders that take the time, amplitude, and phase domains as input. (3) **Pretrain base+F+D(omain)** adds the Cross-domain Aggregator to fuse the information of three domains. (4) **Pretrain base+F+D+S (patial)** adds the Cross-space Aggregator to fuse the spatial information. (5) **Finetune base** directly uses the full pre-training modules and no fine-tune modules to generate predictions. (6) **Finetune base+M** adds the momentum graph module. (7) **FEPCross** is the final version with full modules of pre-training and fine-tuning. It can also be represented as **Pretrain base+F+D+S+C(ontrastive)** or **Finetune base+M+A(ugmented)**. The result is shown in Table 3. By incrementally incorporating the frequency-related modules from the pre-training, a discernible improvement in performance can be observed. This observation demonstrates that the model effectively captures and learns city-invariant patterns in the frequency domain when addressing the cross-city problem, aligning with the intuition that the frequency domain is similar across different cities. Furthermore, the inclusion of the momentum graph and training data enrichment leads to a notable enhancement in performance. This indicates that the incorporation




Fig. 5: The attention map to reconstruct the time domain.

of these modules in the fine-tuning stage imparts robustness to the model and helps mitigate the risk of overfitting the few-shot data.

8 RQ3: Reconstruction Analysis

This section presents an investigation into the enhancement of the expressiveness of the framework through the proposed frequency domain modules. Unlike our previous study in RQ2, we shift our focus to evaluating the time-domain reconstruction error during the pre-training stage rather than the final few-shot forecasting error. If the encoder exhibits more expressiveness, the time-domain reconstruction error should be lower. Here, we consider four pre-training variations mentioned in RQ2 and evaluate their reconstruction performance in the time domain. The result is shown in Fig. 4. We can see that sequentially adding the modules could lead to a better reconstruction error. Moreover, adding the cross-domain (**D**) aggregator leads to a significant improvement, which demonstrates that the fusion of information of frequency domains contributes to the reconstruction of the time domain.

9 RQ4: Cross Domain Analysis

In this section, we aim to investigate the effect that the amplitude and phase domain plays in the reconstruction of the time domain. To visualize this effect, we present the attention map of the Cross-Domain Aggregator during the inference stage, as depicted in Fig. 5. Upon analyzing the attention map, we observe that the amplitude domain, particularly the low-frequency component, plays a significant role in the reconstruction of the time domain. In contrast, the time domain itself contributes less to the process.

These findings indicate that the Cross-Domain Aggregator effectively incorporates information from the frequency domain during the reconstruction pre-training. Consequently, the time domain receives a substantial amount of valuable information from the frequency domain. This wealth of information enables the training of a robust and effective encoder, which in turn facilitates accurate traffic forecasting in the target city.

10 Conclusion

In this paper, we investigate the frequency-enhanced pre-training for the cross-city few-shot traffic forecasting problem. We propose the FEPCross framework to incorporate the frequency domain into the forecasting in the target city. Extensive experiments show that our proposed method not only outperforms other baselines but also exhibits good characteristics such as cross-domain auxiliary information sharing.

Acknowledgment

This work was sponsored by National Key Research and Development Program of China under Grant No.2022YFB3904204, National Natural Science Foundation of China under Grant No. 62102246, 62272301, and Provincial Key Research and Development Program of Zhejiang under Grant No. 2021C01034. Part of the work was done when the students were doing internships at Yunqi Academy of Engineering.

Ethical Statement

The data used in this paper (Chengdu, Shenzhen, PEMS-BAY, and METR-LA) are public open-sourced benchmark datasets, which are widely used in academic research. No personally identifiable information was obtained and people can not infer personal information through the data. The potential use of this work is traffic knowledge transfer across cities, and this work could benefit the downstream applications of developing cities with little traffic data. This work is not potentially a part of policing or military work. The authors of this paper are committed to ethical principles and guidelines in conducting research and have taken measures to ensure the integrity and validity of the data. The use of the data in this study is in accordance with ethical standards and is intended to advance knowledge in the field of cross-city few-shot traffic forecasting.

References

1. Akagi, Y., Nishimura, T., Kurashima, T., Toda, H.: A fast and accurate method for estimating people flow from spatiotemporal population data. In: IJCAI. pp. 3293–3300 (2018)
2. Bai, L., Yao, L., Li, C., Wang, X., Wang, C.: Adaptive graph convolutional recurrent network for traffic forecasting. arXiv preprint arXiv:2007.02842 (2020)
3. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International conference on machine learning. pp. 1597–1607. PMLR (2020)
4. Choi, J., Choi, H., Hwang, J., Park, N.: Graph neural controlled differential equations for traffic forecasting. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 36, pp. 6367–6374 (2022)
5. Cirstea, R.G., Yang, B., Guo, C., Kieu, T., Pan, S.: Towards spatio-temporal aware traffic time series forecasting. In: 2022 IEEE 38th International Conference on Data Engineering (ICDE). pp. 2900–2913. IEEE (2022)
6. Didi, G.i.: Didi chuxing data. <https://gaia.didichuxing.com> (2020)
7. Duan, W., He, X., Zhou, Z., Thiele, L., Rao, H.: Localised adaptive spatial-temporal graph neural network. arXiv preprint arXiv:2306.06930 (2023)
8. Fan, Y., Yeh, C.C.M., Chen, H., Zheng, Y., Wang, L., Wang, J., Dai, X., Zhuang, Z., Zhang, W.: Spatial-temporal graph boosting networks: Enhancing spatial-temporal graph neural networks via gradient boosting. In: Proceedings of the 32nd ACM International Conference on Information and Knowledge Management. pp. 504–513 (2023)
9. Fang, S., Zhang, Q., Meng, G., Xiang, S., Pan, C.: Gstnet: Global spatial-temporal network for traffic flow prediction. In: IJCAI. pp. 2286–2293 (2019)
10. Gupta, M., Kodamana, H., Ranu, S.: Frigate: Frugal spatio-temporal forecasting on road networks. arXiv preprint arXiv:2306.08277 (2023)
11. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16000–16009 (2022)
12. Huang, R., Huang, C., Liu, Y., Dai, G., Kong, W.: Lsgcn: Long short-term traffic prediction with graph convolutional networks. In: IJCAI. pp. 2355–2361 (2020)
13. Ji, J., Wang, J., Jiang, Z., Jiang, J., Zhang, H.: STDEN: Towards physics-guided neural networks for traffic flow prediction. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 36, pp. 4048–4056 (2022)
14. Jiang, J., Wu, B., Chen, L., Zhang, K., Kim, S.: Enhancing the robustness via adversarial learning and joint spatial-temporal embeddings in traffic forecasting. In: Proceedings of the 32nd ACM International Conference on Information and Knowledge Management. pp. 987–996 (2023)
15. Jin, Y., Chen, K., Yang, Q.: Selective cross-city transfer learning for traffic prediction via source city region re-weighting. In: Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. pp. 731–741 (2022)
16. Jin, Y., Chen, K., Yang, Q.: Transferable graph structure learning for graph-based traffic forecasting across cities. In: Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. pp. 1032–1043 (2023)
17. Kang, Z., Liu, Z., Pan, S., Tian, L.: Fine-grained attributed graph clustering. In: Proceedings of the 2022 SIAM International Conference on Data Mining (SDM). pp. 370–378. SIAM (2022)
18. Lan, S., Ma, Y., Huang, W., Wang, W., Yang, H., Li, P.: Dstagnn: Dynamic spatial-temporal aware graph neural network for traffic flow forecasting. In: International Conference on Machine Learning. pp. 11906–11917. PMLR (2022)

19. Li, M., Zhu, Z.: Spatial-temporal fusion graph neural networks for traffic flow forecasting. In: Proceedings of the AAAI conference on artificial intelligence. vol. 35, pp. 4189–4196 (2021)
20. Li, Y., Yu, R., Shahabi, C., Liu, Y.: Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. arXiv preprint arXiv:1707.01926 (2017)
21. Liang, C., Huang, Z., Liu, Y., Liu, Z., Zheng, G., Shi, H., Du, Y., Li, F., Li, Z.: Cblab: Scalable traffic simulation with enriched data supporting. arXiv preprint arXiv:2210.00896 (2022)
22. Lippi, M., Bertini, M., Frasconi, P.: Short-term traffic flow forecasting: An experimental comparison of time-series analysis and supervised learning. *IEEE Trans. Intelligent Transportation Systems* **14**(2), 871–882 (2013)
23. Liu, Z., Liang, C., Zheng, G., Wei, H.: Fdti: Fine-grained deep traffic inference with roadnet-enriched graph. arXiv preprint arXiv:2306.10945 (2023)
24. Liu, Z., Zheng, G., Yu, Y.: Cross-city few-shot traffic forecasting via traffic pattern bank. In: Proceedings of the 32nd ACM International Conference on Information and Knowledge Management. pp. 1451–1460 (2023)
25. Liu, Z., Zheng, G., Yu, Y.: Multi-scale traffic pattern bank for cross-city few-shot traffic forecasting. arXiv preprint arXiv:2402.00397 (2024)
26. Lu, B., Gan, X., Zhang, W., Yao, H., Fu, L., Wang, X.: Spatio-temporal graph few-shot learning with cross-city knowledge transfer. arXiv preprint arXiv:2205.13947 (2022)
27. Ma, Q., Zhang, Z., Zhao, X., Li, H., Zhao, H., Wang, Y., Liu, Z., Wang, W.: Rethinking sensors modeling: Hierarchical information enhanced traffic forecasting. In: Proceedings of the 32nd ACM International Conference on Information and Knowledge Management. pp. 1756–1765 (2023)
28. Nichol, A., Achiam, J., Schulman, J.: On first-order meta-learning algorithms. arXiv preprint arXiv:1803.02999 (2018)
29. Nie, Y., Nguyen, N.H., Sinthong, P., Kalagnanam, J.: A time series is worth 64 words: Long-term forecasting with transformers. arXiv preprint arXiv:2211.14730 (2022)
30. Nikraves, A.Y., Ajila, S.A., Lung, C.H., Ding, W.: Mobile network traffic prediction using mlp, mlpwd, and svm. In: 2016 IEEE International Congress on Big Data (BigData Congress). pp. 402–409. IEEE (2016)
31. Ouyang, X., Yang, Y., Zhou, W., Zhang, Y., Wang, H., Huang, W.: Citytrans: Domain-adversarial training with knowledge transfer for spatio-temporal prediction across cities. *IEEE Transactions on Knowledge and Data Engineering* (2023)
32. Pan, Z., Liang, Y., Wang, W., Yu, Y., Zheng, Y., Zhang, J.: Urban traffic prediction from spatio-temporal data using deep meta learning. In: Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining. pp. 1720–1730 (2019)
33. Rao, X., Wang, H., Zhang, L., Li, J., Shang, S., Han, P.: Fogs: First-order gradient supervision with learning-based graph for traffic flow forecasting. In: Proceedings of International Joint Conference on Artificial Intelligence, IJCAI. ijcai.org (2022)
34. Shao, Z., Zhang, Z., Wang, F., Xu, Y.: Pre-training enhanced spatial-temporal graph neural network for multivariate time series forecasting. In: Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. pp. 1567–1577 (2022)
35. Shao, Z., Zhang, Z., Wei, W., Wang, F., Xu, Y., Cao, X., Jensen, C.S.: Decoupled dynamic spatial-temporal graph neural network for traffic forecasting. arXiv preprint arXiv:2206.09112 (2022)
36. Wang, L., Geng, X., Ma, X., Liu, F., Yang, Q.: Cross-city transfer learning for deep spatio-temporal prediction. arXiv preprint arXiv:1802.00386 (2018)
37. Wang, Y., Jin, H., Zheng, G.: Ctrl: Cooperative traffic tolling via reinforcement learning. In: Proceedings of the 31st ACM International Conference on Information & Knowledge Management. pp. 3545–3554 (2022)

38. Wei, H., Zheng, G., Yao, H., Li, Z.: Intellilight: A reinforcement learning approach for intelligent traffic light control. In: Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. pp. 2496–2505 (2018)
39. Wei, Y., Zheng, Y., Yang, Q.: Transfer knowledge between cities. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp. 1905–1914 (2016)
40. Wu, Z., Pan, S., Long, G., Jiang, J., Zhang, C.: Graph wavenet for deep spatial-temporal graph modeling. arXiv preprint arXiv:1906.00121 (2019)
41. Yao, H., Liu, Y., Wei, Y., Tang, X., Li, Z.: Learning from multiple cities: A meta-learning approach for spatial-temporal prediction. In: The World Wide Web Conference. pp. 2181–2191 (2019)
42. Yu, B., Yin, H., Zhu, Z.: Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. arXiv preprint arXiv:1709.04875 (2017)
43. Yuan, Y., Ding, J., Feng, J., Jin, D., Li, Y.: Unist: A prompt-empowered universal model for urban spatio-temporal prediction. arXiv preprint arXiv:2402.11838 (2024)
44. Zeng, A., Chen, M., Zhang, L., Xu, Q.: Are transformers effective for time series forecasting? In: Proceedings of the AAAI conference on artificial intelligence. vol. 37, pp. 11121–11128 (2023)
45. Zhang, X., Gong, Y., Zhang, X., Wu, X., Zhang, C., Dong, X.: Mask-and contrast-enhanced spatio-temporal learning for urban flow prediction. In: Proceedings of the 32nd ACM International Conference on Information and Knowledge Management. pp. 3298–3307 (2023)
46. Zhang, Y., Li, Y., Zhou, X., Kong, X., Luo, J.: Strans-gan: Spatially-transferable generative adversarial networks for urban traffic estimation. In: 2022 IEEE International Conference on Data Mining. pp. 743–752. IEEE (2022)
47. Zheng, C., Fan, X., Wang, C., Qi, J.: Gman: A graph multi-attention network for traffic prediction. In: Proceedings of the AAAI conference on artificial intelligence. vol. 34, pp. 1234–1241 (2020)
48. Zhou, T., Ma, Z., Wen, Q., Wang, X., Sun, L., Jin, R.: Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting. In: International Conference on Machine Learning. pp. 27268–27286. PMLR (2022)