

# ST-LoRA: Low-rank Adaptation for Spatio-Temporal Forecasting

Weilin Ruan<sup>1</sup>, Wei Chen<sup>1</sup>, Xilin Dang<sup>2</sup>, Jianxiang Zhou<sup>1</sup>, Weichuang Li<sup>1</sup>,

Xu Liu<sup>3</sup>, and Yuxuan Liang<sup>1,4</sup> (✉)

<sup>1</sup> The Hong Kong University of Science and Technology (Guangzhou), China  
{rwlinno, onedeanxxx}@gmail.com, {jzhou814, wli043}@connect.hkust-gz.edu.cn,  
yuxliang@outlook.com

<sup>2</sup> The Chinese University of Hong Kong, Hong Kong xldang23@cse.cuhk.edu.hk

<sup>3</sup> National University of Singapore, Singapore liuxu@comp.nus.edu.sg

<sup>4</sup> State Key Lab of Resources and Environmental Information System, Chinese Academy of Sciences, China yuxliang@outlook.com

**Abstract.** Spatio-temporal forecasting is essential for understanding future dynamics within real-world systems by leveraging historical data from multiple locations. Existing methods often prioritize the development of intricate neural networks to capture the complex dependencies of the data. These methods neglect node-level heterogeneity and face over-parameterization when attempting to model node-specific characteristics. In this paper, we present a novel low-rank adaptation framework for existing spatio-temporal prediction models, termed ST-LoRA, which alleviates the aforementioned problems through node-level adjustments. Specifically, we introduce the node-adaptive low-rank layer and node-specific predictor, capturing the complex functional characteristics of nodes while maintaining computational efficiency. Extensive experiments on multiple real-world datasets demonstrate that our method consistently achieves superior performance across various forecasting models with minimal computational overhead, improving performance by 7% with only 1% additional parameter cost. The source code is available at <https://github.com/RWLinno/ST-LoRA>.

## 1 Introduction

With the rapid advancement of data acquisition technologies and mobile computing, vast spatio-temporal data are being generated for urban analysis and related applications [51,6]. Spatio-temporal forecasting aims to predict future changes based on dynamic temporal observations recorded at static locations with spatial associations [52]. Modeling and analyzing these spatio-temporal dynamic systems can be applied to various prediction scenarios, such as traffic speed forecasting [48,45], taxi demand prediction [47], and air quality prediction [23,25]. Early research primarily focused on traditional time-series models, such as the Historical Average (HA) [36] method and the Auto-Regressive Integrated Moving Average (ARIMA) [26] model, as well as machine learning-based

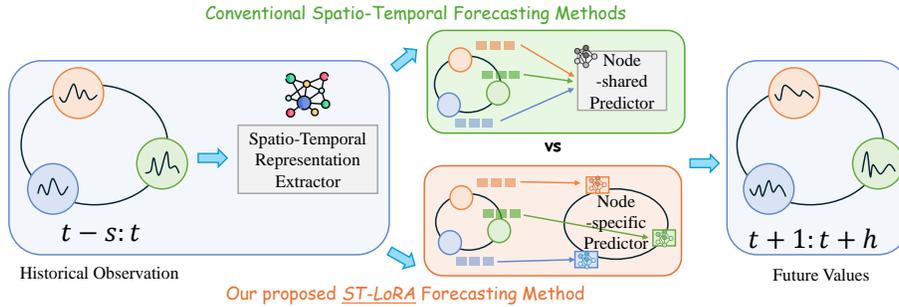


Fig. 1: Comparison between conventional STF methods and our proposed ST-LoRA framework featuring node-specific predictors for node-level fine-tuning.

models [38], including Vector Auto-Regression (VAR) [55] and Artificial Neural Networks (ANN) [13]. These methods were applied directly to spatio-temporal forecasting (STF) without considering spatial dependencies, leading to suboptimal performance. With the accumulation of spatio-temporal big data, recent approaches have shifted towards data-driven deep learning models, which are capable of capturing the inherent spatio-temporal dependencies within dynamic systems. Simple yet effective strategies involve using convolutional neural networks (CNNs) [9] to capture spatial dependencies and recurrent neural networks (RNNs) [49,11,50,4] for temporal dependencies, thereby improving performance.

Given the non-Euclidean nature of spatial dependencies, deep learning methods have evolved to combine sophisticated temporal models with Graph Neural Networks (GNNs) [44,19] for capturing both global temporal dependencies and regional patterns. Spatio-temporal graph neural networks (STGNNs) [20,48] have emerged as powerful tools for learning robust high-level spatio-temporal representations through local information aggregation [16]. Recent years have witnessed significant advances in this field, including innovations in graph convolution architectures [7,42], dynamic graph structure learning [45,15], and efficient attention mechanisms [53,43]. Furthermore, researchers have explored integrating advanced techniques such as self-supervised learning [34] and large language models [54] into spatio-temporal prediction tasks. While these sophisticated approaches have achieved remarkable performance improvements, they often come at the cost of increased computational complexity and memory requirements, making it challenging to balance model effectiveness with operational efficiency.

To better understand the landscape of spatio-temporal forecasting, we present a systematic analysis of existing architectures, revealing their common structural patterns as illustrated in Figure 1. Existing spatio-temporal forecasting methods typically consist of two main components. The first component is the *Spatio-Temporal Representation Extractor*, which serves as the core framework and is responsible for capturing high-order complex spatio-temporal relationships. This component can be implemented using various architectures, such as

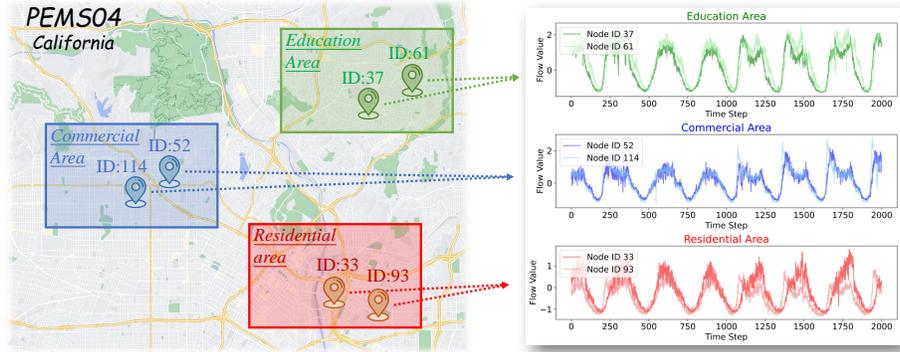


Fig. 2: Traffic flow visualization of each node in different areas on the PEMS04.

CNNs, RNNs, and STGNNs. The second component is the *Node-Shared Predictor*, which takes the advanced spatio-temporal representations extracted by the first component and predicts future changes for each location. This predictor typically consists of parameter-sharing fully connected layers.

While this two-component architecture has been widely adopted in existing methods, it suffers from a significant limitation: the parameter-sharing node predictor fundamentally struggles to address **node-level heterogeneity**. This shared parameterization approach assumes all nodes exhibit similar behavior patterns that can be modeled using identical parameters. However, in real-world scenarios, nodes frequently display distinct temporal dynamics and behavioral characteristics that shared predictors cannot adequately capture. This heterogeneity is particularly pronounced in urban traffic networks, where sensors distributed across diverse functional areas exhibit unique patterns influenced by their specific contexts and surrounding environments. To illustrate this problem intuitively, we analyze the PEMS04 traffic flow dataset [3] from California, examining nodes across different urban zones (Figure 2). Our analysis reveals that even nodes 33 and 93, situated within the same residential area, exhibit substantially different temporal evolution patterns between time steps 1250 and 2000. This finding underscores the limitations of current parameter-sharing approaches in modeling node-specific behavioral characteristics. Consequently, existing spatio-temporal forecasting models attempting to address heterogeneity through extensive shared parameterization inevitably confront the **over-parameterization dilemma**. Maintaining separate parameters for each node not only incurs prohibitive computational and memory costs, especially for large-scale networks with hundreds or thousands of nodes, but also significantly increases the risk of model overfitting and compromised generalization performance.

To address these challenges, we draw inspiration from low-rank matrix factorization techniques [33]. Specifically, we first customize a node-adaptive low-rank layer containing multiple trainable matrices, utilizing low-rank decomposition techniques to effectively reduce computational complexity and enhance model

training efficiency. Subsequently, we propose a novel lightweight and efficient low-rank adaptation framework named ST-LoRA. This framework seamlessly integrates the low-rank layer into existing spatio-temporal forecasting models through a multi-layer fusion residual stacking approach, thereby achieving node-specific predictors and mitigating the effects of overparameterization. Experimental results on real-world traffic datasets show that our framework significantly improves performance over various baseline methods in spatio-temporal forecasting tasks. Our major contributions can be summarized as follows:

- 1) A node-level heterogeneity perspective for STF.** We are the first to introduce low-rank adaptation techniques to the spatio-temporal domain to explicitly account for the node-level heterogeneity. Our proposed Node Adaptive Low-rank Layers capture diverse node-level patterns and distributions by leveraging low-rank matrix factorization while maintaining computational efficiency.
- 2) A general low-rank adaptation method for existing ST models.** We developed node-specific predictors along with a framework called ST-LoRA, which, in a lightweight and efficient manner, allows existing spatio-temporal prediction models to serve as backbone networks to enhance overall performance.
- 3) Extensive empirical studies.** We rigorously evaluate our proposed method on various models and six public traffic datasets. The experimental results demonstrate that our method significantly enhances prediction accuracy across all baseline models while requiring less than **1%** additional learnable parameters, achieving remarkable more than **7%** performance improvements in terms of average RMSE across prediction horizons.

## 2 Preliminaries

### 2.1 Formulation

The objective of STF is to predict future values based on previously observed time series data from  $N$  correlated sensors. This sensor network can be represented as a weighted directed graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{W})$ , where  $\mathcal{V}$  is the node set with  $|\mathcal{V}| = N$ ,  $\mathcal{E}$  is the edge set, and  $\mathcal{W} \in \mathbb{R}^{N \times N}$  is a weighted adjacency matrix that encodes the relationships between nodes. The spatio-temporal data observed on  $\mathcal{G}$  can be represented as a graph signal  $X \in \mathbb{R}^{N \times F}$ , where  $F$  is the number of features associated with each node. Let  $X^{(t)}$  denote the graph signal observed at time  $t$ . The spatio-temporal forecasting problem aims to learn a function  $\mathcal{F}(\cdot)$  that maps  $s$  historical graph signals to  $h$  future graph signals, given a graph  $\mathcal{G}$ :

$$[X^{(t-s+1)}, \dots, X^{(t)}; \mathcal{G}] \xrightarrow{\mathcal{F}(\cdot)} [X^{(t+1)}, \dots, X^{(t+h)}]. \quad (1)$$

### 2.2 Related Work

**Spatio-temporal Forecasting** has evolved into a foundational paradigm for predicting future states by leveraging historical observations across spatial and temporal dimensions. Traditional methods grounded in statistical and time series analysis achieved modest success but exhibited significant limitations in

modeling complex spatial structures and intricate ST relationships [42,16]. To address these shortcomings, deep learning frameworks have increasingly been embraced, which demonstrate superior capability in extracting latent feature representations, including non-linear spatial and temporal correlations from historical data [41,29,40]. Among these advanced frameworks, Spatio-Temporal Graph Neural Networks (STGNNs) have emerged as powerful tools for prediction tasks. By integrating Graph Neural Networks (GNNs) [19] with sophisticated temporal modeling techniques [49], these architectures effectively capture complex spatio-temporal dynamics through local information aggregation. Over the past decade, several influential STGNN architectures have been proposed, including GWNet [45], STGCN [48], DCRNN [20], and AGCRN [2], each demonstrating remarkable performance across diverse spatio-temporal prediction tasks. Complementing these developments, attention mechanisms [39,30] have gained substantial traction due to their effectiveness in modeling dynamic dependencies inherent in spatio-temporal data. Despite the proliferation and diversification of STGNN architectures, performance improvements have begun to plateau, prompting researchers to explore integrating Self-Supervised Learning (SSL) [34,22] and Large Language Models (LLMs) [54,17,46]. Recent studies have further investigated methods to capture spatio-temporal heterogeneity through techniques such as spatial-temporal decoupled masked pre-training [8] and heterogeneity-informed learning approaches [5].

However, these methods often introduce substantial computational overhead and model complexity. Against this backdrop, we introduce a parameter-efficient node-specific adaptation method that significantly enhances existing forecasting frameworks while maintaining minimal parameter and computational costs.

**Low-rank Adaptation** is the technique that decomposes high-dimensional parameter spaces into products of low-rank matrices, reducing computational complexity while preserving essential information. The foundational work on LoRA [12,10] demonstrated that injecting trainable low-rank matrices into pre-trained models enables efficient adaptation with minimal parameter overhead. This approach has been refined through variants like DyLoRA [37] with dynamic adaptation mechanisms and Compacter [18] leveraging parameterized complex multiplication layers for task-specific optimization. The extension of low-rank adaptation to multi-modal and spatio-temporal domains represents a significant advancement for complex applications. MTLORA [1] adapted this approach for multi-task learning scenarios, while robust low-rank reconstruction techniques [14] have shown effectiveness for preserving invariant features across domains. In spatio-temporal forecasting, MFSTN [32] and DeepLGN [24] use matrix or tensor factorization for region-specific parameter decomposition, but their grid-based modeling limits node-level granularity. ST-Adapter [31] introduced low-rank adaptation for cross-modality transfer in spatio-temporal tasks but focused on modality transfer rather than node-specific adaptation.

However, these approaches rely predominantly on region-level predictions, limiting their ability to capture fine-grained node-level heterogeneity and com-

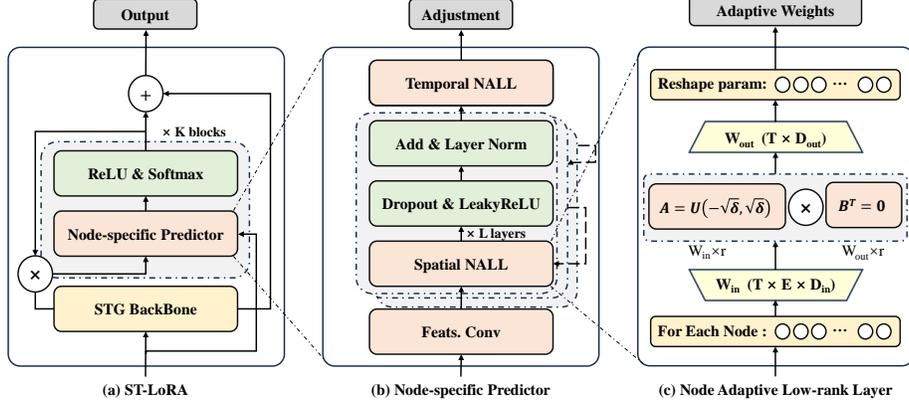


Fig. 3: Overview of the proposed ST-LoRA framework.

promising generalizability across diverse spatial configurations. Our proposed framework addresses these limitations by adapting low-rank optimization for node-level heterogeneity in spatio-temporal forecasting.

### 3 Methodology

In this section, we present ST-LoRA, a lightweight yet effective framework that enhances spatio-temporal models by addressing the node heterogeneity and over-parameterization challenges. As illustrated in Figure 3, our framework consists of three key components: (a) the gated integration mechanism that enhances existing backbone STGNNs, (b) Node-specific Predictors (NSPs) that capture patterns with node-level fine-tuning, and (c) Node-Adaptive Low-rank Layers (NALLs) that leverage low-rank adaptation for parameter customization.

#### 3.1 Node-Adaptive Low-rank Layers

The core innovation of ST-LoRA lies in the Node-Adaptive Low-rank Layers (NALL), which introduce learnable low-rank matrices to efficiently customize the base model parameters for different nodes. Let  $\mathbf{W} \in \mathbb{R}^{d_{out} \times d_{in}}$  represent the base weight matrix, where  $d_{in}$  and  $d_{out}$  are the input and output dimensions, respectively. Traditional approaches would require a complete set of parameters for each node, leading to  $\mathcal{O}(N \cdot d_{in} \cdot d_{out})$  parameters for  $N$  nodes, which is computationally prohibitive for large-scale spatio-temporal networks. Drawing inspiration from low-rank matrix factorization theory [21], NALL decomposes the parameter matrix  $\Delta\mathbf{W}$  into the product of two low-rank matrices. For node  $v_i$  with input  $\mathbf{x}$ , the adaptation process can be formally expressed as:

$$\Delta\mathbf{W}_{v_i} = \mathbf{B}\mathbf{A}_{v_i} \cdot \frac{\alpha}{r}, \quad \hat{\mathbf{y}}_i = \sigma(\mathbf{W}\mathbf{x} + \Delta\mathbf{W}_{v_i}\mathbf{x} + \mathbf{b}), \quad (2)$$

where  $\mathbf{B} \in \mathbb{R}^{d_{out} \times r}$  and  $\mathbf{A}_i \in \mathbb{R}^{r \times d_{in}}$  are learnable low-rank matrices with rank  $r$ ,  $\alpha$  is a scaling factor controlling the adaptation magnitude, and  $\sigma$  is a non-linear activation function (e.g., LeakyReLU). This formulation reduces the parameter complexity from  $\mathcal{O}(N \cdot d_{in} \cdot d_{out})$  to  $\mathcal{O}(r \cdot (d_{in} + d_{out}) + N \cdot r^2)$ , representing a significant reduction when  $r \ll \min(d_{in}, d_{out})$ . By constraining adaptations to lower-dimensional node-specific patterns, NALL efficiently captures meaningful behavioral differences while preventing overfitting to node-specific noise patterns. During inference, while  $\mathbf{W}$  remains frozen to preserve the pre-trained knowledge,  $\mathbf{A}_i$  and  $\mathbf{B}$  continue to adapt and capture node-specific patterns. Our implementation further enhances NALL with dropout with  $p = 0.3$  and Kaiming initialization to stabilize training dynamics and improve convergence stability.

### 3.2 Node-specific Predictors

While a single NALL layer provides efficient parameter adaptation, capturing complex spatio-temporal dependencies requires a more sophisticated architectural design. We propose Node-Specific Predictors (NSP) that hierarchically stack multiple NALL layers with residual connections to model heterogeneous patterns across different nodes. Given an input sequence  $\mathbf{X}_{t-T:t} \in \mathbb{R}^{T \times N \times D}$ , where  $T$  is the sequence length,  $N$  is the number of nodes, and  $D$  is the feature dimension, NSP processes the input through following integrated components:

$$\begin{aligned} \mathbf{H}^{(0)} &= \text{Conv2D}(\mathbf{X}_{t-T:t}), \\ \mathbf{H}^{(l)} &= \mathbf{H}^{(l-1)} + \text{NALL}^{(l)}(\sigma(\mathbf{H}^{(l-1)})), \quad l = 1, 2, \dots, L \\ \hat{\mathbf{Y}}_t &= \mathcal{G}_t(\mathbf{H}^{(L)}), \end{aligned} \quad (3)$$

where  $\mathbf{H}^{(l)} \in \mathbb{R}^{T \times N \times D}$  represents the hidden feature representations at layer  $l$ ,  $\text{Conv2D}(\cdot)$  first extracts temporal features using convolution operations with optimized kernel configurations, multiple NALL layers then process spatial information through residual connections, and finally  $\mathcal{G}_t$  projects the features for temporal prediction. The activation function  $\sigma$  incorporates RMSNorm and dropout to control model complexity and enhance training stability.

The key innovation of NSP lies in its unified approach to spatio-temporal modeling. Unlike traditional GNNs that rely on fixed graph convolutions, NSP leverages the low-rank structure of NALLs to achieve both spatial and temporal adaptability. Specifically, the spatial patterns are captured through the rank-constrained weight matrices in NALLs, which naturally aggregate node-specific information, while the temporal dependencies are modeled through the sequential application of adapted transformations. This design enables NSP to efficiently capture complex node-specific patterns while maintaining computational efficiency through the low-rank structure of NALL layers. The effectiveness of this architecture is particularly evident in scenarios with heterogeneous node behaviors, where traditional fixed-parameter approaches often struggle to capture diverse patterns simultaneously. Importantly, while NSP provides node-specific adaptations, it preserves spatial dependencies captured by the backbone model

by operating on representations that already encode inter-node relationships, thus complementing rather than replacing existing spatial modeling mechanisms.

### 3.3 ST-LoRA Integration

To enhance model generalization and adaptability to diverse spatio-temporal patterns, we propose ST-LoRA, a novel low-rank adaptation framework for spatio-temporal data that seamlessly integrates existing backbone models with our NSP module. Let  $f_\theta$  denote a pre-trained spatio-temporal backbone model and  $\mathcal{H}(\cdot)$  represent our NSP-based enhancement operator. The ST-LoRA framework employs a hierarchical architecture with multiple NSP blocks to capture complex spatio-temporal dependencies at different scales. Given input sequence  $\mathbf{X}_{t-T:t}$ , the integration process can be formulated as:

$$\begin{aligned} \mathbf{Y}_{\text{base}} &= f_\theta(\mathbf{X}_{t-T:t}), \\ \mathbf{Z}^{(1)} &= \mathcal{H}^{(1)}(\sigma(\mathbf{Y}_{\text{base}})), \\ \mathbf{Z}^{(k)} &= \mathcal{H}^{(k)}(\sigma([\mathbf{X}_{t-T:t}, \mathbf{Z}^{(k-1)}])), \quad k = 2, \dots, K \\ \mathbf{R} &= \sigma(\mathcal{F}([\mathbf{X}_{t-T:t}, \frac{1}{K} \sum_{k=1}^K \mathbf{Z}^{(k)}])), \\ \mathbf{Y}_{\text{final}} &= \mathbf{R} \odot \mathbf{Y}_{\text{base}} + (1 - \mathbf{R}) \odot \frac{1}{K} \sum_{k=1}^K \mathbf{Z}^{(k)}, \end{aligned} \tag{4}$$

where  $K$  is the number of NSP blocks,  $\sigma$  is a non-linear activation function (e.g., ReLU),  $[\cdot, \cdot]$  denotes feature concatenation,  $\mathcal{F}$  is a learnable fusion layer that generates adaptive blending weights, and  $\mathbf{R} \in [0, 1]$  is a node-specific gating tensor that controls the contribution of the adaptation mechanism. This allows NSP blocks to access both original historical context and refined node-specific representations, enabling the framework to capture diverse spatio-temporal patterns at different scales while maintaining stability through averaging.

The framework is optimized end-to-end using a temporal Mean Absolute Error (MAE) loss with L2 regularization over the prediction horizon  $T'$ :

$$\mathcal{L} = \frac{1}{T'} \sum_{i=1}^{T'} \|\mathbf{X}_{t+i} - \mathbf{Y}_{\text{final}}^{(t+i)}\|_1 + \lambda \|\alpha\|_2, \tag{5}$$

where  $\lambda$  controls the strength of the regularization term on the gating parameter. This formulation ensures accurate predictions through the MAE term while preventing over-reliance on either the backbone model or the adaptation.

The effectiveness of ST-LoRA stems from its ability to preserve the backbone model’s general prediction capability while introducing node-specific adaptations through the NSP blocks, making it suitable for complex real-world scenarios with heterogeneous spatio-temporal patterns. By integrating NALLs and NSPs, ST-LoRA provides an efficient solution to the fundamental trade-off between modeling node-level heterogeneity and managing computational complexity.

## 4 Experiments

In this section, we conduct extensive experiments to investigate the following Research Questions (RQ):

- **RQ1:** Can ST-LoRA be seamlessly integrated with various spatio-temporal prediction models?
- **RQ2:** How effectively does our framework improve prediction performance across different scenarios?
- **RQ3:** What are the computational overhead and parameter costs of our framework?
- **RQ4:** How do different architectural choices affect the model’s performance?

### 4.1 Experimental Setup

**Datasets.** We evaluate our approach on six public traffic datasets that are widely used in spatio-temporal forecasting research. As summarized in Table 1, these datasets encompass both traffic speed measurements (METR-LA, PEMS-BAY) and traffic flow records (PEMS03/04/07/08), featuring diverse spatial scales ranging from 170 to 883 sensor nodes and temporal ranges spanning 16,992 to 52,116 timestamps. Following standard practice, we split each dataset chronologically into training, validation, and testing sets with ratios of 7:1:2 for speed datasets and 6:2:2 for flow datasets. All above datasets are divided along the time axis into three non-overlapping parts, including training, validation, and test sets. METR-LA and PEMS-BAY are divided in a fraction of 7:1:2 while PEMS03, PEMS04, PEMS07, and PEMS08 are divided in a fraction of 6:2:2.

**Evaluation Protocol.** We adopt three standard metrics for evaluation: Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and Mean Absolute Percentage Error (MAPE). For comprehensive assessment, we examine model performance across different prediction horizons (15-min, 30-min, and 60-min) and report both the horizon-specific and average results. Each experiment is repeated five times with different random seeds to ensure statistical reliability.

Table 1: Statistics and description of datasets we used.

Dataset	#Nodes	#Edges	#Frames	Time Range	Type
METR-LA	207	1515	34,272	03/01/2012 – 06/27/2012	Traffic speed
PEMS-BAY	325	2369	52,116	01/01/2017 – 06/30/2017	Traffic speed
PEMS03	358	547	26208	09/01/2018 – 11/30/2018	Traffic flow
PEMS04	307	340	16992	01/01/2018 – 02/28/2018	Traffic flow
PEMS07	883	866	28224	05/01/2017 – 08/06/2017	Traffic flow
PEMS08	170	295	17856	07/01/2016 – 08/31/2016	Traffic flow

**Baselines.** We evaluate six representative spatio-temporal prediction models as backbone networks to validate the effectiveness of the ST-LoRA framework. Long Short-Term Memory (LSTM) [11] controls the flow of information by introducing a gating mechanism to selectively keep and forget temporal data. Spatio-temporal Graph Convolution Network (STGCN) [48] combines graph convolution and 1D convolution to process spatio-temporal data. Graph WaveNet (GWN) [45] utilizes adaptive adjacency matrices and dilation convolution to capture spatial and temporal correlations in traffic data. Adaptive Graph Convolutional Recurrent Network (AGCRN) [2] infers dependencies between streaming time series through node adaptive parameter learning and data-adaptive graph generation modules. Decoupled Dynamic Spatio-Temporal Graph Neural Network (D2STGNN) [35] is able to separate diffuse and intrinsic traffic information, thus enhancing dynamic graph learning. Finally, Spatio-Temporal Adaptive Embedding transformer (STAE) [27] encodes nodal, spatial, and temporal features through linear layers and multiple embedding layers respectively.

#### 4.2 Model Settings of ST-LoRA (RQ1)

All experiments are conducted using the PyTorch framework on a Linux server equipped with NVIDIA RTX A6000 GPUs. For model training, we employ Adam optimizer with an initial learning rate of 0.001 and weight decay of 0.0005. The evaluation metrics include MAE, RMSE, and MAPE, which are standard in traffic forecasting tasks. For comprehensive assessment, we reported the average performance across all 12 prediction horizons on four traffic flow datasets (PEMS03, PEMS04, PEMS07, and PEMS08).

We directly converted the original baselines to ST-LoRA framework as backbones. The learning rate is set to be adjusted in step 10, and the ratio is 0.1. Various models are implemented concerning the benchmark LargeST [28] and their official source code. For hyper-parameters in the framework, such as the number of NSPs is usually taken as 1, the number of node adaptive low-rank layers is usually taken as 4, and the maximum rank of the low-rank space is usually taken as 16, which is generally adjusted according to the original model. This setup ensures that we seamlessly integrate existing methods into our framework. Experiments show that our proposed framework supports multiple spatio-temporal prediction models, including the aforementioned baselines.

#### 4.3 Performance Comparisons (RQ2)

To validate the effectiveness of our framework, we conducted comprehensive experiments across multiple models and datasets. Specifically, we analyze the performance improvements of different baseline models on the PEMS04 dataset (Table 2), and the generalization capability of our framework across multiple traffic datasets (Table 3). For statistical reliability, each experiment was repeated five times, with models enhanced by our framework denoted with a "+" suffix.

As shown in Table 2, ST-LoRA demonstrates consistent performance improvements across various baseline models, ranging from traditional models to

Table 2: The improvement of different models in the PEMS04 dataset. Here, lower values indicate better performance. All six baselines have achieved significant improvements, denoted by  $\Delta$ . The subscripts indicate standard deviations.

Model	15min			30min			60min			Average		
	MAE ↓	RMSE ↓	MAPE% ↓	MAE ↓	RMSE ↓	MAPE% ↓	MAE ↓	RMSE ↓	MAPE% ↓	MAE ↓	RMSE ↓	MAPE% ↓
HA	28.92±1.25	42.69±1.82	20.31±0.89	33.73±1.28	49.37±1.85	24.01±0.91	46.97±1.31	67.43±1.89	35.11±0.92	38.03±1.28	59.24±1.85	27.88±0.91
VAR	21.94±0.62	34.30±1.02	16.42±0.48	23.72±0.71	36.58±1.08	18.02±0.52	26.76±0.82	40.28±1.23	20.94±0.64	23.51±0.72	36.39±1.11	17.85±0.55
SVR	22.52±0.68	35.30±1.12	14.71±0.45	27.63±0.78	42.23±1.25	18.29±0.49	37.86±1.15	56.01±1.70	26.72±0.82	28.66±0.87	44.59±1.36	19.15±0.59
LSTM	21.94±0.59	33.37±0.93	15.32±0.40	25.83±0.66	39.10±1.04	20.35±0.43	36.41±0.82	50.73±1.28	29.92±0.56	27.14±0.69	41.59±1.08	18.20±0.46
LSTM+	18.89±0.58	29.96±0.91	13.02±0.40	21.31±0.65	34.22±1.04	13.96±0.43	26.34±0.80	41.30±1.26	18.26±0.56	22.18±0.68	35.16±1.07	15.08±0.46
$\Delta$	-3.05±0.18	-3.41±0.20	-2.30±0.14	-4.52±0.27	-4.88±0.29	-6.39±0.38	-10.07±0.60	-9.43±0.57	-11.66±0.70	-4.96±0.30	-6.43±0.39	-3.12±0.19
STGCN	19.45±0.59	30.12±0.92	14.21±0.43	21.85±0.62	34.43±0.97	14.13±0.44	26.97±0.68	41.11±1.06	16.84±0.48	22.70±0.63	35.55±0.98	14.59±0.45
STGCN+	19.12±0.58	29.72±0.91	13.89±0.42	19.92±0.61	31.63±0.96	13.77±0.42	22.07±0.67	34.47±1.05	15.42±0.47	20.37±0.62	31.94±0.97	14.36±0.44
$\Delta$	-0.33±0.02	-0.40±0.02	-0.32±0.02	-1.93±0.12	-2.80±0.17	-0.36±0.02	-4.90±0.29	-6.64±0.40	-1.42±0.09	-2.33±0.14	-3.61±0.22	-0.23±0.01
GWNet	18.65±0.57	29.24±0.89	13.82±0.42	19.57±0.60	30.62±0.92	13.28±0.39	23.07±0.70	35.35±1.08	17.34±0.53	25.45±0.62	39.70±0.97	17.29±0.45
GWNet+	17.89±0.55	28.52±0.87	12.64±0.39	18.88±0.58	29.38±0.89	13.06±0.40	20.89±0.64	32.96±1.00	14.92±0.46	19.22±0.59	30.62±0.93	13.54±0.41
$\Delta$	-0.76±0.05	-0.72±0.04	-1.18±0.07	-0.69±0.04	-1.24±0.07	-0.22±0.01	-2.18±0.13	-2.39±0.14	-2.42±0.15	-6.23±0.37	-9.08±0.54	-3.75±0.23
AGCRN	18.12±0.55	29.45±0.90	12.85±0.39	18.77±0.57	30.08±0.92	12.97±0.40	20.41±0.62	32.87±1.00	14.38±0.44	19.83±0.58	32.26±0.94	13.40±0.41
AGCRN+	17.83±0.54	29.16±0.89	12.55±0.38	18.63±0.57	29.99±0.91	12.82±0.39	19.97±0.61	32.37±0.99	13.78±0.42	18.81±0.57	30.51±0.93	13.05±0.40
$\Delta$	-0.29±0.02	-0.29±0.02	-0.30±0.02	-0.14±0.01	-0.09±0.01	-0.15±0.01	-0.44±0.03	-0.50±0.03	-0.60±0.04	-1.02±0.06	-1.75±0.11	-0.35±0.02
STAE	17.95±0.55	29.12±0.89	12.65±0.39	18.92±0.58	30.09±0.92	13.35±0.41	21.06±0.64	33.37±1.02	15.55±0.47	19.31±0.59	30.86±0.94	13.85±0.42
STAE+	17.65±0.54	28.73±0.88	12.45±0.38	18.62±0.57	29.55±0.90	13.29±0.41	20.40±0.62	32.38±0.99	15.00±0.46	18.89±0.58	30.22±0.92	13.58±0.41
$\Delta$	-0.30±0.02	-0.39±0.02	-0.20±0.01	-0.30±0.02	-0.54±0.03	-0.06±0.01	-0.66±0.04	-0.99±0.06	-0.55±0.03	-0.42±0.03	-0.64±0.04	-0.27±0.02
D2STGNN	18.95±0.58	29.85±0.91	14.82±0.45	19.96±0.61	31.34±0.95	15.52±0.47	23.34±0.71	35.89±1.09	17.39±0.53	20.75±0.63	32.36±0.99	15.91±0.49
D2STGNN+	18.25±0.56	28.92±0.88	14.12±0.43	19.21±0.59	30.50±0.93	13.46±0.41	21.73±0.66	33.73±1.03	17.00±0.52	19.73±0.60	31.05±0.95	14.86±0.45
$\Delta$	-0.70±0.04	-0.93±0.06	-0.70±0.04	-0.75±0.05	-0.84±0.05	-2.06±0.12	-1.61±0.10	-2.16±0.13	-0.39±0.02	-1.02±0.06	-1.31±0.08	-1.05±0.06

Table 3: Performance Improvements of one of the backbone STGNNs on Multiple Traffic Datasets. We use STGCN in the table as an example to illustrate the significant enhancement of our method from the perspective of the dataset.

Dataset	15min			30min			60min			Average		
	MAE	RMSE	MAPE%									
PEMS04	-0.33±0.02	-0.40±0.03	-0.07±0.01	-1.45±0.09	-2.50±0.15	-0.10±0.01	-4.94±0.30	-6.61±0.40	-1.28±0.08	-2.24±0.14	-3.17±0.19	-0.48±0.03
PEMS08	-0.20±0.01	-1.09±0.07	-0.57±0.03	-1.37±0.08	-1.79±0.11	-0.30±0.02	-7.63±0.46	-5.86±0.35	-1.28±0.08	-3.07±0.18	-2.91±0.17	-0.72±0.04
PEMS03	-0.31±0.02	-0.51±0.03	-0.72±0.04	-0.22±0.01	-0.23±0.01	-0.60±0.04	-0.45±0.03	-0.35±0.02	-1.25±0.08	-0.33±0.02	-0.36±0.02	-0.86±0.05
PEMS07	-0.26±0.02	-0.29±0.02	-0.16±0.01	-0.38±0.02	-0.35±0.02	-0.31±0.02	-0.59±0.04	-0.59±0.04	-0.32±0.02	-0.41±0.03	-0.41±0.03	-0.26±0.02
METR-LA	-0.10±0.01	-0.36±0.02	-0.03±0.00	-0.33±0.02	-0.83±0.05	-0.80±0.05	-1.05±0.06	-1.97±0.12	-2.30±0.14	-0.49±0.03	-1.05±0.06	-1.04±0.06
PEMSBAY	-0.04±0.00	-0.06±0.00	-0.11±0.01	-0.10±0.01	-0.45±0.03	-0.23±0.01	-0.50±0.03	-1.14±0.07	-1.02±0.06	-0.21±0.01	-0.55±0.03	-0.45±0.03

state-of-the-art approaches. For traditional methods, LSTM achieves remarkable improvements with MAE reductions of 3.05, 4.52, and 10.07 at 15-min, 30-min, and 60-min horizons, respectively. The enhancement is particularly evident in long-term predictions, where STGCN shows MAE reductions of 0.33, 1.93, and 4.90 across different horizons. Even sophisticated models like D2STGNN and AGCRN benefit from our framework, with D2STGNN achieving consistent MAE reductions of 0.70, 0.75, and 1.61, while AGCRN shows stable improvements with an average MAE reduction of 1.02. These results validate the effectiveness of our framework in enhancing STGNNs through node-level low-rank adaptations.

The results in Table 3 demonstrate our framework’s generalization capability across diverse datasets. Using STGCN as an example, we observe substantial improvements on all datasets. On PEMS04, MAE reductions of 0.33, 1.45, and 4.94 are achieved for 15-min, 30-min, and 60-min horizons respectively, with an av-

erage reduction of 2.24. PEMS08 shows even stronger improvements with MAE reductions up to 7.63 for 60-min predictions. The effectiveness varies systematically across different prediction horizons, with longer-term predictions (60-min) consistently exhibiting more significant improvements. On datasets with distinct spatial structures (PEMS03, PEMS07, METR-LA, and PEMS-BAY), the framework maintains robust performance gains. These improvements across datasets with varying node counts confirm the framework’s versatile adaptability to different types of spatio-temporal data.

The superior performance can be attributed to two key aspects. First, our node-adaptive approach provides a precisely calibrated parameter space for fine-tuning predictions while maintaining computational efficiency through low-rank matrix factorization. Second, the additional parameters from low-rank matrices effectively capture complex spatio-temporal dependencies, including regional characteristics, temporal dynamics, and node interactions, all achieved within a compressed parameter space. This synergistic combination of adaptive capacity and computational efficiency enables our framework to enhance various baseline models consistently across diverse prediction scenarios.

#### 4.4 Efficiency and scalability Studies (RQ3)

**Time Efficiency.** In Figure 4(a), we analyze the computational overhead when integrating our framework with existing models. The results demonstrate that ST-LoRA introduces minimal additional training time while delivering substantial performance gains. With 16 NALLs and 4 MLRFs to enhance, the time increase remains remarkably efficient, adding a mere 0.68 seconds for D2STGNN. This efficiency stems from our low-rank adaptation strategy, which maintains stable training times even with multiple node-specific predictors.

**Framework Scalability.** We evaluate parameter efficiency across six baseline models using consistent node adaptive low-rank layer configurations. As shown

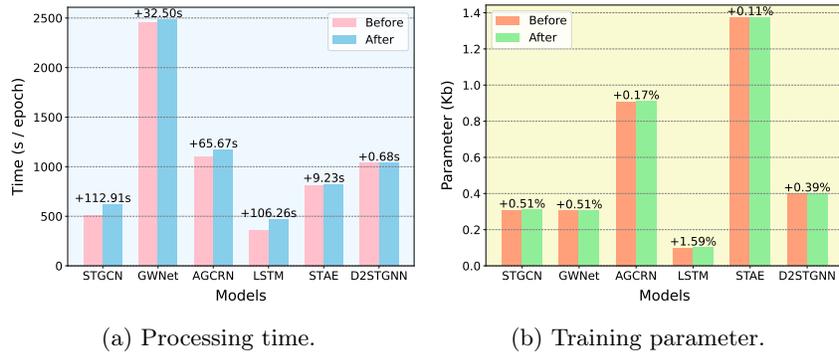


Fig. 4: Efficiency study comparing time and parameter cost of ST-LoRA.

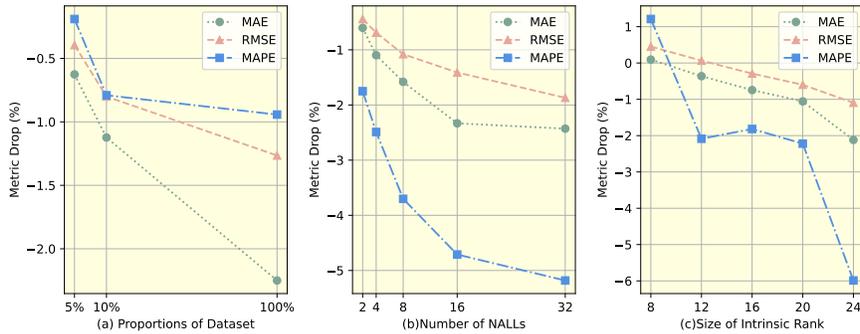


Fig. 5: Parameter sensitivity analysis examining the impact of varying dimensions, placements, and quantities of adaptive components on model performance.

in Figure 4(b), our approach achieves remarkable parameter efficiency, requiring less than 1% additional trainable parameters for most models while delivering 4.2% to 7.3% improvements in average RMSE. Notably, the smallest LSTM model maintains overhead below 2% while achieving an impressive 15% reduction in Average RMSE. This demonstrates ST-LoRA’s ability to effectively capture and adapt to node heterogeneity with minimal computational overhead.

#### 4.5 Parameter Sensitivity Analysis (RQ4)

In our experiments, we fixed the number of layers at 4 and the node embedding dimension at 12, which are the two most critical parameters. The node embedding dimension represents the rank of the low-rank matrix required for the additional parameters of the nodes, which tends to increase as more feature information is included in the data. After stacking multiple layers of NALL, the fine-tuning effect of these additional parameter spaces is amplified. It is important to design these two hyper-parameters in a balanced way because a larger parameter space does not necessarily mean it is easier to learn. We again use the STGCN model and the PEMS04 dataset as an example to explore the relationship between the size of this dataset as well as the two hyper-parameters and the lifting effect, and the results are shown in Figure 5.

## 5 Conclusion

In this paper, we focus on the challenges of node-level heterogeneity and over-parameterization in spatio-temporal modeling. We introduce ST-LoRA, a novel framework featuring Node-Adaptive Low-rank Layers and Node-Specific Predictors that efficiently customize parameters while maintaining computational efficiency. Our proposed approach enhances existing models by effectively capturing heterogeneous features and distributional changes within independent

nodes. The improvements demonstrate that efficient low-rank adaptation can significantly enhance forecasting in domains with heterogeneous node behaviors.

## 6 Limitation and Future work

Despite the promising improvements brought by ST-LoRA, the development and application of Low-rank adaptation techniques have so far been mostly concentrated within large models. When applied to specialized spatio-temporal models, adapting our method often requires substantial manual intervention, which collectively hinder the straightforward and scalable adoption of LoRA across heterogeneous spatio-temporal forecasting models. To overcome these limitations, we plan to continuously update and maintain a comprehensive benchmark that includes a wide variety of spatio-temporal models and datasets. Specifically, our future work will focus on: 1) developing a unified framework that enables flexible and automated adaptation of LoRA structures to different baseline models, supporting multiple LoRA variants for systematic and reproducible comparison; and 2) conduct extensive experiments across diverse real-world spatio-temporal scenarios, thereby providing the research community with a fair and comprehensive baseline for future studies.

## Acknowledgment

This work is supported by the National Natural Science Foundation of China (No. 62402414), the Guangdong Basic and Applied Basic Research Foundation (No. 2025A1515011994), and a grant from State Key Laboratory of Resources and Environmental Information System. This study is also supported by the Guangzhou Municipal Science and Technology Project (No. 2023A03J0011), the Guangzhou Industrial Information and Intelligent Key Laboratory Project (No. 2024A03J0628), and Guangdong Provincial Key Lab of Integrated Communication, Sensing and Computation for Ubiquitous Internet of Things (No. 2023B1212010007).

## References

1. Ahmed Agiza, Marina Neseem, and Sherief Reda. Mtlora: Low-rank adaptation approach for efficient multi-task learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16196–16205, 2024.
2. Lei Bai, Lina Yao, Can Li, Xianzhi Wang, and Can Wang. Adaptive graph convolutional recurrent network for traffic forecasting. *Advances in neural information processing systems*, 33:17804–17815, 2020.
3. Chao Chen, Karl Petty, Alexander Skabardonis, Pravin Varaiya, and Zhanfeng Jia. Freeway performance measurement system: mining loop detector data. *Transportation Research Record*, 1748(1):96–102, 2001.
4. Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling, 2014.

5. Zheng Dong, Renhe Jiang, Haotian Gao, Hangchen Liu, Jinliang Deng, Qingsong Wen, and Xuan Song. Heterogeneity-informed meta-parameter learning for spatiotemporal time series forecasting. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2024.
6. Rong Du, Paolo Santi, Ming Xiao, Athanasios V Vasilakos, and Carlo Fischione. The sensible city: A survey on the deployment and management for smart city monitoring. *IEEE Communications Surveys & Tutorials*, 21(2):1533–1560, 2018.
7. Zheng Fang, Qingqing Long, Guojie Song, and Kunqing Xie. Spatial-temporal graph ode networks for traffic flow forecasting. In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, pages 364–373, 2021.
8. Haotian Gao, Renhe Jiang, Zheng Dong, Jinliang Deng, Yuxin Ma, and Xuan Song. Spatial-temporal-decoupled masked pre-training for spatiotemporal forecasting. *arXiv preprint arXiv:2312.00516*, 2023.
9. Jiuxiang Gu, Zhenhua Wang, Jason Kuen, Lianyang Ma, Amir Shahroudy, Bing Shuai, Ting Liu, Xingxing Wang, Gang Wang, Jianfei Cai, et al. Recent advances in convolutional neural networks. *Pattern recognition*, 77:354–377, 2018.
10. Soufiane Hayou, Nikhil Ghosh, and Bin Yu. Lora+: Efficient low rank adaptation of large models. *arXiv preprint arXiv:2402.12354*, 2024.
11. Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
12. Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models, 2021.
13. Wenhao Huang, Guojie Song, Haikun Hong, and Kunqing Xie. Deep architecture for traffic flow prediction: deep belief networks with multitask learning. *IEEE Transactions on Intelligent Transportation Systems*, 15(5):2191–2201, 2014.
14. I-Hong Jhuo, Dong Liu, DT Lee, and Shih-Fu Chang. Robust visual domain adaptation with low-rank reconstruction. In *2012 IEEE conference on computer vision and pattern recognition*, pages 2168–2175. IEEE, 2012.
15. Renhe Jiang, Zhaonan Wang, Jiawei Yong, Puneet Jeph, Quanjun Chen, Yasumasa Kobayashi, Xuan Song, Shintaro Fukushima, and Toyotaro Suzumura. Spatio-temporal meta-graph learning for traffic forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pages 8078–8086, 2023.
16. Guangyin Jin, Yuxuan Liang, Yuchen Fang, Zezhi Shao, Jincai Huang, Junbo Zhang, and Yu Zheng. Spatio-temporal graph neural networks for predictive learning in urban computing: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 2023.
17. Ming Jin, Yifan Zhang, Wei Chen, Kexin Zhang, Yuxuan Liang, Bin Yang, Jindong Wang, Shirui Pan, and Qingsong Wen. Position: What can large language models tell us about time series analysis. In *Forty-first International Conference on Machine Learning*, 2024.
18. Rabeeh Karimi Mahabadi, James Henderson, and Sebastian Ruder. Compacter: Efficient low-rank hypercomplex adapter layers. *Advances in Neural Information Processing Systems*, 34:1022–1035, 2021.
19. Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks, 2016.
20. Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting, 2017.
21. Yuanzhi Li, Tengyu Ma, and Hongyang Zhang. Algorithmic regularization in over-parameterized matrix sensing and neural networks with quadratic activations. In *Conference On Learning Theory*, pages 2–47. PMLR, 2018.

22. Zhonghang Li, Chao Huang, Lianghao Xia, Yong Xu, and Jian Pei. Spatial-temporal hypergraph self-supervised learning for crime prediction. In *2022 IEEE 38th international conference on data engineering (ICDE)*, pages 2984–2996. IEEE, 2022.
23. Yuxuan Liang, Songyu Ke, Junbo Zhang, Xiuwen Yi, and Yu Zheng. Geoman: Multi-level attention networks for geo-sensory time series prediction. In *IJCAI*, volume 2018, pages 3428–3434, 2018.
24. Yuxuan Liang, Kun Ouyang, Yiwei Wang, Ye Liu, Junbo Zhang, Yu Zheng, and David S Rosenblum. Revisiting convolutional neural networks for citywide crowd flow analytics. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2020, Ghent, Belgium, September 14–18, 2020, Proceedings, Part I*, pages 578–594. Springer, 2021.
25. Yuxuan Liang, Yutong Xia, Songyu Ke, Yiwei Wang, Qingsong Wen, Junbo Zhang, Yu Zheng, and Roger Zimmermann. Airformer: Predicting nationwide air quality in china with transformers. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 14329–14337, 2023.
26. Marco Lippi, Matteo Bertini, and Paolo Frasconi. Short-term traffic flow forecasting: An experimental comparison of time-series analysis and supervised learning. *IEEE Transactions on Intelligent Transportation Systems*, 14(2):871–882, 2013.
27. Hangchen Liu, Zheng Dong, Renhe Jiang, Jiewen Deng, Jinliang Deng, Qianjun Chen, and Xuan Song. Spatio-temporal adaptive embedding makes vanilla transformer sota for traffic forecasting. In *Proceedings of the 32nd ACM international conference on information and knowledge management*, pages 4125–4129, 2023.
28. Xu Liu, Yutong Xia, Yuxuan Liang, Junfeng Hu, Yiwei Wang, Lei Bai, Chao Huang, Zhenguang Liu, Bryan Hooi, and Roger Zimmermann. Largest: A benchmark dataset for large-scale traffic forecasting. *Advances in Neural Information Processing Systems*, 36, 2024.
29. Yisheng Lv, Yanjie Duan, Wenwen Kang, Zhengxi Li, and Fei-Yue Wang. Traffic flow prediction with big data: A deep learning approach. *IEEE Transactions on Intelligent Transportation Systems*, 16(2):865–873, 2014.
30. Ying Ma, Haijie Lou, Ming Yan, Fanghui Sun, and Guoqi Li. Spatio-temporal fusion graph convolutional network for traffic flow forecasting. *Information Fusion*, 104:102196, 2024.
31. Junting Pan, Ziyi Lin, Xiatian Zhu, Jing Shao, and Hongsheng Li. St-adapter: Parameter-efficient image-to-video transfer learning. *Advances in Neural Information Processing Systems*, 35:26462–26477, 2022.
32. Zheyi Pan, Zhaoyuan Wang, Weifeng Wang, Yong Yu, Junbo Zhang, and Yu Zheng. Matrix factorization for spatio-temporal neural networks with applications to urban flow prediction. In *Proceedings of the 28th ACM international conference on information and knowledge management*, pages 2683–2691, 2019.
33. Tara N Sainath, Brian Kingsbury, Vikas Sindhwani, Ebru Arisoy, and Bhuvana Ramabhadran. Low-rank matrix factorization for deep neural network training with high-dimensional output targets. In *2013 IEEE international conference on acoustics, speech and signal processing*, pages 6655–6659. IEEE, 2013.
34. Zezhi Shao, Zhao Zhang, Fei Wang, and Yongjun Xu. 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*, pages 1567–1577, 2022.
35. Zezhi Shao, Zhao Zhang, Wei Wei, Fei Wang, Yongjun Xu, Xin Cao, and Christian S Jensen. Decoupled dynamic spatial-temporal graph neural network for traffic forecasting. *Proceedings of the VLDB Endowment*, 15(11):2733–2746, 2022.

36. Brian L Smith and Michael J Demetsky. Traffic flow forecasting: comparison of modeling approaches. *Journal of transportation engineering*, 123(4):261–266, 1997.
37. Mojtaba Valipour, Mehdi Rezagholizadeh, Ivan Kobyzev, and Ali Ghodsi. Dylora: Parameter efficient tuning of pre-trained models using dynamic search-free low-rank adaptation. *arXiv preprint arXiv:2210.07558*, 2022.
38. JWC Van Lint and CPlJ Van Hinsbergen. Short-term traffic and travel time prediction models. *Artificial Intelligence Applications to Critical Transportation Issues*, 22(1):22–41, 2012.
39. Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
40. Hanqiu Wang, Rongqing Zhang, Xiang Cheng, and Liuqing Yang. Hierarchical traffic flow prediction based on spatial-temporal graph convolutional network. *IEEE Transactions on Intelligent Transportation Systems*, 23(9):16137–16147, 2022.
41. Senzhang Wang, Jiannong Cao, and S Yu Philip. Deep learning for spatio-temporal data mining: A survey. *IEEE transactions on knowledge and data engineering*, 34(8):3681–3700, 2020.
42. Xiaoyang Wang, Yao Ma, Yiqi Wang, Wei Jin, Xin Wang, Jiliang Tang, Caiyan Jia, and Jian Yu. Traffic flow prediction via spatial temporal graph neural network. In *Proceedings of the web conference 2020*, pages 1082–1092, 2020.
43. Haixu Wu, Jiehui Xu, Jianmin Wang, and Mingsheng Long. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. *Advances in neural information processing systems*, 34:22419–22430, 2021.
44. Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems*, 32(1):4–24, 2020.
45. Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, and Chengqi Zhang. Graph wavenet for deep spatial-temporal graph modeling. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 1907–1913, 2019.
46. YiBo Yan, Haomin Wen, Siru Zhong, Wei Chen, Haodong Chen, Qingsong Wen, Roger Zimmermann, and Yuxuan Liang. Urbanclip: Learning text-enhanced urban region profiling with contrastive language-image pretraining from the web. In *Proceedings of the ACM on Web Conference 2024*, pages 4006–4017, 2024.
47. Huaxiu Yao, Fei Wu, Jintao Ke, Xianfeng Tang, Yitian Jia, Siyu Lu, Pinghua Gong, Jieping Ye, and Zhenhui Li. Deep multi-view spatial-temporal network for taxi demand prediction. In *Proceedings of the AAAI conference on artificial intelligence*, 2018.
48. Bing Yu, Haoteng Yin, and Zhanxing Zhu. Spatio-temporal graph convolutional networks: a deep learning framework for traffic forecasting. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 3634–3640, 2018.
49. Yong Yu, Xiaosheng Si, Changhua Hu, and Jianxun Zhang. A review of recurrent neural networks: Lstm cells and network architectures. *Neural computation*, 31(7):1235–1270, 2019.
50. Da Zhang and Mansur R Kabuka. Combining weather condition data to predict traffic flow: a gru-based deep learning approach. *IET Intelligent Transport Systems*, 12(7):578–585, 2018.
51. Junping Zhang, Fei-Yue Wang, Kunfeng Wang, Wei-Hua Lin, Xin Xu, and Cheng Chen. Data-driven intelligent transportation systems: A survey. *IEEE Transactions on Intelligent Transportation Systems*, 12(4):1624–1639, 2011.

52. Qi Zhang, Jianlong Chang, Gaofeng Meng, Shiming Xiang, and Chunhong Pan. Spatio-temporal graph structure learning for traffic forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 1177–1185, 2020.
53. Tian Zhou, Ziqing Ma, Qingsong Wen, Xue Wang, Liang Sun, and Rong Jin. Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting. In *International conference on machine learning*, pages 27268–27286. PMLR, 2022.
54. Tian Zhou, Peisong Niu, Liang Sun, Rong Jin, et al. One fits all: Power general time series analysis by pretrained lm. *Advances in neural information processing systems*, 36, 2024.
55. Eric Zivot and Jiahui Wang. Vector autoregressive models for multivariate time series. *Modeling financial time series with S-PLUS®*, pages 385–429, 2006.