



GraphSparseNet: a Novel Method for Large Scale Traffic Flow Prediction

Weiyang Kong
Sun Yat-Sen University
Guangzhou, China
kongwy3@mail2.sysu.edu.cn

Sen Zhang
Sun Yat-Sen University
Guangzhou, China
zhangs7@mail2.sysu.edu.cn

Kaiqi Wu
Sun Yat-Sen University
Guangzhou, China
wukq5@mail2.sysu.edu.cn

Yubao Liu
Sun Yat-Sen University
Guangdong Key Laboratory of Big Data Analysis and
Processing
Guangzhou, China
liuyubao@mail.sysu.edu.cn

ABSTRACT

Traffic flow forecasting is a critical spatio-temporal data mining task with wide-ranging applications in intelligent route planning and dynamic traffic management. Recent advancements in deep learning, particularly through Graph Neural Networks (GNNs), have significantly enhanced the accuracy of these forecasts by capturing complex spatio-temporal dynamics. However, the scalability of GNNs remains a challenge due to their exponential growth in model complexity with increasing nodes in the graph. Existing methods to address this issue, including sparsification, decomposition, and kernel-based approaches, either do not fully resolve the complexity issue or risk compromising predictive accuracy. This paper introduces GraphSparseNet (GSNet), a novel framework designed to improve both the scalability and accuracy of GNN-based traffic forecasting models. GraphSparseNet is comprised of two core modules: the Feature Extractor and the Relational Compressor. These modules operate with linear time and space complexity, thereby reducing the overall computational complexity of the model to a linear scale. Our extensive experiments on multiple real-world datasets demonstrate that GraphSparseNet not only significantly reduces training time by 3.51x compared to state-of-the-art linear models but also maintains high predictive performance.

PVLDB Reference Format:

Weiyang Kong, Kaiqi Wu, Sen Zhang, and Yubao Liu. GraphSparseNet: a Novel Method for Large Scale Traffic Flow Prediction. PVLDB, 18(7): 2295 - 2307, 2025.
doi:10.14778/3734839.3734862

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/PolynomeK/GSNet>.

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.
Proceedings of the VLDB Endowment, Vol. 18, No. 7 ISSN 2150-8097.
doi:10.14778/3734839.3734862

1 INTRODUCTION

Traffic flow forecasting represents a quintessential spatio-temporal data mining challenge, with profound utility in a spectrum of real-world applications, including intelligent route planning, dynamic traffic management, and smart location-based services [34]. The objective of this endeavor is to prognosticate future traffic patterns by leveraging historical traffic data, typically gleaned from the sensors embedded within transportation networks. The advent of deep learning has revolutionized this domain, with deep learning-based techniques, particularly those grounded in graph neural networks (GNNs), emerging as a dominant force. GNNs excel in capturing the intricate nonlinear dynamics inherent in spatio-temporal datasets, with their efficacy underpinned by the natural alignment between traffic data structures and graph-theoretic principles [14, 32]. In these models, graph nodes correspond to traffic sensors, while edges delineate the interconnections among these sensors.

Despite the predictive prowess of GNN-based methodologies, they are not without their drawbacks, most notably the exponential growth in model complexity. The number of edges in a graph tends to increase exponentially with the number of nodes, posing a significant challenge for the scalability of GNNs to larger datasets. This challenge is exacerbated by the proliferation of sensors within traffic networks, driven by urban development and the pervasive integration of Internet of Things (IoT) technologies. The vast amounts of spatio-temporal data generated by these sensors pose a challenge in applying higher-accuracy GNN methods to larger-scale datasets.

Several strategies have been proposed to expedite the computational efficiency of GNNs. For instance, the AGS [5] introduces a sparsification technique that prunes a trained model, thereby reducing its complexity during the inference phase. However, this approach still incurs substantial computational overhead during the training phase. Other methods, such as GWNet and AGCRN [1, 36], employ Tucker decomposition to construct graph adjacency matrices with a reduced parameter count. Yet, these techniques merely curtail the number of trainable parameters in the adjacency matrix, without alleviating the computational burden associated with graph operations. The computational complexity of adjacency matrix multiplication, for example, remains $O(N^2)$, where N denotes the number of nodes. While some methods have successfully simplified GNNs to enhance scalability, they are not without their




Figure 1: Visualizations correspond to the well-trained adjacency matrix obtained from existing graph neural network method. It can be observed that matrix is highly sparse. Statistical data on the weighted degree also demonstrates the sparsity of the adjacency relationships.

limitations. The BigST [9], which employs a kernel-based approach to simplify GNNs, is susceptible to aberrant gradient values during backpropagation, potentially leading to suboptimal training outcomes and diminished predictive accuracy.

This paper aims to improve the scalability of GNNs for large-scale traffic spatio-temporal data by designing a high-precision, scalable model. This is a non-trivial problem because simplifying GNN models faces two challenges: First, most approaches treat accuracy and scalability in traffic prediction models as two orthogonal issues. Accurately capturing all node relationships leads to an exponential rise in computational demands. As the number of graph nodes increases, generating the adjacency matrix and performing related operations result in a sharp increase in computational complexity. Second, existing methods for simplifying GNNs have significant limitations. Decomposition techniques fail to reduce overall model complexity, sparsification methods cannot be applied during training, and kernel-based approaches risk compromising accuracy. This highlights the need for new strategies to simplify GNNs effectively.

Upon reviewing and observing existing studies, we find that in well-trained adjacency matrices, the relationships between nodes tend to be sparse, with only a small subset of nodes being connected. To support this claim, we visualized a trained adaptive adjacency matrix [1] and computed the weighted degree of each node, as shown in Figures 1 (a) and 1 (b). This trend is consistent across several models based on adaptive matrices [10, 35, 36]. In Figure 1 (a), higher values in the matrix indicate stronger associations between the corresponding node pairs. It is evident that only a small fraction of the matrix entries have significant values, suggesting that the relationships between nodes are highly sparse. In Figure 1 (b), we compute the weighted degree of each node, which reflects both the number and strength of connections for each node in the graph. As observed, with the exception of a few nodes with higher weighted degrees, most nodes maintain relatively low weighted degrees, indicating that the graph has a low connection density, with the majority of nodes having small numbers and weights of adjacent edges. Existing graph methods learn one entire adjacency matrix to express relationships between nodes, leading to the use of a large number of parameters and computations to learn a highly sparse matrix. We believe that the generation of graphs and related

operations can be compressed onto a much smaller space, thereby "circumventing" the computation of all node relationships in the complete graph. We theoretically prove the rationality of this motivation, providing a theoretical derivation to show that learning one entire adjacency matrix can be replaced by learning two very small-scale matrices.

Drawing on our comprehensive analysis and key insights, we introduce a groundbreaking model, GraphSparseNet (GSNet), designed to deliver high predictive accuracy concurrently with scalability to handle large-scale datasets. This innovative model is comprised of two synergistic modules: the Feature Extractor, which is tasked with capturing and encoding the characteristics of graph nodes, and the Relational Compressor, which is responsible for modeling the sparse relationships between nodes. Both modules are engineered to operate with linear time and space complexity, scaling efficiently with an increase in the number of graph nodes. This design enables the model to effectively manage and analyze large-scale data without compromising performance. GraphSparseNet has reduced the complexity from $O(N^2)$ in most existing methods to $O(N)$. Even compared to the state of the art linear models, our method has improved training time by 3.51×. We evaluated the effectiveness and scalability of our proposed framework on multiple real-world datasets from different regions. The results of these experiments consistently demonstrate the superior performance of GraphSparseNet, achieving commendable results across all datasets evaluated. The principal contributions of our work are encapsulated in the following points:

- By conducting a thorough theoretical analysis, we have identified key factors that restrict existing methods and propose innovative solutions to address these limitations. Our approach not only refines the predictive capabilities of GNNs but also significantly enhances their ability to scale with the increasing nodes of traffic data.
- We present GraphSparseNet, a novel framework specifically tailored to address the scalability challenges of large-scale traffic spatio-temporal data. This framework is underpinned by two core modules: the Feature Extractor and the Relational Compressor. The Feature Extractor capture and represent the features of each node within the graph, while the Relational Compressor innovatively models the sparse relationships between nodes. Both modules are designed to operate with linear time and space complexity.
- We conducted experiments on four real-world datasets from different regions, and the results demonstrate the effectiveness and efficiency of our approach.

2 PRELIMINARY

2.1 Traffic Network

Traffic network is represented by a undirected graph $G = (V, E)$, where V is the set of nodes (sensors), $N = |V|$ denotes the number of nodes, and E is the set of edges between two nodes. In our problem, we assume that each node records its traffic flow data as graph signal. A graph signal is $X^t \in \mathbb{R}^N$, where t denotes the t -th time step. The graph signal represents the traffic flow values at the t -th time step.

Table 1: Notation

Notation	Meaning
G	The undirected graph.
V	The set of nodes (sensors).
E	The set of edges between nodes.
N	The number of nodes.
A	The adjacency matrix of graph G .
X	The graph signal.
S	The number of historical time step of X .
T	The number of predict time step of output.
H	The hidden states in model.
C	The dimension of the compressed space.
K	The adjacency matrix in low-dimensional space.
U	The coefficient matrix in low-dimensional space.
W, B	The trainable parameter of the pivotal graph convolution.
V	The transformation matrix.

2.2 Problem Definition

Given a traffic network G and its historical S step graph signal matrix $X^{1:S} = (X^1, X^2, \dots, X^S) \in \mathbb{R}^{N \times S}$, our problem is to predict its next T step graph signals, namely $X^{S+1:S+T} = (X^{S+1}, X^{S+2}, \dots, X^{S+T}) \in \mathbb{R}^{N \times T}$. We equate the problem as finding a function \mathcal{F} to forecast the next T steps data based on the past S steps data:

$$(X^{S+1}, X^{S+2}, \dots, X^{S+T}) = \mathcal{F}((X^1, X^2, \dots, X^S)). \quad (1)$$

3 ANALYSIS

In this section, we discuss the constraints that hinder the scalability of graph neural network methods. We also explain our motivations for simplifying models, supported by theoretical evidence and empirical findings.

Graph neural network methodologies are highly effective in traffic prediction, as they adeptly capture node features while facilitating the integration of features across different nodes in accordance with the graph's topology. The most prevalent graph neural networks can be represented in the following form:

$$H = \text{Agg}(A, X, \Theta) \quad (2)$$

Here, H represents the hidden state of the model, A denotes the adjacency matrix, X represents the input data, and Θ refers to the trainable parameters. The function $\text{Agg}(\cdot)$ is a specific method for aggregating elements in X based on the relationships defined in A . A commonly employed approach for this aggregation function is the use of spatial domain graph convolution operations in the following form:

$$H = \sigma(AWX) \quad (3)$$

Here, $\sigma(\cdot)$ represents arbitrary activation function, and W denotes the trainable parameters. It leverages matrix multiplication to integrate input data X in accordance with the relationships specified by the adjacency matrix A . In the majority of current research [1, 36], to achieve superior predictive performance, matrix A is often composed of trainable parameters, a practice referred to as the adaptive adjacency matrix. This approach allows the model to reduce its

reliance on prior knowledge while enhancing its representational capacity for potential node associations. To prevent over-fitting due to the extensive use of parameters, this adaptive matrix is typically factorized into the product of two smaller matrices:

$$A = \text{SoftMax}(E_1 E_2) \quad (4)$$

Here, $E_1 \in \mathbb{R}^{N \times C}$ and $E_2 \in \mathbb{R}^{C \times N}$ are both trainable parameters, C is a hyperparameter. This technique reduces the number of parameters required to construct the adjacency matrix from N^2 to $2CN$. Employing this technique, a series of methods have empirically demonstrated that the effectiveness of the adaptive adjacency matrix can be maintained even when the parameter C is set to a value significantly smaller than N . When N is sufficiently large, employing a decomposition approach can significantly reduce the number of training parameters needed to construct the adjacency matrix A . However, the complexity of the matrix multiplication for AX in graph convolution is still $O(N^2)$, which greatly limits the scalability of this operation on large-scale data.

AGS [5] proposed the use of a mask matrix to simplify the adjacency matrix A . This approach introduces a Boolean mask matrix β , which has the same dimensions as the adjacency matrix A , after A has been trained. By performing the Hadamard product of A and β , the final adjacency matrix is determined based on the binary gating within β , deciding whether to retain the connections. However, this method can only theoretically reduce the operational speed during the model's inference phase, as it requires the retraining of β to achieve a sparse adjacency matrix after A has been trained. Simultaneously training A and β can lead to difficulties in model convergence, thereby diminishing the model's predictive accuracy.

BigST [9] designed a graph convolutional model with linear complexity. However, BigST can lead to the generation of anomalous gradient values during the model's training process, which affects the training effectiveness and the ultimate predictive accuracy. For graph convolution operations under normal circumstances, The Equation 4 can be rewritten as

$$A_{ij} = \frac{f(s_{ij})}{\sum_{k=1}^N f(s_{ik})} \quad (5)$$

$$s_{ij} = E_{1i} E_{2j}, \quad f(x) = \exp(x) \quad (6)$$

The gradients of the matrix A is derived as [27]:

$$\frac{\partial A_{ij}}{\partial s_{ik}} = \frac{f'(s_{ik})}{f(s_{ik})} (1_{j=k} A_{ij} - A_{ij} A_{ik}) \quad (7)$$

If Equation 6 is substituted into Equation 7, we can deduce that:

$$\left| \frac{\partial A_{ij}}{\partial s_{ik}} \right| \leq \frac{1}{4} \quad (8)$$

It is evident that the gradients of the parameters do not exhibit any significant anomalies. BigST employs a kernel function mapping approach to simplify the graph convolution operation into a linear product of multiple matrices, as follows:

$$s_{ij} = \phi(E_{1i}) \phi(E_{2j}), \quad f(x) = x \quad (9)$$

Here $\phi(\cdot)$ is the kernel function. If Equation 9 is substituted into Equation 7, we can deduce that:

$$\left| \frac{\partial A_{ij}}{\partial s_{ik}} \right| \leq \frac{1}{4|s_{ik}|} \quad (10)$$

There is always a probability that the gradients during back propagation will produce anomalous values since s_{ij} is set from 0 to positive infinity in the BigST.

Our goal is to approximate the graph convolution operation in Equation 3 with a computational complexity of $O(N)$, while maintaining the model's accuracy. This is feasible because, as shown in Figure 1, well-trained adaptive matrices tend to be highly sparse. For the factorized adjacency matrix $A = \text{SoftMax}(E_1 E_2)$, we consider the properties of matrices E_1 and E_2 that:

$$\text{Rank}(E_1) \leq C \quad (11)$$

$$\text{Rank}(E_2) \leq C \quad (12)$$

$$\text{Rank}(E_1 E_2) \leq \min(\text{Rank}(E_1), \text{Rank}(E_2)) \leq C \quad (13)$$

This implies that the column (and row) space of the matrix $E_1 E_2$ can both be spanned by a set of C linearly independent vectors. In other words, we have the following theorem:

THEOREM 3.1. *Let $M \in \mathbb{R}^{N \times N}$ be a matrix with rank C . There always exists a non-unique matrix $K \in \mathbb{R}^{C \times C}$ such that matrix M can be constructed via some matrix multiplication transformations involving K .*

Due to space limitations, the proof of this theorem can be found in the full version [17].

For the sake of convenience in subsequent descriptions, we assume, based on the aforementioned theorem, that:

$$E_1 E_2^{ij} = \sum_{m=1}^C \sum_{n=1}^C U_{mn}^{ij} K_{mn}, \quad (14)$$

here U represents the coefficient of the linear combination, and K_{mn} denotes the elements of matrix K with rank C . Thus, we can consider that matrix $E_1 E_2$ can be derived from matrix K and coefficient matrix U through a series of linear transformations. In other words, the process of the model learning matrix $E_1 E_2$ is equivalent to learning matrix K and the coefficient matrix U .

Since the well-trained adjacency matrix A is sparse, we propose that a significant number of values in U can be replaced with a fixed number, meaning that the number of trainable parameters used to express the combination coefficients in matrix U can be much smaller than its shape. The adjacency matrix A used in existing model is obtained by applying the activation function to $E_1 E_2$. Taking the most commonly used form, $\text{SoftMax}()$, as an example, we have:

$$\text{SoftMax}(E_1 E_2^{ij}) = \prod_{i=1}^C \prod_{j=1}^C \text{SoftMax}(K_{ij})^{u_{ij}} \quad (15)$$

If A_{ij} in the adjacency matrix is zero, we only need to set the corresponding coefficient in U to a fixed negative value (for example, negative 3) to ensure that the calculated A_{ij} is close to zero. For a sparse adjacency matrix A , most elements in U would be set to negative 3, with only a small portion of elements needing to be determined through training. The proportion of trainable parameters in matrix U is positively correlated with the sparsity of matrix A . Similarly, if other activation functions are chosen, we simply need

to set the corresponding elements in U to other fixed values since the computation of the activation function is fixed.

Based on the aforementioned analysis, we can consider that in the adaptive matrix based methods, the technique of setting A as a trainable parameter can be replaced by allowing the model to learn matrix K and the linear combination coefficient matrix U . Intuitively, the information contained in matrix A is equivalent to the information contained in K and U .

Furthermore, in the graph convolution operation of Equation 3, the intention of the matrix multiplication for AX is to perform feature fusion of the input data X according to the adjacency relationship A through linear transformation. The complexity of this operation is $O(N^2)$, which limits the scalability of the graph convolution operation. Since we can learn matrix K and the combination coefficients U separately, in this paper, we propose a new framework to replace the graph convolution operation in Equation 3. The following sections will introduce the specific details of our newly proposed method.

4 PROPOSED MODEL

we commence by elucidating the overarching framework of the model. Subsequently, we introduce the specific module designs within the model.

4.1 Framework of the model

As depicted in Figure 2, our model comprises two distinct modules: Feature Extractor and Relational Compressor. The Feature Extractor compresses the input data into a low-dimensional space and subsequently reconstructs it, while simultaneously updating an embedding that represents the node features. The Relational Compressor transforms the input data into a low-dimensional space, models the implicit adjacency relationships between different nodes in the low-dimensional space, and ultimately reconstructs the data back into its original space. The low-dimensional spaces involved in the two modules are mutually aligned. Both the Feature Extractor and Relational Compressor are stacked in a serial manner, and the outputs from each layer are concatenated with skip connections, which are then transformed into prediction through an output layer.

In our model, there are two types of embedding matrices. One is the input embedding $P \in \mathbb{R}^{N \times d}$, and the other is the node embedding $Q \in \mathbb{R}^{N \times d}$. d denotes the number of channel in the embedding. The node embedding is composed of trainable parameters, the elements of which are updated exclusively during the training process of the model and do not vary with the input data. In contrast, the input embedding is generated from the model's input data and automatically adjusts according to different inputs. The purpose of the node embedding is to model the characteristics of each node in the transportation network during the training process of the model. The input embedding encompasses not only the node feature information but also the spatial information (such as the relationship between nodes) and temporal information present in the input data. The objectives of transforming the input data into input embeddings are twofold: first, to obtain a dense representation for spatio-temporal features, and second, to enable the model to automatically learn the most salient features of the data.

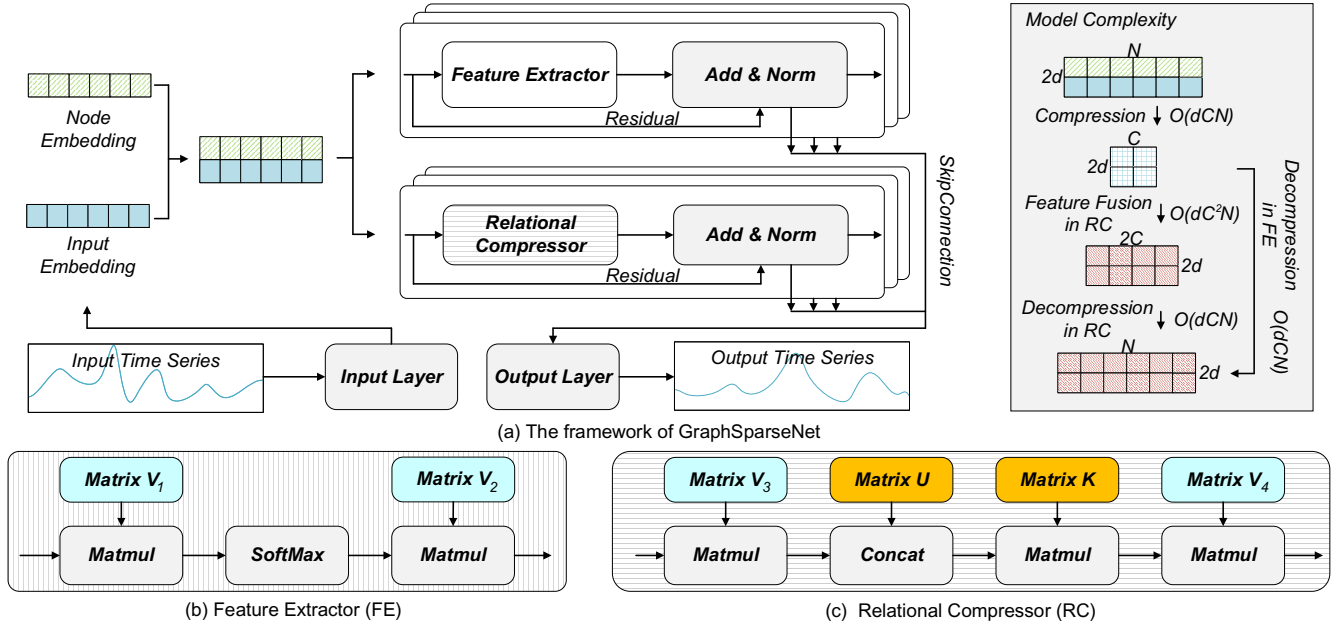


Figure 2: An illustration of the framework.

4.2 Feature Extractor

The advantages of Graph Neural Networks stem from two aspects: one is their powerful feature learning capability, which can accurately capture the features of each node, and the other is their ability to handle graph-structured data, allowing node features to be integrated based on adjacency relationships. By aggregating information from neighboring nodes to update the representation of a node, it can capture complex relationships between nodes, thereby better understanding graph-structured data. In the Feature Extractor module, our primary goal is to enable the model to learn the features of different nodes.

In the design of the Feature Extractor, we have two objectives: one is to update the node embeddings through this module to model the characteristics of each node in the traffic data, and the other is to attempt to compress the input embeddings, which contain spatio-temporal information, to a low-dimensional space, thereby preparing for subsequent feature fusion in low-dimensional space within the model.

To achieve both of these objectives, we have designed a module that compresses and then decompresses the input information based on node features. The input embeddings encompass only the local spatio-temporal characteristics of the input time series, as the data presented to the model at each instance constitutes a small slice of the entire dataset. Concurrently, the node embeddings remain invariant over time, thereby capturing and representing global information. Therefore, we concatenate the two types of embeddings as the input to the module.

$$X_{FE} = P || Q \quad (16)$$

Here, X_{FE} represents the input of Feature Extractor. The concatenated input X_{FE} contains both the local spatio-temporal information of the input data and the global features of each node. A reasonable approach to compressing data from an N -dimensional space to a low-dimensional space is to cluster based on node features. To this end, we use a matrix generated from the node features to compress the module input, as follows:

$$H_{FE} = V_1 X_{FE} \quad (17)$$

$$V_1 = W_1 Q + B_1 \quad (18)$$

Here, H_{FE} represents the hidden state, V_1 is the compressing matrix. W_1 and B_1 are trainable parameter matrices. Subsequently, we designed the decompression process, which is similar to the compression. To introduce non-linear characteristics and control in the Feature Extractor, we added an activation function between the compression and decompression processes. We utilize the SoftMax function to present the hidden layer in the form of a probability distribution, thereby enhancing the training effectiveness of this module during the decompression process.

$$O_{FE} = V_2 \text{SoftMax}(H_{FE}) \quad (19)$$

$$V_2 = W_2 Q + B_2 \quad (20)$$

On one hand, the process of compression and decompression can update the a node embedding Q to learn the features of different nodes. The information contained within this embedding will be shared with Relational Compressor. On the other hand, the model is designed to learn a method of compressing the input to a size of C dimensions, which serves as an auxiliary function for Relational Compressor to map the input into a low-dimensional space. This alignment of the compression ratio in this module with that in Relational Compressor is crucial for the seamless integration and effective operation of the overall model.

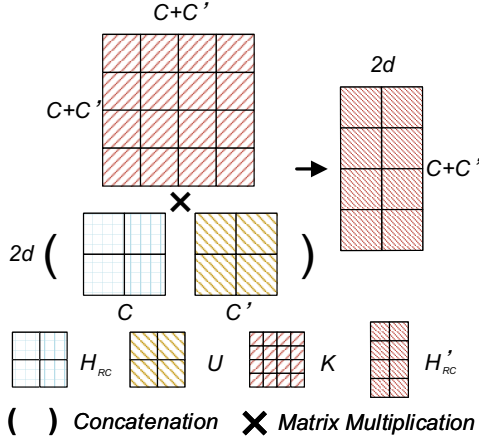


Figure 3: An illustration in Relational Compressor.

4.3 Relational Compressor

As we discussed in the previous analysis, learning the adjacency matrix A can be equivalent to learning matrices K and U . In the design of the Relational Compressor, our objective is to learn the adjacency matrix K and the coefficient matrix U of the traffic graph in the low-dimensional space, and to achieve the fusion of node features in the low-dimensional space to avoid the $O(N^2)$ complexity in spatial graph convolutions. To this end, we also design the model with a compression and decompression manner, which is similar to Feature Extractor. Consistent with the concept in the Feature Extractor, this module also takes the concatenation of node embeddings and input embeddings as input:

$$H_{RC} = V_3(P||Q) \quad (21)$$

Here H_{RC} is the hidden state in Relational Compressor. Unlike the Feature Extractor, in the compression process, to enable the compression to consider the local spatio-temporal features in the input embedding P , rather than compressing solely based on global node features, we consider using two embeddings to generate the compression matrix:

$$V_3 = W_3(P||Q) + B_3 \quad (22)$$

After compressing the input into the low-dimensional space, we let the model learn the coefficient matrix U in the hidden state by means of concatenation.

$$H'_{RC} = H_{RC}||U \quad (23)$$

As shown in Figure 3, the shape of the coefficient matrix U can be controlled by C' . Subsequently, we fuse features according to the low-dimensional adaptive adjacency matrix K through matrix multiplication in the hidden state, as follows:

$$H''_{RC} = KH'_{RC} \quad (24)$$

K is the adaptive adjacency matrix in the $C + C'$ -dimensional space. Both K and U are composed of trainable parameters, and in Section 2, we discussed the feasibility of using K and U to replace the adaptive matrix A . Through concatenation and matrix multiplication, Relational Compressor can achieve the purpose of learning adjacency relationships and feature fusion in the low-dimensional

space. The number of trainable parameters contained in matrices K and U is far less than that in A . Moreover, the computational complexity of the module is also far less than that of spatial GCN and its existing improved schemes, which we will analyze in detail later. Subsequently, we architected the decompression process to be analogous to the compression process, ensuring a symmetrical approach to feature representation.

$$O_{RC} = V_4 \text{SoftMax}(H''_{RC}) \quad (25)$$

$$V_4 = W_4(P||Q) + B_4 \quad (26)$$

The outputs of the two modules are summarized through a skip-connection approach and the final prediction result is obtained through the output layer. The final output is:

$$O = O_{FE} || O_{RC} \quad (27)$$

The output layer consists of an activation function and a fully connected layer. The overall training process of our model is outlined in Algorithm 1.

4.4 Complexity Analysis

In this section, we analyze the complexity of various modules within the proposed model. For the sake of simplicity, we assume that $C = C' = d$ in the subsequent discussions.

Regarding time complexity: The complexity of converting the input graph signal to input embeddings in the input layer is $O(NSC)$, where N is the number of nodes, S is the number of historical time steps, and C is the dimensionality of the low-dimensional space. In the Feature Extractor, the complexity of compression and decompression is $O(C^2N + C^2N)$. In the Relational Compressor, the complexity of compression and decompression is $O(C^2N + C^2N)$, and the complexity of feature fusion in the low-dimensional space is $O(C^3)$. The complexity of the output layer is $O(NTC)$, where T is the time step of the output graph signal. Since C , S , and T are constants, the overall time complexity of the model is $O(N)$.

Regarding space complexity: The space complexity for input embeddings and node embeddings is $O(SN + CN)$. The space complexity for generating the compression and decompression transformation matrices is $O(CN)$. The space complexity for generating K and U is $O(CN + C^2)$. The overall space complexity of the model is $O(N)$.

It can be observed that our approach exhibits a space complexity and time complexity of $O(N)$, which endows our method with superior scalability.

5 EXPERIMENTS

In this section, we evaluated our proposed model by empirically examining on four real-world datasets with the state-of-the-art models for traffic forecasting. To support the reproducibility of the results in this paper, we have released our code on website.¹

5.1 Datasets and Pre-processing

We conduct experiments on four widely used real-world public traffic datasets from: (1)PEMS²(PEMS07 and PEMS08), (2)England³,

¹<https://github.com/PolynomeK/GSNet>

²<http://pems.dot.ca.gov/>

³<http://tris.highwaysengland.co.uk/detail/trafficflowdata>

Algorithm 1 Training Schema of GraphSparseNet

```

1: Input: Graph  $G$ , Graph signal  $X$ 
2: Output: Model prediction  $\hat{X}$ 
3:
4:  $P \leftarrow \text{InputLayer}(X)$   $\triangleright$  Transform  $X$  to input embedding  $P$ 
5:  $Q \leftarrow \text{NodeEmbedding}(G)$   $\triangleright$  Obtain node embeddings  $Q$ 
6:
7:  $V_1, V_2 \leftarrow \text{GenerateTransformationMatrices}(Q)$   $\triangleright$  Generate transformation matrices
8:  $V_3, V_4 \leftarrow \text{GenerateTransformationMatrices}(P||Q)$   $\triangleright$  Generate transformation matrices from concatenation
9:
10: Modules: FE (Feature Extractor), RC (Relational Compressor)
11: for  $i \leftarrow 1$  to  $L$  do  $\triangleright$  Loop for  $L$  layers
12:   FE Module:
13:    $H_{FE} \leftarrow V_1 \cdot (P||Q)$   $\triangleright$  Compress the concatenation with  $V_1$ 
14:    $H_{FE} \leftarrow \text{ActivationFunction}(H_{FE})$   $\triangleright$  Apply activation function
15:    $H_{FE} \leftarrow V_2 \cdot H_{FE}$   $\triangleright$  Decompress with  $V_2$ 
16:    $O_{FE} \leftarrow \text{ResidualConnection}(P||Q, H_{FE})$   $\triangleright$  Residual connection and normalization
17:   RC Module:
18:    $H_{RC} \leftarrow V_3 \cdot (P||Q)$   $\triangleright$  Compress the concatenation with  $V_3$ 
19:    $H'_{RC} \leftarrow H_{RC}||U$   $\triangleright$  Concatenate with matrix  $U$ 
20:    $H''_{RC} \leftarrow K \cdot H'_{RC}$   $\triangleright$  Transform with matrix  $K$ 
21:    $H''_{RC} \leftarrow \text{ActivationFunction}(H''_{RC})$   $\triangleright$  Apply activation function
22:    $H''_{RC} \leftarrow V_4 \cdot H''_{RC}$   $\triangleright$  Decompress with  $V_4$ 
23:    $O_{RC} \leftarrow \text{ResidualConnection}(P||Q, H''_{RC})$   $\triangleright$  Residual connection and normalization
24:    $O \leftarrow \text{UpdateInput}(O_{FE}, O_{RC})$   $\triangleright$  Update input for next layer
25: end for
26:  $Output \leftarrow \text{FullyConnectedLayer}(O)$   $\triangleright$  Final fully connected layer
27:  $\hat{X} \leftarrow \text{ActivationFunction}(Output)$   $\triangleright$  Apply activation function for prediction
28: Backpropagation:  $\text{Compare}(\hat{X}, X_{\text{true}})$   $\triangleright$  Compare prediction with true value

```

and (3)CA [23]. The biggest dataset is California (CA), including a total number of 8,600 sensors. To the best of our knowledge, CA is currently the publicly available dataset with the highest number of nodes recorded by loop detectors. It contain three representative areas: Greater Los Angeles, Greater Bay Area, and San Diego. A brief description are given in Table 2.

Z-score normalization is applied to inputs as follows.

$$X_{\text{input}} = \frac{X - \text{mean}(X)}{\text{std}(X)} \quad (28)$$

Here $\text{mean}(X)$ and $\text{std}(X)$ denote the mean and the standard deviation of the historical time series, respectively.

Table 2: Datasets description.

Dataset	Node	Samples	Sample Rate	Traffic Record
PEMS08	170	17,856	5min	3,035,520
England	314	17,353	15min	5,448,842
PEMS07	883	28,224	5min	24,921,792
CA	8,600	17,280	5min	148,608,000

5.2 Baselines

We compare the proposed method in this paper with current state-of-the-art approaches in the field. The selected baselines are categorized into five groups based on their model frameworks. The first category includes classical statistical methods, specifically ARIMA [33]. The second category consists of CNN-based approaches, such as STResNet [41], ACFM [22] and STGCN [39]. The third category focuses on GNN-based methods, with some incorporating RNN structures, represented by DCRNN [19], GWNet [36], AGCRN [1], AGS [5] and BigST [9]. The fourth category includes transformer-based methods, denoted by PDFormer [13] and Bi-STAT [3]. Lastly, the fifth category includes large model (pretrained neural networks) approaches, denoted as GPT-ST [21] and UniST [40].

5.3 Experimental Setup

Our implementation is based on python 3.12.1, torch 2.2.0, and numpy 1.26.4. tested on Ubuntu 20.04.3 (CPU: Intel(R) Xeon(R) CPU E5-2620 v4 @ 2.10GHz, GPU: GeForce RTX 3080 10GB). Tests on PEMS07, PEMS08 and CA, we adopt 60 minutes as the history time window, a.k.a. 12 observed data points ($S = 12$) are used to forecast traffic in the next 60 minutes ($T = 12$). Tests on England, we adopt 180 minutes as the history time window, a.k.a. 12 observed data points ($S = 12$) and forecast the traffic in the next 180 minutes ($T = 12$). In our model, both the Feature Extractor and Relational Compressor are implemented as three-layer stacked structures. Adam optimizer is used for training our model with an initial learning rate of 0.0001. The best parameters are chosen through a carefully parameter-tuning process on the validation set. The parameters used in the baselines are set to the default values from the original author’s publicly available code. On certain datasets, due to the increased data scale, we adjusted parameters such as the batch size to ensure the model could run without memory constraints. PEMS07 and PEMS08 are split by ratio of 7:1:2 in chronological order while England and CA are split by ratio of 6:2:2. For all datasets, the missing values are filled by the linear interpolation. In our model, the values for C , d , and C' are set to be equal. The batch size for the input data is set to 16. Should we encounter issues with insufficient GPU memory, we will reduce the batch size of the model incrementally until the model can operate without such constraints. We adopt three evaluation metrics commonly used in the field of traffic forecasting to evaluate the performance of models: Mean Absolute Errors (MAE) represents the average absolute value of prediction errors. Mean Absolute Percentage Errors (MAPE) measures the percentage of prediction errors. Root Mean Squared Errors (RMSE) is the square root of the average squared difference between the actual and predicted values for each data point, providing an overall measure of prediction error.

Table 3: Performance comparison of different approaches on all the datasets.

Dataset	PEMS08			England			PEMS07			CA		
Metric	MAE	MAPE	RMSE	MAE	MAPE	RMSE	MAE	MAPE	RMSE	MAE	MAPE	RMSE
ARIMA	31.23	19.25	33.47	4.23	5.72	7.68	33.89	17.60	46.38	34.26	21.35	44.68
STResNet	23.25	15.58	32.25	4.03	5.68	7.55	29.36	15.24	42.46	29.25	20.33	40.63
ACFM	15.86	10.13	25.34	3.52	5.28	7.86	25.86	11.83	39.03	26.38	19.24	37.86
STGCN	17.50	11.29	27.03	3.55	5.30	7.38	25.32	11.16	39.27	25.68	18.43	36.44
DCRNN	17.86	11.45	27.83	3.59	4.90	7.42	25.3	11.66	38.58	25.75	18.62	36.91
GWNet	19.13	12.68	31.05	3.53	4.93	7.57	26.85	12.12	42.78	24.73	17.46	35.86
AGCRN	15.95	10.09	25.22	3.32	5.01	7.33	22.37	9.12	36.55	25.03	17.93	36.23
PDFormer	13.58	9.05	23.51	3.46	4.97	7.46	20.42	8.86	32.87	25.22	18.96	35.43
Bi-STAT	<u>13.62</u>	<u>9.43</u>	<u>23.17</u>	3.42	4.89	7.54	21.13	8.98	33.86	25.36	19.03	35.27
GPT-ST	14.85	9.63	24.32	3.43	4.97	7.42	21.43	9.32	34.76	24.93	17.52	35.86
UniST	16.01	10.23	25.37	3.49	4.88	7.51	24.62	11.56	38.21	25.53	18.87	35.46
AGS	15.50	9.71	25.01	3.15	4.56	7.52	21.56	9.03	34.90	24.96	17.88	36.02
BigST	16.35	10.05	24.37	3.63	5.13	7.76	22.86	10.02	34.03	23.36	16.81	34.98
GSNet(Ours)	14.76	9.48	22.52	<u>3.20</u>	<u>4.82</u>	6.60	<u>20.70</u>	8.77	32.66	19.76	14.39	30.99

For each data point, the absolute error between the actual and predicted values is calculated, and then the average is taken across all data points.

$$MAE = \frac{1}{TN} \sum_{i=1}^T \sum_{j=1}^N |\hat{X}_j^{(t+i)} - X_j^{(t+i)}| \quad (29)$$

Here \hat{X} denotes output of models.

Mean Absolute Percentage Errors (MAPE) measures the percentage of prediction errors. It calculates the relative error between the actual and predicted values for each data point, takes the average across all data points, and expresses the result as a percentage.

$$MAPE = \frac{100\%}{TN} \sum_{i=1}^T \sum_{j=1}^N \left| \frac{\hat{X}_j^{(t+i)} - X_j^{(t+i)}}{X_j^{(t+i)}} \right| \quad (30)$$

Root Mean Squared Errors (RMSE) is the square root of the average squared difference between the actual and predicted values for each data point, providing an overall measure of prediction error.

$$RMSE = \sqrt{\frac{1}{TN} \sum_{i=1}^T \sum_{j=1}^N (\hat{X}_j^{(t+i)} - X_j^{(t+i)})^2} \quad (31)$$

5.4 Comparison with Existing Models

Tables 3, 4 and Figure 4 present the effectiveness and efficiency of multiple models across various datasets.

5.4.1 Performance Comparison. As shown in the table 3, our method demonstrates strong predictive performance across multiple datasets. Notably, on larger-scale datasets such as CA, it achieves the best results across all metrics. On the PEMS07 and England dataset, our approach consistently achieves either the best or second-best performance across various evaluation metrics. Additionally, On the PEMS08 dataset, our method outperforms others specifically in the RMSE metric. Traditional statistical methods such as ARIMA are not effective in handling complex spatio-temporal data and often

exhibit poor performance, whereas graph convolution methods significantly enhance model accuracy. Methods based on CNN, such as STResNet and ACFM, utilize deep learning techniques. Compared to traditional statistical methods, these approaches capture nonlinear features and significantly improve the model’s prediction accuracy. However, the design framework of CNNs inherently limits their ability to consider adjacency relationships in a more detailed manner. Specifically, CNNs are restricted to using a regular receptive field, which makes it challenging to account for higher-order relationships between different nodes in a road network. As a result, there is still room for improvement in prediction accuracy for these methods. DCRNN and STGCN improve upon GCN by employing bidirectional propagation processes, achieving good results compared to traditional statistical methods. GWNet and AGCRN optimizes the generation of adaptive adjacency matrices, leading to better performance than other graph neural network methods. AGS integrates the advantages of multiple spatio-temporal graph neural network methods, resulting in improved inference accuracy compared to other approaches. This method achieves the best accuracy on some metrics in the England dataset and performs slightly less optimally on the PEMS07 and PEMS08 datasets. BigST is designed with scalability in mind for graph neural networks; its proposed linear graph convolution offers high scalability, allowing the model to perform well on larger datasets (e.g., CA), but its predictive accuracy on smaller datasets is not as competitive as other methods. Transformer-based methods, as seen in PDFormer and Bi-STAT, have shown promising performance in prediction tasks. Compared to other approaches, the attention mechanism in Transformers excels at capturing temporal information within time series data, leading to significant improvements in prediction accuracy. However, the complexity of the model design limits its scalability. As a result, these methods struggle to maintain high prediction accuracy when applied to large-scale datasets. GPT-ST improves the training effectiveness through a pretraining strategy. Compared to other GNN-based methods, it significantly enhances the model’s prediction accuracy. On the other hand, UniST is designed with a large model framework that supports training on multiple types

Table 4: Computation cost during training and inference on different models.

Model	Training Time(s/epoch)				Inference Time(s/epoch)			
	PEMS08	England	PEMS07	CA	PEMS08	England	PEMS07	CA
STResNet	32.35	48.32	44.17	153.35	5.93	8.21	7.02	23.26
ACFM	35.44	51.32	62.37	172.64	6.53	8.74	6.96	27.64
GWNet	114.23	130.75	235.32	1757.35	10.35	14.86	12.12	437.53
AGCRN	121.35	147.36	264.16	1921.66	11.72	16.74	14.35	463.23
PDFormer	270.25	325.37	446.36	2653.82	23.42	35.28	29.64	653.24
Bi-STAT	160.34	203.25	352.33	2325.24	15.37	20.03	19.86	502.51
GPT-ST	125.32	150.25	242.37	1869.32	12.01	16.07	15.33	467.53
UniST	703.25	859.36	1560.21	6358.25	36.21	40.35	38.39	1036.72
AGS	135.74	153.25	282.37	2135.12	12.13	16.93	14.76	470.03
BigST	<u>27.95</u>	<u>42.52</u>	<u>37.28</u>	<u>95.72</u>	<u>5.86</u>	<u>6.75</u>	<u>3.27</u>	<u>16.37</u>
GSNet(Ours)	10.37	17.62	11.80	27.26	0.99	2.12	1.21	5.50
Improvement	2.70×	2.41×	3.16×	3.51×	5.91×	3.18×	4.84×	2.98×

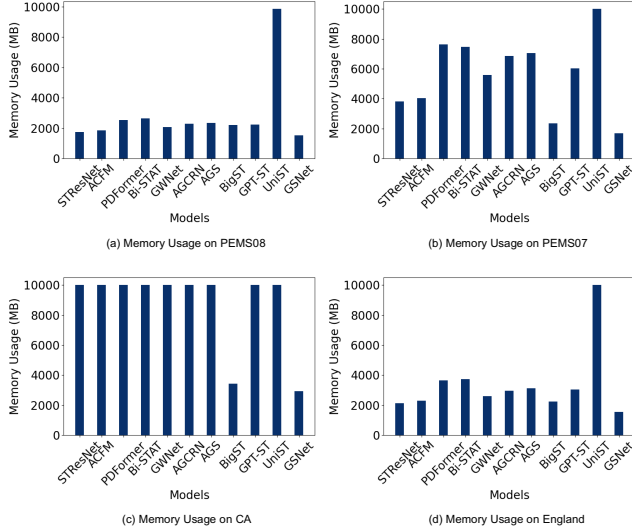


Figure 4: Memory usage on all the datasets.

of data. Its main advantage lies in its ability to perform zero-shot and few-shot learning. However, in the experiments presented in this paper, we maintained fairness in comparison by not utilizing multiple data types for training, which resulted in less favorable performance.

5.4.2 Efficiency Comparison. Figure 4 illustrates the memory usage of various methods. It is evident that GSNet and BigST are the best and second-best performers, respectively. Additionally, as the number of nodes in the dataset increases, the advantage of the linear space complexity of GSNet and BigST becomes increasingly apparent. Apart from GSNet and BigST, all other methods exhibit a cache occupancy of 10GB on the CA dataset. Furthermore, UniST consistently consumed 10GB of memory cache across all datasets except for PEMS08. This uniformity is attributed to the excessive number of nodes in the dataset, which leads to an 'out of memory' error when attempting to run the programs normally. Consequently,

we reduced the batch size of them incrementally until it could be executed under the condition of a 10GB GPU memory constraint.

Table 4 lists the computational efficiency of various graph neural network models, including training and inference times. Our method and BigST achieve the best and second-best runtime results across all datasets. Although AGS performs sparsification on the adaptive matrices in the model, the original paper and code do not provide methods for efficient computation using these sparse matrices, resulting in longer training and inference time compared to other methods. STResNet and ACFM demonstrate competitive efficiency, ranking just behind the BigST model. As CNN-based methods, they do not need to account for adjacency relationships between all nodes, which contributes to their relatively fast runtime. GWNet, AGCRN and GPT-ST do not show a significant difference in runtime compared to our method on smaller datasets (e.g., PEMS08), but their computational times increase dramatically as the number of nodes in the dataset grows. In contrast, PDFormer and Bi-STAT require significantly longer training and inference times. This is due to the substantial computational overhead inherent in transformer-based methods. Similarly, UniST exhibits even greater computational demands, with both its training and inference times far exceeding those of other models, primarily due to its complex framework design.

On the CA dataset, GSNet demonstrates exceptional training efficiency. Specifically, it achieves a training acceleration that is 64 to 70 \times faster compared to GWNet and AGCRN. BigST, despite designing a graph convolution operation with a complexity of $O(N)$, includes a kernel function based on random mapping, which makes the overall operational efficiency of the model slightly lower than ours. On the CA dataset, our method achieves 3.51 \times higher training acceleration. And On the PEMS07 dataset, our method achieves 4.84 \times higher inference acceleration.

In summary, among existing graph neural network methods, some methods, such as GWNet and AGCRN, sacrifice model scalability for higher predictive accuracy, leading to a decline in both accuracy and efficiency on large-scale datasets. Other methods, like BigST enhance model scalability at the expense of predictive accuracy. Our method balances the effectiveness and efficiency of

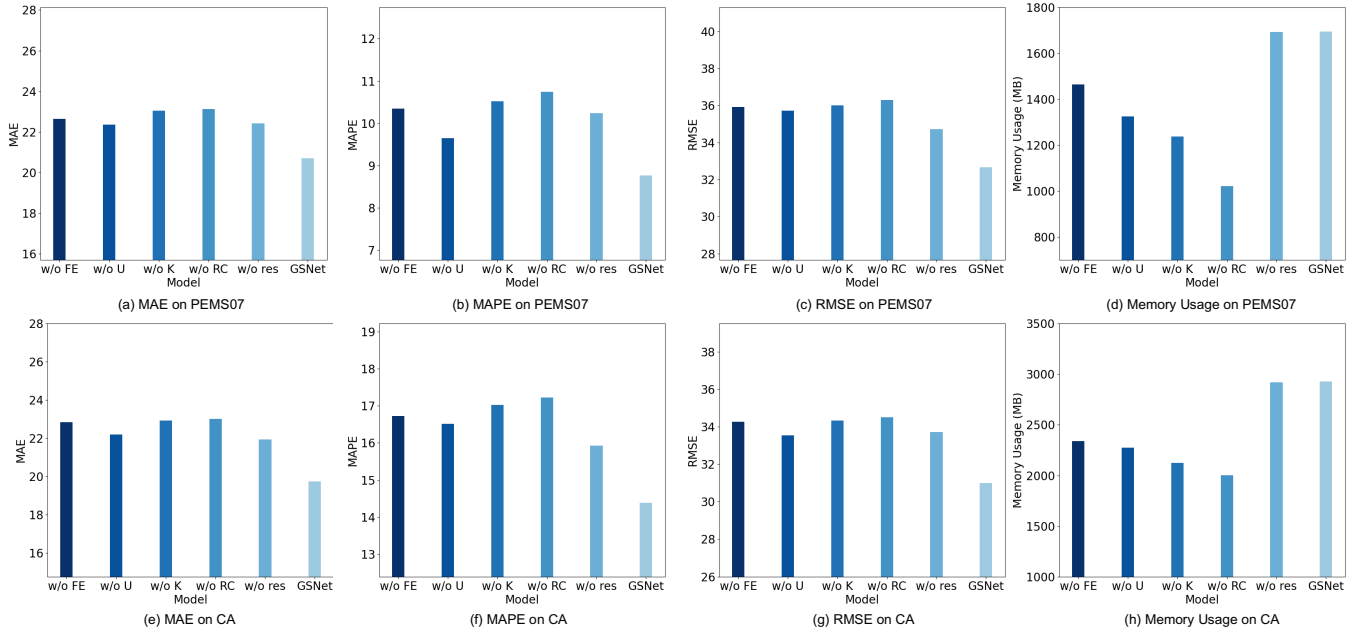


Figure 5: Prediction performance of different variants on PEMS07 and CA datasets.

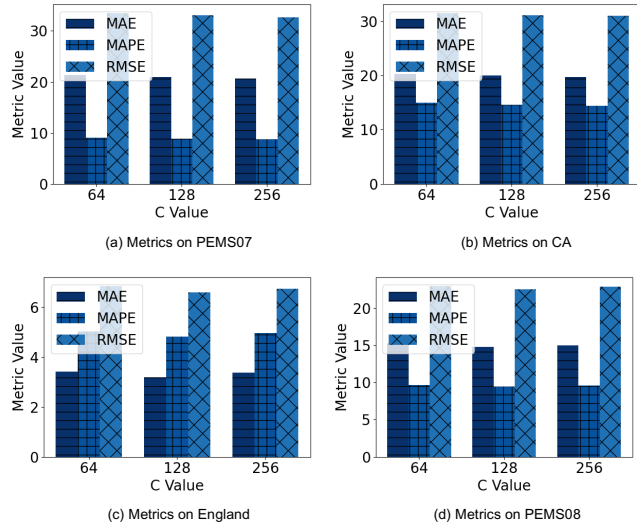


Figure 6: Prediction performance of different C value on all the datasets.

the model, achieving the scalability of high-accuracy graph neural networks on large-scale datasets.

5.5 Ablation Studies

To further verify effectiveness of the proposed method, we conduct ablation studies on several datasets. In particular, we design five variants of our model:

- (1) *w/o FE*: This variant eliminates the Feature Extractor from the model and constructs the model solely with the Relational Compressor.
- (2) *w/o U*: This variant removes the source concatenation from Relational Compressor but retains the adjacency matrix within Relational Compressor. The matrix K has also been appropriately adjusted in terms of its shape to ensure that the matrix multiplication is well-defined.
- (3) *w/o K*: This variant removes the adjacency matrix from Relational Compressor while retaining the source concatenation. The adjacency matrix K has been replaced with an identity matrix of the same shape.
- (4) *w/o RC*: This variant removes Relational Compressor, retaining only the Feature Extractor.
- (5) *w/o res*: This variant removes the residual connections from both Feature Extractor and Relational Compressor.

As illustrated in Figure 5, we first compared the predictive accuracy of different variants across multiple datasets. It can be observed that removing Relational Compressor or the adjacency matrix within Relational Compressor has the most significant impact on the model’s accuracy. This indicates that the fusion of graph node features is a crucial step in enhancing predictive accuracy. Removing Feature Extractor from the model also greatly affects accuracy, demonstrating that accurately learning node features is a prerequisite for accurate prediction. Removing the residual connections from the model significantly reduces its accuracy as well. This is because Feature Extractor and Relational Compressor are combined in a sequential manner, and residual connections effectively mitigate the vanishing or exploding gradient issues in the stacked structure. It is noteworthy that almost all variants result

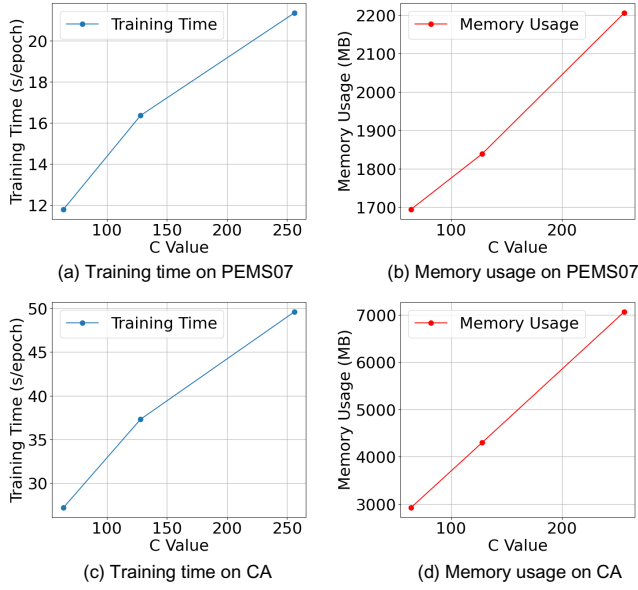


Figure 7: Memory usage and training time of different C value on PEMS07 and CA datasets.

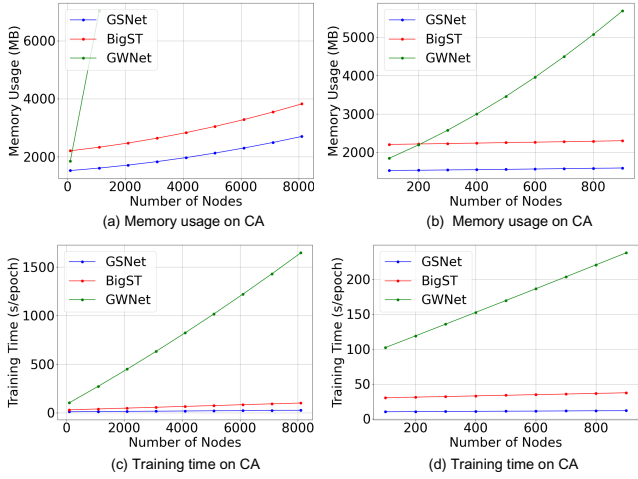


Figure 8: Memory usage and training time of different models with varying N .

in a substantial decrease in the model's predictive accuracy, highlighting the collaborative relationship between different modules in our approach and validating the effectiveness of our overall model framework design.

We also compared the memory usage of different variants. It can be observed that the variant *w/o* RC has the most significant reduction in memory usage. This indicates that the Relational Compressor in our model consumes a considerable amount of memory usage, whereas the Feature Extractor, due to its simpler structure, uses less memory.

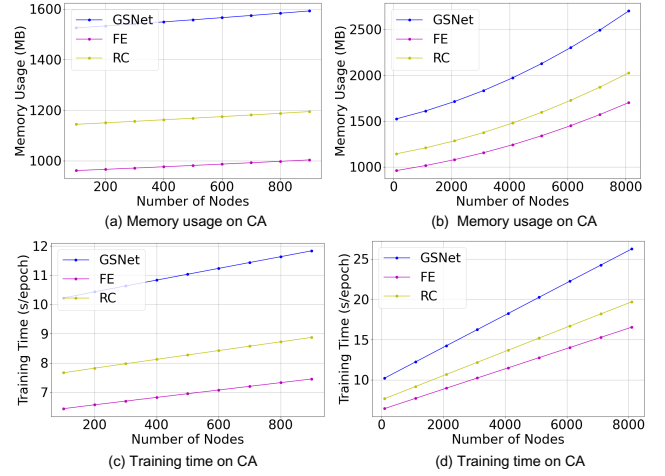


Figure 9: Memory usage and training time of modules in the GSNet with varying N .

5.6 Parameter Sensitivity

In the model, we compress the input data into a low-dimensional space to model the implicit associations between nodes, where the dimension of K is also a hyperparameter. As shown in Figure 6, we conducted a series of experiments to test the model's performance under different values of C . It can be seen that on the PEMS08 dataset, the model's accuracy first increases and then decreases with the increase of C . On the PEMS07 and CA datasets, the model's predictive accuracy increases with the increase of C . We speculate that on larger datasets, the model's performance might also decrease with the increase of C , but due to equipment limitations, we have not tested this hypothesis on larger values of C . The experimental results indicate that for large-scale datasets, appropriately increasing the value of C can enhance the model's predictive accuracy. As illustrated in Figure 7, we also tested the model's operational efficiency and memory usage under different values of C . It can be observed that as the value of C increases, the model's training time, inference time, and memory usage also increase. The increase in training and inference time is relatively slow, while memory usage almost linearly increases with the increase of C . In summary, when using GSNet on larger-scale datasets, the value of C can serve as a trade-off coefficient between effectiveness and efficiency, adjusting the balance between the model's predictive accuracy and operational efficiency.

5.7 Efficiency Analysis

To further assess the computational efficiency of our model across different data scales, we segmented the CA dataset based on sensor ID order and compared our method against several baselines on these varying scales. As shown in Figures 8, both the running time and memory usage of our method incrementally increase with the dataset size. Notably, the rate of this increase is more gradual compared to existing methods, demonstrating our model's superior scalability. With GWNet, we observed that as the number of nodes surpasses 1,000, it consumes the entire memory capacity

of our GPU. To mitigate this, we adapted the model to utilize a smaller batch size, thereby reducing the volume of data processed per batch. It is evident that, among methods with linear complexity, GraphSparseNet (GSNet) significantly outperforms BigST in terms of storage requirements and training efficiency. This comprehensive analysis underscores the robustness and adaptability of our model, offering valuable insights into its performance and efficiency across a spectrum of data scales and settings.

We further examine the efficiency of the two core modules, the Feature Extractor and Relational Compressor, as the data scale increases. To do so, we experiment with two model variants: FE, which retains only the Feature Extractor module, and RC, which retains only the Relational Compressor module. As illustrated in the figure 9, both modules exhibit a roughly linear increase in spatio-temporal complexity with respect to the number of nodes. It is worth noting that the Relational Compressor includes an additional feature fusion process, which leads to higher training time and increased memory usage.

6 RELATED WORK

Traffic flow forecasting is a key challenge in spatio-temporal prediction, similar to tasks like forecasting demand for shared bikes, buses, taxis, and crowd flows [2, 11, 20, 43]. Traditional statistical approaches such as ARIMA [33] and SVM [4], while prevalent in time series forecasting, often fall short due to their inability to account for spatial dimensions, making them less effective for complex spatio-temporal data sets. The advent of deep learning has introduced methods adept at handling the intricacies and non-linearities inherent in traffic data. Convolutional Neural Networks (CNNs), in particular, have become a staple in traffic flow forecasting [22, 24, 37, 41, 42]. By leveraging techniques originally developed for image recognition [32], CNNs can effectively model the spatial relationships between different grid regions. Recurrent Neural Networks (RNNs) have also been instrumental in the analysis of sequence data, bringing their sequence memorization capabilities to bear on traffic flow forecasting [29, 38, 44]. More recent advancements have seen Graph Neural Networks (GNNs) rise to prominence for their ability to manage the spatio-temporal correlations present in traffic flow data, achieving state-of-the-art results [8, 26, 28, 31]. GNNs, initially designed for graph structure analysis, have found widespread application in node embedding [25] and node classification [15]. In the realm of transportation systems, DCRNN [19] employs a bidirectional diffusion process to emulate real-world road conditions and utilizes gated recurrent units to capture temporal dynamics. ASTGCN [7], on the other hand, utilizes dual attention layers to discern the dynamics of spatial dependencies and temporal correlations. STGCN, Graph WaveNet, LSGCN, and AGCRN [1, 12, 36, 39] represent a lineage of methods that build upon Graph Convolutional Networks (GCNs) to extract spatio-temporal information. Notably, Graph WaveNet introduces a self-adaptive matrix to factor in the influence between nodes and their neighbors, while LSGCN employs an attention layer to achieve a similar end. STSGCN, STFGNN, and STGODE [6, 18, 30] propose GCN methodologies designed to capture spatio-temporal information in a synchronous manner. MTGNN [35] introduces a graph learning

module that constructs a dynamic graph by calculating the similarity between learnable node embeddings. DMSTGCN [10] captures spatio-temporal characteristics by forging dynamic associations between nodes. STPGNN [16] enhances predictive accuracy by taking into account special nodes within the road network. In addition to GNN-based methods, several approaches leveraging Transformers [3, 13] and pre-training [21, 40] have also been proposed. However, both Transformer-based and pre-training approaches, along with GNNs, still encounter significant computational challenges when identifying nodes relationships, which limits the scalability of these models on large-scale datasets. Existing adaptive graph neural network methods often rely on fully connected graphs to bolster the model's learning capabilities. Yet, the exponential growth in the number of edges with an increase in nodes poses a challenge for these methods when generalizing to larger-scale datasets. To counter this, AGS [5] has proposed a method for significantly simplifying the adaptive matrix, thereby reducing the model's computational load. However, this approach is limited to the inference stage. In practical applications, the computational cost during the training phase often dwarfs that of the inference phase. A method that maintains linear complexity during training can significantly enhance the model's operational efficiency. BigST [9] introduces a method that employs kernel functions to linearly approximate graph convolution operations, yielding a graph prediction model with linear complexity. However, kernel-based methods can sometimes result in anomalous gradient values during training, impacting model convergence and, by extension, the model's performance. Amidst the ever-expanding scale of traffic data, there is an urgent need for graph neural network methods capable of delivering high-precision predictions at scale.

7 CONCLUSION

In this paper, we addressed the challenge of enhancing the scalability of GNN-based methods for large-scale traffic spatio-temporal data prediction. We identified that existing methods either do not fully address the computational complexity of graph operations or compromise model accuracy when simplifying GNN structures. By introducing GraphSparseNet (GSNet), we have provided a novel approach that leverages two module—the Feature Extractor and the Relational Compressor—to effectively manage and analyze large-scale traffic data. Our theoretical analysis and empirical evaluations highlight the significant advantages of GSNet over existing methods. By reducing model complexity to $O(N)$, GSNet addresses the scalability issue while maintaining high predictive accuracy. Our extensive experiments on real-world traffic datasets demonstrate that GSNet not only achieves high predictive accuracy but also significantly improves training efficiency, outperforming state-of-the-art methods in both aspects. The model's efficiency is further validated by a substantial improvement in training time, outperforming current linear models by 3.51x.

ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers for their helpful comments. This work was supported by the NSFC 61572537, and the CCF-Huawei Populus Grove Challenge Fund 202305. Yubao Liu is the corresponding author.

REFERENCES

- [1] Lei Bai, Lina Yao, Can Li, Xianzhi Wang, and Can Wang. 2017. Adaptive graph convolutional recurrent network for traffic forecasting. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, Vol. 33. 17804–17815.
- [2] Di Chai, Leye Wang, and Qiang Yang. 2018. Bike Flow Prediction with Multi-Graph Convolutional Networks. In *Proceedings of the 26th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. 379–400.
- [3] Changlu Chen, Yanbin Liu, Ling Chen, and Chengqi Zhang. 2023. Bidirectional Spatial-Temporal Adaptive Transformer for Urban Traffic Flow Forecasting. *IEEE Transactions on Neural Networks and Learning Systems* 34, 10 (2023), 6913–6925.
- [4] Harris Drucker, Christopher J Burges, Linda Kaufman, Alex Smola, and Vladimir Vapnik. 1996. Support vector regression machines. In *Advances in neural information processing systems*, Vol. 9.
- [5] Wenying Duan, Xiaoxi He, Zimu Zhou, Lothar Thiele, and Hong Rao. 2023. Localised Adaptive Spatial-Temporal Graph Neural Network. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 448–458.
- [6] Zheng Fang, Qingqing Long, Guojie Song, and Kunqing Xie. 2021. Spatial-Temporal Graph ODE Networks for Traffic Flow Forecasting. (2021), 364–373.
- [7] Shengnan Guo, Youfang Lin, Ning Feng, Chao Song, and Huaiyu Wan. 2019. Attention based spatial-temporal graph convolutional networks for traffic flow forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 33. 922–929.
- [8] Shengnan Guo, Youfang Lin, Huaiyu Wan, Xiucheng Li, and Gao Cong. 2022. Learning Dynamics and Heterogeneity of Spatial-Temporal Graph Data for Traffic Forecasting. *IEEE Transactions on Knowledge and Data Engineering* 34, 11 (2022), 5415–5428.
- [9] Jindong Han, Weijia Zhang, Hao Liu, Tao Tao, Naiqiang Tan, and Hui Xiong. 2024. BigST: Linear Complexity Spatio-Temporal Graph Neural Network for Traffic Forecasting on Large-Scale Road Networks. In *Proceedings of the VLDB Endowment*, Vol. 17. 1081–1090.
- [10] Liangzhe Han, Bowen Du, Leilei Sun, Yanjie Fu, Yisheng Lv, and Hui Xiong. 2021. Dynamic and Multi-Faceted Spatio-Temporal Deep Learning for Traffic Speed Forecasting. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 547–555.
- [11] Qi Hu, Lingfeng Ming, Ruijie Xi, Lu Chen, Christian S. Jensen, and Bolong Zheng. 2021. SOUP: A Fleet Management System for Passenger Demand Prediction and Competitive Taxi Supply. In *2021 IEEE 37th International Conference on Data Engineering*. 2657–2660.
- [12] Rongzhou Huang, Chuyin Huang, Yubao Liu, Genan Dai, and Weiyang Kong. 2020. LSGCN: long short-term traffic prediction with graph convolutional networks. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, Vol. 7. 2355–2361.
- [13] Jiawei Jiang, Chengkai Han, Wayne Xin Zhao, and Jingyuan Wang. 2023. PDFormer: propagation delay-aware dynamic long-range transformer for traffic flow prediction. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 37. 4365–4373.
- [14] Guangyin Jin, Yuxuan Liang, Yuchen Fang, Jincai Huang, Junbo Zhang, and Yu Zheng. 2023. Spatio-Temporal Graph Neural Networks for Predictive Learning in Urban Computing: A Survey. *arXiv preprint arXiv:2303.14483* (2023).
- [15] Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016).
- [16] Weiyang Kong, Ziyu Guo, and Yubao Liu. 2024. Spatio-Temporal Pivotal Graph Neural Networks for Traffic Flow Forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 8627–8635.
- [17] Weiyang Kong, Kaiqi Wu, Sen Zhang, and Yubao Liu. 2025. GraphSparseNet: a Novel Method for Large Scale Traffic Flow Prediction. *arXiv preprint arXiv:2502.19823* (2025).
- [18] Mengzhang Li and Zhanxing Zhu. 2021. Spatial-temporal fusion graph neural networks for traffic flow forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 35. 4189–4196.
- [19] Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. 2018. Diffusion Convolutional Recurrent Neural Network: Data-Driven Traffic Forecasting. In *International Conference on Learning Representations*.
- [20] Yexin Li, Yu Zheng, Huichu Zhang, and Lei Chen. 2015. Traffic Prediction in a Bike-Sharing System. In *Proceedings of the 23rd SIGSPATIAL international conference on advances in geographic information systems*. 1–10.
- [21] Zhonghang Li, Lianghao Xia, Yong Xu, and Chao Huang. 2023. GPT-ST: Generative Pre-Training of Spatio-Temporal Graph Neural Networks. In *Advances in neural information processing systems*, Vol. 36. 70229–70246.
- [22] Lingbo Liu, Jiajie Zhen, Guanbin Li, Geng Zhan, and Liang Lin. 2019. ACFM: A Dynamic Spatial-Temporal Network for Traffic Prediction. *arXiv preprint arXiv:1909.02902* (2019).
- [23] Xu Liu, Yutong Xia, Yuxuan Liang, Junfeng Hu, Yiwei Wang, Lei Bai, Chao Huang, Zhengguang Liu, Bryan Hooi, and Roger Zimmermann. 2023. LargeST: A Benchmark Dataset for Large-Scale Traffic Forecasting. In *Advances in Neural Information Processing Systems*, Vol. 36. 75354–75371.
- [24] Kun Ouyang, Yuxuan Liang, Ye Liu, Zekun Tong, Sijie Ruan, Yu Zheng, and David S. Rosenblum. 2022. Fine-Grained Urban Flow Inference. *IEEE Transactions on Knowledge and Data Engineering* 34, 6 (2022), 2755–2770.
- [25] Shirui Pan, Ruiqi Hu, Sai-fu Fung, Guodong Long, Jing Jiang, and Chengqi Zhang. 2019. Learning graph embedding with adversarial training methods. *IEEE transactions on cybernetics* 50, 6 (2019), 2475–2487.
- [26] Zheyi Pan, Wentao Zhang, Yuxuan Liang, Weinan Zhang, Yong Yu, Junbo Zhang, and Yu Zheng. 2022. Spatio-Temporal Meta Learning for Urban Traffic Prediction. *IEEE Transactions on Knowledge and Data Engineering* 34, 3 (2022), 1462–1476.
- [27] Zhen Qin, Xiaodong Han, Weixuan Sun, Dongxu Li, Lingpeng Kong, Nick Barnes, and Yiran Zhong. 2022. The Devil in Linear Transformer. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*. 7025–7041.
- [28] Yibin Shen, Cheqing Jin, Jiaxun Hua, and Dingjiang Huang. 2022. TTPNet: A Neural Network for Travel Time Prediction Based on Tensor Decomposition and Graph Embedding. *IEEE Transactions on Knowledge and Data Engineering* 34, 9 (2022), 4514–4526.
- [29] Xingjian Shi, Zhourong Chen, Hao Wang, DiYan Yeung, Waikin Wong, and Wangchun Woo. 2015. Convolutional LSTM Network: A Machine Learning Approach for Precipitation Nowcasting. In *Advances in neural information processing systems*, Vol. 28. 802–810.
- [30] Chao Song, Youfang Lin, Shengnan Guo, and Huaiyu Wan. 2020. Spatial-temporal synchronous graph convolutional networks: A new framework for spatial-temporal network data forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 34. 914–921.
- [31] Junkai Sun, Junbo Zhang, Qiaofei Li, Xiuwen Yi, Yuxuan Liang, and Yu Zheng. 2022. Predicting Citywide Crowd Flows in Irregular Regions Using Multi-View Graph Convolutional Networks. *IEEE Transactions on Knowledge and Data Engineering* 34, 5 (2022), 2348–2359.
- [32] David Alexander Tedjopurnomo, Zhifeng Bao, Baihua Zheng, Farhana Murtaza Choudhury, and A. K. Qin. 2022. A Survey on Modern Deep Neural Network for Traffic Prediction: Trends, Methods and Challenges. *IEEE Transactions on Knowledge and Data Engineering* 34, 4 (2022), 1544–1561.
- [33] Billy M Williams and Lester A Hoel. 2003. Modeling and forecasting vehicular traffic flow as a seasonal ARIMA process: Theoretical basis and empirical results. *Journal of transportation engineering* 129, 6 (2003), 664–672.
- [34] Yuankai Wu and Huachun Tan. 2016. Short-term traffic flow forecasting with spatial-temporal correlation in a hybrid deep learning framework. *arXiv preprint arXiv:1612.01022* (2016).
- [35] Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, Xiaojuan Chang, and Chengqi Zhang. 2020. Connecting the Dots: Multivariate Time Series Forecasting with Graph Neural Networks. In *Proceedings of the 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 753–763.
- [36] Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, and Chengqi Zhang. 2019. Graph wavenet for deep spatial-temporal graph modeling. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*. 1907–1913.
- [37] Huaxiu Yao, Xianfeng Tang, Hua Wei, Guanjie Zheng, Yanwei Yu, and Zhenhui Li. 2018. Modeling spatial-temporal dynamics for traffic prediction. *arXiv preprint arXiv:1803.01254* (2018).
- [38] Junchen Ye, Leilei Sun, Bowen Du, Yanjie Fu, Xinran Tong, and Hui Xiong. 2019. Co-Prediction of Multiple Transportation Demands Based on Deep Spatio-Temporal Neural Network. In *Proceedings of the 25th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 305–313.
- [39] Bing Yu, Haoteng Yin, and Zhanxing Zhu. 2018. Spatio-temporal graph convolutional networks: a deep learning framework for traffic forecasting. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*. 3634–3640.
- [40] Yuan Yuan, Jingtao Ding, Jie Feng, Depeng Jin, and Yong Li. 2024. Unist: A prompt-empowered universal model for urban spatio-temporal prediction. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 4095–4106.
- [41] Junbo Zhang, Yu Zheng, and Dekang Qi. 2017. Deep Spatio-Temporal Residual Networks for Citywide Crowd Flows Prediction. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 31. 1655–1661.
- [42] Junbo Zhang, Yu Zheng, Dekang Qi, Ruiyuan Li, and Xiuwen Yi. 2016. DNN-Based Prediction Model for Spatio-Temporal Data. In *Proceedings of the 24th ACM SIGSPATIAL international conference on advances in geographic information systems*. 1–4.
- [43] Boming Zhao, Pan Xu, Yexuan Shi, Yongxin Tong, Zimu Zhou, and Yuxiang Zeng. 2019. Preference-aware task assignment in on-demand taxi dispatching: An online stable matching approach. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 33. 2245–2252.
- [44] Ali Zonozi, Jung-Jae Kim, Xiaoli Li, and Gao Cong. 2018. Periodic-CRN: A Convolutional Recurrent Model for Crowd Density Prediction with Recurring Periodic Patterns. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*, Vol. 18. 3732–3738.