

Spatial-temporal Forecasting for Regions without Observations

Xinyu Su
The University of Melbourne
Australia
suxs3@student.unimelb.edu.au

Jianzhong Qi
The University of Melbourne
Australia
jianzhong.qi@unimelb.edu.au

Egemen Tanin
The University of Melbourne
Australia
etanin@unimelb.edu.au

Yanchuan Chang
The University of Melbourne
Australia
yanchuan.chang@unimelb.edu.au

Majid Sarvi
The University of Melbourne
Australia
majid.sarvi@unimelb.edu.au

ABSTRACT

Spatial-temporal forecasting plays an important role in many real-world applications, such as traffic forecasting, air pollutant forecasting, crowd-flow forecasting, and so on. State-of-the-art spatial-temporal forecasting models take data-driven approaches and rely heavily on data availability. Such models suffer from accuracy issues when data is incomplete, which is common in reality due to the heavy costs of deploying and maintaining sensors for data collection. A few recent studies attempted to address the issue of incomplete data. They typically assume some data availability in a region of interest either for a short period or at a few locations. In this paper, we further study spatial-temporal forecasting for a region of interest without any historical observations, to address scenarios such as unbalanced region development, progressive deployment of sensors or lack of open data. We propose a model named STSM for the task. The model takes a contrastive learning-based approach to learn spatial-temporal patterns from adjacent regions that have recorded data. Our key insight is to learn from the locations that resemble those in the region of interest, and we propose a selective masking strategy to enable the learning. As a result, our model outperforms adapted state-of-the-art models, reducing errors consistently over both traffic and air pollutant forecasting tasks. The source code is available at <https://github.com/suzy0223/STSM>.

1 INTRODUCTION

Spatial-temporal forecasting is an important component of many real-world applications, e.g., Intelligent Transportation Systems (ITS) and air quality monitoring systems. Current state-of-the-art methods for spatial-temporal forecasting are data-driven. They leverage sequential models, e.g., 1-D convolutional neural networks (CNN) [20, 22, 31] or recurrent neural networks (RNN) [16], to capture temporal features, along with spatial models, e.g., graph neural networks (GNN), which model spatial relations [20, 22, 31]. However, *data scarcity* is ubiquitous due to high deployment and maintenance costs of sensors and unstable transmission mediums. Hence, developing a model that enables accurate forecasting without complete historical data is critical.

Prior attempts to address the data scarcity issue for spatial-temporal tasks fall into two categories: (1) *data missing at times* [14, 17, 25, 32, 36]: observations at locations of interest are incomplete for all time due to complex environments and/or faulty sensors, or a short sensor deployment time; (2) *data missing at*

(*scattering*) *locations* [1, 30, 39]: observations are missing because the locations of interest do not have historical data recorded at all (e.g., no sensors have been deployed at those locations). For the latter category, recent studies revisit Kriging [7]. The aim is to generate fine-grained records through coarse-grained observations by inserting derived data for the unobserved locations. The state-of-the-art models [30, 39] adopt neural networks as a solution. *Both categories assume some data available for the region of interest, shown in Fig. 1 (a) and 1 (b).*




Figure 1: Problem setting comparison. Coloured maps and grey maps indicate data observed and unobserved, respectively. Our focus is Case (c).

Existing studies fail to consider a "continuous" data scarcity scenario, where all locations without any data observations are skewed towards a continuous sub-region (*i.e., the region of interest has no available data*) while the region boundary is adjacent to the locations with data observations. Fig. 1 (c) illustrates this scenario, which may occur when (1) sensors are deployed progressively from one region to another (one such scenario has been observed in Hong Kong [25]), (2) some regions do not have resources to deploy the sensors (e.g., Shanghai traffic congestion data only covers the core city region [37]), or (3) regions are not willing to open their data (e.g., the Tom-Tom traffic index [27] does not have data from mainland China).

To fill this gap, we propose a new problem – *spatial-temporal forecasting for a region of interest without historical data*.

This new problem is challenging, and existing Kriging-based approaches [30, 39] do not address the problem directly. This is because the core idea of the Kriging-based approach is to perform data interpolation for the locations without observations based on nearby locations (or time) with observations, which is unavailable under our setting.

For example, an earlier model, IGNNK [30], represents the region of interest as a graph and exploits GNNs for Kriging. It reports substantial performance drops in our setting (detailed in Section 5.2) because GNNs become ineffective when the local neighbourhood of a location does not have historical observations to help infer the data. The state-of-the-art Kriging model, INCREASE [39], also uses a graph-based region representation. It aggregates information from nearest neighbours in advance

© 2024 Copyright held by the owner/author(s). Published in Proceedings of the 27th International Conference on Extending Database Technology (EDBT), 25th March-28th March, 2024, ISBN 978-3-89318-095-0 on OpenProceedings.org. Distribution of this paper is permitted under the terms of the Creative Commons license CC-by-nc-nd 4.0.

and then uses GRUs (which is a type of RNN) to capture the temporal correlation of the data. This model fails to utilise the global features of the graph as it only considers the nearest neighbours.

We propose a spatial-temporal forecasting model with a selective masking strategy, named STSM, to enable learning from the locations that resemble those in the region of interest and forecasting for the region of interest without observations. We mask observed data from sub-regions (which are near the region of interest and have historical data) and train STSM to make predictions for the sub-regions. Then, at testing time, we exploit the similarity between the region of interest and the sub-regions used at training to make predictions for the region of interest.

Unlike existing studies that use random masking [30], STSM incorporates a *selective masking module* to mask sub-regions that are more similar to the region of interest. This strategy makes it easier for STSM to extend its forecasting capability from the masked sub-regions to the region of interest. The selective masking module fuses regional features, road network features and spatial distance to compute similarity scores between the masked sub-region (i.e., masked locations) and the region of interest (i.e., unobserved locations). The similarity scores are normalised into the range of [0,1] and serve as the probabilities for drawing the sub-regions (or locations) to be masked.

Further, we use the observed locations' historical data to generate pseudo-observations for the unobserved locations and the masked locations, which enables computing a temporal similarity based adjacency matrix, i.e., we build links between observed locations and unobserved locations that have high temporal similarities. This helps identify more similar neighbours.

Overall, this paper makes the following contributions:

- We propose a new spatial-temporal forecasting task - forecasting for regions without historical observations. This task can be applied to addressing issues of unbalanced region development and lack of open data.
- We design a selective masking module to guide our model STSM to mask observed locations that have high similarity with the unobserved region to enable STSM to generalise predictions to the unobserved region.
- We design an efficient pseudo-observation generating strategy and compute a temporal adjacency matrix based on it to help identify the more informative neighbours and enhance model learning efficacy.
- Extensive experiments show that our model outperforms the state-of-the-art model that we adapt to this new problem, in terms of forecasting accuracy.

2 RELATED WORK

2.1 Spatial-temporal Forecasting

Current state-of-the-art spatial-temporal forecasting models are mainly based on deep neural networks. DCRNN [16] introduces a diffusion convolutional recurrent neural network to model the spatial correlation between locations and adapts gated recurrent units (GRU) to model the temporal dependency. GRUs and other RNN models have a recurrent structure that suffers in model running time and in the effectiveness of modelling longer sequences. To overcome this limitation, GraphWaveNet [31] utilises 1-D temporal convolutional modules to capture the temporal dependency. In addition, the attention mechanism [33] is widely used in spatial-temporal forecasting [6, 8, 10, 38]. A series of recent studies further embed heterogeneous relations into adjacency matrices, including temporal similarity [15, 22] and embedding

similarity [20]. Some studies [11–13, 18] adopt self-supervised learning to enhance spatial-temporal pattern representations. Meanwhile, DeepSTUQ [21] considers prediction uncertainty when forecasting traffic. These models assume fully available historical data and suffer in learning capability when data is incomplete.

2.2 Spatial-temporal Forecasting with Incomplete Data

Existing spatial-temporal forecasting methods for incomplete data can be divided into two categories from the data perspective: *data missing at times* and *data missing at (scattering) locations*.

Data missing at times: Observations at locations of interest are incomplete for all time. For this category, one dominant class of studies focuses on random or continuous data missing at times caused by harsh environments [14, 17], e.g., extreme weather or transmission device issues. Generative adversarial networks (GAN) are applied to address this issue [32, 36]. Another study [25] uses transfer learning for settings where sensors have only been deployed for a short period of time (e.g., 10 days).

Data missing at (scattering) locations: Some locations of interest do not have observations at all. Problems considering this setting, i.e., Kriging [7], aim to impute fine-grained records via coarse-grained records, which is to recover signals for unobserved locations. Gaussian process regression [29] is a classic solution, while it suffers from low efficiency and poor scalability. Tensor/matrix completion algorithms [2, 24, 41] show better efficiency on large datasets. They combine the low-rank structure and regularisers to maintain local and global consistency. Most tensor/matrix completion algorithms are transductive. They cannot process a new location of interest without re-training. Recent studies [30, 39] propose inductive structures. For example, IGNNK [30] utilises the inductive nature of GNNs together with 1-D CNN to recover records for unobserved locations. This model struggles to handle high data missing ratios where there is little information to learn from the neighbourhood. INCREASE [39] adopts an RNN for inductive imputation. This method uses heterogeneous relations for more accurate estimation, while it struggles in capturing global spatial-temporal patterns.

2.3 Graph Contrastive Learning

Our proposed STSM is based on contrastive learning, in particular, graph contrastive learning (GCL), which applies contrastive learning to graph data. The basic idea of contrastive learning is to maximize the similarity between positive samples while minimizing that between negative samples.

A series of studies [34, 35, 42] focuses on graph augmentation modules that generate positive and negative samples. For example, GraphCL [35] introduces four augmentation methods, such as node dropping and edge perturbation, to create positive graph pairs. Later, GCA [42] and JOAO [34] improve the augmentation strategies by taking the node weights and the edge weights into consideration. Besides, some studies [23, 28] aim to maximize the mutual information between graph inputs at different scales, e.g., nodes vs. graphs.

A few studies [5, 18] introduce GCL into spatial learning tasks. For example, SARN [5] learns road network embeddings via contrastive learning. STGCL [18] applies GCL to predict traffic flow assuming complete data. Unlike STGCL, our model can process incomplete data because of the proposed selective masking module that can adaptively augment spatial-temporal graphs based on

the heterogeneous similarities between observed and unobserved locations.

3 PROPOSED BASE MODEL

We start with a basic version of our proposed model, named STSM-RNC (Fig. 2). We first define basic concepts and our studied problem. Then, we present our base model and its training and testing procedures.

3.1 Concepts and Problem Statement

Region and Region Graph. We represent a region as a graph $G = (V, E)$. The set of graph vertices V represents N locations of interest in the region, and the set of graph edges E represents the connection between the locations. Graph G has a feature matrix $\mathbf{L} \in \mathbb{R}^{N \times F}$ for the locations, where F is the dimensionality of the location features. The features of a location consist of two parts, i.e., region information and road network information, which will be detailed in Section 4.1.

For each location $v_i \in V$, $x_i^t \in \mathbb{R}^C$ represents the *observations* at v_i at time step t , where C is the number of different types of observations, e.g., traffic speed, PM2.5, etc.

Observed and unobserved regions. Region graph G can be further divided into two adjacent sub-regions based on the availability of observations for the sub-regions, i.e., the *observed region* and the *unobserved region*. The locations in these two regions are referred to as *observed locations* (i.e., with observations) and *unobserved locations* (i.e., without observations), respectively. We use R_o to denote the region that contains *all and only the* observed locations. Similarly, we use R_u to denote the region that contains *all and only the* unobserved locations. These two regions have no overlap with each other, i.e., $R_o \cap R_u = \phi$. We use $G_o = (V_o, E_o)$ to represent the graph on the observed locations, where $V_o \subset V$ denotes the set of the observed locations, and $E_o \subset E \cap (V_o \times V_o)$ denotes the set of the edges over V_o . We use N_o to denote the size of V_o . Similarly, we define the graph on the unobserved locations, which is denoted as G_u , and the nodes V_u (and its size N_u) and edges E_u on this graph. Note that, $N = N_o + N_u$; $\mathbf{X}_{G_o}^t = (x_1^t, \dots, x_{N_o}^t) \in \mathbb{R}^{N_o \times C}$ denotes the observations of the observed locations in G_o at time steps t ; and $\hat{\mathbf{X}}_{G_u}^t = (\hat{x}_1^t, \dots, \hat{x}_{N_u}^t) \in \mathbb{R}^{N_u \times C}$ denotes estimated values for the unobserved locations in G_u at time step t .

Problem definition. Given a region graph G as defined above, with location features \mathbf{L} and past observations on the observed locations for a time window of length T , we aim to learn a function f to predict the values for the unobserved locations over the next T' time steps:

$$\hat{\mathbf{X}}_{G_u}^{t+1}, \dots, \hat{\mathbf{X}}_{G_u}^{t+T'} = f(\mathbf{X}_{G_o}^{t-T+1}, \dots, \mathbf{X}_{G_o}^t; \mathbf{L}) \quad (1)$$

3.2 Overview of Our Base Model

We next present a basic model (i.e., STSM-RNC in Section 5.2.2) that directly combines *spatial-temporal modelling* with *random sub-graph masking* to forecast for regions without observations as Fig. 2 shows. The main idea of STSM-RNC is to learn a model that can forecast observations (e.g., traffic speed or PM2.5) for masking sub-graphs and to extend this capacity to forecast for unobserved locations.

We denote the full graph with those unobserved locations as G . We mask a subset of the locations of G_o to generate a masked view G_o^m (Section 3.3). Following previous studies [9, 16], we use temporal similarity-based adjacency matrix and spatial-based

adjacency matrix for spatial correlation modelling. To compute the temporal similarity for the unobserved locations, we first compute pseudo-observations for all unobserved locations. Then, we use dynamic time warping (DTW) [3, 15] to compute temporal similarities among all observed locations, and the temporal similarities between the observed and the unobserved locations. Besides, for the masked locations in each model training epoch, we compute pseudo-observations for them and compute the temporal similarities between masked locations and observed locations. After these steps, we obtain $\mathbf{X}_{G_o^m}^{t-T+1:t}$ and $\mathbf{X}_{G^m}^{t-T+1:t}$ for training and testing, respectively, where both $\mathbf{X}_{G_o^m}^{t-T+1:t}$ and $\mathbf{X}_{G^m}^{t-T+1:t}$ contain the pseudo-observations for the masked or unobserved locations in G_o or G .

We feed $\mathbf{X}_{G_o^m}^{t-T+1:t}$ into a *spatial-temporal modelling module* to generate the prediction result $\mathbf{X}_{G_o^m}^{t+1:t+T'}$ (Section 3.4) and compute the mean squared error between the prediction and the ground truth as prediction loss to optimise the spatial-temporal model. After the model is trained, we feed $\mathbf{X}_{G^m}^{t-T+1:t}$ into the model to obtain predictions for the unobserved locations (Section 3.5).

3.3 Sub-graph Masking

STSM-RNC learns to predict values for masked locations during training and then extends this capability to predict values for unobserved locations at testing. To simulate the setting of a continuous region without data observations, which we focus on, we mask the sub-graph formed by each selected location and its 1-hop neighbours instead of a set of scattering locations.

Defining a sub-graph. The sub-graph of an observed location is formed by its 1-hop neighbours. We compute a location's 1-hop neighbours based on a spatial adjacency matrix \mathbf{A}_{sg} , which is defined by Eq. 2, where ϵ_{sg} is a hyper-parameter, and the $dist(c_i, c_j)$ denotes the distance between locations i and j (c_i and c_j are their geo-coordinates.) We use Euclidean distance in this paper for efficiency considerations, though road network distance can also be used.

$$A_{sg,ij} = \begin{cases} 1 & \exp(-\frac{dist(c_i, c_j)^2}{\sigma^2}) \geq \epsilon_{sg}, \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

Sub-graph masking. We use a masking ratio δ_m to define the percentage of observed locations to be masked. The number of locations masked is expected to be $N_o \cdot \delta_m$. Since the sub-graph of each location may have a different size, STSM-RNC iteratively and randomly selects a location and masks the location and its 1-hop neighbours until the number of masked locations reaches $N_o \cdot \delta_m$.

3.4 Spatial-temporal Modelling

The spatial-temporal modelling module of our base model STSM-RNC contains the 1-D convolution networks for temporal correlation modelling and the graph convolutional networks (GCNs) for spatial correlation modelling. The spatial-temporal modelling module stacks multiple blocks to compute the final output. Fig. 3 shows the structure of the spatial-temporal modelling module and details the l th block of the module. Each block contains a temporal correlation modelling module and a spatial correlation modelling module. These two modules are parallel in each block. We first describe the input features and the adjacency matrices used in the spatial-temporal modelling module. Then, we detail the temporal and spatial correlation modelling modules.




Figure 2: Model architecture of STSM-RNC. The model contains a sub-graph masking module and a spatial-temporal modelling module.

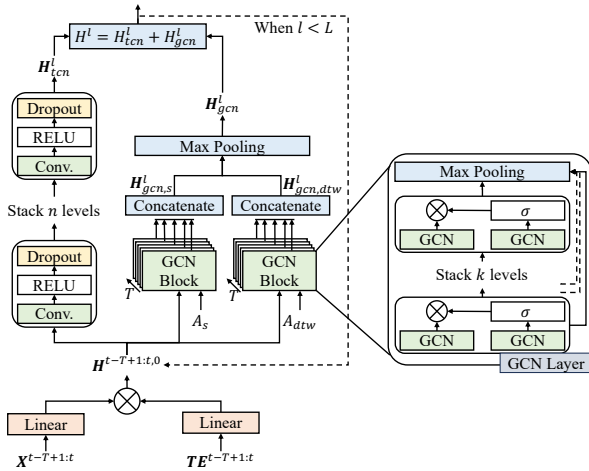


Figure 3: The structure of the spatial-temporal model

3.4.1 Input Features and Adjacency Matrices. The input features of STSM-RNC consist of historical observations for the observed locations, pseudo-observations for the unobserved and the masked locations, and temporal attention which is used to indicate the time of a day. Besides, we compute two types of adjacency matrices for the GCNs in our model.

We compute pseudo-observations for the i -th unobserved or masked location from real observations $x_i^t = \sum_{j \in N_o} \alpha_{i,j} x_j^t$. The weight for each observed location is determined by its spatial distance to location i , as defined by Eq. 3:

$$\alpha_{i,j} = \frac{\text{dist}(c_i, c_j)^{-1}}{\sum_{l \in N_o} \text{dist}(i, l)^{-1}} \quad (3)$$

This step can introduce more information into an unobserved or masked location based on that of its neighbours.

Then, we build a temporal similarity-based adjacency matrix. We follow a prior work [15] and adopt DTW to compute the temporal similarity. Since pseudo-observations can be regarded as real observations with noises, we only build links among the observed locations and from observed locations to unobserved (masked) locations (i.e., unobserved locations cannot send messages to observed locations directly during GCN training). This way, we avoid polluting the embeddings of the observed locations by those of the unobserved (masked) locations. We compute q_{kk} and q_{ku} most similar pairs of observed locations and pairs of observed and unobserved (or masked) locations, respectively. We establish a temporal similarity-based adjacency matrix

($\mathbf{A}_{dtw}^{train} \in \mathbb{R}^{N_o \times N_o}$ in the training process and $\mathbf{A}_{dtw} \in \mathbb{R}^{N \times N}$ in the testing process) by assigning an edge weight of 1 for these pairs, and 0 for the rest of the location pairs. Since the locations are masked dynamically in each training epoch, \mathbf{A}_{dtw}^{train} is updated in each model training epoch.

Further, we use temporal attention to capture periodic patterns, which can significantly impact the observed values of spatial-temporal data, e.g., rush hours. Given the length of a time interval that the spatial-temporal observations are recorded (e.g., 5 minutes), we can compute the number of intervals in a day, denoted as Td . Now each observation interval in a day gets an interval ID in $[0, Td - 1]$. Given an input of length T , we compute a time-of-day embedding $\mathbf{TE} \in \mathbb{R}^T$, which stores the interval IDs in the input time window. For example, $\mathbf{TE} = [0, 1, 2, 3]$ represents an input observation sequence that starts at the first interval of a day and ends at the fourth interval of a day.

To attach the time embedding \mathbf{TE} to the model input, we first project $\mathbf{TE}^{t-T+1:t} \in \mathbb{R}^{N_o \times T \times 1}$ and the input features $\mathbf{X}^{t-T+1:t} \in \mathbb{R}^{N_o \times T \times C}$ into the same latent space and then multiply them as formulated by Eq. 4:

$$\mathbf{H}^{t-T+1:t,0} = \phi_1(\mathbf{X}^{t-T+1:t}) \otimes \phi_2(\mathbf{TE}^{t-T+1:t}) \quad (4)$$

Here, $\mathbf{X}^{t-T+1:t} \in \mathbb{R}^{N_o \times T \times C}$ is the input observation sequence of the observed graph (i.e., G_o or G_o^m), while $\mathbf{TE}^{t-T+1:t}$ is the corresponding time embedding; $\phi_1(\cdot)$ and $\phi_2(\cdot)$ are linear functions that project the input observation sequence and time embedding into the same latent space for element-wise multiplication. We now obtain the features $\mathbf{H}^{t-T+1:t,0} \in \mathbb{R}^{N_o \times T \times C'}$, which serve as the inputs to the spatial-temporal model.

3.4.2 Temporal Correlation Modelling. 1-D convolution neural networks have shown strong performance in temporal feature modelling. We adopt 1-D dilated convolution neural networks to embed the temporal features. For ease of presentation, we simplify the notation of the input features of the first layer of the spatial-temporal model, $\mathbf{H}^{t-T+1:t,0}$, to \mathbf{H}^0 .

$$\mathbf{H}_{tcn}^l = *_{d_j^l} \sigma(W_{d_j^l}^l \mathbf{H}^{l-1}) \quad (5)$$

Here, $\mathbf{H}^{l-1} \in \mathbb{R}^{N \times T \times C'}$ is the output of the previous layer; $\mathbf{H}_{tcn}^l \in \mathbb{R}^{N \times T \times C'}$, $l = 1, 2, \dots, L$ is the output of the l -th layer's 1-D temporal convolution networks; $*_{d_j^l}$ means to stack 1-D dilated temporal convolution networks, where d_j^l represents the exponential dilation rate, $d_j^l = 2^j$. To keep the same dimensionality for the time sequence representation, we use zero-padding. Function $\sigma(\cdot)$ is the activation function (e.g., *ReLU* or *sigmoid*).

3.4.3 *Spatial Correlation Modelling.* We use a graph convolutional network (GCN) to model spatial correlations. The basic idea of GCN is to aggregate features from neighbours:

$$GCN(\mathbf{A}, \mathbf{H}) = \tilde{\mathbf{D}}^{-1/2} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-1/2} \mathbf{Z} \mathbf{W} \quad (6)$$

where $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}$, and $\tilde{\mathbf{D}}$ is a diagonal matrix. Matrix $\mathbf{Z} \in \mathbb{R}^{N \times C}$ is the input graph node features. Matrix $\mathbf{W} \in \mathbb{R}^{C \times C'}$ contains the parameters to be learned by the model, where C is the input dimensionality and C' is the output dimensionality. Next, we define GCN layers with two parallel GCNs, denoted as *GCNL*:

$$GCNL(\mathbf{A}, \mathbf{Z}) = GCN(\mathbf{A}, \mathbf{Z}) * \text{sigmoid}(GCN(\mathbf{A}, \mathbf{Z})) \quad (7)$$

We stack GCN layers to build a GCN block. The output of each GCN layer is the input of the next GCN layer, as shown in Eq. 8, where $q \in [1, k]$. The first layer input is $\mathbf{H}_{gcn,r}^{l,t-p,0} = \mathbf{H}^{l-1,t-p}$, where $p \in [T-1, 0]$ and $r \in \{s, dtw\}$ represents two types of adjacency matrices (i.e., spatial proximity-based matrix and a temporal similarity-based matrix).

$$\mathbf{H}_{gcn,r}^{l,t-p,q} = GCNL_{l,q-1}(\mathbf{A}_r, \mathbf{H}_{gcn,r}^{l,t-p,q-1}) \quad (8)$$

We use $\max(\cdot)$ to aggregate the outputs of the GCN layers to obtain the output of l -th GCN block (Eq. 9)

$$\mathbf{H}_{gcn,r}^{l,t-p} = \max_{q=1,\dots,k} (GCNL_{l,q}(\mathbf{A}_r, \mathbf{H}_{gcn,r}^{l,t-p,q})) \quad (9)$$

Then, we concatenate the output from each time slot:

$$\mathbf{H}_{gcn,r}^l = \parallel_{p \in [T-1, 0]} \mathbf{H}_{gcn,r}^{l,t-p} \quad (10)$$

After that, we use $\max(\cdot)$ again to aggregate the outputs corresponding to the two different adjacency matrices, as shown in Eq. 11, to obtain the output of the l -th layer:

$$\mathbf{H}_{gcn}^l = \max_{r \in \{s, dtw\}} (\mathbf{H}_{gcn,r}^l) \quad (11)$$

We follow previous studies [9, 16] and adopt Eq. 2 with different threshold - ϵ_s to compute the spatial-based adjacency matrix. Meanwhile, we follow another work [15] and adopt DTW [3] to compute \mathbf{A}_{dtw} , as described earlier.

We combine the output of TCN and GCN to obtain the output of the l -th layer:

$$\mathbf{H}^l = \mathbf{H}_{gcn}^l + \mathbf{H}_{tcn}^l \quad (12)$$

Finally, we obtain the output $\mathbf{H}^{t+1:t+T',L}$ at the L -th layer following the steps described above and linear functions to project $\mathbf{H}^{t+1:t+T',L}$ to a lower dimension (as Eq. 13 shows).

$$\hat{\mathbf{X}}^{t+1:t+T'} = \sigma(\phi_4(\sigma(\phi_3(\mathbf{H}^{t+1:t+T',L})))) \quad (13)$$

Here, ϕ_3 and ϕ_4 are linear functions, and σ is an activation function. $\hat{\mathbf{X}}^{t+1:t+T'} \in \mathbb{R}^{N_o \times T' \times C}$ represents the prediction values.

3.5 Model Training and Testing

Model Training: We obtain predicted values $\hat{\mathbf{X}}_{G_o^m}^{t+1:t+T'}$ on the graph view G_o^m that is generated by sub-graph masking. Then, we compute the mean squared error between the prediction values $\hat{x}_i^{t+p'}$ and the ground truth $x_i^{t+p'}$ as the prediction loss (Eq. 18).

$$L_{pred} = \frac{1}{N_o T} \sum_{i=1}^{N_o} \sum_{p'=1}^T \|\hat{x}_i^{t+p'} - x_i^{t+p'}\|_2^2 \quad (14)$$

Model Testing: During model testing, we first compute pseudo-observations for the unobserved locations, and let the graph G with pseudo-observations be G_m . Then, we build the temporal similarity-based adjacency matrix utilising these pseudo-observations. After that, we feed the features $\mathbf{X}_{G_m}^{t-T+1:t} \in \mathbb{R}^{N \times T \times C}$

into the trained model to produce the predicted observations $\hat{\mathbf{X}}_{G_m}^{t+1:t+T'}$ for the unobserved locations.

4 PROPOSED FULL MODEL

Section 3 introduced our base model. In this section, we will introduce two modules - *selective masking module* and *contrastive learning module* that enhance our proposed model performance. These two modules together with our base model STSM-RNC form our full model STSM. Figure 4 shows its overall structure.

Recall that our core idea is to learn a model that can extend the forecasting capability for masked locations to unobserved locations.

The generalisability of STSM on performance on the full graph G depends on the similarity between the masked locations in G_o^m and the unobserved locations in G . To mask the locations that have higher similarity with the unobserved locations, we propose a *selective masking module*, to enhance the *sub-graph masking*, exploiting the similarity between the observed locations and the unobserved locations to help forecast the values (e.g., traffic speed or PM2.5) at the unobserved locations. This module leverages the regional information and the road network information surrounding the locations, as well as the spatial distance to compute the masking probabilities of the locations. We mask the locations based on such probabilities at each model training epoch to generate G_o^m . This module can guide STSM to learn to predict values for the locations that have higher similarities with the unobserved locations, thus enhancing the generalisability of the model.

We further design a contrastive learning module that takes a graph contrastive learning-based approach and constructs two views of the graph - one view contains complete spatial-temporal data (the original view), and the other view contains incomplete spatial-temporal data (the augmentation view). The view with complete data is used to guide the prediction for the view with incomplete data. The augmentation view G_o^m is generated by the selective masking module. Using contrastive learning, we learn a model that generates similar predictions for both graph views. The trained model is then applied on the full graph G to make predictions for the unobserved locations.

We feed $\mathbf{X}_{G_o}^{t-T+1:t}$ and $\mathbf{X}_{G_o^m}^{t-T+1:t}$ into the proposed *spatial-temporal modelling module* (described in Section 3.4) to obtain the graph representations $\mathbf{Z}_{G_o}^{t+T'}$ and $\mathbf{Z}_{G_o^m}^{t+T'}$ for contrastive learning and generating the prediction result $\mathbf{X}_{G_o^m}^{t+1:t+T'}$.

When the model is trained, we feed $\mathbf{X}_{G_m}^{t-T+1:t}$ into the model to obtain predictions for the unobserved locations (Section 3.5).

4.1 Selective Masking

STSM learns to predict values for masked locations during training, and then extends this capability to predict values for unobserved locations at testing. Intuitively, the higher the similarity between the masked locations and unobserved locations, the easier it is for the trained model to make predictions for the unobserved locations. We compute similarities between those 1-hop sub-graphs (defined in Section 3.3) and the unobserved region. After that, we use our proposed selective masking module to guide STSM to mask a sub-region formed by the sub-graphs of observed locations that are the most similar to the unobserved region. Heuristically, this strategy leads to more accurate forecasting results for the locations in the unobserved region.

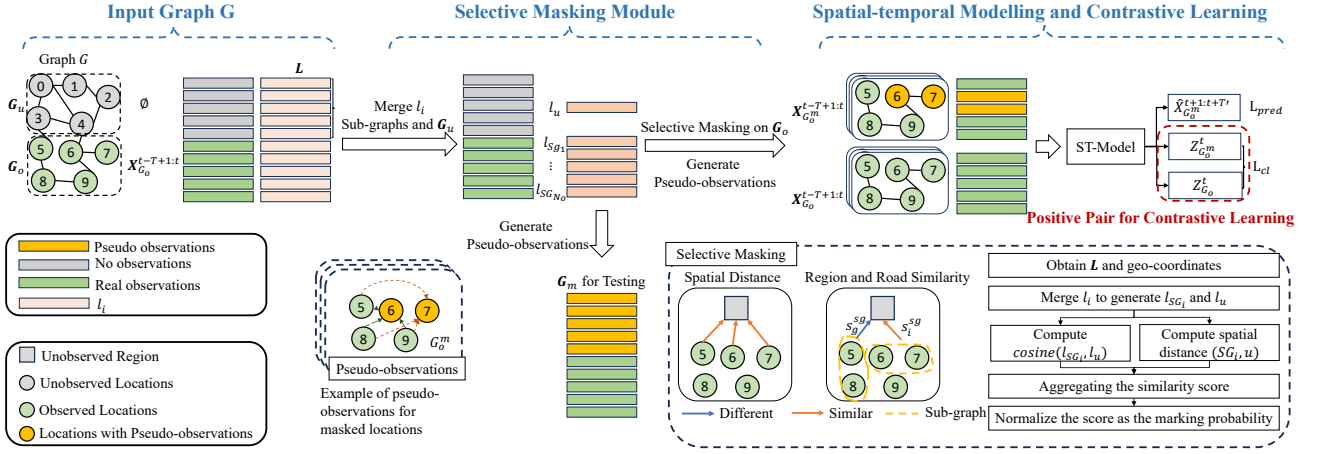


Figure 4: Model architecture of STSM. The model contains three main parts. (1) *The selective masking module* leverages the regional and road network representations and the spatial distances to compute the similarity between observed locations (i.e., their sub-graphs) and the unobserved region. Masking probabilities are assigned based on the similarity scores. (2) *The contrastive learning module* guides STSM to make similar predictions for location graphs with complete data and graphs with incomplete data. (3) *The spatial-temporal modelling module* (as described in Section 3.4) utilises GCNs and 1-D TCNs to model spatial and temporal features, together with a contrastive learning loss to optimise the model. To enhance model performance, STSM generates pseudo-observations for unobserved locations and computes a temporal similarity-based adjacency matrix. During the testing process, STSM fills unobserved locations with pseudo-observations and then feeds the graph into ST-Model to obtain the prediction results.

Table 1: Categories of Points of Interest

Categories	Subcategories	Categories	Subcategories
#1	university, college, school, kindergarten, research institute, language school, childcare	#2	commercial, office, studio
#3	retail, supermarket	#4	hotel, motel, guest house, hostel
#5	arts centre, library, events venue, community centre, conference centre, theatre, exhibition centre, planetarium, music venue, gallery, artwork, museum, social centre, zoo, aquarium, theme park	#6	clinic, hospital, veterinary, pharmacy, doctors, nursing home, dentist, social facility, baby hatch
#7	bridges	#8	cinema
#9	fountain, garden, park, trampoline park, water park, ranger station, dog park, viewpoint, nature reserve, attraction, bbq	#10	casino, gambling, nightclub, stripclub, dance
#11	church, chapel, cathedral, kingdom hall, monastery, mosque, presbytery, religious, shrine, synagogue, temple, place of worship	#12	cafe, ice cream, restaurant, pub, bar, food court, fast food, biergarten
#13	parking, parking entrance, parking exit, parking space, carport, motorcycle parking	#14	taxi, bus station, transportation, stop position, stop area, train station, platform, station
#15	warehouse	#16	industrial
#17	residential, apartment, apartments	#18	construction
#19	marketplace	#20	caravan site, camp site, camp pitch, picnic site, picnic table
#21	pitch, sports centre, sports hall, stadium, swimming area, swimming pool, track, grandstand, pavilion, riding hall, sports, fitness centre, fitness station	#22	civic, government, public
#23	fuel, car wash, car repair, vehicle inspection, car rental, car sharing	#24	atm, bank, bureau de change
#25	boat rental, ferry terminal, boat sharing	#26	barn, conservatory, cowshed, farm auxiliary, greenhouse, slurry tank, stable, sty

Sub-graph representation. To measure the similarity between a sub-graph and the region formed by the unobserved

locations, we need to first compute a representation for the sub-graph of each observed location. We form such a representation (i.e., an embedding) with three components:

- (1) *Point of interest (POI) features.* For each observed location, we draw a circle centered at the location with radius r (a system parameter) and collect all POIs inside the circle from OpenStreetMap [19]. We classify the POIs into Γ categories (cf. Table 1). The POI feature component of the sub-graph embedding, denoted by $l_i^{poi} \in \mathbb{R}^\Gamma$, is a vector that keeps a count of the POIs of each category. We further obtain the number of floors of the buildings and the areas of the parks in the circular region from OpenStreetMap [19] to indicate the prosperity of the sub-graph, as $l_i^{scale} \in \mathbb{R}^1$. For example, a sub-graph (i.e., a local region) with a 60-level building is more prosperous than a sub-graph with only a 4-level building. We concatenate l_i^{poi} and l_i^{scale} to obtain a sub-graph’s region embedding, denoted as $l_i^{region} = [l_i^{poi} || l_i^{scale}] \in \mathbb{R}^{\Gamma+1}$. Here, $||$ denotes concatenation.
- (2) *Road network features.* We select the nearest road of the location. To represent the road network corresponding to the sub-graph, we use a 4-dimensional vector $l_i^{road} \in \mathbb{R}^4$ where the dimensions represent *highway_level*, *maxspeed*, *is_oneway* and *number of lanes*.

Finally, we concatenate the regional representation and the road network representation of location i to form its embedding, i.e., $l_i = [l_i^{region} || l_i^{road}] \in \mathbb{R}^{\Gamma+5}$. The embedding of the sub-graph of location i , denoted by l_{SG_i} , is computed as the average embedding of the embeddings of all locations in the sub-graph, i.e., $l_{SG_i} = 1/|V_{SG_i}| \sum_{j \in V_{SG_i}} l_j$.

Similarity between sub-graphs and the unobserved region. Following the same strategy, we can compute an embedding l_u for the full unobserved region by averaging the embeddings of all unobserved location. Then, the similarity between the sub-graph of location i and the unobserved region is computed as the cosine similarity of the two embeddings, i.e., $s_i^{sg} = \text{cosine}(l_{SG_i}, l_u)$, combined with the spatial proximity $sp_i^{sg} = 1/\text{dist}(c_i, c_u)$ to guide the masking process.

We compute the embedding similarities between all sub-graphs and the unobserved region, denoted as $S^{sg} = [s_1^{sg}, \dots, s_{N_o}^{sg}]$, and the spatial proximity $SP^{sg} = [sp_1^{sg}, \dots, sp_{N_o}^{sg}]$.

Sub-graph masking. We use a masking ratio δ_m to define the percentage of observed locations to be masked. Since the sub-graph of each location may have a different size, we compute the average size of all sub-graphs, denoted as $\delta_s = \frac{1}{N_o} \sum_{i \in N_o} |V_{SG_i}|$. If we mask sub-graphs with the same probability $\delta_{ms} = \delta_m/\delta_s$, the final number of locations masked is expected to be $N_o \cdot \delta_m$. Since we want to use similarity to guide STSM to mask the observed locations, we combine the similarities S^{sg} , the spatial proximity SP^{sg} and the sub-graph masking ratio δ_{ms} to compute the masking probability for each location as Eq. 15 shows. This equation normalises S^{sg} and SP^{sg} to make them contribute equally.

$$P = (S^{sg} \cdot \frac{\delta_{ms}}{\frac{1}{N_o} \sum_{i \in N_o} s_i^{sg}} + SP^{sg} \cdot \frac{\delta_{ms}}{\frac{1}{N_o} \sum_{i \in N_o} sp_i^{sg}}) / 2 \quad (15)$$

The sub-graph size and the size of graph G impact the value $p_i \in P$. When they are large, p_i can become very small, such that all sub-graphs have very close probability values. To address this issue, we only keep the top- K most similar sub-graphs and set the similarity values as 0 for the remaining sub-graphs, where K is a hyper-parameter. This strategy reduces the number of sub-graphs that can be masked in a graph. Then, we mask the

sub-graphs more similar to the unobserved region, based on a masking probability ρ_i , drawn from a Bernoulli distribution $\rho_i \sim \text{Bern}(p_i)$ to generate a graph with masked locations (G_o^m).

4.2 Graph Contrastive Learning

Based on Section 4.1, we use G_o (i.e., the graph with complete data) to generate a graph G_o^m with incomplete data (i.e., a graph with masked locations). Graph G_o^m can be regarded as G_o with perturbations (i.e., G_o^m and G_o are two views of the observed graph, and G_o^m is an augmentation of G_o). To guide STSM to produce similar predictions on G_o^m and G_o , we apply contrastive learning in the training process on these two views of the graph.

STSM adopts graph-level contrastive learning. We use the original graph G_o to explain our graph representation generating steps. First, $[\mathbf{X}_{G_o}^{t-T+1}, \dots, \mathbf{X}_{G_o}^t]$ is inputted into the spatial-temporal model (as described in Section 3.4) to generate an output for each time slot, denoted as $\mathbf{H}_{G_o}^{t+T',L}$, where L is the number of layers in the spatial-temporal model. Then, STSM takes the last layer output of the spatial-temporal model for the last time step, i.e., $\mathbf{H}_{G_o}^{t+T',L}$, to obtain a graph representation $\mathbf{Z}_{G_o}^{t+T'}$. We aggregate all locations’ representations and project them into a new latent space, formulated as Eq. 16, where $\phi(\cdot)$ is a *linear function*. We follow the same steps to generate the representation $\mathbf{Z}_{G_o^m}^{t+T'}$ of G_o^m .

$$\mathbf{Z}_{G_o}^{t+T'} = \phi(\text{ReLU}(\phi(\sum_{i \in N_o} h_{i,G_o}^{t+T',L}))) \quad (16)$$

A batch of M input time windows are sampled at training, which form $2M$ representations, where $(\mathbf{Z}_{G_o}^{t+T'}, \mathbf{Z}_{G_o^m}^{t+T'})$ is a positive pair (i.e., graph G_o and graph G_o^m from the same time slot form a positive pair). Negative pairs are generated from the other $M-1$ graphs in the batch (i.e., graph G_o and graph G_o^m from different time slots in a batch form negative pairs), denoted as $(\mathbf{Z}_{G_o}^{t+T'}, \mathbf{Z}_{G_o^m}^{t'+T'})$. We adopt the contrastive loss to maximise the mutual information of the sample pairs:

$$L_{cl} = -\frac{1}{M} \sum_{t \in M} \log \frac{\exp(\text{sim}(\mathbf{Z}_{G_o}^{t+T'}, \mathbf{Z}_{G_o^m}^{t+T'})/\tau)}{\sum_{t' \in M, t' \neq t} \exp(\text{sim}(\mathbf{Z}_{G_o}^{t+T'}, \mathbf{Z}_{G_o^m}^{t'+T'})/\tau)} \quad (17)$$

The final loss function to optimise STSM is:

$$L = L_{pred} + \lambda L_{cl} \quad (18)$$

where λ is a coefficient to balance the prediction loss and the contrastive learning loss.

5 EXPERIMENTS

5.1 Experimental Setup

5.1.1 Datasets. We conduct experiments on three highway traffic datasets, an urban traffic dataset and an air quality dataset.

- **PEMS-Bay** [16] contains traffic speed data collected from 325 sensors on highways in the Bay Area, California, between January and June 2017.
- **PEMS-07** [4] contains traffic speed data collected by sensors on highways in Los Angeles. Following a previous study [20], we randomly sample 400 sensors and use their collected data between September and December 2022 as the dataset.
- **PEMS-08** [4] contains traffic speed data collected by sensors on highways in the San Bernardino area, California.

Similarly, we use data collected by 400 randomly sampled sensors between September and December 2022.

- **Melbourne** contains traffic speed data collected by 182 sensors in Melbourne City, Australia, between July and September 2022, from the AIMES project [26].
- **AirQ** [40] contains pollutant concentration data (PM2.5) collected by 63 sensors in Beijing and Tianjin, which are two adjacent cities in China, between May 2014 and April 2015.

All traffic records collected from PEMS are given in 5-minute windows, i.e., 288 time slots per day, while the traffic records of the Melbourne dataset are given in 15-minute windows, i.e., 96 time slots per day. The air quality records are given in hourly windows, i.e., 24 time slots per day. Table 2 summarises the dataset statistics and Fig. 5 visualises the sensor distribution among all datasets. The region and road network information used for selective masking is obtained from OpenStreetMap [19].

Following a baseline work [30], we use the records in the past two hours to make predictions for the next two hours, i.e., $T = T' = 2hours$ in Eq. 1 for the traffic datasets. Following another baseline work [39], we use the records in the past 24 hours to make predictions for the next 24 hours, i.e., $T = T' = 24hours$ in Eq. 1 for the air quality dataset. We split each dataset into three sets by 4:1:5 for training, validation and testing, where the locations within each set are adjacent to each other. Note that, the locations in the training set and the validation set are considered as observed locations, while those in the test set are considered as unobserved locations. Dataset split is space-based, where the sensors are divided horizontally or vertically into three sets based on the sensor geo-coordinates. For each dataset, we create four different splits and report the average performance on each dataset. We use the data recorded during the first 70% of the time for the training, and the last 30% of the time for testing. Fig. 6 shows a dataset partitioning on PEMS-Bay and its temporal partitioning.

Table 2: Dataset Statistics

Dataset	Time period	Interval	#Sensors
PEMS-Bay	01/01/2017 - 30/06/2017	5 min	325
PEMS-07	01/09/2022 - 31/12/2022	5 min	400
PEMS-08	01/09/2022 - 31/12/2022	5 min	400
Melbourne	01/07/2022 - 30/09/2022	15 min	182
AirQ	01/05/2014 - 30/04/2015	1 hr	63

5.1.2 Competitors. There are no existing models for our proposed problem. We adopt the following adapted models for an empirical comparison with our proposed model STSM:

- **GE-GAN** [32] is a transductive data imputation method based on generative adversarial networks (GAN) that utilise a generator to generate estimating values and a discriminator to classify the real and generated values.
- **IGNNK** [30] is an inductive graph neural network for spatial-temporal Kriging.
- **INCREASE** [39] is an inductive graph representation learning network based on GRUs and the state-of-the-art for spatial-temporal Kriging.

5.1.3 Implementation Details. We use the default settings of the baseline models from their source code. The baseline models were proposed for data imputation, while we aim for prediction.

We change their ground truth to the future time window rather than the past time window to train the models and obtain the prediction.

We train our model using the Adam optimiser with the learning rate starting at 0.01. The batch size is 32. For the hyper-parameters in our models, τ is 0.5, σ_m is 0.5, ϵ_s is 0.05 and q_{kk} and q_{ku} are set to 1. We leave the details of the other model hyper-parameters (i.e., λ , σ_{sg} , r_{poi} and K) in Table 3. These parameter values are obtained through grid search on the validation set, except r_{poi} which is fine-tuned only based on the similarity between sub-graphs and the unobserved regions (i.e., S^{sg}). Also, parameter values can be shared among datasets with similar distributions, e.g., when only the number or the density of sensors is changed (cf., Table 6 and Table 7). In the experiments, we use different thresholds for the spatial-based matrices A_s and A_{sg} . Fig. 7 visualises the two adjacency matrices on PEMS-Bay.

The experiments are run on an NVIDIA Tesla V100 GPU.

Table 3: Parameter settings

Parameter	PEMS-Bay	PEMS-07	PEMS-08	Melbourne	AirQ
λ	0.01	1	0.5	0.5	1
ϵ_{sg}	0.5	0.7	0.5	0.4	0.6
r_{poi} (m)	200	500	500	50	500
K	35	35	35	45	5

We adopt four commonly used metrics to evaluate model performance, including root mean squared error (RMSE), mean absolute error (MAE), mean absolute percentage error (MAPE) and R-square (R2). The first three measure the forecasting errors, while R2 measures how much better the model prediction results are compared with just using average observations as results [39].

5.2 Experimental Results

We first compare the overall performance of our model with those of the baseline methods. Then, we report the results of an ablation study to verify the effectiveness of each module in STSM. Finally, we study the impact of parameters to test model robustness.

5.2.1 Model Performance Comparison. We first compare STSM with the baseline methods.

(1) **Overall Results.**: Table 4 summarises the overall performance results. STSM and its variants including the base model STSM-RNC (detailed in Section 5.2.2) outperform all the competitors on all four datasets, except for the MAPE measure on AirQ.

GE-GAN is a transductive method that generates values for unobserved locations utilising their similar locations based on graph embeddings. It is difficult to find similar locations when there are many unobserved locations in a large area, resulting in poor forecasting accuracy. In urban areas such as Melbourne City, GE-GAN outperforms the other two baseline models because this area is relatively small.

IGNNK is an inductive model with GNNs to model spatial correlations and 1-D convolution neural networks to capture temporal correlations. It struggles in our task because data missing at continuous locations makes it difficult for the GNNs to learn the spatial correlation patterns. Although it has a slightly low MAPE on AirQ, its MAE and RMSE are still much larger than those of our model STSM on this dataset. This can be explained




Figure 5: Visualisations of sensor distribution

Table 4: Overall model performance. “↓” (and “↑”) indicates that lower (and larger) values are better. The best baseline results are underlined, and the best results of the proposed model STSM are in **bold**. Improvement computes the errors reduced by the best variant of the proposed model compared with the best baseline model, where N/A indicates that the improvement results cannot be calculated, because of negative measurement values on the baseline methods.

Dataset	Metric	GE-GAN	IGNNK	INCREASE	STSM-RNC	STSM-NC	STSM-R	STSM	Improvement
PEMS-Bay	RMSE↓	31.184	9.611	<u>8.820</u>	8.626	8.642	8.628	8.610	+2.38%
	MAE↓	26.020	5.616	<u>5.243</u>	5.352	5.352	5.120	5.237	+0.11%
	MAPE↓	0.432	0.145	<u>0.133</u>	0.130	0.131	0.128	0.130	+2.26%
	R2↑	-9.042	0.063	<u>0.195</u>	0.228	0.226	0.230	0.231	+18.46%
PEMS-07	RMSE↓	21.103	11.556	<u>8.465</u>	8.516	8.512	8.444	8.390	+0.89%
	MAE↓	15.647	9.167	<u>5.495</u>	5.352	5.411	5.417	5.111	+6.99%
	MAPE↓	0.273	0.181	<u>0.125</u>	0.125	0.126	0.125	0.123	+1.60%
	R2↑	-4.351	-0.715	<u>0.155</u>	0.156	0.142	0.157	0.169	+9.03%
PEMS-08	RMSE↓	23.409	10.599	<u>8.275</u>	8.021	7.937	8.018	7.925	+4.23%
	MAE↓	17.611	7.934	<u>5.019</u>	4.982	4.962	5.006	4.899	+2.39%
	MAPE↓	0.299	0.158	<u>0.116</u>	0.115	0.115	0.116	0.114	+1.72%
	R2↑	-6.531	-0.678	<u>0.056</u>	0.114	0.134	0.115	0.136	+142.86%
Melbourne	RMSE↓	<u>10.064</u>	14.635	10.321	9.844	9.884	10.07	9.175	+8.83%
	MAE ↓	<u>7.780</u>	12.511	8.302	7.803	7.188	7.882	7.308	+7.61%
	MAPE ↓	<u>0.369</u>	0.746	0.453	0.399	0.366	0.389	0.388	+0.81%
	R2↑	<u>-0.175</u>	-2.355	-0.266	-0.120	0.027	-0.171	0.027	N/A
AirQ	RMSE↓	295.579	74.873	<u>73.977</u>	69.956	68.126	68.968	67.571	+8.66%
	MAE ↓	244.824	<u>52.726</u>	56.165	50.301	48.451	49.302	48.141	+14.29%
	MAPE ↓	9.142	<u>1.459</u>	2.168	1.789	1.643	1.737	1.692	-12.61%
	R2↑	-17.917	-0.067	<u>-0.024</u>	0.074	0.123	0.099	0.141	N/A

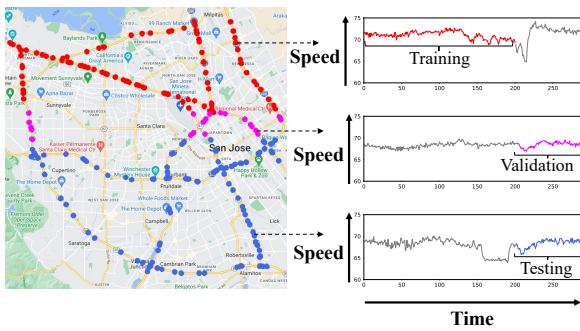


Figure 6: Data partitioning on PEMS-Bay from spatial (left; horizontal partitioning) and temporal (right) perspectives. The red, pink and blue dots on the map represent the observed locations for training, the observed locations for validation and the unobserved locations for testing, respectively.

by that a lower MAE happens on smaller observations while a higher MAE happens on larger observations.

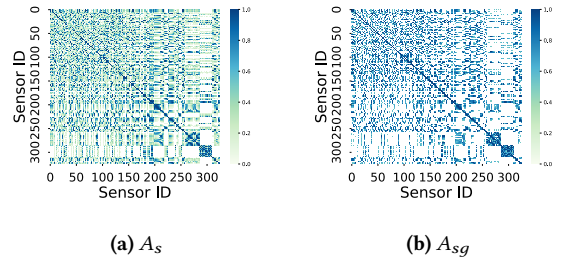


Figure 7: Visualization of adjacency matrices. The blank density in these figures reflects the sparsity of the adjacency matrix. The right figure has more blank space due to we use a larger threshold to limit sub-graph size.

INCREASE, the state-of-the-art spatial-temporal Kriging model, learns heterogeneous spatial relations and diverse temporal patterns, which presents the best performance among the baseline models. However, it is still outperformed by our model STSM. Our model reduces forecasting errors by up to 14% on the AirQ dataset and increases R2 by up to 142% on the PEMS-08 dataset

because of our temporal adjacency matrix to model the temporal similarity, our selective masking module to model spatial and semantic similarity, and contrastive learning to enhance model robustness.

We also report the model training and testing time over the traffic datasets. We omit the running time results on AirQ due to its small scale. The training time of all models is at the same scale. GE-GAN requires more training epochs to converge. However, when it comes to testing time, GE-GAN and STSM are faster than IGNNK and INCREASE.

Table 5: Model training time

Model	Time	PEMS-Bay	PEMS-07	PEMS-08	Mel.
GE-GAN	Train (h)	4.4	4.1	4.1	0.3
	Test (s)	0.9	0.7	0.8	0.1
IGNNK	Train (h)	0.3	0.2	0.2	0.1
	Test (s)	8.3	7.8	8.8	2.4
INCREASE	Train (h)	0.3	0.2	0.2	0.2
	Test (s)	9.5	7.3	7.5	4.0
STSM	Train (h)	1.1	1.9	2.2	0.3
	Test (s)	1.6	1.3	1.2	0.3

(2) *Varying the unobserved ratio.* We vary the unobserved ratio from 0.2 to 0.5 on all datasets, i.e., from 20% to 50% of all the sensor locations in a dataset are treated as unobserved locations. Same as before, we split each dataset horizontally or vertically and report the average performance on the four setups (each split creates two alternative settings of training and testing sets). Fig. 8 presents the results. Since INCREASE has the best performance among the baselines when varying the unobserved ratio, we only show its results for this set of experiments. STSM outperforms INCREASE in all settings, except when 20% of the locations are considered unobserved on PEMS-08. We notice that sometimes the prediction errors drop even with a higher ratio of unobserved locations. This is because some unobserved locations are easier than others to predict. Including such unobserved locations reduces the mean prediction error. Here, we have only shown results in RMSE. Results on other metrics show similar patterns, which are omitted for conciseness. The same applies to the experiments below.

(3) *Varying the number of sensors.* We merge PEMS08 and PEMS07 into a larger region, such that we can vary the number of sensors from 200 to 800 by vertically splitting the space (and hence the sensor locations) into four equal-sized partitions (i.e., each partition contains 200 sensors) based on geo-coordinates. Table 6 shows reports how the model prediction errors change with more sensors are added into the dataset. We see that our model STSM consistently outperforms all three baseline models in terms of RMSE and R2.

(4) *Varying the density of sensors.* We further vary the number of sensors from 200 to 964 (which is the maximum number of sensors) on PEMS-08, to test the impact of the density of the sensors. The results in Table 7 show that STSM again outperforms all baseline models in almost all cases (i.e., 19 out of 20), further confirming the robustness of the model.

5.2.2 Ablation Study. We conduct an ablation study with three variants of STSM:

STSM-NC disables the contrastive learning module.

STSM-R replaces the selective masking module with a random masking module that randomly chooses a location and its 1-hop neighbours to be masked until the target masking ratio is reached.

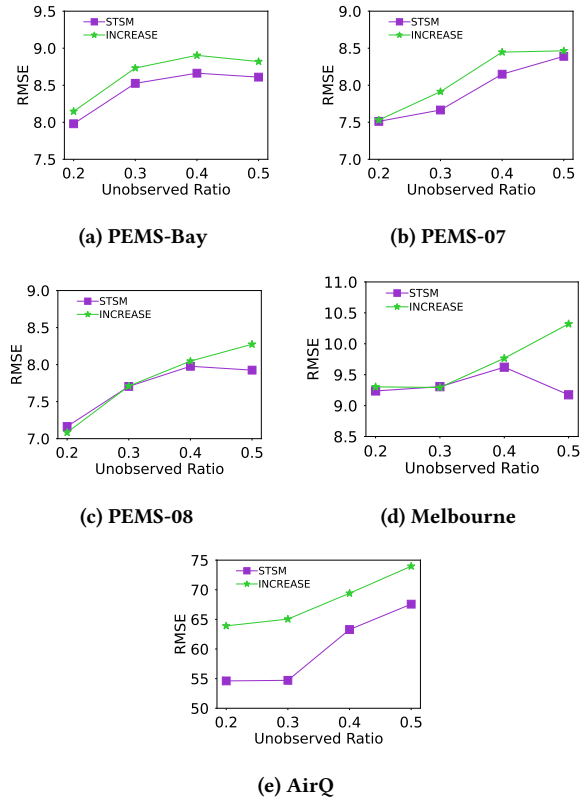


Figure 8: Model performance vs. unobserved ratio

Table 6: Varying the number of sensors. “↓” (and “↑”) indicates that lower (and larger) values are better. The best baseline results are underlined, and the best model results are in **bold.**

#Sensor number	Model	RMSE↓	MAE↓	MAPE↓	R2↑
200	GE-GAN	24.442	18.348	0.309	-7.828
	IGNNK	10.845	7.774	0.157	-0.950
	INCREASE	<u>7.800</u>	<u>4.816</u>	<u>0.110</u>	<u>0.129</u>
	STSM	7.775	4.649	0.109	0.133
400	GE-GAN	21.929	16.244	0.283	-4.702
	IGNNK	15.896	12.709	0.238	-2.998
	INCREASE	<u>8.880</u>	<u>5.115</u>	<u>0.130</u>	<u>0.115</u>
	STSM	8.718	5.384	0.131	0.141
600	GE-GAN	23.105	17.043	0.296	-5.243
	IGNNK	11.492	9.428	0.229	-1.089
	INCREASE	<u>8.658</u>	<u>5.494</u>	<u>0.129</u>	<u>0.142</u>
	STSM	8.629	5.548	0.130	0.148
800	GE-GAN	21.801	15.867	0.280	-4.349
	IGNNK	12.113	8.530	0.172	-1.295
	INCREASE	<u>8.370</u>	<u>5.055</u>	<u>0.119</u>	<u>0.113</u>
	STSM	8.134	5.008	0.118	0.165

STSM-RNC (our base model as described in Section 3) replaces the selective masking module with the random masking module and disables contrastive learning.

(1) *Impact of selective masking.* As Table 4 shows, STSM outperforms STSM-R among all datasets, except for the MAPE and MAE on PEMS-Bay. We further compare the similarity between the masked sub-graphs and the unobserved region in the training process. Table 8 presents the results, which shows that selective

Table 7: Varying the density of sensors. “↓” (and “↑”) indicates that lower (and larger) values are better. The best baseline results are underlined, and the best model results are in **bold.**

#Sensor number	Model	RMSE↓	MAE↓	MAPE↓	R2↑
200	GE-GAN	24.293	18.102	0.306	-7.201
	IGNNK	13.847	9.981	0.194	-2.623
	INCREASE	<u>8.391</u>	<u>5.188</u>	<u>0.117</u>	0.016
	STSM	7.889	4.707	0.111	0.134
400	GE-GAN	23.409	17.611	0.299	-6.531
	IGNNK	10.599	7.934	0.158	-0.678
	INCREASE	<u>8.275</u>	<u>5.019</u>	<u>0.116</u>	<u>0.056</u>
	STSM	7.925	4.899	0.114	0.136
600	GE-GAN	23.933	17.695	0.298	-7.400
	IGNNK	12.348	9.198	0.175	-1.745
	INCREASE	<u>7.982</u>	<u>5.109</u>	<u>0.110</u>	<u>0.061</u>
	STSM	7.708	4.936	0.110	0.127
800	GE-GAN	22.711	17.016	0.288	-6.405
	IGNNK	13.172	10.203	0.191	-2.006
	INCREASE	<u>8.055</u>	<u>4.847</u>	<u>0.114</u>	<u>0.072</u>
	STSM	7.841	4.942	0.114	0.119
964	GE-GAN	22.180	16.685	0.283	-6.186
	IGNNK	12.476	9.735	0.184	-1.639
	INCREASE	<u>8.052</u>	<u>4.851</u>	<u>0.114</u>	<u>0.064</u>
	STSM	7.831	4.757	0.111	0.113

masking can guide the model to mask sub-graphs with higher similarities to the unobserved regions.

We also compare the performance of STSM-NC and STSM-RNC. STSM-NC outperforms STSM-RNC on PEMS-08, Melbourne and AirQ. On PEMS-Bay and PEMS-07, STSM-NC and STSM-RNC yield similar performance. These results confirm the importance of the selective masking module.

(2) *Impact of contrastive learning.* Table 4 shows that STSM outperforms STSM-NC over all freeway traffic datasets. On the urban datasets (i.e., Melbourne and AirQ), STSM has better RMSE while STSM-NC is better at MAE or MAPE. Besides, STSM-R outperforms STSM-RNC in most cases (14 out of 20). These results confirm that contrastive learning is also an important component that contributes to the strong model performance.

Table 8: Similarity gain compared with random masking

Dataset	PEMS-Bay	PEMS-07	PEMS-08	Mel.	AirQ
Sim. Gain (%)	9.35	14.76	5.39	6.87	19.66

5.2.3 *Parameter Study.* We test the impact of K and ϵ_{sg} .

(1) *Impact of the number of top similar locations (sub-graphs) K .* This parameter determines the number of sub-graphs that may be masked (i.e., impacts the value of P in Eq. 15). It impacts the performance of STSM and STSM-NC because these two model variants use the selective masking module. The results in Fig. 9 show that the performance of both STSM and STSM-NC is more stable on freeway traffic datasets than on the other datasets. The reason is that the freeway datasets contain more sensors than the Melbourne and AirQ datasets (i.e., the parameter changing ratio impacts models’ sensitivity).

(2) *Impact of the spatial-based matrix’s threshold ϵ_{sg} .* This parameter is used to control the size of sub-graphs. When ϵ_{sg} becomes larger, the sub-graph size becomes smaller (i.e., fewer locations in each sub-graph) because of fewer links in the graph.

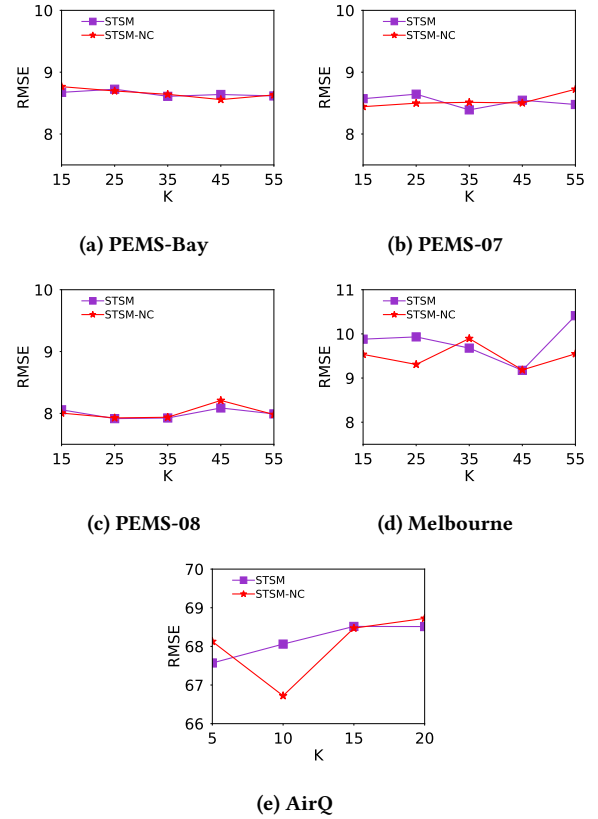


Figure 9: Model performance vs. K

STSM and its variants all mask locations based on the sub-graphs, and we test the impact of ϵ_{sg} on all of them. Fig. 10 shows the results. STSM and its variants are again robust about this parameter, especially on the freeway traffic datasets. For Melbourne and AirQ, their smaller numbers of sensors and the complex urban road network information lead to more sensitivities. Note that the fluctuations are quite small compared with the traffic and air pollutant observation values.

5.2.4 *Impact of Space Splits.* The relative position of the unobserved and observed regions can impact the model result. Our experiments above use horizontal or vertical space splits. To further verify the robustness of STSM, considering the circular nature of many city layouts, we study another space splitting strategy (i.e., “ring” splitting) as shown in Fig. 11 - the centre region is the observed region (red dots) for training, the region in the middle ring (pink dots) is for validation, and the outer regions are unobserved (blue dots) for testing. We conduct experiments on PEMS-Bay with this strategy. As shown in Table 9, STSM again outperforms all baseline models consistently, with an advantage of up to 9% in terms of R2.

5.2.5 *Impact of Temporal Correlation Learning.* The techniques used to capture temporal correlation patterns can impact the model effectiveness. We have used 1-D CNN in STSM for its simplicity. In this set of experiments, we further explore the extensibility of STSM to incorporate advanced temporal correlation modelling techniques. We replace 1-D CNN with a transformer encoder (which is an advanced sequence learning model) and a gated fusion module [38] to fuse each block’s spatial and temporal embeddings. We denote this variant as *STSM-trans*. Table 10

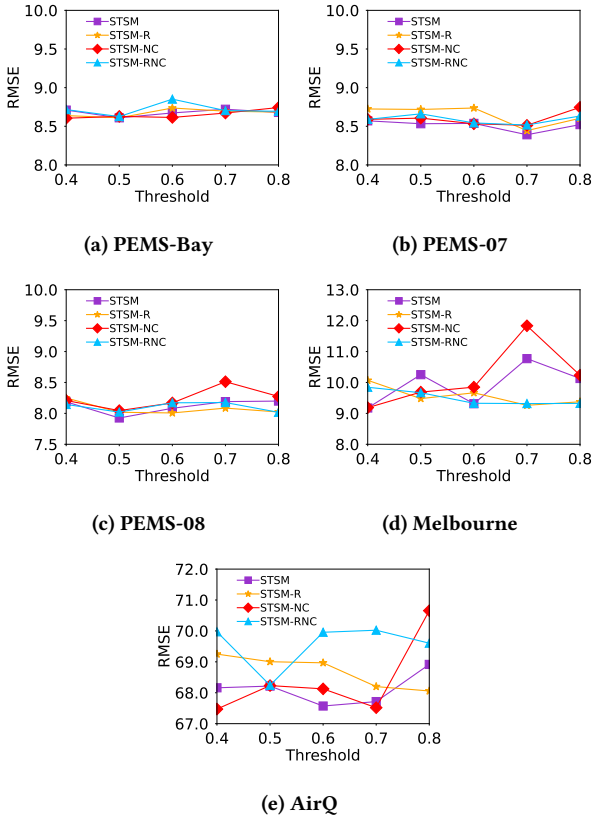


Figure 10: Model performance vs. ϵ_{sg}

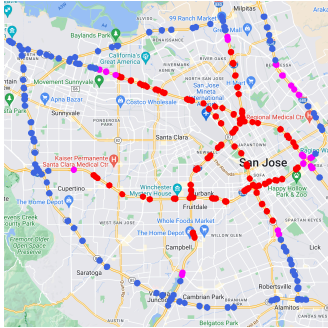


Figure 11: Visualization of different types of sensor distribution on PEMS-Bay

Table 9: Model performance on PEMS-Bay (Ring Split)

Model	RMSE	MAE	MAPE	R2
GE-GAN	32.674	26.740	0.447	-10.612
IGNNK	12.625	10.233	0.2	-0.737
INCREASE	8.558	4.807	0.125	0.190
STSM	8.462	4.757	0.124	0.208
Improvement	+1.1%	+1.0%	+0.8%	+9.5%

presents the experimental results on PEMS-Bay. Overall, STSM-trans outperforms STSM, which verifies the extensibility of STSM to incorporate advanced correlation pattern learning models.

Table 10: Model performance with advanced temporal correlation learning modules on PEMS-Bay

Model	RMSE	MAE	MAPE	R2
STSM	8.610	5.237	0.130	0.231
STSM-trans	8.562	5.251	0.129	0.240

5.2.6 *Impact of Distance Functions.* We used Euclidean distance in our model for efficiency consideration. Road network distance is an alternative choice. To study the impact of the distance function, we compare STSM (using the Euclidean distance) with two variants: STSM-rd-a uses road network distances for computing the adjacency matrices (i.e., A_s and A_{sg}) and pseudo-observations, while STSM-rd-m uses road network distance for computing the adjacency matrices (i.e., A_s and A_{sg}) only. Table 11 shows that STSM has the best performance among all variants, which verifies that Euclidean distance is efficient and effective for our model. STSM-rd-m performs better than STSM-rd-a because Euclidean distance leads to better quality of the pseudo-observations.

Table 11: Model performance when using different distance functions on PEMS-Bay

Model	RMSE	MAE	MAPE	R2
STSM	8.610	5.237	0.130	0.231
STSM-rd-a	8.993	5.426	0.134	0.158
STSM-rd-m	8.708	5.339	0.132	0.213

6 CONCLUSIONS AND FUTURE WORK

We proposed a new task - spatial-temporal forecasting for a region of interest without historical observations while this region’s adjacent region has such data. We design a novel model named STSM for the task. We propose a selective masking module based on region, road network and spatial distance features. This module can guide STSM to mask locations in the adjacent region that have higher similarity with those in the region of interest, which is beneficial for extending the forecasting capability of STSM to the region of interest. Besides, STSM exploits contrastive learning to enhance model forecasting efficacy. Extensive experimental results on real-world datasets, including traffic data and air quality data, show that STSM outperforms adapted state-of-the-art models consistently in forecasting accuracy. This advantage benefits from (1) the selective masking module which guides the model to mask regions more similar to the region of interest, resulting in better generalization of predictions, (2) contrastive learning which enhances model accuracy over incomplete data, and (3) temporal similarity based adjacency matrix computation which strengthens the learning capability of GCNs, allowing messages passing from observed locations to unobserved locations.

We only considered one unobserved region. In the future, we plan to extend STSM to deal with multiple unobserved regions at the same time.

ACKNOWLEDGMENTS

This work is partially supported by Australian Research Council (ARC) Discovery Project DP230101534.

REFERENCES

- [1] Gabriel Appleby, Linfeng Liu, and Li-Ping Liu. 2020. Kriging Convolutional Networks. In *AAAI*, Vol. 34. 3187–3194.
- [2] Mohammad Taha Bahadori, Qi Rose Yu, and Yan Liu. 2014. Fast Multivariate Spatio-Temporal Analysis via Low Rank Tensor Learning. In *NeurIPS*, Vol. 27. 3491–3499.
- [3] Donald J Berndt and James Clifford. 1994. Using Dynamic Time Warping to Find Patterns in Time Series. In *KDD workshop*. 359–370.
- [4] California Transportation Agencies (CalTrans) Performance Measurement System (PeMS). 2001. <https://pems.dot.ca.gov>.
- [5] Yanchuan Chang, Egemen Tanin, Xin Cao, and Jianzhong Qi. 2023. Spatial Structure-Aware Road Network Embedding via Graph Contrastive Learning. In *EDBT*, Vol. 26. 144–156.
- [6] Razvan-Gabriel Cirstea, Bin Yang, Chenjuan Guo, Tung Kieu, and Shirui Pan. 2022. Towards Spatio-temporal Aware Traffic Time Series Forecasting. In *ICDE*. 2900–2913.
- [7] Noel Cressie and Christopher K Wikle. 2015. *Statistics for Spatio-Temporal Data*.
- [8] Yuchen Fang, Yanjun Qin, Haiyong Luo, Fang Zhao, Bingbing Xu, Liang Zeng, and Chenxing Wang. 2023. When Spatio-Temporal Meet Wavelets: Disentangled Traffic Forecasting via Efficient Spectral Graph Attention Networks. In *ICDE*. 517–529.
- [9] Zheng Fang, Qingqing Long, Guojie Song, and Kunqing Xie. 2021. Spatial-Temporal Graph ODE Networks for Traffic Flow Forecasting. In *KDD*. 364–373.
- [10] Shengnan Guo, Youfang Lin, Ning Feng, Chao Song, and Huaiyu Wan. 2019. Attention based Spatial-Temporal Graph Convolutional Networks for Traffic Flow Forecasting. In *AAAI*, Vol. 33. 922–929.
- [11] Shengnan Guo, Youfang Lin, Letian Gong, Chenyu Wang, Zeyu Zhou, Zekai Shen, Yiheng Huang, and Huaiyu Wan. 2023. Self-Supervised Spatial-Temporal Bottleneck Attentive Network for Efficient Long-term Traffic Forecasting. In *ICDE*. 1585–1596.
- [12] Jiahao Ji, Jingyuan Wang, Chao Huang, Junjie Wu, Boren Xu, Zhenhe Wu, Junbo Zhang, and Yu Zheng. 2022. Spatio-Temporal Self-Supervised Learning for Traffic Flow Prediction. In *AAAI*, Vol. 37. 4356–4364.
- [13] Jiawei Jiang, Chengkai Han, Wayne Xin Zhao, and Jingyuan Wang. 2023. PDFFormer: Propagation Delay-aware Dynamic Long-range Transformer for Traffic Flow Prediction. In *AAAI*, Vol. 37. 4365–4373.
- [14] Ibai Laña, Ignacio Iñaki Olabarrieta, Manuel Velez, and Javier Del Ser. 2018. On the Imputation of Missing Data for Road Traffic Forecasting: New Insights and Novel Techniques. *Transportation Research part C: emerging technologies* 90 (2018), 18–33.
- [15] Mengzhang Li and Zhanxing Zhu. 2021. Spatial-Temporal Fusion Graph Neural Networks for Traffic Flow Forecasting. In *AAAI*, Vol. 35. 4189–4196.
- [16] Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. 2018. Diffusion Convolutional Recurrent Neural Network: Data-Driven Traffic Forecasting. In *ICLR*.
- [17] Yuebing Liang, Zhan Zhao, and Lijun Sun. 2022. Memory-Augmented Dynamic Graph Convolution Networks for Traffic Data Imputation with Diverse Missing Patterns. *Transportation Research Part C: Emerging Technologies* 143 (2022), 103826.
- [18] Xu Liu, Yuxuan Liang, Chao Huang, Yu Zheng, Bryan Hooi, and Roger Zimmermann. 2022. When Do Contrastive Learning Signals Help Spatio-Temporal Graph Forecasting?. In *SIGSPATIAL*. 1–12.
- [19] OpenStreetMap US Northeast data dump. 2018. <https://download.geofabrik.de/>.
- [20] Jianzhong Qi, Zhuwei Zhao, Egemen Tanin, Tingru Cui, Neema Nassir, and Majid Sarvi. 2022. A Graph and Attentive Multi-Path Convolutional Network for Traffic Prediction. *TKDE* 35 (2022), 6548–6560.
- [21] Weizhu Qian, Dalin Zhang, Yan Zhao, Kai Zheng, and James J.Q. Yu. 2023. Uncertainty quantification for traffic forecasting: A unified approach. In *ICDE*. 992–1004.
- [22] 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 *AAAI*, Vol. 34. 914–921.
- [23] Fan-Yun Sun, Jordon Hoffman, Vikas Verma, and Jian Tang. 2020. InfoGraph: Unsupervised and Semi-supervised Graph-Level Representation Learning via Mutual Information Maximization. In *ICLR*.
- [24] Koh Takeuchi, Hisashi Kashima, and Naonori Ueda. 2017. Autoregressive Tensor Factorization for Spatio-Temporal Predictions. In *ICDM*. 1105–1110.
- [25] Yihong Tang, Ao Qu, Andy HF Chow, William HK Lam, SC Wong, and Wei Ma. 2022. Domain Adversarial Spatial-Temporal Network: A Transferable Framework for Short-Term Traffic Forecasting Across Cities. In *CIKM*. 1905–1915.
- [26] The Australian Integrated Multimodal EcoSystem. 2016. <https://eng.unimelb.edu.au/industry/aimes>.
- [27] TomTom International BV. 2023. <https://www.tomtom.com/traffic-index/>.
- [28] Petar Velickovic, William Fedus, William L Hamilton, Pietro Liò, Yoshua Bengio, and R Devon Hjelm. 2019. Deep Graph Infomax. In *ICLR*.
- [29] Christopher KI Williams and Carl Edward Rasmussen. 2006. *Gaussian Processes for Machine Learning*. Vol. 2.
- [30] Yuankai Wu, Dingyi Zhuang, Aurelie Labbe, and Lijun Sun. 2021. Inductive Graph Neural Networks for Spatiotemporal Kriging. In *AAAI*, Vol. 35. 4478–4485.
- [31] Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, and Chengqi Zhang. 2019. Graph WaveNet for Deep Spatial-Temporal Graph Modeling. In *IJCAI*. 1907–1913.
- [32] Dongwei Xu, Chenchen Wei, Peng Peng, Qi Xuan, and Haifeng Guo. 2020. GEGAN: A Novel Deep Learning Framework for Road Traffic State Estimation. *Transportation Research Part C: Emerging Technologies* 117 (2020), 102635.
- [33] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*. 2048–2057.
- [34] Yuning You, Tianlong Chen, Yang Shen, and Zhangyang Wang. 2021. Graph Contrastive Learning Automated. In *ICML*. 12121–12132.
- [35] Yuning You, Tianlong Chen, Yongduo Sui, Ting Chen, Zhangyang Wang, and Yang Shen. 2020. Graph Contrastive Learning with Augmentations. In *NeurIPS*, Vol. 33. 5812–5823.
- [36] Weibin Zhang, Pulin Zhang, Yinghao Yu, Xiyang Li, Salvatore Antonio Biancardo, and Junyi Zhang. 2021. Missing Data Repairs for Traffic Flow with Self-Attention Generative Adversarial Imputation Net. *TITS* 23 (2021), 7919–7930.
- [37] Yang Zhang. 2013. Data Imputation for Regional Traffic Index Based on GSO-NN. *Advanced Materials Research*. 658 (2013), 587–591.
- [38] Chuanpan Zheng, Xiaoliang Fan, Cheng Wang, and Jianzhong Qi. 2020. GMAN: A Graph Multi-Attention Network for Traffic Prediction. In *AAAI*, Vol. 34. 1234–1241.
- [39] Chuanpan Zheng, Xiaoliang Fan, Cheng Wang, Jianzhong Qi, Chaochao Chen, and Longbiao Chen. 2023. INCREASE: Inductive Graph Representation Learning for Spatio-Temporal Kriging. In *WWW*. 673–683.
- [40] Yu Zheng, Xiuwen Yi, Ming Li, Ruiyuan Li, Zhangqing Shan, Eric Chang, and Tianrui Li. 2015. Forecasting Fine-grained Air Quality Based on Big Data. In *KDD*. 2267–2276.
- [41] Tinghui Zhou, Hanhui Shan, Arindam Banerjee, and Guillermo Sapiro. 2012. Kernelized Probabilistic Matrix Factorization: Exploiting Graphs and Side Information. In *SDM*. 403–414.
- [42] Yanqiao Zhu, Yichen Xu, Feng Yu, Qiang Liu, Shu Wu, and Liang Wang. 2021. Graph Contrastive Learning with Adaptive Augmentation. In *WWW*. 2069–2080.