

Spatially Constrained Transformer with Efficient Global Relation Modelling for Spatio-Temporal Prediction

Ashutosh Sao^{a,*} and Simon Gottschalk^a

^aL3S Research Center, Leibniz University Hannover


Abstract. Accurate spatio-temporal prediction is crucial for the sustainable development of smart cities. However, current approaches often struggle to capture important spatio-temporal relationships, particularly overlooking global relations among distant city regions. Most existing techniques predominantly rely on Convolutional Neural Networks (CNNs) to capture global relations. However, CNNs exhibit neighbourhood bias, making them insufficient for capturing distant relations. To address this limitation, we propose ST-SAMPLENET, a novel transformer-based architecture that combines CNNs with self-attention mechanisms to capture both local and global relations effectively. Moreover, as the number of regions increases, the quadratic complexity of self-attention becomes a challenge. To tackle this issue, we introduce a lightweight region sampling strategy that prunes non-essential regions and enhances the efficiency of our approach. Furthermore, we introduce a spatially constrained position embedding that incorporates spatial neighbourhood information into the self-attention mechanism, aiding in semantic interpretation and improving the performance of ST-SAMPLENET. Our experimental evaluation on three real-world datasets demonstrates the effectiveness of ST-SAMPLENET. Additionally, our efficient variant achieves a 40% reduction in computational costs with only a marginal compromise in performance, approximately 1%.

1 Introduction

Spatio-temporal prediction aims to predict future spatio-temporal dynamics based on historical data. This prediction task is commonly categorised into two domains: road network-level predictions [26, 14, 6], often modelled as graph-based problems, and region-level predictions [39, 18], typically approached as grid-based problems. This paper addresses region-level prediction, emphasising its significance for city planners, administrators, and ride-hailing companies. Focusing on this level is crucial for informed decision-making in smart city planning and development [9]. Accurate region-level prediction enables one to view city-wide traffic dynamics, which helps in the sustainable development of cities to provide better access to transport and public services.

For accurate spatio-temporal prediction, three essential spatio-temporal relations need to be addressed:

- *Local Spatial Relations* between closed neighbouring regions. For example, congestion in one region often affects adjacent regions due to the interconnectedness of road networks.



(a) Hannover Traffic 8 – 9 am. (b) Corresponding Attention Map.

Figure 1. Hannover traffic between 8 – 9 am of a working day and the corresponding attention map from ST-SAMPLENET. The sparse attention demonstrates the dependency on only a few regions for prediction.

- *Global Spatial Relations* between distant regions. For instance, distant residential areas often share similar traffic patterns, such as peak traffic during morning and evening rush hours.
- *Temporal Relations* between input time intervals. For example, traffic decreases in office areas in late hours as people return home.

Existing approaches predominantly focus on capturing local spatial relations by utilising CNN-based networks [38, 7, 12] and temporal relations through LSTMs/GRUs [33, 18] or Transformers [39, 2], and often overlook global relations by solely relying on CNNs to capture them. CNNs inherently exhibit a neighbour bias, meaning they are better at capturing local patterns and relationships within a spatial context. However, their ability to capture global relations is limited due to their localised receptive fields and parameter sharing, which prioritise nearby information over distant contexts. To overcome these limitations, we propose ST-SAMPLENET, which, besides capturing local and temporal relations effectively, also captures global relations by employing the self-attention mechanism [30].

However, as the number of regions increases either by expanding the observation area (larger city or state) or increasing the granularity (smaller region size), the quadratic complexity of self-attention becomes the major bottleneck for training and inference.

We observe that different regions hold varying importance at different time intervals, making it unnecessary to include all regions in each prediction. For example, if we know the traffic conditions on a particular highway, we can infer the traffic conditions for other regions on the same highway, thus eliminating the need to include them all. Figure 1a illustrates the traffic in the German city Hannover between 8 – 9 am on a weekday. High traffic volumes towards the east,

* Corresponding Author. Email: sao@L3S.de

north-west, and south correspond to highways, while moderate to low traffic in central areas correspond to residential, commercial, and recreational areas. Figure 1b shows the corresponding attention map generated by ST-SAMPLENET. It can be seen that ST-SAMPLENET requires only a limited amount of regions of varying traffic levels (from low to high) to predict the city-wide traffic conditions indicating that other regions are redundant, i.e., their traffic can be inferred using the information from the important regions.

Therefore, we propose a lightweight region sampling strategy that prunes non-essential regions at different intervals and addresses the quadratic complexity inherent in self-attention mechanisms. For sampling, we utilised Gumbel-Softmax [16], which has a two-fold advantage. Firstly, it is differentiable, thus making end-to-end training possible. Secondly, it introduces stochasticity to the sampling process, occasionally selecting non-important regions, thereby enhancing the model's generalisation capability through exploratory behaviour.

Additionally, we introduce a spatially constrained and learnable position embedding (SCPE) that integrates neighbourhood information in a hierarchical manner, such as locality, city, and state levels and enhances semantic interpretability. This approach constrains the self-attention mechanism to prioritise immediate neighbours, while simultaneously the learnable structure of the embedding also provides the flexibility to capture relations among distant neighbours. Consequently, SCPE not only comprehends spatial dependencies across different distances but also offers semantic interpretability.

Thus, overall, our contributions are as follows:

- We propose ST-SAMPLENET¹, a novel spatio-temporal prediction model that effectively captures crucial spatio-temporal relations, including global relations, often neglected in prior approaches.
- We propose a lightweight region sampling strategy that prunes the non-important regions and makes ST-SAMPLENET efficient.
- We propose a learnable and interpretable position embedding that embeds spatial neighbourhood information.
- Experiments on three real-world datasets prove the efficiency and effectiveness of ST-SAMPLENET.

2 Preliminaries

In this section, we formally define the problem of spatio-temporal prediction and the external features used in ST-SAMPLENET.

2.1 Problem Statement

Region Following [37], we divide a city into $H \times W$ uniform grids, each representing a distinct region denoted as r^n , $n \in [1, H \times W]$.

Spatio-Temporal Image During each time interval t , we collect spatio-temporal measurements [37] (e.g., inflow/outflow) for individual regions. The measurements for each region r^n are combined to create a single-channel image, and multiple such images are stacked together to construct a spatio-temporal image (referred to as an image) [40], denoted as $\mathcal{X}_t \in \mathbb{R}^{M \times H \times W}$. Here, M signifies the total number of distinct measurements.

Spatio-Temporal Prediction The objective is to forecast the image at the subsequent time interval \mathcal{X}_{t+1} given a sequence of historical images $\langle \mathcal{X}_t, \mathcal{X}_{t-1}, \dots \rangle$.

2.2 External Features

Spatio-temporal measurements depend on the region's semantic attributes (e.g., residential or commercial nature) and temporal factors. For instance, peak traffic in residential and office areas can be seen during morning and evening rush hours. Hence, we incorporate these two factors as external features in our model.

Semantic Features To capture the semantic aspects of a region, we utilise Points of Interest (POIs). For each region r^n , the count of POIs in specific categories p (e.g. count of residential buildings) is quantified as follows:

$$x^{p,r^n} = |\{q \mid q \in \mathcal{Q}^p \text{ s.t. } g_q \in r^n\}|, \quad (1)$$

where $g_q \in r^n$ indicates that the POI with location g_q lies within region r^n and \mathcal{Q}^p is the set of all POIs of category p in the city. Aggregated POI counts form single-channel images for each category, stacked to create the multi-dimensional POI image $\mathcal{X}^{poi} \in \mathbb{R}^{P \times H \times W}$, with P denoting the total number of distinct categories.

Temporal Features Regarding temporal information, we employ one-hot encoding to capture day-of-the-week information x_t^{dow} and weekend status x_t^{we} . Due to the cyclic nature of the time of day (tod) and to reduce the length of the feature vector, we apply sinusoidal encoding to encode them:

$$x_t^{tod} = [\sin\left(\frac{2\pi t'}{1440}\right), \cos\left(\frac{2\pi t'}{1440}\right)], \quad (2)$$

where t' represents the time elapsed in minutes from midnight until the start of time interval t . The three-time encodings are concatenated to form the feature vector $\mathcal{X}_t^{time} = [x_t^{dow}; x_t^{we}; x_t^{tod}] \in \mathbb{R}^{10}$, encompassing comprehensive temporal features in our model.

3 Approach

In this section, we present the ST-SAMPLENET approach. We start by describing the architecture and explaining each of its components. Next, we provide a detailed explanation of our spatially constrained position embedding. Following this, we describe our region sampling strategy in detail. Finally, we outline the training procedure.

3.1 Architecture

Figure 2 illustrates the architecture of ST-SAMPLENET. It consists of three primary components: the *Spatial Encoder*, the *Temporal Encoder*, and the *Predictor*. For a given timestamp $t - i$, the *Spatial Encoder* takes three inputs: the spatio-temporal image \mathcal{X}_{t-i} , the semantic features of the city \mathcal{X}^{poi} , and the temporal features \mathcal{X}_{t-i}^{time} . It processes these inputs to generate the spatial representation \mathcal{Z}_{t-i}^S .

Subsequently, the *Temporal Encoder* operates on the spatial representations of input intervals $\langle \mathcal{Z}_t^S, \mathcal{Z}_{t-1}^S, \dots, \mathcal{Z}_{t-n}^S \rangle$, producing the final representation \mathcal{Z}_{t+1}^{ST} of the subsequent time interval $t + 1$. This final representation is then utilised by the *Predictor* to make predictions ($\hat{\mathcal{X}}_{t+1}$). Further details regarding each component are provided below.

3.1.1 Spatial Encoder

The Spatial Encoder is responsible for capturing the spatial dependencies between regions at any time interval $t - i$. It has four sub-components namely *Local Feature Encoder (LFE)*, *Semantic Feature Encoder (SFE)*, *Temporal Feature Encoder (TFE)* and *Global Feature Encoder (GFE)*.

¹ <https://github.com/ashusao/ST-SampleNet>

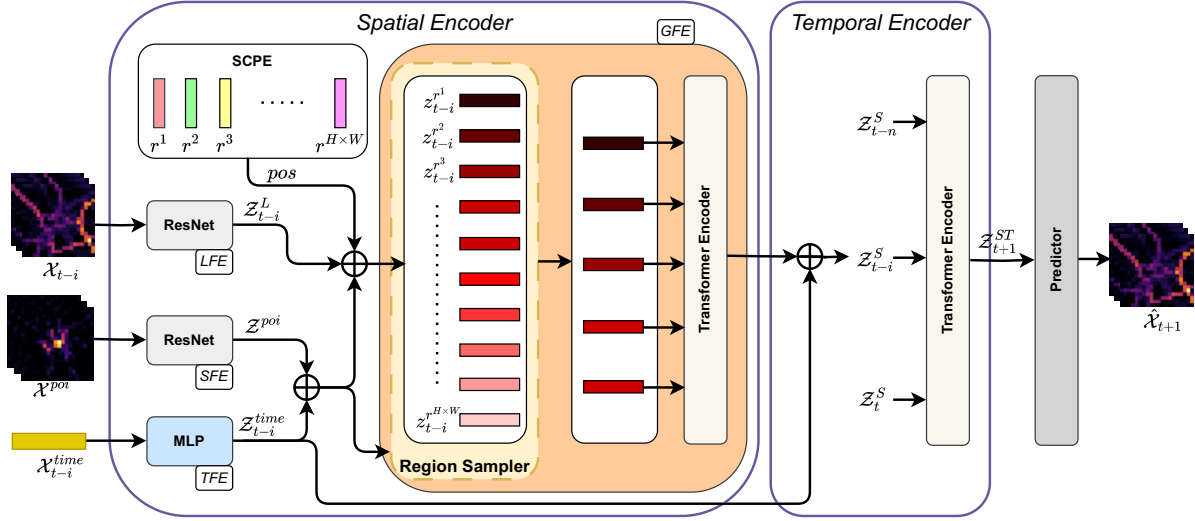


Figure 2. ST-SAMPLENET architecture has three main components: (i) *Spatial Encoder* – learns the spatial dependency amongst different regions, (ii) *Temporal Encoder* – learns the temporal dependency amongst input time intervals and (iii) *Predictor* – makes the final predictions.

Local Feature Encoder (LFE) – focuses on capturing local spatial dependencies. It utilises a RESNET [10] to achieve this. Starting with the input image \mathcal{X}_{t-i} , it begins with an initial convolutional layer, increasing feature dimensions to d . This is followed by m residual blocks designed to learn local neighbourhood relations. Each block consists of two convolutional layers and a skip connection. Finally, a 1×1 convolution merges feature maps, yielding the representation \mathcal{Z}_{t-i}^L of the image:

$$\mathcal{Z}_{t-i}^L = \text{RESNET}(\mathcal{X}_{t-i}) \in \mathbb{R}^{N \times d}, \quad (3)$$

where $N = H \times W$. Batch Normalisation [15], and GELU [13] activation are applied after each convolution operation.

Semantic Feature Encoder (SFE) – Regions with similar semantic attributes exhibit spatial correlations similar to traffic features often occurring in proximity. For instance, recreational amenities like bars or pubs are often clustered in city centres for convenience. Therefore, another RESNET model, similar to *LFE*, is employed to capture these semantic relationships and generate semantic representations:

$$\mathcal{Z}^{poi} = \text{RESNET}(\mathcal{X}^{poi}) \in \mathbb{R}^{N \times d}, \quad (4)$$

where $N = H \times W$.

The residual blocks m were set to 3 for both RESNETs.

Temporal Feature Encoder (TFE) – generates the representation for temporal features utilising a multi-layer perceptron (MLP) containing two linear layers with GELU activation:

$$\mathcal{Z}_{t-i}^{time} = \text{MLP}(\mathcal{X}_{t-i}^{time}) \in \mathbb{R}^{N \times d}, \quad (5)$$

where $N = H \times W$.

Global Feature Encoder (GFE) – aims to capture the global dependency amongst distant regions utilising the Transformer’s Encoder module. We first fuse all three features (\mathcal{Z}_{t-i}^L , \mathcal{Z}^{poi} and \mathcal{Z}_{t-i}^{time}) along with our spatially constrained position embedding (described in Section 3.2) that provides region position information through element-wise addition. The fused representation corresponding to each region r^n is indicated by $z_{t-i}^{r^n}$.

Since not all regions hold equal significance at different intervals, our region sampling module (elaborated in Section 3.3) then assigns importance weights to the fused representation of each region indicated by the colour darkness in the figure. Using these weights, only the top- k regions are selected and utilised by the self-attention mechanism of the Transformer Encoder [30] to learn the global relation amongst the region and generate the final spatial representation \mathcal{Z}_{t-i}^S for the image at time interval $t - i$:

$$\mathcal{Z}_{t-i}^S = \text{ENCODER}(z_{t-i}^{r^1}, \dots, z_{t-i}^{r^k}) \in \mathbb{R}^d. \quad (6)$$

3.1.2 Temporal Encoder

The Temporal Encoder captures the temporal dependency between images of different time intervals. Time is a crucial factor in determining the traffic trend; for instance, if it’s late evening, traffic generally decreases as people commute back from work and vice versa in the early morning. Therefore, we first fuse the temporal features \mathcal{Z}_{t-i}^{time} with the spatial representation \mathcal{Z}_{t-i}^S of images of input time intervals through element-wise addition and then learn the temporal dependencies between them using another Transformer Encoder:

$$\mathcal{Z}_{t+1}^{ST} = \text{ENCODER}(\mathcal{Z}_{t-i}^S, \dots, \mathcal{Z}_t^S) \in \mathbb{R}^d, \quad (7)$$

where \mathcal{Z}_{t+1}^{ST} is the spatio-temporal representation for the next time interval $t + 1$.

3.1.3 Predictor

The Predictor makes the final prediction utilising the spatio-temporal representation \mathcal{Z}_{t+1}^{ST} . It consists of a linear layer followed by $\tanh(\cdot)$ activation:

$$\hat{\mathcal{X}}_{t+1} = \tanh(\mathcal{Z}_{t+1}^{ST} \cdot \mathbf{W} + \mathbf{B}), \quad (8)$$

where \mathbf{W} and \mathbf{B} are the predictor’s learnable parameters.

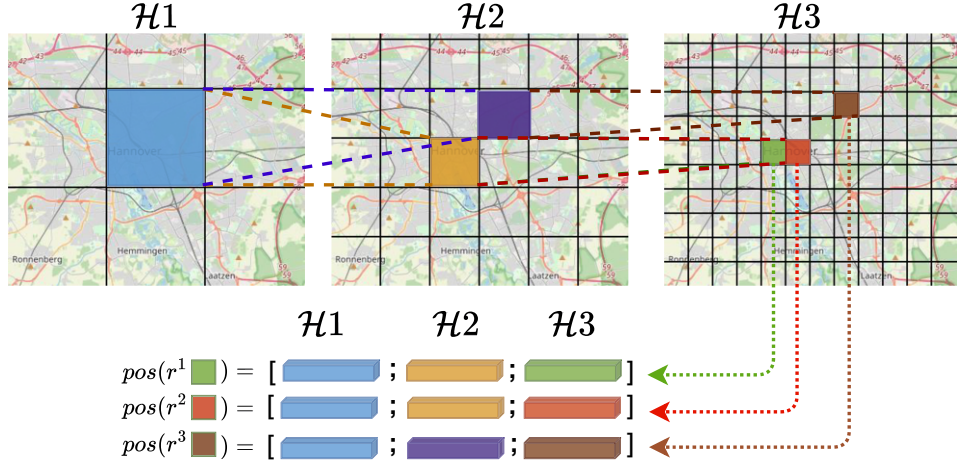


Figure 3. The city is divided into multiple levels of granularities ($\mathcal{H}_1, \mathcal{H}_2, \mathcal{H}_3$). Embeddings of different levels are concatenated to generate the position embedding (SCPE) of a region r^n . Map data: © OpenStreetMap contributors, ODbL.

3.2 Spatially Constrained Position Embedding

We employ a hierarchical method for embedding the position information of regions, as illustrated in Figure 3. To encode the position of a region r^n , we embed its position at multiple levels ($\mathcal{H}_1, \mathcal{H}_2, \mathcal{H}_3, \dots$). The lowest level captures the position at the highest granularity (e.g., $100\text{m} \times 100\text{m}$), while the highest level captures the position at the lowest granularity (e.g., $5\text{km} \times 5\text{km}$). The embeddings at different levels are concatenated to form the position embedding of the region r^n :

$$\text{pos}(r^n) = [\mathcal{H}_1(r^n); \mathcal{H}_2(r^n); \dots, \mathcal{H}_\ell(r^n)] \in \mathbb{R}^d, \quad (9)$$

where $\mathcal{H}_\ell(\cdot)$ represents the learnable embedding dictionary at level ℓ . We set ℓ to 3 in our model.

Semantically, SCPE captures hierarchical position information, including locality, city, state, country, etc. Crucially, it informs the self-attention mechanism about each region's relative position, indicating proximity or distance. For example, in Figure 3, regions r^1 and r^2 are immediate neighbours at two levels (\mathcal{H}_1 and \mathcal{H}_2), resulting in similar embedding vectors at those levels (depicted in blue and yellow). In contrast, r^3 , a distant region, exhibits similarity at only one level \mathcal{H}_1 (depicted in blue). Consequently, this structure encourages the self-attention mechanism to prioritise attention to immediately adjacent neighbours while at the same time, its learnable design also allows it to learn the dependencies amongst distant neighbours. Thus, our SCPE not only captures spatial dependencies at varying distances but also enhances semantic interpretation.

Our proposed position embedding has important application in diverse spatio-temporal tasks, specifically, when dealing with multiple geospatial points spanning various distances, from close (e.g., 10m) to very distant (e.g., 100km) such as trajectory representation learning [17, 8], trajectory similarity comparison [31, 20] etc.

3.3 Region Sampling

The importance of a region depends on its semantic attributes and the time of day. For instance, commercial regions are crowded during working hours and deserted during non-working hours. Therefore, we utilise semantic and temporal representations to predict an importance score for each region. Specifically, we first fuse \mathcal{Z}^{poi}

and \mathcal{Z}_{t-i}^{time} and project them using a Multi-Layer Perceptron (MLP) consisting of a linear layer and GELU activation:

$$\mathcal{Z}^{sampler} = \text{MLP}(\mathcal{Z}^{poi} + \mathcal{Z}_{t-i}^{time}) \in \mathbb{R}^{N \times d}, \quad (10)$$

where $N = H \cdot W$.

To assign importance scores and prune regions, a naive approach involves mixing the representations of all regions using an MLP to generate a score for each region. This method first flattens the features, resulting in a feature vector of length $N \cdot d$. Then, a single-layer MLP is applied with a weight matrix of shape $N \times N \cdot d = \mathcal{O}(N^2)$, making the approach computationally inefficient.

To address this, we introduce a lightweight module that applies the MLP individually to each region. This method does not involve flattening, so the feature vector length remains d and the weight matrix shape is $N \times d = \mathcal{O}(N)$. However, this method misses the global information necessary to compute relative importance score. To incorporate global information into the MLP, we split the sampler's representation ($\mathcal{Z}^{sampler}$) of each region r^n into two parts: the first half is denoted as local features ($z^{r^n, local} \in \mathbb{R}^{d'}$), where $d' = \frac{d}{2}$. We then compute the mean across all regions from the second half to obtain the global feature ($z^{global} \in \mathbb{R}^{d'}$). We concatenate the local and global features and apply the MLP to compute the probabilities to keep or drop a region:

$$\begin{aligned} z'^{r^n} &= \text{MLP}([z^{r^n, local}, z^{global}]) \\ \rho &= \text{SOFTMAX}(\mathcal{Z}') \in \mathbb{R}^{N \times 2}, \end{aligned} \quad (11)$$

where $\rho^{r^n, 0}$ and $\rho^{r^n, 1}$ indicate the probabilities to drop and keep a region r^n , respectively.

Sampling directly from the distribution ρ presents a challenge for end-to-end training due to its non-differentiability. To overcome this, we employ the Gumbel-Softmax technique [16]. This method is applied to the probability distribution of the regions to be kept, allowing us to sample important regions while ensuring the differentiability of the operation. Furthermore, Gumbel-Softmax introduces randomness to the sampling process in a controlled manner, enabling the occasional selection of less important regions. This stochastic element enhances the model's robustness and its ability to generalize.

3.4 Training

To avoid training being dominated by large deviations, we utilise the mean squared error (MSE) together with the mean absolute error (MAE) as the loss function, which are defined as follows:

$$\mathcal{L}_{MSE} = \frac{1}{N} \sum_1^N (\mathcal{X}_{t+1} - \hat{\mathcal{X}}_{t+1})^2 \quad (12)$$

$$\mathcal{L}_{MAE} = \frac{1}{N} \sum_1^N |\mathcal{X}_{t+1} - \hat{\mathcal{X}}_{t+1}| \quad (13)$$

Furthermore, we use self-distillation to mitigate the information loss from region pruning at each time interval t . To achieve this, we align the behavior of the region-pruned student model with that of the teacher model, which contains all regions. This alignment is achieved by minimizing the KL divergence [19] between their spatial representations at each input time interval t :

$$\mathcal{L}_{KL} = \frac{1}{T} \sum_{t=1}^T \text{KL}(\mathcal{Z}_t^{S,teacher} \parallel \mathcal{Z}_t^{S,student}), \quad (14)$$

where T is the total number of input time intervals.

The overall training loss is the combination of the above three losses:

$$\mathcal{L} = \mathcal{L}_{MSE} + \mathcal{L}_{MAE} + \alpha \cdot \mathcal{L}_{KL}, \quad (15)$$

where we set $\alpha = 0.3$ in all our experiments.

4 Evaluation Setup

In this section, we describe the datasets used in our experiments and the baselines we compare to, followed by an overview of our experimental settings.

4.1 Datasets

We utilised floating car data from two German cities, Hannover and Dresden, as well as publicly available taxi data² from New York City (NYC). Each city was divided into $500\text{m} \times 500\text{m}$ regions. For the floating car data, we computed three types of measurements ($M = 3$): *Inflow* and *Outflow* [37], representing the number of vehicles entering and exiting the region, respectively, and *Density* [18], measuring the number of vehicles within a region during a specific time interval. For NYC taxi data, we computed two types of measurements ($M = 2$): *Pickup* and *DropOff* [34], indicating the number of taxis picked up and dropped off in a region, respectively. Dataset details are provided in Table 1.

Additionally, we extracted 10 categories of POIs from OpenStreetMap [24] to capture semantic characteristics. The POI categories and their counts for the three cities are described in Table 2.

4.2 Baselines

We compare ST-SAMPLENET to the following baselines:

- **HISTORICAL AVERAGE (HA)** – predicts by computing the average value of historical measurements.
- **DEEPST** [37] – utilises CNNs to process the images corresponding to closeness, period, and trend, respectively. Subsequently, it employs element-wise addition to combine all the information.

- **ST-RESNET** [38] – updates the CNN backbone of DEEPST with that of the ResNet architecture.
- **HCONVLSTM** [36] – utilises ConvLSTM layers to jointly capture the spatio-temporal dependencies by only considering the closest historical time intervals to make the prediction.
- **DEEPSTN+** [7] – is an upgrade of ST-RESNET using a *ResPlus* unit that conducts two types of convolution: one to capture local neighbourhood relations with a smaller kernel size, and another to capture relations among remote regions with a larger kernel size.
- **BDSTN** [2] – is a transformer-based architecture that first applies a multi-layer perception to learn spatial dependencies and then employs a transformer encoder to learn temporal dependencies.
- **ST-GSP** [39] – enhances BDSTN by replacing the MLP with that of a ResNet architecture for learning spatial dependencies.
- **GRIDFORMER** [28] – adapts the Video Swin Transformer [22] for spatio-temporal prediction. It arranges the 3D Swin Transformer blocks in a pyramid structure to capture spatio-temporal dependencies and uses attention-based fusion to account for shifts in periodicity.
- **DEEPCROWD** [18] – stacks CONVLSTM layers in a pyramid architecture and fuses low and high-resolution feature maps to learn better feature representations.
- **ST-3DGMR** [11] – utilises a dilated 3D Convolution [29] along with the residual connection to jointly capture the spatio-temporal relations.
- **ST-3DMDDN** [12] – updates ST-3DGMR by breaking the 3D Convolution into 2D Convolution in spatial dimension and 1D Convolution in temporal dimension. This significantly lowers the number of parameters and increases the model's efficiency.

4.3 Experimental Settings

We designate December 2019 as the test set for Hannover and Dresden and December 2023 for New York. The remaining data of each city serves as the training set. To ensure model robustness, 20% of the training data is used for validation. ST-SAMPLENET's feature dimension d is set to 128. Following prior work [7, 38], the historical time interval's closeness, period, and trend are set to 4, 3, and 2, respectively to select historical spatio-temporal images.

All models are trained on NVIDIA RTX A6000 GPUs using the AdamW optimizer [23]. The learning rate is tuned within [0.005, 0.001, 0.0005, 0.0001]. Training is conducted for 500 epochs or until the validation error stagnates for 30 consecutive epochs. Models achieving the best performance on the respective validation set are selected and evaluated on the test set. The effectiveness of the models is measured using root mean squared error (RMSE) and mean absolute error (MAE), whereas efficiency is measured using Giga Floating-Point Operations Per Second (GFLOPS)³.

5 Result

In this section, we present the results of ST-SAMPLENET. We begin by comparing its performance with various baselines. Next, we analyze the effect of the region sampling strategy. We then examine the contributions of different components of ST-SAMPLENET, followed by an analysis of the learned spatial position embeddings. Finally, we demonstrate the region sampling behavior with a case study.

² <https://www.nyc.gov/site/tlc/about/tlc-trip-record-data.page>

Table 1. Dataset descriptions.

	Hannover	Dresden	NYC
Latitude	[52.3290, 52.4189]	[51.0026, 51.0749]	[40.7085, 40.8344]
Longitude	[9.6605, 9.8076]	[13.6593, 13.8606]	[-74.0214, -73.9146]
Grid size	500m \times 500m	500m \times 500m	500m \times 500m
#Regions	400 = 20 \times 20	448 = 16 \times 28	504 = 18 \times 28
Time span	Jul - Dec 2019	Jul - Dec 2019	Jul - Dec 2023
Time Interval	1 hour	1 hour	1 hour

Table 2. Considered POI categories and their counts.

Category	Count		
	Hannover	Dresden	NYC
Commercial	1,008	1,618	752
Culture	557	418	1,100
Education	306	339	655
Health	513	698	891
Public Service	486	572	823
Recreation	1,314	1,256	7,362
Residential	6,265	858	218
Sports	882	923	1,740
Tourism	869	1,178	1,687
Transport	2,028	2,960	1,518

5.1 Comparison with Baselines

The performance comparison of ST-SAMPLENET and baselines on three datasets is detailed in Table 3. The naive historical average (HA) performs the worst, emphasising the limitations of a simple heuristic-based approach in capturing the complex spatio-temporal dynamics of a city.

Among the deep learning baselines, HCONVLSTM performs the poorest due to its limited consideration of only the closest historical information, disregarding periodicity and trend information. CNN-based approaches (DEEPST, ST-RESNET, and DEEPSTN+) perform better than HCONVLSTM but exhibit inferior performance compared to DEEPCROWD, BDSTN, and ST-GSP, attributed to their negligence of capturing temporal dependencies.

Among the Transformer-based baselines, ST-GSP performs better than BDSTN because it better captures local spatial relationships using RESNET while the BDSTN employs an MLP. GRIDFORMER outperforms both ST-GSP and BDSTN by addressing the global spatial relationships ignored by the others. However, GRIDFORMER still falls short of ST-SAMPLENET because it does not adequately capture local spatial relationships, relying solely on linear projection of patches.

The recently introduced 3D convolution-based models ST-3DGMR and ST-3DMDDN exhibit the best performances among all baselines. However, the local receptive field of 3D convolutions limits their ability to capture distant dependencies, making them perform inferior compared to ST-SAMPLENET.

Overall, ST-SAMPLENET outperforms the best baseline (ST-3DMDDN) by an average of 6.84% w.r.t. RMSE and 6.75% w.r.t. MAE across all the datasets, highlighting the importance of careful modelling all the three types of spatio-temporal relations in the architecture.

5.2 Effect of Region Sampling

Figure 4 depicts the effect of sampling on the performance of ST-SAMPLENET. The values on the x-axis indicate the ratio of the

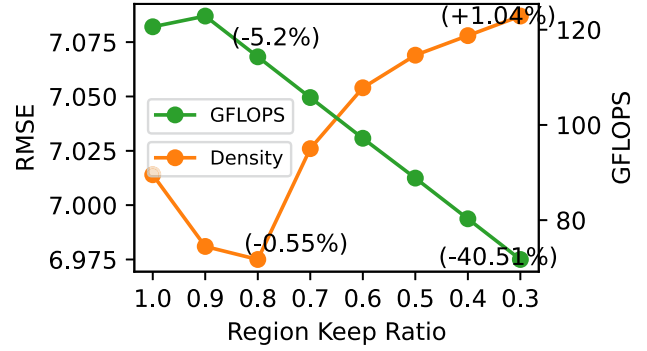


Figure 4. Effect of sampling on density prediction and corresponding computational cost for Hannover city. On x-axis is the ratio of region kept after sampling, on y-axis (left) is the RMSE for density prediction and on y-axis (right) is the corresponding GFLOPS.

regions kept after sampling. The values on the left and right axes indicate the RMSE in the density prediction of Hannover and GLOPS taken by the model.

It can be seen that the computational cost monotonically decreases as the region keep ratio decreases except at the keep ratio of 0.9. This is due to the addition of region sampling module parameters, which is not present on the base model when full attention is applied at a keep ratio of 1.0.

Regarding performance, it can be seen that at first, the model gets better by approximately 0.55% and then the performance degrades. This is because until a certain threshold (0.8 in this case), redundant and non-essential regions are pruned, thus, there is no loss of information. Additionally, distillation from the base model further regularises the model and makes it robust. Afterwards, the loss of information becomes significant, and performance degrades. To investigate this further, we also perform a case study which is described in Section 5.5.

Overall, comparing effectiveness with efficiency, ST-SAMPLENET can reduce the computational cost by approximately 40% at the marginal drop in performance of around 1%.

5.3 Contribution of Different Components

In this study, we aim to investigate the contribution ST-SAMPLENET's different components for which we evaluated the following variants of ST-SAMPLENET on the Hannover dataset:

- **w/o SCPE** – We replace our SCPE with a learnable position embedding similar to BERT [3] (details in Section 5.4).
- **w/o LFE** – In this variant, we replace RESNET of LFE with that of a linear layer similar to BERT or ViT [5].
- **w/o SFE** – We remove the RESNET of SFE, and semantic features are not fed to the model.

³ We use the `calflops` [35] module to compute GFLOPS of our model.

Table 3. ST-SAMPLENET’s performance compared to the baselines in Hannover, Dresden and New York City. IMPROV. (%) indicates the percentage improvement of ST-SAMPLENET over the best-performing baseline (underlined).

Approach	Hannover						Dresden						NYC			
	Density		Inflow		Outflow		Density		Inflow		Outflow		Pickup		Dropoff	
	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE
HA	20.22	11.69	19.61	11.12	19.66	11.11	20.39	10.44	17.92	9.78	18.03	9.78	21.04	4.60	19.07	4.49
HCONVLSTM	9.27	5.54	9.28	5.50	9.19	5.43	9.30	5.29	9.00	5.16	8.88	5.10	8.36	1.67	6.92	1.49
DEEPST	9.13	4.95	9.23	4.99	9.15	4.91	9.05	5.23	8.71	5.08	8.65	5.03	8.20	1.63	6.85	1.47
ST-RESNET	8.49	4.86	8.45	4.84	8.59	4.87	8.61	4.68	7.84	4.38	8.08	4.47	7.87	1.57	6.77	1.44
DEEPSTN+	8.32	4.46	8.18	4.45	8.17	4.42	8.18	4.13	7.70	4.05	7.68	4.01	7.76	1.45	6.66	1.34
DEEPCROWD	8.29	4.74	8.23	4.70	8.21	4.68	8.04	4.51	7.67	4.39	7.61	4.38	7.66	1.52	6.57	1.32
BDSTN	8.08	4.54	8.05	4.50	8.01	4.46	8.11	4.10	7.67	4.05	7.70	4.04	7.50	1.42	6.25	1.28
ST-GSP	7.90	4.46	7.93	4.43	7.90	4.41	7.93	4.01	7.66	4.01	7.69	4.00	7.24	1.40	6.07	1.29
GRIDFORMER	7.87	4.39	7.88	4.36	7.84	4.34	7.89	4.36	7.62	4.27	7.53	4.25	7.21	2.23	5.97	1.91
ST-3DGMR	7.76	4.40	7.81	4.39	7.78	4.37	7.93	4.02	7.63	3.90	7.64	3.97	7.37	1.41	5.91	1.26
ST-3DMDDN	7.62	4.31	7.72	4.32	7.67	4.29	7.75	3.91	7.59	3.88	7.60	3.94	7.03	1.37	5.84	1.25
ST-SAMPLENET	7.01	4.00	7.11	4.03	7.10	4.01	7.22	3.66	7.08	3.63	7.06	3.62	6.64	1.28	5.48	1.17
IMPROV. (%)	8.01	7.19	7.90	6.71	7.43	6.52	6.84	6.39	6.72	6.44	7.10	8.12	5.54	6.57	6.16	6.4

- **w/o TFE** – Temporal features are not fed to the model in this variant.
- **w/o GFE** – We remove the transformer encoder of *GFE* responsible for capturing global spatial dependency.
- **w/o TEMP. ENC.** – We remove the transformer encoder responsible for capturing temporal dependency, and the mean spatial representation of input intervals is fed to the predictor.

Table 4. Contribution of different components of ST-SAMPLENET on the Hannover dataset.


	Volume		Inflow		Outflow	
	RMSE	MAE	RMSE	MAE	RMSE	MAE
ST-SAMPLENET	7.01	4.00	7.11	4.03	7.10	4.01
w/o SCPE	7.17	4.07	7.27	4.09	7.25	4.07
w/o LFE	7.28	4.12	7.36	4.13	7.33	4.11
w/o SFE	7.16	4.01	7.29	4.06	7.24	4.03
w/o TFE	7.12	4.05	7.20	4.07	7.18	4.05
w/o GFE	7.22	4.10	7.34	4.13	7.32	4.11
w/o TEMP. ENC.	7.28	4.12	7.38	4.15	7.38	4.13

The results in Table 4 demonstrate that each component contributes to the effectiveness of ST-SAMPLENET. Specifically, the contributions of *LFE*, *GFE*, and *Temporal Encoder* are the most important, emphasising the significance of capturing the local, global and temporal relation in spatio-temporal modelling.

5.4 Learned Spatial Position Embedding

To further investigate the impact of Spatially Constrained Position Embedding (SCPE), we compare it with a learnable position embedding akin to BERT [3] as described in Section 5.3 (w/o SCPE). Results detailed in Table 4 demonstrate that passing neighbourhood information and constraining the model to pay more attention to closed neighbours ST-SAMPLENET’s effectiveness.

Furthermore, Figure 5 depicts the correlation of the learned position embedding, where Figure 5a depicts how each region in the Hannover city correlates to another. The small square structures along the diagonal indicate that each region primarily attends to its immediate neighbours. Given that the regions are arranged row by row in this



(a) Each row shows the correlation between the spatial position embedding of a specific region and all other regions⁴.

(b) The heat map shows the correlation between the spatial position embeddings of various regions and that of a specific bordered region.

Figure 5. Visualisation of learned position embeddings showing the correlation of spatial positions between different regions in Hannover.

figure, the spaces before and after these squares represent the neighbouring regions of the rows above and below, respectively. To better understand this, we also visualise the correlation of a specific region, illustrated in Figure 5b. This figure demonstrates how one region is related to the others in the city. It demonstrates the hierarchical nature where regions close to each other are more related than others.

Thus, our SCPE, offering both semantic interpretability and improved predictive capability, proves beneficial for ST-SAMPLENET.

5.5 Case Study: Sampling Behaviour

To understand the sampling behaviour of our ST-SAMPLENET, we performed a case study as illustrated in Figure 6. Figure 6a portrays the traffic patterns in Hannover between 8 and 9 am on a typical working day. Meanwhile, Figure 6b and Figure 6c exhibit the regions pruned to generate the representation of the same time interval at a keep ratio of 0.8 and 0.5, respectively.

Looking at the regions pruned at a ratio of 0.8, it can be observed that ST-SAMPLENET prunes two types of regions. The first category comprises regions with consistently negligible traffic, such as forested areas or regions containing lakes. The second category includes regions characterised by deterministic traffic patterns, exemplified by highways in the northeast of the city. Since these regions often contain redundant information, with traffic patterns mirroring

⁴ For better visualisation, we only show the correlation between 100 regions in the north-western corner of the city.



Figure 6. Case Study: Hannover traffic at 8-9 am on weekdays and the pruned regions (in white) at different keep ratios.

those in other parts of the city, ST-SAMPLENET considers them non-essential and prunes them. For instance, the traffic on pruned highways in the northeast can be accurately predicted based on the traffic of interconnected highways in the southeast and east.

However, at the keep ratio of 0.5, ST-SAMPLENET tends to over-prune, neglecting important non-redundant regions. Notably, regions towards the city centre, where traffic exhibits significant variability due to factors such as events, weather, etc., are overlooked. Consequently, the model’s performance degrades, as discussed in Section 5.2 and illustrated in Figure 4 earlier.

6 Related Work

Spatio-temporal prediction models have gained significant attention in the realm of smart city applications, leading to the development of numerous spatio-temporal models in recent years.

Early endeavours in this domain primarily focused on capturing local spatial relations. They utilised fully convolutional architectures where spatio-temporal images of input time intervals are stacked as channels and convolutions are used to capture both spatial and temporal relations, such as by DEEPST [37], ST-RESNET [38], and DEEPSTN+ [7]. However, these approaches tend to be ineffective as they ignore the modelling of temporal and global spatial relations inherent in spatio-temporal data, relying solely on CNNs to capture them.

Consequently, subsequent works proposed a two-component approach to spatio-temporal modelling, separating spatial and temporal aspects. Notable examples include DMVST-NET [32] and STDN [33], which use local CNNs to capture local spatial relations and RNNs (LSTMs and GRUs) to capture temporal relations. However, these methods ignore the relationships between distant regions, making them less effective.

The advent of the CONVLSTM [27] and CONVGRU [1] architectures further advanced spatio-temporal prediction techniques, as seen in works like PERIODIC-CRN [40] and DEEPCROWD [18]. Despite their advancements, the local receptive field of convolution in these architectures limits their ability to capture global spatial relations. To address this, PERIODIC-CRN and DEEPCROWD stack multiple layers in a pyramid structure to capture distant relations. However, in such architectures, information from distant regions must pass through multiple layers and recurrent steps, which can cause the information to degrade or get diluted.

As Transformers [30] gained popularity in Natural Language Processing [3, 4] and Computer Vision [5, 21, 22] tasks, they have also been employed in the spatio-temporal domain. BDSTN [2] uses an MLP to capture spatial relations, followed by a Transformer en-

coder to capture temporal relations. ST-GSP [39] improves upon BDSTN by replacing the MLP with a RESNET module. However, both models overlook global spatial dependencies. GRIDFORMER [28] adapts the Video Swin Transformer [22] for spatio-temporal prediction by arranging 3D Swin Transformer blocks in a pyramid structure. GRIDFORMER is similar to our approach in that it uses a Transformer’s encoder to capture global relations, but it differs significantly in other aspects. Firstly, it uses patches of multiple regions as tokens, whereas ST-SAMPLENET uses individual regions as tokens, making our model more interpretable. Secondly, GRIDFORMER divides the patches in three dimensions and then applies linear transformation, which does not adequately capture local spatial and temporal relations whereas ST-SAMPLENET specifically employs a RESNET and Transformer’s encoder to capture these relations.

More recent approaches draw inspiration from video representation learning [25], treating spatio-temporal data as sequences of frames and applying 3D convolution to learn spatio-temporal representations jointly. Models such as ST-3DGMR [11] and ST-3DMDN [12] exemplify this paradigm. However, 3D convolution-based approaches operate over local regions, giving them a limited receptive field, which makes them inappropriate for capturing distant relationships.

Unlike previous approaches that struggle to incorporate all three essential spatio-temporal relationships needed for accurate predictions, ST-SAMPLENET effectively addresses all of them. Additionally, our region sampling strategy reduces the quadratic complexity associated with the self-attention mechanism, enhancing efficiency.

7 Conclusion

In conclusion, we propose ST-SAMPLENET, a novel architecture for spatio-temporal modelling, effectively capturing essential spatio-temporal relations, including global spatial dependency often neglected in prior approaches. We propose a spatial position embedding which is semantically interpretable and enhances model performance. Additionally, our region sampling strategy in ST-SAMPLENET enhances efficiency by significantly reducing computation costs while maintaining the performance.

Acknowledgements

This work was partially funded by the Federal Ministry for Economic Affairs and Climate Action (BMWK), Germany (“ATTENTION!”, 01MJ22012D), and by the Federal Ministry for Digital and Transport (BMDV), Germany (“MoToRes”, 19F2271A).

References

- [1] N. Ballas, L. Yao, C. Pal, and A. C. Courville. Delving deeper into convolutional networks for learning video representations. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016.
- [2] D. Cao, K. Zeng, J. Wang, P. K. Sharma, X. Ma, Y. Liu, and S. Zhou. Bert-based deep spatial-temporal network for taxi demand prediction. *IEEE Transactions on Intelligent Transportation Systems*, 2022.
- [3] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 2019.
- [4] L. Dong, N. Yang, W. Wang, F. Wei, X. Liu, Y. Gao, M. Zhou, and H.-W. Hon. Unified language model pre-training for natural language understanding and generation. *Advances in neural information processing systems*, 2019.
- [5] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. D ehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021.
- [6] S. Elmi. Deep stacked residual neural network and bidirectional lstm for speed prediction on real-life traffic data. In *Proceedings of ECAI*, 2020.
- [7] J. Feng, Y. Li, Z. Lin, C. Rong, F. Sun, D. Guo, and D. Jin. Context-aware spatial-temporal neural network for citywide crowd flow prediction via modeling long-range spatial dependency. *ACM Trans. Knowl. Discov. Data*, 2021.
- [8] T.-Y. Fu and W.-C. Lee. Trembr: Exploring road networks for trajectory representation learning. *ACM Trans. Intell. Syst. Technol.*, 2020.
- [9] A. A. Hameed. Smart city planning and sustainable development. *IOP Conference Series: Materials Science and Engineering*, may 2019.
- [10] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [11] R. He, Y. Xiao, X. Lu, S. Zhang, and Y. Liu. St-3dgm: Spatio-temporal 3d grouped multiscale resnet network for region-based urban traffic flow prediction. *Information Sciences*, may 2023.
- [12] R. He, C. Zhang, Y. Xiao, X. Lu, S. Zhang, and Y. Liu. Deep spatio-temporal 3d dilated dense neural network for traffic flow prediction. *Expert Systems with Applications*, 2024.
- [13] D. Hendrycks and K. Gimpel. Bridging nonlinearities and stochastic regularizers with gaussian error linear units. *CoRR*, 2016.
- [14] R. Huang, C. Huang, Y. Liu, G. Dai, and W. Kong. Lsgcn: Long short-term traffic prediction with graph convolutional networks. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization, 2020.
- [15] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37*, 2015.
- [16] E. Jang, S. Gu, and B. Poole. Categorical reparameterization with gumbel-softmax. In *5th International Conference on Learning Representations, ICLR 2017*, 2017.
- [17] J. Jiang, D. Pan, H. Ren, X. Jiang, C. Li, and J. Wang. Self-supervised trajectory representation learning with temporal regularities and travel semantics. In *2023 IEEE 39th International Conference on Data Engineering (ICDE)*, 2023.
- [18] R. Jiang, Z. Cai, Z. Wang, C. Yang, Z. Fan, Q. Chen, K. Tsubouchi, X. Song, and R. Shibasaki. Deepcrowd: A deep model for large-scale citywide crowd density and flow prediction. *IEEE Transactions on Knowledge and Data Engineering*, 2023.
- [19] T. Kim, J. Oh, N. Y. Kim, S. Cho, and S.-Y. Yun. Comparing kullback-leibler divergence and mean squared error loss in knowledge distillation. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, 2021.
- [20] X. Li, K. Zhao, G. Cong, C. S. Jensen, and W. Wei. Deep representation learning for trajectory similarity computation. In *2018 IEEE 34th International Conference on Data Engineering (ICDE)*, 2018.
- [21] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, 2021.
- [22] Z. Liu, J. Ning, Y. Cao, Y. Wei, Z. Zhang, S. Lin, and H. Hu. Video swin transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022.
- [23] I. Loshchilov and F. Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019.
- [24] OpenStreetMap contributors. Planet dump retrieved from <https://planet.osm.org> . <https://www.openstreetmap.org>, 2023.
- [25] R. Qian, T. Meng, B. Gong, M.-H. Yang, H. Wang, S. Belongie, and Y. Cui. Spatiotemporal contrastive video representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [26] X. Rao, H. Wang, L. Zhang, J. Li, S. Shang, and P. Han. Fogs: First-order gradient supervision with learning-based graph for traffic flow forecasting. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization, 2022.
- [27] X. Shi, Z. Chen, H. Wang, D.-Y. Yeung, W.-k. Wong, and W.-c. Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*. MIT Press, 2015.
- [28] C. Su, C. Wu, and D. Lian. Gridformer: Spatial-temporal transformer network for citywide crowd flow prediction. In *ECAI 2023 - 26th European Conference on Artificial Intelligence, September 30 - October 4, 2023, Krak w, Poland - Including 12th Conference on Prestigious Applications of Intelligent Systems (PAIS 2023)*. IOS Press, 2023.
- [29] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3d convolutional networks. In *2015 IEEE International Conference on Computer Vision (ICCV)*, 2015.
- [30] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin. Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2017.
- [31] P. Yang, H. Wang, Y. Zhang, L. Qin, W. Zhang, and X. Lin. T3s: Effective representation learning for trajectory similarity computation. In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*, 2021.
- [32] H. Yao, F. Wu, J. Ke, X. Tang, Y. Jia, S. Lu, P. Gong, J. Ye, D. Chuxing, and Z. Li. Deep multi-view spatial-temporal network for taxi demand prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*. AAAI Press, 2018.
- [33] H. Yao, X. Tang, H. Wei, G. Zheng, and Z. Li. Revisiting spatial-temporal similarity: A deep learning framework for traffic prediction. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence and Thirty-First Innovative Applications of Artificial Intelligence Conference and Ninth AAAI Symposium on Educational Advances in Artificial Intelligence*. AAAI Press, 2019.
- [34] J. Ye, L. Sun, B. Du, Y. Fu, X. Tong, and H. Xiong. Co-prediction of multiple transportation demands based on deep spatio-temporal neural network. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2019.
- [35] X. Ye. callops: a FLOPs and params calculate tool for neural networks in PyTorch framework, 2023. URL <https://github.com/MrYxJ/calculate-flops.pytorch>.
- [36] Z. Yuan, X. Zhou, and T. Yang. Hetero-convlstm: A deep learning approach to traffic accident prediction on heterogeneous spatio-temporal data. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. Association for Computing Machinery, 2018.
- [37] J. Zhang, Y. Zheng, D. Qi, R. Li, and X. Yi. Dnn-based prediction model for spatio-temporal data. In *Proceedings of the 24th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. Association for Computing Machinery, 2016.
- [38] J. Zhang, Y. Zheng, and D. Qi. Deep spatio-temporal residual networks for citywide crowd flows prediction. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*. AAAI Press, 2017.
- [39] L. Zhao, M. Gao, and Z. Wang. St-gsp: Spatial-temporal global semantic representation learning for urban flow prediction. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining*. Association for Computing Machinery, 2022.
- [40] A. Zonozi, J.-J. Kim, X. Li, and G. Cong. Periodic-crn: A convolutional recurrent model for crowd density prediction with recurring periodic patterns. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*. AAAI Press, 2018.