



Traffic forecasting on new roads using spatial contrastive pre-training (SCPT)

Arian Prabowo^{1,2,3} · Hao Xue³ · Wei Shao² · Piotr Koniusz^{2,4} · Flora D. Salim³

Received: 29 November 2022 / Accepted: 4 September 2023 / Published online: 27 September 2023
© The Author(s) 2023

Abstract

New roads are being constructed all the time. However, the capabilities of previous deep forecasting models to generalize to new roads not seen in the training data (unseen roads) are rarely explored. In this paper, we introduce a novel setup called a spatio-temporal split to evaluate the models' capabilities to generalize to unseen roads. In this setup, the models are trained on data from a sample of roads, but tested on roads not seen in the training data. Moreover, we also present a novel framework called Spatial Contrastive Pre-Training (SCPT) where we introduce a spatial encoder module to extract latent features from unseen roads during inference time. This spatial encoder is pre-trained using contrastive learning. During inference, the spatial encoder only requires two days of traffic data on the new roads and does not require any re-training. We also show that the output from the spatial encoder can be used effectively to infer latent node embeddings on unseen roads during inference time. The SCPT framework also incorporates a new layer, named the spatially gated addition layer, to effectively combine the latent features from the output of the spatial encoder to existing backbones. Additionally, since there is limited data on the unseen roads, we argue that it is better to decouple traffic signals to trivial-to-capture periodic signals and difficult-to-capture Markovian signals, and for the spatial encoder to only learn the Markovian signals. Finally, we empirically evaluated SCPT using the ST split setup on four real-world datasets. The results showed that adding SCPT to a backbone consistently improves forecasting performance on unseen roads. More importantly, the improvements are greater when forecasting further into the future. The codes are available on GitHub: <https://github.com/cruiserese/archgroup/forecasting-on-new-roads>.

Keywords Cyber-physical systems · Intelligent transport systems · Spatio-temporal · Sensor networks

1 Introduction

Traffic forecasting is a critical component of intelligent transport systems. It enables the proactive management of traffic congestion and the efficient utilization of limited resources such as road space and public transportation. Road networks also changes through time as newly constructed roads are being added to the existing networks. Despite its importance, traffic forecasting on these new roads are not well explored (Roth and Liebig 2022).

The naive approach is to re-train the models with new data whenever new roads are constructed. This is not ideal for a two reasons. Firstly, re-training for an entire network is expensive. For example, training a model on one third of the entire California highway network took nearly 8 GPU-days (Mallick et al. 2020). Secondly, new roads, by their nature, have limited data, complicating the re-training process. Therefore, there is a need for methods to evaluate and extend the capabilities of a trained model to generalize to new roads not seen in the training data (unseen roads).

For evaluation, we introduced a novel setup to split data into train-validate-test sets to evaluate a model's capability to generalize to unseen roads. We refer to this setup as a spatio-temporal (ST) split because it incorporates both the temporal and spatial aspects of the data. In addition to reserving future time for validation and testing, we also allocate a portion of the roads (representing the spatial aspect) for evaluation. This allows us to assess the model's performance in terms of both time-based forecasting accuracy and its ability to generalize to new spatial locations within the network. This setup is data-driven, meaning it does not require contextual information about these new roads, e.g., road types, speed limits. Instead, ST splits only requires minimal traffic data on the new roads and does not require any re-training. The traffic data required is very short, two days as opposed to few months that is typical to the traffic forecasting tasks (Jiang et al. 2021).

For methods aimed at extending the capabilities of trained models to generalize to new, unseen scenarios, unsupervised pretraining such as contrastive learning has emerged as a promising approach. Unsupervised pretraining involves training an encoder on a large amount of unlabeled data to learn useful representations without relying on explicit labels. One popular technique within unsupervised pretraining is contrastive learning, which encourages the encoder to capture meaningful features by contrasting positive and negative samples.

In domains such as natural language processing (Brown et al. 2020), audio (Oord et al. 2018), and images (Chen et al. 2020a, b) the adoption of unsupervised pretraining has led to the development of powerful encoders capable of achieving generalized performance on diverse downstream tasks. Surprisingly, the application of unsupervised frameworks to traffic forecasting remains limited (Liu et al. 2022a) (Li et al. 2022b) with zero exploration of leveraging contrastive learning techniques to enhance model generalization to previously unseen roads.

In this paper we proposed a novel framework called **Spatial Contrastive Pre-Training (SCPT)** illustrated in Fig. 1. Here, we added a new pre-training stage

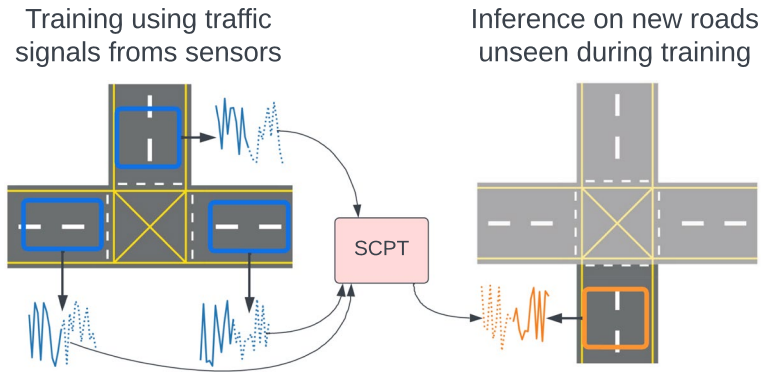


Fig. 1 Our novel traffic forecasting framework, Spatial Contrastive Pre-Training (SCPT), enables accurate forecasts on new roads (*orange*) that were not seen during training (Color figure online)

where the spatial encoder is exposed to the entire historical traffic signal of a road in the training set and then tasked to minimize the contrastive loss between latent representation of different roads (blue boxes). The models are only trained on the data from the roads in the training set (blue timeseries). They also take as input the output of the spatial encoders. During inference, forecasting is performed on traffic signals (orange timeseries) from new sensors (orange box) on roads which are not previously seen during training.

Besides the contrastive pre-training, we also introduce three techniques in the SCPT framework to increase its effectiveness. The first technique is a spatially gated addition (SGA) layer that uses a gating mechanism to integrate the output of the spatial encoders with the backbone forecasting models more effectively. Furthermore, it is important to consider the decomposition of traffic signals into two distinct components: trivial-to-capture periodic signals and difficult-to-capture Markovian signals. If only the raw traffic signal is used for contrastive learning, the model may primarily rely on the easily captured periodic signals, while disregarding the more challenging Markovian signals. This is why we have decoupled these two types of signals and employed a separate method to model the periodic signals, ensuring that the model effectively learns the intricate Markovian signals. Finally, many forecasting models use latent node embeddings that are learned during training (Wu et al. 2019; Bai et al. 2020), making it impossible to generalize these node embeddings on unseen roads during inference time. We show that the output of the spatial encoders can be used to effectively infer the node embeddings on unseen roads during inference time.

To empirically evaluate our proposed framework, we use the current state-of-the-art model, called Graph WaveNet (GWN) (Wu et al. 2019) as our forecasting model backbone. We then implemented the SCPT framework, using the ST split setup, on four real-world traffic datasets, including PeMS-11k, the largest publicly available dataset used for deep traffic forecasting. The results showed that the SCPT framework consistently improved the efficiency of the backbone.

The main contributions of this paper are:

- A new data splitting strategy called a **ST split**. This allows the evaluation of a framework's capability to perform traffic forecasting on unseen roads.
- A novel framework called **SCPT** that uses contrastive pre-training to allow forecasting models to generalize to unseen roads during inference time.
- Empirical evidences from extensive experiments on four real world datasets to gain insights to the effectiveness of SCPT and its components.

2 Related works

2.1 Statistical and machine learning approaches

The earliest work in traffic forecasting can be classified as data-driven and statistical approaches to machine learning, starting from the Box-Jenkins technique (Ahmed and Cook 1979). Others included the autoregressive integrated moving average (ARIMA) (Hamed et al. 1995) and ARIMA-like approaches such as KARIMA (Van Der Voort et al. 1996), subset ARIMA (Lee and Fambro 1999), ARIMAX (Williams 2001), VARMA (Kamarianakis and Prastacos 2003), and SARIMA (Williams and Hoel 2003). There are also classical machine learning methods such as the support vector regression (Jeong et al. 2013; Lippi et al. 2013; Chen et al. 2012).

2.2 Early deep learning models

In the context of traffic forecasting, the advances in deep learning mainly focused on coming up with more sophisticated architectures to improve only the models' accuracy, neglecting other research goals such as actionability and explainability (Manibardo et al. 2021).

The early architectures used to better capture the temporal dynamics included stacked autoencoders (Lv et al. 2014), gated recurrent units (GRU) (Fu et al. 2016), and long short-term memory (LSTM) (Shao 2020; Cui et al. 2020).

Starting from STGCN (Yu et al. 2018) and DCRNN (Li et al. 2018), the general architecture of choice to capture the traffic spatiotemporal dynamic is to alternate various spatial and temporal modules. STGCN (Yu et al. 2018) alternated between convolutional neural networks (CNN) and spectral graph convolution networks (GCN) (Defferrard et al. 2016). DCRNN (Li et al. 2018) used GRU and diffusion convolution.

Notice that similar the spatiotemporal analysis has also been applied to problems such as prediction of road networks (Prabowo et al. 2019), flight delay prediction (Shao et al. 2019, 2022), and energy use forecasting (Prabowo et al. 2023).

2.3 Current deep learning models

In the subsequent discussion, we will delve into the latest progress in deep learning models for traffic forecasting and their applicability to our present study.

GWN (Wu et al. 2019) used a CNN called WaveNet (Van Den Oord et al. 2016) and also diffusion convolution from DCRNN. Moreover, GWN argued that there are hidden spatial dependencies that are not captured by adjacency matrices constructed from physical road networks. Instead, it argued that the latent topological connectivity should be learned from data. This remained an open challenge until today (Wu et al. 2020; Shang et al. 2021; Li et al. 2022b). Our proposed framework, SCPT, extended this capability by being able to infer the latent topological connectivity of new roads which are unseen in the training data during inference time. GWN is also noteworthy because it remained to be the state-of-the-art according to the most recent benchmark study (Jiang et al. 2021). For this reason we chose GWN as the backbone forecasting model in this paper.

In our prior studies, G-SWaN (Prabowo et al. 2023; Prabowo 2022), we observed that each sensor exhibits unique dynamics. Building upon this insight, we enhanced GWN by introducing a novel graph neural network called spatial graph transformer, which effectively captures these unique dynamics. Motivated by this, our current work directly addresses the question of how to learn these individual dynamics for new roads. We also further investigated the limitations of graph attention mechanism in traffic forecasting and explored message-passing mechanism instead (Prabowo et al. 2023).

MTGNN (Wu et al. 2020) and GTS (Shang et al. 2021) extended the idea of learning the latent topological connectivity from the data to any multivariate timeseries, instead of only traffic-related timeseries. Conversely, because we are modelling traffic as multivariate timeseries in this paper, SCPT could easily be extended beyond traffic and spatiotemporal data, to any multivariate timeseries. An example of such use case is for electricity and gas retailers. It is important for them to be able to accurately forecast the usage of new clients with limited historical data and SCPT would be very well equipped to tackle this problem.

GWN, AGCRN (Bai et al. 2020), GTS and SPGCL (Li et al. 2022b) used latent road embeddings that are learned during training. Although these embeddings improved the model's forecasting accuracy, they prohibits inferences on new roads unseen in the training data because the latent embeddings for those roads have not been learned. However, this paper shows that under the SCPT framework we can effectively infer the latent road embeddings on unseen roads during inference time using the spatial encoder.

ASTGCN and D²STGNN (Shao et al. 2022) decoupled the traffic signals into multiple components and used specialized modules in their architecture to handle each signal component. ASTGCN decoupled them to recent, daily-periodic, and weekly-periodic components. Similarly, our SCPT framework decouples them to Markovian and periodic signals. We call it Markovian, instead of recent, to highlight the property of the behaviors that our model is trying to capture. Because the unseen roads setup is data scarce and the periodic signals is easier to model, we used discrete cosine transform (DCT) (Ahmed et al. 1974) to model the period signals and to let the deep learning models to focus on the Markovian signals.

Noteworthy is also few-shot learning (Snell et al. 2017; Zhang et al. 2020; Lin et al. 2018; Wang et al. 2021; Zhang et al. 2022; Zhu and Koniusz 2022; Zhang et al. 2022; Wang and Koniusz 2022a, b; Zhu and Koniusz 2023) which can adapt

to novel class concepts but it remains unclear how to apply few-shot learning on graphs.

2.4 Contrastive learning

There exist numerous approaches for the general-purpose graph contrastive learning (Wang and Isola 2020; Zhu et al. 2021; Zhu and Koniusz 2021; Zhang et al. 2022, 2023; Zhu and Koniusz 2022). However, to the best of our knowledge, there are only two works that implemented contrastive learning for traffic forecasting. Both still used performance accuracy as their primary objective instead of other downstream tasks. Liu et al. (2022a) argued for the need of contrastive loss due to data scarcity in traffic forecasting. Our unseen roads setup exacerbated this data scarcity problem and our results agree with them regarding the importance of contrastive learning.

SPGCL (Li et al. 2022b) used contrastive learning for neighbour connectivity selections, arguing that only roads that are similar in the contrastive embedding space should be topologically connected in the latent space. This is like our integration of SCPT with our backbone where we use the output of the contrastively pre-trained spatial encoder to construct the latent topological connectivity. The primary difference is that we allow this construction to be done on unseen roads during inference time.

FUNS-N (Roth and Liebig 2022) is the only prior work tackling unseen roads according to our knowledge although it was primarily tested on simulated dataset. In their paper, they called ‘unseen roads’ as ‘unobserved nodes’. Instead of motivating via newly constructed roads, they formulated their task as spatial imputation during spatio-temporal forecasting on a sensor network. Another important difference with SCPT was that they chose a context-driven method (e.g., speed limits and road types) instead of a data-driven method.

3 Methods: SCPT

In this paper, we first introduce a novel setup where forecasting is performed on new roads unseen in the training data (Fig. 1). Here, the model (including the spatial encoder) is exposed to only a sample of sensor data in a traffic network. Our main contribution is SCPT, a novel framework where a spatial encoder is pre-trained using contrastive loss, such that when it is attached to a forecasting model backbone, the latter generalizes well to new roads unseen in the training data.

This problem is formally defined in Sect. 3.1. The new setup requires a novel pre-training/training/validation/testing sets splitting strategy. We call this the spatio-temporal split and it is described in Sect. 3.2. The pre-training of spatial encoder using contrastive loss is described in Sect. 3.3. The architecture of the spatial encoder itself, as implemented in this paper, is described in Sect. 3.4. With that being said, note that our framework is agnostic to the architecture of the spatial encoder. The way the SCPT framework decouples traffic signal is described in Sect. 3.6 Next,

we introduce a new layer, called SGA, to effectively integrate the encoder output to existing traffic forecasting architectures. This is described in Sect. 3.5.

Finally, to demonstrate the capabilities of the spatial encoder trained via the spatial contrastive pre-training framework, we combine it with GWN (Wu et al. 2019) as the forecasting model backbone. The details of how to integrate SCPT frameworks to GWN and other backbones in general are described in Sect. 3.7.

3.1 Problem definition: forecasting on unseen roads

Figure 1 shows how traffic signals are generated by sensors installed on roads in a traffic network. This can be formulated as a multivariate timeseries where the traffic signals from the sensors forms a timeseries. An alternative abstraction is for a sensor to be formulated as a node in a spatio-temporal graph. The edges in this graph capture the spatial connections and relationships between sensors, reflecting either the physical proximity between sensors and road segments or the similarity of traffic signals in the latent space. From these perspectives, we can use the terms node, sensor, and road interchangeably depending on the context.

The traffic dataset, denoted as $\mathbf{X} \in \mathbb{R}^{M \times K}$, is represented as a tensor. Here, M represents the number of traffic sensors in the dataset, and \mathcal{M} is the set of all road segments or sensors in the traffic network. Thus, $M = |\mathcal{M}|$. The dimension K corresponds to the number of timesteps in the dataset. For brevity, we assume there is only one traffic metric per timestep per sensor. However, our method generalizes to multiple traffic metrics. Each data point $\mathbf{x}_{\mathcal{N},k} = \mathbf{X}_{\mathcal{N},k:k+L}$ is a tensor, where $N = |\mathcal{N}|$ is the number of traffic sensors in the data point and where $\mathcal{N} \in \mathcal{M}$ and, thus, $N \leq M$, and k is the index of the timestep, and L is the number of timesteps in the data point. The tensor \mathbf{X} is the traffic metric at a particular road, in a particular timestep.

As road networks are abstracted to spatio-temporal graphs with weighted edges, the topological connectivity between nodes is represented by a sparse adjacency matrix (\mathbf{A}). The adjacency matrix is normalized between zero and one. An edge with higher weight means that the nodes are closer together, while lower weight means that the nodes are further apart. This topological formulation works both in the physical and latent spaces.

The classical traffic forecasting task is a multi-step forecasting problem formalized as follows: $\mathbf{h}(\mathbf{x}_{\mathcal{M},k}, \mathbf{A}) = \mathbf{x}_{\mathcal{M},k:L+F}$ where F is the forecasting horizon.

However, since our model has not been trained on the data from all the roads in \mathcal{M} , we augment the forecast using a latent representation of the historical data produced by a spatial encoder $\mathbf{E}(\cdot)$. Thus, our final formulation is as follows: $\mathbf{h}(\mathbf{x}_{\mathcal{M},k}, \mathbf{A}, \mathbf{E}(\mathbf{X}_{\mathcal{M},:k-1})) = \mathbf{x}_{\mathcal{M},k:L+F}$

3.2 Spatio-temporal split

When it comes to training/validation/testing splitting strategies, classical traffic forecasting where all roads are seen during training follows the setup in a typical timeseries. The usual random shuffling is not appropriate for timeseries as

training can only be done on historical data, and inferences are typically made on future data. Temporally, the validation set must be at the future of training set, and the testing set is at the future of the validation set. Spatially, the whole set contains all of the sensors. This is shown in the top row of the right hand side of Fig. 2.

However, in our setup (bottom row) we are performing inference (testing) on roads that the model has never seen. We call this new strategy a spatio-temporal (ST) split. Here, we splits the dataset into nine subsets as shown at the left. The top row shows the typical split a traffic forecasting setup where all roads are seen during training. Meanwhile, the bottom row shows the new setup where not all roads are seen during training and forecasting is done on unseen nodes during inference (testing) time. The encoder and forecasting models are exposed to different splits of the dataset during different stages. When the encoder and model are colored red, they are being trained; but when they are blue, their weights are frozen for inference. Note that the index of each timestep is in sequence, but the road index is random.

During pre-training, the encoder is trained on set *A*. Then, it is validated on set $A \cup B \cup D \cup E$ (not shown in the figure). This is to ensure that the encoder can generalize when it receives data from previously unseen sensors. From this point onward, the encoder’s weights are frozen. The next subsection explains the pre-training procedures in detail.

During training, the model is trained on set *A* with additional input from the encoder. From this point onward, the model’s weights are frozen. During validation, the model is validated on set *E*, which is temporally separate from sets *A* and *D* to ensure that the model is not overfitted on past data, but generalizes into the future. However, the inputs to the encoder is set *D*, which represent the past historical data. During testing, the model is tested on set *I*, which only includes roads that neither the encoder, nor the model, has seen before. The historical data $G \cup H$ serve as the inputs to the encoder.

Regarding the sizes of each set, for the temporal split, we follow the typical non-random split with the ratio of 7/1/2 (Jiang et al. 2021). For the spatial split, we roughly follow the same ratio. In contrast to the non-random temporal split, the spatial split is random. In this work we use a uniform distribution. Further analysis in Sect. 4.7 shows that even uniform sampling produces acceptable results.

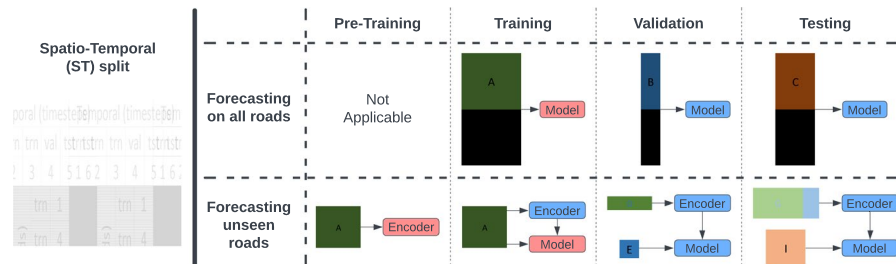


Fig. 2 The ST splitting strategy divides the dataset into nine subsets (left side), while the right side illustrates the usage of different subset combinations at different stages

3.3 Pre-training using contrastive loss

This section describes the pre-training procedure for the spatial encoder. We use set A (Fig. 2) as the pre-training data. The pre-training using contrastive loss, as well as the architecture of the encoder, is shown in Fig. 3. Adapted from SimCLR (Chen et al. 2020b) to increase training efficiency of traffic forecasting, it learns spatial representations by maximizing agreement of representations generated by a stochastic spatial encoder via a contrastive loss.

We randomly sample a minibatch of N sensors from set A . The spatial encoder takes each sensor’s historical data twice and outputs two different representations of the sensor. The representations are different because the spatial encoder is stochastic. More formally:

$$\mathbf{e}_n \sim \mathbf{E}(\mathbf{X}_{n,:}); \mathbf{e}'_n \sim \mathbf{E}(\mathbf{X}_{n,:})$$

where $\mathbf{E}(\cdot)$ is the spatial encoder, $\mathbf{X}_{n,:} \in \mathbb{R}^K$ is the historical data of sensor n , K is the number of timesteps in the set, $\mathbf{e}_n, \mathbf{e}'_n \in \mathbb{R}^D$ are the two different representations of sensor n , and D is the embedding dimension. The details of the spatial encoder will be described in the next subsection.

Although we want the two representations $\mathbf{e}_n, \mathbf{e}'_n$ to be similar to each other, following (Chen et al. 2020b), we define the contrastive loss on the output of the projection head $p(\cdot)$. We use a fully connected layer $FC(\cdot)$ as the projection head (proj. in Fig. 3):

$$\mathbf{z}_n = p(\mathbf{e}_n) = FC(\mathbf{e}_n)$$

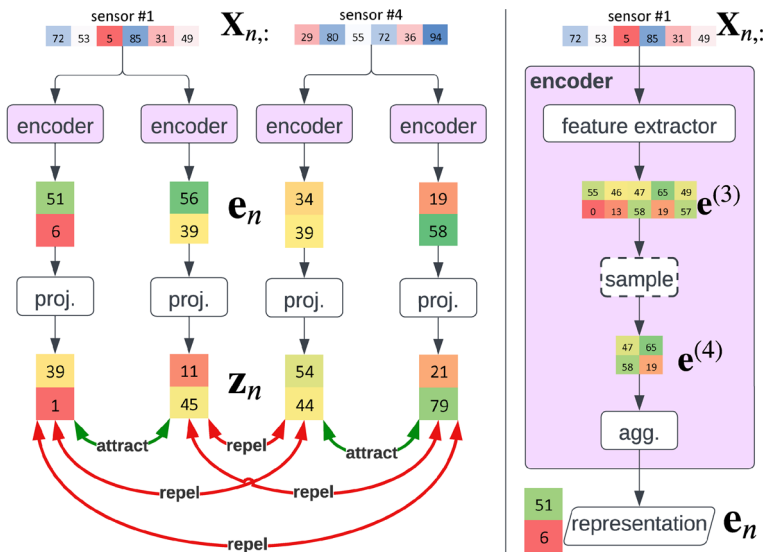


Fig. 3 On the left, the use of contrastive loss to pre-train the spatial encoder is depicted, while on the right, the framework of the (spatial) encoder is illustrated

Finally, we use the normalized temperature-scaled cross entropy (NT-Xent) loss function (Chen et al. 2020b), which is also known as InfoNCE (Oord et al. 2018). The goal of this loss function is to ensure that the representations coming from the same sensor are similar to each other, while representations from different sensors are different. Temperature is a hyperparameter of this loss function. In section B we show our framework is robust across a wide range of temperature.

3.4 Spatial encoder

The spatial encoder $\mathbf{E}(\mathbf{X}_{n,:})$ extracts a latent representation of a road \mathbf{e}_n from its historical data $\mathbf{X}_{n,:}$ in a stochastic manner. Our spatial contrastive pre-training framework is agnostic to the choice of the spatial encoder architecture. In this work, we pick simple architecture roughly based on the feature extractor for graph for time series (GTS) (Shang et al. 2021). In the rest of this section, we describe our implementation of the spatial encoder as shown on the right in Fig. 3.

First, we pass the sensor's historical data $\mathbf{X}_{n,:}$ to a feature extractor. The feature extractor consists of a sequence of three dilated convolutional layers (van den Oord et al. 2016). The windows are of sizes: 13, 12, and 24; while the strides are: 1, 12, and 24, respectively. Since each timestep in all of the datasets are 5 min long, these choices are made such that the output of each layer has the receptive field of sizes: 1 h, 2 h, and 1 day, respectively, and the final convolution output has no overlap between days. The first convolution projects the data to a latent space of size D , which remains constant for the rest of the framework unless noted otherwise. In between the convolutions there is a ReLU activation and also a batch normalization layer. Formally:

$$\mathbf{e}^{(1)} = \text{BN}^{(1)}(\text{ReLU}(\text{conv}_{\text{size}=13, \text{stride}=1}(\mathbf{X}_{n,:}))) \quad (1)$$

$$\mathbf{e}^{(2)} = \text{BN}^{(2)}(\text{ReLU}(\text{conv}_{\text{size}=12, \text{stride}=12}(\mathbf{e}^{(1)}))) \quad (2)$$

$$\mathbf{e}^{(3)} = \text{conv}_{\text{size}=24, \text{stride}=24}(\mathbf{e}^{(2)}) \quad (3)$$

Note that the sizes of $\mathbf{e}^{(i)}$ depends on the size of set A and they are also different for every layer due to the convolution dilation.

To extract multiple representations of sensor n , we sample from the latent space $\mathbf{e}^{(3)}$. Additionally, this sampling process also reduces the dimension from the thousands of timesteps in the raw data to a smaller latent space. Unlike the feature extractor in Shang et al. (2021) that uses a fully connected layer, this sampling process allows the encoder to extract representations from sensors with different lengths of historical data. This offers a practical advantage when we are inferring on newly installed sensors on newly constructed roads with limited historical data. Thus it is possible to perform inferences on new roads with only 2 days of data. We uniformly sampled half of the days:

$$\mathbf{e}^{(4)} = \text{sample}(\mathbf{e}^{(3)})$$

such that $\mathbf{e}^{(3)} \in \mathbb{R}^{D \times K^{(3)}}$, $\mathbf{e}^{(4)} \in \mathbb{R}^{D \times K^{(4)}}$, and $K^{(4)} = K^{(3)}/2$.

Finally, to aggregate the sampled representation, inspired from graph isomorphism network (GIN) (Xu et al. 2019), we take the 1st, 2nd, and ∞^{th} statistical moments of $\mathbf{e}^{(4)}$. Then, we vectorize along the latent dimension and apply a fully connected layer to reduce the dimension back to D . The output is a latent vector representation \mathbf{e}_n of road n . There are also ReLU activation and batch normalization layers as follows:

$$\mathbf{e}_n = \text{BN}^{(4)}(\text{ReLU}(\text{FC}(\text{BN}^{(3)}(\mu(\mathbf{e}^{(4)}) \oplus \text{std}(\mathbf{e}^{(4)}) \oplus \text{max_pool}(\mathbf{e}^{(4)}))))))$$

where $\mu(\cdot)$ is a mean pooling layer, $\text{std}(\cdot)$ is a standard deviation pooling layer, $\text{max_pool}(\cdot)$ is a maximum pooling layer, and \oplus is a concatenation operation along the latent dimension.

3.5 SGA layer

There are many different ways to integrate the output of the encoder into different layers in the forecasting architecture. Here, for simplicity, we opt for summation. However, a naive summation is problematic because it injects the same amount of spatial information all the time. The amount of spatial information that a layer needs might differ between layers. For this reason, we used gating mechanism with a scalar per sensor to decide how much spatial information should be added at each layer.

In brief, $\text{SGA}(\cdot)$ layer adds the latent representation vector \mathbf{e}_n to the activation $\mathbf{h}_n^{(l)}$ of the l^{th} layer of the model, weighted by a coefficient $c_n(\mathbf{h}_n^{(l)}, \mathbf{e}_n)$ that is unique for every sensor n . More formally:

$$\mathbf{h}_n^{(l+1)} = \text{SGA}(\mathbf{h}_n^{(l)}, \mathbf{e}_n) = \mathbf{h}_n^{(l)} + c_n(\mathbf{h}_n^{(l)}, \mathbf{e}_n)\mathbf{e}_n$$

where $\mathbf{h}_n^{(l)} \in \mathbb{R}^D$ is the activation of the l^{th} layer for sensor n in the forecasting model, $\mathbf{e}_n \in \mathbb{R}^D$ is the latent representation of sensor n (the output of the frozen spatial encoder), and $c_n(\cdot) \in \mathbb{R}$ is the weight for sensor n at layer l . We calculate the weight $c_n(\cdot)$ using a multi-layer perceptron (MLP) with one hidden layer, ReLU activation, and wrap it under a sigmoid $\sigma(\cdot)$ to ensure that the weight is between 0 and 1:

$$c_n(\mathbf{h}_n^{(l)}, \mathbf{e}_n) \in \mathbb{R} = \sigma(\text{FC}^{(2)}(\text{ReLU}(\text{FC}^{(1)}(\mathbf{h}_n^{(l)} \oplus \mathbf{e}_n))))$$

The final fully connected layer $\text{FC}^{(2)}$ has an output size of 1 to ensure that the output is a scalar.

3.6 Traffic signal decoupling

Data scarcity has been identified as an essential issue in traffic forecasting (Liu et al. 2022a). The forecasting on unseen roads setup only exacerbated the data scarcity

issue. To address this, we decouple traffic signals to periodic Markovian signals. The periodic signals are easy to model, i.e. traffic on Wednesday morning rush hours are similar to traffic from other Wednesday mornings, but is very different than traffic during Friday evening. In contrast, the Markovian signals capture the complex spatio-temporal correlations between the recent past traffic from nearby roads.

Typically, deep learning models are able to learn the periodic signals effectively without compromising their capabilities to learn the Markovian signals at the same time. However, due to the data scarcity issue in the new setup, we argue that it is better for the spatial encoders to only learn the Markovian signals. We enforced this by decoupling the traffic signals.

Formally, we can decouple the traffic signals as follows: $\mathbf{x}_k = \mathbf{s}(k) + \hat{\mathbf{x}}_k$ where $\mathbf{s}(k)$ is the periodic signals since they are only dependant on the timestep index alone and $\hat{\mathbf{x}}_k$ is the remaining Markovian signals. Through this decoupling, the problem formulation becomes: $\mathbf{h}(\mathbf{x}_k, \hat{\mathbf{x}}_k, \mathbf{A}, \mathbf{E}(\hat{\mathbf{X}}_{:k-1})) + \mathbf{s}(k) = \mathbf{x}_{k+L+F}$ as we ignore the spatial dimension in the notations for brevity.

We determined the periodic signals $\mathbf{s}(k)$ by using road-wise DCT (Ahmed et al. 1974) transform on the training set $A \cup D \cup G$, keeping the low-frequency coefficients by setting the high-frequency coefficients to zeroes, and then inverting the transform. The number of low-frequency coefficients to be kept depended upon the minimization of MAE of the reconstruction against the validation set $B \cup E \cup H$.

3.7 Integration with GWN

To demonstrate the capabilities of the spatial contrastive pre-training framework, we combine it with GWN (Wu et al. 2019) as the forecasting model. We picked GWN because it is the current state-of-the-art architecture according to the latest benchmark study (Jiang et al. 2021). The architecture of GWN, as well as the details of the integration with our framework, are illustrated in Fig. 4. We are using the code implementation by the benchmark study (Jiang et al. 2021) that made them publicly available at github.com/deepkashiwa20/DL-Traff-Graph.

In brief, GWN consist of an initial MLP encoder, a sequence of ST-layers (in yellow), and a final MLP decoder. Each ST layer consists of a G-TCN layer (Van Den Oord et al. 2016) and GCN layer (Li et al. 2018). There are also residual connections that go into the ST layers, and skip connections from the output of ST layers. The skip connections are aggregated with concatenation before going into the final MLP decoder. Integration with adaptive adjacency matrix is not shown.

The left hand side of Fig. 4 shows that to integrate the outputs of the spatial encoder (in blue) trained using spatial contrastive pre-training to GWN, the outputs are given to ST-layers (in yellow) of GWN. The right side shows that the components of an ST-layer are made up of gated temporal convolutional (G-TCN) layers and GCN layers. We use SGA (in green) described in the previous subsection to inject the spatial information to the input of the G-TCN and the GCN layers. Note that the weights of the spatial encoder are frozen during the training and inference of the GWN.

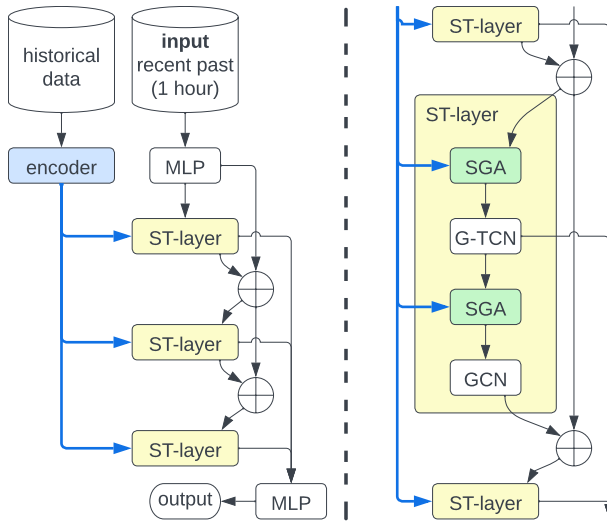


Fig. 4 The left side of the figure illustrates the flow of outputs from the spatial encoder (blue) into the spatio-temporal (ST) layers (yellow). On the right side, the usage of Spatially Gated Addition (SGA) to integrate spatial information from the spatial encoder into the input of the G-TCN and GCN layers within the ST-layer is depicted (Color figure online)

There are many other alternatives to our strategy of injecting spatial information into the spatial and temporal modules. One naive approach is to treat it as a positional encoding at the start with the initial MLP. However, we picked this strategy instead so that the models are not encumbered by forcing it to "remember" the entire spatial information that is only given once at the beginning. Moreover, different layers might also require different "amount" of spatial information. The SGA layer allows the network to modulate the "amount" of spatial information that gets injected at each layer.

Additionally, GWN used an adaptive adjacency matrix \mathbf{A}_{adp} constructed via learned road embeddings $\mathbf{A}_{adp} = SoftMax(ReLU(\mathbf{r}_1, \mathbf{r}_2^T))$ where \mathbf{r}_1 and \mathbf{r}_2 are the road embeddings. Because the road embeddings are learned during training, GWN does not allow the use of adaptive adjacency matrix in an unseen roads setup. However, under the SCPT framework, we can use the outputs of the spatial encoder to infer the node embeddings via an MLP with one hidden layer $\mathbf{r} = FC(ReLU(FC(\mathbf{e})))$. Each FC had their own unique set of weights and were not shared.

4 Experiments

4.1 Dataset

We utilized four real-world datasets for our experiments: **METR-LA**, **METR-LA**, **METR-LA**, and **METR-LA**. The first three mentioned are popular datasets from the latest benchmark study (Jiang et al. 2021). The last one is the largest in the deep

traffic forecasting literature (Mallick et al. 2020). For more detailed description these datasets, including the detailed statistics, please refer to appendix A.

4.2 Setups

We use the same hyperparameters as the benchmark paper (Jiang et al. 2021), including the use of Adam as the optimizer and Mean Absolute Error (MAE) as the loss function of the forecasting model. The forecasting on unseen roads setup is more prone to overfitting on the trained roads. Therefore, we use weight decay which value we treat as a hyperparameter. Since we are comparing training time, we are training (and pre-training) for 100 epochs without the use of early stopping. The latent spaces of the spatial encoder output has 32 dimensions, the same size as the GWN. The middle layer of the MLP in SGA has 128 hidden nodes; the same amount as the hidden layer in the final MLP of GWN. Similarly, the middle layers of the MLP in the construction of adaptive adjacency matrix also have 128 hidden nodes. We used a standard scaler on the input of both the spatial encoder and GWN. When the traffic signals are decoupled, the standard scaler is only applied after the decoupling. All experiments were run on either NVIDIA Tesla P100 or V100 graphic cards.

4.3 Results

To empirically evaluate the capabilities of SCPT, our proposed framework, on forecasting on new roads unseen in the training data, we use the ST split setup and tested it on three real-world data. We replicated each experiment 10 times with different seeds for model weight initialization, sampling in the spatial encoder, and the selection of roads in the ST splits. The results are displayed in Table 1, showing the averages and the standard deviations across the 10 seeds.

The results shows that the SCPT framework consistently improve the backbone baselines across all three datasets and all three metrics. The performance gain are more pronounced in the PeMS-BAY and PeMS-D7(m) datasets when compared to the METR-LA dataset. This can be attributed to the fact that the METR-LA dataset has the widest speed distribution as indicated in the standard deviation in Table 4. More importantly, the SCPT framework also consistently reduce the standard deviations across all three datasets and all three metrics. Thus, the SCPT framework, not only offers better performances, but also higher reliability.

4.4 Performance across forecasting horizons

To further analyze the performance of the SCPT framework, we breakdown the results from the PeMS-D7(m) dataset in Table 1 based on forecasting horizons. This is shown in Fig. 5. The average performance over the 10 seeds are show by the solid lines. The thicker red line shows the percentage improvement of SCPT. As expected, we see decreases in performance (increase in MAE) as

Table 1 Performances evaluation of the SCPT framework using ST split

Dataset	Methods	RMSE	MAE	MAPE
METR-LA	GWN	10.3405 ± 0.2634	4.7373 ± 0.1618	12.2677 ± 0.8058
	GWN+SCPT	10.0385 ± 0.2112	4.5645 ± 0.1556	11.5002 ± 0.8007
	Δ(%)	3%	4%	6%
PeMS-BAY	GWN	4.5059 ± 0.1613	2.0126 ± 0.1037	4.7779 ± 0.4303
	GWN+SCPT	3.9658 ± 0.1266	1.8163 ± 0.0875	4.1358 ± 0.2740
	Δ(%)	12%	10%	13%
PeMS-D7(m)	GWN	6.4635 ± 0.3103	3.4327 ± 0.1974	8.6896 ± 0.7844
	GWN+SCPT	5.6893 ± 0.2552	3.0794 ± 0.1448	7.6770 ± 0.6678
	Δ(%)	12%	10%	12%

In this setup, the models are trained on only 70%, validated on 10%, and tested on 20% of the roads. This table shows the average performance across 12 timesteps (1 h) on the 20% of the roads that are unseen during the training with the best performance marked in bold. Δ(%) denotes the percentage of error reduction

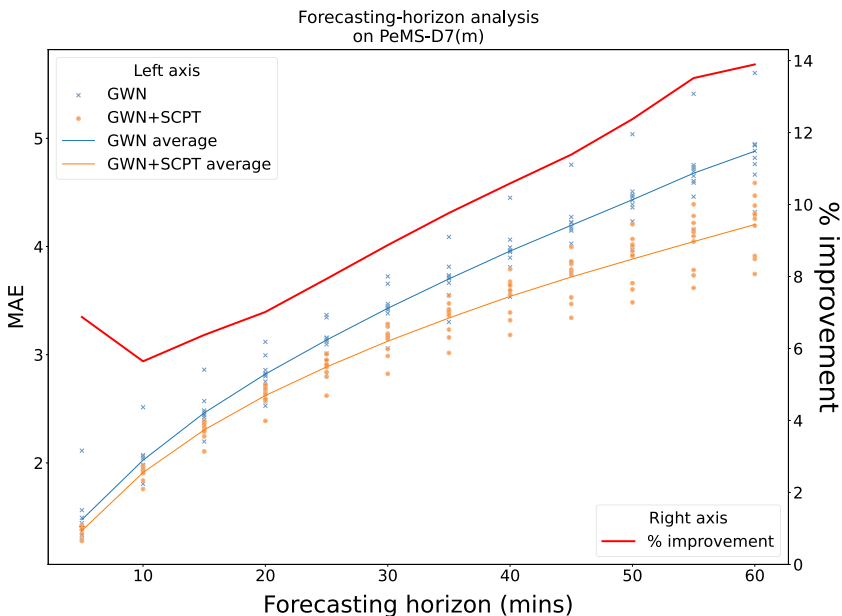


Fig. 5 Performance across forecasting horizons

the forecasting horizon increases i.e. it is harder to forecast further to the future. The SCPT framework always improve the performance across all forecasting horizons. More importantly, the performance gain brought by the SCPT framework also widens as forecasting horizon increases. Although the average MAE improvement is 10% as shown in Table 1, at the furthest forecasting horizon (1 h ahead), the improvement is nearly 14%. These results show the superiority of the SCPT framework.

4.5 Ablation study

Here, we performed a complete ablation study to evaluate each component of the SCPT framework by iterating through all the possible combinations. Table 2 shows the results of the ablation study. Firstly, all of the entry of the + column is positive, meaning that all the individual components and combinations of them improve the baseline.

Based on this analysis, the most important component is AdpAdj, that is, the use of the output of the spatial encoders to construct the node embeddings in GWN. GWN uses these node embeddings to construct the a adaptive adjacency. Ablating AdpAdj effectively removed the adaptive adjacency from GWN as it has no mechanism to infer the node embeddings of new roads unseen in the training data. The biggest lost of performance (- column) is found when ablating the AdpAdj, the second last row.

This finding agrees with the growing literature on the importance of inferring the latent topological connectivity of the traffic networks (Wu et al. 2020; Shang et al. 2021; Li et al. 2022b). This further highlights the capabilities of the SCPT frameworks to be easily integrated to various backbone models that require learned node embeddings and learning the latent topological connectivity of the traffic networks.

Additionally, our analysis revealed that the SGA component plays a critical role in integrating the output of the spatial encoder into the backbone network. Notably, the addition of SGA alone exhibited comparable improvements in performance to the addition of AdpAdj alone. Specifically, both configurations exhibited a gain of 0.083 in the third and fifth row of the + column. This finding highlights the non-trivial nature of combining the output of the spatial encoder

Table 2 Ablation study on the PeMS-D7(m) dataset

Methods				MAE			
SCPT	SGA	Decoupling	AdpAdj	mean	std	+	-
0	0	0	0	3.433	0.197	0.000	0.353
✓	0	0	0	3.366	0.181	0.066	0.287
✓	✓	0	0	3.349	0.161	0.083	0.270
✓	0	✓	0	3.398	0.236	0.035	0.319
✓	0	0	✓	3.350	0.234	0.083	0.271
✓	0	✓	✓	3.249	0.255	0.184	0.169
✓	✓	0	✓	3.101	0.141	0.332	0.022
✓	✓	✓	0	3.406	0.187	0.026	0.327
✓	✓	✓	✓	3.079	0.145	0.353	0.000

The best performance are marked in bold

Every experiment is replicated five times (except the first and last ones which are replicated ten times). The first row is the backbone baseline GWN and the last row is the full GWN+SCPT. The + column shows the MAE reduction when compared to the GWN baseline (first row). The - column shows the performance reduction when compared to the full SCPT framework

with the backbone network and emphasizes the importance of the SGA layer in effectively incorporating spatial information.

4.6 Scalability to large dataset

The capability to generalize forecasting performance well to unseen roads also opens up new avenues for more efficient traffic forecasting. Training on large traffic network is costly. For example, training a model on one third of the entire California highway network took nearly 8 GPU-days (Mallick et al. 2020). To make deep traffic forecasting more applicable for traffic management, a new scaling paradigm is required. Instead of training a model on the entire traffic network, SCPT frameworks provide traffic managers with an attractive trade-off between forecasting accuracy and lower training cost by training only on a sample of roads while generalizing performance to the rest of the network.

To explore this new direction towards a more efficient traffic forecasting, we run a set of experiments where we train only on 1% of the roads in PeMS-11k(s), but tested on the entire traffic network. The results are shown in Table 3 We also use a metric called medianMAE12 to compare our results with (Mallick et al. 2020). In medianMAE12, the MAE per road is calculated only for the 1 h ahead forecasting horizon (12 timesteps). Then, instead of averaging the MAE across all the roads, we take the median.

The results shows that training on a smaller sample of roads is a feasible way to increase the efficiency of traffic forecasting on larger traffic networks. The GWN baseline achieved comparable performance to PeMS-BAY. When combined with the SCPT framework, the performances increase even further. The additional training cost is only 6 GPU-minutes. The small degradation in performance is tolerable compared to more than hundreds fold savings in training time.

Table 3 Performance comparison on using the SCPT framework to train on a small sample (1%) of roads to scale to a large dataset PeMS-11k(s) with the best marked in bold. $\Delta(\%)$ denotes the percentage of error reduction

Method	GWN	GWN+SCPT	$\Delta(\%)$	GP-DCRNN
RMSE	5.6345 \pm 0.7469	4.6741 \pm 0.2089	17%	
MAE	2.8241 \pm 0.2840	2.4273 \pm 0.2171	14%	
MAPE	5.6345 \pm 0.7469	4.6741 \pm 0.2089	17%	
medianMAE12	3.4554 \pm 0.2343	3.2442 \pm 0.3071	6%	2.0200
Training time	00:16:39	00:22:28		7 days, 22:34:53
Roads seen in training (count)	111			11160
Roads seen in training (%)	1%			100%

4.7 Effect of randomness on sensors selection during the spatio-temporal split

In this set of experiments, we show that the impact of randomness in the ST split (Fig. 2) is greater than the effect of randomness in the models’ weight initialization. To this end, we separated the random seed used for weight initialization of the model and the random seed used for the ST split, and compared the impact. We replicated this 36 times, each with a unique combination of model and splitting seeds. The METR-LA dataset was used.

The result is shown in Fig. 6. Each value in a cell is the MAE of a run where the seed for the weight initialization is the heading of the row and the seed for sensor splitting is the heading of the column. In every row, the model initialization seed is kept constant, so the spread captured by the standard deviation (on the right hand side of the figure) is solely due to randomness in the sensor selection process. Similarly, for every column, the ST split is kept constant, so the spread is solely due to the randomness in the weight initialization process. The average standard deviation due to randomness in the spatio-temporal splitting process (0.060) is similar to that of the model initialization process (0.057). Notably, the combination of both sources of randomness (bottom right) does not significantly increase the spread of the performance (0.068). This shows that the uniform random distribution over the sensors employed in the spatio-temporal split strategy does not affect the framework’s performance significantly more than the randomness due to model’s weight initialization.

MAE		spatial seed						std. (0.060)
		1	2	3	4	5	6	
model initialization seed	1	4.713	4.660	4.773	4.698	4.634	4.718	0.044
	2	4.679	4.778	4.727	4.698	4.702	4.756	0.034
	3	4.598	4.684	4.727	4.620	4.788	4.692	0.063
	4	4.815	4.706	4.777	4.654	4.668	4.766	0.059
	5	4.647	4.723	4.724	4.707	4.577	4.731	0.056
	6	4.872	4.608	4.698	4.676	4.535	4.612	0.106
std. (0.057)		0.095	0.053	0.028	0.031	0.083	0.051	0.068

Fig. 6 Analyzing the proposed framework’s performance variance based on randomness in sensor selections in comparison with randomness in model weight initialization

4.8 Computational cost of pre-training

Spatial contrastive pre-training has a minimal impact on additional computational cost. For the METR-LA dataset, when trained on all of the sensors, pre-training took less than 33 s, a negligible amount when compared to the training time of the forecasting model, which took 2 h, 17 min. Similarly, on PeMS-BAY, pre-training added an extra 58 s to the 5 h and 27 min of the training process.

5 Limitations and future works

In the context of ST splits, our current approach of using uniform sampling has shown acceptable results (Sect. 4.7). However, we acknowledge the limitation of its naivety and the potential for further improvement. In future work, we aim to explore more sophisticated sampling strategies, such as clustering-based approaches where we cluster the sensors based on the traffic signals, and sample only the centroids, or sampling based on the network topology. These strategies have the potential to enhance the representation of the road network and improve the forecasting performance.

Our proposed SCPT framework is agnostic to the exact architecture of the spatial encoder. The only requirement is that it behaves in a stochastic manner. In this work, we proposed a simplistic WaveNet (Van Den Oord et al. 2016) based encoder with statistical moments (Xu et al. 2019) aggregations. It is yet to be explored if performance can be improved by using more sophisticated domain generic encoders such as r-drop (Liang et al. 2021) and data2vec (Baevski et al. 2022), or something more tailored to time series such as ts2vec (Yue et al. 2022).

In this paper, we employ different architectures for the spatial encoder and the backbone network. Although using a single architecture for both components would be an elegant solution, it is not a straightforward task. The spatial encoder's role is to summarize contrastive features from long time series (months to years), while the backbone network is designed to capture complex dynamics from shorter time series (minutes to hours). Balancing these distinct requirements and developing a unified architecture would be an intriguing direction for future research.

We chose GWN (Wu et al. 2019) to be our backbone architecture because it is currently the best architecture available. It is yet to be confirmed if the efficiency improvements brought by our framework can generalize to future state-of-the-art architectures.

Finally, we mentioned that forecasting using the SCPT framework can be generalized to any multivariate timeseries beyond just traffic. More research are required in this direction.

6 Conclusion

In this paper, we proposed a novel task, that is, to perform traffic forecasting on new roads unseen in the training data. To perform evaluations on this new task, we propose a novel setup called ST split. Then we introduced the SCPT framework to

train a spatial encoder on sensors' historical data. Additionally, we implemented a simple spatial encoder to showcase our framework. Next, we introduced an SGA layers, traffic signal decoupling, and a method to infer node embeddings using the output of the spatial encoder. Finally, we evaluated our framework using GWN as the backbone forecasting model, on four real world datasets, showing consistent increases in performance and extensively analysed all of the components in the SCPT framework.

A: Datasets details

We utilized four real-world datasets for our experiments. The **METR-LA** dataset was collected from loop detectors in Los Angeles, United States county highways. The **PeMS-BAY** dataset was collected by the California Transportation Agencies (CalTrans) Performance Measurement System (PeMS) using loop detectors located around the San Francisco Bay Area. Similarly, the **PeMS-D7(m)** dataset was also collected by CalTrans PeMS using loop detectors, but only during weekdays. The **PeMS-11k** dataset, which is the largest in the deep traffic forecasting literature, was also collected by CalTrans PeMS using loop detectors. However, to focus on spatial generalization and optimize resource utilization, we only utilized a two-month period from the original one-year-long dataset, referred to as **PeMS-11k(s)**.

The first three datasets mentioned are from the latest benchmark study (Jiang et al. 2021), and they have been made publicly available on their GitHub repository github.com/deepkashiwa20/DL-Traff-Graph. The last dataset is from (Mallick et al. 2020) and is also available on their GitHub page https://github.com/tanwimallick/graph_partition_based_DCRNN. For detailed statistics about these datasets, please refer to Table 4.

B: Temperature sensitivity

One of the main hyperparameters introduced in our framework is the NT-Xent temperature. We conducted an experiment to show that our framework is robust across a large range of temperature values. This experiment was run using the METR-LA dataset. Since temperature is a hyperparameter, we are analyzing the framework's performance on the validation set.

Figure 7 shows the result of this set of experiments. It shows that the validation performance is stable across a wide range of temperature, from around 10^1 to approximately 10^4 , which is about three orders of magnitude. Additionally, the validation MAE is correlated with the test MAE over changes in temperature, as shown in Fig. 8, which shows that validation and test MAE are correlated against variation in NT-Xent temperature. These results suggest that the robustness of the framework across different temperatures during hyperparameter optimization can reasonably transfer to the test performance.

Table 4 Detailed statistics on the real world datasets

Dataset:		METRLA	PeMSBAY	PeMS-D7(m)	PeMS-11k(s)
Spatial	Nodes	207	325	228	11,160
	Edges	1,515	2,694	7,304	234,966
Temporal	Duration (timesteps)	34,272	52,116	12,672	25,632
	Duration (days)	121	150	61	89
	Time start	01-Mar-12	01-Jan-17	01-May-12	01-Feb-18
	Time end	30-Jun-12	31-May-17	30-Jun-12	30-Apr-18
	Granularity (mins)	5	5	5	5
Speed (mph)	Min	0.00	0.00	3.00	3.00
	Q1	57.13	62.10	57.50	62.60
	Median	63.22	65.30	64.10	65.10
	Mean	58.46	62.62	58.89	63.14
	Q3	66.50	67.50	66.70	67.80
	Max	70.00	85.10	82.60	99.30
	Standard Deviation	20.26	9.59	13.48	9.01
	Missing values	8.82%	0.00%	0.00%	0.00%
Size	Entry	7,094,304	16,937,700	2,889,216	286,053,120
	Compressed (MB)	54	130	6	2,235

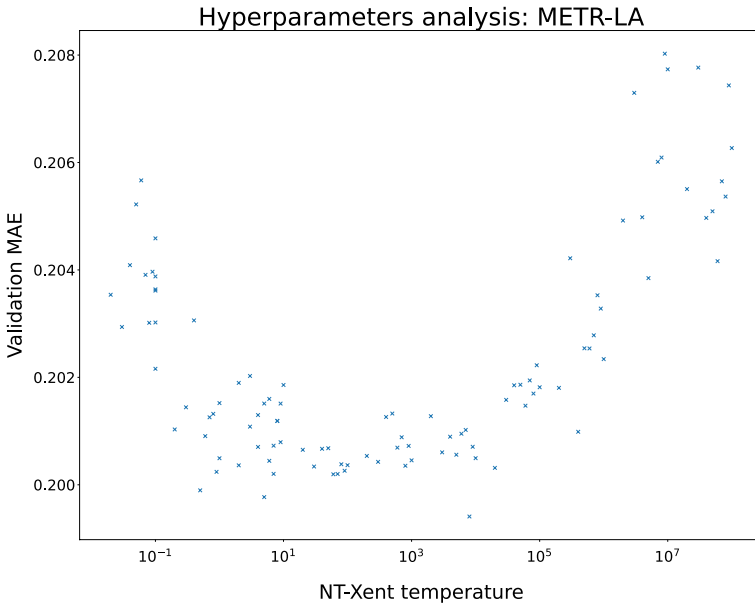


Fig. 7 Analyzing the framework’s sensitivity against NT-Xent temperature hyperparameter

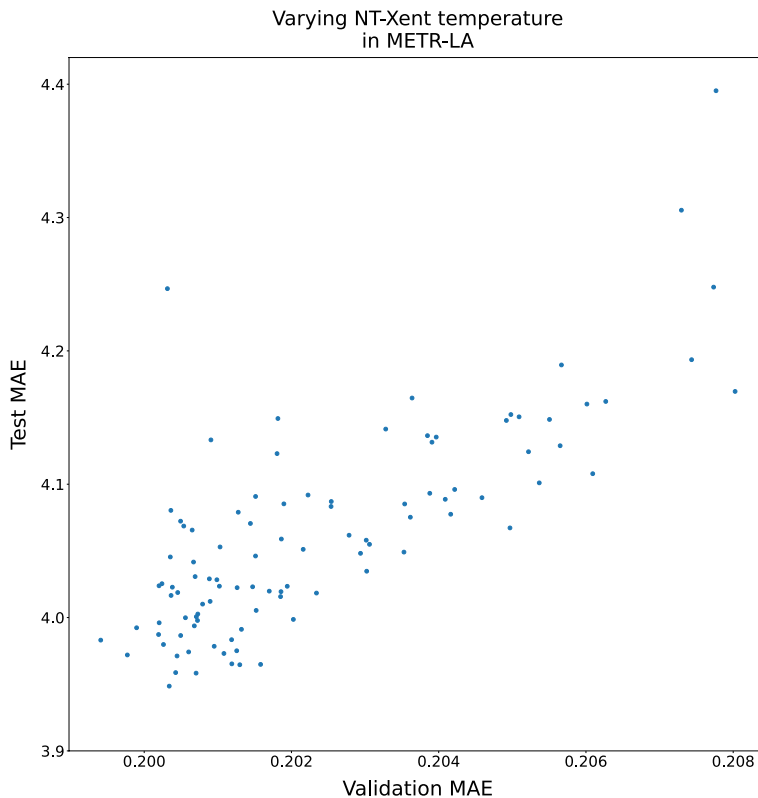


Fig. 8 Correlation validation and test MAE when varying NT-Xent temperature

Funding Open Access funding enabled and organized by CAUL and its Member Institutions.

Declarations

Conflict of interest The authors declare that this research was supported by the Data61/CSIRO PhD scholarship, RMIT Research Stipend Scholarship (RRSS), and the Australian Government RTP Scholarship. We would like to also acknowledge the support of the Investigative Analytics team (Data61/CSIRO) and Cisco's National Industry Innovation Network (NIIN) Research Chair Program. The research utilized computing resources and services provided by Gadi, supercomputer of the National Computational Infrastructure (NCI) supported by the Australian Government, and Bracewell, supercomputer of the Commonwealth Scientific and Industrial Research Organisation (CSIRO).

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Ahmed MS, Cook AR (1979) Analysis of Freeway Traffic Time-series Data by Using Box-Jenkins Techniques vol. 722. Transportation Research Record
- Ahmed N, Natarajan T, Rao KR (1974) Discrete cosine transform. *IEEE Trans Comput* 100(1):90–93
- Baeviski A, Hsu W-N, Xu Q, Babu A, Gu J, Auli M (2022) Data2vec: a general framework for self-supervised learning in speech, vision and language. arXiv preprint [arXiv:2202.03555](https://arxiv.org/abs/2202.03555)
- Bai L, Yao L, Li C, Wang X, Wang C (2020) Adaptive graph convolutional recurrent network for traffic forecasting. In: *NeurIPS*
- Brown T, Mann B, Ryder N, Subbiah M, Kaplan JD, Dhariwal P, Neelakantan A, Shyam P, Sastry G, Askell A et al (2020) Language models are few-shot learners. *NeurIPS* 33:1877–1901
- Chen C, Wang Y, Li L, Hu J, Zhang Z (2012) The retrieval of intra-day trend and its influence on traffic prediction. *Transp Res Part C Emerg Technol* 22:103–118
- Chen T, Kornblith S, Norouzi M, Hinton G (2020) A simple framework for contrastive learning of visual representations. In: *ICML*, pp. 1597–1607. PMLR
- Chen M, Radford A, Child R, Wu J, Jun H, Luan D, Sutskever I (2020) Generative pretraining from pixels. In: *ICML*, pp. 1691–1703. PMLR
- Cui Z, Lin L, Pu Z, Wang Y (2020) Graph markov network for traffic forecasting with missing data. *Transp Res Part C Emerg Technol* 117:102671
- Defferrard M, Bresson X, Vandergheynst P (2016) Convolutional neural networks on graphs with fast localized spectral filtering. *NeurIPS* 29:3844–3852
- Fu R, Zhang Z, Li L (2016) Using lstm and gru neural network methods for traffic flow prediction. In: *YAC*, pp 324–328. IEEE
- Hamed MM, Al-Masaeid HR, Said ZMB (1995) Short-term prediction of traffic volume in urban arterials. *J Transp Eng* 121(3):249–254
- Jeong Y-S, Byon Y-J, Castro-Neto MM, Easa SM (2013) Supervised weighting-online learning algorithm for short-term traffic flow prediction. *TITS* 14(4):1700–1707
- Jiang R, Yin D, Wang Z, Wang Y, Deng J, Liu H, Cai Z, Deng J, Song X, Shibasaki R (2021) DI-traffic: survey and benchmark of deep learning models for urban traffic prediction. *CIKM*
- Kamarianakis Y, Prastacos P (2003) Forecasting traffic flow conditions in an urban network: comparison of multivariate and univariate approaches. *Transp Res Record* 1857(1):74–84
- Lee S, Fambro DB (1999) Application of subset autoregressive integrated moving average model for short-term freeway traffic volume forecasting. *Transp Res Record* 1678(1):179–188
- Liang X, Wu L, Li J, Wang Y, Meng Q, Qin T, Chen W, Zhang M, Liu T-Y (2021) R-drop: regularized dropout for neural networks. In: *NeurIPS*
- Lin T-Y, Maji S, Koniusz P (2018) Second-order democratic aggregation. In: *ECCV*
- Lippi M, Bertini M, Frasconi P (2013) Short-term traffic flow forecasting: An experimental comparison of time-series analysis and supervised learning. *TITS* 14(2):871–882
- Liu X, Liang Y, Huang C, Zheng Y, Hooi B, Zimmermann R (2022) When do contrastive learning signals help spatio-temporal graph forecasting? In: *SIGSPATIAL*
- Li Y, Yu R, Shahabi C, Liu Y (2018) Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. In: *ICLR*
- Li R, Zhong T, Jiang X, Trajcevski G, Wu J, Zhou F (2022) Mining spatio-temporal relations via self-paced graph contrastive learning. In: *SIGKDD*, pp. 936–944
- Lv Y, Duan Y, Kang W, Li Z, Wang F-Y (2014) Traffic flow prediction with big data: a deep learning approach. *T-ITS* 16(2):865–873
- Mallick T, Balaprakash P, Rask E, Macfarlane J (2020) Graph-partitioning-based diffusion convolutional recurrent neural network for large-scale traffic forecasting. *Transp Res Record* 2674(9):473–488
- Manibardo EL, Laña I, Del Ser J (2021) Deep learning for road traffic forecasting: Does it make a difference? *T-ITS*
- Oord A.v.d, Li Y, Vinyals O (2018) Representation learning with contrastive predictive coding. arXiv preprint [arXiv:1807.03748](https://arxiv.org/abs/1807.03748)
- Prabowo A (2022) Spatiotemporal deep learning. PhD thesis, RMIT University
- Prabowo A, Chen K, Xue H, Sethuvenkatraman S, Salim FD (2023) Continually learning out-of-distribution spatiotemporal data for robust energy forecasting. In: *ECML PKDD*. Springer

- Prabowo A, Koniusz P, Shao W, Salim F (2019) Coltrane: convolutional trajectory network for deep map inference. *BuildSys*, p 10. Association for Computing Machinery, New York, USA
- Prabowo A, Shao W, Xue H, Koniusz P, Salim FD (2023) Because every sensor is unique, so is every pair: Handling dynamicity in traffic forecasting. In: *IoTDI*, pp 93–104. Association for Computing Machinery, New York, NY, USA
- Prabowo A, Xue H, Shao W, Koniusz P, Salim FD (2023) Message Passing Neural Networks for Traffic Forecasting
- Roth A, Liebig T (2022) Forecasting unobserved node states with spatio-temporal graph neural networks. In: *Data Mining Workshops ICDMW'22*
- Shang, C., Chen, J., Bi, J (2021) Discrete graph structure learning for forecasting multiple time series. In: *ICLR*
- Shao H (2020) Deep learning approaches for traffic prediction. PhD thesis, Nanyang Technological University, Nanyang
- Shao W, Prabowo A, Zhao S, Koniusz P, Salim FD (2022) Predicting flight delay with spatio-temporal trajectory convolutional network and airport situational awareness map. *Neurocomputing* 472:280–293
- Shao Z, Zhang Z, Wei W, Wang F, Xu Y, Cao X, Jensen CS (2022) Decoupled dynamic spatial-temporal graph neural network for traffic forecasting. *Proc VLDB Endow* 15(11):2733–2746
- Shao W, Prabowo A, Zhao S, Tan S, Koniusz P, Chan J, Hei X, Feest B, Salim FD (2019) Flight delay prediction using airport situational awareness map. *SIGSPATIAL '19*, pp 432–435. Association for Computing Machinery, New York, NY, USA
- Snell J, Swersky K, Zemel RS (2017) Prototypical networks for few-shot learning. In: Guyon I, von Luxburg U, Bengio S, Wallach HM, Fergus R, Vishwanathan SVN, Garnett R (eds) *NeurIPS*, pp 4077–4087
- Van Den Oord A, Dieleman S, Zen H, Simonyan K, Vinyals O, Graves A, Kalchbrenner N, Senior AW, Kavukcuoglu K (2016) Wavenet: A generative model for raw audio. *SSW* 125:2
- van den Oord A, Dieleman S, Zen H, Simonyan K, Vinyals O, Graves A, Kalchbrenner N, Senior A, Kavukcuoglu K (2016) WaveNet: A Generative Model for Raw Audio. In: *Proc. 9th ISCA workshop on speech synthesis workshop (SSW 9)*, p 125
- Van Der Voort M, Dougherty M, Watson S (1996) Combining kohonen maps with arima time series models to forecast traffic flow. *Transp Res Part C Emerg Technol* 4(5):307–318
- Wang T, Isola P (2020) Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In: *ICML*
- Wang L, Koniusz P (2022) Uncertainty-dtw for time series and sequences. In: *ECCV*, pp 176–195. Springer
- Wang L, Koniusz P (2022) Temporal-viewpoint transportation plan for skeletal few-shot action recognition. In: *ACCV*
- Wang L, Liu J, Koniusz P (2021) 3D skeleton-based few-shot action recognition with JEANIE is not so naïve. *arXiv preprint arXiv: 2112.12668*
- Williams BM (2001) Multivariate vehicular traffic flow prediction: evaluation of arimax modeling. *Transp Res Record* 1776(1):194–200
- Williams BM, Hoel LA (2003) Modeling and forecasting vehicular traffic flow as a seasonal arima process: theoretical basis and empirical results. *J Transp Eng* 129(6):664–672
- Wu Z, Pan S, Long G, Jiang J, Chang X, Zhang C (2020) Connecting the dots: multivariate time series forecasting with graph neural networks. In: *SIGKDD*
- Wu Z, Pan S, Long G, Jiang J, Zhang C (2019) Graph wavenet for deep spatial-temporal graph modeling. *IJCAI*
- Xu K, Hu W, Leskovec J, Jegelka S (2019) How powerful are graph neural networks? In: *ICLR*
- Yue Z, Wang Y, Duan J, Yang T, Huang C, Tong Y, Xu B (2022) Ts2vec: towards universal representation of time series
- Yu B, Yin H, Zhu Z (2018) Spatio-temporal graph convolutional networks: a deep learning framework for traffic forecasting. In: *IJCAI*
- Zhang S, Luo D, Wang L, Koniusz P (2020) Few-shot object detection by second-order pooling. In: *ACCV*
- Zhang S, Murray N, Wang L, Koniusz P (2022) Time-reversed diffusion tensor transformer: a new TENET of few-shot object detection. In: *ECCV*. Springer
- Zhang S, Wang L, Murray N, Koniusz P (2022) Kernelized few-shot object detection with efficient integral aggregation. In: *CVPR*

- Zhang Y, Zhu H, Song Z, Koniusz P, King I (2022) Costa: covariance-preserving feature augmentation for graph contrastive learning. In: SIGKDD
- Zhang Y, Zhu H, Song Z, Koniusz P, King I (2023) Spectral feature augmentation for graph contrastive learning and beyond. In: AAAI
- Zhu H, Koniusz P (2021) Simple spectral graph convolution. In: ICLR
- Zhu H, Koniusz P (2022) Generalized laplacian eigenmaps. NeurIPS
- Zhu H, Koniusz P (2023) Transductive few-shot learning with prototype-based label propagation by iterative graph refinement. CVPR
- Zhu H, Koniusz, P (2022) EASE: Unsupervised discriminant subspace learning for transductive few-shot learning. CVPR
- Zhu H, Sun K, Koniusz P (2021) Contrastive laplacian eigenmaps. NeurIPS

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

Arian Prabowo^{1,2,3}  · Hao Xue³ · Wei Shao² · Piotr Koniusz^{2,4} · Flora D. Salim³

✉ Flora D. Salim
flora.salim@unsw.edu.au

Arian Prabowo
arian.prabowo@gmail.com

Hao Xue
hao.xue1@unsw.edu.au

Wei Shao
wei.shao@data61.csiro.au

Piotr Koniusz
piotr.koniusz@data61.csiro.au

¹ Computing Technologies, RMIT, Melbourne, VIC 3000, Australia

² Data61, CSIRO, Canberra, ACT 2601, Australia

³ Computer Science and Engineering, UNSW, Sydney, NSW 2052, Australia

⁴ Engineering, Computing and Cybernetics, ANU, Canberra, ACT 2600, Australia