



# Graph Neural Processes for Spatio-Temporal Extrapolation

Junfeng Hu  
National University of Singapore  
junfengh@u.nus.edu

Yuxuan Liang\*  
Hong Kong University of Science and  
Technology (Guangzhou)  
yuxuanliang@hkust-gz.edu.cn

Zhencheng Fan  
University of Technology Sydney  
zhencheng.fan@student.uts.edu.au

Hongyang Chen  
Zhejiang Lab  
dr.h.chen@ieee.org

Yu Zheng  
JD Intelligent Cities Research & JD  
iCity, JD Technology  
msyuzheng@outlook.com

Roger Zimmermann  
National University of Singapore  
rogerz@comp.nus.edu.sg

## ABSTRACT

We study the task of spatio-temporal extrapolation that generates data at target locations from surrounding contexts in a graph. This task is crucial as sensors that collect data are sparsely deployed, resulting in a lack of fine-grained information due to high deployment and maintenance costs. Existing methods either use learning-based models like Neural Networks or statistical approaches like Gaussian Processes for this task. However, the former lacks uncertainty estimates and the latter fails to capture complex spatial and temporal correlations effectively. To address these issues, we propose Spatio-Temporal Graph Neural Processes (STGNP), a neural latent variable model which commands these capabilities simultaneously. Specifically, we first learn deterministic spatio-temporal representations by stacking layers of causal convolutions and cross-set graph neural networks. Then, we learn latent variables for target locations through vertical latent state transitions along layers and obtain extrapolations. Importantly during the transitions, we propose Graph Bayesian Aggregation (GBA), a Bayesian graph aggregator that aggregates contexts considering uncertainties in context data and graph structure. Extensive experiments show that STGNP has desirable properties such as uncertainty estimates and strong learning capabilities, and achieves state-of-the-art results by a clear margin.

## CCS CONCEPTS

• **Information systems** → *Spatial-temporal systems*.

## KEYWORDS

Spatio-temporal data mining, neural processes, data extrapolation

### ACM Reference Format:

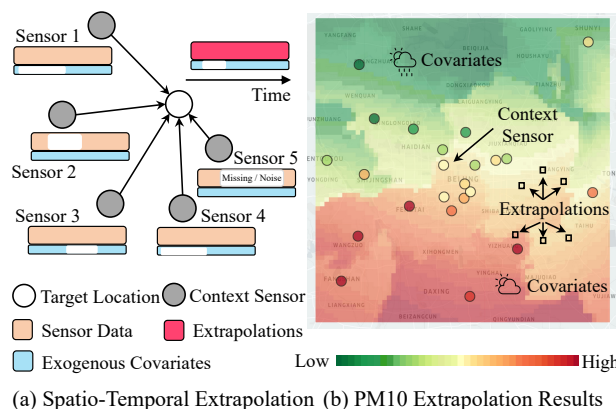
Junfeng Hu, Yuxuan Liang, Zhencheng Fan, Hongyang Chen, Yu Zheng, and Roger Zimmermann. 2023. Graph Neural Processes for Spatio-Temporal Extrapolation. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '23)*, August 6–10, 2023, Long Beach, CA, USA. ACM, Long Beach, CA, 12 pages. <https://doi.org/10.1145/3580305.3599372>

\*Y. Liang is the corresponding author of this paper. Email: yuxliang@outlook.com



This work is licensed under a Creative Commons Attribution International 4.0 License.

KDD '23, August 6–10, 2023, Long Beach, CA, USA.  
© 2023 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-0103-0/23/08.  
<https://doi.org/10.1145/3580305.3599372>



**Figure 1:** (a) Context sensors 1-5 are utilized to generate data of a target location, considering exogenous covariates and the graph structure. (b) Extrapolations of our STGNP.

## 1 INTRODUCTION

Spatio-temporal graph data, such as air quality readings [24, 51], traffic flow data [23, 31, 46], and climate data [25] reported by deployed sensors, are ubiquitous in the physical world. Analyzing such data fosters a variety of applications for smart cities, enhancing people's lives and helping decision-making [15]. Ideally, the data should be fine-grained to realize its benefits but it is often impractical to deploy and maintain sufficient sensors in an area of interest because of the high expenditure [36, 43]. For example, a professional station to measure air quality can cost around \$200,000 to build and \$30,000 per year to maintain [50]. As a result, many applications have to rely on sparse data, leading to suboptimal solutions. Thus, finding ways to approximate the data in areas with no sensors has become a pressing issue.

In this paper, we study the problem of *spatio-temporal extrapolation*, which involves estimating a function that predicts the spatio-temporal data at target locations of interest areas based on the surrounding context nodes and related exogenous covariates, operated in a graph structure, as shown in Figure 1(a). As an example, we use air quality extrapolation [50] illustrated in Figure 1(b). We utilize the air quality index from context sensors to extrapolate the indexes at the target locations, taking into account covariates such as weather conditions that can also impact air quality.

To achieve our goal, one pivotal property that needs to be considered is spatio-temporal correlations, i.e., the spatial dependencies within the graph and the temporal dependencies along the time axis. To capture these correlations, Neural Networks (NNs), especially Spatio-Temporal Graph Neural Networks (STGNNs), nowadays have become a favorite due to their tempting learning competence [12, 46]. However, NNs have two main limitations: (i) *They lack the sought-after ability to estimate uncertainties.* Sensors often produce signals with different levels of ambiguity, such as noisy or missing observations due to unpredictable factors such as network disruptions. Explicitly capturing these uncertainties has proved to be a boon for making reliable decisions [42, 44]. However, most NNs are deterministic and unable to account for these uncertainties. (ii) *Their ability to generalize to new data is limited.* NNs require a large amount of training data to parameterize a function. Nevertheless, their reliance on parameters hinders their adaptability in unpredictable real-world environments, as the model needs to be retrained whenever the environment changes. Additionally, their sensitivity to hyperparameters demands a hyperparameter search to attain optimal performance.

The limitations of NNs have led researchers to draw inspiration from probability models, with Gaussian Processes (GPs) being one potential approach [35]. GPs define a stochastic process in which the spatio-temporal relations are modeled by various kernels [32]. Their Bayesian principle and non-parametric nature enable them to handle uncertainty and generalize well to a wide range of functions [29]. However, GPs suffer from limited expressivity of kernels, which can be disadvantageous. To address these issues, Neural Processes (NPs) [9] have emerged as a promising approach. NPs learn to construct stochastic processes by parameterizing neural networks in which an aggregator is introduced to integrate contexts. By combining the strengths of both NNs and GPs, NPs offer a compelling framework for spatio-temporal extrapolation.

Unfortunately, NPs cannot be applied directly to spatio-temporal graph data due to the following factors: (i) *They struggle to learn temporal dependencies effectively.* Existing NPs [33, 38] utilize latent state transitions to capture temporal relations recurrently. However, transitions tend to only focus on learned latent variables, disregarding the contextual information at later recurrent steps. This phenomenon, known as transition collapse, can impede learning over long sequences [38]. (ii) *They are incapable of modeling spatial relationships in a graph.* Existing NPs' aggregation operations [11, 16, 41] lack the ability to model complex spatial relations defined in a graph. In addition, their deterministic nature makes them suboptimal for aggregating data with varying levels of ambiguity, such as missing values and noise in Figure 1(a).

To tackle these challenges, we propose Spatio-Temporal Graph Neural Processes (STGNP) for spatio-temporal extrapolation over graphs. STGNP has two stages: the first stage leverages a deterministic network to learn spatio-temporal representations of nodes. Instead of relying on a recurrent structure, the temporal dynamics are modeled by stacking convolution layers in a bottom-up way [1], while the spatial relations are captured by cross-set graph neural networks. In the second stage, we employ state transitions to aggregate latent variables of target nodes following a top-down manner. Here, the transition assumes horizontal time independence and incorporates long-range temporal evolution from the upper

layers. As the number of transitions only relates to the number of stacked layers, much smaller than the sequential length, our model naturally exhibits resistance to transition collapse.

For the aggregator in each transition, we claim that different context nodes have different levels of importance. Motivated by [41], we propose Graph Bayesian Aggregation (GBA) that directly aggregates distributions over latent variables regulated by the graph. Intuitively, a context node contributes less to the latent distribution if it is far from the target location or exhibits a high degree of ambiguity recognized by the module. This design explicitly considers graph structure into NP's aggregator and enhances the model's capability to capture node uncertainties.

In summary, our main contributions are summarized as follows:

- We propose Spatio-Temporal Graph Neural Processes. To the best of our knowledge, this is the first work that generalizes NPs to spatio-temporal graph modeling. STGNP captures uncertainties explicitly and can generalize to different functions robustly, which is a major advantage over NNs approaches. Additionally, STGNP is able to learn temporal relations and graph data effectively, which sets it apart from classical NPs models.
- We introduce Graph Bayesian Aggregation, a Bayesian method for aggregating context nodes, which allows the aggregator to model graph structure and uncertainties for context nodes.
- We conduct comprehensive experiments to evaluate the performance and properties of STGNP. Our results demonstrate that STGNP outperforms state-of-the-art baselines by a significant margin and exhibits compelling properties, such as uncertainty estimates, high generalizability, and robustness to noisy data.

## 2 PRELIMINARIES

We first define concepts and notations of spatial-temporal data on graphs. Then, we introduce the basic concepts of neural processes.

### 2.1 Definitions and Notations

**Definition 1 (Graph).** We represent nodes as a graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ , where  $\mathcal{V}$  is the vertex set and  $\mathcal{E}$  is the edge set defining the weights between nodes. The  $K$ -hop neighborhood of a node  $v \in \mathcal{V}$  denoted by  $\mathcal{N}_K(v)$  is the set of nodes that are reachable from  $v$  with  $K$  steps. Based on  $\mathcal{E}$  and  $K$ , an  $K$ -hop adjacency matrix  $A^K$  is derived to measure the non-Euclidean distances between connected neighbors.

**Definition 2 (Spatio-Temporal Data).** Observed signals are retrieved from each node in the graph. We use  $Y_i = (y_{i,1}, \dots, y_{i,t}, \dots, y_{i,T}) \in \mathbb{R}^{T \times d_y}$  to denote data of node  $i$  that is measured over a time window  $T$ , where  $d_y$  is the number of features.  $Y = (Y_1, \dots, Y_n, \dots, Y_N) \in \mathbb{R}^{N \times T \times d_y}$  is denoted as a signal tensor of all nodes over the window  $T$ , where  $N$  is the total number of observable nodes in the graph.

**Definition 3 (Exogenous Covariates).** Exogenous covariates benefit the learning process as they usually have notable correlations with node data. These covariates are readily available from different sources. For instance, weather conditions can affect air pollutant data and they are collected from weather stations. We denote these factors as a tensor  $X \in \mathbb{R}^{N \times T \times d_x}$  and consider them explicitly.

## 2.2 Neural Processes

Neural Processes [9] construct stochastic processes that map  $x \in \mathbb{R}^{d_x}$  in an input domain to  $y \in \mathbb{R}^{d_y}$  in an output space, conditioned on a context set  $C = \{(x_n, y_n)\}_{n=1}^N$  of observed input-output pairs. NPs follow the same principle as Gaussian Processes except the stochastic process is learned implicitly by neural networks. Specifically, NPs define a conditional latent variable framework, where the distribution of a latent variable  $z$  is described by a learned conditional prior  $p(z|C)$  from the context set. Then, with inputs of target variables  $X_D = \{x_m\}_{m=1}^M$  in a target set  $D$ , a likelihood module  $p(Y_D|X_D, z)$  is trained to predict the corresponding output variables  $Y_D$ . The following posterior predictive likelihood formulates the generative process of NPs:

$$p(Y_D|X_D, C) = \int p(Y_D|X_D, z)p(z|C)dz. \quad (1)$$

In practice, NPs assume the target variables are independent, decomposing the likelihood such that  $p(Y_D|X_D, z)$  is factorized as  $\prod_{m=1}^M p(y_m|x_m, z)$ , where  $M = |D|$ . NPs organize a meta-learning framework where each pair of  $\{C, D\}$  constructs its own stochastic process, making it less parametric dependent with strong generalizability. Note that conditions  $C$  in the prior should be aggregated by a permutation-invariant function (e.g., mean, attention) to define a stochastic process, according to Kolmogorov Extension Theorem [30]. As the marginalization of the latent variable  $z$  is normally intractable, the model is usually trained either by Monte-Carlo (MC) sampling to estimate Equation 1 directly [8] or by variational approximation of maximizing the evidence lower bound (ELBO) [9]:

$$\log p(Y_D|X_D, C) \geq \mathbb{E}_{q(z|C \cup D)} \left[ \sum_{m=1}^M \log \frac{p(y_m|x_m, z)p(z|C)}{q(z|C \cup D)} \right], \quad (2)$$

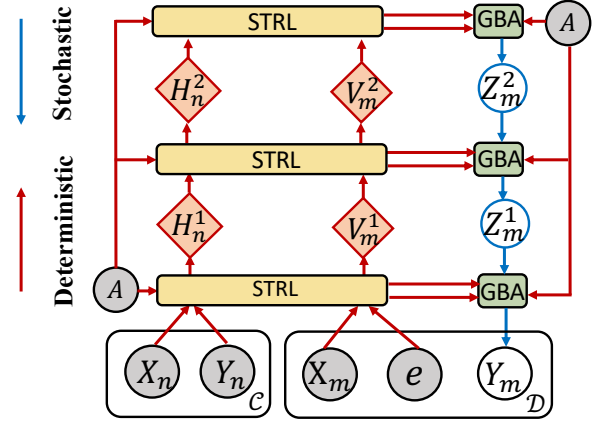
where  $q(z|C \cup D)$  and  $p(y_m|x_m, z)$  are the approximated posterior and the likelihood learned by neural networks. As the true conditional prior  $p(z|C)$  in the numerator is also intractable, the same module  $q(\cdot)$  is employed to approximate  $p(z|C) \approx q(z|C)$ .

## 3 METHODOLOGY

In this section, we propose STGNP, a neural latent variable model to enhance spatio-temporal extrapolation. As its graphical model illustrates in Figure 2, the key pipeline is to learn deterministic representations (STRL) and stochastic latent variables (GBA) in two stages. We first introduce the problem of spatio-temporal extrapolation. Then, we describe the deterministic stage for learning spatio-temporal representations and derive Graph Bayesian Aggregation to aggregate contexts in the stochastic stage. Finally, we introduce the generative process and the optimization procedure. Note that as target nodes are independent, we only discuss a single target node  $m$  in the following sections for brevity.

### 3.1 Problem Statement

We formulate spatio-temporal extrapolation in the NPs framework and first define the context set  $C$  containing nodes with exogenous covariates and observed data  $\{(X_n, Y_n)\}_{n=1}^N \in \mathbb{R}^{N \times T \times (d_x + d_y)}$ . Our goal is to learn a posterior predictive distribution  $p(Y_D|X_D, C, A)$  to generate  $Y_D \in \mathbb{R}^{M \times T \times d_y}$  over the same time period in the



**Figure 2: Graphical model with three layers for illustration.**  $V_m^l, Z_m^l, H_n^l \in \mathbb{R}^{T \times d_l}$  are deterministic representations, latent variables of a target node  $m$  and representations of a context node.  $e \in \mathbb{R}^{d_0}$  is the learnable target token. The shadow circle, STRL, GBA denote an observed variable, a spatio-temporal representation learning, and Graph Bayesian Aggregation.

target set  $D$  where  $M$  is the total number of target nodes, given the covariates  $X_D$ , the context set  $C$ , and the adjacency matrix  $A$ . In this paper, we use subscript  $m$  and  $n$  to index target and context nodes respectively, and adopt the terms location, node, and sensor interchangeably.

### 3.2 Spatio-Temporal Representation Learning

The deterministic stage has three building blocks to capture spatial and temporal correlations: learnable target token, dilated causal convolution, and cross-set graph convolution. We introduce them individually and then illustrate the overall learning framework.

**Learnable Target Token.** Our network takes sensor data and covariates as inputs; however, the data  $Y_m$  of a target node is unknown. Existing methods typically preprocess it by either filling it with zeros [2, 45] or employing linear interpolation to estimate its values [14]. However, using zero to represent target variables can be confusing, as it may be interpreted as the semantic zero within the dataset. Moreover, interpolation incurs large errors, which also hampers model performance. Inspired by Masked AutoEncoder [13], we leverage a shared learnable token  $e \in \mathbb{R}^{d_0}$  as embeddings for target nodes, while using an embedding layer with the parameter  $W \in \mathbb{R}^{d_y \times d_0}$  to get embeddings for context nodes. The token is optimized by the network to identify a position in the feature space that represents the target node, which avoids inferior preprocessing.

**Cross-Set Graph Convolution Layer.** Graph convolution is a seminal operation to learn spatial relations in a graph structure. Existing GCNs methods typically treat dependencies over all nodes equally [14, 45]. However, in our task, relations across the target and context sets take priority due to their direct influence on the target nodes. Based on this insight, we argue that disregarding internal relations within the two sets does not adversely affect performance and propose cross-set graph convolution (CSGCN), in which only



with its variance  $\sigma_{R_n^l}^2$ , which guarantees valid Gaussian conditioning during inference [41]:

$$p(R_n^l | Z_m^l, A_{m,n}) = \mathcal{N}(R_n^l | A_{m,n} Z_m^l, \text{diag}(\sigma_{R_n^l}^2)), \quad (7)$$

$$(R_n^l, \sigma_{R_n^l}^2) = \text{Enc}_R^l(H_n^l),$$

where  $R_n$  and  $\sigma_{R_n}^2$  are parameterized by  $\text{Enc}_R^l(\cdot)$ . The Gaussian assumption avoids an intractable computation of the marginal likelihood of the posterior's denominator. In fact, we can calculate it easily by Gaussian conditioning in a closed-form solution. (see proof in Appendix B):

$$\bar{\sigma}_{Z_m^l}^2 = \left[ (\sigma_{Z_m^l}^2)^{-2} + \sum_{n \in \mathcal{N}_1^c(m)} (\sigma_{R_n^l}^2 / A_{m,n})^{-2} \right]^{-1}, \quad (8)$$

$$\bar{\mu}_{Z_m^l} = \bar{\sigma}_{Z_m^l}^2 \left( \mu_{Z_m^l} / \sigma_{Z_m^l}^2 + \sum_{n \in \mathcal{N}_1^c(m)} A_{m,n} R_n / \sigma_{R_n^l}^2 \right), \quad (9)$$

where  $\bar{\sigma}_{Z_m^l}^2$  and  $\bar{\mu}_{Z_m^l}$  are updated parameters and the operations are conducted in an element-wise manner. With factorization, the conditioning is efficient to compute, avoiding costly matrix inversion. In addition, all the calculations are differentiable so that GBA can be optimized in an end-to-end way by stochastic gradient descent.

There are significant implications behind the aggregation. First, it incorporates the graph structure into the model by applying a linear transformation through the adjacency matrix, which is equivalent to GCNs when neglecting uncertainty terms. This suggests that GBA has similar learning abilities to GCNs. Second, the aggregation takes the uncertainties of nodes into consideration, which is a compelling property against previous methods. From the equations, the contribution of a context node is determined by its learned observation  $R_n^l$ , the variance  $\sigma_{R_n^l}^2$ , and the distance weight  $A_{m,n}$ . Specifically, Equation 8 gives a reasonable assumption that a context node located at a greater distance would provide less confident information. Additionally, Equation 9 implicates a node's contribution diminishes when its associated variance  $\sigma_{R_n^l}^2$  is large, signifying higher ambiguity. This theoretically guarantees the model's robustness when dealing with noisy data.

### 3.4 Generative Process

The target latent variable  $Z_m^l$  depends on its representation  $V_m^l$  and those of the context nodes  $H^l$ . The longer-range temporal dependencies are transited by conditioning  $Z_m^l$  on  $Z_m^{l+1}$ , forming a vertical time hierarchy. In practice, given  $V_m^l$  and a sample from  $p(Z_m^{l+1})$ , the network  $\text{Enc}_Z^l(Z_m^{l+1}, V_m^l)$  first learns a prior  $p(Z_m^l)$  over the target node in Equation 6. Then, the deterministic representations of context nodes are adopted to learn their latent observations by  $\text{Enc}_R^l(H_n^l)$  in Equation 7. Next, parameters of  $p(Z_m^l)$  are updated according to Equation 8 and 9. After the bottom layer  $l = 1$ , a likelihood model concatenates samples  $Z_m = (Z_m^1, \dots, Z_m^L)$  from all layers and the target node's exogenous covariates  $X_m$  to predict its extrapolations  $Y_m$ . Formally, the generative process of STGNP is summarized as:

$$p(Y_m, Z_m | X_m, C, A) = p(Y_m | X_m, Z_m) \prod_{l=1}^L p(Z_m^l | Z_m^{l+1}, V_m^l, H^l, A), \quad (10)$$

where the first term is a likelihood; the second term denotes a conditional prior aggregated through the GBA. Note that at the top layer  $L$ ,  $Z_m^{L+1} = \mathbf{0}$  and the likelihood is assumed to be a factorized Gaussian distribution.

### 3.5 Inference and Optimization

Typically, closed-form solutions for non-linear transitions and likelihood do not exist; thus we train the model through variational approximation. The approximated posterior  $q(Z_m | C \cup \mathcal{D}, A)$  has the same structure as the conditional prior but takes target node data  $Y_m$  as inputs. Then the deterministic and stochastic modules can be optimized together by the evidence lower-bound (ELBO):

$$\log p(Y_m | X_m, C, A) \geq \mathbb{E}_{q(Z_m)} [\log p(Y_m | X_m, Z_m)] - \mathbb{KL}(q(Z_m | C \cup \mathcal{D}, A) || p(Z_m | X_m, C, A)). \quad (11)$$

Given the hierarchical structure of Equation 10, the Kullback–Leibler divergence term  $\mathbb{KL}$  can be further decomposed as:

$$\mathbb{KL}(\cdot || \cdot) = \sum_{l=1}^L \mathbb{E}_{q(Z_m^{l+1})} [\mathbb{KL}(q(Z_m^l | Z_m^{l+1}, V_m^l, H^l, A) || p(Z_m^l | Z_m^{l+1}, V_m^l, H^l, A))], \quad (12)$$

where unlike using the learned token,  $V_m^{l+1}$  is the feature embeddings of  $Y_m$ . Following [9], we use the same variational module to approximate the conditional prior so that  $p(\cdot) = q(\cdot)$  in Equation 12. During optimization, ELBO can be minimized using stochastic gradient descent with the reparameterization trick [19].

## 4 EXPERIMENTS

To evaluate the performance and properties of STGNP, we conducted experiments to answer the following questions:

- **Q1:** How does STGNP perform on real-world datasets compared to other baselines?
- **Q2:** What is the quality of the uncertainty estimates?
- **Q3:** What is the effect of each component in our model, e.g., the learnable token, CSGCN, DCCConv and GBA.
- **Q4:** Is STGNP resistant to different missing ratios in datasets?
- **Q5:** Is it sensitive to hyperparameters and prone to overfitting?
- **Q6:** Does STGNP show robust generalizability when the domain of sensors changes?

### 4.1 Experimental Setup

**4.1.1 Dataset Descriptions.** We carry out experiments on three real-world spatio-temporal datasets.

- **Beijing** [51]: Beijing contains air quality indexes (AQI) from 36 stations and district-level meteorological attributes. Following [5, 12], we aim to extrapolate the AQI of PM2.5, PM10, and NO2, using meteorological attributes such as temperature, humidity, pressure, wind speed, direction, and weather as covariates.
- **London**<sup>1</sup>: We adopt London to evaluate the performance on different domains to answer **Q6**. It collects signals from 24 stations. Note that some baselines use non-publicly available covariates like POIs in the above two datasets. To ensure a fair comparison, we do not utilize them for all models in this work.

<sup>1</sup>[https://www.biendata.xyz/competition/kdd\\_2018/](https://www.biendata.xyz/competition/kdd_2018/)



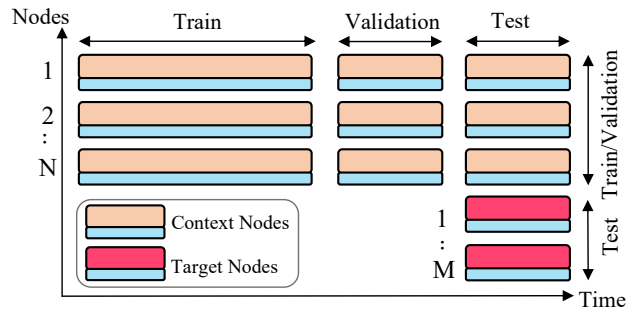
- **UrbanWater** [22, 26, 27]: The urban water quality data comes from 15 monitoring stations in Shenzhen, which collects 3 water measures of water quality: residual chlorine (RC), turbidity (TU), and power of hydrogen (pH). Following [26], we extrapolate RC given exogenous covariates TU and pH.

**4.1.2 Baselines.** We consider eight baselines which can be categorized into three classes.

- **Statistical models:** We utilize KNN, IDW [28], RF [7], and ANCL [32]. ANCL is a GPs-based framework that designs specialized kernels for different data attributes.
- **Neural Network methods:** We take ADAIN [5] and MCAM [12] as NNs baselines. ADAIN uses MLP and RNN layers for static and dynamic data, followed by an attention mechanism to aggregate features, whereas MCAM introduces multi-channel attention blocks for static and dynamic correlations.
- **Neural Processes approaches:** We modify SNP [38] and name the variant as SGNP. As SNP cannot deal with graph data, we add a cross-set graph network at each recurrent step before the aggregation. SGANP is an advanced version of SGNP modified from [33], which utilizes attention mechanisms as the aggregator.

Our model, STGNP, is based on a probabilistic framework and is a general method that is not tailored to specific tasks. It can also be applied in situations where exogenous covariates are not available by simply removing the corresponding causal convolution blocks. In contrast, ANCL relies on periodic and categorical kernels and MCAM utilizes horizontal and vertical wind speeds that are specific to the air quality task. SGNP and SGANP have a similar NP framework to ours, but we abandon the recurrent structure and propose GBA to aggregate nodes under a Bayesian framework.

**4.1.3 Evaluation Strategy and Hyperparameters.** In Figure 5, we illustrate the evaluation strategy for experiments. As we lack fine-grained data for the areas, we manually set aside 30% of existing stations in the dataset as target nodes for reporting performance. The remaining 70% is used for training purposes, ensuring a ratio of  $3N = 7M$ . This ensures that target nodes will not involve in the training phase. The data is sequentially divided into three segments: an 80% training set, a 10% validation set, and a 10% test set. During training and validation, we use  $N - 3$  nodes to extrapolate the remaining 3 randomly selected nodes. We report extrapolation metrics of MAE, MSE, and MAPE of the target stations. Since our



**Figure 5: Evaluation strategy and dataset division.** During testing, we use context nodes to extrapolate target nodes.

model, SGNP, and SGANP contain stochastic modules, we adopt the mean of distributions directly, instead of sampling, to report the results. The time window  $T$  is 24 and the adjacency matrix is constructed based on the locations of stations, normalized by a thresholded Gaussian kernel [37]. It is worth noting that, in the case of NNs, data processing is necessary to handle missing values in the dataset. For this reason, we employ linear interpolation to fill in the missing values for ADAIN and MCAM, while for the other baselines, missing values are left as zero. Hyperparameters are the same on all datasets for STGNP. The deterministic learning stage has 3 layers with a kernel size  $k = 3$  and channel numbers [16, 32, 64]. The stochastic stage is a 3-layer  $1 \times 1$  convolutions module and the channels of latent variables are [16, 32, 64]. The likelihood function is a 3-layer  $1 \times 1$  convolution with 128 channels in each layer. STGNP, ADAIN, MCAM, SGNP, and SGANP are implemented with PyTorch and trained on an NVIDIA A100 GPU. We repeat each experiment 5 times and report the average and variance of the metrics. Our implementations<sup>2</sup> are publicly available.

## 4.2 Overall Performance

To answer **Q1**, we report the overall results of baselines over Beijing and Water datasets, as shown in Table 1. Note that ANCL and MCAM require meteorological features, so they cannot be applied to the Water dataset. From the table, we see that STGNP consistently outperforms other baseline models with a notable margin, with the lowest errors on all datasets. Moreover, we have the following observations. First, the GPs model ANCL outperforms the other statistical models, indicating that GPs have the essential ability to capture complex dependencies with dedicated kernels. Second, NNs models surpass the above methods on all datasets because of their strong learning capability. Third, we find that SGNP and SGANP cannot consistently outperform NNs, possibly due to the transition collapse of the recurrent structure which may cause them to ignore the input contexts. Comparing these two models, although the attention aggregator performs better than the mean aggregator on tasks like computer vision, this is not always the case for graph data. This is because the data is constrained by a graph, which should be considered explicitly. Our STGNP has the best results due to the use of causal convolutions to alleviate transition collapse and the GBA to take into fact node uncertainties.

Another interesting discovery is that our model, SGNP, SGANP, and ANCL exhibit lower metric variances compared to other NN models. Evidently, for PM10 concentrations with large extrapolation errors, the metric variances for these models are small whereas for ADAIN and MCAM, their MAE and RMSE reach high values of 1.3, 3.4, and 3.2, 5.5, separately. The primary reason for the models with low variances is their ability to model data in a probabilistic manner and their awareness of uncertainty. This characteristic enables the models to be robust to parameter initialization and reduces the chance of getting trapped in local minima.

We also visualize extrapolation results in Figure 6 which illustrates the PM2.5 extrapolations of the five best models on the Beijing dataset, together with the ground truth. We observe that our model produces the most accurate results toward the ground truth data. This is particularly evident during the sudden change in the data

<sup>2</sup><https://github.com/hjf1997/STGNP>

**Table 1: Performances of STGNP and the baselines on two datasets. We denote the metric variance as  $\Delta_{var} = 0.1 \times var$ .**

Model	Beijing-PM2.5			Beijing-PM10			Beijing-NO2			Water-RC		
	MAE	RMSE	MAPE	MAE	RMSE	MAPE	MAE	RMSE	MAPE	MAE	RMSE	MAPE
KNN	28.08 $\pm\Delta_0$	39.87 $\pm\Delta_0$	0.62	61.11 $\pm\Delta_0$	99.20 $\pm\Delta_0$	0.61	24.18 $\pm\Delta_0$	31.36 $\pm\Delta_0$	0.95	0.21 $\pm\Delta_0$	0.25 $\pm\Delta_0$	0.48
IDW	39.11 $\pm\Delta_0$	48.90 $\pm\Delta_0$	0.73	72.72 $\pm\Delta_0$	116.28 $\pm\Delta_0$	0.69	24.15 $\pm\Delta_0$	29.76 $\pm\Delta_0$	0.96	0.16 $\pm\Delta_0$	0.21 $\pm\Delta_0$	0.31
RF	24.20 $\pm\Delta_0$	35.31 $\pm\Delta_0$	0.54	48.35 $\pm\Delta_0$	80.75 $\pm\Delta_0$	0.51	23.49 $\pm\Delta_0$	30.33 $\pm\Delta_0$	0.92	0.17 $\pm\Delta_0$	0.22 $\pm\Delta_0$	0.35
ANCL	19.72 $\pm\Delta_5$	30.87 $\pm\Delta_8$	0.44	32.43 $\pm\Delta_4$	53.15 $\pm\Delta_6$	0.31	20.95 $\pm\Delta_3$	26.50 $\pm\Delta_6$	0.79	—	—	—
ADAIN	16.81 $\pm\Delta_{16}$	27.00 $\pm\Delta_{26}$	0.32	31.25 $\pm\Delta_{13}$	55.08 $\pm\Delta_{34}$	0.28	16.86 $\pm\Delta_4$	22.87 $\pm\Delta_{11}$	0.54	0.15 $\pm\Delta_0$	0.19 $\pm\Delta_0$	0.27
MCAM	16.40 $\pm\Delta_{10}$	26.47 $\pm\Delta_{11}$	0.34	32.17 $\pm\Delta_{32}$	56.41 $\pm\Delta_{55}$	0.30	17.89 $\pm\Delta_1$	23.89 $\pm\Delta_1$	0.63	—	—	—
SGNP	17.82 $\pm\Delta_1$	28.51 $\pm\Delta_9$	0.37	33.76 $\pm\Delta_2$	63.96 $\pm\Delta_9$	0.29	17.26 $\pm\Delta_1$	22.74 $\pm\Delta_0$	0.59	0.14 $\pm\Delta_0$	0.17 $\pm\Delta_0$	0.26
SGANP	17.06 $\pm\Delta_2$	26.74 $\pm\Delta_3$	0.33	31.52 $\pm\Delta_3$	55.96 $\pm\Delta_6$	0.28	17.42 $\pm\Delta_2$	23.01 $\pm\Delta_0$	0.62	0.15 $\pm\Delta_0$	0.18 $\pm\Delta_0$	0.26
STGNP	<b>14.75<math>\pm\Delta_3</math></b>	<b>25.20<math>\pm\Delta_4</math></b>	<b>0.28</b>	<b>27.82<math>\pm\Delta_3</math></b>	<b>49.20<math>\pm\Delta_7</math></b>	<b>0.26</b>	<b>15.37<math>\pm\Delta_1</math></b>	<b>21.98<math>\pm\Delta_3</math></b>	<b>0.45</b>	<b>0.11<math>\pm\Delta_0</math></b>	<b>0.16<math>\pm\Delta_0</math></b>	<b>0.23</b>

after 9 AM on March 29, as highlighted in the red box. This demonstrates the effectiveness of our model in capturing both spatial information from surrounding sensors and temporal dependencies.

### 4.3 Uncertainty Estimates Analysis

One of the key benefits of our STGNP model is its capability to provide high-quality uncertainty estimates. To answer Q2, we first evaluate the qualitative results to show that they provide valuable information. In Figure 6, our model consistently has more accurate extrapolations compared with baselines. However, at 3–6 AM, March 28 shown in the blue box, all methods fail to generate precise data. Significantly, our model renders higher uncertainties (orange shadow), indicating possible inaccurate extrapolations. This ability to accurately estimate uncertainty can be extremely useful in practical decision-making. For instance, if the uncertainty for a location is consistently high, researchers could use this information to prioritize the deployment of sensors in order to achieve more accurate data analysis with minimal expenditure.

Next, to evaluate the quality of our estimates quantitatively, we follow the approach used in [4] and examine the number of extrapolations that fall in the uncertainty intervals. Assume a Gaussian likelihood with variance  $\sigma^2$  to suggest uncertainty, the intervals  $1\sigma$ ,  $2\sigma$ ,  $3\sigma$  centered at the predicted data cover  $\approx 68.3\%$ ,  $95.5\%$ ,  $99.7\%$  of the probability density. We posit that a better model should have a higher proportion of points falling in the intervals. Table 2 reports

the statistical results of the proportions for ANCL, SGNP, SGANP, and STGNP. It shows that all NPs methods have superior performance compared to ANCL and that our STGNP outperforms SGNP and SGANP in most metrics. This is likely due to the fact that our model also considers uncertainties in context nodes through GBA.

**Table 2: Proportions (%) of data falling in the intervals of  $1\sigma - 3\sigma$ , where  $\sigma$  is the standard deviation of the likelihood.**


Model	Beijing-PM2.5	Beijing-PM10	Beijing-NO2	Water-RC
	$1\sigma / 2\sigma / 3\sigma$	$1\sigma / 2\sigma / 3\sigma$	$1\sigma / 2\sigma / 3\sigma$	$1\sigma / 2\sigma / 3\sigma$
ANCL	44 / 67 / 86	16 / 30 / 49	35 / 70 / 88	—
SGNP	71 / 92 / 96	72 / 91 / 97	71 / <b>95 / 98</b>	62 / 90 / 95
SGANP	73 / 92 / <b>97</b>	78 / 93 / 97	69 / 91 / 96	61 / 92 / 97
STGNP	<b>77 / 92 / 98</b>	<b>85 / 96 / 98</b>	<b>76 / 94 / 98</b>	<b>66 / 94 / 99</b>

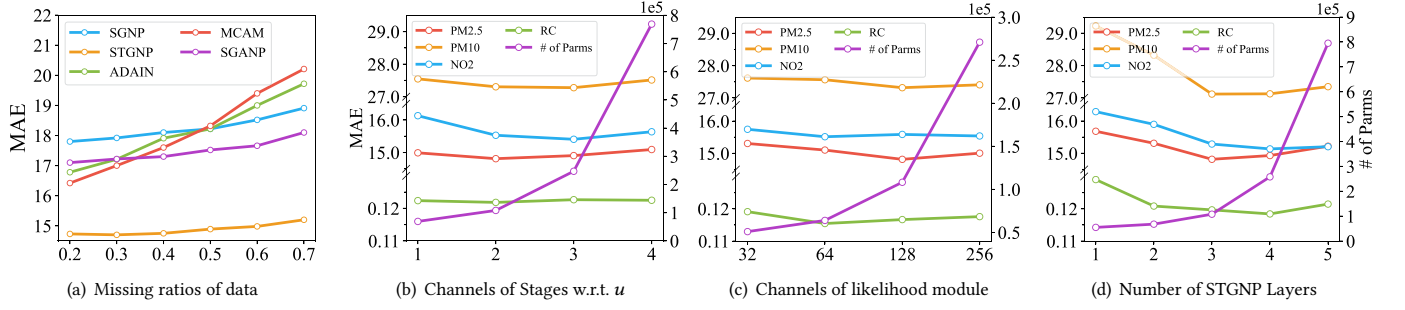
### 4.4 Ablation Study

To assess the contribution of each component to the overall performance of our model and answer Q3, we conduct ablation studies and present the results in Figure 8. In each study, we only modify the corresponding part while leaving other settings unchanged.

**Effect of learnable token:** We remove the learned token and use linear interpolation to preprocess target nodes (w/o TK). The results unveil that the token has a positive impact on the model’s performance. We believe this is because using simple interpolation to represent all signals of the target nodes leads to significant errors. Mingled with context nodes’ signals, this hobbles the model’s learning ability, and even the uncertainty estimates struggle to rectify them. In contrast, the learned token identifies a suitable position in the feature space that accurately indicates the target nodes, thereby avoiding this problem.

**Effect of CSGCN:** To evaluate the effectiveness of cross-set graph convolutional network, we compare it to three variants: (a) w/o CSGCN: removing CSGCN in our model; (b) r/p GCN: replacing CSGCN with a standard GCN. (c) r/p RGCN: replacing it with an advance relational GCN [34]. Our results show that removing the spatial learning module CSGCN leads to a degradation in performance, underscoring the importance of capturing spatial dependencies. Additionally, we find our CSGCN achieves performance on

**Figure 6: PM2.5 extrapolation results of the Beijing dataset during March 28 – 31, 2015.**



**Figure 7: Data missing ratio and hyperparameter study. In (b), channel numbers are  $[8u, 16u, 32u]$  for causal convolutions and latent variables with  $u$  ranging from 1 to 4.**

par with the standard GCN, which learns dependencies among all nodes in two sets, and even slightly better than the RGCN, which explicitly characterizes categorical relations among nodes. This observation suggests that modeling dependencies across two sets are sufficient to achieve satisfactory performance, thereby validating our insight.

**Effect of GBA:** Results from experiments where GBA is discarded (w/o GBA) or replaced with a max aggregator (r/p MAX), a mean aggregator (r/p MEAN), or an attention aggregator (r/p ATTN) highlight the critical role of GBA. Removing or replacing it leads to a significant decline in results on most datasets. This is likely because context nodes have varying levels of ambiguities. Without uncertainty estimates, the model has difficulty extracting valuable context information, which hampers the performance of the likelihood module.


**Effect of causal convolution:** We replace our temporal learning framework with an RNN structure (r/p RNN) and the results show a strong deterioration in the model’s performance. The outcome indicates transition collapse, which occurs in an RNN because of a large number of transitions, can have a substantial impact on the model’s capability. In contrast, our STGNP model effectively mitigates this issue by significantly reducing the number of transitions.

#### 4.5 Missing Ratio Study

Evaluating the model’s performance on data with missing values is important, as real-world sensors might lose signals due to unpredictable factors. To answer Q4, we train models using manually corrupted data by randomly replacing a ratio of data with zero. Following the same procedure, we preprocess them with linear interpolation for NNs, while leaving the missing values as zero for NPs models. The results, which illustrate the MAE performance of various baselines for ratios ranging from 0.2 to 0.7, are shown in Figure 7(a). We observe significant performance degradation in ADAIN and MCAM, which can be similarly attributed to large errors caused by interpolation. In contrast, STGNP and SGNP demonstrate better performance, with STGNP performing the best results. The key factor behind their success lies in the ability of their likelihood modules to capture uncertainties associated with the target nodes. This capability effectively mitigates the impact of missing data. Compellingly, the GBA module in our STGNP is able to further enhance its robustness by accurately modeling uncertainties in the signals of individual sensors.

#### 4.6 Hyperparameter Study

To answer Q5, we study the performance of STGNP under various hyperparameter settings. We first explore the impact of the number of channels in two stages. Following the default setting of the channel numbers  $[8u, 16u, 32u]$  where  $u = 2$ , we vary  $u$  from 1 to 4. From Figure 7(b), we observe that even with the drastically increasing number of learnable parameters (violet curves), the MAE metrics (green, blue, red, and yellow curves) keep stable. Then, we examine the effect of changing the number of channels of the likelihood module ranging from 32 to 256 and also observe stable performances, as shown in Figure 7(c). Finally, we investigate the effect of the number of layers ranging from 1 to 5. From Figure 7(d), the performances initially improve as it increases, but start to oscillate since 3 layers. This can be explained by the perceptive field of the model, where 1 and 2 layers correspond to fields of 3 and 7, respectively. These limited perceptive fields constrain the model’s performance. However, once the perceptive field becomes sufficiently large (i.e., at 3 layers and above), the model captures long enough temporal dependencies, leading to improved and more stable performance. These findings suggest STGNP is insensitive to hyperparameters, as long as the perceptive field is adequately large.



**Figure 8: Effectiveness of different modules in STGNP.**



## 4.7 Cross-Domain Evaluation

To assess the generalization ability of models and answer **Q6**, we train them on the Beijing dataset and evaluate results using the London dataset. Table 3 reports the performances of PM2.5, where the first and second columns of MAE mean the performance of training models using London data directly and training on Beijing while evaluating on London. We confirm that our STGNP has the best results in both training approaches. In addition, when measuring the performance discrepancy, we also discover that STGNP has the smallest performance degeneration, especially compared to NNs models. This is likely due to STGNP’s strong spatio-temporal learning capability as well as its non-parametric design inherited from GPs. This allows the model to learn high-level spatio-temporal principles in the feature space and to instantiate these dependencies during testing by constructing a stochastic process from the sets.

**Table 3: Performances of cross-domain evaluation.**

Model	Beijing→London-PM2.5		
	MAE	RMSE	MAPE
ADAIN	4.35/3.57 (↓0.78)	5.79/4.69 (↓1.10)	0.62/0.42 (↓0.20)
MCAM	4.38/3.78 (↓0.60)	5.38/4.80 (↓0.58)	0.68/0.49 (↓0.19)
SGNP	4.31/4.11 (↓0.20)	5.41/5.12 (↓0.29)	0.66/0.60 (↓0.06)
SGANP	3.85/3.61 (↓0.24)	5.15/4.82 (↓0.33)	0.52/0.45 (↓0.07)
STGNP	<b>3.20/3.03 (↓0.17)</b>	<b>4.30/4.26 (↓0.04)</b>	<b>0.40/0.33 (↓0.07)</b>

## 5 RELATED WORK

### 5.1 Spatio-Temporal Extrapolation

Spatio-temporal extrapolation is a task that involves predicting the state of a surrounding environment based on known information. Early works in this area use statistical machine learning methods, such as  $K$ -Nearest Neighbors (KNN) and Random Forest (RF) [7], to solve this problem. KNN approaches rely on linear dependencies between data points, while RF can capture non-linear dependencies. However, these methods only consider spatial relations and struggle to model more complex and dynamic correlations. Gaussian Processes [35] learn to construct stochastic processes, in which the spatio-temporal dependencies are captured by flexible kernels that are designed to handle different types of features [21]. For instance, Patel *et al.* [32] used periodical and Hamming distance kernels for temporal and categorical features respectively. However, these kernels are tailored to specific scenarios and the cubic time complexity limits their applicability. Some approaches view extrapolation as a tensor completion task [49]. Based on a low-rank matrix assumption, these methods capture the spatio-temporal patterns while being efficient to optimize. However, they are transductive, meaning that they can only infer the state of nodes involved in the training process. They are not able to generalize to new nodes that were not present in the training data. Recently, Neural Networks (NNs) have become the dominant paradigm. Cheng *et al.* [5] proposed an attention model for air quality inference, where the dynamic and static data are encoded by RNNs and MLPs, and the attention mechanism integrates the features of nodes. Han *et al.* [12] combined GCNs with a multi-channel attention module to improve

performance. However, NNs tend to struggle with learning uncertainties and can overfit on datasets with low amounts of data. Some NNs also view extrapolation as a kriging problem. For instance, Appleby *et al.* [2] interpolate a node given its neighbors with time information as extra features, and We *et al.* [45] learn the temporal dynamics, explicitly. However, the first only captures spatial relations and the second cannot incorporate exogenous covariates. On the contrary, our approach is able to achieve both goals.

### 5.2 Neural Processes Family

Neural Processes (NPs) combine the merits of both NNs and GPs [9], which possess strong learning ability and uncertainty estimates. Basically, it induces latent variables over the context set, forming a conditional latent variable model. Then, a likelihood module is utilized to generate the target outputs. Le *et al.* [20] found that NPs suffer an underfitting problem because of the incapable aggregation function (e.g., mean or sum). Then, Kim *et al.* [16] proposed Attentive Neural Processes (ANP) to calculate the importance within/across the context set and target set. Kim *et al.* [17] further proposed a stochastic attention mechanism for aggregation where the attention weights are inferred using Bayesian inference and Volpp *et al.* [41] introduced a stochastic aggregator to aggregate context variables directly. However, these works only consider the spatial domain and cannot handle graph data. Singh *et al.* [38] first focused on sequential data and proposed Sequential Neural Processes (SNP). With a latent state transition function from a variational recurrent neural network (VRNN) [6], it constructs a sequential stochastic process for a timeline. Then, Yoon *et al.* [47] introduced Recurrent Memory Reconstruction to compensate for the distribution shift in a sequence. However, these recurrent structures have the transition collapse problem, which can make it difficult to learn temporal relations over long sequences. In contrast, our method utilizes casual convolutions to alleviate the challenge.

## 6 CONCLUSIONS AND FUTURE WORK

We introduce Spatio-Temporal Graph Neural Processes, the first framework for spatio-temporal extrapolation in the Neural Processes family. Our model captures temporal relations and addresses the transition collapse problem using causal convolutions while effectively learning spatial dependencies using the cross-set graph network. The Graph Bayesian Aggregation aggregates context nodes in a way that takes into account their uncertainties and enhances the learning ability of NPs on graph data. Experimental results demonstrate the superiority of STGNP in terms of extrapolation accuracy, uncertainty estimates, robustness, and generalizability. In the future, an intriguing direction would be to explore STGNP for spatio-temporal forecasting, which is another fundamental task in the area. By aggregating historical representations, the model could provide predictions about future time steps.

## ACKNOWLEDGMENTS

This research is supported by Singapore Ministry of Education Academic Research Fund Tier 2 under MOE’s official grant number T2EP20221-0023. It is also supported by Guangzhou Municipal Science and Technology Project 2023A03J0011.

## REFERENCES

- [1] Emre Aksan and Otmar Hilliges. 2019. STCN: Stochastic Temporal Convolutional Networks. In *International Conference on Learning Representations, ICLR*.
- [2] Gabriel Appleby, Linfeng Liu, and Li-Ping Liu. 2020. Kriging convolutional networks. In *The Thirty-Fourth Conference on Artificial Intelligence AAAI*, Vol. 34. 3187–3194.
- [3] Christopher M Bishop and Nasser M Nasrabadi. 2006. *Pattern recognition and machine learning*. Springer.
- [4] Di Chai, Leye Wang, and Qiang Yang. 2018. Bike flow prediction with multi-graph convolutional networks. In *International Conference on Advances in Geographic Information Systems, SIGSPATIAL*. 397–400.
- [5] Weiyu Cheng, Yanyan Shen, Yanmin Zhu, and Linpeng Huang. 2018. A neural attention model for urban air quality inference: Learning the weights of monitoring stations. In *AAAI*. 2151–2158.
- [6] Junyoung Chung, Kyle Kastner, Laurent Dinh, Kratarth Goel, Aaron C Courville, and Yoshua Bengio. 2015. A recurrent latent variable model for sequential data. *Annual Conference on Neural Information Processing Systems, NeurIPS* (2015), 2980–2988.
- [7] Khaled Fawagreh, Mohamed Medhat Gaber, and Eyad Elyan. 2014. Random forests: from early developments to recent advancements. *Systems Science & Control Engineering: An Open Access Journal* 2, 1 (2014), 602–609.
- [8] Andrew Foong, Wessel Bruinsma, Jonathan Gordon, Yann Dubois, James Requeima, and Richard Turner. 2020. Meta-learning stationary stochastic process prediction with convolutional neural processes. *NeurIPS* 33 (2020), 8284–8295.
- [9] Marta Garnelo, Jonathan Schwarz, Dan Rosenbaum, Fabio Viola, Danilo J Rezende, SM Eslami, and Yee Whye Teh. 2018. Neural processes. In *ICML Workshop*.
- [10] Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*. 249–256.
- [11] Jonathan Gordon, Wessel P. Bruinsma, Andrew Y. K. Foong, James Requeima, Yann Dubois, and Richard E. Turner. 2020. Convolutional Conditional Neural Processes. In *ICLR*.
- [12] Qilong Han, Dan Lu, and Rui Chen. 2021. Fine-Grained Air Quality Inference via Multi-Channel Attention Model. In *the International Joint Conference on Artificial Intelligence, IJCAI*. 2512–2518.
- [13] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. 2022. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 16000–16009.
- [14] Junfeng Hu, Yuxuan Liang, Zhencheng Fan, Yifang Yin, Ying Zhang, and Roger Zimmermann. 2021. Decoupling Long- and Short-Term Patterns in Spatiotemporal Inference. *arXiv preprint arXiv:2109.09506* (2021).
- [15] Guangyin Jin, Yuxuan Liang, Yuchen Fang, Jincui Huang, Junbo Zhang, and Yu Zheng. 2023. Spatio-Temporal Graph Neural Networks for Predictive Learning in Urban Computing: A Survey. *arXiv preprint arXiv:2303.14483* (2023).
- [16] Hyunjik Kim, Andriy Mnih, Jonathan Schwarz, Marta Garnelo, S. M. Ali Eslami, Dan Rosenbaum, Oriol Vinyals, and Yee Whye Teh. 2019. Attentive Neural Processes. In *ICLR*.
- [17] Mingyu Kim, Kyeongryeol Go, and Se-Young Yun. 2022. Neural Processes with Stochastic Attention: Paying more attention to the context dataset. In *ICLR*.
- [18] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *ICLR*.
- [19] Diederik P. Kingma and Max Welling. 2014. Auto-Encoding Variational Bayes. In *ICLR*.
- [20] Tuan Anh Le, Hyunjik Kim, Marta Garnelo, Dan Rosenbaum, Jonathan Schwarz, and Yee Whye Teh. 2018. Empirical evaluation of neural process objectives. In *NeurIPS workshop on Bayesian Deep Learning*. 71.
- [21] Naiqi Li, Wenjie Li, Jifeng Sun, Yinghua Gao, Yong Jiang, and Shu-Tao Xia. 2020. Stochastic deep gaussian processes over graphs. *NeurIPS* 33 (2020), 5875–5886.
- [22] Yuxuan Liang, Songyu Ke, Junbo Zhang, Xiuwen Yi, and Yu Zheng. 2018. Geo-MAN: Multi-Level Attention Networks for Geo-sensory Time Series Prediction. In *International Joint Conference on Artificial Intelligence*, Vol. 2018. 3428–3434.
- [23] Yuxuan Liang, Kun Ouyang, Lin Jing, Sijie Ruan, Ye Liu, Junbo Zhang, David S Rosenblum, and Yu Zheng. 2019. Urbanfm: Inferring Fine-Grained Urban Flows. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*. 3132–3142.
- [24] Yuxuan Liang, Yutong Xia, Songyu Ke, Yiwei Wang, Qingsong Wen, Junbo Zhang, Yu Zheng, and Roger Zimmermann. 2022. AirFormer: Predicting Nationwide Air Quality in China with Transformers. *arXiv preprint arXiv:2211.15979* (2022).
- [25] Haitao Lin, Zhangyang Gao, Yongjie Xu, Lirong Wu, Ling Li, and Stan Z. Li. 2022. Conditional Local Convolution for Spatio-Temporal Meteorological Forecasting. In *The Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI*. 7470–7478.
- [26] Ye Liu, Yuxuan Liang, Kun Ouyang, Shuming Liu, David Rosenblum, and Yu Zheng. 2020. Predicting urban water quality with ubiquitous data—a data-driven approach. *IEEE Transactions on Big Data* (2020).
- [27] Ye Liu, Yu Zheng, Yuxuan Liang, Shuming Liu, and David S Rosenblum. 2016. Urban Water Quality Prediction based on Multi-Task Multi-View Learning. In *Proceedings of the 25th international joint conference on artificial intelligence*.
- [28] George Y Lu and David W Wong. 2008. An adaptive inverse-distance weighting spatial interpolation technique. *Computers & geosciences* (2008), 1044–1055.
- [29] Jaakko Luttinen and Alexander Ilin. 2012. Efficient Gaussian process inference for short-scale spatio-temporal modeling. In *Artificial Intelligence and Statistics*. PMLR, 741–750.
- [30] Bernt Øksendal. 2003. Stochastic differential equations. *Stochastic differential equations* (2003), 65–84.
- [31] Zheyi Pan, Yuxuan Liang, Weifeng Wang, Yong Yu, Yu Zheng, and Junbo Zhang. 2019. Urban Traffic Prediction from Spatio-Temporal Data using Deep Meta Learning. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*. 1720–1730.
- [32] Zeel B. Patel, Palak Purohit, Harsh M. Patel, Shivam Sahni, and Nipun Batra. 2022. Accurate and Scalable Gaussian Processes for Fine-Grained Air Quality Inference. In *AAAI*. 12080–12088.
- [33] Shenghao Qin, Jiacheng Zhu, Jimmy Qin, Wenshuo Wang, and Ding Zhao. 2019. Recurrent attentive neural process for sequential data. *NeurIPS Workshop* (2019).
- [34] Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max Welling. 2018. Modeling relational data with graph convolutional networks. In *The Semantic Web: 15th International Conference, ESWC*.
- [35] Matthias Seeger. 2004. Gaussian processes for machine learning. *International journal of neural systems* 14, 02 (2004), 69–106.
- [36] Kangjia Shao, Yang Wang, Zhengyang Zhou, Xike Xie, and Guang Wang. 2021. TrajForesee: How limited detailed trajectories enhance large-scale sparse information to predict vehicle trajectories? In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*. IEEE, 2189–2194.
- [37] David I Shuman, Sunil K Narang, Pascal Frossard, Antonio Ortega, and Pierre Vandergheynst. 2013. The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. *IEEE signal processing magazine* (2013).
- [38] Gautam Singh, Jaesik Yoon, Youngsung Son, and Sungjin Ahn. 2019. Sequential Neural Processes. In *NeurIPS*. 10254–10264.
- [39] Yusuke Tashiro, Jiaming Song, Yang Song, and Stefano Ermon. 2021. CSDI: Conditional score-based diffusion models for probabilistic time series imputation. *Advances in Neural Information Processing Systems* 34 (2021), 24804–24816.
- [40] Aäron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew W. Senior, and Koray Kavukcuoglu. 2016. WaveNet: A Generative Model for Raw Audio. In *The Speech Synthesis Workshop, ISCA*. 125.
- [41] Michael Volpp, Fabian Flürenbrock, Lukas Grossberger, Christian Daniel, and Gerhard Neumann. 2020. Bayesian context aggregation for neural processes. In *ICLR*.
- [42] Bin Wang, Jie Lu, Zheng Yan, Huaishao Luo, Tianrui Li, Yu Zheng, and Guangquan Zhang. 2019. Deep uncertainty quantification: A machine learning approach for weather forecasting. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2087–2095.
- [43] Pengkun Wang, Chaochao Zhu, Xu Wang, Zhengyang Zhou, Guang Wang, and Yang Wang. 2022. Inferring intersection traffic patterns with sparse video surveillance information: An st-gan method. *IEEE Transactions on Vehicular Technology* 71, 9 (2022), 9840–9852.
- [44] Haomin Wen, Youfang Lin, Yutong Xia, Huaiyu Wan, Roger Zimmermann, and Yuxuan Liang. 2023. Diffstg: Probabilistic spatio-temporal graph forecasting with denoising diffusion models. *arXiv preprint arXiv:2301.13629* (2023).
- [45] Yuankai Wu, Dingyi Zhuang, Aurelie Labbe, and Lijun Sun. 2021. Inductive Graph Neural Networks for Spatiotemporal Kriging. In *AAAI*, Vol. 35. 4478–4485.
- [46] Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, and Chengqi Zhang. 2019. Graph WaveNet for Deep Spatial-Temporal Graph Modeling. In *IJCAI*. 1907–1913.
- [47] Jaesik Yoon, Gautam Singh, and Sungjin Ahn. 2020. Robustifying sequential neural processes. In *ICML*. 10861–10870.
- [48] Fisher Yu and Vladlen Koltun. 2016. Multi-Scale Context Aggregation by Dilated Convolutions. In *4th International Conference on Learning Representations, ICLR*.
- [49] Hsiang-Fu Yu, Nikhil Rao, and Inderjit S Dhillon. 2016. Temporal regularized matrix factorization for high-dimensional time series prediction. *Advances in neural information processing systems* 29 (2016).
- [50] Yu Zheng, Furui Liu, and Hsun-Ping Hsieh. 2013. U-air: When urban air quality inference meets big data. In *International Conference on Knowledge Discovery and Data Mining, SIGKDD*. 1436–1444.
- [51] Yu Zheng, Xiuwen Yi, Ming Li, Ruiyuan Li, Zhangqing Shan, Eric Chang, and Tianrui Li. 2015. Forecasting fine-grained air quality based on big data. In *International Conference on Knowledge Discovery and Data Mining, SIGKDD*. 2267–2276.

## A MATHEMATICAL NOTATION

We define the major mathematical notations in the paper in Table 4 for better understanding.

**Table 4: Major notations used in the paper.**

Notation	Dimension	Description
$N, M$	$\mathbb{R}^1$	number of context, target nodes
$T$	$\mathbb{R}^1$	time length of a sequence
$A_{m,n}$	$\mathbb{R}^1$	weight between node $m$ and $n$
$n, m$	$\mathbb{R}^1$	index of a context and target node
$d_l, d_x, d_y$	$\mathbb{R}^1$	feature dimensionalities
$\mathcal{C}$	$\mathbb{R}^{N \times T \times (d_x + d_y)}$	context set
$\mathcal{D}$	$\mathbb{R}^{M \times T \times (d_x + d_y)}$	target set
$H$	$\mathbb{R}^{M \times T \times \sum_{l=1}^L d_l}$	representations of context nodes
$X_n, Y_n$	$\mathbb{R}^{T \times (d_x + d_y)}$	covariates, data of a context node
$X_m, Y_m$	$\mathbb{R}^{T \times (d_x + d_y)}$	those of a target node
$H_n$	$\mathbb{R}^{T \times \sum_{l=1}^L d_l}$	representations of context node $n$
$V_m$	$\mathbb{R}^{T \times \sum_{l=1}^L d_l}$	representations of target node $m$
$R_n$	$\mathbb{R}^{T \times \sum_{l=1}^L d_l}$	latent observations of node $n$
$Z_m$	$\mathbb{R}^{T \times \sum_{l=1}^L d_l}$	latent variables of target node $m$

## B DERIVATION OF GRAPH BAYESIAN AGGREGATION

In this section, we give formal derivations of the proposed Graph Bayesian Aggregation. We first derive the general GBA without factorization or specific graph structure. We assume a latent prior  $Z$  over a target node (omitting subscript  $m$  for brevity). The latent observation functions of all context nodes are an independent linear transformation of  $Z$  following Gaussian distributions:

$$p(Z) = \mathcal{N}(Z | \mu, \Lambda^{-1}), \quad (13)$$

$$p(R_n | Z) = \mathcal{N}(R_n | A_n Z, L_n^{-1}) \quad \text{For } n \in [1, N], \quad (14)$$

where we use the precision matrix  $\Lambda$  and  $L$  for convenience;  $A_n$  is a transformation matrix, representing the graph structure. The logarithmic joint probability over  $Z$  and  $[R_1, \dots, R_N]$  is:

$$\begin{aligned} \ln p(Z, R_1, \dots, R_N) &= \ln p(Z) + \ln p(R_1 \dots R_N | Z) \\ &= -\frac{1}{2} (Z - \mu)^T \Lambda (Z - \mu) \\ &\quad - \frac{1}{2} \sum_{n=1}^N (R_n - A_n Z)^T L_n (R_n - A_n Z) + \text{Const}, \end{aligned} \quad (15)$$

where the second order terms can be decomposed as:

$$-\frac{1}{2} \begin{bmatrix} Z \\ R_1 \\ \vdots \\ R_N \end{bmatrix}^T \underbrace{\begin{bmatrix} \Lambda + \sum_{n=1}^N A_n^T L_n A_n & -A_1^T L_1 & \cdots & -A_N^T L_N \\ -L_1 A_1 & L_1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ -L_N A_N & 0 & \cdots & L_N \end{bmatrix}}_P \begin{bmatrix} Z \\ R_1 \\ \vdots \\ R_N \end{bmatrix}.$$

According to [3], matrix  $P$  above is the precision matrix of the joint distribution, where the covariance matrix can be calculated as:

$$\text{cov}[Z, R_1, \dots, R_N] = P^{-1}. \quad (16)$$

Next, the linear terms in Equation 15 can be decomposed as:

$$Z^T \Lambda \mu = [Z, R_1, \dots, R_N]^T [\Lambda \mu, 0, \dots, 0]^T. \quad (17)$$

Then from [3], the mean of the joint distribution is computed by:

$$\mathbb{E}[Z, R_1, \dots, R_N] = P^{-1} [\Lambda \mu, 0, \dots, 0]^T = [\mu, A_1 \mu, \dots, A_N \mu]^T. \quad (18)$$

The mean and covariance of the marginal distribution of  $p(R_1, \dots, R_N)$  can be extracted from Equation 17 and 16:

$$\mathbb{E}[R_1 \quad \cdots \quad R_N] = [A_1 \mu, \dots, A_N \mu]^T, \quad (19)$$

$$\text{cov}[R_1 \cdots R_N] = \begin{bmatrix} L_1^{-1} + A_1 \Lambda^{-1} A_1^T & & \\ & \ddots & \\ L_N^{-1} + A_N \Lambda^{-1} A_N^T & & \end{bmatrix}. \quad (20)$$

Finally, with  $p(Z, R_1, \dots, R_N)$ , and  $p(R_1, \dots, R_N)$ , we could obtain the probability of  $p(Z | R_1, \dots, R_N)$  through Gaussian conditioning, which is also a Gaussian:

$$\Sigma_{Z | R_1 \dots R_N} = \left( L + \sum_{n=1}^N A_n \Lambda A_n^T \right)^{-1}, \quad (21)$$

$$\mu_{Z | R_1 \dots R_N} = \Sigma_{Z | R_1 \dots R_N} \left( \Lambda \mu + \sum_{n=1}^N A_n^T L_n R_n \right). \quad (22)$$

In GBA, the Gaussian distributions assume to be factorized; thus  $\Lambda^{-1} = \text{diag}(\sigma_Z)$  and  $L_n^{-1} = \text{diag}(\sigma_{R_n})$ . The transformation  $A_n = I(a_n)$ , where  $a_n$  represents the distance between a target node and the context node  $n$ . Then, the Gaussian can be further weight as:

$$\bar{\sigma}_Z^2 = \left[ (\sigma_Z)^{-2} + \sum_{n=1}^N (\sigma_{R_n} / a_n)^{-2} \right]^{-1}, \quad (23)$$

$$\bar{\mu}_Z = \bar{\sigma}_Z^2 \left( \mu_Z / \sigma_Z^2 + \sum_{n=1}^N a_n R_n / \sigma_{R_n}^2 \right). \quad (24)$$

## C DERIVATION OF ELBO FOR STGNP

Here, we derive the evidence lower-bound (ELBO) for STGNP. For brevity, we still omit the subscript  $m$ .

$$\begin{aligned} \log p(Y | X, C, A) &= \log \mathbb{E}_{q(Z | C \cup \mathcal{D}, A)} \frac{p(Y, Z | X, C, A)}{q(Z | C \cup \mathcal{D}, A)} \\ &\geq \mathbb{E}_{q(\cdot)} \left[ \log \frac{p(Y | X, Z) p(Z | X, C, A)}{q(Z | C \cup \mathcal{D}, A)} \right] \\ &= \mathbb{E}_{q(\cdot)} \left[ \log p(Y | X, Z) - \log \frac{q(Z | C \cup \mathcal{D}, A)}{p(Z | X, C, A)} \right] \\ &= \mathbb{E}_{q(\cdot)} \left[ \log p(Y | X, Z) - \log \frac{\prod_{l=1}^L q(Z^l | Z^{l+1}, V^l, H^l, A)}{\prod_{l=1}^L q(Z^l | Z^{l+1}, V^l, H^l, A)} \right] \\ &= \mathbb{E}_{q(Z | C \cup \mathcal{D}, A)} [\log p(Y | X, Z)] - \\ &\quad \sum_{l=1}^L \mathbb{E}_{q(Z^{l+1})} \left[ \mathbb{KL} \left( q(Z^l | Z^{l+1}, V^l, H^l, A) \| q(Z^l | Z^{l+1}, V^l, H^l, A) \right) \right]. \end{aligned} \quad (25)$$

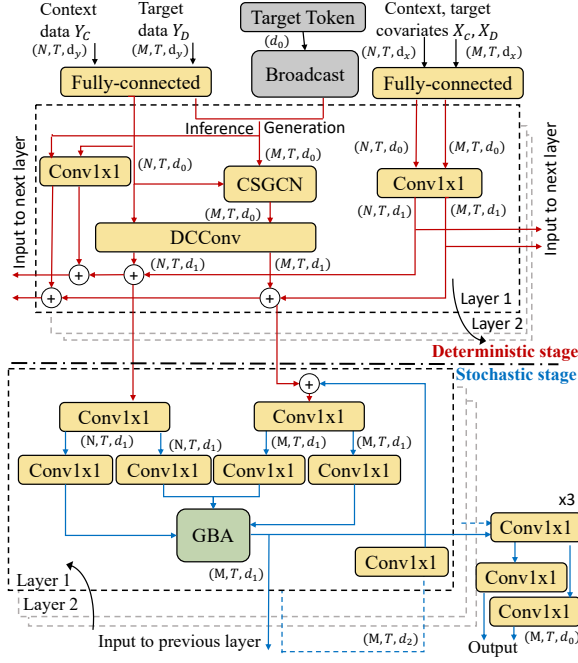


Figure 9: Detailed Architecture of STGNP.

## D EXPERIMENTAL DETAILS

### D.1 STGNP Architectures

Our STGNP has two major architectural components for deterministic and stochastic learning as shown in Figure 9:

- **Deterministic spatio-temporal stage:** The core component of this stage is the spatio-temporal learning module consisting of dilated causal convolutions and cross-set graph neural networks. Each layer contains a CSGCN to learn spatial dependencies and a DCCConv to model temporal relations of sensor data. A DCCConv has a kernel size of  $k = 3$  and the number of output channels is denoted by  $d$ . A CSGCN includes a multilayer perceptron layer with  $d$  neurons. We stack 3 layers with skip connections to capture spatial-temporal dependencies with  $d = [16, 32, 64]$ . The dilation factor is set to have an exponentially increasing rate of 2 w.r.t the layers so the receptive field reaches 15. Although this field is smaller than the length of the sequence ( $T = 24$ ), the ablation study shows that it still produces satisfactory performance.
- **Stochastic generative stage:** The Graph Bayesian Aggregation and the likelihood function are major modules in this stage. The GBA includes a prior and a latent observation module. Both two modules contain a 1-layer  $1 \times 1$  convolution with  $d$  kernels followed by  $2 \times 1 \times 1$  convolutions to obtain the mean and variance. We stack 3 GBAs corresponding to the 3 blocks in the first stage with  $d = [16, 32, 64]$ . The likelihood function generating extrapolation results is a 3-layer  $1 \times 1$  convolution with 128 channels.

### D.2 Training Settings

All parameters are initialized with Xavier normalization [10] and optimized by the Adam optimizer [18] with a learning rate of  $10^{-3}$ .

We train each model for 150 epochs. At each iteration, we randomly sample  $N - 3$  nodes to extrapolate the remaining 3 nodes, with the time length  $T = 24$ . Note that the number of target nodes has an impact on the performance of the trained models. We conducted experiments to determine the optimal number of target nodes and found that using 3 nodes generally resulted in the best performance across all baseline models.

### D.3 Cross-Domain Evaluation

We first learn models using the Beijing dataset and then evaluate their performances on the London dataset. As the London dataset lacks weather information, we remove this attribute in Beijing during training. The other training procedures remain the same. Note that we only investigate the performance of PM2.5 concentration and exclude stations BX1, and HR1 due to their large portion of missing values. This is because, for the London dataset, both PM10 and NO2 have a significant amount of missing data (5/7 stations without any signal, and 2/1 stations with missing rates larger than 60%), which makes training unstable and the performance of the models largely depends on the training and testing data split.

## E ADDITIONAL VISUALIZATIONS

We present additional visualization results of STGNP and the baselines on the Beijing dataset. As shown in Figure 10, our model consistently outperforms others on all time steps. Additionally, on less accurate extrapolations, our model is able to yield large uncertainties, which facilitates decision-making.

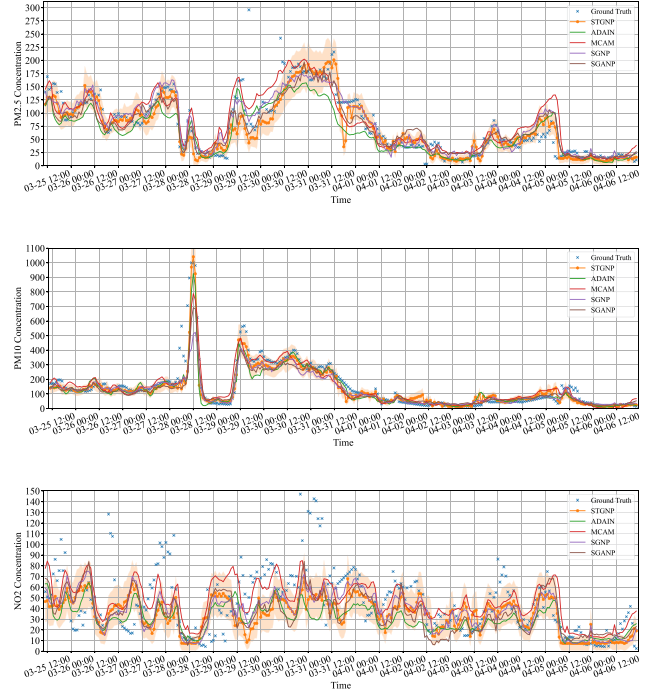


Figure 10: Visualizations of PM2.5, PM10, and NO2 extrapolations of station 1019 on the Beijing dataset.