
Conditional Temporal Neural Processes with Covariance Loss

Boseon Yoo¹ Jiwoo Lee² Janghoon Ju² Seijun Chung¹ Soyeon Kim¹ Jaesik Choi^{1,3}

Abstract

We introduce a novel loss function, *Covariance Loss*, which is conceptually equivalent to conditional neural processes and has a form of regularization so that is applicable to many kinds of neural networks. With the proposed loss, mappings from input variables to target variables are highly affected by dependencies of target variables as well as mean activation and mean dependencies of input and target variables. This nature enables the resulting neural networks to become more robust to noisy observations and recapture missing dependencies from prior information. In order to show the validity of the proposed loss, we conduct extensive sets of experiments on real-world datasets with state-of-the-art models and discuss the benefits and drawbacks of the proposed Covariance Loss.

1. Introduction

The goal of regression problems is to find relationships or functions from input features to continuous outputs (or target variables), such as stock price prediction or weather forecasting. For such problems, Gaussian processes (GPs) (Rasmussen, 2003) are promising solutions. In theory, GPs can represent any arbitrary smooth functions that satisfy mean square continuity and differentiability (Williams & Rasmussen, 2006). With the Gaussian kernel function, GPs map input variables to an infinite polynomial space and estimate the covariance of target variables. The optimization of GPs is a procedure to find a feature space in which all of the target variables can be properly explained by the covariance of target variables estimated by corresponding points in the feature space.

Recently, conditional neural processes (CNPs) are intro-

duced in (Garnelo et al., 2018a), which are the novel efforts that implement such GPs with neural networks. In CNPs, neural networks experience various target points for the same input points to learn distributions of target variables. CNPs predict the mean and variance of each target variable and share the optimization scheme of GPs which find the feature space (the space of basis function) in which all the target variables have linear relations with the covariance matrix of the target variables, which is a direct measure of the dependencies among the target variables. This is the key component behind the success of GPs and CNPs. This optimization strategy of GPs and CNPs enables more accurate prediction even under situations in which only a small volume of training data is available.

To take the main principle of GPs and CNPs into the learning of neural networks, we propose a novel loss function, called *Covariance Loss*. Covariance Loss consists of i) the traditional mean square error (MSE) term for minimizing prediction error and ii) a regularization term for minimizing MSE between the covariance matrix of the basis functions and empirical covariance matrix of target variables. Thus, by optimizing a neural network with the proposed Covariance Loss, we aim the network to find the basis function space that reflects dependencies of target variables. The resulting network takes benefit from the main principle of GPs and CNPs as well as mean dependencies between input and target variables. Theoretically, the Wishart distance is the correct mean for optimizing the covariance matrix of target variables, but we decide to use MSE to reflect dependencies of target variables more explicitly for finding basis function spaces and to make the proposed loss function more applicable to neural network based approaches by minimizing computational complexity.

With extensive sets of experiments, we show that considering dependencies of target variables for finding feature mapping allows for predictions that are more robust to not only noisy observations but also missing dependencies from prior information. To demonstrate the validity of the proposed Covariance Loss, we employ neural networks that are designed to explicitly consider the spatial and temporal dependencies, namely, spatio-temporal graph convolutional network (STGCN) (Yu et al., 2018) and graph wavenet (GWN) (Wu et al., 2019), and conduct extensive sets of experiment on well-known benchmark datasets.

¹Graduate School of AI, Korea Advanced Institute of Science and Technology, Republic of Korea. ²Department of Computer Science and Engineering, Ulsan National Institute of Science and Technology, Republic of Korea. ³Ineeji Inc., Republic of Korea. Correspondence to: Jaesik Choi <jaesik.choi@kaist.ac.kr>.

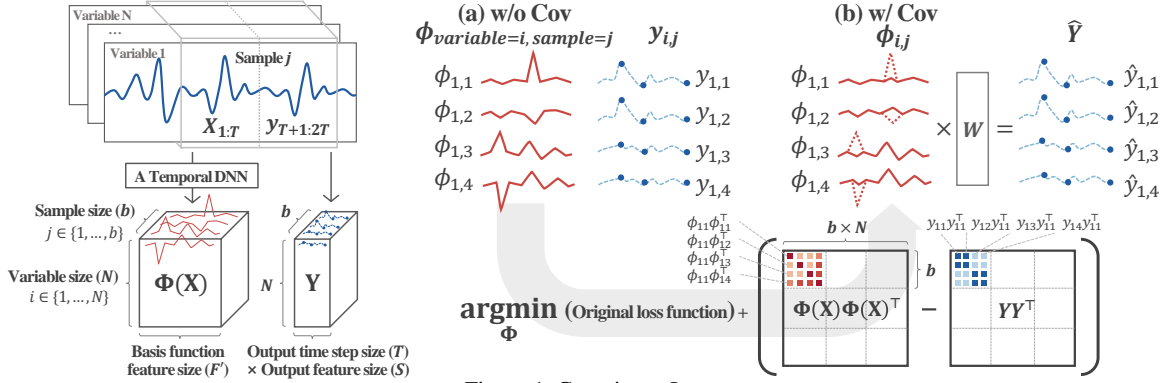


Figure 1. Covariance Loss

2. Background

2.1. Conditional Neural Processes

Conditional neural processes (CNPs) (Garnelo et al., 2018a) are indirect implementations of traditional machine learning technique, which are Gaussian processes (GPs) (Rasmussen, 2003) with neural networks. To implement GPs, in CNPs, neural networks experience various target values for the same input points and predict mean and variance. To learn distributions of mappings from input variables to corresponding target variables, both GPs and CNPs are optimized by maximizing the likelihood of the target variables on the Gaussian distribution specified by predicted mean and variance as in Equation 1.

$$\log p(\mathbf{y}|\mathbf{X}, \Theta) = -\frac{1}{2}\mathbf{y}^\top \mathbf{K}^{-1}\mathbf{y} - \frac{1}{2}\log |\mathbf{K}| - \frac{n}{2}\log(2\pi) \quad (1)$$

Thus, to maximize the likelihood of target variables in Equation 1, the key factor is to find feature (basis) space such that all target variables can be explained by the covariance matrix of target variables estimated from the basis function which is indicated by the term $\mathbf{K}^{-\frac{1}{2}}\mathbf{y}$ in Equation 1. Please note that $\mathbf{K}^{-\frac{1}{2}}\mathbf{y}$ means a relationship between target variables and their dependencies estimated from the corresponding pairs of basis functions. This is conceptually equivalent to our optimization goal that is to make basis functions reflect the dependencies of target variables. We show the equivalence in Section 3.3.

2.2. Spatio Temporal Graph Convolutional Network

The goal of this paper is to show that predictions with the proposed loss become more accurate by considering dependencies of target variables for learning basis functions. To analyze the effect of considering dependencies of target variables, we choose spatio-temporal graph convolutional network (STGCN) (Yu et al., 2018) that explicitly discovers and utilizes dependencies of input variables. STGCN is a variant of graph convolutional network (GCN) (Kipf & Welling, 2017) that generalizes convolution operator on

graph-structured data using adjacency matrix which specifies neighboring variables being considered together for feature extraction. k -hop neighbors' information can be utilized with polynomial approximations such as Chebyshev polynomials (Defferrard et al., 2016) for feature extractions. GCN (Kipf & Welling, 2017) introduces the first-order approximation on Chebyshev expansion.

With the mathematically well-established graph convolution operation, GCN has achieved successful performance, while being computationally efficient. STGCN extends GCN by stacking 1-dimensional convolution layers and graph convolution layers to extract features considering the temporal and spatial locality of input variables. Thus, accommodating STGCN with Covariance Loss that is related to spatial and temporal dependencies of output variables facilitates analyzing the characteristics imposed by the dependencies.

3. Covariance Loss

In this section, we explain the proposed Covariance Loss with definitions and give implementation details. In addition, we analyze the constraint imposed by the proposed loss and discuss the equivalence between Covariance Loss and CNPs.

3.1. Key Idea

Figure 1 presents the proposed Covariance Loss with STGCN. For a given dataset, $\mathbf{D} = \{\mathbf{X}_l, \mathbf{Y}_l | l = 1, \dots, n\}$, $\mathbf{X}_l \in \mathbb{R}^{T \times N \times F}$, $\mathbf{Y}_l \in \mathbb{R}^{T \times N \times S}$, where n denotes the size of \mathbf{D} . T and N indicate the size of time window, the number of variables. F and S are the numbers of features of input and target variables, respectively. For given \mathbf{X}_l , STGCN generates a basis function, $\Phi(\mathbf{X}_l) \in \mathbb{R}^{1 \times N \times F'}$, where F' indicates the dimension of the basis function. We use notations of \mathbf{X} , \mathbf{Y} to denote randomly chosen batch dataset of size b and $\phi_{i,j}$ indicates a set of basis functions for i^{th} variable and j^{th} sample. The key idea of the proposed loss is to find a basis function space that minimizes MSE between the covariance matrix of basis function, \mathbf{K} , which is defined as the inner product of $\Phi(\mathbf{X})$ and the empirical covariance

matrix of \mathbf{Y} as shown in Figure 1. $\mathbf{Y}\mathbf{Y}^\top$ measures how each pair of target variables moves together from their mean, and indicates underlying dependencies of the target variables. Meanwhile, the inner product of the basis functions, $\Phi(\mathbf{X})\Phi(\mathbf{X})^\top$ indicates covariance measures among each pair of basis functions for predicting \mathbf{Y} . This results in the basis functions altered from (a) to (b) as shown in Figure 1. In Figure 1 (a), while each pair $(\mathbf{y}_{1,1}, \mathbf{y}_{1,2})$, $(\mathbf{y}_{1,3}, \mathbf{y}_{1,4})$ has high covariance, the corresponding pairs of basis functions, $(\phi_{1,1}, \phi_{1,2})$, $(\phi_{1,3}, \phi_{1,4})$ have low covariance since they covariate with opposite directions. With Covariance Loss our aim is that the basis functions $(\phi_{1,1}, \phi_{1,2})$, $(\phi_{1,3}, \phi_{1,4})$ in (a) move toward each other as shown in Figure 1 (b) so that they have the covariance of target variables. We give implementation details with formal definitions of key components for Covariance Loss.

Definition 1 Given basis function $\Phi(\mathbf{X})$, $\Phi(\mathbf{X})$ and $f(\mathbf{X})$ are in a relation of a linear regression as $f(\mathbf{X}) = \Phi(\mathbf{X})\mathbf{w} + \epsilon$. Thus, $\Phi(\mathbf{X})$ is equivalent to the basis function expansion for linear regression problems.

Definition 2 We define the Covariance $\tilde{\Sigma}$ of two target variables, \mathbf{Y}_i and \mathbf{Y}_j , such that $\tilde{\Sigma}_{i,j} = (\mathbf{Y}_i - \bar{\mathbf{Y}}_i)(\mathbf{Y}_j - \bar{\mathbf{Y}}_j)$.

$\bar{\mathbf{Y}}_i$ and $\bar{\mathbf{Y}}_j$ denote the means of \mathbf{Y}_i and \mathbf{Y}_j . The covariance represents temporal and spatial dependencies of target variable when $i = j$ and $i \neq j$, respectively.

Definition 3 We define Covariance Loss such that

$$\frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n (\tilde{\Sigma}_{i,j} - \sigma^2 \Phi(\mathbf{X}_i)\Phi(\mathbf{X}_j)^\top)^2, \quad (2)$$

where σ denotes the variance of the weights of the last hidden layer. With the proposed Covariance Loss, we measure dependencies of target variables and reflect it to the corresponding basis function as shown in Figure 1. Finally, we combine Covariance Loss with the general MSE loss with the importance factor λ as shown in Equation 3.

$$\frac{1}{n} (\mathbf{Y} - \hat{\mathbf{Y}})^2 + \lambda \left(\frac{1}{n^2} (\tilde{\Sigma}_{\mathbf{Y}} - \sigma^2 \Phi \Phi^\top)^2 \right), \quad (3)$$

where $\hat{\mathbf{Y}}$ is the prediction from the model.

3.2. Constraint Analysis

In this section, we analyze the resulting constraint imposed by Covariance Loss. In regression problems, the goal of the optimization with MSE is to learn the basis function of \mathbf{X} , $\Phi(\mathbf{X})$ and \mathbf{w} such that,

$$\mathbf{Y} = \sum_{i=1}^{F'} w_i \Phi^i(\mathbf{X}). \quad (4)$$




Figure 2. The distribution of the effect of cross-terms for prediction (i.e. contribution) in log-scale, measured by Equation 8 on PeMSD7(M) dataset. The fraction of cross-terms that have zero effect: 33% (STGCN), 87% (STGCN-Cov)

Meanwhile, optimization goal of Covariance Loss is to minimize MSE between covariance matrices of empirical target variables and corresponding basis functions, each of which has a form of an inner product such that,

$$\mathbf{Y} \cdot \mathbf{Y}' = \sum_{i=1}^{F'} w_i^2 \Phi^i(\mathbf{X}) \Phi^i(\mathbf{X}'). \quad (5)$$

This imposes an additional constraint on the optimization that, to make Equations 4 and 5 hold at the same time, $\sum_i \sum_j w_i w_j \Phi(X_i) \Phi(X_j)$, for $(i \neq j)$ should converge to zero. For the case that the basis function feature size $F' = 3$, replacing the left terms of Equation 5 with Equation 4 reveals the constraint in a matrix form such that

$$\begin{bmatrix} w_1 \Phi^1(\mathbf{X}) \\ w_2 \Phi^2(\mathbf{X}) \\ w_3 \Phi^3(\mathbf{X}) \end{bmatrix}^\top \begin{bmatrix} 0 & w_2 & w_3 \\ w_1 & 0 & w_3 \\ w_1 & w_2 & 0 \end{bmatrix} \begin{bmatrix} \Phi^1(\mathbf{X}') \\ \Phi^2(\mathbf{X}') \\ \Phi^3(\mathbf{X}') \end{bmatrix} = 0. \quad (6)$$

Decomposition after multiplying the first two matrices gives the final form of the constraint on Φ and \mathbf{w} imposed by the proposed loss function as shown in Equation 7.

$$\begin{bmatrix} \Phi^1(\mathbf{X}) & \Phi^2(\mathbf{X}) & \Phi^3(\mathbf{X}) \end{bmatrix} \begin{bmatrix} 0 & w_1 w_2 & w_1 w_3 \\ w_1 w_2 & 0 & w_2 w_3 \\ w_1 w_3 & w_2 w_3 & 0 \end{bmatrix} \begin{bmatrix} \Phi^1(\mathbf{X}') \\ \Phi^2(\mathbf{X}') \\ \Phi^3(\mathbf{X}') \end{bmatrix} = 0 \quad (7)$$

Figure 2 shows that the distribution of the cross-terms' contributions on predictions of STGCN and STGCN-Cov which has an architecture of STGCN but is optimized with Covariance Loss for PeMSD7(M) dataset. We measure the contributions using the Equation 8.

$$\sum_{i=1}^{F'} \sum_{j=1}^{F'} w_i w_j \Phi^i(\mathbf{X}) \Phi^j(\mathbf{X}') - \sum_{i=1}^{F'} w_i^2 \Phi^i(\mathbf{X}) \Phi^i(\mathbf{X}') \quad (8)$$

As shown in the Figure 2, 87% of the cross-terms over the overall predictions of STGCN-Cov have zero effect while STGCN reports only 33%. In addition, as λ in Equation 5 increases, the fraction of cross-terms that have zero effect increases (Appendix 3). In Section 5.3.1, we visualize that this constraint results in learning w_i and $\Phi^i(\mathbf{X})$ of Equation 4 such that $w_i \Phi^i(\mathbf{X}) = \alpha_i \mathbf{Y}$, where $\sum_i \alpha_i = 1$.

3.3. CNPs and Covariance Loss

In CNPs, a target variable \mathbf{Y} for a given point \mathbf{X} follows the Gaussian distribution with a mean $\boldsymbol{\mu}$ and a variance $\boldsymbol{\Sigma}_Y$. Then, the joint probability of \mathbf{Y} and \mathbf{Y}' follows the joint Gaussian distributions as shown in Equation 9.

$$p(\mathbf{Y}, \mathbf{Y}') \sim \mathcal{N}\left(\begin{bmatrix} \boldsymbol{\mu} \\ \boldsymbol{\mu}' \end{bmatrix}, \mathbb{E}[\mathbf{Y}\mathbf{Y}'^\top] - \mathbb{E}[\mathbf{Y}]\mathbb{E}[\mathbf{Y}']\right) \quad (9)$$

Meanwhile, a neural network based estimation can be seen as a linear regression operation accommodated with the learning of non-linear mapping of \mathbf{x} , referred as basis function with a notation, $\Phi(\mathbf{X})$, such that $f(\mathbf{X}) = \Phi(\mathbf{X})\mathbf{w} + \epsilon$, $\epsilon \sim \mathcal{N}(0, \boldsymbol{\Sigma}_\epsilon)$. Thus, the $f(\mathbf{X})$ follows the Gaussian distribution of $\mathcal{N}(0, \boldsymbol{\Sigma}_f)$ as well under the limitation of infinite number of hidden nodes with the i.i.d. assumption. This further indicates that the joint probability of $f(\mathbf{X})$ and $f(\mathbf{X}')$ follows multivariate Gaussian distribution as in Equation 10.

$$p(f(\mathbf{X}), f(\mathbf{X}')) \sim \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \Phi(\mathbf{X})\boldsymbol{\Sigma}_f\Phi(\mathbf{X}')^\top\right) \quad (10)$$

The learning of neural networks with MSE is a usual procedure that guides the prior zero mean in Equation 10 to the true mean $\boldsymbol{\mu}$ of the target variables as in Equation 9. While most of neural network based regression algorithms are interested in finding the basis function space that gives the best predictions on the mean and ignore dependencies (i.e. the second moment in Equations 9 and 10) of target variables and in basis functions, GPs and CNPs assume zero mean in some cases, focus on considering the dependencies, \mathbf{K} for finding basis function space as discussed in 2.1.

Similarly, our Covariance Loss guides the learning of the basis function according to MSE between the covariance matrices in Equation 9 and Equation 10, directly. Without loss of generality, by assuming zero mean, Covariance Loss can be seen as $\Phi(\mathbf{X})\Phi(\mathbf{X}')^\top - \mathbf{Y}\mathbf{Y}'^\top = 0$. By defining the covariance matrix of target variables, \mathbf{K} , as estimated by the basis function, $\Phi(\mathbf{X})\Phi(\mathbf{X}')^\top$, our optimization goal becomes finding a basis function space that satisfies $\mathbf{Y}\mathbf{Y}'^\top = \mathbf{K}$. Our optimization goal is aligned to that of GPs and CNPs from the perspective that the dependencies of target variables are reflected in basis functions. However, the difference lies in that CNPs learn such a basis function space indirectly by means of maximizing the likelihood of given dataset while optimizations with Covariance Loss exploit MSE between the second moments of Equations 9 and 10.

Please note that, instead of modifying network topology or exchanging learning scheme, we bring the main principle of GPs and CNPs into the learning of neural networks with simple modifications on the usual MSE based optimizations. Furthermore, since Covariance Loss has a form of regularization, it is applicable to several kinds of networks.

4. Related Work

In this section, we briefly review existing work that exploit main principle of GPs or covariance matrix for neural networks. In (Neal, 1995), it is introduced that a fully-connected neural network is equivalent to GPs if the network consists of a single-layer with the infinite number of hidden units with an i.i.d. prior over its parameters. This study is extended to deep convolutional neural networks (CNNs) with multiple layers (Lee et al., 2018; Matthews et al., 2018; Novak et al., 2018; Garriga-Alonso et al., 2018) and other architecture (Yang, 2019). Especially, the exact equivalence between infinitely wide deep networks and GPs is derived in (Lee et al., 2018). The authors further extend their work by deriving GPs kernels for multi-hidden-layer neural networks with general non-linearity based on signal propagation theory in (Cho & Saul, 2009). A neural network-based stochastic model is introduced in (Garnelo et al., 2018b), which maximizes the likelihood of the given dataset assuming Gaussian prior. Embedding Gaussian distributions modeling into CNNs is proposed in (Wang et al., 2017a), which introduces a unique transformation of Gaussian distribution to symmetric positive definite matrix.

The utilization of the feature covariance matrix is actively applied to vision tasks (Lin & Maji, 2017; Wang et al., 2017a; Kafle & Kanan, 2017; Wang et al., 2017b; Li et al., 2017). The studies in (Lin et al., 2015; Gao et al., 2016) attempt to generate representations considering covariance matrix of features generated from CNNs. This work utilizes pairwise correlations of the information about features and locations of objects. A global pooling approach that utilizes the covariance matrix of features extracted by CNNs and generates global representation for pattern recognition is introduced in (Wang et al., 2020a;b). In time-series task, the matrix of features that represents the correlation between each pair of input variables is utilized. In (Zhang et al., 2019), signature matrices is introduced. Signature matrix consists of pairwise inner-product of input variables and implicitly contains correlation information of the variables, which helps the model to capture the spatial and temporal patterns of time-series data. This study is similar to our work as it utilizes spatial-temporal dependencies of input variables, but the difference is that Covariance Loss forces basis functions to have dependencies of corresponding target variables while (Zhang et al., 2019) utilizes dependencies of input variables for finding mapping functions. These studies are similar to our work in the sense that they exploit the main principle of GPs or covariance matrix for generating representative features, while the differences lie in that the proposed Covariance Loss uses the covariance matrix for learning the basis function space highly affected by the dependencies of target variables. Furthermore, the proposed Covariance Loss does not require any derivation from the covariance functions or matrices of the GPs. Also note that




Figure 3. Covariance matrix of \mathbf{X} , basis function of DNN, basis function of DNN-Cov, and one-hot-encoded label. Brighter color indicates higher covariance value. (b) is similar to (a) and some basis functions have high covariance with basis function that belongs to other classes. In contrast, (c) is similar to (d) and basis functions in each class are mutually exclusive.




Figure 4. An ambiguous sample in MNIST dataset. The sample (left) is similar to samples of class ‘0’ and the prediction with mean activation of input variable is incorrect (middle). In contrast, the optimization with Covariance Loss shows correct prediction for the ambiguous sample (right).

we do not set prior of the given data pair $(\mathbf{x}_i, \mathbf{y}_i)$ to find the representation. Rather than changing the network or model itself, Covariance Loss can replace the conventional loss functions with proper hyper-parameter.

5. Experimental Evaluations

This section demonstrates validity of Covariance Loss on classification and regression problems with extensive sets of experiments. For experiments, we employ state-of-the-art models such as STGCN and GWNET and optimize the models with original objective function and Covariance Loss to analyze the characteristics of resulting basis function spaces. Note that we use a notation of *model-Cov* or *Cov* to indicate models optimized with Covariance Loss.

5.1. Experimental Setups

Our experiments are conducted in an environment with Intel(R) Xeon(R) Gold 6226 CPU @ 2.70GHz and NVIDIA Quadro RTX 600 GPU cards. To generate empirical covariance matrix of target variables in a given batch dataset, we use zero mean and residuals of predictions and target variables. We employ mean absolute error (MAE), mean absolute percent error (MAPE) and root mean square error (RMSE) to evaluate the prediction performance and only valid data points are involved for the evaluations. For experiments, we share training and test datasets used by the original work.

5.2. The Effect of Covariance Loss on Classification

To analyze the effect of Covariance Loss on classification problem, we employ a simple deep neural network (DNN) and compare characteristics of resulting basis function spaces for MNIST dataset. To measure Covariance Loss we assume zero mean for estimating the empirical covariance matrix of target variables.

Figure 3 shows the covariance matrices of input variables (a), basis functions of DNN (b) and DNN-Cov (c) and one-hot-encoded labels (d) that are sorted in an ascending order of the labels. As shown in the figure, the covariance matrix of the basis function generated by DNN (b) seems to have a similar pattern with that of input variables (a) while the covariance matrix from DNN-Cov (c) has a similar pattern with the covariance matrix of the target variables (d). From Figure 3 (a) and (b), we can see that the basis functions for each sample are highly affected by mean activations of input variables and thus, even though the samples belong to different classes, the corresponding basis functions have high covariance. Meanwhile, Figure 3 (c) shows that, regardless of whether the input samples have high covariance or not, the optimization with Covariance Loss focuses more on finding a space in which only the basis functions that belongs to the same class have high covariance while the basis functions for different classes have small covariance to reflect the dependencies of target variables.

Figure 4 shows the effect caused by the change in basis function space. In the figure, for a given sample (left), we evaluate the covariance of the sample with the mean of samples for each class (middle) while the right chart of the figure shows covariance of basis functions of them, which are resulting from the optimization with Covariance Loss. As shown in the figure, the prediction of DNN follows the mean activations of input variable and ends up with an incorrect classification (middle). In contrast, DNN-Cov successfully finds a basis function space that basis functions for the same class have high covariance and achieves a correct prediction (right).

For performance evaluation, we employ state-of-the-art models such as GMNN (Qu et al., 2019) and SSP (Izadi et al., 2020) and compare classification accuracy as shown in Table 1. For given datasets, the predictions of the models optimized with Covariance Loss (w Cov) have achieved more accurate classifications compared to those of models

Table 1. Classification accuracy comparison (%)

Model	Dataset	w/o Cov	w/ Cov
DNN	MNIST	98.39	98.75
CNN	CIFAR-10	88.73	91.37
GMNN	PubMed	81.40	81.74
SSP	Cora	89.88	90.02




Figure 5. While variance of basis function of STGCN is irrelevant to that of target variables (left), basis function of STGCN-Cov has variance of target variables (right).




Figure 6. Regardless of the correlation of variables, either independent (left, Node 52 and 129) or correlated (right, Node 1 and 7), the dependencies of target variables are well-reflected to the covariance of corresponding basis functions from STGCN-Cov.

optimized with RMSE (w/o Cov) as shown in Table 1.

5.3. The Effect of the Covariance Loss on Regression

This section analyzes the effect of the Covariance Loss on regression problems which is highly unlikely to the classification case that dependencies of target variables are given as a form of mutually exclusive one-hot-encoding. For the sets of experiments, we use residuals between the target variables and the predictions, for estimating the empirical covariance matrix.

5.3.1. CASE STUDY: STGCN ON PEMS7(M) DATASET

We first show that basis functions from optimizations with Covariance Loss can reflect dependencies of target variables in continuous domain. Then, we analyze the effect of Covariance Loss on basis function spaces. For this set of experiments, we use STGCN and PeMSD7(M) dataset in (Yu et al., 2018). PeMSD7(M) is a highway traffic dataset from California which collects the velocity of cars for every 5 minutes. It consists of 12,612 velocity data over 228 stations. The optimization goal is to predict the velocities of cars on the stations for the next 15~60 minutes using histories of the recent 60 minutes.

To validate that the resulting basis functions from neural networks optimized with Covariance Loss reflect the dependencies of target variables, we compare the diagonal and non-diagonal elements of $\Phi(\mathbf{X})\Phi(\mathbf{X})^\top$ obtained from STGCN and STGCN-Cov with those of $\mathbf{Y}\mathbf{Y}^\top$. In Figures 5 and 6, the diagonal elements (variance) and non-diagonal elements (covariance) of $\Phi(\mathbf{X})\Phi(\mathbf{X})^\top$ and the corresponding $\mathbf{Y}\mathbf{Y}^\top$ are plotted as a function of time. As shown in the figure, the variance and covariance of the basis functions obtained from STGCN are irrelevant to those of the




Figure 7. Prediction and basis function (dashed color lines) from STGCN (left) and STGCN-Cov (right) on Node 7. The proposed Covariance Loss function encourages the STGCN to have the better prediction and detect the beginning of rush hour much earlier.




Figure 8. Prediction and basis function (dashed color lines) of STGCN (left) and STGCN-Cov (right) on Node 110. In this negative example, the proposed Covariance Loss function makes prediction worse since there is a significant difference in distribution between the train and test dataset.

empirical target variables (left). In contrast, STGCN-Cov successfully learns the basis functions that the variance and covariance of empirical target variables are well-reflected, which is consistent with the classification case in Figure 3. Now we analyze the effect of Covariance Loss on basis function space.

In Figures 7 and 8, we plot true labels (black line), predictions (red line) and basis functions (dashed color lines) as a function of time for STGCN (left) and STGCN-Cov (right). In the figures, we can see the characteristic of basis function space imposed by Covariance Loss. Even though the summation of the weighted basis functions at time t is a prediction at time t in both STGCN and STGCN-Cov, each basis function of STGCN-Cov has the shape of the prediction but only differs in their scales while each basis function from STGCN has a shape irrelevant to the prediction. This is because the constraint in Section 3.2. The most straightforward solution satisfying the constraint is to learn \mathbf{w} and $\Phi(\cdot)$ such that $w_i^2 \Phi^i(\mathbf{X})^2 = \lambda_i \mathbf{Y}^2$ and $\sum_{i=1}^{F'} \lambda_i = 1$, where λ and F' denotes a scale factor and the dimension of the basis functions, respectively. This explains that basis function of STGCN-Cov has a shape of the prediction as colored dash lines. This fact also leads to a high covariance between basis functions for target variables with high covariance, which is consistent with the case of classification problem. Now, we analyze the effect of the changed basis function space on predictions.

Optimizing with Covariance Loss result in changing basis function space from (b) to (c) of Figure 3. This leads to predictions that are more accurate and robust to noisy obser-

Table 2. Performance comparison on PeMSD7(M)

Model	MAE (15/30/45 min)	MAPE (15/30/45 min)	RMSE (15/30/45 min)
HA	4.01	10.61	7.2
LSVR	2.50/3.63/4.54	5.81/8.88/11.50	4.55/6.67/8.28
ARIMA	5.55/5.86/6.27	12.92/13.94/15.20	9.00/9.13/9.38
FNN	2.74/4.02/5.04	6.38/9.72/12.38	4.75/6.98/8.58
FC-LSTM	3.57/3.94/4.16	8.60/9.55/10.10	6.20/7.03/7.51
DCRNN	2.37/3.31/4.01	5.54/8.06/9.99	4.21/5.96/7.13
STGCN(Cheb)	2.25/3.03/3.57	5.26/7.33/8.69	4.04/5.70/6.77
STGCN(1st)	2.26/3.09/3.79	5.24/7.39/9.12	4.07/5.77/7.03
STGCN-Cov	2.20/2.97/3.51	5.14/7.26/8.74	4.02/5.64/6.70

vation or mis-connected dependencies by a graph structure as shown in Figure 7. In Figure 7, predictions of STGCN are highly affected by noisy observations and become incorrect as shown in the middle of rush hour (left). In contrast, STGCN-Cov (right) not only detects the beginning of rush hour earlier than STGCN (left) but also alleviates the volatility in prediction during rush hour. However, such a change in basis function space may result in performance degradation as shown in Figure 8. Optimization with Covariance Loss forces the corresponding basis functions to have a similar degree of covariance even though the mapping from input variables to such basis functions may conflict with the mean activation and the mean dependencies. When the conflict is severe, Covariance Loss may bring quantitatively marginal performance improvement or even worse predictions.

Table 2 shows prediction accuracy evaluations for short, middle and long term predictions on PeMSD7(M) dataset. For this set of experiments, we employ STGCN, STGCN(1st) with chebyshev expansions order of 1, and STGCN(Cheb) with order of 3, and use parameters presented in (Yu et al., 2018). As baseline algorithms, we employ Historical Average (HA), Linear Support Vector Regression (LSVR), Feed-Forward Neural Network (FNN), ARIMA_{kal} (Hamilton, 1994), FC-LSTM (Sutskever et al., 2014) and DCRNN (Li et al., 2018). As shown in the table, STGCN-Cov successfully improve the accuracy almost all of the evaluation

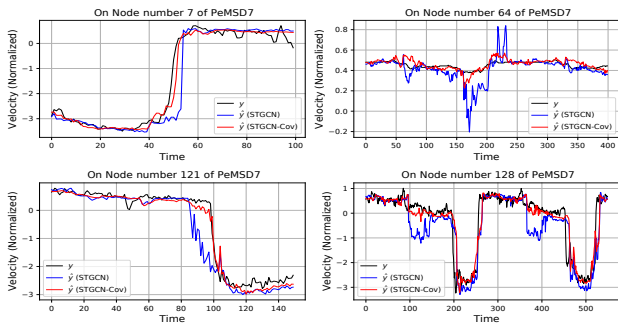


Figure 9. Predictions and labels of STGCN (red) and STGCN-Cov (blue) in PeMSD7




Figure 10. Prediction and basis functions (dashed color lines) of GWNET (left) and GWNET-Cov (right) on Node 6 of METR-LA

metric. Figure 9 shows how the proposed loss achieve such performance improvement. As shown in the figure, while noisy input variables and miss-connected dependencies deteriorate the overall performance of STGCN, STGCN-Cov has achieved correct predictions by learning basis function space that dependencies of target variables are considered.

5.3.2. CASE STUDY: GWNET ON METR-LA DATASET

In this section, we show the validity of Covariance Loss with a set of experiments for GWNET on METR-LA dataset in (Li et al., 2018). METR-LA dataset contains records of statistics on traffic speed in the highway of Los Angeles, collected from 207 sensors and 35,000 records from each sensor.

Figure 10 shows predictions and basis functions of GWNET and GWNET-Cov. The optimization with Covariance Loss results in learning basis functions that dependencies of target variables are reflected to. Each basis function of GWNET-Cov has a shape of predictions but differs in the scale, which are consistent with the case in Section 5.3.1. As shown in the figure, learning such basis functions leads to more earlier detection of the end of rush hours.

Table 3 shows the prediction results of traffic conditions over the next 15, 30, and 60 minutes of baseline models. From the table, we can see that GWNET-Cov outperforms all the others for the short term predictions and successfully improves the RMSE in every case. In Figure 11, we present how the proposed Covariance Loss improves GWNET for

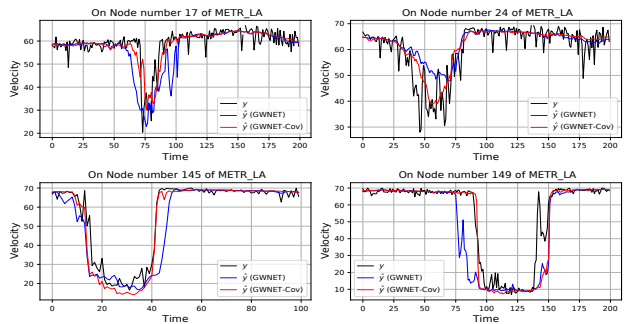


Figure 11. Predictions and labels of GWNET (red) and GWNET-Cov (blue) in METR-LA

Table 3. Performance comparison for METR-LA and PeMS-BAY datasets.

DATASET	METR-LA			PEMS-BAY		
	MAE (15/30/60 MIN)	MAPE (15/30/60 MIN)	RMSE (15/30/60 MIN)	MAE (15/30/60 MIN)	MAPE (15/30/60 MIN)	RMSE (15/30/60 MIN)
HA	4.16	13.0	7.8	2.88	6.84	5.59
ARIMA	3.99/5.15/6.90	9.6/12.7/17.4	8.1/10.5/13.2	1.62/2.33/3.38	3.5/5.4/8.3	3.30/4.76/6.50
FC-LSTM	3.44/3.77/4.37	9.60/10.9/13.2	6.30/7.23/8.69	2.05/2.20/2.37	4.8/5.2/5.7	4.19/4.55/4.69
WAVENET	2.99/3.59/4.45	8.04/10.3/13.6	5.89/7.28/8.93	1.39/1.83/2.35	2.91/4.16/5.87	3.01/4.21/5.43
DCRNN	2.77/3.15/3.60	7.30/8.80/10.5	5.38/6.45/7.60	1.38/1.74/2.07	2.9/3.9/4.9	2.95/3.97/4.74
STGCN	2.88/3.47/4.59	7.6/9.6/12.7	5.74/7.24/9.40	1.46/2.00/2.67	2.9/4.1/5.4	3.01/4.31/5.73
GWNET	2.69/3.07/3.53	6.90/8.23/9.8	5.15/6.22/7.37	1.30/1.63/1.95	2.7/3.7/4.6	2.74/3.70/4.52
GWNET-Cov	2.69/3.07/3.53	6.83/8.26/9.85	5.14/6.17/7.27	1.30/1.62/1.91	2.69/3.59/4.47	2.73/3.67/4.40

the long-term predictions (60 minutes in Table 3). Even though GWNET is able to learn a graph structure and capture spatio-temporal dependencies automatically, still suffers from inaccurate predictions caused by noisy observations and mis-connected dependencies. Figure 11 shows that the proposed Covariance Loss may become a promising solution for such cases since it has a chance to scrutinize the dependencies between every pair of target variables.

5.3.3. CASE STUDY: GWNET ON PEMS-BAY DATASET

We conduct another set of experiments with GWNET and PeMS-BAY dataset (Wu et al., 2019). PeMS-BAY contains the velocity of cars collected from 325 traffic sensor stations in the Bay Area in California, which consists of 16,937,179 samples. For performance comparison, we employ baseline models stated in Section 5.3.2 and the optimization goal is to predict velocities over the next 60 minutes from the velocities over the past 60 minutes. Table 3 shows performance evaluations of short, middle, and long-term predictions. As shown in the table, GWNET-Cov has achieved the most successful performance in almost all the metrics. Figure 12 shows how GWNET-Cov improves the performance. As shown in the figure, GWNET is suffering from noisy observations and mean dependencies. In contrast, GWNET-Cov successfully detects the beginning and end of rush hours and sudden changes in velocity in many cases.

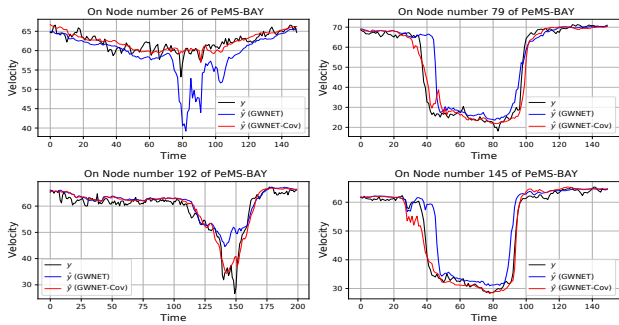


Figure 12. Predictions and labels of GWNET (red) and GWNET-Cov (blue) in PeMS-BAY.




Figure 13. Comparison of CNPs (left) and CNPs-Cov (right).




Figure 14. The effect of Covariance Loss on 1D regression Task of CNPs: CNP-Cov (left), CNP only (right)

5.3.4. CASE STUDY: CNPs, 1D LINE CURVE DATASET

In this section, we compare the performance of CNP and CNP-Cov for 1-dimensional curve fitting. For this set of experiments, we use line curve dataset which is employed in original CNPs (Garnelo et al., 2018a) and METR-LA dataset. Figure 13 shows interpolation and extrapolation performance of CNPs (left) and CNPs-Cov (right). As shown in the figure, while CNPs outperform CNPs-Cov in terms of NLL, CNPs-Cov shows more competitive RMSE. This is because with Covariance Loss, the altered basis function space brings more accurate predictions and small variances around contexts (black dots) by considering dependencies of target variables. In Figure 14, we plot 1D regression result of CNP and CNP-Cov on METR-LA dataset. As shown in the figure, predictions with Covariance Loss pursue dynamics rather than the mean dependencies which are consistent with the experiments in previous sections.

5.3.5. NOISE ROBUSTNESS

This section demonstrates the improved robustness of models optimized with Covariance Loss on noisy observations. For this set of experiments, we add noisy signals to sen-




Figure 15. The effect of noisy nodes on prediction results: As the number of noisy nodes increases from 0 to 15, the accuracy of STGCN (blue color) deteriorates severely while STGCN-Cov (red color) shows robustness on the noisy signals.




Figure 16. RMSE degeneration of STGCN v.s. STGCN-Cov: STGCN-Cov shows predictions more robust to noisy observations before loosing the global dependencies ($\#$ noisy nodes ≤ 15) in contrast to STGCN.

sor nodes of PeMSD7(M) and compare prediction errors of STGCN and STGCN-Cov as the number of noisy nodes increases. Figure 15 depicts predictions on node 1 of STGCN (blue) and STGCN-Cov (red) as the number of noisy nodes increases from 0 to 15. Note that we use a notation of Noise- k to indicate the number of noisy nodes. As shown in the figure, the prediction results of STGCN severely deteriorate when the number of noisy nodes is more than 5 while STGCN-Cov still shows relatively accurate predictions.

Figure 16 shows that RMSE degeneration of STGCN and STGCN-Cov as a function of the number of noisy nodes. We can see that STGCN-Cov is more robust than STGCN when the number of noisy nodes is less than 15 due to the spatial dependencies. However, as the number of nodes increases, the performance of STGCN-Cov follows global dependencies.




Figure 17. The effect of batch size on CIFAR10 classification.




Figure 18. Memory usage and running time comparison.

5.3.6. SCALABILITY ANALYSIS

In this section, we analyze the scalability of Covariance Loss. Since we employ a mini-batch dataset based learning scheme with many iterations, we evaluate the effect of batch size on accuracy and time complexity. To identify the effect of batch size, we compare classification accuracy on CIFAR10 dataset by varying batch size as shown in Figure 17. We can see that the effect of batch size is not significant. This is because under the random batch based training scheme with many iterations, mini-batch based Covariance Loss evaluation is eventually converged to Gram matrix based evaluation. We also test the required memory usage and running time of GWNET and GWNET-Cov with increasing batch size in Figure 18. Since Covariance Loss is random mini-batch based, it is applicable to larger dataset with mini-batch based training scheme. However, as shown in the figure, memory requirements and running time are proportional to the batch size.

6. Conclusion

In this paper, we propose a novel objective function, Covariance Loss that is applicable to several kinds of neural networks with a simple modification on the traditional MSE based optimization. Optimizations with Covariance Loss enable neural networks to learn basis function spaces that dependencies of target variables are reflected explicitly which is a main reason for success of GPs and CNPs. We analyze the form of constraint imposed by considering the dependencies for optimizations and the effect of the constraint on both classification and regression problems. With extensive sets of experiments, we demonstrate that optimizations with Covariance Loss enable more competitive predictions.

Acknowledgement

This work is supported by Institute of Information & communications Technology Planning & (IITP) grant funded by the Korea government (MSIT) (No.2017-0-01779, A machine learning and statistical inference framework for explainable artificial intelligence) and (No.2019-0-00075, Artificial Intelligence Graduate School Program (KAIST)). This work is also supported by Korea Technology and Information Promotion Agency (TIPA) grant funded by the Korea government (SEMs) (No.S2969711).

References

- Cho, Y. and Saul, L. K. Kernel methods for deep learning. In *Proceedings of the Advances in Neural Information Processing Systems*, pp. 342–350, 2009.
- Defferrard, M., Bresson, X., and Vandergheynst, P. Convolutional neural networks on graphs with fast localized spectral filtering. In *Proceedings of the Advances in Neural Information Processing Systems*, pp. 3844–3852, 2016.
- Gao, Y., Beijbom, O., Zhang, N., and Darrell, T. Compact bilinear pooling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, June 2016.
- Garnelo, M., Rosenbaum, D., Maddison, C., Ramalho, T., Saxton, D., Shanahan, M., Teh, Y. W., Rezende, D., and Eslami, S. A. Conditional neural processes. In *International Conference on Machine Learning*, pp. 1704–1713, 2018a.
- Garnelo, M., Schwarz, J., Rosenbaum, D., Viola, F., Rezende, D. J., Eslami, S., and Teh, Y. W. Neural processes. *arXiv preprint arXiv:1807.01622*, 2018b.
- Garriga-Alonso, A., Rasmussen, C. E., and Aitchison, L. Deep convolutional networks as shallow gaussian processes. In *International Conference on Learning Representations*, 2018.
- Hamilton, J. D. *Time series analysis*, volume 2. Princeton New Jersey, 1994.
- Izadi, M. R., Fang, Y., Stevenson, R., and Lin, L. Optimization of graph neural networks with natural gradient descent. *arXiv preprint arXiv:2008.09624*, 2020.
- Kafle, K. and Kanan, C. An analysis of visual question answering algorithms. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1965–1973, 2017.
- Kipf, T. N. and Welling, M. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017.
- Lee, J., Bahri, Y., Novak, R., Schoenholz, S. S., Pennington, J., and Sohl-Dickstein, J. Deep neural networks as gaussian processes. In *International Conference on Learning Representations*, 2018.
- Li, P., Xie, J., Wang, Q., and Zuo, W. Is second-order information helpful for large-scale visual recognition? In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2070–2078, 2017.
- Li, Y., Yu, R., Shahabi, C., and Liu, Y. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. In *International Conference on Learning Representations*, 2018.
- Lin, T.-Y. and Maji, S. Improved bilinear pooling with cnns. *arXiv preprint arXiv:1707.06772*, 2017.
- Lin, T.-Y., RoyChowdhury, A., and Maji, S. Bilinear cnn models for fine-grained visual recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1449–1457, 2015.
- Matthews, A. G. d. G., Hron, J., Rowland, M., Turner, R. E., and Ghahramani, Z. Gaussian process behaviour in wide deep neural networks. In *International Conference on Learning Representations*, 2018.
- Neal, R. M. *Bayesian learning for neural networks*. PhD thesis, Citeseer, 1995.
- Novak, R., Xiao, L., Bahri, Y., Lee, J., Yang, G., Hron, J., Abolafia, D. A., Pennington, J., and Sohl-dickstein, J. Bayesian deep convolutional networks with many channels are gaussian processes. In *International Conference on Learning Representations*, 2018.
- Qu, M., Bengio, Y., and Tang, J. Gmnn: Graph markov neural networks. In *International Conference on Machine Learning*, pp. 5241–5250, 2019.
- Rasmussen, C. E. Gaussian processes in machine learning. In *Summer School on Machine Learning*, pp. 63–71. Springer, 2003.
- Sutskever, I., Vinyals, O., and Le, Q. V. Sequence to sequence learning with neural networks. In *Proceedings of the Advances in Neural Information Processing Systems*, pp. 3104–3112, 2014.
- Wang, Q., Li, P., and Zhang, L. G2denet: Global gaussian distribution embedding network and its application to visual recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2730–2739, 2017a.
- Wang, Q., Xie, J., Zuo, W., Zhang, L., and Li, P. Deep cnns meet global covariance pooling: Better representation and generalization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020a.
- Wang, Q., Zhang, L., Wu, B., Ren, D., Li, P., Zuo, W., and Hu, Q. What deep cnns benefit from global covariance pooling: An optimization perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10771–10780, 2020b.

- Wang, Y., Long, M., Wang, J., and Yu, P. S. Spatiotemporal pyramid network for video action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1529–1538, 2017b.
- Williams, C. K. and Rasmussen, C. E. *Gaussian processes for machine learning*, volume 2. MIT Press Cambridge, MA, 2006.
- Wu, Z., Pan, S., Long, G., Jiang, J., and Zhang, C. Graph wavenet for deep spatial-temporal graph modeling. In *International Joint Conference on Artificial Intelligence 2019*, pp. 1907–1913, 2019.
- Yang, G. Scaling limits of wide neural networks with weight sharing: Gaussian process behavior, gradient independence, and neural tangent kernel derivation. *arXiv preprint arXiv:1902.04760*, 2019.
- Yu, B., Yin, H., and Zhu, Z. Spatio-temporal graph convolutional networks: a deep learning framework for traffic forecasting. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pp. 3634–3640, 2018.
- Zhang, C., Song, D., Chen, Y., Feng, X., Lumezanu, C., Cheng, W., Ni, J., Zong, B., Chen, H., and Chawla, N. V. A deep neural network for unsupervised anomaly detection and diagnosis in multivariate time series data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 1409–1416, 2019.