

---

# Conditional Diffusion Model with Nonlinear Data Transformation for Time Series Forecasting

---

J Rishi<sup>\* 1</sup> GVS Mothish<sup>\* 1</sup> Deepak Subramani<sup>1</sup>

## Abstract

Time-series forecasting finds application across domains such as finance, climate science, and energy systems. We introduce the Conditional Diffusion with Nonlinear Data Transformation Model (CN-Diff), a generative framework that employs novel nonlinear transformations and learnable conditions in the forward process for time series forecasting. A new loss formulation for training is proposed, along with a detailed derivation of both forward and reverse process. The new additions improve the diffusion model’s capacity to capture complex time series patterns, thus simplifying the reverse process. Our novel condition facilitates learning an efficient prior distribution. This also reduces the gap between the true negative log-likelihood and its variational approximation. CN-Diff is shown to perform better than other leading time series models on nine real-world datasets. Ablation studies are conducted to elucidate the role of each component of CN-Diff.

## 1. Introduction

Time series forecasting has become an essential element in modern data-driven decision-making processes, spanning a wide range of fields including financial pricing analysis (Kim, 2003), economics (Henrique et al., 2019), transportation planning (Huang et al., 2023), energy (Dumas et al., 2022) and various other fields (Li et al., 2024a; Rasul et al., 2022). There has been substantial advancement in time series modeling, beginning with traditional statistical techniques such as ARIMA and state-space models, which initially dominated forecasting efforts. Later, various architectures of deep neural networks, such as recurrent neural

networks (Hewamalage et al., 2021), convolutional neural networks (Yue et al., 2022), and transformers (Vaswani, 2017a) have been used for time-series modeling.

In addition to traditional deep learning architectures, diffusion models have recently gained prominence as effective tools for generative tasks. Their outstanding performance has been evidenced in various domains, exceeding the capabilities of conventional generative approaches in image generation (Ho et al., 2020; Dhariwal & Nichol, 2021), video synthesis (Harvey et al., 2022; Blattmann et al., 2023), and cross-modal applications (Saharia et al., 2022). Given this impressive success, there is a growing research interest in leveraging the potential of diffusion models for time series forecasting, leading to the development of several diffusion-based frameworks for time series modeling. Each of these frameworks is specifically designed to tackle the complexities and difficulties that arise in the context of temporal forecasting (Tashiro et al., 2021; Li et al., 2024b; Cao et al., 2024; Shen & Kwok, 2023; Shen et al., 2024; Lopez Alcaraz & Strodtzoff, 2023). These approaches largely rely on conditional diffusion models. For instance, TimeDiff (Shen & Kwok, 2023) employs future mixup as a condition-guided diffusion model. TimeGrad (Rasul et al., 2021) merges a standard diffusion model with the recurrent neural network’s hidden states. CSDI (Tashiro et al., 2021) uses self-supervised masking to steer a non-autoregressive denoising process. While effective, these time-series diffusion models do not completely exploit diffusion model properties.

Conventional training for conditional diffusion models involves gradually adding a fixed linear Gaussian noise at each forward stage. The reverse process is then optimized to closely replicate the forward process but with conditions only applied to the reverse process. Firstly, this approach confines diffusion models in the non-learnable forward process. Secondly, simplifying into an isotropic Gaussian prior complicates the data generation process in the reverse process for time series applications.

This paper introduces the Conditional Diffusion with Nonlinear Data Transformation Model (CN-Diff), a framework incorporating a nonlinear time-dependent learnable data transformation into the forward process alongside the learnable

---

<sup>\*</sup>Equal contribution <sup>1</sup>Department of Computational and Data Sciences, Indian Institute of Science, Bengaluru, India. Correspondence to: J Rishi <rishij@iisc.ac.in>, GVS Mothish <mothishg@iisc.ac.in>, Deepak Subramani <deepakns@iisc.ac.in>.

condition. This modification results in a non-Markovian series of latent variables, each formed by transforming the data and subsequently adding noise, which is learned through the reverse process.

The main contributions of the present paper are as follows.

1. We introduce a nonlinear, time-dependent data transformation along with a learnable condition for data generation in the forward process for time series forecasting.
2. We develop an objective function based on a non-Markovian learnable reverse generative process to train CN-Diff.
3. Experimental results demonstrate that our model surpasses or matches the performance of other leading time series forecasting models at the time of writing.

## 2. Preliminaries

### 2.1. Diffusion Model

Diffusion models are generative frameworks employing latent variables. For a data sample  $x \sim q(x)$ , a forward noising process produces latent variables  $x^0, x^1, x^2, \dots, x^T$ . The reverse process aims to reconstruct  $x$  by generating the same latent variables in reverse order.

In the standard diffusion model, the forward process follows a linear Gaussian Markov chain (Sohl-Dickstein et al., 2015; Ho et al., 2020). In denoising diffusion probabilistic models (DDPM), at step  $t$ ,  $x^t$  is generated by modifying the previous state  $x^{t-1}$  (multiplied by  $\sqrt{\alpha_t}$ ) with zero-mean Gaussian noise and variance  $(1 - \alpha_t)$ , that is,  $q(x^t|x^{t-1}) = \mathcal{N}(x^t; \sqrt{\alpha_t}x^{t-1}, (1 - \alpha_t)I)$ . We derive the marginal distribution as  $q(x^t|x) = \mathcal{N}(x^t; \sqrt{\bar{\alpha}_t}x, (1 - \bar{\alpha}_t)I)$ , where  $\bar{\alpha}_t$  is defined as  $\prod_{s=1}^t \alpha_s$ . Hence,  $x^t = \sqrt{\bar{\alpha}_t}x + \sqrt{1 - \bar{\alpha}_t}\epsilon$ , with  $\epsilon \sim \mathcal{N}(0, I)$ . Leveraging these marginal distributions, the joint distribution for the latent variables  $x^0, x^1, x^2, \dots, x^T$  is

$$q(x^{0:T}|x) = \prod_{t=1}^T q(x^t|x^{t-1}).$$

The forward process is typically fixed, lacking trainable parameters, and it is designed so that  $q(x^0|x) \approx \delta(x^0 - x)$  and  $q(x^T|x) \approx \mathcal{N}(x^T; 0, I)$ . If accessing  $q(x^{t-1}|x^t)$  were feasible, we could sample from  $x^T \sim \mathcal{N}(x^T; 0, I)$  and reverse it to yield  $x^0 \sim q(x^0) \approx q(x)$ . The distribution  $q(x^{t-1}|x^t)$  depends implicitly on  $q(x)$ , forming a complex relationship. Hence, we approximate the reverse process through a Markov chain adopting the form

$$p_\theta(x^{0:T}) = p(x^T) \prod_{t=1}^T p_\theta(x^{t-1}|x^t),$$

with  $p(x^T) = \mathcal{N}(x^T; 0, I)$ . The integration of the forward process  $q$  with the reverse process  $p_\theta$  is akin to a (hierarchical) variational autoencoder (Kingma, 2013a; Rezende et al., 2014). During training, the standard variational bound of the negative log-likelihood is minimized. In the case of DDPM (Ho et al., 2020), the loss to be minimized is

$$\mathcal{L} = \mathbb{E}_q \left[ \underbrace{D_{\text{KL}}(q(x^T|x) \| p(x^T))}_{\mathcal{L}_{\text{prior}}} - \underbrace{\log p_\theta(x|x^0)}_{\mathcal{L}_{\text{rec}}} + \sum_{t=1}^T \underbrace{D_{\text{KL}}(q(x^{t-1}|x^t, x) \| p_\theta(x^{t-1}|x^t))}_{\mathcal{L}_{\text{diff}}} \right].$$

Given the fixed nature of process  $q$  and distribution  $p_\theta(x^T) = p(x^T)$ , the prior term  $\mathcal{L}_{\text{prior}}$  can be disregarded as it does not depend on parameters  $\theta$ . Since  $\log p_\theta(x|x^0)$  is often modeled by a Gaussian distribution with low variance, the reconstruction term  $\mathcal{L}_{\text{rec}}$  also remains unaffected by  $\theta$ . Consequently, the diffusion term  $\mathcal{L}_{\text{diff}}$  is the only part that the model parameters  $\theta$  influence.  $\mathcal{L}_{\text{diff}}$  is the sum of Kullback–Leibler (KL) divergences between the posterior distribution in the forward process  $q(x^{t-1}|x^t, x)$  and the assumed normal distribution  $p_\theta(x^{t-1}|x^t) = \mathcal{N}(x^{t-1}; \mu_\theta(x^t, t), \Sigma_\theta(x^t, t))$  in the reverse process. Here, the variance  $\Sigma_\theta(x^t, t)$  is set to  $\sigma_t^2 I$ , while the mean  $\mu_\theta(x^t, t)$  is learned by a neural network with parameters  $\theta$ . This process is typically viewed as a noise estimation or data prediction problem (Benny & Wolf, 2022). To estimate noise, the network  $\epsilon_\theta$  forecasts the noise in the diffused input  $x^t$  and then computes  $\mu_\theta(x^t, t)$  using the formula  $\left( \frac{1}{\sqrt{\alpha_t}}x^t - \frac{1-\alpha_t}{\sqrt{1-\bar{\alpha}_t}\sqrt{\alpha_t}}\epsilon_\theta(x^t, t) \right)$ . The parameter  $\theta$  is learned by minimizing the loss function  $\mathcal{L}_\epsilon = \mathbb{E}_{x, \epsilon, t} [\|\epsilon - \epsilon_\theta(x^t, t)\|^2]$ . Alternatively, in the data prediction approach, a denoising network  $x_\theta$  is employed to derive an estimate  $x_\theta(x^t, t)$  of the clean data  $x^0$  from  $x^t$ , and then set to

$$\mu_\theta(x^t, t) = \frac{\sqrt{\alpha_t}(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t}x^t + \frac{\sqrt{\bar{\alpha}_{t-1}}\beta_t}{1 - \bar{\alpha}_t}x_\theta(x^t, t).$$

Here, the parameter  $\theta$  is learned by minimizing the loss

$$\mathcal{L}_x = \mathbb{E}_{x, \epsilon, t} \|x - x_\theta(x^t, t)\|^2.$$

For time series diffusion models, forecasting  $x_\theta$  has been found to be more effective than predicting  $\epsilon_\theta$ , as shown in (Feng et al., 2024; Shen & Kwok, 2023). Our work adopts this method.

### 2.2. Conditional Diffusion models for time series forecasting

In time series forecasting, the aim is to predict future values  $x_{1:H} \in \mathbb{R}^{d \times H}$  based on past observations  $x_{-L+1:0} \in \mathbb{R}^{d \times L}$ .

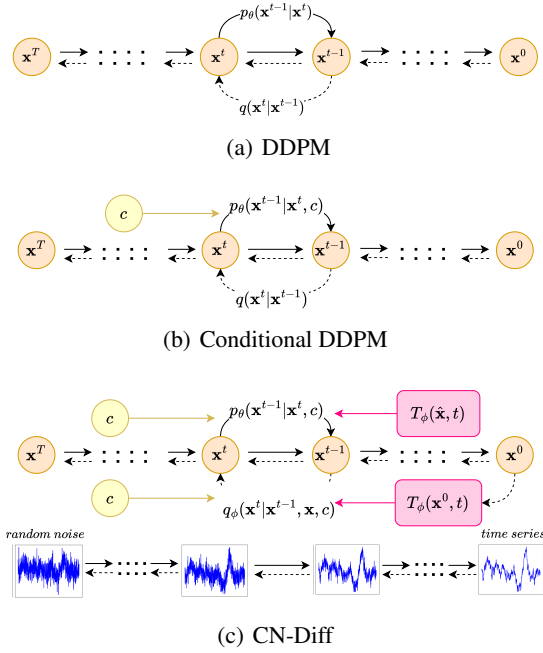


Figure 1. Directed graphical models of (a)DDPM (b)Conditional DDPM for time series and (c) CN-Diff

Here,  $L$  is the lookback window length,  $H$  is the forecast window length, and  $d$  is the count of variables, with superscripts and subscripts indicating diffusion time steps and time series values, respectively. When using conditional diffusion models for time series prediction, the following distribution is considered.

$$p_\theta(x_{1:H}^{0:T}|c) = p_\theta(x_{1:H}^T) \prod_{t=1}^T p_\theta(x_{1:H}^{t-1}|x_{1:H}^t, c).$$

where  $x_{1:H}^T \sim \mathcal{N}(0, I)$  and  $c$  represent the condition, which changes depending on the specific models used (Rasul et al., 2021; Shen & Kwok, 2023; Shen et al., 2024; Li et al., 2024b). The procedure for denoising at step  $t$  is

$$p_\theta(x_{1:H}^{t-1}|x_{1:H}^t, c) = \mathcal{N}(x_{1:H}^{t-1}; \mu_\theta(x_{1:H}^t, t|c), \sigma_t^2 I)$$

During the inference process, the generated sample linked to  $x_{1:H}^t$  is denoted as  $\hat{x}_{1:H}^t$ . We start by setting  $\hat{x}_{1:H}^T$  as a noise vector sampled from  $\mathcal{N}(0, I)$ . By repeatedly executing the denoising step outlined in the above equation until reaching  $t = 1$ , we derive the final generated sample,  $\hat{x}_{1:H}$ .

### 3. Methodology

Consider time series data  $x_{0:H} = x_0, x_1, x_2, \dots, x_H$ , where each  $x_i$  represents a vector with one or several variables with significant correlation between them at time  $i$  (Liu et al., 2023). Current conditional diffusion models

for time series convert these correlated data distributions into an isotropic Gaussian prior by adjusting data points and progressively adding linear Gaussian noise to generate latent variables. Conditional generation is performed only during the reverse process to help in learning these latent variables. This method presents two issues: first, it confines diffusion models during the forward process, rendering them fixed and untrainable. Second, converting to a simple isotropic Gaussian prior complicates data generation during the reverse process. To address these challenges, we propose CN-Diff, a nonlinear data transformation framework that learns time-dependent distributions of latent space and incorporates conditional distributions in the forward process.

In what follows, first, we describe the nonlinear, time-dependent data transformation devised for the forward process. By integrating conditional information into the forward formulation, we create a diffusion model specifically designed for time series forecasting. Subsequently, we derive the variational bound loss pertinent to this model. We hypothesize that incorporating a more adaptable distribution in the forward process can markedly decrease the disparity between the log-likelihood and the variational bound, thereby enabling the model to more effectively capture temporal dynamics and variable interrelationships. Furthermore, embedding conditional information in the formulation aids in synchronizing the prior distribution with the intrinsic temporal dynamics seen in real-world data.

#### 3.1. CN-Diff formulation and objective

Let us define the nonlinear time dependent data transformation of data for marginal distributions as

$$q_\phi(x_{1:H}^t|x_{1:H}) = \mathcal{N}\left(x_{1:H}^t; \sqrt{\alpha_t}T_\phi(x_{1:H}, t), (1 - \bar{\alpha}_t)I\right),$$

where  $T_\phi(x_{1:H}, t) : \mathbb{R}^{d \times H} \times [0, T] \mapsto \mathbb{R}^{d \times H}$  is a nonlinear neural operator parameterized by  $\phi$  that applies a time-dependent transformation to the data point  $x_{1:H}$ . For ease of notation, we replace  $x_{1:H}$  with  $\mathbf{x}$  in the following discussion.

We now introduce a learnable condition to the forward process, denoted  $c$ , inspired by a similar formulation used in image diffusion (Pandey et al., 2022). Consequently, the marginal distribution is (Figure 1):

$$q_\phi(\mathbf{x}^t|\mathbf{x}, c) = \mathcal{N}(\mathbf{x}^t; \sqrt{\alpha_t}T_\phi(\mathbf{x}, t) + (1 - \sqrt{\alpha_t})c, (1 - \bar{\alpha}_t)I). \quad (1)$$

For  $t = T$  along with an appropriately regulated noise schedule  $\alpha_t$ ,  $\bar{\alpha}_T \approx 0$  results in  $q_\phi(\mathbf{x}^T|\mathbf{x}, c) \approx \mathcal{N}(c, I)$ . Simply put, the Gaussian  $\mathcal{N}(c, I)$  serves as our learnable prior distribution and inference requires executing our reverse process on it.

**Algorithm 1** Training

---

**Input:**  $\mathbf{x}(x_{1:H}), x_{-L+1:0}$   
**repeat**  
 $\epsilon \sim \mathcal{N}(0, I), t \sim U[1, T]$   
 $\mathbf{x}^t \sim q_\phi(\mathbf{x}^t | \mathbf{x}, c)$   
 Compute the loss in Eq. (4)  
 Take numerical optimization step on:  $\nabla \mathcal{L}_{\text{CN-Diff}}$   
**until** until converged

---

We get the following posterior distribution ( $q_\phi(\mathbf{x}^{t-1} | \mathbf{x}^t, \mathbf{x}, c)$ ) that satisfies Eq.(1):

$$\begin{aligned} \mathbf{x}^{t-1} &= \zeta_1 \mathbf{x}^t + \zeta_2 c \\ &\quad - T_\phi(\mathbf{x}, t) \zeta_1 \sqrt{\bar{\alpha}_t} \\ &\quad + T_\phi(\mathbf{x}, t-1) (\zeta_1 \sqrt{\bar{\alpha}_t} + \zeta_0) + \sigma_{t-1|t}^2 \epsilon \end{aligned}$$

See Appendix A.1 for the detailed derivation.

Given that our forward process exhibits non-Markovian characteristics(Eqn.7), employing a Markovian reverse process is inadequate for accurately reconstructing the forward process. We therefore define a framework that takes the reverse process as non-Markovian which is one of many can be taken. Our non-Markovian trainable reverse generative process  $p_\theta(\mathbf{x}^{0:T})$  is as follows:

$$p_\theta(\mathbf{x}^{0:T}) = p(\mathbf{x}^T) \prod_{t=1}^T p_\theta(\mathbf{x}^{t-1} | \mathbf{x}^t) p_\theta(\mathbf{x}^0 | \mathbf{x}^t), \quad (2)$$

$$\text{where } p(\mathbf{x}^T) = \mathcal{N}(\mathbf{x}^T; c, I). \quad (3)$$

For this choice, our **Lemma A.1** shows that, this approach results in a training objective analogous to that of DDPM even without non linear transform.

Refer to Appendix A.1 for proof.

Consequently, the corresponding variational objective of CN-Diff is represented by following form (Appendix A.2).

$$\begin{aligned} \mathcal{L}_{\text{CN-Diff}} &= \mathbb{E}_{q_\phi} \left[ \underbrace{D_{\text{KL}}(q_\phi(\mathbf{x}^T | \mathbf{x}, c) \| p(\mathbf{x}^T | c))}_{\mathcal{L}_{\text{prior}}} \right. \\ &\quad \left. - \underbrace{\sum_{t=1}^T \log p_\theta(\mathbf{x} | \mathbf{x}^t, c)}_{\mathcal{L}_{\text{rec}}} \right. \\ &\quad \left. + \underbrace{\sum_{t=1}^T D_{\text{KL}}(q_\phi(\mathbf{x}^{t-1} | \mathbf{x}^t, \mathbf{x}, c) \| p_\theta(\mathbf{x}^{t-1} | \mathbf{x}^t, c))}_{\mathcal{L}_{\text{diff}}} \right]. \end{aligned} \quad (4)$$

Note that forward process is parametrized by  $\phi$  by the transformation  $T_\phi(\cdot)$ , making the objective different from (Ho

**Algorithm 2** Inference

---

$\mathbf{x}^T \sim \mathcal{N}(c, I)$   
**for**  $t = T$  **to** 1 **do**  
 $\hat{\mathbf{x}} = \hat{\mathbf{x}}_\theta(\mathbf{x}^t, t, c)$   
 $\mathbf{x}^{t-1} = \zeta_1 \mathbf{x}^t + \zeta_2 c - T_\phi(\hat{\mathbf{x}}, t) \zeta_1 \sqrt{\bar{\alpha}_t} + T_\phi(\hat{\mathbf{x}}, t-1) (\zeta_1 \sqrt{\bar{\alpha}_t} + \zeta_0) + \sigma_{t-1|t}^2 \epsilon$   
**end for**  
 $\hat{\mathbf{x}} \sim p(\mathbf{x} | \mathbf{x}_0, c)$

---

et al., 2020). Consequently, the prior and reconstruction terms depend on  $\phi$ , and therefore must be included in the optimization process.

In the reverse process we approximate posteriors as,  $p_\theta(\mathbf{x}^{t-1} | \mathbf{x}^t, c) \approx q_\phi(\mathbf{x}^{t-1} | \mathbf{x}^t, \hat{\mathbf{x}}_\theta(\mathbf{x}^t, t), c)$ . Thus,  $\mathcal{L}_{\text{diff}}$  and  $\mathcal{L}_{\text{prior}}$  are

$$\begin{aligned} \mathcal{L}_{\text{diff}} &= D_{\text{KL}}(q_\phi(\mathbf{x}^{t-1} | \mathbf{x}^t, \mathbf{x}, c) \| p_\theta(\mathbf{x}^{t-1} | \mathbf{x}^t, c)) \\ &= \frac{1}{2\sigma_{t-1|t}^2} \left\| \zeta_1 \sqrt{\bar{\alpha}_t} (T_\phi(\mathbf{x}, t) - T_\phi(\hat{\mathbf{x}}_\theta(\mathbf{x}^t, t), t)) - \right. \\ &\quad \left. (\zeta_1 \sqrt{\bar{\alpha}_t} + \zeta_0) (T_\phi(\hat{\mathbf{x}}_\theta(\mathbf{x}^t, t-1), t) - T_\phi(\mathbf{x}, t-1)) \right\|_2^2, \\ \mathcal{L}_{\text{prior}} &= D_{\text{KL}}(q_\phi(\mathbf{x}^T | \mathbf{x}) \| p(\mathbf{x}^T)) \\ &= \frac{1}{2} \bar{\alpha}_T \|T_\phi(\mathbf{x}, T) - c\|_2^2. \end{aligned}$$

The complete derivation is given in Appendix A.3. It is essential to distinguish between the objectives of DDPM and CN-Diff. For DDPM, the objective is to accurately predict only the original data point  $x$ . On the other hand, CN-Diff is designed to predict the transformed data point  $T_\phi(x, t)$  along with the original data point  $x$ . Despite this modification, the optimization process for CN-Diff remains simulation-free, thereby facilitating its efficient training through KL divergence computed on samples at all time steps. The procedures for both training and inference are detailed in Algorithms 1 and 2.

## 4. Experiments

This section presents time series forecasting experiments using CN-Diff and compares its performance against recent deep neural models across nine popular real-world datasets. We begin by describing the datasets, architectural framework, and hyperparameters. We then present results benchmarked against existing forecasting outcomes. An ablation study follows to assess the unique components of our CN-Diff model.

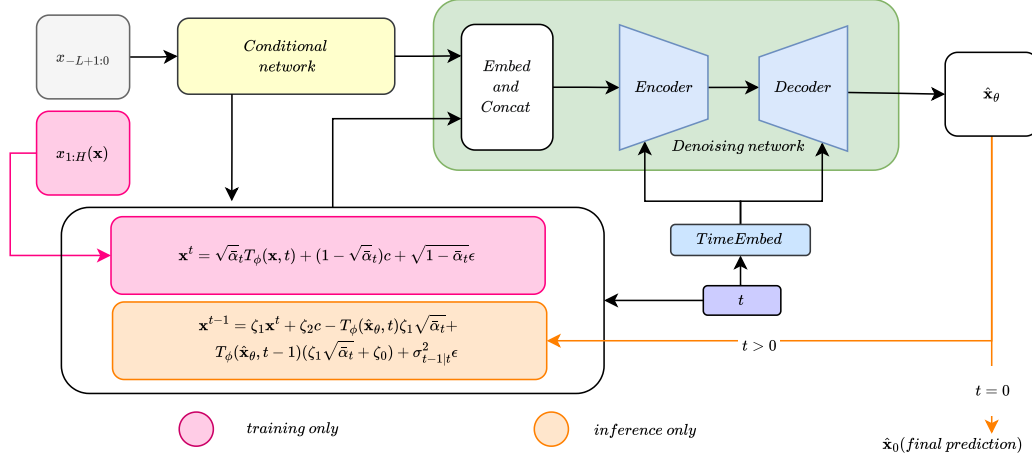


Figure 2. An illustration of the CN-Diff architecture. During the training phase, historical data ( $x_{-L+1:0}$ ) is processed through a conditional network, concatenated with noised data, and subsequently passed through the denoising network. During the inference phase, the model initiates from  $(\mathcal{N}(c, I))$  and performs denoising to obtain  $(\hat{x}_0)$ .

## 4.1. Experimental Setup

### 4.1.1. DATASET

We performed an experimental study on nine real-world time series datasets for daily, weekly, and monthly forecasts (Shen et al., 2024; Fan et al., 2022; Zhou et al., 2021; Wu et al., 2021). These datasets are (1) Wind<sup>1</sup> - wind power data for 2020-2021 every 15 minutes; (2) Caiso<sup>2</sup> - hourly electricity loads over eight years from various California regions; (3) Traffic<sup>3</sup> - hourly road occupancy rates from San Francisco Bay area sensors; (4) Electricity<sup>4</sup> - hourly electricity use of 321 clients over two years; (5) Weather<sup>5</sup> - 10-minute meteorological data for 2020-2021; (6) NorPool<sup>6</sup> - eight years of hourly energy production data from various European countries; (7) Exchange<sup>7</sup> - daily exchange rates for eight countries: Australia, United Kingdom, Canada, Switzerland, China, Japan, New Zealand and Singapore (Lai et al., 2018); (8) ETTh1 and (9) ETTm1<sup>8</sup> are benchmarks for China’s electricity transformer temperature data, over two years, ETTh1 reported hourly and ETTm1 every 15

<sup>1</sup><https://github.com/PaddlePaddle/PaddleSpatial/tree/main/paddlespatial/datasets/WindPower>

<sup>2</sup><https://www.energyonline.com/Data/>

<sup>3</sup><https://pems.dot.ca.gov/>

<sup>4</sup><https://archive.ics.uci.edu/ml/datasets/ElectricityLoadDiagrams20112014>

<sup>5</sup><https://www.bgc-jena.mpg.de/wetter/>

<sup>6</sup><https://data.nordpoolgroup.com/power-system/production/>

<sup>7</sup><https://github.com/laiguokun/multivariate-time-series-data>

<sup>8</sup><https://github.com/zhouhaoyi/ETDataset>

minutes (Zhou et al., 2021). Additional details regarding the dataset’s characteristics are available in Table 2.

Following (Shen et al., 2024), we adopt a train:validation:test ratio of 7:1:2 for Exchange, Weather, Wind, Traffic and Electricity data sets, and 6:2:2 for ETTh1 and ETTm1 data sets. For Norpool, the training data consists of observations before April 1, 2020; validation data spans from April 1 to October 1, 2020; and testing data is after October 1, 2020. For Caiso, the training set includes data before January 1, 2020, the validation period extends from January 1 to October 1, 2020, and the test set is post-October 1, 2020. Given that data sets exhibit different sampling intervals, we focus on prediction tasks with suitable prediction lengths based on the characteristics of the dataset (Shen & Kwok, 2023; Shen et al., 2024).

### 4.1.2. IMPLEMENTATION DETAILS

Our CN-Diff model is trained using the Adam optimizer with a learning rate of  $10^{-3}$ . The training employs a batch size of 64 and incorporates early stopping (maximum of 100 epochs). We use  $T = 100$  diffusion steps with a quadratic variance schedule beginning at  $\beta_1 = 10^{-4}$  and progressing to  $\beta_T = 10^{-1}$ . The length of the look-back window is selected from the set  $\{96, 192, 336, 720, 1440\}$ , determined by performance evaluations on the validation dataset averaged over ten runs. All experiments are executed on a single Nvidia RTX A6000 GPU with 48GB of vRAM. Figure 2 shows the architecture employed for the CN-Diff model. For condition  $c$ , a single dense layer is used. For  $T_\phi$ , two dense layers with tanh activation are used, one layer handling the feature dimension and the other handling the



Table 1. Multivariate prediction of MSEs on nine real-world time series datasets(subscripts is the rank). CSDI runs out of memory on Traffic and Electricity. Results of all baselines are from (Shen et al., 2024)

Method	NorPool	Caiso	Traffic	Electricity	Weather	Exchange	ETTh1	ETTm1	Wind	Rank
Ours	<b>0.531</b> <sub>(1)</sub>	<b>0.094</b> <sub>(1)</sub>	<u>0.374</u> <sub>(2)</sub>	<b>0.145</b> <sub>(1)</sub>	<b>0.296</b> <sub>(1)</sub>	<b>0.016</b> <sub>(1)</sub>	<b>0.405</b> <sub>(1)</sub>	<u>0.340</u> <sub>(2)</sub>	0.988 <sub>(5)</sub>	1.7
mr-Diff	0.645 <sub>(3)</sub>	0.127 <sub>(4)</sub>	0.474 <sub>(7)</sub>	0.155 <sub>(4)</sub>	<b>0.296</b> <sub>(1)</sub>	<b>0.016</b> <sub>(1)</sub>	0.411 <sub>(4)</sub>	<u>0.340</u> <sub>(2)</sub>	<b>0.881</b> <sub>(1)</sub>	3.0
TimeDiff	0.665 <sub>(5)</sub>	0.136 <sub>(7)</sub>	0.564 <sub>(8)</sub>	0.193 <sub>(6)</sub>	<u>0.311</u> <sub>(3)</sub>	0.018 <sub>(7)</sub>	<u>0.407</u> <sub>(2)</sub>	<b>0.336</b> <sub>(1)</sub>	<u>0.896</u> <sub>(2)</sub>	4.6
TimeGrad	1.152 <sub>(21)</sub>	0.258 <sub>(20)</sub>	1.745 <sub>(23)</sub>	0.736 <sub>(22)</sub>	0.392 <sub>(15)</sub>	0.079 <sub>(21)</sub>	0.993 <sub>(23)</sub>	0.874 <sub>(22)</sub>	1.209 <sub>(22)</sub>	21.0
CSDI	1.011 <sub>(20)</sub>	0.253 <sub>(19)</sub>	-	-	0.356 <sub>(10)</sub>	0.077 <sub>(20)</sub>	0.497 <sub>(8)</sub>	0.529 <sub>(18)</sub>	1.066 <sub>(10)</sub>	15.0
SSSD	0.872 <sub>(13)</sub>	0.195 <sub>(11)</sub>	0.642 <sub>(12)</sub>	0.255 <sub>(13)</sub>	0.349 <sub>(9)</sub>	0.061 <sub>(17)</sub>	0.726 <sub>(19)</sub>	0.464 <sub>(14)</sub>	1.188 <sub>(20)</sub>	14.2
D <sup>3</sup> VAE	0.745 <sub>(10)</sub>	0.241 <sub>(18)</sub>	0.928 <sub>(18)</sub>	0.286 <sub>(16)</sub>	0.375 <sub>(12)</sub>	0.200 <sub>(23)</sub>	0.504 <sub>(10)</sub>	0.362 <sub>(9)</sub>	1.118 <sub>(16)</sub>	14.7
CPF	1.613 <sub>(24)</sub>	0.383 <sub>(22)</sub>	1.625 <sub>(22)</sub>	0.793 <sub>(23)</sub>	1.390 <sub>(24)</sub>	<b>0.016</b> <sub>(1)</sub>	0.730 <sub>(20)</sub>	0.482 <sub>(16)</sub>	1.140 <sub>(18)</sub>	18.9
PSA-GAN	1.501 <sub>(23)</sub>	0.510 <sub>(24)</sub>	1.614 <sub>(21)</sub>	0.535 <sub>(21)</sub>	1.220 <sub>(22)</sub>	0.018 <sub>(7)</sub>	0.623 <sub>(18)</sub>	0.537 <sub>(19)</sub>	1.127 <sub>(17)</sub>	19.1
N-Hits	0.716 <sub>(8)</sub>	0.131 <sub>(5)</sub>	0.386 <sub>(3)</sub>	0.152 <sub>(3)</sub>	0.323 <sub>(5)</sub>	<u>0.017</u> <sub>(6)</sub>	0.498 <sub>(9)</sub>	0.353 <sub>(7)</sub>	1.033 <sub>(7)</sub>	5.9
FiLM	0.723 <sub>(9)</sub>	0.179 <sub>(9)</sub>	0.628 <sub>(11)</sub>	0.210 <sub>(9)</sub>	0.327 <sub>(6)</sub>	<b>0.016</b> <sub>(1)</sub>	0.426 <sub>(6)</sub>	0.347 <sub>(5)</sub>	0.984 <sub>(4)</sub>	6.7
Depts	0.662 <sub>(4)</sub>	0.106 <sub>(3)</sub>	1.019 <sub>(20)</sub>	0.319 <sub>(18)</sub>	0.761 <sub>(20)</sub>	0.020 <sub>(10)</sub>	0.579 <sub>(14)</sub>	0.380 <sub>(11)</sub>	1.082 <sub>(13)</sub>	12.6
NBeats	0.832 <sub>(11)</sub>	0.141 <sub>(8)</sub>	<b>0.373</b> <sub>(1)</sub>	0.269 <sub>(14)</sub>	1.344 <sub>(23)</sub>	<b>0.016</b> <sub>(1)</sub>	0.586 <sub>(16)</sub>	0.391 <sub>(12)</sub>	1.069 <sub>(11)</sub>	10.8
Scaleformer	0.983 <sub>(16)</sub>	0.207 <sub>(14)</sub>	0.618 <sub>(10)</sub>	0.195 <sub>(7)</sub>	0.462 <sub>(17)</sub>	0.036 <sub>(13)</sub>	0.613 <sub>(17)</sub>	0.481 <sub>(15)</sub>	1.359 <sub>(23)</sub>	14.7
PatchTST	0.851 <sub>(12)</sub>	0.193 <sub>(10)</sub>	0.831 <sub>(17)</sub>	0.225 <sub>(11)</sub>	0.782 <sub>(21)</sub>	0.047 <sub>(15)</sub>	0.526 <sub>(12)</sub>	0.372 <sub>(10)</sub>	1.077 <sub>(12)</sub>	13.3
FedFormer	0.873 <sub>(14)</sub>	0.205 <sub>(12)</sub>	0.591 <sub>(9)</sub>	0.238 <sub>(12)</sub>	0.342 <sub>(8)</sub>	0.133 <sub>(22)</sub>	0.541 <sub>(13)</sub>	0.426 <sub>(13)</sub>	1.113 <sub>(15)</sub>	13.1
Autoformer	0.940 <sub>(15)</sub>	0.226 <sub>(16)</sub>	0.688 <sub>(16)</sub>	0.201 <sub>(8)</sub>	0.360 <sub>(11)</sub>	0.056 <sub>(16)</sub>	0.516 <sub>(11)</sub>	0.565 <sub>(20)</sub>	1.083 <sub>(14)</sub>	14.1
Pyraformer	1.008 <sub>(19)</sub>	0.273 <sub>(21)</sub>	0.659 <sub>(13)</sub>	0.273 <sub>(15)</sub>	0.394 <sub>(16)</sub>	0.032 <sub>(12)</sub>	0.579 <sub>(14)</sub>	0.493 <sub>(17)</sub>	1.061 <sub>(9)</sub>	15.1
Informr	0.985 <sub>(17)</sub>	0.251 <sub>(17)</sub>	0.664 <sub>(14)</sub>	0.298 <sub>(17)</sub>	0.385 <sub>(13)</sub>	0.073 <sub>(19)</sub>	0.775 <sub>(22)</sub>	0.673 <sub>(21)</sub>	1.168 <sub>(19)</sub>	17.7
Transformer	1.005 <sub>(18)</sub>	0.206 <sub>(13)</sub>	0.671 <sub>(15)</sub>	0.328 <sub>(19)</sub>	0.388 <sub>(14)</sub>	0.062 <sub>(18)</sub>	0.759 <sub>(21)</sub>	0.992 <sub>(23)</sub>	1.201 <sub>(21)</sub>	18.0
SCINet	<u>0.613</u> <sub>(2)</sub>	<u>0.095</u> <sub>(2)</sub>	0.434 <sub>(6)</sub>	0.171 <sub>(5)</sub>	0.329 <sub>(7)</sub>	0.036 <sub>(13)</sub>	0.465 <sub>(7)</sub>	0.359 <sub>(8)</sub>	1.055 <sub>(8)</sub>	6.4
NLinear	0.707 <sub>(7)</sub>	0.135 <sub>(6)</sub>	0.430 <sub>(5)</sub>	<u>0.147</u> <sub>(2)</sub>	0.313 <sub>(4)</sub>	0.019 <sub>(9)</sub>	0.410 <sub>(3)</sub>	0.349 <sub>(6)</sub>	0.989 <sub>(5)</sub>	5.2
DLinear	0.670 <sub>(6)</sub>	0.461 <sub>(23)</sub>	0.389 <sub>(4)</sub>	0.215 <sub>(10)</sub>	0.488 <sub>(18)</sub>	0.022 <sub>(11)</sub>	0.415 <sub>(5)</sub>	0.345 <sub>(4)</sub>	0.899 <sub>(3)</sub>	9.3
LSTMa	1.481 <sub>(22)</sub>	0.217 <sub>(15)</sub>	0.966 <sub>(19)</sub>	0.414 <sub>(20)</sub>	0.662 <sub>(19)</sub>	0.403 <sub>(24)</sub>	1.149 <sub>(24)</sub>	1.030 <sub>(24)</sub>	1.464 <sub>(24)</sub>	21.2

Table 2. Dataset characteristics

dataset	dim	#observations	freq.	$H$ (steps)
NorPool	18	70,128	1 hour	1 month (720)
Caiso	10	74,472	1 hour	1 month (720)
Weather	21	52,696	10 mins	1 week (672)
ETTm1	7	69,680	15 mins	2 days (192)
Wind	7	48,673	15 mins	2 days (192)
Traffic	862	17,544	1 hour	1 week (168)
Electricity	321	26,304	1 hour	1 week (168)
ETTh1	7	17,420	1 hour	1 week (168)
Exchange	8	7,588	1 day	2 weeks (14)

forecast window dimension. More details of the code and hyperparameters are in the Appendix B.

## 4.2. Main Results

Table 1 presents the mean squared errors (MSEs) in various multivariate time series datasets. The data reveal that the proposed CN-Diff method achieves superior performance in 6 out of 9 datasets. In particular, the enhancement is particularly pronounced in more complex datasets, namely Norpool and ETTh1. In the remaining three datasets, CN-Diff has Rank 2 in two cases, but only with a very narrow margin with respect to Rank 1 for those data sets. For Wind, CN-Diff achieves the fourth rank. The wind data set represents the highly volatile wind power data. The superior performance of the multiresolution diffusion (mrDiff (rank

1 for wind)) model suggests that wind patterns need to be analyzed at multiple time scales simultaneously. We believe that incorporating CN-Diff at multiple resolutions to build multiresolution CN-Diff could help us to model such datasets better.

Overall, the average ranking of the CN-Diff model is 1.7 showing its state-of-the-art performance. Section 5 has brief descriptions of the other leading models.

Table 3 displays the mean squared errors (MSEs) for multiple univariate time series datasets. The findings indicate that the CN-Diff excels in four of nine datasets. In the other five datasets, it secures second place in two instances. Overall, the average ranking of the CN-Diff surpasses all other leading models. Additional results for the mean absolute error (MAEs) are provided in Appendix D. Appendix F provides a qualitative comparison with various models and includes visualizations for different data sets. The visualizations show that CN-Diff is capable of capturing seasonal patterns without explicitly being trained to do so.

## 4.3. Qualitative analysis

Figure 3 illustrates the prediction results on ETTh1 generated by CN-Diff alongside three competitive models: Informr, DLinear, and FiLM. It is evident that CN-Diff yields predictions of superior quality compared to the other models.

Table 3. Univariate prediction of MSEs on nine real-world time series datasets (subscript is rank). Results of all other models are from (Shen et al., 2024).

Method	NorPool	Caiso	Traffic	Electricity	Weather	Exchange	ETTh1	ETTm1	Wind	Rank
Ours	<u>0.606</u> <sub>(2)</sub>	<u>0.118</u> <sub>(2)</sub>	0.126 <sub>(4)</sub>	0.234 <sub>(3)</sub>	<b>0.002</b> <sub>(1)</sub>	<b>0.015</b> <sub>(1)</sub>	<b>0.066</b> <sub>(1)</sub>	<b>0.038</b> <sub>(1)</sub>	2.171 <sub>(3)</sub>	2.0
mr-Diff	0.667 <sub>(5)</sub>	0.122 <sub>(4)</sub>	<b>0.119</b> <sub>(1)</sub>	0.234 <sub>(3)</sub>	<b>0.002</b> <sub>(1)</sub>	<u>0.016</u> <sub>(2)</sub>	<b>0.066</b> <sub>(1)</sub>	0.039 <sub>(3)</sub>	2.182 <sub>(5)</sub>	2.8
TimeDiff	0.636 <sub>(3)</sub>	0.122 <sub>(4)</sub>	<u>0.121</u> <sub>(2)</sub>	<u>0.232</u> <sub>(2)</sub>	<b>0.002</b> <sub>(1)</sub>	0.017 <sub>(5)</sub>	<b>0.066</b> <sub>(1)</sub>	0.040 <sub>(6)</sub>	2.407 <sub>(14)</sub>	4.2
TimeGrad	1.129 <sub>(23)</sub>	0.325 <sub>(23)</sub>	1.223 <sub>(24)</sub>	0.920 <sub>(24)</sub>	<b>0.002</b> <sub>(1)</sub>	0.041 <sub>(21)</sub>	0.078 <sub>(11)</sub>	0.048 <sub>(12)</sub>	2.530 <sub>(19)</sub>	17.6
CSDI	0.967 <sub>(22)</sub>	0.192 <sub>(15)</sub>	0.393 <sub>(21)</sub>	0.520 <sub>(19)</sub>	<b>0.002</b> <sub>(1)</sub>	0.071 <sub>(24)</sub>	0.083 <sub>(16)</sub>	0.050 <sub>(16)</sub>	2.434 <sub>(16)</sub>	16.7
SSSD	1.145 <sub>(24)</sub>	0.176 <sub>(13)</sub>	0.151 <sub>(9)</sub>	0.370 <sub>(12)</sub>	0.004 <sub>(12)</sub>	0.023 <sub>(17)</sub>	0.097 <sub>(21)</sub>	0.049 <sub>(14)</sub>	3.149 <sub>(23)</sub>	16.1
D <sup>3</sup> VAE	0.964 <sub>(21)</sub>	0.521 <sub>(24)</sub>	0.151 <sub>(9)</sub>	0.535 <sub>(20)</sub>	0.003 <sub>(10)</sub>	0.019 <sub>(13)</sub>	0.078 <sub>(11)</sub>	0.044 <sub>(10)</sub>	2.679 <sub>(21)</sub>	15.4
CPF	0.855 <sub>(16)</sub>	0.260 <sub>(22)</sub>	0.279 <sub>(19)</sub>	0.609 <sub>(22)</sub>	<b>0.002</b> <sub>(1)</sub>	<u>0.016</u> <sub>(2)</sub>	0.080 <sub>(14)</sub>	0.041 <sub>(7)</sub>	2.430 <sub>(15)</sub>	13.1
PSA-GAN	0.658 <sub>(4)</sub>	0.150 <sub>(8)</sub>	0.250 <sub>(18)</sub>	0.273 <sub>(8)</sub>	0.035 <sub>(22)</sub>	0.020 <sub>(14)</sub>	0.084 <sub>(17)</sub>	0.051 <sub>(17)</sub>	2.510 <sub>(18)</sub>	14.0
N-Hits	0.739 <sub>(10)</sub>	0.170 <sub>(12)</sub>	0.147 <sub>(8)</sub>	0.346 <sub>(9)</sub>	<b>0.002</b> <sub>(1)</sub>	0.017 <sub>(5)</sub>	0.089 <sub>(18)</sub>	0.043 <sub>(9)</sub>	2.406 <sub>(13)</sub>	9.4
FILM	0.707 <sub>(8)</sub>	0.185 <sub>(14)</sub>	0.198 <sub>(14)</sub>	0.260 <sub>(6)</sub>	0.007 <sub>(15)</sub>	0.018 <sub>(10)</sub>	0.070 <sub>(4)</sub>	<b>0.038</b> <sub>(1)</sub>	<u>2.143</u> <sub>(2)</sub>	8.2
Depts	0.668 <sub>(6)</sub>	<b>0.107</b> <sub>(1)</sub>	0.151 <sub>(9)</sub>	0.380 <sub>(15)</sub>	0.024 <sub>(20)</sub>	0.020 <sub>(14)</sub>	0.070 <sub>(4)</sub>	0.046 <sub>(11)</sub>	3.457 <sub>(24)</sub>	11.6
NBeats	0.768 <sub>(12)</sub>	0.125 <sub>(6)</sub>	0.142 <sub>(7)</sub>	0.378 <sub>(14)</sub>	0.137 <sub>(23)</sub>	<u>0.016</u> <sub>(2)</sub>	0.095 <sub>(20)</sub>	0.048 <sub>(12)</sub>	2.434 <sub>(16)</sub>	12.4
Scaleformer	0.778 <sub>(13)</sub>	0.232 <sub>(17)</sub>	0.286 <sub>(20)</sub>	0.361 <sub>(10)</sub>	0.009 <sub>(18)</sub>	0.035 <sub>(20)</sub>	0.150 <sub>(23)</sub>	0.078 <sub>(23)</sub>	2.646 <sub>(20)</sub>	18.2
PatchTST	<b>0.595</b> <sub>(1)</sub>	0.193 <sub>(16)</sub>	0.177 <sub>(13)</sub>	0.450 <sub>(18)</sub>	0.026 <sub>(21)</sub>	0.020 <sub>(14)</sub>	0.106 <sub>(22)</sub>	0.052 <sub>(19)</sub>	2.698 <sub>(22)</sub>	16.2
FedFormer	0.891 <sub>(17)</sub>	0.164 <sub>(10)</sub>	0.173 <sub>(12)</sub>	0.376 <sub>(13)</sub>	0.005 <sub>(13)</sub>	0.050 <sub>(23)</sub>	0.076 <sub>(8)</sub>	0.065 <sub>(22)</sub>	2.351 <sub>(12)</sub>	14.4
Autoformer	0.946 <sub>(20)</sub>	0.248 <sub>(18)</sub>	0.473 <sub>(22)</sub>	0.659 <sub>(23)</sub>	0.003 <sub>(10)</sub>	0.041 <sub>(21)</sub>	0.081 <sub>(15)</sub>	0.051 <sub>(17)</sub>	2.349 <sub>(11)</sub>	17.4
Pyraformer	0.933 <sub>(19)</sub>	0.165 <sub>(11)</sub>	0.136 <sub>(5)</sub>	0.389 <sub>(16)</sub>	0.020 <sub>(19)</sub>	0.017 <sub>(8)</sub>	0.076 <sub>(8)</sub>	0.054 <sub>(20)</sub>	2.279 <sub>(7)</sub>	12.6
Informer	0.804 <sub>(14)</sub>	0.250 <sub>(19)</sub>	0.213 <sub>(16)</sub>	0.363 <sub>(11)</sub>	0.007 <sub>(15)</sub>	0.023 <sub>(17)</sub>	0.076 <sub>(8)</sub>	0.049 <sub>(14)</sub>	2.297 <sub>(8)</sub>	13.6
Transformer	0.928 <sub>(19)</sub>	0.250 <sub>(19)</sub>	0.238 <sub>(17)</sub>	0.430 <sub>(15)</sub>	0.007 <sub>(15)</sub>	0.018 <sub>(10)</sub>	0.092 <sub>(19)</sub>	0.058 <sub>(21)</sub>	2.306 <sub>(10)</sub>	16.0
SCINet	0.746 <sub>(11)</sub>	0.154 <sub>(9)</sub>	0.212 <sub>(15)</sub>	0.272 <sub>(7)</sub>	<b>0.002</b> <sub>(1)</sub>	0.018 <sub>(10)</sub>	0.071 <sub>(7)</sub>	0.039 <sub>(3)</sub>	<b>2.063</b> <sub>(1)</sub>	7.1
NLinear	0.708 <sub>(9)</sub>	0.147 <sub>(7)</sub>	0.124 <sub>(3)</sub>	<b>0.231</b> <sub>(1)</sub>	<b>0.002</b> <sub>(1)</sub>	0.017 <sub>(5)</sub>	0.070 <sub>(4)</sub>	0.039 <sub>(3)</sub>	2.193 <sub>(6)</sub>	4.3
DLinear	0.671 <sub>(7)</sub>	<u>0.118</u> <sub>(2)</sub>	0.139 <sub>(6)</sub>	0.244 <sub>(5)</sub>	0.168 <sub>(24)</sub>	0.017 <sub>(5)</sub>	0.078 <sub>(11)</sub>	0.041 <sub>(7)</sub>	2.171 <sub>(3)</sub>	7.8
LSTMa	0.836 <sub>(15)</sub>	0.253 <sub>(21)</sub>	1.032 <sub>(23)</sub>	0.596 <sub>(21)</sub>	0.005 <sub>(13)</sub>	0.031 <sub>(19)</sub>	0.167 <sub>(24)</sub>	0.091 <sub>(24)</sub>	2.299 <sub>(9)</sub>	18.8

#### 4.4. Ablation Study

To evaluate the individual effects of each component of the CN-Diff model, we conducted various experiments. Initially, we illustrate the impact of the introduction of the condition  $c$  (Eqn.1), followed by an examination of the effects of the nonlinear transformation.

In the condition ablation study (Table 4), we present the forecast performance in terms of MSE and MAE when using only the diffusion model (DDPM (Ho et al., 2020)) and when applying the condition only in the backward process, which represents the standard conditional diffusion setting for time series analysis). These numbers are compared with the results with our CN-Diff conditional method, where the condition is used in both forward and backward processes. The MSE and MAE improve by 8% and 6% for ETTh1 (feature dimension of 7), 16% and 9% for electricity (feature dimension of 321) and 21% and 16% for traffic (feature dimension of 862). This clearly demonstrates that with an increase in the feature dimension, the effect of our conditioning is enhanced while maintaining the same forecast window. We can infer that, as the feature dimension increases, learning from an isotropic Gaussian prior becomes considerably more challenging and CN-Diff solves this challenge.

For the ablation study on the time-dependent nonlinear data transformation  $T_\phi$ , we consider four variations (Table 5): (i) CN-Diff without introducing any nonlinear transformation; nonlinear transformations (ii) only along the feature dimension, (iii) only along the look-back window dimension, and (iv) along both dimensions. From variation (i) to (ii), the MSE and MAE improved by 54% and 36% for ETTh1, 72% and 52% for Electricity, 57% and 35% for Caiso and 51% and 30% for Wind. From (i) to (iii), the improvements are similar at 54% and 38% for ETTh1, 70% and 49% for Electricity, 25% and 8% for Caiso and 48% and 27% for Wind. From (i) to (iv), an additional improvement is observed at 57% and 40% for ETTh1, 73% and 53% for Electricity, 61% and 38% for Caiso and 53% and 31% for Wind. Clearly, the incorporation of nonlinear data transformation in both dimensions effectively captures the significant correlations present within both the sequence and feature dimensions of time series data.

We present a joint ablation study in Appendix E.1, examining the contributions due to variations of nonlinear transformation and CN-Diff condition. Additional ablation studies for different diffusion parameters are found in Appendix E.3, and for various look-back windows in Appendix E.2. The results indicate a notable enhancement with each component added to the diffusion model. Optimal performance with varying diffusion parameters is attained at diffusion




Figure 3. Visualizations of ETTh1 dataset predictions by (a)Informer(Zhou et al., 2021) (b)DLinear(Zeng et al., 2023) (c)FiLM(Zhou et al., 2022a) and (d) CN-Diff

Table 4. Predicting MSEs and MAEs for DDPM, conditional DDPM and CN-Diff condition

Methods	Diff		Cond Diff		CN-Diff Cond	
Datasets	MSE	MAE	MSE	MAE	MSE	MAE
ETTh1	1.748	0.927	1.030	0.746	<b>0.943</b>	<b>0.704</b>
Electricity	1.784	1.018	0.637	0.564	<b>0.536</b>	<b>0.512</b>
Traffic	1.794	0.885	1.071	0.573	<b>0.848</b>	<b>0.479</b>

noise steps  $T = 100$  and an upper variance schedule of  $\beta_T = 0.1$ .

## 5. Related work

Time series forecasting is utilized in numerous aspects of daily life, ranging from essential applications like traffic and electricity monitoring, and healthcare (Liu et al., 2022b), to more intricate uses such as predicting stock prices, weather, and wind patterns (Shen & Kwok, 2023). Due to its wide significance, research has advanced from traditional state space models to contemporary deep learning methodologies.

Basis expansion has been utilized by several models for time-series forecasting. FiLM (Zhou et al., 2022a) employed Fourier analysis and low-rank matrix approximation to reduce noise, along with Legendre polynomial projection to maintain historical data representations. NBeats (Oreshkin et al., 2019) is an interpretable architecture that combines polynomial trend modeling with Fourier techniques for detecting seasonality. Depts (Fan et al., 2022) builds on NBeats by adding a periodicity module for periodic series, while N-Hits (Challu et al., 2023) enhances the approach using multi-scale hierarchical interpolation.

Models such as SCINet (Liu et al., 2022a) utilize a recursive

strategy involving downsampling, convolution, and interaction, to harness temporal dependencies in downsampled subsequences. NLinear (Zeng et al., 2023) implements a basic approach by normalizing the time series before using a linear layer for prediction. DLinear (Zeng et al., 2023) employs a seasonal-trend decomposition akin to Autoformer (Wu et al., 2021).

Several models utilized the transformer (Vaswani, 2017b) and its variations for time series prediction. Informer (Zhou et al., 2021) used sparse attention to minimize computational load and employed a generative-style decoder for rapid long-sequence forecasts in one forward pass. Autoformer (Wu et al., 2021) substituted traditional self-attention with an autocorrelation layer, while Fedformer (Zhou et al., 2022b) incorporated frequency domain mapping through a frequency-enhanced module. Pyraformer (Liu et al., 2022b) introduced pyramidal attention for multi-resolution representation, and Scaleformer (Shabani et al., 2023) adopted a shared-weight multilevel forecasting approach, from broad to fine scales. Inspired by vision transformers (Dosovitskiy et al., 2020), PatchTST (Nie et al., 2023) partitions time series data into subseries patches and uses self-supervised pre-training to extract local semantic features, thus improving long-term prediction.



Table 5. Predicting MSEs and MAEs by applying Nonlinear Transformations across different configurations

Methods	Without $T_\phi$		Feature dim		Forecast window		Both	
Datasets	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
Electricity	0.536	0.512	0.150	0.246	0.160	0.260	<b>0.145</b>	<b>0.243</b>
ETTh1	0.943	0.704	0.429	0.447	0.430	0.439	<b>0.405</b>	<b>0.421</b>
Caiso	0.242	0.307	0.104	0.200	0.182	0.283	<b>0.094</b>	<b>0.191</b>
Wind	2.090	1.023	1.015	0.715	1.067	0.743	<b>0.988</b>	<b>0.701</b>

Recently, diffusion models have emerged as a promising method for time series forecasting. TimeGrad (Rasul et al., 2021) pioneered the use of conditional diffusion models, leveraging autoregressive prediction informed by RNN hidden states. CSDI (Tashiro et al., 2021), applies non-autoregressive generation through self-supervised masking but relies on dual transformers, facing boundary inconsistencies and computational challenges with large datasets. SSSD (Lopez Alcaraz & Strodthoff, 2023) aims to reduce CSDI’s computational load using structured state-space models, but it continues with masking-based conditioning, preserving boundary issues. TimeDiff (Shen & Kwok, 2023) introduced future mixup and autoregressive initialization within an encoder-decoder scheme for denoising. TMDM (Li et al., 2024b) integrates transformers with a diffusion process for probabilistic multivariate time series forecasting. TimeDiT (Cao et al., 2024) is a foundational model for time series, which employed a transformer type architecture to capture temporal dependencies and employs diffusion processes to generate samples. mr-Diff (Shen et al., 2024) is a recent study utilizing a multi-resolution strategy, incorporating fine-to-coarse patterns as latent variables to aid in denoising.

One step diffusion (Frans et al., 2024), a family of generative models that use a single network and training phase to produce high-quality samples in a single or multiple sampling steps in image diffusion. NDM (Bartosh et al., 2024) applies a non-linear time-dependent transformation for image diffusion models with the Markovian reverse process. Polydiffuse (Chen et al., 2023) uses a guided diffusion model for polygonal shape reconstruction, where the prior is learned via guidance networks  $\mu(x, t, i)$  and  $\sigma(x, t, i)$ , which are independent of the condition. Also they have not incorporated the condition in the forward process (ref. Fig. 3 [(Chen et al., 2023)]). In contrast, our approach incorporates a condition in forward formulation results in condition-dependent prior. And our training procedure and loss function differs significantly, as polydiffuse uses separate stages for prior learning and denoising. PriorGrad (Lee et al., 2021) is diffusion-based generative model for speech synthesis, which introduces a prior distribution based on data statistics rather than a learned prior. While their reverse process starts by sampling from  $\mathcal{N}(0, \Sigma)$ , our method samples from

$\mathcal{N}(c, \Sigma)$  as we include a learnable condition in the forward process. These fundamental differences set our work apart from the above.

## 6. Conclusion and Future work

This paper presents CN-Diff, a new framework for time series forecasting that combines non-linear data transformation with conditional information in the forward process. We have derived the variational loss for the diffusion model from this framework. Our experiments indicate that CN-Diff matches or surpasses the performance of top time series models across nine real-world datasets. The avenues of future work include adding an explicit seasonal trend diffusion model, exploring different conditions using varied resolutions, and metadata-based conditional generation. We also plan to extend our model for imputation, telemetry-based anomaly detection, and other engineering applications.

## Acknowledgment

This research is partially supported by a research grant (MoES/36/OOIS/Extra/84/2022) from the Ministry of Earth Sciences, Government of India. We are also grateful for the TCS PhD Fellowship to J Rishi.

## Impact Statement

This article presents new research exploring a novel diffusion-based framework for modeling time series data. The proposed model exhibits excellent performance when compared to existing diffusion-based time series models for forecasting tasks. Our model shows great promise for improving forecast accuracy in diverse applications such as environmental, finance, climate, and industrial domains. Our research is centered on forecasting time series data, which on its own does not have any major ethical concerns. The downstream application developers must exercise caution to ensure that no ethical concerns arise when using our model for societal applications.

## References

- Bahdanau, D., Cho, K. H., and Bengio, Y. Neural machine translation by jointly learning to align and translate. In *3rd International Conference on Learning Representations, ICLR 2015*, 2015.
- Bartosh, G., Vetrov, D., and Naesseth, C. A. Neural diffusion models. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 3073–3095, 2024.
- Benny, Y. and Wolf, L. Dynamic dual-output diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11482–11491, 2022.
- Blattmann, A., Rombach, R., Ling, H., Dockhorn, T., Kim, S. W., Fidler, S., and Kreis, K. Align your latents: High-resolution video synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 22563–22575, 2023.
- Cao, D., Ye, W., and Liu, Y. Timedit: General-purpose diffusion transformers for time series foundation model. In *ICML 2024 Workshop on Foundation Models in the Wild*, 2024.
- Challu, C., Olivares, K. G., Oreshkin, B. N., Ramirez, F. G., Canseco, M. M., and Dubrawski, A. Nhits: Neural hierarchical interpolation for time series forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pp. 6989–6997, 2023.
- Chen, J., Deng, R., and Furukawa, Y. Polydiffuse: Polygonal shape reconstruction via guided set diffusion models. *Advances in Neural Information Processing Systems*, 36: 1863–1888, 2023.
- Das, A., Kong, W., Leach, A., Mathur, S., Sen, R., and Yu, R. Long-term forecasting with tide: Time-series dense encoder. *arXiv preprint arXiv:2304.08424*, 2023a.
- Das, A., Kong, W., Sen, R., and Zhou, Y. A decoder-only foundation model for time-series forecasting. *arXiv preprint arXiv:2310.10688*, 2023b.
- Dhariwal, P. and Nichol, A. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2020.
- Dumas, J., Wehenkel, A., Lanaspeze, D., Cornélusse, B., and Sutura, A. A deep generative model for probabilistic energy forecasting in power systems: normalizing flows. *Applied Energy*, 305:117871, 2022.
- Esser, P., Kulal, S., Blattmann, A., Entezari, R., Müller, J., Saini, H., Levi, Y., Lorenz, D., Sauer, A., Boesel, F., et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first International Conference on Machine Learning*, 2024.
- Fan, W., Zheng, S., Yi, X., Cao, W., Fu, Y., Bian, J., and Liu, T.-Y. Depts: Deep expansion learning for periodic time series forecasting. In *International Conference on Learning Representations*, 2022.
- Feng, S., Miao, C., Zhang, Z., and Zhao, P. Latent diffusion transformer for probabilistic time series forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 11979–11987, 2024.
- Frans, K., Hafner, D., Levine, S., and Abbeel, P. One step diffusion via shortcut models. *arXiv preprint arXiv:2410.12557*, 2024.
- Ge, S., Nah, S., Liu, G., Poon, T., Tao, A., Catanzaro, B., Jacobs, D., Huang, J.-B., Liu, M.-Y., and Balaji, Y. Preserve your own correlation: A noise prior for video diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 22930–22941, 2023.
- Harvey, W., Naderiparizi, S., Masrani, V., Weilbach, C., and Wood, F. Flexible diffusion modeling of long videos. *Advances in Neural Information Processing Systems*, 35: 27953–27965, 2022.
- Henrique, B. M., Sobreiro, V. A., and Kimura, H. Literature review: Machine learning techniques applied to financial market prediction. *Expert Systems with Applications*, 124: 226–251, 2019.
- Hewamalage, H., Bergmeir, C., and Bandara, K. Recurrent neural networks for time series forecasting: Current status and future directions. *International Journal of Forecasting*, 37(1):388–427, 2021.
- Ho, J., Jain, A., and Abbeel, P. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Huang, X., Wu, D., and Boulet, B. Metaprobformer for charging load probabilistic forecasting of electric vehicle charging stations. *IEEE Transactions on Intelligent Transportation Systems*, 24(10):10445–10455, 2023.
- Jeha, P., Bohlke-Schneider, M., Mercado, P., Kapoor, S., Nirwan, R. S., Flunkert, V., Gasthaus, J., and Januschowski, T. Psa-gan: Progressive self attention gans for synthetic

- time series. In *The Tenth International Conference on Learning Representations*, 2022.
- Jia, Y., Anaissi, A., and Suleiman, B. Resnls: An improved model for stock price forecasting. *Computational Intelligence*, 40(1):e12608, 2024.
- Kim, K.-j. Financial time series forecasting using support vector machines. *Neurocomputing*, 55(1-2):307–319, 2003.
- Kingma, D. P. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013a.
- Kingma, D. P. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013b.
- Lai, G., Chang, W.-C., Yang, Y., and Liu, H. Modeling long-and short-term temporal patterns with deep neural networks. In *The 41st international ACM SIGIR conference on research & development in information retrieval*, pp. 95–104, 2018.
- Langley, P. Crafting papers on machine learning. In Langley, P. (ed.), *Proceedings of the 17th International Conference on Machine Learning (ICML 2000)*, pp. 1207–1216, Stanford, CA, 2000. Morgan Kaufmann.
- Lee, S.-g., Kim, H., Shin, C., Tan, X., Liu, C., Meng, Q., Qin, T., Chen, W., Yoon, S., and Liu, T.-Y. Priorgrad: Improving conditional denoising diffusion models with data-dependent adaptive prior. *arXiv preprint arXiv:2106.06406*, 2021.
- Li, L., Carver, R., Lopez-Gomez, I., Sha, F., and Anderson, J. Generative emulation of weather forecast ensembles with diffusion models. *Science Advances*, 10(13):eadk4489, 2024a.
- Li, Y., Lu, X., Wang, Y., and Dou, D. Generative time series forecasting with diffusion, denoise, and disentanglement. *Advances in Neural Information Processing Systems*, 35: 23009–23022, 2022a.
- Li, Y., Lu, X., Wang, Y., and Dou, D. Generative time series forecasting with diffusion, denoise, and disentanglement. *Advances in Neural Information Processing Systems*, 35: 23009–23022, 2022b.
- Li, Y., Chen, W., Hu, X., Chen, B., Zhou, M., et al. Transformer-modulated diffusion models for probabilistic multivariate time series forecasting. In *The Twelfth International Conference on Learning Representations*, 2024b.
- Liu, M., Zeng, A., Chen, M., Xu, Z., Lai, Q., Ma, L., and Xu, Q. Scinet: Time series modeling and forecasting with sample convolution and interaction. *Advances in Neural Information Processing Systems*, 35:5816–5828, 2022a.
- Liu, S., Yu, H., Liao, C., Li, J., Lin, W., Liu, A. X., and Dustdar, S. Pyraformer: Low-complexity pyramidal attention for long-range time series modeling and forecasting. In *# PLACEHOLDER\_PARENT\_METADATA\_VALUE#*, 2022b.
- Liu, Y., Guo, H., Zhang, L., Liang, D., Zhu, Q., Liu, X., Lv, Z., Dou, X., and Gou, Y. Research on correlation analysis method of time series features based on dynamic time warping algorithm. *IEEE Geoscience and Remote Sensing Letters*, 20:1–5, 2023.
- Lopez Alcaraz, J. M. and Strodthoff, N. Diffusion-based time series imputation and forecasting with structured atate apace models. *Transactions on machine learning research*, pp. 1–36, 2023.
- Nie, Y., Nguyen, N. H., Sinthong, P., and Kalagnanam, J. A time series is worth 64 words: Long-term forecasting with transformers. In *The Eleventh International Conference on Learning Representations*, 2023.
- Oreshkin, B. N., Carpo, D., Chapados, N., and Bengio, Y. N-beats: Neural basis expansion analysis for interpretable time series forecasting. In *International Conference on Learning Representations*, 2019.
- Pandey, K., Mukherjee, A., Rai, P., and Kumar, A. Diffusevae: Efficient, controllable and high-fidelity generation from low-dimensional latents. *Transactions on Machine Learning Research*, 2022.
- Peebles, W. and Xie, S. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4195–4205, 2023.
- Rangapuram, S. S., Kapoor, S., Nirwan, R. S., Mercado, P., Januschowski, T., Wang, Y., and Bohlke-Schneider, M. Coherent probabilistic forecasting of temporal hierarchies. In *International Conference on Artificial Intelligence and Statistics*, pp. 9362–9376. PMLR, 2023.
- Rasul, K., Seward, C., Schuster, I., and Vollgraf, R. Autoregressive denoising diffusion models for multivariate probabilistic time series forecasting. In *International Conference on Machine Learning*, pp. 8857–8868. PMLR, 2021.
- Rasul, K., Park, Y.-J., Ramström, M. N., and Kim, K.-M. Vq-ar: Vector quantized autoregressive probabilistic time series forecasting. *arXiv preprint arXiv:2205.15894*, 2022.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. In *International conference on machine learning*, pp. 1278–1286. PMLR, 2014.

- Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E. L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems*, 35: 36479–36494, 2022.
- Shabani, M. A., Abdi, A. H., Meng, L., and Sylvain, T. Scaleformer: Iterative multi-scale refining transformers for time series forecasting. In *The Eleventh International Conference on Learning Representations*, 2023.
- Shen, L. and Kwok, J. Non-autoregressive conditional diffusion models for time series prediction. In *International Conference on Machine Learning*, pp. 31016–31029. PMLR, 2023.
- Shen, L., Chen, W., and Kwok, J. Multi-resolution diffusion models for time series forecasting. In *The Twelfth International Conference on Learning Representations*, 2024.
- Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., and Ganguli, S. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pp. 2256–2265. PMLR, 2015.
- Song, J., Meng, C., and Ermon, S. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.
- Taga, E. O., Ildiz, M. E., and Oymak, S. Timepfn: Effective multivariate time series forecasting with synthetic data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 20761–20769, 2025.
- Tashiro, Y., Song, J., Song, Y., and Ermon, S. Csd: Conditional score-based diffusion models for probabilistic time series imputation. *Advances in Neural Information Processing Systems*, 34:24804–24816, 2021.
- Vaswani, A. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017a.
- Vaswani, A. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017b.
- Wu, H., Xu, J., Wang, J., and Long, M. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. *Advances in neural information processing systems*, 34:22419–22430, 2021.
- Yue, Z., Wang, Y., Duan, J., Yang, T., Huang, C., Tong, Y., and Xu, B. Ts2vec: Towards universal representation of time series. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pp. 8980–8987, 2022.
- Zeng, A., Chen, M., Zhang, L., and Xu, Q. Are transformers effective for time series forecasting? In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pp. 11121–11128, 2023.
- Zhang, Z., Meng, L., and Gu, Y. Sageformer: Series-aware framework for long-term multivariate time-series forecasting. *IEEE Internet of Things Journal*, 11(10):18435–18448, 2024.
- Zhou, H., Zhang, S., Peng, J., Zhang, S., Li, J., Xiong, H., and Zhang, W. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pp. 11106–11115, 2021.
- Zhou, T., Ma, Z., Wen, Q., Sun, L., Yao, T., Yin, W., Jin, R., et al. Film: Frequency improved legendre memory model for long-term time series forecasting. *Advances in neural information processing systems*, 35:12677–12690, 2022a.
- Zhou, T., Ma, Z., Wen, Q., Wang, X., Sun, L., and Jin, R. Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting. In *International conference on machine learning*, pp. 27268–27286. PMLR, 2022b.

## A. Formal Derivations

### A.1. Forward Posterior

Given:

$$q_\phi(\mathbf{x}^t | \mathbf{x}, c) = \mathcal{N}(\mathbf{x}^t; \sqrt{\bar{\alpha}_t} T_\phi(\mathbf{x}, t) + (1 - \sqrt{\bar{\alpha}_t})c, (1 - \bar{\alpha}_t)I) \quad (5)$$

From (5), We can write,

$$\begin{aligned} \mathbf{x}^t &= \sqrt{\bar{\alpha}_t} T_\phi(\mathbf{x}, t) + (1 - \sqrt{\bar{\alpha}_t})c + \sqrt{1 - \bar{\alpha}_t} \epsilon \\ &= \sqrt{\bar{\alpha}_t} T_\phi(\mathbf{x}, t) + (1 - \sqrt{\bar{\alpha}_t})c + \sqrt{1 - \alpha_t \bar{\alpha}_{t-1}} \epsilon \\ &= \sqrt{\bar{\alpha}_t} T_\phi(\mathbf{x}, t) + (1 - \sqrt{\bar{\alpha}_t})c + \sqrt{1 - \alpha_t} \epsilon + \sqrt{\alpha_t - \alpha_t \bar{\alpha}_{t-1}} \epsilon \\ &= \sqrt{\bar{\alpha}_t} T_\phi(\mathbf{x}, t) + (1 - \sqrt{\bar{\alpha}_t})c + \sqrt{1 - \alpha_t} \epsilon + \sqrt{\alpha_t} \sqrt{1 - \bar{\alpha}_{t-1}} \epsilon \\ &= \sqrt{\bar{\alpha}_t} T_\phi(\mathbf{x}, t) + (1 - \sqrt{\bar{\alpha}_t})c + \sqrt{1 - \alpha_t} \epsilon + \sqrt{\alpha_t} (\mathbf{x}^{t-1} - \sqrt{\bar{\alpha}_{t-1}} T_\phi(\mathbf{x}, t-1) - (1 - \sqrt{\bar{\alpha}_{t-1}})c) \\ &= \sqrt{\alpha_t} \mathbf{x}^{t-1} + \sqrt{\bar{\alpha}_t} (T_\phi(\mathbf{x}, t) - T_\phi(\mathbf{x}, t-1)) + (1 - \sqrt{\alpha_t})c + \sqrt{1 - \alpha_t} \epsilon \end{aligned} \quad (6)$$

By introducing nonlinear, time-dependent transformation of data, along with condition, results in a forward process that is non-Markovian as:

$$q_\phi(\mathbf{x}^t | \mathbf{x}^{t-1}, \mathbf{x}, c) = \mathcal{N}(\mathbf{x}^t; \sqrt{\alpha_t} \mathbf{x}^{t-1} + \sqrt{\bar{\alpha}_t} (T_\phi(\mathbf{x}, t) - T_\phi(\mathbf{x}, t-1)) + (1 - \sqrt{\alpha_t})c, (1 - \alpha_t)I) \quad (7)$$

From (7 and 5), Posterior distribution can be derived as:

$$\begin{aligned} q_\phi(\mathbf{x}^{t-1} | \mathbf{x}^t, \mathbf{x}, c) &\propto \frac{q_\phi(\mathbf{x}^{t-1}, \mathbf{x}^t, \mathbf{x}, c)}{q_\phi(\mathbf{x}^t, \mathbf{x}, c)} \\ &\propto \frac{q_\phi(\mathbf{x}^t | \mathbf{x}^{t-1}, \mathbf{x}, c)}{q_\phi(\mathbf{x}^t, \mathbf{x}, c)} q_\phi(\mathbf{x}^{t-1}, \mathbf{x}, c) \\ &\propto q_\phi(\mathbf{x}^t | \mathbf{x}^{t-1}, \mathbf{x}, c) q_\phi(\mathbf{x}^{t-1} | \mathbf{x}, c) \\ &\propto \exp \left( -\frac{1}{2} \left( \frac{(\mathbf{x}^t - \sqrt{\alpha_t} \mathbf{x}^{t-1} - \sqrt{\bar{\alpha}_t} (T_\phi(\mathbf{x}, t) - T_\phi(\mathbf{x}, t-1)) - (1 - \sqrt{\alpha_t})c)^2}{1 - \alpha_t} \right. \right. \\ &\quad \left. \left. + \frac{(\mathbf{x}^{t-1} - \sqrt{\bar{\alpha}_{t-1}} T_\phi(\mathbf{x}, t-1) - (1 - \sqrt{\bar{\alpha}_{t-1}})c)^2}{1 - \bar{\alpha}_{t-1}} \right) \right) \\ &\propto \exp \left( -\frac{1}{2} \left( \left( \frac{\alpha_t}{1 - \alpha_t} + \frac{1}{1 - \bar{\alpha}_{t-1}} \right) (\mathbf{x}^{t-1})^2 \right. \right. \\ &\quad \left. \left. - \frac{(2\sqrt{\alpha_t} \mathbf{x}^{t-1} (\mathbf{x}^t - \sqrt{\bar{\alpha}_t} (T_\phi(\mathbf{x}, t) - T_\phi(\mathbf{x}, t-1)) - (1 - \sqrt{\alpha_t})c)}{1 - \alpha_t} \right. \right. \\ &\quad \left. \left. - \frac{2\mathbf{x}^{t-1}}{1 - \bar{\alpha}_{t-1}} (\sqrt{\bar{\alpha}_{t-1}} T_\phi(\mathbf{x}, t-1) + (1 - \sqrt{\bar{\alpha}_{t-1}})c) \right) \right) \end{aligned} \quad (8)$$

By reformulating the above propotional form, akin to normal distribution, we can see that,

$$q_\phi(\mathbf{x}^{t-1} | \mathbf{x}^t, \mathbf{x}, c) = \mathcal{N}(\mathbf{x}^{t-1}; \mu_{t-1|t}, \sigma_{t-1|t}^2 I) \quad (10)$$

Where,

$$\begin{aligned} \mu_{t-1|t} &= \left( \frac{\sqrt{\alpha_t} \mathbf{x}^t - \sqrt{\bar{\alpha}_t} \sqrt{\alpha_t} (T_\phi(\mathbf{x}, t) - T_\phi(\mathbf{x}, t-1)) - \sqrt{\alpha_t} (1 - \sqrt{\alpha_t})c}{1 - \alpha_t} + \right. \\ &\quad \left. \frac{\sqrt{\bar{\alpha}_{t-1}} T_\phi(\mathbf{x}, t-1) + (1 - \sqrt{\bar{\alpha}_{t-1}})c}{1 - \bar{\alpha}_{t-1}} \right) \sigma_{t-1|t}^2 \end{aligned} \quad (11)$$



$$\sigma_{t-1|t}^2 = \frac{(1 - \alpha_t)(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t} \quad (12)$$

Mean( $\mu_{t-1|t}$ ) can be simplified as,

$$\begin{aligned} \mu_{t-1|t} &= \left( \frac{\sqrt{\alpha_t} \mathbf{x}^t - \sqrt{\bar{\alpha}_t} \sqrt{\alpha_t} (T_\phi(\mathbf{x}, t) - T_\phi(\mathbf{x}, t-1)) - \sqrt{\alpha_t} (1 - \sqrt{\alpha_t}) c}{(1 - \alpha_t)} \frac{(1 - \alpha_t)(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t} \right. \\ &\quad \left. + \frac{\sqrt{\bar{\alpha}_{t-1}} T_\phi(\mathbf{x}, t-1) + (1 - \sqrt{\bar{\alpha}_{t-1}}) c}{(1 - \bar{\alpha}_{t-1})} \frac{(1 - \alpha_t)(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t} \right) \\ &= \left( \sqrt{\alpha_t} \mathbf{x}^t - \sqrt{\bar{\alpha}_t} \sqrt{\alpha_t} (T_\phi(\mathbf{x}, t) - T_\phi(\mathbf{x}, t-1)) - \sqrt{\alpha_t} (1 - \sqrt{\alpha_t}) c \right) \frac{(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t} \\ &\quad + \left( \sqrt{\bar{\alpha}_{t-1}} T_\phi(\mathbf{x}, t-1) + (1 - \sqrt{\bar{\alpha}_{t-1}}) c \right) \frac{(1 - \alpha_t)}{1 - \bar{\alpha}_t} \end{aligned} \quad (13)$$

$$\mu_{t-1|t} = \zeta_1 \mathbf{x}^t + \zeta_2 c - T_\phi(\mathbf{x}, t) \zeta_1 \sqrt{\bar{\alpha}_t} + T_\phi(\mathbf{x}, t-1) (\zeta_1 \sqrt{\bar{\alpha}_t} + \zeta_0)$$

Where,

$$\zeta_0 = \frac{1 - \alpha_t}{1 - \bar{\alpha}_t} \sqrt{\bar{\alpha}_{t-1}}; \quad \zeta_1 = \frac{(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t} \sqrt{\alpha_t}; \quad \zeta_2 = \left( 1 + \frac{(\sqrt{\bar{\alpha}_t} - 1)(\sqrt{\alpha_t} + \sqrt{\bar{\alpha}_{t-1}})}{1 - \bar{\alpha}_t} \right);$$

## A.2. Loss Formulation

Although our intention was not to explicitly make the forward process non-Markovian, our formulation has resulted in a non-Markovian forward process Eq.7. So, we can make use of joint distribution from DDIM (Song et al., 2020) as:

$$q(x^{0:T}|x) = q(x^T|x) \prod_{t=1}^T q(x^{t-1}|x^t, x),$$

Next, we define Non-Markovian trainable generative process  $p_\theta(x_{0:T})$  as:

$$p_\theta(x^{0:T}) = p(x^T) \prod_{t=1}^T p_\theta(x^{t-1}|x^t) p_\theta(x^0|x^t), \quad (14)$$

$$\text{where } p(x^T) = \mathcal{N}(x^T; 0, I). \quad (15)$$

This formulation is similar to Equation (55) in DDIM (Song et al., 2020). However, the objective in the referenced work is to minimize the number of sampling time steps; therefore, certain time steps are employed for image generation, and others are included in the variational objective, ultimately demonstrating that the loss formulation aligns with DDPM.

Our definition of trainable reverse generative process (15) is a result of introducing a nonlinear transformation in the marginal distribution, which makes it challenging to learn from the reverse Markovian process. However, even without non-linearity, the resulting loss will be similar to that of DDPM, (Refer Lemma A.1)

Now, by incorporating this into our variational loss term, the variational loss is given as:

$$\begin{aligned} \mathbb{E}_{q_\phi} \left[ \underbrace{D_{\text{KL}}(q_\phi(\mathbf{x}^T|\mathbf{x}, c) \| p(\mathbf{x}^T|c))}_{\mathcal{L}_{\text{prior}}} - \underbrace{\sum_{t=1}^T \log p_\theta(\mathbf{x}|\mathbf{x}^t, c)}_{\mathcal{L}_{\text{rec}}} \right. \\ \left. + \underbrace{\sum_{t=1}^T D_{\text{KL}}(q_\phi(\mathbf{x}^{t-1}|\mathbf{x}^t, \mathbf{x}, c) \| p_\theta(\mathbf{x}^{t-1}|\mathbf{x}^t, c))}_{\mathcal{L}_{\text{diff}}} \right]. \end{aligned} \quad (16)$$

**Lemma A.1.** For  $p_\theta(x^{0:T})$  defined in Eq. 15, variational lower bound results in DDPM training objective.

*Proof.* Based on (Ho et al., 2020), the variational objective ( $J_{x_\theta}$ ) can be expressed as follows:

$$\begin{aligned}
 J_{x_\theta} &= \mathbb{E}_q [\log q(x^{1:T} | x_0) - \log p_\theta(x^{0:T})] \\
 &= \mathbb{E}_q \left[ \log q(x^T | x^0) + \sum_{t=1}^T \log q(x^{t-1} | x^t, x^0) - \sum_{t=1}^T \log p_\theta(x^{t-1} | x^t) - \sum_{t=1}^T \log p_\theta(x^0 | x^t) - \log p_\theta(x^T) \right] \\
 &= \mathbb{E}_q \left[ \underbrace{D_{\text{KL}}(q(x^T | x^0) \| p(x^T))}_{\mathcal{L}_{\text{prior}}} + \underbrace{\sum_{t=1}^T D_{\text{KL}}(q(x^{t-1} | x^t, x^0) \| p_\theta(x^{t-1} | x^t))}_{\mathcal{L}_{\text{diff}}} \right] + \underbrace{\mathbb{E}_q \left[ - \sum_{t=1}^T \log p_\theta(x^0 | x^t) \right]}_{\mathcal{L}_{\text{rec}}} \quad (17) \\
 &= \mathcal{L}_{\text{prior}} + \sum_{t=1}^T \mathbb{E}_q \left[ \frac{1}{2\bar{\sigma}_t^2} \|x^0 - x_\theta(x^t, t)\|^2 \right] + \sum_{t=1}^T \mathbb{E}_q \left[ \frac{1}{2\sigma_t^2} \|x^0 - x_\theta(x^t, t)\|^2 \right] \\
 &\approx \underbrace{\|x^0 - x_\theta(x^t, t)\|^2}_{\mathcal{L}_{\text{DDPM}}}
 \end{aligned}$$

Therefore, When predicting  $x^0$ , with a trainable non markovian reverse process defined in Eq. 15 will result in DDPM training objective.

### A.3. Variational Objective

To calculate the diffusion term  $\mathcal{L}_{\text{diff}}$  of the objective (4), we need to compute the KL divergence between the forward posterior distribution  $q_\phi(\mathbf{x}^{t-1} | \mathbf{x}^t, \mathbf{x}, c)$  and the reverse distribution  $p_\theta(\mathbf{x}^{t-1} | \mathbf{x}^t, c)$ . Since we use parameterization  $p_\theta(\mathbf{x}^{t-1} | \mathbf{x}^t, c) = q_\phi(\mathbf{x}^{t-1} | \mathbf{x}^t, \hat{\mathbf{x}}_\theta(x^t, t), c)$ , both of these distributions are normal distributions with the same variance, so we can evaluate the KL divergence between them analytically as follows:

$$\begin{aligned}
 D_{\text{KL}}(q_\phi(\mathbf{x}^{t-1} | \mathbf{x}^t, \mathbf{x}, c) \| p_\theta(\mathbf{x}^{t-1} | \mathbf{x}^t, c)) &= \\
 \frac{1}{2\sigma_{t-1|t}^2} \left\| \zeta_1 \sqrt{\bar{\alpha}_t} (T_\phi(\mathbf{x}, t) - T_\phi(\hat{\mathbf{x}}_\theta(\mathbf{x}^t, t), t)) - (\zeta_1 \sqrt{\bar{\alpha}_t} + \zeta_0) (T_\phi(\hat{\mathbf{x}}_\theta(\mathbf{x}^t, t-1), t) - T_\phi(\mathbf{x}, t-1)) \right\|_2^2. \quad (18)
 \end{aligned}$$

We can compute the prior term as follows:

$$\begin{aligned}
 D_{\text{KL}}(q_\phi(\mathbf{x}^T | \mathbf{x}, c) \| p(\mathbf{x}^T | c)) &= \frac{1}{2} \left[ \log \frac{|I|}{\sigma_T^2 I} - d + \text{Tr}\{I^{-1} \sigma_T^2 I\} + \|\sqrt{\bar{\alpha}_T} T_\phi(\mathbf{x}, T) + (1 - \sqrt{\bar{\alpha}_T})c - c\|_2^2 \right] \\
 &= \frac{1}{2} \left[ -d \log \sigma_T^2 + d \sigma_T^2 + \|\sqrt{\bar{\alpha}_T} T_\phi(\mathbf{x}, T) - \sqrt{\bar{\alpha}_T} c\|_2^2 \right] \\
 &= \frac{1}{2} \left( d (\sigma_T^2 - \log \sigma_T^2 - 1) + \bar{\alpha}_T \|T_\phi(\mathbf{x}, T) - c\|_2^2 \right). \\
 &= \frac{1}{2} \bar{\alpha}_T \|T_\phi(\mathbf{x}, T) - c\|_2^2. \quad (19)
 \end{aligned}$$

The reconstruction term is considered for all latent variables similar to VAE(Kingma, 2013b)

## B. Implementation Details

The MLP modules within the embeddings of the condition and noised input, as well as the decoder, are from TiDE (Das et al., 2023a). This MLP framework is acknowledged as an effective component for universal time series analysis models (Das et al., 2023b). The encoder employs the diffusion transformer block (DiT), with the diffusion time step adjusted by adaptive layer normalization (Peebles & Xie, 2023; Esser et al., 2024). Hyperparameters are tuned specifically regarding embedding layers, encoder layers, and hidden dimensions while other hyperparameters remain constant as detailed in Section 4.1.2. Optimal hyperparameters are selected via cross-validation and summarized for various datasets in Table B.

Table 6. Best hyperparameters for different datasets

datasets	hidden dimension	embedding layers	encoder layers
NorPool	256	1	2
Caiso	128	1	2
Weather	64	1	1
ETTm1	512	1	1
Wind	512	1	1
Traffic	64	2	2
Electricity	256	2	1
ETTh1	256	2	3
Exchange	64	2	1

### C. Baselines

We have thoroughly benchmarked our CN-Diff model against 23 baseline models using a variety of methodologies, from basis expansion techniques to diffusion models. These include: (i) diffusion-based models (Shen et al., 2024; Shen & Kwok, 2023; Rasul et al., 2021; Tashiro et al., 2021; Lopez Alcaraz & Strodthoff, 2023); (ii) basis expansion methods such as those detailed in (Challu et al., 2023; Zhou et al., 2022a; Fan et al., 2022; Oreshkin et al., 2019); (iii) alternative generative techniques like GAN and VAE as noted in (Li et al., 2022a; Rangapuram et al., 2023; Jeha et al., 2022). We also evaluated CN-Diff against transformers (Shabani et al., 2023; Nie et al., 2023; Zhou et al., 2022b; Wu et al., 2021; Liu et al., 2022b; Zhou et al., 2021; Vaswani, 2017b) and other deep learning approaches (Liu et al., 2022a; Zeng et al., 2023; Bahdanau et al., 2015). Note that we don’t implement these baseline models. Instead, we rely on the reported numbers in literature.

### D. Results of MAE

Tables 8 and 9 present the mean absolute error (MAE) results for multivariate and univariate time series forecasting tasks. It is evident that CN-Diff continues to demonstrate a superior overall performance for MAE as the evaluation metric. The average rank difference for multivariate is significant, but for univariate, it is almost comparable to mr-Diff. Furthermore, acknowledging that the performance of deep models in time series forecasting can be affected by various random initializations, Table 7 illustrates the prediction results for univariate time series in five distinct random runs.

Table 7. Univariate prediction errors of CN-Diff obtained on five runs.

	Exchange		ETTh1		Weather		Electricity	
	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
0	0.0152	0.0933	0.0660	0.1991	0.0025	0.0342	0.2370	0.3465
1	0.0154	0.0934	0.0649	0.1975	0.0021	0.0341	0.2268	0.3387
2	0.0153	0.0936	0.0656	0.1985	0.0025	0.0343	0.2289	0.3426
3	0.0161	0.0950	0.0670	0.2001	0.0027	0.0349	0.2393	0.3431
4	0.0157	0.0943	0.0644	0.1977	0.0024	0.0346	0.2392	0.3429
mean	0.0155	0.0939	0.0656	0.1986	0.0024	0.0344	0.2342	0.3427
std deviation	0.0004	0.0007	0.0010	0.0010	0.0002	0.0003	0.0059	0.0027

Table 8. Multivariate prediction of MAEs on nine real-world time series datasets(subscripts is the rank). CSDI runs out of memory on Traffic and Electricity. Results of all baselines are from (Shen et al., 2024)

	NorPool	Caiso	Traffic	Electricity	Weather	Exchange	ETTh1	ETTm1	Wind	Rank
Ours	<b>0.554</b> <sub>(1)</sub>	<b>0.191</b> <sub>(1)</sub>	0.270 <sub>(4)</sub>	<u>0.243</u> <sub>(2)</sub>	<u>0.324</u> <sub>(2)</sub>	<b>0.079</b> <sub>(1)</sub>	<u>0.421</u> <sub>(2)</sub>	0.378 <sub>(5)</sub>	0.701 <sub>(4)</sub>	2.4
mr-Diff	0.604 <sub>(3)</sub>	0.219 <sub>(4)</sub>	0.320 <sub>(6)</sub>	0.252 <sub>(4)</sub>	<u>0.324</u> <sub>(2)</sub>	0.082 <sub>(4)</sub>	0.422 <sub>(3)</sub>	<u>0.373</u> <sub>(2)</sub>	<b>0.675</b> <sub>(1)</sub>	3.2
TimeDiff	0.611 <sub>(4)</sub>	0.234 <sub>(7)</sub>	0.384 <sub>(9)</sub>	0.305 <sub>(7)</sub>	<b>0.312</b> <sub>(1)</sub>	0.091 <sub>(8)</sub>	0.430 <sub>(4)</sub>	<b>0.372</b> <sub>(1)</sub>	0.687 <sub>(3)</sub>	4.9
TimeGrad	0.821 <sub>(20)</sub>	0.339 <sub>(19)</sub>	0.849 <sub>(23)</sub>	0.630 <sub>(22)</sub>	0.381 <sub>(15)</sub>	0.193 <sub>(20)</sub>	0.719 <sub>(23)</sub>	0.605 <sub>(23)</sub>	0.793 <sub>(22)</sub>	20.8
CSDI	0.777 <sub>(18)</sub>	0.345 <sub>(20)</sub>	-	-	0.374 <sub>(13)</sub>	0.194 <sub>(21)</sub>	0.438 <sub>(6)</sub>	0.442 <sub>(16)</sub>	0.741 <sub>(11)</sub>	15.0
SSSD	0.753 <sub>(14)</sub>	0.295 <sub>(11)</sub>	0.398 <sub>(14)</sub>	0.363 <sub>(13)</sub>	0.350 <sub>(9)</sub>	0.127 <sub>(13)</sub>	0.561 <sub>(18)</sub>	0.406 <sub>(11)</sub>	0.778 <sub>(19)</sub>	13.6
D <sup>3</sup> VAE	0.692 <sub>(10)</sub>	0.331 <sub>(17)</sub>	0.483 <sub>(18)</sub>	0.372 <sub>(15)</sub>	0.380 <sub>(14)</sub>	0.301 <sub>(23)</sub>	0.502 <sub>(15)</sub>	0.391 <sub>(9)</sub>	0.779 <sub>(20)</sub>	15.7
CPF	0.889 <sub>(22)</sub>	0.424 <sub>(22)</sub>	0.714 <sub>(22)</sub>	0.643 <sub>(23)</sub>	0.781 <sub>(24)</sub>	0.082 <sub>(4)</sub>	0.597 <sub>(21)</sub>	0.472 <sub>(17)</sub>	0.757 <sub>(16)</sub>	19.0
PSA-GAN	0.890 <sub>(23)</sub>	0.477 <sub>(23)</sub>	0.697 <sub>(21)</sub>	0.533 <sub>(21)</sub>	0.578 <sub>(23)</sub>	0.087 <sub>(7)</sub>	0.546 <sub>(17)</sub>	0.488 <sub>(19)</sub>	0.756 <sub>(14)</sub>	18.7
N-Hits	0.643 <sub>(8)</sub>	0.221 <sub>(5)</sub>	<u>0.268</u> <sub>(2)</sub>	0.245 <sub>(3)</sub>	0.335 <sub>(5)</sub>	0.085 <sub>(6)</sub>	0.480 <sub>(9)</sub>	0.388 <sub>(7)</sub>	0.734 <sub>(9)</sub>	6.0
FILM	0.646 <sub>(9)</sub>	0.278 <sub>(9)</sub>	0.398 <sub>(14)</sub>	0.320 <sub>(9)</sub>	0.336 <sub>(6)</sub>	<b>0.079</b> <sub>(1)</sub>	0.436 <sub>(5)</sub>	0.374 <sub>(3)</sub>	0.717 <sub>(6)</sub>	6.9
Depts	0.611 <sub>(4)</sub>	0.204 <sub>(3)</sub>	0.568 <sub>(20)</sub>	0.401 <sub>(18)</sub>	0.394 <sub>(17)</sub>	0.100 <sub>(10)</sub>	0.491 <sub>(13)</sub>	0.412 <sub>(13)</sub>	0.751 <sub>(13)</sub>	12.3
NBeats	0.832 <sub>(21)</sub>	0.235 <sub>(8)</sub>	<b>0.265</b> <sub>(1)</sub>	0.370 <sub>(14)</sub>	0.420 <sub>(18)</sub>	<u>0.081</u> <sub>(3)</sub>	0.521 <sub>(16)</sub>	0.409 <sub>(12)</sub>	0.741 <sub>(11)</sub>	11.6
Scaleformer	0.769 <sub>(17)</sub>	0.310 <sub>(13)</sub>	0.379 <sub>(8)</sub>	0.304 <sub>(6)</sub>	0.438 <sub>(19)</sub>	0.138 <sub>(15)</sub>	0.579 <sub>(20)</sub>	0.475 <sub>(18)</sub>	0.864 <sub>(23)</sub>	15.4
PatchTST	0.710 <sub>(11)</sub>	0.293 <sub>(10)</sub>	0.411 <sub>(17)</sub>	0.348 <sub>(12)</sub>	0.555 <sub>(22)</sub>	0.147 <sub>(16)</sub>	0.489 <sub>(12)</sub>	0.392 <sub>(10)</sub>	0.720 <sub>(7)</sub>	13.0
FedFormer	0.744 <sub>(12)</sub>	0.317 <sub>(14)</sub>	0.385 <sub>(10)</sub>	0.341 <sub>(11)</sub>	0.347 <sub>(8)</sub>	0.233 <sub>(22)</sub>	0.484 <sub>(10)</sub>	0.413 <sub>(14)</sub>	0.762 <sub>(17)</sub>	13.1
Autoformer	0.751 <sub>(13)</sub>	0.321 <sub>(15)</sub>	0.392 <sub>(13)</sub>	0.313 <sub>(8)</sub>	0.354 <sub>(10)</sub>	0.167 <sub>(17)</sub>	0.484 <sub>(10)</sub>	0.496 <sub>(20)</sub>	0.756 <sub>(14)</sub>	13.3
Pyraformer	0.781 <sub>(19)</sub>	0.371 <sub>(21)</sub>	0.390 <sub>(11)</sub>	0.379 <sub>(16)</sub>	0.385 <sub>(16)</sub>	0.112 <sub>(12)</sub>	0.493 <sub>(14)</sub>	0.435 <sub>(15)</sub>	0.735 <sub>(10)</sub>	14.9
Informer	0.757 <sub>(15)</sub>	0.336 <sub>(18)</sub>	0.391 <sub>(12)</sub>	0.383 <sub>(17)</sub>	0.364 <sub>(11)</sub>	0.192 <sub>(19)</sub>	0.605 <sub>(22)</sub>	0.542 <sub>(21)</sub>	0.772 <sub>(18)</sub>	17.0
Transformer	0.765 <sub>(16)</sub>	0.321 <sub>(15)</sub>	0.410 <sub>(16)</sub>	0.405 <sub>(19)</sub>	0.370 <sub>(12)</sub>	0.178 <sub>(18)</sub>	0.567 <sub>(19)</sub>	0.592 <sub>(22)</sub>	0.785 <sub>(21)</sub>	17.6
SCINet	<u>0.601</u> <sub>(2)</sub>	<u>0.193</u> <sub>(2)</sub>	0.335 <sub>(7)</sub>	0.280 <sub>(5)</sub>	0.344 <sub>(7)</sub>	0.137 <sub>(14)</sub>	0.463 <sub>(8)</sub>	0.389 <sub>(8)</sub>	0.732 <sub>(8)</sub>	6.8
NLinear	0.636 <sub>(6)</sub>	0.223 <sub>(6)</sub>	0.293 <sub>(5)</sub>	<b>0.239</b> <sub>(1)</sub>	0.328 <sub>(4)</sub>	0.091 <sub>(8)</sub>	<b>0.418</b> <sub>(1)</sub>	0.375 <sub>(4)</sub>	0.706 <sub>(5)</sub>	4.4
DLinear	0.640 <sub>(7)</sub>	0.497 <sub>(24)</sub>	<u>0.268</u> <sub>(2)</sub>	0.336 <sub>(10)</sub>	0.444 <sub>(20)</sub>	0.102 <sub>(11)</sub>	0.442 <sub>(7)</sub>	0.378 <sub>(5)</sub>	<u>0.686</u> <sub>(2)</sub>	9.8
LSTMa	0.974 <sub>(24)</sub>	0.305 <sub>(12)</sub>	0.510 <sub>(19)</sub>	0.444 <sub>(20)</sub>	0.501 <sub>(21)</sub>	0.534 <sub>(21)</sub>	0.782 <sub>(24)</sub>	0.699 <sub>(24)</sub>	0.897 <sub>(24)</sub>	21.0

Table 9. Univariate prediction of MAEs on nine real-world time series datasets(subscripts is the rank). Results of all baselines are from (Shen et al., 2024)

Method	NorPool	Caiso	Traffic	Electricity	Weather	Exchange	ETTh1	ETTm1	Wind	Rank
Ours	<b>0.590</b> <sub>(1)</sub>	<u>0.206</u> <sub>(2)</sub>	0.206 <sub>(3)</sub>	0.343 <sub>(4)</sub>	0.034 <sub>(4)</sub>	<b>0.094</b> <sub>(1)</sub>	0.199 <sub>(2)</sub>	<b>0.149</b> <sub>(1)</sub>	<u>1.168</u> <sub>(2)</sub>	2.2
mr-Diff	0.609 <sub>(3)</sub>	0.212 <sub>(5)</sub>	<u>0.197</u> <sub>(2)</sub>	<b>0.332</b> <sub>(1)</sub>	<b>0.032</b> <sub>(1)</sub>	<b>0.094</b> <sub>(1)</sub>	<b>0.196</b> <sub>(1)</sub>	<b>0.149</b> <sub>(1)</sub>	<u>1.168</u> <sub>(2)</sub>	1.9
TimeDiff	0.613 <sub>(4)</sub>	0.209 <sub>(4)</sub>	0.207 <sub>(4)</sub>	0.341 <sub>(3)</sub>	0.035 <sub>(5)</sub>	0.102 <sub>(8)</sub>	0.202 <sub>(3)</sub>	0.154 <sub>(7)</sub>	1.209 <sub>(6)</sub>	4.9
TimeGrad	0.841 <sub>(24)</sub>	0.386 <sub>(23)</sub>	0.894 <sub>(24)</sub>	0.898 <sub>(24)</sub>	0.036 <sub>(7)</sub>	0.155 <sub>(22)</sub>	0.212 <sub>(9)</sub>	0.167 <sub>(13)</sub>	1.239 <sub>(12)</sub>	17.6
CSDI	0.763 <sub>(21)</sub>	0.282 <sub>(15)</sub>	0.468 <sub>(21)</sub>	0.540 <sub>(20)</sub>	0.037 <sub>(8)</sub>	0.200 <sub>(24)</sub>	0.221 <sub>(13)</sub>	0.170 <sub>(15)</sub>	1.218 <sub>(8)</sub>	16.1
SSSD	0.770 <sub>(22)</sub>	0.263 <sub>(13)</sub>	0.226 <sub>(7)</sub>	0.403 <sub>(9)</sub>	0.041 <sub>(12)</sub>	0.118 <sub>(17)</sub>	0.250 <sub>(21)</sub>	0.169 <sub>(14)</sub>	1.356 <sub>(23)</sub>	15.3
D <sup>3</sup> VAE	0.774 <sub>(23)</sub>	0.613 <sub>(24)</sub>	0.237 <sub>(10)</sub>	0.539 <sub>(19)</sub>	0.039 <sub>(10)</sub>	0.107 <sub>(14)</sub>	0.221 <sub>(13)</sub>	0.160 <sub>(10)</sub>	1.321 <sub>(20)</sub>	15.9
CPF	0.710 <sub>(16)</sub>	0.338 <sub>(19)</sub>	0.385 <sub>(20)</sub>	0.592 <sub>(22)</sub>	0.035 <sub>(5)</sub>	<b>0.094</b> <sub>(1)</sub>	0.221 <sub>(13)</sub>	0.153 <sub>(6)</sub>	1.256 <sub>(13)</sub>	12.8
PSA-GAN	0.623 <sub>(6)</sub>	0.250 <sub>(9)</sub>	0.355 <sub>(18)</sub>	0.373 <sub>(7)</sub>	0.139 <sub>(23)</sub>	0.109 <sub>(15)</sub>	0.225 <sub>(17)</sub>	0.174 <sub>(17)</sub>	1.287 <sub>(17)</sub>	14.3
N-Hits	0.646 <sub>(8)</sub>	0.276 <sub>(14)</sub>	0.232 <sub>(8)</sub>	0.419 <sub>(10)</sub>	<u>0.033</u> <sub>(2)</sub>	0.100 <sub>(6)</sub>	0.228 <sub>(18)</sub>	0.157 <sub>(9)</sub>	1.256 <sub>(13)</sub>	9.8
FILM	0.654 <sub>(10)</sub>	0.290 <sub>(16)</sub>	0.315 <sub>(15)</sub>	0.362 <sub>(6)</sub>	0.069 <sub>(15)</sub>	0.104 <sub>(11)</sub>	0.210 <sub>(7)</sub>	<b>0.149</b> <sub>(1)</sub>	1.189 <sub>(4)</sub>	9.4
Depts	0.616 <sub>(5)</sub>	<b>0.205</b> <sub>(1)</sub>	0.241 <sub>(11)</sub>	0.434 <sub>(13)</sub>	0.102 <sub>(20)</sub>	0.106 <sub>(13)</sub>	0.202 <sub>(3)</sub>	0.165 <sub>(11)</sub>	1.472 <sub>(24)</sub>	11.2
NBeats	0.671 <sub>(11)</sub>	0.228 <sub>(6)</sub>	0.225 <sub>(6)</sub>	0.439 <sub>(14)</sub>	0.130 <sub>(22)</sub>	0.096 <sub>(4)</sub>	0.242 <sub>(19)</sub>	0.165 <sub>(11)</sub>	1.236 <sub>(10)</sub>	11.4
Scaleformer	0.687 <sub>(13)</sub>	0.320 <sub>(17)</sub>	0.375 <sub>(19)</sub>	0.430 <sub>(11)</sub>	0.083 <sub>(18)</sub>	0.148 <sub>(20)</sub>	0.302 <sub>(23)</sub>	0.210 <sub>(23)</sub>	1.348 <sub>(22)</sub>	18.4
PatchTST	<b>0.590</b> <sub>(1)</sub>	0.260 <sub>(12)</sub>	0.269 <sub>(12)</sub>	0.478 <sub>(18)</sub>	0.098 <sub>(19)</sub>	0.111 <sub>(16)</sub>	0.260 <sub>(22)</sub>	0.174 <sub>(17)</sub>	1.338 <sub>(21)</sub>	15.3
FedFormer	0.725 <sub>(18)</sub>	0.254 <sub>(10)</sub>	0.278 <sub>(13)</sub>	0.453 <sub>(15)</sub>	0.057 <sub>(14)</sub>	0.168 <sub>(23)</sub>	0.212 <sub>(9)</sub>	0.195 <sub>(21)</sub>	1.271 <sub>(15)</sub>	15.3
Autoformer	0.755 <sub>(20)</sub>	0.339 <sub>(20)</sub>	0.495 <sub>(22)</sub>	0.623 <sub>(23)</sub>	0.040 <sub>(11)</sub>	0.152 <sub>(21)</sub>	0.220 <sub>(12)</sub>	0.174 <sub>(17)</sub>	1.319 <sub>(19)</sub>	18.3
Pyraformer	0.747 <sub>(19)</sub>	0.257 <sub>(11)</sub>	0.215 <sub>(5)</sub>	0.455 <sub>(16)</sub>	0.107 <sub>(21)</sub>	0.104 <sub>(11)</sub>	0.211 <sub>(8)</sub>	0.179 <sub>(20)</sub>	1.284 <sub>(16)</sub>	13.1
Informer	0.698 <sub>(14)</sub>	0.345 <sub>(21)</sub>	0.308 <sub>(14)</sub>	0.433 <sub>(12)</sub>	0.069 <sub>(15)</sub>	0.118 <sub>(17)</sub>	0.212 <sub>(9)</sub>	0.172 <sub>(16)</sub>	1.236 <sub>(10)</sub>	14.1
Transformer	0.723 <sub>(17)</sub>	0.345 <sub>(21)</sub>	0.336 <sub>(17)</sub>	0.469 <sub>(17)</sub>	0.071 <sub>(17)</sub>	0.103 <sub>(10)</sub>	0.247 <sub>(20)</sub>	0.196 <sub>(22)</sub>	1.212 <sub>(7)</sub>	16.4
SCINet	0.653 <sub>(9)</sub>	0.244 <sub>(8)</sub>	0.322 <sub>(16)</sub>	0.377 <sub>(8)</sub>	0.037 <sub>(8)</sub>	0.101 <sub>(7)</sub>	0.205 <sub>(6)</sub>	0.150 <sub>(5)</sub>	<b>1.167</b> <sub>(1)</sub>	7.6
NLinear	0.637 <sub>(7)</sub>	0.238 <sub>(7)</sub>	<b>0.192</b> <sub>(1)</sub>	<u>0.334</u> <sub>(2)</sub>	<u>0.033</u> <sub>(2)</sub>	0.097 <sub>(5)</sub>	0.203 <sub>(5)</sub>	<b>0.149</b> <sub>(1)</sub>	1.197 <sub>(5)</sub>	3.9
DLinear	0.671 <sub>(11)</sub>	<u>0.206</u> <sub>(2)</sub>	0.236 <sub>(9)</sub>	0.348 <sub>(5)</sub>	0.310 <sub>(24)</sub>	0.102 <sub>(8)</sub>	0.222 <sub>(16)</sub>	0.155 <sub>(8)</sub>	1.221 <sub>(9)</sub>	10.2
LSTMa	0.707 <sub>(15)</sub>	0.333 <sub>(18)</sub>	0.757 <sub>(22)</sub>	0.557 <sub>(21)</sub>	0.053 <sub>(13)</sub>	0.136 <sub>(19)</sub>	0.332 <sub>(24)</sub>	0.239 <sub>(24)</sub>	1.298 <sub>(18)</sub>	19.3

## E. Ablations

In this section, we present ablation studies for various hyperparameters, beginning with an examination of different components within CN-Diff, as well as the look-back window, diffusion time steps, and Gaussian variance schedule.

### E.1. Architecture components

Table E.1 presents an analysis of the various components integrated into CN-Diff. Beginning with DDPM (Ho et al., 2020), we provide results that illustrate the effects of incorporating only non-linearity, the conditional aspect independently, and the combined integration of both into the DDPM, evaluated across multiple datasets.

Table 10. Predicting MSE and MAE for different components of CN-Diff for Multivariate forecasting (improvement relative to Diffusion is denoted within brackets.)

Methods	Diffusion		+ $T_\phi$		+ CN-Diff Cond		All (CN-Diff)	
Datasets	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
ETTh1	1.748	0.927	1.320(25%)	0.730(21%)	0.943(46%)	0.704(24%)	<b>0.405</b> (77%)	<b>0.421</b> (55%)
Exchange	0.090	0.219	0.016(82%)	0.080(63%)	0.021(60%)	0.101(54%)	<b>0.015</b> (83%)	<b>0.079</b> (64%)
Caiso	0.586	0.553	0.403(31%)	0.433(22%)	0.242(59%)	0.307(44%)	<b>0.094</b> (84%)	<b>0.191</b> (65%)
Norpool	1.819	1.047	1.441(42%)	0.918(12%)	1.142(37%)	0.814(22%)	<b>0.531</b> (71%)	<b>0.554</b> (47%)
Electricity	1.824	1.030	1.597(12%)	0.951(8%)	0.536(71%)	0.512(50%)	<b>0.145</b> (92%)	<b>0.243</b> (76%)
Traffic	1.794	0.885	1.397(22%)	0.799(10%)	0.848(53%)	0.479(46%)	<b>0.374</b> (79%)	<b>0.270</b> (70%)

### E.2. Look back Window

Table E.2 presents the prediction MSE and MAE for CN-Diff utilizing various lengths of the lookback window. The values considered are  $L = \{96, 192, 336, 720, 1440\}$ . It is evident that on the datasets Electricity, ETTh1, and weather (corresponding to 168-step, 168-step, and 672-step-ahead predictions, respectively), satisfactory performance is achieved when  $L$  is set to 336 or higher.

Table 11. Predicting MSE and MAE for different Lookback windows( $L$ ) for Multivariate forecasting

Look back window	96		192		336		720		1440	
Datasets	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
ETTh1	0.433	0.437	0.426	0.437	<b>0.405</b>	<b>0.421</b>	0.434	0.448	0.440	0.452
Weather	0.530	0.503	0.347	0.367	0.339	0.354	0.307	0.335	<b>0.296</b>	<b>0.324</b>
Electricity	0.167	0.262	0.152	0.248	0.149	0.246	0.148	0.246	<b>0.145</b>	<b>0.243</b>

### E.3. Diffusion Parameters

Diffusion steps, denoted as  $T$ , is a critical parameter. Table E.3 presents the MSE and MAE of CN-Diff as a function of  $T$ . It is evident that establishing  $T = 100$  results in consistent performance in all three datasets examined. This observation is consistent with the findings in (Li et al., 2022b), which suggest that a minimal value of  $T$  may result in incomplete diffusion, while an excessively large  $K$  might lead to superfluous computational overhead.(For all datasets analyzed, the best length of the look-back window was considered.)

The variance of Gaussian noise  $\beta_k$ , is formulated through the implementation of a quadratic variance schedule. In the present study, the variable  $\beta_T$  is varied within the set  $\{0.001, 0.01, 0.1, 0.9\}$ , while  $\beta_1$  is maintained at a constant value of  $10^{-4}$ . As evidenced in Table E.3, optimizing the parameter setting  $\beta_T = 0.1$  consistently results in enhanced performance.



Table 12. Predicting MSE and MAE for different diffusion noise steps( $T$ ) for Multivariate forecasting

Noise Steps	50		100		200		500	
Datasets	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
ETTh1	0.415	0.431	<b>0.405</b>	<b>0.421</b>	0.409	0.426	0.422	0.437
Weather	0.299	0.330	<b>0.296</b>	<b>0.324</b>	0.325	0.360	0.362	0.378
Electricity	0.145	0.244	<b>0.145</b>	<b>0.243</b>	0.145	0.243	0.147	0.244

 Table 13. Predicting MSE and MAE for different variance upper bound ( $\beta_T$ ) for Multivariate forecasting

$\beta_T$	0.001		0.01		0.1		0.9	
Datasets	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
ETTh1	0.807	0.656	0.462	0.463	<b>0.405</b>	<b>0.421</b>	0.425	0.438
Weather	0.312	0.347	0.340	0.375	<b>0.296</b>	<b>0.324</b>	0.357	0.373
Electricity	1.745	0.986	0.243	0.345	<b>0.145</b>	<b>0.243</b>	0.146	0.243

## F. Qualitative analysis




Figure 4. Visualization of CN-Diff forecasts across various datasets

The qualitative analysis presented in Figure 4 demonstrates the forecasting capabilities of our CN-Diff model across four diverse datasets: Weather, Caiso, Electricity, and Norpool. These visualizations provide strong evidence of the model’s effectiveness in capturing temporal patterns and generating accurate predictions across varying domains. In Figure 4(a), the Weather dataset results show CN-Diff successfully tracking the overall trend and seasonal patterns, with predictions (green line) closely following the ground truth (blue line). For the Caiso dataset in Figure 4(b), the model effectively captures the high-frequency oscillations that characterize this time series, with particularly strong alignment between the predicted and actual values in the latter portions of the sequence. The Electricity dataset results in Figure 4(c) further illustrate CN-Diff’s ability to model periodic patterns with consistent amplitude and frequency matching the ground truth. Finally, Figure 4(d) showcases results on the Norpool dataset, where our model accurately predicts the irregular fluctuations and demonstrates robustness to the dataset’s inherent volatility.

## G. Computational analysis

We present the computational cost analysis to compare our model with other diffusion-based time series forecasting methods. Refer to the tables for training (Table:14) and inference (Table:15) comparisons (Results for other baselines are taken from (Shen et al., 2024)). Note that, SSSD utilizes structural state-space-based diffusion layers along with multiple dense layers; while TimeGrad leverages RNN hidden states; TimeDiff incorporates future mean and autoregressive initialization; and mr-Diff employs multiple diffusion models. In contrast, our innovation lies in the diffusion formulation rather than in the architectural modifications. Both key components we introduced, the nonlinear transform ( $T_\phi(\cdot)$ ) and the condition network, consist of simple linear layers with an activation function. This design choice allows us to achieve greater computational efficiency compared to existing diffusion models for time series forecasting, and this can be realized by the table below.

Table 14. Training time (ms) for different models and sequence lengths ( $H$ ) for ETTh1 univariate.

	$H = 96$	$H = 168$	$H = 192$	$H = 336$	$H = 720$
<b>CN-Diff</b>	<b>0.21</b>	<b>0.26</b>	<b>0.27</b>	<b>0.31</b>	<b>0.39</b>
mr-Diff	0.59	0.69	0.71	0.74	0.82
TimeDiff	0.71	0.75	0.77	0.82	0.85
TimeGrad	2.11	2.42	3.21	4.22	5.93
CSDI	5.72	7.09	7.59	10.59	17.21
SSSD	16.98	19.34	22.64	32.12	52.93

Table 15. Inference time (ms) for different models and sequence lengths ( $H$ ) for ETTh1 univariate.

Model	Trainable Params	$H = 96$	$H = 168$	$H = 192$	$H = 336$	$H = 720$
<b>CN-Diff</b>	<b>1.1M</b>	<b>6.2</b>	<b>6.7</b>	<b>6.9</b>	<b>7.8</b>	<b>9.1</b>
mr-Diff	1.4M	12.5	14.3	14.9	16.8	27.5
TimeDiff	1.7M	16.2	17.3	17.6	26.5	34.6
TimeGrad	3.1M	870.2	1620.9	1854.5	3119.7	6724.1
CSDI	10M	90.4	128.3	142.8	398.9	513.1
SSSD	32M	418.6	590.2	645.4	1054.2	2516.9

## H. Additional Results

### H.1. Additional Datasets

In order to further validate our approach, we have carried out supplementary experiments employing multivariate synthetic data. This data utilizes Gaussian process kernels as described by (Taga et al., 2025). Additionally, we incorporated the CSI-300 dataset, which represents the price-weighted average index of 300 prominent companies listed on the Shanghai and Shenzhen Stock Exchanges, as cited in (Jia et al., 2024). The results of these experiments, detailed in Table 16 below, indicate that our model achieves a higher level of performance when benchmarked against other existing models.

Table 16. Performance comparison of CN-Diff across models on Stock and Synthetic datasets using MSE and MAE metrics.

Model	Stock (MSE)	Stock (MAE)	Synthetic (MSE)	Synthetic (MAE)
<b>CN-Diff</b>	<b>0.020</b>	<b>0.109</b>	<b>0.722</b>	<b>0.308</b>
Autoformer	0.058	0.190	0.767	0.339
PatchTST	0.028	0.129	0.771	0.345
FiLM	0.040	0.159	0.752	0.319
DLinear	0.044	0.170	0.747	0.311
CSDI	0.037	0.169	0.809	0.506

## H.2. Additional Models

In our experiments, we have compared our method with CPF(Rangapuram et al., 2023), a model using Graph Neural Networks. To further strengthen our analysis, we now incorporate a comparison with SageFormer(Zhang et al., 2024) as well, which is a series-aware graph-enhanced Transformer model for time series forecasting. The table below 17 demonstrates CN-Diff’s superiority over graph-based methods in time series forecasting.

Table 17. Comparison of Model Performance Metrics across different Datasets

Dataset	CN-Diff (MSE/MAE)	Sage Former (MSE/MAE)
ETTm1	<b>0.340 / 0.378</b>	0.370 / 0.388
Electricity	<b>0.145 / 0.243</b>	0.159 / 0.255
Stock	<b>0.020 / 0.109</b>	0.030 / 0.135
Exchange	<b>0.016 / 0.079</b>	0.016 / 0.081
ETTh1	<b>0.405 / 0.421</b>	0.421 / <b>0.419</b>
Weather	<b>0.296 / 0.324</b>	0.324 / 0.344

## H.3. Analysis of Non-linear Transformation( $T_\phi(\cdot)$ )

In Table 5 and E.1, we have shown empirical evidence through the ablation results presented that CN-Diff performs better than the ablation run without  $T_\phi$ . To understand what is learned by  $T_\phi$ , we take a trained model and explore the correlation between the features in the learned latent representation space of  $T_\phi(x, t)$  at different diffusion model timesteps. We observed that there are increasing correlations between the features in learned latent space at different timesteps as shown in tables below. Thus, we hypothesize that this increased correlation and time-dependent adaptability perhaps facilitate a more effective diffusion process for time series forecasting. Similar observations of correlation in latent space that help diffusion models have been made in image and video diffusion works (Ge et al., 2023).

In the space of diffusion models for time series, paper (Tashiro et al., 2021) has noted that learning the correlation between feature and temporal space is necessary for time series imputation tasks. We will add these to the discussion in the ablation study section in the revision.

Table 18. Feature Correlation Matrix for Actual input(ETTh1 dataset)

Features	0	1	2	3	4	5	6
0	<b>1.00</b>	0.08	<b>0.99</b>	0.07	0.22	-0.20	-0.19
1	0.08	<b>1.00</b>	0.08	<b>0.94</b>	0.05	0.22	-0.06
2	<b>0.99</b>	0.08	<b>1.00</b>	0.09	0.12	-0.27	-0.19
3	0.07	<b>0.94</b>	0.09	<b>1.00</b>	-0.10	-0.06	-0.07
4	0.22	0.05	0.12	-0.10	<b>1.00</b>	<b>0.56</b>	-0.06
5	-0.20	0.22	-0.27	-0.06	<b>0.56</b>	<b>1.00</b>	0.11
6	-0.19	-0.06	-0.19	-0.07	-0.06	0.11	<b>1.00</b>

Table 19. Feature Correlation Matrix for Transformed Input  $T_\phi(\cdot)$ (ETTh1 dataset)

Features	0	1	2	3	4	5	6
0	<b>1.00</b>	<b>-0.52</b>	<b>-0.60</b>	<b>0.99</b>	<b>0.93</b>	<b>-0.96</b>	<b>-0.56</b>
1	<b>-0.52</b>	<b>1.00</b>	-0.31	<b>-0.58</b>	<b>-0.68</b>	<b>0.60</b>	<b>0.63</b>
2	<b>-0.60</b>	-0.31	<b>1.00</b>	<b>-0.57</b>	-0.33	<b>0.55</b>	-0.08
3	<b>0.99</b>	<b>-0.58</b>	<b>-0.57</b>	<b>1.00</b>	<b>0.94</b>	<b>-0.97</b>	<b>-0.59</b>
4	<b>0.93</b>	<b>-0.68</b>	-0.33	<b>0.94</b>	<b>1.00</b>	<b>-0.86</b>	<b>-0.80</b>
5	<b>-0.96</b>	<b>0.60</b>	<b>0.55</b>	<b>-0.97</b>	<b>-0.86</b>	<b>1.00</b>	0.43
6	<b>-0.56</b>	<b>0.63</b>	-0.08	<b>-0.59</b>	<b>-0.80</b>	0.43	<b>1.00</b>