

# *TSGDiff*: Rethinking Synthetic Time Series Generation from a Pure Graph Perspective

Lifeng Shen<sup>1\*</sup>, Xuyang Li<sup>1</sup>, Lele Long<sup>1</sup>

<sup>1</sup>Chongqing Key Laboratory of Computational Intelligence, Key Laboratory of Cyberspace Big Data Intelligent Security, Ministry of Education, Sichuan-Chongqing Co-construction Key Laboratory of Digital Economy Intelligence and Key Laboratory of Big Data Intelligent Computing, Chongqing University of Posts and Telecommunications  
shenlf@cqupt.edu.cn, lixuyang.lee@foxmail.com, longlele0208@foxmail.com

## Abstract

Diffusion models have shown great promise in data generation, yet generating time series data remains challenging due to the need to capture complex temporal dependencies and structural patterns. In this paper, we present *TSGDiff*, a novel framework that rethinks time series generation from a graph-based perspective. Specifically, we represent time series as dynamic graphs, where edges are constructed based on Fourier spectrum characteristics and temporal dependencies. A graph neural network-based encoder-decoder architecture is employed to construct a latent space, enabling the diffusion process to model the structural representation distribution of time series effectively. Furthermore, we propose the Topological Structure Fidelity (Topo-FID) score, a graph-aware metric for assessing the structural similarity of time series graph representations. Topo-FID integrates two sub-metrics: Graph Edit Similarity, which quantifies differences in adjacency matrices, and Structural Entropy Similarity, which evaluates the entropy of node degree distributions. This comprehensive metric provides a more accurate assessment of structural fidelity in generated time series. Experiments on real-world datasets demonstrate that *TSGDiff* generates high-quality synthetic time series data generation, faithfully preserving temporal dependencies and structural integrity, thereby advancing the field of synthetic time series generation.

**Code** — <https://github.com/jvaeylee/TSGDiff>

## 1 Introduction

Multivariate time series generation is of great importance in various domains, such as energy management (Chen et al. 2017), financial market forecasting (Bao, Yue, and Rao 2017), and medical monitoring (Zammel et al. 2023). Traditional generative models, including generative adversarial networks (GANs) (Esteban, Hyland, and Ratsch 2017) and variational autoencoders (VAEs) (Fortuin et al. 2020), have made significant strides in this area. However, they still face notable challenges. First, these models often struggle to effectively capture the complex spatio-temporal dependencies among variables (e.g., TimeGAN (Yoon, Jarrett, and der Schaar 2019) and TimeVAE (Desai et al. 2022)). Second,

their reliance on Euclidean space assumptions limits their ability to represent the topological and structural characteristics inherent in time series data.

Effectively modeling temporal dependency structures is a key challenge in time series analysis. To address this, various techniques have been developed, including random historical skip connections (Kieu et al. 2019), multi-scale temporal links (Shen et al. 2021), and autoregressive temporal structures (Yu, Rao, and Dhillon 2016). These approaches are designed to capture intricate and long-range temporal relationships within time series data.

In addition to capturing temporal dependencies, understanding and implementing generative mechanisms for time series data has garnered significant attention. Generative models such as TimeGAN (Yoon, Jarrett, and der Schaar 2019), TimeVAE (Desai et al. 2022), and diffusion-based approaches like CSDI (Tashiro et al. 2021) and TimeDiff (Shen and Kwok 2023) have shown promise in synthesizing time series data by learning the underlying data distributions. For instance, diffusion-based models like Diffusion-TS (Yuan and Qiao 2024) integrate seasonal-trend decomposition with denoising diffusion probabilistic models to learn temporal characteristics using Fourier-based loss terms. However, these methods often focus solely on raw data domains and fail to fully capture dynamic multivariate interactions. Additionally, architectures such as the encoder-decoder transformer in Diffusion-TS process spatial and temporal information separately, which limits their ability to model complex interdependencies between variables.

More recently, graph-based methods (Han et al. 2020; Yi et al. 2024; Cheng et al. 2020) have emerged as a powerful alternative in spatiotemporal prediction tasks. For example, STGCN (Han et al. 2020) combines graph and temporal convolutions to model spatio-temporal patterns, while FourierGNN (Yi et al. 2024) uses spectral graph theory and Fourier transforms to analyze time series in frequency domains. By representing temporal dependencies as graph structures, these methods excel at capturing complex dependencies of temporal data. Despite their success in prediction tasks, their potential for generative tasks in synthetic time series generation has been largely overlooked.

To address these limitations, we propose *TSGDiff*, a novel framework for synthetic time series generation that rethinks the problem from a pure graph perspective. The core in-

\*Corresponding Author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

novation of *TSGDiff* lies in its unified approach, which combines dynamic graph construction and a diffusion-based generative process within the latent graph space. Nodes in the graph represent variables, and edges are constructed based on Fourier spectrum characteristics, enabling the model to capture intricate temporal dependencies and encode structural relationships in a flexible and adaptive manner. The diffusion-based process operates in the latent graph space, effectively modeling temporal semantic distributions and enabling the generation of time series with realistic dynamics and structural coherence. By unifying Fourier-based graph construction, latent representation learning, and diffusion-based modeling, *TSGDiff* offers a robust and flexible solution for structured time series generation, ensuring high fidelity, temporal accuracy, and semantic coherence. Furthermore, we propose the Topological Structure Fidelity (Topo-FID) score, a novel graph-aware metric for assessing the similarity of time series graph structures. Topo-FID integrates two complementary sub-metrics: Graph Edit Similarity, which quantifies differences in adjacency matrices, and Structural Entropy Similarity, which evaluates the entropy of node degree distributions. This comprehensive metric bridges gaps in traditional evaluation methods, providing a more accurate and meaningful assessment of time series generation quality.

To summarize, our main contributions include:

- We propose *TSGDiff*, the first framework to rethink synthetic time series generation from a graph-based perspective. By integrating a diffusion model within the latent graph space, it effectively captures and models the structural representation distribution of time series data.
- We introduce the Topological Structure Fidelity (Topo-FID), a graph-aware metric that quantifies the structural fidelity of the generated timeseries distribution.
- We validate the effectiveness of *TSGDiff* on commonly-used real-world datasets, demonstrating its ability to generate high-quality time series data with realistic temporal and structural characteristics.

## 2 Background

### 2.1 Synthetic Time Series Generation

In recent years, a variety of methods based on different generative paradigms have emerged for time series generation (Naiman et al. 2024; Crabbé et al. 2024; Zhou et al. 2023; Park et al. 2024; Alaa, Chan, and van der Schaar 2021). Generative Adversarial Network (GAN)-based approaches (Esteban, Hyland, and Ratsch 2017), such as TimeGAN (Yoon, Jarrett, and der Schaar 2019), capture temporal dynamics by jointly optimizing supervised and adversarial objectives. Variational Autoencoder (VAE)-based methods (Fortuin et al. 2020) have been developed with specialized decoder structures tailored for time series data, incorporating trend and seasonal decomposition.

With the rise of diffusion models in generative tasks, approaches like CSDI (Tashiro et al. 2021) and TimeDiff (Shen and Kwok 2023) have demonstrated remarkable

performance in generating time series through iterative denoising processes. Specifically, CSDI introduces a self-supervised masking condition designed for imputation tasks, while TimeDiff develops an efficient future-mixup mechanism tailored for forecasting. For unconditional generation, Diffusion-TS mentioned in Section 1 is one of the recent popular works. It models time series data distribution by a diffusion model enhanced by seasonal-trend decomposition and a frequency-enhanced objective. Despite their strengths, the above methods directly process raw time series data and fail to account for the temporal structured relationships from a pure graph perspective.

### 2.2 Denoising Diffusion Probabilistic Models

A diffusion probabilistic model (Ho, Jain, and Abbeel 2020) is designed to learn the reversal of a Markov chain process, known as the diffusion process, which incrementally adds noise to data, ultimately obliterating the original signal.

**Forward Diffusion.** Consider a data sample  $\mathbf{x} \in \mathbb{R}^d \sim p(\mathbf{x})$  and a set of latent variables  $\{\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_K\}$  ( $\mathbf{x}_0 = \mathbf{x}$ ) that interpolate between the data distribution and a Gaussian distribution as the diffusion steps progress. The forward process is formally defined as a Markov chain, parameterized by a sequence of variances  $\beta_k$  and  $\alpha_k := 1 - \beta_k$ , and can be expressed as:  $q(\mathbf{x}_{1:K} | \mathbf{x}_0) = \prod_{k=1}^K q(\mathbf{x}_k | \mathbf{x}_{k-1})$ , where  $q(\mathbf{x}_k | \mathbf{x}_{k-1}) = \mathcal{N}(\mathbf{x}_k; \sqrt{1 - \beta_k}\mathbf{x}_{k-1}, \beta_k\mathbf{I})$ .

As the number of diffusion steps increases, more noise is added to the data. Consequently,  $q(\mathbf{x}_k | \mathbf{x})$  has a closed-form solution, which can be described in a general form:

$$\mathbf{x}_k = \sqrt{\bar{\alpha}_k}\mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_k}\epsilon, \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}), \quad (1)$$

where  $\bar{\alpha}_k = \prod_{i=1}^k (1 - \beta_i) \in (0, 1)$ ,  $\mathbf{x}_0 = \mathbf{x}$ , and  $\mathbf{x}_K \sim \mathcal{N}(0, \mathbf{I})$ . As the diffusion process progresses, the latent variable  $\mathbf{x}_k$  becomes increasingly noisy, eventually converging to  $\mathbf{x}_K$ , which approximates a Gaussian distribution and becomes independent of the initial data sample  $\mathbf{x}$ .

**Backward Denoising.** The denoising process is learnable and defined by the inverted Markov chain:

$$p_\theta(\mathbf{x}_{0:K}) = p(\mathbf{x}_K) \prod_{k=1}^K p_\theta(\mathbf{x}_{k-1} | \mathbf{x}_k), \quad (2)$$

where  $p(\mathbf{x}_K) = \mathcal{N}(0, \mathbf{I})$  is a known prior. The conditional distribution  $p_\theta(\mathbf{x}_{k-1} | \mathbf{x}_k)$  is approximated by a Gaussian distribution:

$$q(\mathbf{x}_{k-1} | \mathbf{x}_k, \mathbf{x}) = \mathcal{N}(\mu_k(\mathbf{x}_k, \mathbf{x}), \sigma_k^2\mathbf{I}), \quad (3)$$

where  $\mu_k(\mathbf{x}_k, \mathbf{x})$  has a closed-form solution and  $\sigma_k$  is a hyperparameter. Instead of predicting the clean data sample  $\mathbf{x}$  directly, the denoising network  $\epsilon_\theta$  is trained to estimate the noise  $\epsilon$  added during the forward process to the latent variable  $\mathbf{x}_k$ . The noise-prediction loss function is:

$$L = \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x}), k \sim \mathcal{U}\{1, \dots, K\}, \epsilon \sim \mathcal{N}(0, \mathbf{I})} [\|\epsilon - \epsilon_\theta(\mathbf{x}_k, k)\|_2^2], \quad (4)$$

where  $\mathbf{x}_k = \sqrt{\bar{\alpha}_k}\mathbf{x} + \sqrt{1 - \bar{\alpha}_k}\epsilon$ , and  $\epsilon_\theta(\mathbf{x}_k, k)$  predicts the noise component  $\epsilon$ .

The conditional distribution  $p_\theta(\mathbf{x}_{k-1} | \mathbf{x}_k)$  is approximated as:  $p_\theta(\mathbf{x}_{k-1} | \mathbf{x}_k) \approx q(\mathbf{x}_{k-1} | \mathbf{x}_k, \hat{\mathbf{x}}_\theta(\mathbf{x}_k, k))$ , where

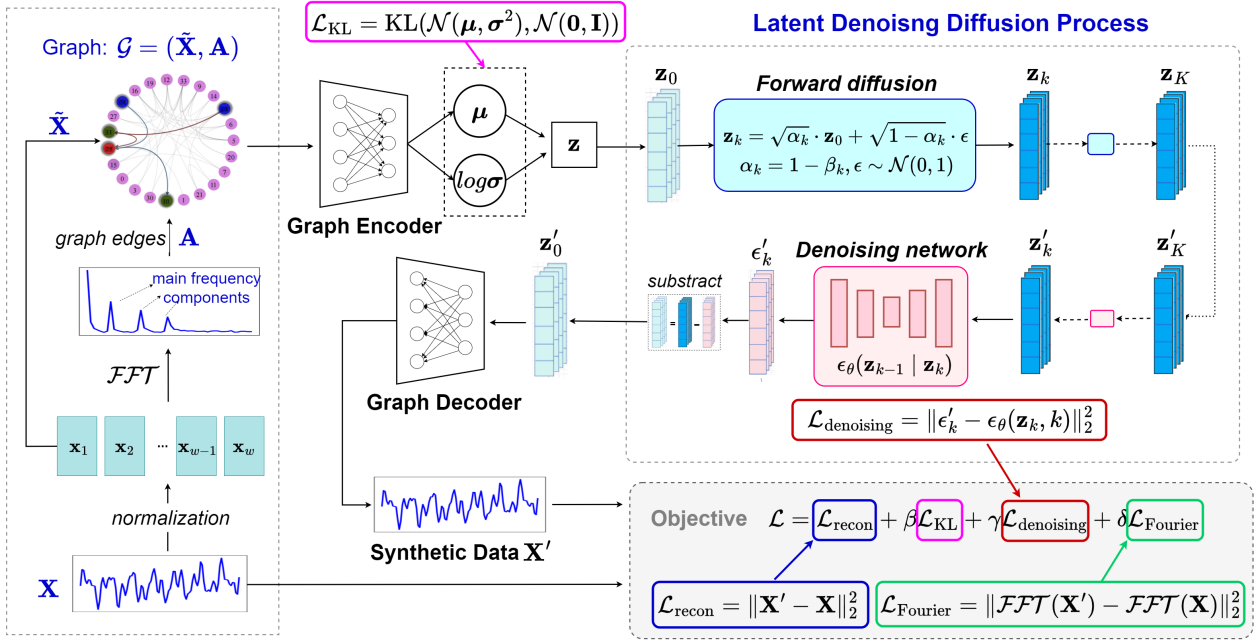


Figure 1: Overview of the proposed *TSGDiff* model.

$\hat{\mathbf{x}}_\theta(\mathbf{x}_k, k)$  is the implied clean sample estimate based on the predicted noise  $\epsilon_\theta(\mathbf{x}_k, k)$ :

$$\hat{\mathbf{x}}_\theta(\mathbf{x}_k, k) = \frac{\mathbf{x}_k - \sqrt{1 - \alpha_k} \epsilon_\theta(\mathbf{x}_k, k)}{\sqrt{\alpha_k}}. \quad (5)$$

**Sampling.** For sampling from the trained diffusion model, we follow the inference distribution proposed in (Song, Meng, and Ermon 2021). The reverse process iteratively generates  $\mathbf{x}_{k-1}$  from  $\mathbf{x}_k$  using a Gaussian distribution parameterized by the mean and variance. The mean depends on both the current noisy variable  $\mathbf{x}_k$  and the predicted clean data  $\mathbf{x}$ , while the variance controls stochasticity during sampling. When the variance is set to zero, the deterministic DDIM sampler is used, enabling efficient and stable sampling by directly denoising  $\mathbf{x}_k$  into a clean sample through a simplified closed-form update. This iterative process continues until the final reconstructed sample  $\mathbf{x}'_0$  is obtained.

### 3 Methodology

This section formally elaborates on the proposed *TSGDiff*, a framework for synthetic time series generation that approaches the problem from a graph-based perspective. As shown in Figure 1, the framework consists of three main stages: graph construction and encoding, latent diffusion-based modeling, and graph decoding. Initially, the time series data  $\mathbf{X}$  is standardized and segmented using a sliding window approach. Each window  $\tilde{\mathbf{X}}$  undergoes a Fourier transform to extract periodicity information, which is then used to construct a dynamic graph  $\mathcal{G}(\tilde{\mathbf{X}}, \mathbf{A})$ . In this graph, nodes represent variables, while edges in the adjacency matrix  $\mathbf{A}$  are established based on characteristics of the Fourier spectrum, enabling the model to capture temporal dependencies and structural relationships effectively. The constructed

graph is encoded into a latent representation using a graph neural network. Within this latent graph space, a diffusion-based generative process progressively refines noisy representations to recover meaningful temporal and structural patterns. This ensures the generated data maintains realistic dynamics and structural coherence. Finally, the denoised latent representation is decoded back into synthetic time series data, completing the generation process.

#### 3.1 Instance Normalization

First, the sliding window technique divides long time series into fixed-length segments to extract local temporal patterns. Given a time series  $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T]$ , the segments are expressed as:  $\mathbf{X}_{i:i+w} = [\mathbf{x}_i, \mathbf{x}_{i+1}, \dots, \mathbf{x}_{i+w-1}]$ , where  $w$  denotes the window size and each segment captures localized trends while maintaining temporal dependencies. To standardize features and enhance model training stability, Min-Max normalization is applied. This technique scales each feature of the time series to a target range, typically  $[0, 1]$ , by linearly transforming the data based on its minimum and maximum values, ensuring consistent feature scales across all samples. For simplicity, we use  $\tilde{\mathbf{X}}_i$  to represent the window  $\mathbf{X}_{i:i+w}$ .

#### 3.2 Graph Construction

To capture temporal dependencies (Yi et al. 2023), each normalized segment  $\tilde{\mathbf{X}}_i$  is transformed into a graph  $\mathcal{G}$ . In this graph, each time step in the slice corresponds to a node, and the node feature is the time series value at that time step.

Fourier transform (Bracewell and Bracewell 1986) is utilized as the foundation for analyzing the frequency-domain characteristics of the time series. By decomposing the time-

domain signal into its frequency components for each variable dimension independently, it enables the identification of periodicities that guide the graph construction process. For a time series slice  $\tilde{\mathbf{X}}_i$ , the discrete Fourier transform (DFT) for a specific variable is computed as:

$$\tilde{X}_i^{(p)} = \sum_{n=0}^{w-1} \tilde{x}_{i+n} \cdot e^{-j\frac{2\pi}{w}pn}, \quad p = 0, 1, \dots, w-1, \quad (6)$$

where  $\tilde{X}_i^{(p)}$  denotes the  $p$ -th frequency component, and  $w$  represents the segment length (window size). The Fourier transform produces a frequency spectrum for each variable in the slice, highlighting the amplitude of its frequency components. The top three frequencies with the highest amplitudes are identified using a peak detection method and converted into their corresponding periods. These periods are then used to define edges in the graph, enabling the model to capture short-term, medium-term, and long-term variation patterns (Yoon, Jarrett, and der Schaar 2019; Segarra et al. 2017; Isufi, Gama, and Ribeiro 2021).

The graph  $\mathcal{G}$  is constructed by defining edges based on both temporal and periodic dependencies. To ensure the capture of local temporal relationships, an edge is always created between consecutive time steps. Additionally, edges are added based on the detected periods from the frequency spectrum, linking nodes that are periodic neighbors. This approach allows the graph to encode both short-term and long-term dependencies. The adjacency matrix  $\mathbf{A}$  is used to represent the graph structure, where  $\mathbf{A}_{ij} = 1$  indicates the presence of an edge between nodes  $i$  and  $j$ , and  $\mathbf{A}_{ij} = 0$  otherwise. The node feature matrix  $\tilde{\mathbf{X}}$  stores the normalized time series values associated with each node. As shown in Figure 1 (left top), this graph construction process integrates temporal and periodic relationships into a cohesive structure, providing a rich representation of the time series to support downstream modeling tasks.

### 3.3 Graph Encoding

The graph convolutional layer (Kipf and Welling 2017; Katznelson 1970) aggregates neighbor information by matrix multiplication between the adjacency matrix  $\mathbf{A}$  and the node feature matrix  $\tilde{\mathbf{X}}$ . A single graph convolution operation is defined as:

$$\mathbf{y}' = \text{Mish}(\text{BN}((W\mathbf{A}\tilde{\mathbf{X}} + b)^T)^T), \quad (7)$$

where  $W$  and  $b$  are learnable parameters,  $\text{BN}(\cdot)$  represents batch normalization, and  $\text{Mish}(\cdot)$  is the non-linear activation function. Residual connections are used to combine the input  $\tilde{\mathbf{X}}$  with the output  $\mathbf{y}'$ , ensuring better gradient flow and preserving original features.

The graph encoder stacks multiple graph convolutional layers to process the input graph (node feature matrix  $\tilde{\mathbf{X}}$  and adjacency matrix  $\mathbf{A}$ ). After multiple layers of feature aggregation, mean pooling is applied to obtain the global graph representation  $\mathbf{X}_{\text{pool}}$ . The encoder then uses two linear layers to output the mean  $\mu$  and logarithmic standard deviation  $\log \sigma$  of the latent distribution:

$$\mu = \text{Linear}_{\mu}(\mathbf{X}_{\text{pool}}), \quad \log \sigma = \text{Linear}_{\log \sigma}(\mathbf{X}_{\text{pool}}). \quad (8)$$

The latent vector  $\mathbf{z}$  is sampled from the latent distribution using the reparameterization trick, ensuring differentiability:

$$\mathbf{z} = \mu + \epsilon \odot \exp\left(\frac{\log \sigma}{2}\right), \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}), \quad (9)$$

where  $\epsilon$  is sampled from a standard normal distribution, and  $\odot$  denotes element-wise multiplication.

The graph decoder reconstructs the input graph from the latent vector  $\mathbf{z}$ . Starting with  $\mathbf{z}$ , fully connected layers with batch normalization and Mish activation progressively increase the feature dimension. The final output  $\hat{\mathbf{X}}$  is computed using a linear transformation followed by the Tanh function to ensure the reconstructed graph matches the normalized range of the input:

$$\hat{\mathbf{X}} = \text{Tanh}(\text{Linear}(\mathbf{z})). \quad (10)$$

This ensures the restoration of node features with dimensions consistent with the input graph data.

The graph encoder-decoder framework follows the principles of a Variational Autoencoder (VAE). A key component of this structure is the KL divergence loss  $\mathcal{L}_{\text{KL}}$ , which ensures that the latent variable distribution approximates a standard normal distribution  $\mathcal{N}(0, \mathbf{I})$ . The KL divergence loss is defined as:

$$\mathcal{L}_{\text{KL}} = -\frac{1}{2} \sum_{j=1}^d (1 + \log \sigma_j^2 - \mu_j^2 - \sigma_j^2), \quad (11)$$

where  $\mu_j$  and  $\sigma_j^2$  are the mean and variance of the latent distribution for dimension  $j$ .

This loss regularizes the latent space by encouraging the learned distribution  $q_{\phi}(\mathbf{z}|\mathbf{X})$  to be close to the prior distribution  $p(\mathbf{z}) = \mathcal{N}(0, \mathbf{I})$ , ensuring smooth and meaningful representations in the latent space.

### 3.4 Latent Diffusion-based Modeling

The KL divergence loss  $\mathcal{L}_{\text{KL}}$  regularizes the latent variable  $\mathbf{z}$  by constraining it to follow a standard normal distribution  $\mathcal{N}(0, \mathbf{I})$ . However, this constraint alone is insufficient for capturing the complex structure of data or enabling high-quality generative modeling. To address this, we introduce a latent diffusion-based approach, which refines and generates  $\mathbf{z}$  through a denoising diffusion process.

**Latent Diffusion Process.** The forward process gradually corrupts  $\mathbf{z}$  by adding Gaussian noise over time:

$$\mathbf{z}_k = \sqrt{\alpha_k} \cdot \mathbf{z} + \sqrt{1 - \alpha_k} \cdot \epsilon, \quad \epsilon \sim \mathcal{N}(0, 1), \quad (12)$$

where  $\alpha_k = \prod_{s=1}^k \alpha_s$ ,  $\alpha_s = 1 - \beta_s$ , and  $\beta_s$  is a noise coefficient. As  $k$  increases,  $\mathbf{z}_k$  approaches pure noise.

The reverse process recovers  $\mathbf{z}$  by learning the posterior distribution  $p_{\theta}(\mathbf{z}_{k-1}|\mathbf{z}_k, k)$ :

$$p_{\theta}(\mathbf{z}_{k-1}|\mathbf{z}_k, k) := \mathcal{N}(\mathbf{z}_{k-1}; \mu_{\theta}(\mathbf{z}_k, k), \Sigma_{\theta}(k)), \quad (13)$$

where  $\mu_{\theta}$  is the predicted mean and  $\Sigma_{\theta}(k)$  is the noise variance. The reverse update is:

$$\mathbf{z}_{k-1} = \sqrt{\frac{1}{\alpha_k}} \cdot \left( \mathbf{z}_k - \frac{1 - \alpha_k}{\sqrt{1 - \alpha_k}} \cdot (\mathbf{z}_k - \hat{\mathbf{z}}) \right) + \sqrt{\beta_k} \cdot \epsilon', \quad (14)$$

where  $\hat{\mathbf{z}} = \mathbf{z}_{\theta}(\mathbf{z}_k, k)$  is the predicted latent variable, and  $\epsilon' \sim \mathcal{N}(0, 1)$  is reverse process noise.

**Latent Diffusion Network.** The reverse process is implemented using a latent diffusion network with an input layer, intermediate blocks, and an output layer:

$$\begin{aligned} \mathbf{h}_0 &= W_{in}(\text{concat}(\mathbf{z}_k, \mathbf{t}_{emb})) + b_{in}, \\ \mathbf{h}_l &= \text{ReLU}(W_l \mathbf{h}_{l-1} + b_l), \quad l = 1, \dots, L, \\ \hat{\mathbf{z}}_{k-1} &= W_{out} \mathbf{h}_L + b_{out}. \end{aligned} \quad (15)$$

Here,  $\mathbf{z}_k$  is the noisy latent variable,  $\mathbf{t}_{emb}$  is the time step embedding, and  $\hat{\mathbf{z}}_{k-1}$  is the denoised prediction. Intuitively, Latent diffusion extends the KL-regularized latent space into a fully generative framework by iteratively denoising  $\mathbf{z}$ . This approach enables expressive modeling of complex data structures while supporting efficient and diverse sampling in the latent space.

**Training Objective** The denoising loss measures the difference between the noisy latent variable and the output of the denoising network  $\epsilon_\theta$ , which predicts clean latent variables:

$$L_{\text{denoising}} = \mathbb{E}_{k, \mathbf{z}_k, \epsilon} [\|\epsilon - \epsilon_\theta(\mathbf{z}_k, k)\|^2], \quad (16)$$

where  $\epsilon_\theta$  is the noise estimated by the denoising network,  $\mathbf{z}_k$  is the noisy latent variable at time step  $k$ , and  $\epsilon$  is the actual noise added during the forward process.

To further enhance consistency, we incorporate a Fourier loss inspired by Diffusion-TS, which enforces alignment in the frequency domain:

$$L_{\text{Fourier}} = \mathbb{E}_{\mathbf{x}_0} \|\mathcal{FFT}(\mathbf{x}_0) - \mathcal{FFT}(\hat{\mathbf{x}}_0)\|^2, \quad (17)$$

where  $\mathcal{FFT}$  represents the Fourier transform,  $\mathbf{x}_0$  is the original data, and  $\hat{\mathbf{x}}_0$  is the reconstruction.

The total loss integrates all components:

$$\mathcal{L} = \mathcal{L}_{\text{recon}} + \beta \mathcal{L}_{\text{KL}} + \gamma \mathcal{L}_{\text{denoising}} + \delta \mathcal{L}_{\text{Fourier}}, \quad (18)$$

where  $\mathcal{L}_{\text{recon}}$  is the reconstruction loss capturing the MSE between the input data and decoder output, and  $\beta, \gamma, \delta$  are weights balancing each term. In practice, we set  $\beta = 0.2$ , while  $\gamma = 1$  and  $\delta = 1$ . By minimizing  $\mathcal{L}$ , the model learns to denoise latent variables, enforce frequency consistency, and reconstruct time series data accurately, ensuring the generated data aligns with the original distribution.

## 4 Experiments

### 4.1 Setup

**Datasets.** We use six public real-world time series datasets from diverse domains: i) `ETTh`: Hourly energy and environmental data (2016–2018) with over 15,000 time steps. ii) `Stocks`<sup>1</sup>: Daily multi-stock data with prices and trading volumes spanning multiple years. iii) `Exchange`<sup>2</sup>: Daily USD exchange rates (1990–2010) against major currencies. iv) `Weather`<sup>3</sup>: Meteorological data recorded every 10 minutes, with over 50,000 time steps and 20+ features like temperature and humidity. v) `Wind`<sup>4</sup>: Hourly wind farm generation data across seasons, spanning over a year. vi) `EEG`<sup>5</sup>:

<sup>1</sup><https://finance.yahoo.com/quote/GOOG>

<sup>2</sup><https://github.com/laiguokun/multivariate-time-series-data>

<sup>3</sup><https://www.bgc-jena.mpg.de/wetter/>

<sup>4</sup><https://github.com/PaddlePaddle/PaddleSpatial/tree/main/paddlespatial/datasets/WindPower>

<sup>5</sup><https://archive.ics.uci.edu/dataset/264/eeg+eye+state>

High-frequency (128Hz) EEG signals with over 15,000 entries.

**Evaluation Metrics.** We evaluate synthetic data using four traditional metrics: Context-FID (Paul et al. 2022) (local context quality), Correlational score (Ni et al. 2020) (temporal dependency via cross-correlation), Discriminative score (Yoon, Jarrett, and der Schaar 2019) (classifier-based similarity), and Predictive score (Yoon, Jarrett, and der Schaar 2019) (TSTR-based utility).

These metrics, however, fail to capture the intrinsic structural and frequency characteristics of time series data. To address this, we propose Topological Structure Fidelity (Topo-FID), a new metric that assesses the similarity of distributions between time series samples by evaluating their graph representations  $\mathcal{G}$ . Formally, Topo-FID is defined as the expected similarity of graph representations derived from pairs of sampled time series:

$$\text{Topo-FID} = \mathbb{E}_{(\mathcal{G}, \hat{\mathcal{G}}) \sim P} [\alpha \cdot S_{\text{edit}} + (1 - \alpha) \cdot S_{\text{entropy}}], \quad (19)$$

where  $\mathcal{G}$  and  $\hat{\mathcal{G}}$  are the graph representations of real and synthetic time series samples,  $P$  is the sampling distribution.  $\alpha$  is fixed to be 0.5 by default. Here,  $S_{\text{edit}}$  is the Graph Edit Similarity, which measures edge structure similarity:

$$S_{\text{edit}} = 1 - \frac{\sum_{i,j} |\mathbf{A}_{i,j} - \hat{\mathbf{A}}_{i,j}|}{N^2}, \quad (20)$$

where  $\mathbf{A}$  and  $\hat{\mathbf{A}}$  are adjacency matrices, and  $N$  is the number of nodes.  $S_{\text{entropy}}$  is the Structural Entropy Similarity, which evaluates node degree distribution similarity:

$$S_{\text{entropy}} = \frac{1}{1 + |H(\mathbf{A}) - H(\hat{\mathbf{A}})|}, \quad (21)$$

with  $H(\mathbf{A}) = -\sum_{d \in D} p(d) \log_2(p(d) + \epsilon)$ , where  $p(d)$  is the degree distribution of  $\mathbf{A}$ , and  $\epsilon = 10^{-10}$  avoids logarithmic errors.

**Baselines.** In experiments, we have selected four popular and strong baselines for comparison, including: i) a diffusion model, Diffusion-TS (Yuan and Qiao 2024); ii) two GAN-based models, TimeGAN (Yoon, Jarrett, and der Schaar 2019) and Cot-GAN (Ni et al. 2020); iii) a variational autoencoder, TimeVAE (NASA 2015).

**Implementation Details.** In the Graph Encoder, we use three stacked GraphConvBlocks, with both the hidden layer and latent space dimensions set to 1600. The Diffusion model employs three DiffusionBlocks with 64 units each (where each DiffusionBlock is a three-layer MLP network), and the timesteps are set to 1000. The Graph Decoder uses four fully connected layers to restore the 1600-dimensional features to their original dimension. The sliding window size is 48 (say, 2 days for `ETTh`, 48 days for `Stocks`, 48 days for `Exchange`, 8h for `Weather`, 6h for `Wind`, 375ms for `EEG`), with a stride of 1. We use a batch size of 128, a learning rate of 0.01, and train for 500 epochs. All experiments are repeated five times, implemented in PyTorch (Paszke et al. 2019), and conducted on an NVIDIA RTX 4090 GPU with 24GB memory.

Metric	Methods	ETTh	Stocks	Exchange	Weather	Wind	EEG
Topo-FID ↗ (higher the better)	TSGDiff	<b>0.986±0.0%</b>	<b>0.826±0.0%</b>	<b>0.869±0.0%</b>	<b>0.867±0.0%</b>	<b>0.869±0.0%</b>	<b>0.852±0.0%</b>
	Diffusion-TS	0.864±0.0%	0.785±0.0%	0.806±0.0%	0.828±0.0%	0.819±0.0%	0.788±0.0%
	TimeGAN	0.798±0.0%	0.795±0.0%	0.801±0.0%	0.816±0.0%	0.824±0.0%	0.828±0.0%
	Cot-GAN	0.894±0.0%	<u>0.802±0.0%</u>	<u>0.809±0.0%</u>	<u>0.859±0.0%</u>	<u>0.825±0.0%</u>	<u>0.841±0.0%</u>
	TimeVAE	<u>0.904±0.0%</u>	0.798±0.0%	0.804±0.0%	0.816±0.0%	<u>0.829±0.0%</u>	0.769±0.0%
Context-FID ↘ (lower the better)	TSGDiff	<b>0.224±0.1%</b>	<b>0.357±1.5%</b>	<b>0.061±0.9%</b>	<b>0.353±0.1%</b>	<b>0.256±1.5%</b>	<b>0.020±0.3%</b>
	Diffusion-TS	<u>0.225±0.4%</u>	0.530±2.5%	0.067±1.0%	<u>1.161±0.1%</u>	<u>0.491±2.4%</u>	0.030±0.6%
	TimeGAN	1.392±0.5%	<u>0.377±1.7%</u>	1.103±0.9%	<u>2.420±2.3%</u>	<u>5.087±2.7%</u>	0.023±0.7%
	Cot-GAN	3.486±0.7%	<u>0.596±0.8%</u>	1.523±2.1%	5.892±1.2%	5.409±0.8%	3.267±0.5%
	TimeVAE	3.452±0.6%	0.830±0.9%	<u>0.065±0.7%</u>	13.952±1.3%	5.174±0.0%	8.874±0.0%
Correlational ↘ (lower the better)	TSGDiff	<b>0.024±0.1%</b>	<b>0.026±0.1%</b>	<b>0.019±0.3%</b>	<b>0.035±0.1%</b>	<b>0.022±0.2%</b>	<b>0.201±0.1%</b>
	Diffusion-TS	<u>0.024±0.4%</u>	<u>0.027±0.1%</u>	0.028±0.5%	0.054±0.3%	<u>0.033±1.7%</u>	0.466±0.2%
	TimeGAN	<u>0.126±0.5%</u>	<u>0.061±0.7%</u>	0.131±0.6%	0.155±0.1%	<u>0.454±0.3%</u>	0.697±0.2%
	Cot-GAN	0.032±0.3%	0.052±1.3%	0.031±0.5%	0.087±0.9%	0.043±0.2%	0.214±0.1%
	TimeVAE	0.025±0.2%	0.053±0.1%	<u>0.022±1.2%</u>	<u>0.043±0.7%</u>	0.064±0.5%	<u>0.209±1.2%</u>
Discriminative ↘ (lower the better)	TSGDiff	<b>0.056±0.6%</b>	<b>0.019±1.3%</b>	<b>0.052±0.9%</b>	<b>0.283±0.1%</b>	<b>0.073±0.3%</b>	<b>0.301±1.5%</b>
	Diffusion-TS	<u>0.111±0.1%</u>	<u>0.075±0.5%</u>	0.172±0.9%	0.398±1.8%	<u>0.161±1.1%</u>	0.492±1.3%
	TimeGAN	0.353±0.7%	<u>0.193±2.5%</u>	0.465±1.5%	0.494±1.8%	<u>0.495±0.5%</u>	0.391±1.7%
	Cot-GAN	0.114±1.5%	0.102±1.6%	0.154±2.5%	<u>0.297±0.8%</u>	0.168±0.7%	0.335±1.1%
	TimeVAE	0.346±1.6%	0.260±3.3%	<u>0.125±2.1%</u>	<u>0.496±1.9%</u>	0.338±0.9%	<u>0.496±2.7%</u>
Predictive ↘ (lower the better)	TSGDiff	<b>0.020±0.1%</b>	<b>0.005±0.1%</b>	<b>0.004±0.0%</b>	<b>0.009±0.0%</b>	<b>0.007±0.0%</b>	<b>0.001±0.1%</b>
	Diffusion-TS	<u>0.026±0.0%</u>	<u>0.005±0.2%</u>	0.005±0.2%	0.011±0.0%	<u>0.009±0.0%</u>	<u>0.002±0.2%</u>
	TimeGAN	<u>0.030±0.2%</u>	<u>0.006±0.1%</u>	<u>0.004±0.1%</u>	<u>0.009±0.1%</u>	<u>0.009±0.1%</u>	<u>0.002±0.3%</u>
	Cot-GAN	0.028±0.2%	0.013±0.1%	<u>0.006±0.3%</u>	<u>0.021±0.2%</u>	0.026±0.2%	0.035±0.3%
	TimeVAE	0.038±0.2%	0.016±0.2%	0.006±0.1%	0.031±0.3%	0.022±0.0%	0.039±0.2%

Table 1: Results of multivariate time-series generation (bold indicates best performance, while underline is the second best).

Method	Topo-FID↑	Context-FID↓	Correlational↓	Discriminative↓	Predictive↓
TSGDiff	<b>0.986±0.0%</b>	<b>0.224±0.1%</b>	<b>0.024±0.1%</b>	<b>0.056±0.6%</b>	<b>0.020±0.1%</b>
w/o KL	0.787±0.0%	44.467±1.5%	0.159±0.1%	0.500±1.3%	0.064±0.1%
w/o Denoising	0.908±0.0%	3.922±0.9%	0.146±0.3%	0.432±0.9%	0.019±0.2%
w/o Fourier	0.883±0.0%	0.407±0.1%	0.024±0.1%	0.061±0.7%	0.021±0.3%

Table 2: Ablation study results on ETTh dataset (Bold indicates the best).

## 4.2 Results for Synthetic Time Series Generation

Table 1 highlights the performance of *TSGDiff* compared to baselines, including Diffusion-TS (Yuan and Qiao 2024), TimeGAN (Yoon, Jarrett, and der Schaar 2019), Cot-GAN (Ni et al. 2020), and TimeVAE (NASA 2015). The results show that *TSGDiff* consistently produces higher-quality synthetic data across nearly all metrics, improving the average Discriminative score by 50% across six datasets.

Compared to Diffusion-TS, *TSGDiff* leverages graph-based latent space modeling to capture structural temporal dependencies better, excelling on complex datasets like *Weather* and *EEG*. Against GAN-based models such as TimeGAN and Cot-GAN, it avoids mode collapse, producing more robust and diverse samples. Relative to TimeVAE, *TSGDiff* preserves structural coherence more effectively

through its unified graph-based framework.

These advantages are especially evident in the Topo-FID metric, where *TSGDiff* significantly outperforms all baselines, demonstrating its ability to generate data with superior structural fidelity and temporal accuracy.

## 4.3 Visualization Results

We used two visualization methods to evaluate time series synthesis performance as suggested by (Yuan and Qiao 2024): i) projecting original and synthetic data into a two-dimensional space using t-SNE (der Maaten and Hinton 2008), and ii) visualizing data distributions via kernel density estimation. As shown in Figure 2, while the first row (t-SNE) does not clearly indicate which method performs better, the second row (kernel density estimation) demonstrates

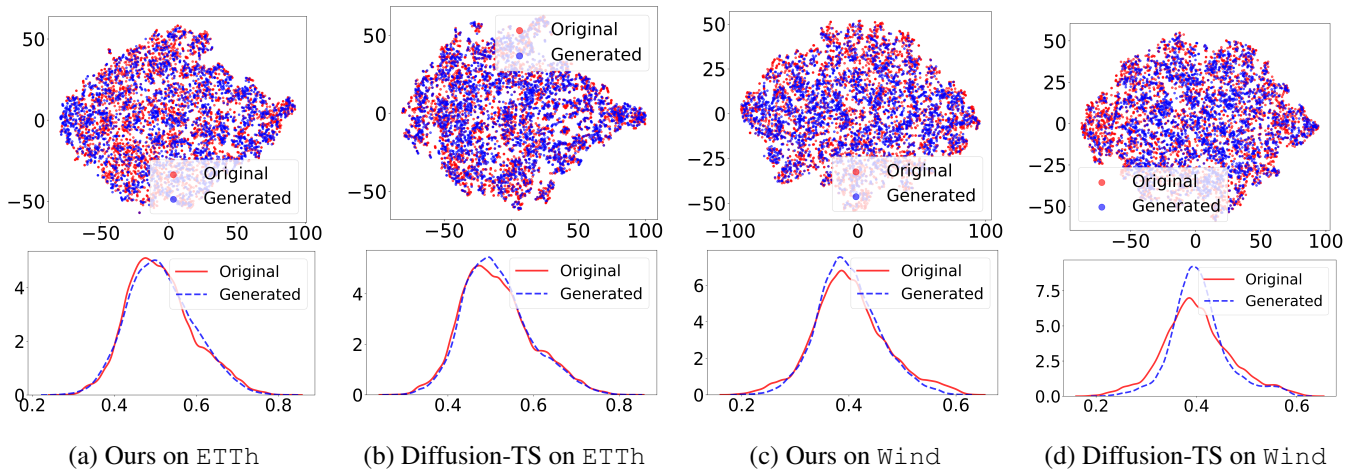


Figure 2: Visualizations of the time series synthesized by *TSGDiff* and Diffusion-TS.

that the synthetic distribution generated by *TSGDiff* closely matches the original one, outperforming Diffusion-TS.

#### 4.4 Ablation Study

To assess the contributions of different components in *TSGDiff*, we compared the full model with three variants: i) *w/o KL divergence loss*: Removes the KL divergence loss to analyze its role in constraining latent variable distributions. ii) *w/o Denoising loss*: Excludes the diffusion model loss, evaluating its contribution to generative capability. iii) *w/o Fourier loss*: Omits the Fourier loss term to examine its impact on capturing frequency-domain features.

As shown in Table 2, the KL divergence loss, denoising loss, and Fourier loss each play a critical role in *TSGDiff*'s performance. The KL divergence loss provides a foundational prior by aligning latent variables with a standard normal distribution, ensuring diversity and stability while preventing mode collapse and unstable training. Building on this foundation, the denoising loss further refines latent variables during the reverse diffusion process, preserving the fidelity of generated samples. Meanwhile, the Fourier loss enhances reconstruction precision by enforcing frequency-domain consistency, preserving key spectral features like periodicity. Together, these components form a cohesive framework that ensures robust performance, high-quality data reconstruction, and effective handling of complex time series synthesis tasks.

#### 4.5 Efficiency Analysis

We analyzed *TSGDiff*'s graph construction time cost to evaluate the efficiency of transforming time series into graphs.

As shown in Table 3, the graph construction process is highly efficient, with time costs consistently under 5 seconds and accounting for less than 0.1% of the total training time. This demonstrates that the process introduces negligible overhead, ensuring smooth integration into the overall pipeline. This efficiency aligns seamlessly with the latent diffusion framework in *TSGDiff*, which operates in a compressed latent space to balance computational cost and

	ETTh	Stocks	Weather
# iterations	68,000	14,500	61,000
graph construction (s)	3.12	0.82	3.79
total training (s)	4,506.15	970.23	4025.31
ratio (%)	0.07	0.08	0.09

Table 3: Efficiency analysis. Here,  $\text{ratio} = \frac{\text{graph construction}}{\text{total training}}$ .

model performance. The lightweight yet expressive graph representations enable the latent diffusion process to focus on capturing complex temporal dependencies, while maintaining high-quality reconstructions. Together, these components highlight the practicality and scalability of *TSGDiff* in handling large-scale datasets and complex time series tasks.

## Conclusion

We propose *TSGDiff*, a novel framework for structured time series generation that rethinks the problem from a graph perspective. By unifying dynamic graph construction, Fourier spectrum-based edge modeling, and diffusion-based generative processes in latent graph space, *TSGDiff* effectively captures intricate temporal dependencies and structural relationships. The model generates time series with realistic dynamics, high temporal accuracy, and structural coherence, addressing key challenges in traditional approaches. Additionally, we introduce the Topological Structure Fidelity (Topo-FID) score, a novel graph-aware metric that combines Graph Edit Similarity and Structural Entropy Similarity to comprehensively evaluate the quality of generated time series. This metric fills critical gaps in traditional evaluation methods by providing a more accurate and meaningful assessment of structural similarity. *TSGDiff* offers a robust and flexible solution for synthetic time series generation, setting a foundation for future research into scalable, adaptive, and high-fidelity generative models.

## Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grant Nos. 62221005, 62450043, 62222601, 62176033 and 62576056.

## References

- Alaa, A. M.; Chan, A. J.; and van der Schaar, M. 2021. Generative time-series modeling with fourier flows. In *ICLR*.
- Bao, W.; Yue, J.; and Rao, Y. 2017. A deep learning framework for financial time series using stacked autoencoders and long-short term memory. *PLoS ONE*, 12(7).
- Bracewell, R. N.; and Bracewell, R. N. 1986. *The fourier transform and its applications*, volume 31999. McGraw - Hill New York.
- Chen, Y.; Xu, P.; Chu, Y.; Li, W.; Wu, Y.; Ni, L.; Bao, Y.; and Wang, K. 2017. Short-term electrical load forecasting using the Support Vector Regression (SVR) model to calculate the demand response baseline for office buildings. *Applied Energy*, 195: 659–670.
- Cheng, Z.; Yang, Y.; Wang, W.; Hu, W.; Zhuang, Y.; and Song, G. 2020. Time2graph: revisiting time series modeling with dynamic shapelets. In *AAAI*.
- Crabbé, J.; Huynh, N.; Stanczuk, J.; and van der Schaar, M. 2024. Time series diffusion in the frequency domain. *ICML*.
- der Maaten, L. V.; and Hinton, G. 2008. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(11).
- Desai, A.; Freeman, C.; Beaver, I.; and Wang, Z. 2022. TimeVAE: a variational auto-encoder for multivariate time series generation. In *ICLR*.
- Esteban, C.; Hyland, S. L.; and Ratsch, G. 2017. Real-valued (medical) time series generation with recurrent conditional GANs. *Preprint arXiv:1706.02633*.
- Fortuin, V.; Baranchuk, D.; Ratsch, G.; and Mandt, S. 2020. GP-VAE: deep probabilistic time series imputation. *AIS-TATS*, 108: 1651–1661.
- Han, H.; Zhang, M.; Hou, M.; Zhang, F.; Wang, Z.; Chen, E.; et al. 2020. STGCN: A Spatial-temporal aware graph learning method for POI recommendation. In *ICDM*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. In *NeurIPS*.
- Isufi, E.; Gama, F.; and Ribeiro, A. 2021. EdgeNets: edge varying graph neural networks. *IEEE*, 44: 7457–7473.
- Katznelson, Y. 1970. *An introduction to harmonic analysis*. Cambridge University Press.
- Kieu, T.; Yang, B.; Guo, C.; and Jensen, C. S. 2019. Outlier detection for time series with recurrent autoencoder ensembles. In *IJCAI*.
- Kipf, T. N.; and Welling, M. 2017. Semi-supervised classification with graph convolutional networks. In *ICLR*.
- Naiman, I.; Berman, N.; Pemper, I.; Arbiv, I.; Fadlon, G.; and Azencot, O. 2024. Utilizing image transforms and diffusion models for generative modeling of short and long time series. In *NeurIPS*.
- NASA. 2015. Pluto: The 'other' red planet. <https://www.nasa.gov/nh/pluto-the-other-red-planet>. Accessed: 2018-12-06.
- Ni, H.; Szpruch, L.; Wiese, M.; Liao, S.; and Xiao, B. 2020. Conditional sig-wasserstein gans for time series generation. *Preprint arXiv:2006.05421*.
- Park, J.; Lee, S.; Jeong, W.; Choi, Y.; and Lee, J. 2024. Leveraging priors via diffusion bridge for time series generation. *Preprint arXiv:2408.06672v1*.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; Desmaison, A.; Köpf, A.; Yang, E.; DeVito, Z.; Raison, M.; Tejani, A.; Chilamkurthy, S.; Steiner, B.; Fang, L.; Bai, J.; and Chintala, S. 2019. PyTorch: an imperative style, high-performance deep learning library. *NeurIPS*, 32.
- Paul, J.; Michael, B.-S.; Pedro, M.; Shubham, K.; Rajbir, S. N.; Valentin, F.; Jan, G.; and Tim, J. 2022. Psa-gan: progressive self attention gans for synthetic time series. In *ICLR*.
- Segarra, S.; Mateos, G.; Marques, A. G.; and Ribeiro, A. 2017. Blind identification of graph filters. *IEEE Trans. Signal Process.*, 65(5): 1146–1159.
- Shen, L.; and Kwok, J. T. 2023. Non-autoregressive conditional diffusion models for time series prediction. *ICML*, 202: 30512–30525.
- Shen, L.; Yu, Z.; Ma, Q.; and Kwok, J. T. 2021. Time series anomaly detection with multiresolution ensemble decoding. In *AAAI*.
- Song, J.; Meng, C.; and Ermon, S. 2021. Denoising diffusion implicit models. *ICLR*.
- Tashiro, Y.; Song, J.; Song, Y.; and Ermon, S. 2021. CSDI: conditional score-based diffusion models for probabilistic time series imputation. In *NeurIPS*.
- Yi, K.; Zhang, Q.; Fan, W.; He, H.; Hu, L.; Wang, P.; An, N.; Cao, L.; and Niu, Z. 2023. FourierGNN: rethinking multivariate time series forecasting from a pure graph perspective. In *NeurIPS*.
- Yi, K.; Zhang, Q.; Fan, W.; He, H.; Hu, L.; Wang, P.; An, N.; Cao, L.; and Niu, Z. 2024. Fourier graph neural network for multivariate time series forecasting. In *ICLR*.
- Yoon, J.; Jarrett, D.; and der Schaar, M. V. 2019. Time-series generative adversarial networks. In *NeurIPS*, volume 32.
- Yu, H.-F.; Rao, N.; and Dhillon, I. S. 2016. Temporal regularized matrix factorization for high-dimensional time series prediction. *NeurIPS*.
- Yuan, X.; and Qiao, Y. 2024. Diffusion-TS: interpretable diffusion for general time series generation. In *ICLR*.
- Zammel, Z.; Khabou, N.; Souifi, L.; and Rodriguez, I. B. 2023. Time series prediction models in healthcare: Systematic literature review. *ReDCAD Laboratory, ENIS, University of Sfax, Tunisia*.
- Zhou, L.; Poli, M.; Xu, W.; Massaroli, S.; and Ermon, S. 2023. Deep latent state space models for time-series generation. *ICML*.