
Utilizing Image Transforms and Diffusion Models for Generative Modeling of Short and Long Time Series

Ilan Naiman* Nimrod Berman* Itai Pemper Idan Arbiv Gal Fadlon Omri Azencot

Department of Computer Science
Ben-Gurion University of The Negev

{naimani, bermann, itaipem, arbivid, galfad}@post.bgu.ac.il
azencot@cs.bgu.ac.il

Abstract

Lately, there has been a surge in interest surrounding generative modeling of time series data. Most existing approaches are designed either to process short sequences or to handle long-range sequences. This dichotomy can be attributed to gradient issues with recurrent networks, computational costs associated with transformers, and limited expressiveness of state space models. Towards a unified generative model for varying-length time series, we propose in this work to transform sequences into images. By employing invertible transforms such as the delay embedding and the short-time Fourier transform, we unlock three main advantages: i) We can exploit advanced diffusion vision models; ii) We can remarkably process short- and long-range inputs within the same framework; and iii) We can harness recent and established tools proposed in the time series to image literature. We validate the effectiveness of our method through a comprehensive evaluation across multiple tasks, including unconditional generation, interpolation, and extrapolation. We show that our approach achieves consistently state-of-the-art results against strong baselines. In the unconditional generation tasks, we show remarkable mean improvements of 58.17% over previous diffusion models in the short discriminative score and 132.61% in the (ultra-)long classification scores. Code is at <https://github.com/azencot-group/ImagenTime>.

1 Introduction

Generative modeling of real-world information such as images [72], texts [13], and other types of data [99, 55, 8] has drawn increased attention recently. In this work, we focus on the setting of generative modeling (GM) of general time series information. There are several factors that govern the complexity required from sequential data generators including the sequence length, its number of features, the appearance of transient vs. long-range effects, and more. Existing generative models for time series are typically designed either for multivariate short-term sequences [44, 19] or univariate long-range data [103], often resulting in separate and completely different neural network frameworks. However, a natural question arises: Can one develop a unified framework equipped to handle both high-dimensional short sequences and low-dimensional long time series?

Earlier approaches for processing time series based on recurrent neural networks (RNNs) handled short sequences well [62, 3, 43, 76], however, modeling long-range dependencies turned out to be significantly more challenging. Particularly, RNNs suffer from the well-known vanishing and exploding gradient problem [9, 70] that prevents them from learning complex patterns and long-range dependencies. To address long-context modeling and memory retention, extensive research is devoted to approaches such as long short-term memory (LSTM) models [42], unitary evolution RNNs [5]

*Equal Contribution

and Lipschitz RNNs [24]. A different approach for processing sequential information is based on the Transformer [93], eliminating any recurrent connections. Recent remarkable results have been obtained with transformers on natural language processing [13] and time series forecasting [96, 104, 68] tasks. Alas, transformers are underexplored as generative models for long-range time series data. This may be in part due to their computational costs that scale quadratically as $\mathcal{O}(L^2)$ with the sequence length L , and in part because transformer forecasters are inferior to linear tools [101].

Beyond RNNs and the Transformer, recent works have considered the state space model (SSM) for modeling long-range time series information. For instance, the structured SSM (S4) [36] employed a parameterization that reduced computational costs via evaluations of Cauchy kernels. Further, the deep linear recurrent unit (LRU) is inspired by the similarities between SSMs and RNNs, and it demonstrated impressive performance in modeling long-range dependencies (LRD). Still, generative modeling of long-range sequential data via state space models remains largely underexplored. Recent work suggested LS4 [103], a latent time series generative model that builds upon linear state space equations. LS4 utilizes autoregressive dependencies to expressively model time series (potentially non-stationary) distributions. However, this model struggles with short-length sequences as we show in our study, potentially due to limited expressivity of linear SSMs.

To overcome gradient issues of recurrent backbones, temporal computational costs of transformers, and expressivity problems of SSMs, we represent time series information via small-sized *images*. Transforming raw sequences to other encodings has been useful for processing audio [34] as well as general time series data [95, 38, 56]. Moreover, a similar approach was employed to generative modeling of time series with generative adversarial networks (GANs) [12, 39]. However, unstable training dynamics and mode collapse negatively affect the performance of GAN-based tools [59]. In contrast, transforming time series to images is underexplored in the context of generative *diffusion* models. There are several fundamental advantages to our approach. First, there have been remarkable advancements in diffusion models for vision data that we can exploit [81, 40, 86, 45]. Second, using images instead of sequences elegantly avoids the challenges of long-term modeling. For instance, a moderately-sized 256×256 image corresponds to a time series of length up to $65k$, as we show in Sec. 3. Finally, there is a growing body of literature dealing with time series as images on generative, classification, and forecasting tasks, whose results can be applied in our work and in future studies.

In this work, we propose a new diffusion-based framework for generative modeling of general time series data, designed to seamlessly process both short-, long-, and *ultra*-long-range sequences. To evaluate our method, we consider standard benchmarks for short to ultra-long time series focusing on unconditional generation. Our approach supports efficient sampling, and it attains state-of-the-art results in comparison to recent generative models for sequential information. As far as we know, there are no existing tools handling both short and long sequence data. In addition to its strong unconditional generation capabilities, our approach is also tested in conditional scenarios involving the interpolation of missing information and extrapolation. Overall, we obtained state-of-the-art results in such cases with respect to existing tools. We further analyze and ablate our technique to motivate some of our design choices. The contributions of our work can be summarized as follows:

1. We view generative modeling of time series as a visual challenge, allowing to harness advances in time series to image transforms as well as vision diffusion models.
2. We develop a novel generative model for time series that scales from short to very long sequence lengths without significant modifications to the neural architecture or training method.
3. Our approach achieves state-of-the-art results in comparison to strong baselines in unconditional and conditional generative benchmarks for time series of lengths in the range $[24, 17.5k]$. Particularly, we attain the best scores on a new challenging benchmark of very long sequences that we introduce.

2 Related work

Time series to image works. Motivated by the success of convolutional neural networks on vision data, several works have transformed time series to images using Gramian Angular Fields [95], Recurrence Plots [38], and Line Graphs [56]. This innovation allows leveraging computer vision techniques, tested on tasks such as time series classification and imputation. In speech analysis and processing, the short-time Fourier transform (STFT) stands out as a widely used method [1, 2, 94, 26]. It tracks the changes in frequency components over time, making it essential for analyzing and understanding audio and speech data. Recent research [71, 17] has explored mel-spectrogram

transforms within diffusion models, including integration with advanced latent diffusion spaces [58]. Furthermore, combining time series images and Wasserstein GANs [12, 39] have been considered for generative modeling. Yet, representing general time series as images within diffusion models for tasks such as unconditional generation, interpolation, and extrapolation, remains largely underexplored. The goal of this work is to make a step toward bridging this gap.

Diffusion models. Both denoising diffusion probabilistic models (DDPM) [81, 40] and score-based generative models [84, 85] have demonstrated their effectiveness across diverse domains including images [74, 41], audio [15, 52], and graphs [69, 97, 10]. Song et al. [86] showed that DDPM and score-based models can be both interpreted as stochastic differential equations (SDE). Further works focused on improving generation quality by using latent diffusion processes in autoencoder architectures [74]. Another research direction deals with lowering the number of neural function evaluations (NFEs), which originally ranged from hundreds to thousands NFEs. For instance, Karras et al. [45] obtain a low Fréchet inception distance (FID) with only 35 evaluations, whereas the recent consistency models [82] achieved comparable results with only a single function evaluation.

Generative modeling of time series. Generative adversarial networks (GANs) [31] have shown remarkable success in generating realistic data across various domains. Specifically, their application to time series information by joint optimization of supervised and adversarial objectives via TimeGAN captured the inherent dynamics of real-world signals [99]. Similarly, GT-GAN [44] utilizes diverse tools including ordinary differential equations (DE) [16], neural controlled DE [51], and continuous time-flow processes to model both regularly- and irregularly-sampled data [21]. Nevertheless, GANs suffer from challenges, primarily due to unstable training dynamics and mode collapse [59]. Beyond GANs, variational autoencoders (VAEs) have been also considered for generative modeling of sequential data [22, 54, 73], where the work [66] achieved strong results using Koopman-based approaches [7, 6, 65, 11]. To process long-range dependencies and stiff dynamics [79], Zhou et al. [103] introduced LS4, a latent generative model based on linear state space equations. Following the success of diffusion models in other domains, there is a growing desire to adapt them for time series data. However, this adaption is not straightforward and entails the design of a suitable backbone [88, 57, 19, 100, 67]. Other approaches focused on regression problems [49], based on manifold learning tools [47, 48]. Instead, we propose a new framework for generative modeling of time series by transforming such data to images and using existing strong diffusion vision models.

3 Background

In what follows, we state the problem, we mention two effective time series to image transformations, and we briefly discuss the essentials of diffusion-based generative modeling.

Problem statement. We address the problem of generating time series (TS), sampled from a learned distribution $\tilde{p}(x)$ that is similar to an unknown distribution $p(x)$, for which we have a set of observed TS data. The given observations include data samples $x \in \mathbb{R}^{L \times K}$, where L represents the sequence length and K denotes the number of features. Formally, the generative modeling task is often termed "unconditional generation" [37], and it entails learning a model M capable of sampling unseen time series \tilde{x} from $\tilde{p}(x)$. Additionally, our work addresses a secondary problem known as "conditional generation". In this setting, given an additional signal c , we learn the unknown (conditional) distribution $p(x|c)$. For example, the signal c can be an observed part from the TS. This conditional modeling proves useful for tasks such as time series interpolation and extrapolation.

Time series to image transforms. We focus in our study on two invertible time series to image transformations: 1) the delay embedding; and 2) the short time Fourier transform. We provide below a brief overview of these transforms and their inverse. We consider additional transforms and we discuss more details in App. A. Fig. 3 illustrates a time series signal and its related images.

Delay embeddings [87] transform a univariate time series $x_{1:L} \in \mathbb{R}^L$ to an image by arranging the information of the series in columns and pad if needed. Let m, n be two user parameters representing

the skip value and the column dimension, respectively. We construct the following matrix X ,

$$X = \begin{bmatrix} x_1 & x_{m+1} & \cdots & x_{L-n} \\ \vdots & \vdots & \cdots & \vdots \\ x_n & x_{n+m+1} & \cdots & x_L \end{bmatrix} \in \mathbb{R}^{n \times q}, \quad (1)$$

where $q = \lceil (L - n)/m \rceil$. The image x_{img} is created by padding with zeros to fit the neural network input constraints. Given x_{img} , the original time series $x_{1:L}$ can be extracted in multiple ways. For instance, if $m = 1$, then $x_{1:L}$ is formed by concatenating the first row and last column of x_{img} . The delay embedding scales naturally to long sequences, e.g., setting $m = n = 256$ allows to encode $65k$ sequences with 256×256 images.

Short Time Fourier Transform (STFT) [35] is a well-known transformation that maps a signal from its original domain into the frequency domain. To preserve the temporal structure, STFT applies a rolling window on the time axis, extracting time series segments for which the fast Fourier transform (FFT) is applied. Given an input signal $x \in \mathbb{R}^{L \times K}$, STFT produces an image $x_{\text{img}} \in \mathbb{R}^{2K \times H \times W}$, where the channels are doubled to store the real and imaginary parts, and H, W are derivatives of user parameters. STFT requires a minimum window length, and thus, short sequences may require a linear interpolation to match length constraints. Remarkably, the short time Fourier transform is invertible via reverse STFT with a negligible loss of information. Importantly, in contrast to the common practice in audio processing, we do not further compute the spectrogram of STFT, avoiding non-trivial inverse transformations.

Diffusion models. Diffusion processes gradually add noise to an image, following a predefined noise scheduling scheme. Generating new images is possible by learning a model that removes noise. The diffusion process $\{\mathbf{x}(t)\}_{t=0}^T$ is the path of a stochastic differential equation (SDE) [86], where an initial sample $\mathbf{x}(0)$ is drawn from the data distribution $p_0(\mathbf{x})$. The initial sample is modified to $\mathbf{x}(T)$, sampled from a simple prior distribution such as a normal Gaussian $\mathcal{N}(0, I)$. Formally, the forward process is governed by an SDE of the form,

$$d\mathbf{x} = f(\mathbf{x}, t)dt + gdw, \quad (2)$$

where $f(\cdot, t) : \mathbb{R}^d \rightarrow \mathbb{R}^d$ represents the drift coefficient, $g \in \mathbb{R}$ is the diffusion scalar, and w denotes a standard Wiener process. To facilitate sampling, we need to derive the reverse SDE. It is well-known that the reverse process [4] is given by,

$$d\mathbf{x} = [f(\mathbf{x}, t) - g^2 \nabla_{\mathbf{x}} \log p_t(\mathbf{x})]d\bar{t} + g d\bar{w}, \quad (3)$$

where \bar{t} denotes reverse time and \bar{w} is a reverse Wiener process. Given Eq. (3), one can derive a deterministic process, characterized by trajectories that share identical marginal probability densities. Formally, we obtain the following ordinary differential equation (ODE),

$$d\mathbf{x} = f(\mathbf{x}, t) - g^2 \nabla_{\mathbf{x}} \log p_t(\mathbf{x}). \quad (4)$$

Diffusion models compute an estimator $s_\theta(\mathbf{x}, t)$ to approximate the infeasible $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ via

$$\min_{\theta} \mathbb{E}_t \{ \mathbb{E}_{x_0, x_t} | s_\theta(\mathbf{x}, t) - \nabla_{\mathbf{x}} \log p_{0t}(\mathbf{x}_t | x_0) |^2 \}, \quad (5)$$

where p_{0t} denotes the joint distribution of the initial data and noisy sample. In practice, minimizing Eq. (5) is done by learning the noise pattern of input images. For a more comprehensive discussion regarding score-based models, we refer to [86, 45]. We specify in Sec. 4 the particular diffusion model employed in this work, along with further additional details.

4 Method

Our approach to generative modeling of time series information is based on the following simple observation and straightforward idea. We observe that diffusion models for vision have demonstrated remarkable progress and results recently [40, 74, 72]. Therefore, our idea is to transform sequences to images, allowing their processing using established diffusion vision models. Fortunately, there are several efficient time series to image maps with effective inverse transforms [87, 35, 95, 56]. Our computational pipeline is composed of three main building blocks: 1) a time series to image module; and 2) a diffusion model; and 3) an image to time series component. The diffusion model is the

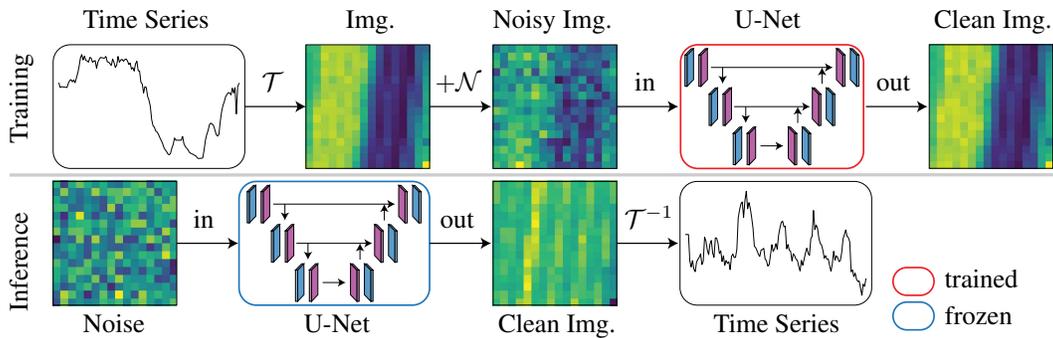


Figure 1: Our training pipeline (top) involves transforming a time series signal to its e.g., delay embedding image, process the image with a diffusion model, and output its cleaned version. During inference (bottom), we sample from a standard normal distribution and obtain a clean image using the trained diffusion model. Finally, we transform the image back to the time series domain.

only learnable parameter-based part of our neural network. We illustrate our generative modeling framework for time series information in Fig. 1, depicting the above building blocks with proper notations for inputs and outputs. Formally, given an input time series $x \in \mathbb{R}^{L \times K}$ with L the sequence length and K the number of features, we transform it to an image $x_{\text{img}} \in \mathbb{R}^{C \times H \times W}$. Noise is added to the latter image yielding the tensor $x_{\text{img}}(t)$ which is processed with our diffusion model, whose output $s(x_{\text{img}}, t) \in \mathbb{R}^{C \times H \times W}$ represents the cleaned image. During inference, noise $x_{\text{img}}(T)$ is sampled from $\mathcal{N}(0, I)$, iterated backward to $x_{\text{img}}(0)$ and transformed to a time series $\tilde{x} \in \mathbb{R}^{L \times K}$.

There are several options to choose from regarding the time series to image (`ts2img`) transform, $\mathcal{T}: \mathbb{R}^{L \times K} \rightarrow \mathbb{R}^{C \times H \times W}$. While all transforms are applicable in our framework, we opt for `ts2img` maps that are efficient to compute, provide informative images, scalable across short and long sequences, and have a closed-form inverse. For instance, line graphs [56] are efficient with a closed-form inverse, however, they produce images that are mostly non-informative as they contain blank pixels. Similarly, the Gramian angular field transform [95] essentially stores the sequence in its main diagonal, and thus, it is not straightforward to apply it to long-range data. In this work, we focus on using the delay embedding and short time Fourier transforms. Both `ts2img` maps satisfy all requirements above. Moreover, our empirical ablation analysis in Sec. 5.5 highlights that these transformations attain the best results on average. The inverse transforms \mathcal{T}^{-1} for delay embedding and STFT are parameter-less, deterministic, and highly efficient. See Sec. 3 and App. A.

At the heart of our **ImagenTime** framework lies the generative diffusion model backbone. Diffusion models for vision data have enjoyed increased attention over the past few years, with strong techniques appearing at an unprecedented rate [81, 84, 40, 85, 86, 74]. One limitation, shared among all diffusion models, is the requirement to iteratively denoise the image during inference, resulting in costly neural function evaluations (NFEs). While multiple works focused on alleviating this issue [77, 83, 28, 82], the work by Karras et al. [45] offers an enhanced score-based model with a good balance between rapid sampling and high-quality generations. Specifically, they presented a clearer design space for the factors that determine the performance of diffusion models, and they suggested EDM that employs a second-order ODE for the reverse process, yielding low FID images in 35 NFEs. Thus, we utilize in this work the EDM diffusion model as our generative backbone.

We conclude by briefly discussing the training and inference procedures. During *training*, we process batches of time series data X , for which we apply \mathcal{T} using either delay embedding or STFT to obtain a batch of images, i.e., $X_{\text{img}} = \mathcal{T}(X)$. Subsequently, we employ the training procedure of EDM to learn the score function $s_{\theta}(X_{\text{img}}, t)$. For *inference*, we use the trained EDM model and we compute the reverse ODE in Eq. (4) for sampling new data points. In practice, we follow the same inference procedure specified in [45]. Finally, given a batch of sampled images, \tilde{X}_{img} , we apply the inverse transform \mathcal{T}^{-1} to achieve a batch of generated time series samples, i.e., $\tilde{X} = \mathcal{T}^{-1}(\tilde{X}_{\text{img}})$.

5 Experiments

We use standard unconditional and conditional quantitative and qualitative benchmarks to extensively validate our framework’s ability to generate high-quality time series samples. First, we test our

Table 1: Error measures for the short time series unconditional discriminative and prediction tasks.

Method	Stocks		Energy		MuJoCo	
	disc↓	pred↓	disc↓	pred↓	disc↓	pred↓
KoVAE	.009 ± .006	.037 ± .000	.143 ± .011	.251 ± .000	.076 ± .017	.038 ± .002
DiffTime	.050 ± .017	.038 ± .001	.101 ± .019	.250 ± .003	.059 ± .009	.042 ± .000
GT-GAN	.077 ± .031	.040 ± .000	.221 ± .068	.312 ± .002	.245 ± .029	.055 ± .000
TimeGAN	.102 ± .021	.038 ± .001	.236 ± .012	.273 ± .004	.409 ± .028	.082 ± .006
RCGAN	.196 ± .027	.040 ± .001	.336 ± .017	.292 ± .004	.436 ± .012	.081 ± .003
C-RNN-GAN	.399 ± .028	.038 ± .000	.449 ± .001	.483 ± .005	.412 ± .095	.055 ± .004
T-Forcing	.226 ± .035	.038 ± .001	.483 ± .004	.315 ± .005	.499 ± .000	.142 ± .014
P-Forcing	.257 ± .026	.043 ± .001	.412 ± .006	.303 ± .005	.500 ± .000	.102 ± .013
WaveNet	.232 ± .028	.042 ± .001	.397 ± .010	.311 ± .006	.385 ± .025	.333 ± .004
WaveGAN	.217 ± .022	.041 ± .001	.363 ± .012	.307 ± .007	.357 ± .017	.324 ± .006
LS4	.199 ± .065	.068 ± .013	.474 ± .003	.251 ± .000	.333 ± .029	.062 ± .006
Ours	.037 ± .006	.036 ± .000	.040 ± .004	.250 ± .000	.007 ± .005	.033 ± .001

framework on short-term and long-term standard time series unconditional generation benchmarks (Sec. 5.1, Sec. 5.2). Then, we introduce a novel benchmark for ultra-long sequences (above 10k steps) and evaluate our method in comparison to strong baselines (Sec. 5.3). Then, we consider interpolation and extrapolation benchmarks, similar to [78], to test our model on conditional generation tasks (Sec. 5.4). Further, we extended these benchmarks with additional short- and ultra-long setups, which test the framework’s robustness to lengths. Finally, we conclude with an extensive ablation of our framework (Sec. 5.5). More details on the experimental settings can be found in App. B.

5.1 Short-Term Unconditional Generation

Data, baselines, and metrics. We employ our framework on the unconditional generation benchmark reported in [19]. The benchmark includes four synthetic and real-world datasets with a fixed length of 24. The first dataset, *Stocks*, consists of daily historical Google stock data from 2004 to 2019, comprising six channels: high, low, opening, closing, and adjusted closing prices, as well as volume. This data lacks periodicity and is dominated by random walks. The second dataset, *Energy*, is a multivariate appliance energy prediction dataset [14], featuring 28 channels with correlated features, and it exhibits noisy periodicity and continuous-valued measurements. The third dataset, *MuJoCo* (Multi-Joint dynamics with Contact), serves as a versatile physics generator for simulating TS data with 14 channels [89]. We report results on the simple synthetic *Sine* dataset of sine functions in App. C.1. Our framework is compared with state-of-the-art short-term time series generative models. KoVAE [66], DiffTime [19], GT-GAN [44], TimeGan [99], RCGAN [25], C-RNN-GAN [64], T-Forcing [33], P-forcing [32], WaveNet [91], WaveGAN [23], and LS4 [103], which is the state-of-the-art generative model for modeling long sequences. The benchmark employs two metrics: 1) The *Predictive (pred)* metric assesses the utility of the generated data. 2) The *Discriminative (disc)* metrics gauge the similarity of distributions using a proxy discriminator. For all experiments, we used the *delay embedding* transform with an embedding of $n = 8$ and a delay of $m = 3$, yielding a 8×8 image. We use 18 sampling steps with the EDM model [45] as the diffusion generative backbone.

Quantitative and qualitative results. The results for the short-term unconditional benchmark are shown in Tab. 1. Our framework achieves state-of-the-art results on all datasets and metrics. Particularly, we note MuJoCo, where we improved the second-best method by 88% and 21% in the discriminative and predictive scores. In general, the second-best approach is DiffTime. Importantly, while LS4 performs well on long sequences, our results indicate that it struggles with short sequences. In comparison, we will show below that in addition to obtaining SOTA results on short-term time series, we also achieve strong results in the long-term case (Sec. 5.2). We also evaluate our method using two common qualitative tests [99]. First, we compute a two-dimensional t-SNE [92] embedding for real and synthetic data. The desired outcome is that both datasets span similar regions and shapes in 2D. We plot the embeddings of the real data, ours and GT-GAN in Fig. 2A, highlighting that our generated point clouds are closer to the real data in comparison GT-GAN. Tab. 11 reports the Wasserstein distances between the t-SNE embeddings of the generated data and the real data, showing that our approach is superior to GT-GAN. Second, we estimate the probability density functions in Fig. 2D. Our approach generates densities similar to the real densities, whereas GT-GAN introduces noticeable errors. The rest of the short-term datasets’ qualitative analysis appears in App. C.3.

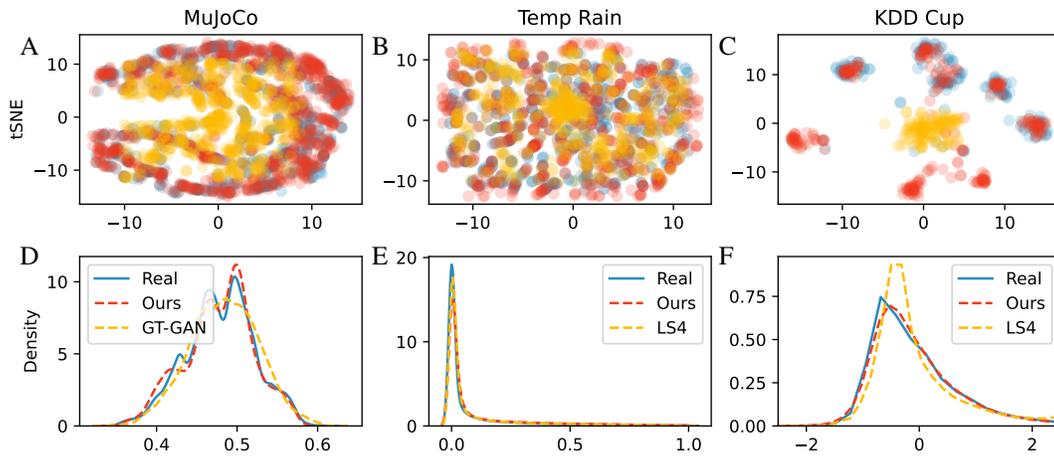


Figure 2: We plot the 2D t-SNE embeddings of synthetic data generated with our method and SOTA tools vs. the real data (top). Then, we compare their probability density functions (bottom).

Table 2: Long time series unconditional marginal, classification, and prediction tasks’ results.

Method	FRED-MD			NN5 Daily			Temp Rain		
	marg↓	class↑	pred↓	marg↓	class↑	pred↓	marg↓	class↑	pred↓
RNN-VAE	.132	.036	1.47	.137	.000	.967	.017	.000	159
GP-VAE	.152	.016	2.05	.117	.002	1.17	.183	.000	2.31
ODE ² VAE	.122	.028	.567	.211	.001	1.19	1.83	.000	1.13
Latent ODE	.042	.327	.013	.107	.000	1.04	.011	.000	145
TimeGAN	.081	.029	.058	.040	.001	1.34	.498	.003	1.96
SDEGAN	.084	.501	.677	.085	.085	1.01	.990	.017	2.46
SaShiMi	.048	.001	.232	.020	.045	.849	.758	.000	2.12
LS4	.022	.544	.037	.007	.636	.241	.083	.976	.521
Ours	.021	.862	.009	.005	1.02	.393	.409	5.80	.377

5.2 Long-Term Unconditional Generation

Data, baselines, and metrics. We utilize the long-term time series benchmark presented in [103]. It includes three long-term real-world time series datasets obtained from the Monash Time Series Forecasting Repository [29]: FRED-MD, NN5 Daily, and Temperature Rain. We omit Solar Weekly as it is short-term. These datasets were chosen based on their average 1-lag autocorrelation metric, measuring their correlation over time. Their 1-lag values range from 0.38 to 0.98, highlighting a diverse range of temporal dynamics that present challenges for generative learning tasks. Each dataset contains approximately 750 time steps. We compare our method with state-of-the-art long-term generative methods: LS4 [103], SaShiMi [30], SDEGAN [50], TimeGAN [99], Latent ODE [75], ODE²VAE [98], GP-VAE [27] and RNN-VAE [18]. Three different metrics are used to evaluate the generative performance: Marginal (marg), Classification (class), and Prediction (pred). Marginal scores measure the absolute difference between the empirical probability density functions of two distributions. Classification scores use a sequence model to classify samples as real or generated; high scores indicate less distinguishable samples. Prediction scores utilize a train-on-synthetic-test-on-real sequence-to-sequence model to predict future steps; lower scores indicate higher predictability. We used the STFT transform in all our experiments, creating a 32×32 size image. The number of sampling steps is 18 and we use the EDM model [45] as the diffusion generative backbone.

Quantitative and qualitative results. We present the results in Tab. 2. LS4 performs well across most datasets and metrics. In comparison, our method outperforms LS4 and the other techniques in almost all cases. On NN5 Daily pred and on Temp Rain marg we achieve inferior results. Notably, Zhou et al. [103] discuss the challenge of measuring the marginal score for the Temp Rain dataset due to frequent zero values. We highlight that our approach substantially improves the classification and prediction scores for the Temp Rain dataset. Additionally, our framework achieves strong results in the classification scores for the FRED-MD and NN5 Daily datasets. We report our results with standard deviations in App. C.4; the results emphasize the statistical significance improvement our framework achieves. Our qualitative results are shown in Fig. 2(B, D) and in App. C.3.

5.3 Ultra-long Term Unconditional Generation

Data, baselines, and metrics. We conclude our unconditional generation evaluation by considering the challenging setting of *ultra-long* sequences. As far as we know, this setup is underexplored in the literature, and moreover, considering short, long, and ultra-long time series for a single framework is novel to our work. Specifically, we use the following real-world datasets from the Monash Time Series Forecasting Repository [29]: San Francisco Traffic (Traffic) [53] and KDD-Cup 2018 (KDD-Cup) [61]. The datasets’ lengths are 17544 and 10920, respectively. Traffic includes an hourly time series detailing the road occupancy rates on the San Francisco Bay Area freeways from 2015 to 2016. KDD-Cup represents the air quality level from 2017 to 2018 estimated by 59 stations across two cities, Beijing (35 stations) and London (24 stations), measured in an hourly rate. We process Traffic with the delay embedding transform ($n = 144, m = 136$), yielding 144×144 images. KDD-Cup is transformed by STFT, resulting in 112×112 images. The sampling steps are 36 in both datasets.

Quantitative and qualitative results. As shown in Tab. 3, our method consistently achieves superior results in all cases. Notably, it attains on KDD-Cup a pred score of .001 compared to LS4’s second-best score of .049. These results highlight our framework’s scalability to very long sequences, demonstrating impressive performance across all sequence lengths as we demonstrated in the previous sections. We also report results with standard deviations in App. C.6, emphasizing the statistical significance of our framework. Finally, our qualitative results for this setting are shown in Fig. 2(C, E) and in App. C.3.

Table 3: Ultra-long unconditional generation.

Method	Traffic			KDD-Cup		
	pred ↓	class ↑	marg ↓	pred ↓	class ↑	marg ↓
Latent-ODE	1.01	.000	.180	.079	.013	.009
LS4	.170	.630	.002	.049	.488	.002
Ours	.138	.684	.001	.001	.842	.001

5.4 Conditional Generation of Time Series

In addition to the unconditional generation benchmark we consider above, we also evaluate our approach on conditional generation tasks. We focus on the imputation (interpolation) and forecasting (extrapolation) tasks, following the experimental setup in [75, 78]. Our approach can be adapted to solve these tasks via a simple modification. For instance, in the interpolation task, the goal is to generate the missing values. Thus, we apply our diffusion model only in the missing locations using a corresponding mask. The rest of the values are left unchanged. A similar mechanism can be applied to extrapolation. Generally, this approach is similar to image inpainting techniques [60]. In the interpolation task, we randomly mask 50% of the sequence values, whereas in the extrapolation challenge, we split the sequence in half, where the second half represents the target values. In this benchmark, we consider short, long, and ultra-long sequences. Our comparison focuses on generative methods that can handle long-range dependencies including ODE-RNN [75], Latent ODE [75], CRU [78], and LS4 [103]. Further details about the experiments can be found in App. B.

Datasets. In the short-term setting, we use ETT* datasets [102], that contain electricity loads of various resolutions (ETTh1, ETTh2, and ETTm1, ETTm2) from two electricity stations. The sequence length is 96. For the long-term case, we utilize an established benchmark [75, 78, 103], including the Physionet and USHCN datasets. The Physionet dataset [80] includes health measurements of 41 sensors collected from 8000 ICU patients within the first 48 hours of admission. The United States Historical Climatology Network (USHCN) [63] consists of daily measurements from 1218 weather stations across the United States, including data on precipitation, snowfall, snow depth, and minimum and maximum temperatures. For the ultra-long setting, we use the datasets mentioned in Sec. 5.3.

Results. The results of the conditional generation benchmark are detailed in Tab. 4. Values represent the mean squared error (MSE), and thus, lower is better. MSE values are multiplied by $\times 10^{-3}$ and $\times 10^{-2}$ for Physionet and USHCN, respectively, in both experiments. We denote in bold the best method per dataset. The short, long, and ultra-long results are placed at the top, middle, and bottom sections of the table. Overall, our method presents stellar results in all settings, except for ETTm1 where it is second-best. Notably, we mention that in the short interpolation, our results are ≈ 4 times better than the second-best method, CRU. Similarly, we improve the SOTA by $\approx 30\%$ in the short extrapolation. Our results are particularly strong in the long interpolation setting, where we

Table 4: Interpolation and extrapolation results on datasets of varying lengths. The asterisk (*) denotes non-converging runs, running for over seven days.

Dataset	Interpolation					Extrapolation				
	ODE-RNN	Latent ODE	CRU	LS4	Ours	ODE-RNN	Latent ODE	CRU	LS4	Ours
ETTh1	.210	.671	.283	.642	.069	-	1.02	1.02	3.42	.701
ETTh2	.182	.712	.368	3.40	.058	-	1.17	1.09	3.83	.667
ETTm1	.762	.502	.086	.114	.038	-	.592	.643	3.05	.634
ETTm2	.116	.247	.179	.488	.049	-	.414	.378	3.64	.348
Physionet	2.30	2.12	1.82	.620	.004	3.01	4.21	6.29	4.94	1.34
USHCN	8.31	17.9	.160	.050	.006	1.96	2.03	1.27	4.36	1.20
Traffic	.404	.985	*	.990	.090	-	1.01	*	2.23	.221
KDD-Cup	.205	.847	.190	.970	.144	-	.696	.723	6.51	.368

improve by two- and one-orders of magnitude on Physionet and USHCN, respectively. Finally, we also highlight our ultra-long interpolation results which are ≈ 4 times better than ODE-RNN. We conclude that our approach shows robustness to varying sequence lengths, presenting extremely strong results across several datasets and in comparison to state-of-the-art generative models.

5.5 Ablation Studies

We conclude our empirical section by thoroughly inspecting different components of our framework. Specifically, we show that our approach is robust to different image resolutions (App. C.8). We also experiment with a range of hyper-parameters (App. C.9), demonstrating the stability of our approach. Our performance evaluation highlights that our method is comparable to LS4 in terms of training and inference time (App. C.10). Below, we ablate the effect of various image transforms on the performance in the unconditional test. We evaluate our model using four different transforms: folding, Gramian angular field (GAF), delay embedding (DE) and STFT, and we detail the results in Tab. 5. While DE and STFT are slightly better on short and long sequences, respectively, we emphasize that all other transforms perform reasonably well across the various datasets and metrics. GAF does not scale to long sequences as it produces huge images, and thus, it is omitted from the long-term test. We conclude that our framework is robust to the choice of image transformation.

Table 5: Short- and long-term ablation of various image transforms using several datasets and metrics.

	Energy		MuJoCo		FRED-MD			NN5 Daily		
	disc↓	pred↓	disc↓	pred↓	marg↓	class↑	pred↓	marg↓	class↑	pred↓
Folding	<u>.074</u>	.250	<u>.017</u>	.031	.012	1.67	<u>.021</u>	.010	.776	.436
GAF	.349	.269	.049	.034	-	-	-	-	-	-
DE	.040	.250	.007	<u>.033</u>	<u>.017</u>	<u>1.65</u>	<u>.021</u>	<u>.007</u>	.871	<u>.394</u>
STFT	.271	<u>.256</u>	.071	<u>.033</u>	.021	.862	.009	.005	<u>.822</u>	.307

6 Conclusion

While new generative models for general time series data appear rapidly, the majority of existing frameworks are specifically designed to process either short or long sequences. The lack of a unified framework for varying lengths time series can be justified by the shortcomings of current available tools: gradient issues of recurrent networks, temporal computational costs of transformers, and limited expressiveness of state space models. In this work, we address this problem by introducing a novel generative model for time series based on signal-to-image invertible transforms and a vision diffusion backbone. The benefit of our approach is threefold: we exploit advanced diffusion models for vision, we seamlessly process short-to-ultra-long sequences, and we can utilize tools from the signal-to-image literature. We extensively evaluate our framework in the unconditional and conditional settings using short, long, and ultra-long sequences, considering multiple datasets, and in comparison to state-of-the-art models. Our experiments show the superiority of our framework, setting new SOTA results. Further, we demonstrate the robustness of our method through several ablation studies. Our approach requires slightly higher computational resources, which we leave for further consideration and future work. Finally, we believe that the proposed framework has the potential to be applicable in additional tasks including classification, anomaly detection, few-shot learning, and more generally, serve as a foundation model.

Acknowledgements

This research was partially supported by the Lynn and William Frankel Center of the Computer Science Department, Ben-Gurion University of the Negev, an ISF grant 668/21, an ISF equipment grant, and by the Israeli Council for Higher Education (CHE) via the Data Science Research Center, Ben-Gurion University of the Negev, Israel.

References

- [1] J. Allen. Short term spectral analysis, synthesis, and modification by discrete Fourier transform. *IEEE transactions on acoustics, speech, and signal processing*, 25(3):235–238, 1977.
- [2] J. B. Allen and L. R. Rabiner. A unified approach to short-time Fourier analysis and synthesis. *Proceedings of the IEEE*, 65(11):1558–1564, 1977.
- [3] S.-I. Amari. Learning patterns and pattern sequences by self-organizing nets of threshold elements. *IEEE Transactions on computers*, 100(11):1197–1206, 1972.
- [4] B. D. Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982.
- [5] M. Arjovsky, A. Shah, and Y. Bengio. Unitary evolution recurrent neural networks. In *International conference on machine learning*. PMLR, 2016.
- [6] O. Azencot, N. B. Erichson, V. Lin, and M. Mahoney. Forecasting sequential data using consistent Koopman autoencoders. In *International Conference on Machine Learning*, pages 475–485. PMLR, 2020.
- [7] O. Azencot, W. Yin, and A. Bertozzi. Consistent dynamic mode decomposition. *SIAM Journal on Applied Dynamical Systems*, 18(3):1565–1585, 2019.
- [8] O. Bar-Tal, H. Chefer, O. Tov, C. Herrmann, R. Paiss, S. Zada, A. Ephrat, J. Hur, Y. Li, T. Michaeli, O. Wang, D. Sun, T. Dekel, and I. Mosseri. Lumiere: A space-time diffusion model for video generation. *arXiv preprint arXiv:2401.12945*, 2024.
- [9] Y. Bengio, P. Simard, and P. Frasconi. Learning long-term dependencies with gradient descent is difficult. *IEEE transactions on neural networks*, 1994.
- [10] N. Berman, E. Kosman, D. Di Castro, and O. Azencot. Generative modeling of graphs via joint diffusion of node and edge attributes. *arXiv preprint arXiv:2402.04046*, 2024.
- [11] N. Berman, I. Naiman, and O. Azencot. Multifactor sequential disentanglement via structured Koopman autoencoders. In *The Eleventh International Conference on Learning Representations, ICLR, 2023*.
- [12] E. Brophy, Z. Wang, and T. E. Ward. Quick and easy time series generation with established image-based GANs. *arXiv preprint arXiv:1902.05624*, 2019.
- [13] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [14] L. M. Candanedo, V. Feldheim, and D. Deramaix. Data driven prediction models of energy use of appliances in a low-energy house. *Energy and buildings*, 140:81–97, 2017.
- [15] N. Chen, Y. Zhang, H. Zen, R. J. Weiss, M. Norouzi, and W. Chan. Wavegrad: Estimating gradients for waveform generation. In *9th International Conference on Learning Representations, ICLR, 2021*.
- [16] R. T. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.
- [17] Z. Chen, Y. Wu, Y. Leng, J. Chen, H. Liu, X. Tan, Y. Cui, K. Wang, L. He, S. Zhao, J. Bian, and D. P. Mandic. ResGrad: Residual denoising diffusion probabilistic models for text to speech. *arXiv preprint arXiv:2212.14518*, 2022.
- [18] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.

- [19] A. Coletta, S. Gopalakrishnan, D. Borrajo, and S. Vyetenko. On the constrained time-series generation problem. *Advances in Neural Information Processing Systems*, 36, 2024.
- [20] E. De Brouwer, J. Simm, A. Arany, and Y. Moreau. GRU-ODE-Bayes: Continuous modeling of sporadically-observed time series. *Advances in neural information processing systems*, 32, 2019.
- [21] R. Deng, B. Chang, M. A. Brubaker, G. Mori, and A. Lehmann. Modeling continuous stochastic processes with dynamic normalizing flows. *Advances in Neural Information Processing Systems*, 33:7805–7815, 2020.
- [22] A. Desai, C. Freeman, Z. Wang, and I. Beaver. TimeVAE: a variational auto-encoder for multivariate time series generation. *arXiv preprint arXiv:2111.08095*, 2021.
- [23] C. Donahue, J. J. McAuley, and M. S. Puckette. Adversarial audio synthesis. In *7th International Conference on Learning Representations, ICLR*, 2019.
- [24] N. B. Erichson, O. Azencot, A. Queiruga, L. Hodgkinson, and M. W. Mahoney. Lipschitz recurrent neural networks. In *International Conference on Learning Representations*, 2021.
- [25] C. Esteban, S. L. Hyland, and G. Rätsch. Real-valued (medical) time series generation with recurrent conditional GANs. *arXiv preprint arXiv:1706.02633*, 2017.
- [26] P. Flandrin, G. Rilling, and P. Goncalves. Empirical mode decomposition as a filter bank. *IEEE signal processing letters*, 11(2):112–114, 2004.
- [27] V. Fortuin, D. Baranchuk, G. Rätsch, and S. Mandt. GP-VAE: Deep probabilistic time series imputation. In *International conference on artificial intelligence and statistics*, pages 1651–1661. PMLR, 2020.
- [28] Z. Geng, A. Pokle, and J. Z. Kolter. One-step diffusion distillation via deep equilibrium models. *Advances in Neural Information Processing Systems*, 36, 2024.
- [29] R. Godahewa, C. Bergmeir, G. I. Webb, R. J. Hyndman, and P. Montero-Manso. Monash time series forecasting archive. *arXiv preprint arXiv:2105.06643*, 2021.
- [30] K. Goel, A. Gu, C. Donahue, and C. Ré. It’s raw! audio generation with state-space models. In *International Conference on Machine Learning*, pages 7616–7633. PMLR, 2022.
- [31] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.
- [32] A. Goyal, A. Lamb, Y. Zhang, S. Zhang, A. C. Courville, and Y. Bengio. Professor forcing: A new algorithm for training recurrent networks. *Advances in neural information processing systems*, 29, 2016.
- [33] A. Graves. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*, 2013.
- [34] S. Greenberg, W. A. Ainsworth, A. N. Popper, R. R. Fay, N. Mogran, H. Bourlard, and H. Hermansky. Automatic speech recognition: An auditory perspective. *Speech processing in the auditory system*, pages 309–338, 2004.
- [35] D. Griffin and J. Lim. Signal estimation from modified short-time Fourier transform. *IEEE Transactions on acoustics, speech, and signal processing*, 32(2):236–243, 1984.
- [36] A. Gu, K. Goel, and C. Re. Efficiently modeling long sequences with structured state spaces. In *International Conference on Learning Representations*, 2021.
- [37] X. Guo and L. Zhao. A systematic survey on deep generative models for graph generation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(5):5370–5390, 2022.
- [38] N. Hatami, Y. Gavet, and J. Debayle. Classification of time-series images using deep convolutional neural networks. In *Tenth international conference on machine vision (ICMV 2017)*, volume 10696, pages 242–249. SPIE, 2018.
- [39] J. Hellermann and S. Lessmann. Leveraging image-based generative adversarial networks for time series generation. *arXiv preprint arXiv:2112.08060*, 2021.
- [40] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [41] J. Ho, C. Saharia, W. Chan, D. J. Fleet, M. Norouzi, and T. Salimans. Cascaded diffusion models for high fidelity image generation. *Journal of Machine Learning Research*, 23(47):1–33, 2022.

- [42] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [43] J. J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 1982.
- [44] J. Jeon, J. Kim, H. Song, S. Cho, and N. Park. GT-GAN: General purpose time series synthesis with generative adversarial networks. *Advances in Neural Information Processing Systems*, 35:36999–37010, 2022.
- [45] T. Karras, M. Aittala, T. Aila, and S. Laine. Elucidating the design space of diffusion-based generative models. *Advances in Neural Information Processing Systems*, 35:26565–26577, 2022.
- [46] T. Karras, S. Laine, and T. Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019.
- [47] I. Kaufman and O. Azencot. Data representations’ study of latent image manifolds. In *International Conference on Machine Learning, ICML*, volume 202 of *Proceedings of Machine Learning Research*, pages 15928–15945. PMLR, 2023.
- [48] I. Kaufman and O. Azencot. Analyzing deep transformer models for time series forecasting via manifold learning. *Transactions on Machine Learning Research, TMLR*, 2024.
- [49] I. Kaufman and O. Azencot. First-order manifold data augmentation for regression learning. In *Forty-first International Conference on Machine Learning, ICML*, 2024.
- [50] P. Kidger, J. Foster, X. Li, and T. J. Lyons. Neural SDEs as infinite-dimensional GANs. In *International conference on machine learning*, pages 5453–5463. PMLR, 2021.
- [51] P. Kidger, J. Morrill, J. Foster, and T. Lyons. Neural controlled differential equations for irregular time series. *Advances in Neural Information Processing Systems*, 33:6696–6707, 2020.
- [52] Z. Kong, W. Ping, J. Huang, K. Zhao, and B. Catanzaro. DiffWave: A versatile diffusion model for audio synthesis. In *9th International Conference on Learning Representations, ICLR*, 2021.
- [53] G. Lai, W.-C. Chang, Y. Yang, and H. Liu. Modeling long-and short-term temporal patterns with deep neural networks. In *The 41st international ACM SIGIR conference on research & development in information retrieval*, pages 95–104, 2018.
- [54] H. Li, S. Yu, and J. C. Príncipe. Causal recurrent variational autoencoder for medical time series generation. In *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI*, pages 8562–8570, 2023.
- [55] R. Li, X. Li, K.-H. Hui, and C.-W. Fu. SP-GAN: Sphere-guided 3d shape generation and manipulation. *ACM Transactions on Graphics (TOG)*, 40(4):1–12, 2021.
- [56] Z. Li, S. Li, and X. Yan. Time series as images: Vision transformer for irregularly sampled time series. *Advances in Neural Information Processing Systems*, 36, 2023.
- [57] H. Lim, M. Kim, S. Park, J. Lee, and N. Park. TSGM: Regular and irregular time-series generation using score-based generative models. *openreview.com*, 2023.
- [58] H. Liu, Z. Chen, Y. Yuan, X. Mei, X. Liu, D. Mandic, W. Wang, and M. D. Plumbley. AudioLDM: Text-to-audio generation with latent diffusion models. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 21450–21474. PMLR, 23–29 Jul 2023.
- [59] M. Lucic, K. Kurach, M. Michalski, S. Gelly, and O. Bousquet. Are GANs created equal? a large-scale study. *Advances in neural information processing systems*, 31, 2018.
- [60] A. Lugmayr, M. Danelljan, A. Romero, F. Yu, R. Timofte, and L. Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11461–11471, 2022.
- [61] Z. Luo, J. Huang, K. Hu, X. Li, and P. Zhang. Accuair: Winning solution to air quality prediction for KDD cup 2018. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1842–1850, 2019.
- [62] W. S. McCulloch and W. Pitts. A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, 1943.

- [63] M. Menne, C. Williams Jr, and R. Vose. Long-term daily climate records from stations across the contiguous united states, 2015.
- [64] O. Mogren. C-RNN-GAN: Continuous recurrent neural networks with adversarial training. *arXiv preprint arXiv:1611.09904*, 2016.
- [65] I. Naiman and O. Azencot. An operator theoretic approach for analyzing sequence neural networks. In *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI*, pages 9268–9276. AAAI Press, 2023.
- [66] I. Naiman, N. B. Erichson, P. Ren, M. W. Mahoney, and O. Azencot. Generative modeling of regular and irregular time series data via Koopman VAEs. In *The Twelfth International Conference on Learning Representations, ICLR*, 2024.
- [67] S. S. Narasimhan, S. Agarwal, O. Akcin, S. Sanghavi, and S. P. Chinchali. Time weaver: A conditional time series generation model. In *Forty-first International Conference on Machine Learning, ICML*, 2024.
- [68] Y. Nie, N. H. Nguyen, P. Sinthong, and J. Kalagnanam. A time series is worth 64 words: Long-term forecasting with transformers. In *The Eleventh International Conference on Learning Representations, ICLR*, 2023.
- [69] C. Niu, Y. Song, J. Song, S. Zhao, A. Grover, and S. Ermon. Permutation invariant graph generation via score-based generative modeling. In *International Conference on Artificial Intelligence and Statistics*, pages 4474–4484. PMLR, 2020.
- [70] R. Pascanu, T. Mikolov, and Y. Bengio. On the difficulty of training recurrent neural networks. In *International conference on machine learning*. PMLR, 2013.
- [71] V. Popov, I. Vovk, V. Gogoryan, T. Sadekova, and M. Kudinov. Grad-TTS: A diffusion probabilistic model for text-to-speech. In *International Conference on Machine Learning*, pages 8599–8608. PMLR, 2021.
- [72] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022.
- [73] P. Ren, R. Nakata, M. Lacour, I. Naiman, N. Nakata, J. Song, Z. Bi, O. A. Malik, D. Morozov, O. Azencot, et al. Learning physics for unveiling hidden earthquake ground motions via conditional generative modeling. *arXiv preprint arXiv:2407.15089*, 2024.
- [74] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022.
- [75] Y. Rubanova, R. T. Chen, and D. K. Duvenaud. Latent ordinary differential equations for irregularly-sampled time series. *Advances in neural information processing systems*, 32, 2019.
- [76] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning internal representations by error propagation. Technical report, California Univ San Diego La Jolla Inst for Cognitive Science, 1985.
- [77] T. Salimans and J. Ho. Progressive distillation for fast sampling of diffusion models. In *The Tenth International Conference on Learning Representations, ICLR*, 2022.
- [78] M. Schirmer, M. Eltayeb, S. Lessmann, and M. Rudolph. Modeling irregular time series with continuous recurrent units. In *International Conference on Machine Learning*, pages 19388–19405. PMLR, 2022.
- [79] L. F. Shampine and S. Thompson. Stiff systems. *Scholarpedia*, 2(3):2855, 2007.
- [80] I. Silva, G. Moody, D. J. Scott, L. A. Celi, and R. G. Mark. Predicting in-hospital mortality of ICU patients: The physionet/computing in cardiology challenge 2012. In *2012 Computing in Cardiology*, pages 245–248. IEEE, 2012.
- [81] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. PMLR, 2015.
- [82] Y. Song and P. Dhariwal. Improved techniques for training consistency models. In *12th International Conference on Learning Representations, ICLR*, 2024.
- [83] Y. Song, P. Dhariwal, M. Chen, and I. Sutskever. Consistency models. In *International Conference on Machine Learning, ICML*, volume 202 of *Proceedings of Machine Learning Research*, pages 32211–32252. PMLR, 2023.

- [84] Y. Song and S. Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.
- [85] Y. Song and S. Ermon. Improved techniques for training score-based generative models. *Advances in neural information processing systems*, 33:12438–12448, 2020.
- [86] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021.
- [87] F. Takens. Detecting strange attractors in turbulence. In *Dynamical Systems and Turbulence, Warwick 1980: proceedings of a symposium held at the University of Warwick 1979/80*, pages 366–381. Springer, 2006.
- [88] Y. Tashiro, J. Song, Y. Song, and S. Ermon. CSDI: Conditional score-based diffusion models for probabilistic time series imputation. *Advances in Neural Information Processing Systems*, 34:24804–24816, 2021.
- [89] E. Todorov, T. Erez, and Y. Tassa. MuJoCo: A physics engine for model-based control. In *2012 IEEE/RSJ international conference on intelligent robots and systems*, pages 5026–5033. IEEE, 2012.
- [90] A. Vahdat and J. Kautz. Nvae: A deep hierarchical variational autoencoder. *Advances in neural information processing systems*, 33:19667–19679, 2020.
- [91] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. W. Senior, and K. Kavukcuoglu. WaveNet: A generative model for raw audio. In *The 9th ISCA Speech Synthesis Workshop, SSW*, page 125. ISCA, 2016.
- [92] L. Van der Maaten and G. Hinton. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11), 2008.
- [93] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 2017.
- [94] M. Vetterli and C. Herley. Wavelets and filter banks: Theory and design. *IEEE transactions on signal processing*, 40(9):2207–2232, 1992.
- [95] Z. Wang and T. Oates. Imaging time-series to improve classification and imputation. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI*, pages 3939–3945. AAAI Press, 2015.
- [96] H. Wu, J. Xu, J. Wang, and M. Long. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. *Advances in neural information processing systems*, 34:22419–22430, 2021.
- [97] Q. Yan, Z. Liang, Y. Song, R. Liao, and L. Wang. SwinGNN: Rethinking permutation invariance in diffusion models for graph generation. *Trans. Mach. Learn. Res.*, 2024.
- [98] C. Yildiz, M. Heinonen, and H. Lahdesmaki. ODE2VAE: Deep generative second order ODEs with bayesian neural networks. *Advances in Neural Information Processing Systems*, 32, 2019.
- [99] J. Yoon, D. Jarrett, and M. Van der Schaar. Time-series generative adversarial networks. *Advances in neural information processing systems*, 32, 2019.
- [100] X. Yuan and Y. Qiao. Diffusion-TS: Interpretable diffusion for general time series generation. In *The Twelfth International Conference on Learning Representations, ICLR*, 2024.
- [101] A. Zeng, M. Chen, L. Zhang, and Q. Xu. Are transformers effective for time series forecasting? In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pages 11121–11128, 2023.
- [102] H. Zhou, S. Zhang, J. Peng, S. Zhang, J. Li, H. Xiong, and W. Zhang. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 11106–11115, 2021.
- [103] L. Zhou, M. Poli, W. Xu, S. Massaroli, and S. Ermon. Deep latent state space models for time-series generation. In *International Conference on Machine Learning*, pages 42625–42643. PMLR, 2023.
- [104] T. Zhou, Z. Ma, Q. Wen, X. Wang, L. Sun, and R. Jin. Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting. In *International conference on machine learning*, pages 27268–27286. PMLR, 2022.

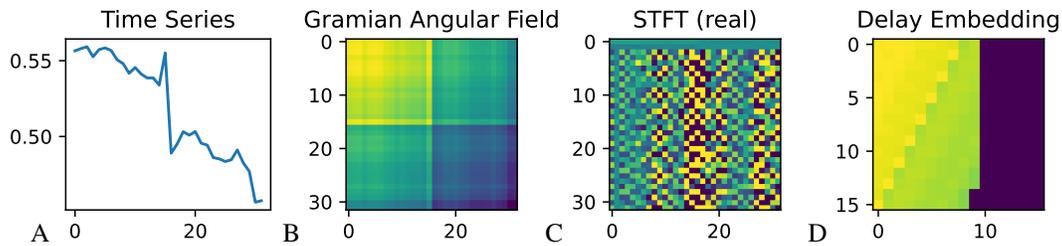


Figure 3: We plot above a time series signal (A), and its image transformations via the Gramian angular field (B), STFT (C), and the delay embedding (D).

A Domain Transformations

We give a detailed description for each domain transformation and its corresponding forward and inverse processes. We present in Fig. 3 a visual example of a time series and its various different image transformations.

Folding is a simple naïve transformation. Given a time series x , we fold it into an image x_{img} by starting from the first row on the left and continuing to the right, jumping to a new row whenever reaching the end of a row. Finally, if needed, we pad with zeroes at the end of the image. The inverse transformation back to the time series is simply taking the in-padded area of the image and unfolding it back to the time series. Although it is simple, this transformation can scale to very long sequences. Folding can be viewed as a particular example of delay embedding, as we detail below.

Gramian Angular Field is introduced in [95] for subsequent imputation and classification tasks. It depicts a time series within a polar coordinate system rather than the usual Cartesian coordinates. In the Gramian matrix, each element corresponds to the cosine of the sum of angles. The inverse action, from the image to the time series, is to take the main diagonal. While being a very informative transformation, a major constraint of this transformation is that the height and the width are linear with the size of the time series, preventing it from scaling to long sequences.

Delay Embedding [87] transforms a univariate time series $x_{1:L} \in \mathbb{R}^L$ into an image by organizing the series' information into columns and padding as necessary. The hyperparameters for this transformation are m and n , where m represents the skip value and n is the column dimension. For a given arbitrary channel of a time series, the transformation constructs the matrix X as follows:

$$X = \begin{bmatrix} x_1 & x_{m+1} & \dots & x_{L-n} \\ \vdots & \vdots & \dots & \vdots \\ x_n & x_{n+m+1} & \dots & x_L \end{bmatrix} \in \mathbb{R}^{n \times q},$$

where $q = \lceil (L - n)/m \rceil$. The image x_{img} is created by padding with zeros to meet the neural network input requirements. The example presented here is for a single channel; scaling to multiple dimensions is straightforward by concatenating each matrix X along another channel. Given an input signal $x \in \mathbb{R}^{L \times K}$, the transformation produces an output $x_{\text{img}} \in \mathbb{R}^{K \times n \times q}$. We pad the image with zeroes right after to create an $x_{\text{img}} \in \mathbb{R}^{K \times n \times n}$.

The original time series $x_{1:L}$ can be reconstructed from x_{img} by taking each marginal progression from the columns of the matrix X . There are various methods to reverse the transformation. For instance, if $m = 1$, $x_{1:L}$ is formed by concatenating the first row and the last column of x_{img} . The delay embedding naturally scales to long sequences; for example, setting $m = n = 256$ allows encoding $65k$ sequences with 256×256 images.

Short Time Fourier Transform (STFT) [35] is a widely used transformation that converts a signal from its original time domain into the frequency domain. The process of computing STFT involves dividing a time-domain signal into shorter, fixed-length segments and then applying the Fourier transform to each segment individually. Given an input signal $x \in \mathbb{R}^{L \times K}$, the STFT produces an output $x_{\text{img}} \in \mathbb{R}^{2K \times H \times W}$. In this output, the number of channels is doubled to store both the

real and imaginary parts of the transformed signal, and H and W are determined by user-defined parameters. These parameters include n_fft , which specifies the size of the Fourier transform, and hop_length , which defines the distance between successive sliding window frames. Unlike typical audio processing practices, we do not compute the magnitude spectrogram from the STFT output. Instead, we retain both the real and imaginary components within the image, thereby avoiding the need for additional complex spectrogram estimation. This approach maintains the integrity of the full spectral information. After obtaining the STFT images, we normalize them to the range $[-1, 1]$, ensuring that the data is scaled appropriately for subsequent processing.

B Experimental Setting

We use the same architectural backbone for all experiments: EDM [45]. We use the AdamW optimizer and train for 1000 epochs, although in practice, all models converged in the range 300-500 epochs. For each task, we elaborate on specific variations in settings and hyperparameters and provide additional information on the training and evaluation protocol.

B.1 Short-term unconditional generation.

Data. For the short-term unconditional generation (Sec. 5.1), we utilize four synthetic and real-world datasets with a fixed length of 24: *Stocks*, consisting of daily historical Google stock data from 2004 to 2019, comprising six channels: high, low, opening, closing, and adjusted closing prices, as well as volume. This data lacks periodicity and is dominated by random walks. The second dataset, *Energy*, is a multivariate appliance energy prediction dataset [14], featuring 28 channels with correlated features, and it exhibits noisy periodicity and continuous-valued measurements. The third dataset, *MuJoCo* (Multi-Joint dynamics with Contact), serves as a versatile physics generator for simulating TS data with 14 channels [89]. The last dataset, *Sine*, is a multivariate simulated dataset, where each sample $x_i^j(j)$ is defined as $\sin(2\pi\eta t + \theta)$, where η is sampled from a uniform distribution $[0, 1]$ and θ is sampled from a uniform distribution $[-\pi, \pi]$, with five channels for j .

Hyperparameters. We describe below in Tab. 6 the different hyperparameters used in our framework. For all datasets, we used the same default sampler of EDM [45], and we mention the hyperparameters of their U-net model that we tune in our work. Please see [45] for further details about the U-net model hyperparameters.

Table 6: Short-term unconditional generation hyperparameters including short time Fourier transform (STFT), delay embedding (DE) hyperparameters and diffusion hyperparameters

	Stocks	Energy	MuJoCo	Sine
General				
<i>image size</i>	8×8	8×8	8×8	8×8
<i>learning rate</i>	10^{-4}	10^{-4}	10^{-4}	10^{-4}
<i>batch size</i>	128	128	128	128
DE				
<i>embedding (n)</i>	8	8	8	8
<i>delay (m)</i>	3	3	3	3
STFT				
<i>n_fft</i>	-	-	-	-
<i>hop_length</i>	-	-	-	-
Diffusion				
<i>U-net channels</i>	128	128	64	128
<i>in channels</i>	[1, 2, 2, 2]	[1, 2, 2, 4]	[1, 2, 2, 2]	[1, 2, 2, 2]
<i>sampling steps</i>	18	18	18	18

Evaluation. We utilize the benchmark proposed in [99] to evaluate short-term unconditional generation and adhere to its evaluation protocol. This protocol comprises two scores: a predictive score and a discriminative score. The predictive score assesses the utility of the generated data by

training an independent prediction model on the generated data; superior generations result in better prediction scores for this model. The discriminative score evaluates the similarity of distributions using a proxy discriminator trained to distinguish between generated and original samples; higher scores indicate that the generative model has accurately captured the underlying distribution of the data. For more details about the evaluation protocol, please refer to [19] or [99].

B.2 Long-term unconditional generation.

Data. In our exploration of long-term unconditional generation, we employ the benchmark for long-term time series data as presented in [103]. This benchmark encompasses three extensive real-world time series datasets from the Monash Time Series Forecasting Repository [29]: FRED-MD, NN5 Daily, and Temperature Rain. These datasets were meticulously selected based on their average 1-lag autocorrelation metric, which quantifies the 1-step correlation over time. The 1-lag values, ranging from 0.38 to 0.98, exemplify a broad spectrum of temporal dynamics, thereby presenting significant challenges for generative learning models. To ensure uniformity in the NN5 Daily and FRED-MD datasets, each sequence within these datasets is normalized such that each trajectory is centered at its mean and adheres to a normal distribution. This normalization approach is advantageous for datasets like NN5 Daily, where the minimum and maximum values can vary substantially across different data points. For the Temperature Rain dataset, sequences are scaled to the $[0, 1]$ range, considering the data's consistently positive values and its tendency to cluster around the x-axis with occasional sharp spikes. Each dataset comprises approximately 750 time steps, providing a robust basis for evaluating long-term generative performance.

Hyperparameters. We describe below in Tab. 7 the different hyperparameters used in our framework. For all datasets, we used the same default sampler of EDM [45], and we mention the hyperparameters of the U-net model that we tune in our work, please see [45] for more details about these hyperparameters. In addition, For all long-term experiments, we use the AdamW optimizer with a weight decay of 10^{-5} .

Table 7: Long-term unconditional generation hyperparameters including short time Fourier transform (STFT), delay embedding (DE) hyperparameters and diffusion hyperparameters

	Fred-MD	Temperature Rain	NN5 Daily
General			
<i>image size</i>	32×32	32×32	32×32
<i>learning rate</i>	10^{-4}	10^{-4}	10^{-4}
<i>batch size</i>	32	64	32
DE			
<i>embedding(n)</i>	—	—	—
<i>delay(m)</i>	—	—	—
STFT			
<i>n_fft</i>	63	63	63
<i>hop_length</i>	23	23	25
Diffusion			
<i>U-net channels</i>	128	128	128
<i>in channels</i>	[1, 2, 4, 4]	[1, 2, 4, 4]	[1, 2, 4, 4]
<i>sampling steps</i>	18	18	18

Evaluation. To assess model performance, we follow the benchmark used in [103]. Our evaluation comprises classification and prediction models, each employing linear encoders and decoders with a single S4 layer having 16 hidden state dimensions. In the classification model, the encoder maps data dimensions to 16 hidden states. The S4 layer's output sequence is averaged before being passed to the decoder, which produces logits for binary classification using cross-entropy loss. Similarly, the prediction model's encoder maps input to a 16-dimensional hidden state, while the decoder maps it back to the original data dimension, predicting $k = 10$ future steps. Both models are trained using the AdamW optimizer, which has a learning rate of 0.01 over 100 epochs and a batch size of 128. The optimizer generates samples equal to testing data points to train the models together.

B.3 Ultra-long-term unconditional generation.

Data. For the ultra-long-term unconditional task, we introduce a novel benchmark consisting of two datasets: San Francisco Traffic (Traffic) [53] and KDD-Cup 2018 (KDD-Cup) [61]. The datasets' lengths are 17544 and 10920, respectively. Traffic includes an hourly time series detailing the road occupancy rates on the San Francisco Bay Area freeways from 2015 to 2016. KDD-Cup represents the air quality level from 2017 to 2018 estimated by 59 stations across two cities, Beijing (35 stations) and London (24 stations), measured in an hourly rate. We follow the same normalization procedure applied to Fred-MD and NN5 daily, as described in B.2.

Hyperparameters. In Tab. 8 below, we outline the various hyperparameters used in our framework. For all datasets, we employed the default sampler of EDM [45], and we specified the U-net model hyperparameters that we tuned in our study. For more information on the U-net model hyperparameters, please refer to [45].

Table 8: Ultra-long-term unconditional generation hyperparameters including short time Fourier transform (STFT), delay embedding (DE) hyperparameters and diffusion hyperparameters

	Traffic	KDD-Cup
General		
<i>image size</i>	144×144	112×112
<i>learning rate</i>	10^{-4}	10^{-4}
<i>batch size</i>	8	16
DE		
<i>embedding(n)</i>	144	–
<i>delay(m)</i>	136	–
STFT		
<i>n_fft</i>	–	223
<i>hop_length</i>	–	98
Diffusion		
<i>U-net channels</i>	128	128
<i>in channels</i>	[1, 2, 4, 4]	[1, 2, 4, 4]
<i>sampling steps</i>	18	18

Evaluation. For the ultra-long-term unconditional generation task, we follow the same procedure outlined in B.2. We use the same classification and prediction models, as they effectively distinguish between low- and high-quality ultra-long-term generations.

B.4 Conditional generation

Data. For short-term interpolation and extrapolation benchmarks (Sec. 5.4), we use the *ETT** datasets [102], each with a fixed length of 96. The *ETT* datasets are crucial indicators for long-term electric power deployment, containing two years of data from two separate counties in China. The datasets are divided into *ETTh1* and *ETTh2* for 1-hour intervals, and *ETTm1* and *ETTm2* for 15-minute intervals. Each data point includes the target value "oil temperature" and six power load features.

For long-term interpolation and extrapolation, we employ a well-established benchmark [75, 78, 103], incorporating the Physionet and USHCN datasets. The data extraction for the USHCN dataset follows the procedure detailed by [20]. Notably, both datasets exhibit sparsity across many features and contain numerous zero values. The Physionet dataset [80] includes health measurements from 41 sensors collected from 8,000 ICU patients within the first 48 hours of admission. The United States Historical Climatology Network (USHCN) [63] provides daily measurements from 1,218 weather stations across the United States, covering precipitation, snowfall, snow depth, and minimum and maximum temperatures. For the conditional tasks, we strictly follow the training and evaluation procedures described by [78], and we refer readers to this work for a comprehensive explanation of the evaluation protocol. Both datasets span approximately 1,000 to 2,000 time steps.

For the ultra-long-term conditional generation task, we utilize the same data used in the ultra-long-term unconditional generation benchmark, and we refer to App. B.3 for more details about the datasets.

Hyperparameters. We describe below in Tab. 9 the different hyperparameters used in our framework. The hyperparameters are similar for both tasks. Therefore, we present them in a unified table. For all datasets, we used the same default sampler of EDM[45]; we also present the hyperparameters of the U-net model; for further details about them, please see [45].

Table 9: Conditional generation hyperparameters including short time Fourier transform (STFT), delay embedding (DE) hyperparameters and diffusion hyperparameters

	ETT _{h1}	ETT _{h2}	ETT _{m1}	ETT _{m2}	Physionet	USHCN	Traffic	KDD-Cup
General								
<i>image size</i>	32 × 32	32 × 32	32 × 32	32 × 32	32 × 32	32 × 32	144	128
<i>learning rate</i>	10 ⁻⁴	10 ⁻⁴	10 ⁻⁵	30 ⁻⁵	10 ⁻⁵	10 ⁻⁴	10 ⁻⁴	10 ⁻⁴
<i>batch size</i>	32	32	16	64	8	8	8	8
DE								
<i>embedding(n)</i>	32	32	32	32	32	32	144	128
<i>delay(m)</i>	3	3	3	3	30	30	122	86
STFT								
<i>n_fft</i>	–	–	–	–	–	–	–	–
<i>hop_length</i>	–	–	–	–	–	–	–	–
Diffusion								
<i>U-net channels</i>	128	128	64	128	128	128	128	128
<i>in channels</i>	[1, 2, 2, 2]	[1, 2, 2, 4]	[1, 2, 4, 8]	[1, 2, 4, 8]	[1, 2, 4, 4]	[1, 2, 4, 4]	[1, 2, 4, 4]	[1, 2, 4, 4]
<i>sampling steps</i>	18	18	18	18	18	18	36	36

Evaluation. For the short-term and ultra-long-term datasets (ETT*, Traffic, KDD-Cup), we follow the next procedure. In the interpolation task, we randomly mask 50% of the input data and train the models to predict the missing masked 50%. In the extrapolation task, we mask the second half of the sequence and train the models to predict this missing half. We measure the distance between the models’ generated outcomes and the ground truth using MSE loss. Accurate generation will lead to a smaller distance between the prediction and the ground truth, thus indicating the models’ interpolation and extrapolation capabilities. For the long-term datasets, we follow a similar procedure as above; however, since the data is sparse and irregularly sampled, the masking is slightly different. We adhere to the exact interpolation and extrapolation processes described in [78] and refer to that source for more details.

C Additional Experiments and Analysis

C.1 Short-term unconditional generation

Due to space constraints in the main paper, we report the rest of the benchmark here. In Tab. 10, we show our model performance on the simple toy **Sine** dataset.

C.2 Wasserstein distance analysis

In Tab. 11, we present the Wasserstein distances calculated between our generated 2D point cloud and the actual data. A lower score indicates greater similarity between the clusters, meaning that a lower score is preferable. Our approach yields the best scores across all datasets in comparison to GT-GAN on short sequences and LS4 on long and ultra-long time series.

C.3 Short-term unconditional generation qualitative analysis

We include the *Stocks*, *Energy* and the *MuJoCo* qualitative t-SNE evaluation (Fig. 4(A,B,C)) and density analysis (Fig. 4(D, E, F)). In addition, we show in Tab. 11 a quantitative evaluation of the t-SNE clusters Wasserstein distance. Both the visual results and the quantitative results demonstrate our framework’s ability to learn the true distribution across multiple datasets.

Table 10: Short time series unconditional generation task on the Sine dataset.

Method	Sine	
	disc↓	pred↓
KoVAE	.005 ± .003	.093 ± .000
DiffTime	.013 ± .006	.093 ± .000
GT-GAN	.012 ± .014	.097 ± .000
TimeGAN	.011 ± .008	.093 ± .019
RCCGAN	.022 ± .0068	.097 ± .001
C-RNN-GAN	.229 ± .040	.127 ± .004
WaveNet	.158 ± .011	.117 ± .008
WaveGAN	.277 ± .013	.134 ± .013
LS4	.342 ± .007	.132 ± .011
Ours	.014 ± .009	.094 ± .000

Table 11: We calculate the Wasserstein distances between the original cluster, our generated samples cluster the other method cluster shown in Figs. 4, 5 and 6.

Method	Stocks	Energy	MuJoCo	Temp Rain	NN5 Daily	Traffic	KDD-Cup
GT-GAN	6.20	3.35	3.83	-	-	-	-
LS4	-	-	-	3.27	5.23	4.63	11.84
Ours	2.84	3.20	1.05	2.85	5.21	3.19	6.59

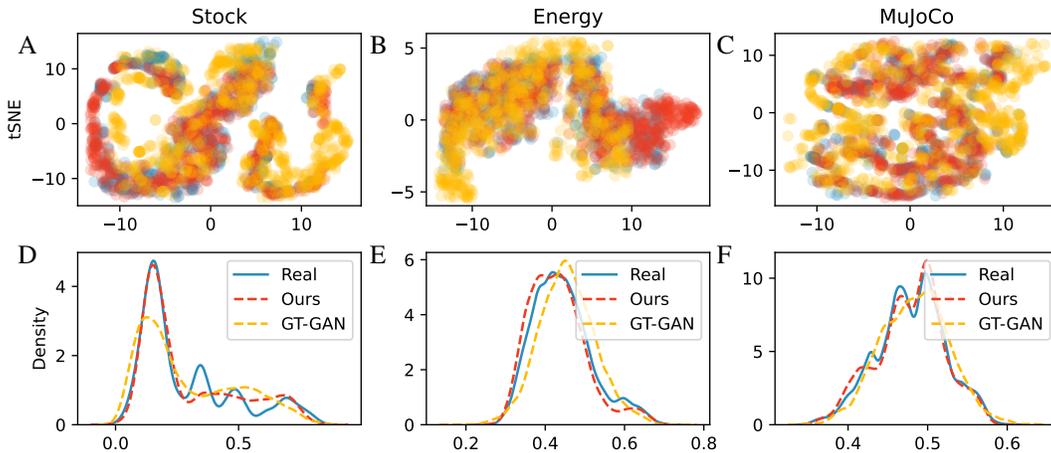


Figure 4: We plot the 2D t-SNE embeddings of synthetic data generated with our method and SOTA tools vs. the real data (top). Then, we compare their probability density functions (bottom).

C.4 Long-term unconditional generation with standard deviation

In Tab. 12, we present the results of our method on unconditional generation of long sequences with standard deviation. The results demonstrate our method’s statistical significance compared to the state-of-the-art method LS4.

Table 12: Long time series unconditional generation task with standard deviation.

	LS4			Ours		
	marg↓	class ↑	pred ↓	marg↓	class ↑	pred ↓
FRED-MD	.022	.544	.037	.021 ± .000	.862 ± .227	.009 ± .003
NN5 Daily	.007	.636	.241	.005 ± .000	.822 ± .157	.307 ± .037
Temp Rain	.083	.976	.521	.409 ± .000	5.80 ± .974	.377 ± .022

C.5 Long-term unconditional generation qualitative analysis

We include the qualitative t-SNE evaluation for *Temp Rain* and *NN5 Daily* (Fig. 5(A, B)) and their density analysis (Fig. 5(D, E)). Additionally, we provide in Tab. 11 the quantitative evaluation of the t-SNE clusters using the Wasserstein distance. Our results indicate the superiority of our approach in comparison to other techniques.

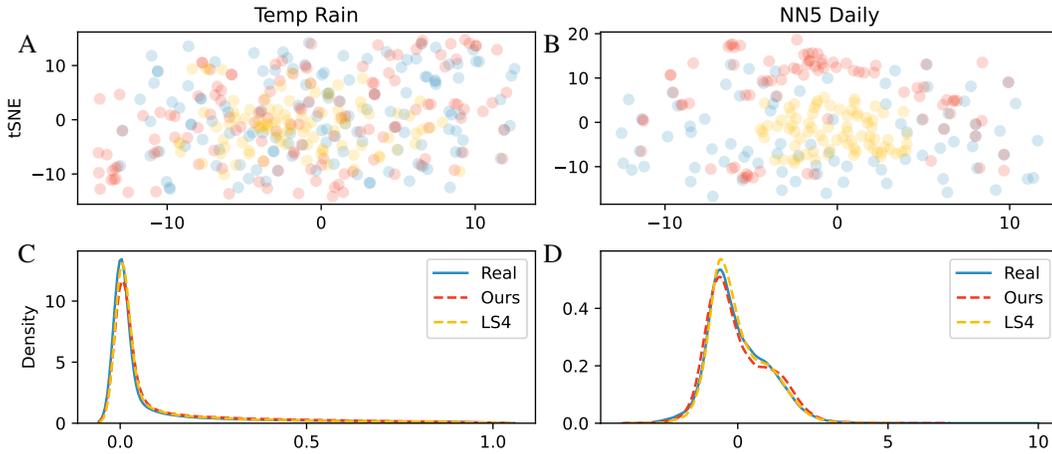


Figure 5: We plot the 2D t-SNE embeddings of synthetic data generated with our method and SOTA tools vs. the real data (top). Then, we compare their probability density functions (bottom).

C.6 Ultra-long time series unconditional generation with standard deviation

In Tab. 13, we present the results of our method for unconditional generation of ultra-long sequences, including standard deviations. These results demonstrate the statistical significance of our method compared to the state-of-the-art competitive methods.

Table 13: Ultra-long unconditional generation with standard deviation

Method	Traffic			KDD-Cup		
	pred ↓	class ↑	marg ↓	pred ↓	class ↑	marg ↓
Latent ODE	1.01 ± .412	.000 ± .000	.180 ± .000	.079 ± .055	.013 ± .020	.009 ± .000
LS4	.170 ± .030	.630 ± .060	.002 ± .000	.049 ± .046	.488 ± .164	.002 ± .000
Ours	.138 ± .014	.684 ± .019	.001 ± .000	.001 ± .000	.842 ± .245	.001 ± .000

C.7 Ultra-long-term unconditional generation qualitative analysis

We include the qualitative t-SNE evaluation for *Traffic* and *KDD-Cup* (Fig. 6(A, B)) and the density analysis (Fig. 6(D, E)). Additionally, we provide in Tab. 11 the quantitative evaluation of the t-SNE clusters using the Wasserstein distance. Our results highlight our method ability to handle very long sequences.

C.8 Image size ablation

Given the significant impact of image size on the computational resources required by our method, we investigate whether varying image size influences different transformations. We explore the effect of different image sizes on our framework, utilizing long-term datasets with short time Fourier transform (STFT) and short-term datasets with delay embedding. For short time series consisting of 24 time steps, we test image sizes of 8 and 16, as we observe that scaling to larger sizes may not yield benefits. For long time series of approximately 750 steps, we experiment with sizes of 32, 64, and 128. Notably, 32 is the minimum size required to contain enough pixels for representing the long sequence adequately. We present the results in Tab. 14. While most results demonstrate high

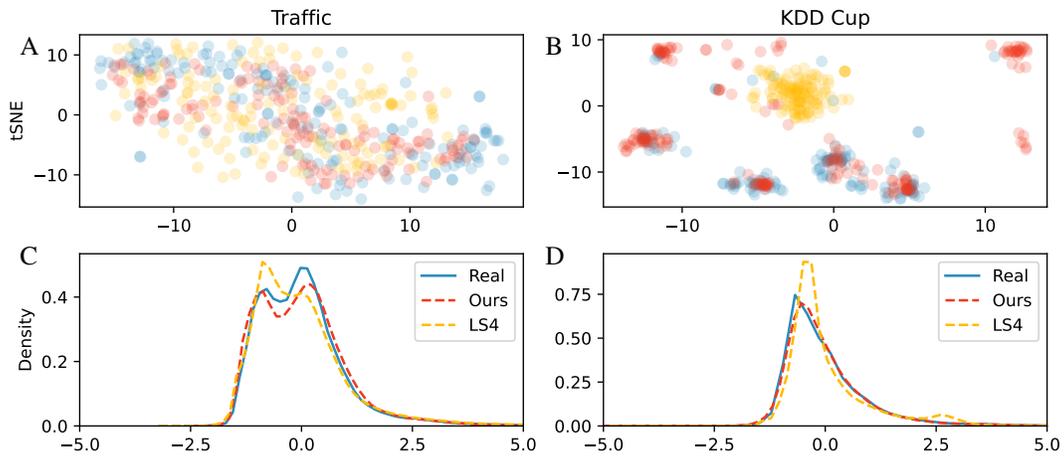


Figure 6: We plot the 2D t-SNE embeddings of synthetic data generated with our method and SOTA tools vs. the real data (top). Then, we compare their probability density functions (bottom).

competitiveness compared to other methods, inconclusive experimental results indicate that image size does not significantly affect generation quality. Therefore, it is reasonable to use the minimum length that a transformation can accommodate to benefit from minimum computational costs.

Table 14: Image size ablation study.

Image Size	FRED-MD			NN5 Daily			Energy		MuJoCo	
	marg↓	class ↑	pred ↓	marg↓	class ↑	pred ↓	disc↓	pred ↓	disc↓	pred ↓
Short Series										
8 x 8	-	-	-	-	-	-	.040	.250	.007	.033
16 x 16	-	-	-	-	-	-	.059	.250	.036	.032
Long Series										
32 x 32	.021	.862	.009	.005	.822	.307	-	-	-	-
64 x 64	.021	1.82	.021	.008	.829	.440	-	-	-	-
128 x 128	.016	1.65	.023	.012	.733	.430	-	-	-	-

C.9 Hyperparameters Ablation

Ablation study on diffusion sampling steps. While [45] demonstrate an improvement in FID score with a larger number of steps in their work, we do not observe the same trend in our framework. The results are presented in Tab. 15 in the first section, indicating an unclear trend across different datasets and metrics.

Ablation study on batch size. The results in Tab. 15 in the middle section demonstrate that our framework is unaffected by different batch sizes. This is a positive indication of our framework’s adaptability to various computational environments, whether with low memory or high memory capabilities.

Ablation study on learning rate. In our examination of the learning rate, we have made an intriguing observation. We have found that when the learning rate equals or exceeds 10^{-3} , the diffusion backbone [45] tends to collapse, resulting in the generation of irrelevant signals. This phenomenon is clearly demonstrated in Tab. 15 in the middle section. With a learning rate of 10^{-3} , the *disc* scores are .256 and 0.499 for the MuJoCo and Energy datasets, respectively, indicating random generation. However, when the learning rate is lowered, we have not observed any such collapse of the backbone on any dataset or task.

Table 15: Hyperparameters ablation study. We study the effect of different hyperparameters on our framework. We utilize four short and long-term datasets.

Hyperparameter	FRED-MD			NN5 Daily			Energy		MuJoCo	
	marg↓	class ↑	pred ↓	marg↓	class ↑	pred ↓	disc↓	pred ↑	disc↓	pred ↑
Diffusion Sampling Steps										
18	.021	.862	.009	.005	.822	.307	.040	.250	.007	.033
36	.015	1.33	.020	.009	.829	.399	.052	.250	.017	.032
72	.017	1.36	.022	.009	.836	.395	.057	.250	.025	.031
144	.018	1.32	.024	.010	.836	.397	.058	.250	.025	.030
Batch Size										
16	.022	1.41	.023	.009	.804	.403	.060	.250	.012	.032
32	.018	1.31	.021	.010	.850	.396	.059	.250	.009	.032
64	.019	1.32	.021	.009	.842	.394	.050	.250	.019	.032
Learning Rate										
10^{-3}	.019	1.02	.025	.012	.827	.415	.499	.252	.256	.043
10^{-4}	.021	1.54	.024	.010	.813	.401	.065	.249	.007	.033
10^{-5}	.021	1.16	.025	.007	1.01	.421	.056	.250	.020	.031

C.10 Computational Resources Comparison

In this section, we compare the computational resources required by our proposed method and the LS4 method, focusing on training and inference wall-clock runtime and model size in terms of parameters, and we analyze the FLOPs used per method. Although our method, which utilizes image transforms and diffusion models, has a larger model size in terms of parameters, it remains comparable to LS4 regarding training and inference time. Despite the larger model size, our method achieves similar training and inference efficiency, making it a viable and scalable solution for large-scale time series generation tasks as shown in Tab.16. Furthermore, the rapid advancements and growing research interest in faster sampling techniques for diffusion models [82] present an opportunity to further enhance our method’s efficiency. Leveraging these developments, our approach can integrate even more optimized diffusion models, potentially reducing the computational time and resources required for training and inference, thus improving scalability for large-scale time series generation tasks. Finally, we analyze the FLOPs used in our method compared to LS4 and DiffTime on the Stock, nn5daily and KDD Cup datasets in Tab. 17.

Table 16: Computational resources in terms of training wall-clock runtime (WCR) in minutes(m) or hours(h), and model parameters (MP) in millions (M)

Method	Stocks		Energy		NN5 Daily		Temp Rain	
	WCR	MP	WCR	MP	WCR	MP	WCR	MP
TimeGAN	2h 59m	48K	3h 37m	1M	-	-	-	-
GT-GAN	12h 20m	41K	10h 39m	57k	-	-	-	-
DiffTime	52m	240k	-	-	46m	32M	-	-
LS4	5h 30m	2.7M	2h	2.1M	53m	2.1M	27h	2.3M
Ours	1h 10m	575K	1h	2M	58m	5.9M	30h	6.4M

C.11 Scaling Laws Analysis

In this section, we investigate how the performance of our proposed method scales with the size of the underlying image diffusion model. Specifically, we evaluate the impact of increasing the model size from a few thousand parameters to several hundred million on various time-series datasets. Additionally, we compare this trend with other state-of-the-art time-series generative models, analyzing how their performance is affected by model size increments. Interestingly, merely increasing the model parameters does not improve their performance. It demonstrates that simply enlarging previous methods does not necessarily enhance their generation capabilities. Moreover, in the case of LS4 on the KDD Cup dataset, increasing the model’s parameters to 100 million results

Table 17: FLOPs analysis on DiffTime, LS4 and Our method

#Params	DiffTime			LS4			Ours		
	Stocks	nn5daily	KDD Cup	Stocks	nn5daily	KDD Cup	Stocks	nn5daily	KDD Cup
500k	0.057G	0.751G	14.485G	0.003G	0.107G	1.480G	0.009G	0.123G	1.480G
1M	0.037G	1.301G	24.394G	0.007G	0.233G	3.225G	0.014G	0.217G	2.624G
5M	0.747G	4.289G	99.402G	0.048G	1.594G	22.015G	0.070G	1.084G	13.21G
25M	3.838G	20.864G	414.15G	0.191G	6.299G	86.966G	0.291G	4.683G	57.25G
50M	7.996G	41.060G	711.04G	0.352G	11.602G	160.172G	0.633G	10.06G	123.0G
100M	14.544G	80.198G	1387.3G	0.759G	25.040G	—	1.335G	21.27G	260.3G
150M	22.407G	118.865G	1874.5G	1.084G	35.739G	—	1.890G	31.58G	386.6G

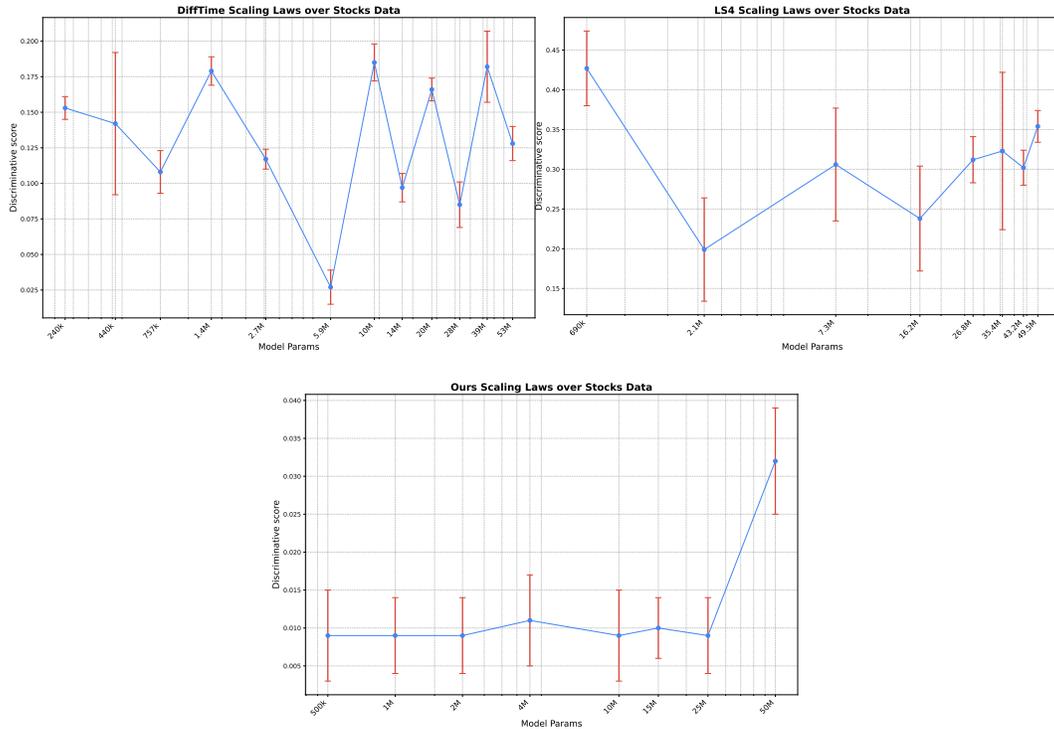


Figure 7: Scaling analysis of different models on Stocks data. A lower discriminative score is better.

in memory collapse, making it infeasible to run a batch size of one with the current resources used for training all models. We present the results for our method in Tab. 18, for LS4 in Tab. 19 and the results for DiffTime, in Tab. 20 and Tab. 21. Finally, we add a visualization of the results for the Stocks dataset in Fig. 7.

Table 18: Our framework scaling laws

#P	Stocks		Energy		marg↓	nn5daily	
	disc↓	pred↓	disc↓	pred↓		class↑	pred↓
0.5M	.009 ± .006	.036 ± .000	.100 ± .009	.253 ± .000	.012	.765 ± .275	.407 ± .640
1M	.009 ± .005	.036 ± .000	.087 ± .005	.251 ± .000	.010	1.10 ± .315	.421 ± .076
2M	.009 ± .005	.036 ± .000	.04 ± .006	.250 ± .000	.006	1.03 ± .263	.389 ± .056
4M	.011 ± .006	.036 ± .000	.053 ± .004	.250 ± .000	.005	.859 ± .368	.443 ± .117
10M	.009 ± .006	.036 ± .000	-	-	.006	.859 ± .153	.428 ± .032
15M	.010 ± .004	.036 ± .000	-	-	-	-	-
25M	.009 ± .005	.036 ± .000	-	-	-	-	-
50M	.032 ± .007	.036 ± .000	-	-	.005	.420 ± 0.021	.884 ± 0.223

Table 19: LS4 scaling laws

#P	Energy		Stocks		marg↓	nn5daily	
	disc↓	pred↓	disc↓	pred↓		class↑	pred↓
0.69M	.498 ± .000	.374 ± .019	.427 ± .047	.047 ± .0045	.061	.002 ± .003	1.97 ± 1.17
2.1M	.474 ± .003	.251 ± .000	.199 ± .065	.068 ± .013	.011	.719 ± .138	.305 ± .054
7.3M	.488 ± .002	.311 ± .003	.306 ± .071	.058 ± .001	.067	.005 ± .009	1.33 ± .288
16.2M	.489 ± .003	.262 ± .001	.238 ± .066	.038 ± .000	.011	.733 ± .344	.335 ± .105
26.8M	.494 ± .002	.311 ± .003	.312 ± .029	.054 ± .001	.019	.213 ± .167	.275 ± .034
35.4M	-	-	.323 ± .099	.071 ± .003	.008	.771 ± .225	.294 ± .069
43.2M	-	-	.302 ± .022	.067 ± .000	-	-	-
49.5M	.497 ± .001	.294 ± .002	.354 ± .020	.037 ± .000	-	-	-
102.4M	.497 ± .001	.276 ± .001	-	-	.090	.006 ± .000	7.07 ± 2.21
152.7M	.498 ± .000	.296 ± .005	-	-	.165	.001 ± .000	2.02 ± .897

Table 20: DiffTime scaling laws

#P	Stocks		marg↓	nn5 daily	
	disc↓	pred↓		class↑	pred↓
240k	.153 ± .008	.037 ± .037	.030	.179 ± .131	4.33 ± 1.54
440k	.142 ± .050	.038 ± .038	.026	.284 ± .252	2.79 ± .591
757k	.108 ± .015	.038 ± .038	.015	.421 ± .316	1.92 ± .941
1.4M	.179 ± .010	.037 ± .037	.021	.525 ± .410	11.43 ± 19.03
2.7M	.117 ± .007	.039 ± .039	.016	.279 ± .217	.981 ± .207
5.9M	.027 ± .012	.037 ± .037	.006	.324 ± .160	.666 ± .199
10M	.185 ± .013	.037 ± .037	.016	.272 ± .135	.939 ± .224
14M	.097 ± .010	.037 ± .037	.011	.308 ± .199	10.53 ± 7.06
20M	.166 ± .008	.037 ± .037	.017	.312 ± .147	.441 ± .084
28M	.085 ± .016	.039 ± .039	.019	.134 ± .097	.583 ± .086
39M	.182 ± .025	.042 ± .042	.021	.182 ± .155	.681 ± .187
53M	.128 ± .012	.037 ± .037	.008	.140 ± .160	.896 ± .598
103M	-	-	.015	.345 ± .132	.584 ± .211
156M	-	-	.018	.675 ± .744	1.700 ± 1.161

C.12 Other Image Generative Models

Our goal in this paper is to leverage recent advancements in computer vision to develop an elegant and robust solution for time-series data, addressing different sequence lengths and setting a baseline for handling short, long, and ultra-long sequences. We aim to take advantage of the fact that image architectures are more thoroughly explored. We hypothesize that the improvements we observe are largely due to the more advanced development of image architectures compared to time-series architectures. To further investigate this, we used NVAE [90], and StyleGAN [46], instead of the diffusion model. We observe the results in Tab. 22. The results imply that using recently better-explored architecture yields better results when using the same transformations. This understanding strengthens our hypothesis for the robustness and efficiency of diffusion models.

Table 21: DiffTime scaling laws, KDD Cup

#P	KDD Cup		
	marg↓	class↑	pred↓
3.4M	.026	.001 ± .000	1103.27 ± 753.21
13.5M	.022	.001 ± .000	902.77 ± 632.74
54M	.024	.000 ± .000	930.83 ± 607.39
104M	.024	.001 ± .000	970.21 ± 693.19
121M	.024	.000 ± .000	968.18 ± 611.15

Table 22: Other image generative models results

Models/Datasets	KDD			NN Daily			Stocks	
	Marginal ↓	Classifier ↑	Predictor ↓	Marginal ↓	Classifier ↑	Predictor ↓	Disc ↓	Pred ↓
Style GAN	0.020	0.001	0.233	0.020	0.091	2.100	0.276	0.042
NVAE	0.008	0.031	0.107	0.020	0.089	0.600	0.081	0.049

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The experiment section, related work and method section support the main claims.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: The conclusion section discusses the shortcomings of our framework and future improvements.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: All models and hyperparameters are extensively reported in the appendix. In addition, the code will be publicly available at the end of the double-blind process.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [NA]

Justification: All datasets are public, and the code will be publicly available at the end of the double-blind process.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Experimental Setting section in the appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: On the main table and if note, reported in the appendix due to space constraints and convince reasons.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Appendix computational resources analysis

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification:

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: Appendix

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.