
TIME SERIES DIFFUSION IN THE FREQUENCY DOMAIN

Jonathan Crabbé*, Nicolas Huynh*, Jan Stanczuk, Mihaela van der Schaar

DAMTP

University of Cambridge

{jc2133, nvth2, js2164, mv472}@cam.ac.uk

ABSTRACT

Fourier analysis has been an instrumental tool in the development of signal processing. This leads us to wonder whether this framework could similarly benefit generative modelling. In this paper, we explore this question through the scope of time series diffusion models. More specifically, we analyze whether representing time series in the frequency domain is a useful inductive bias for score-based diffusion models. By starting from the canonical SDE formulation of diffusion in the time domain, we show that a dual diffusion process occurs in the frequency domain with an important nuance: Brownian motions are replaced by what we call mirrored Brownian motions, characterized by mirror symmetries among their components. Building on this insight, we show how to adapt the denoising score matching approach to implement diffusion models in the frequency domain. This results in frequency diffusion models, which we compare to canonical time diffusion models. Our empirical evaluation on real-world datasets, covering various domains like healthcare and finance, shows that frequency diffusion models better capture the training distribution than time diffusion models. We explain this observation by showing that time series from these datasets tend to be more localized in the frequency domain than in the time domain, which makes them easier to model in the former case. All our observations point towards impactful synergies between Fourier analysis and diffusion models.

1 Introduction

Deep generative modelling leverages the inductive bias of neural networks to learn complex, high-dimensional probability distributions from real-world datasets. Among other applications, generative models allow for generation of new synthetic samples consistent with the distribution of the training data, yet distinct from the actual data encountered during training. Recently this field has seen tremendous progress in various modalities including image [1, 2], audio [3, 4], video [5] and text [6] generation, as well as addressing inverse problems such as in-painting [7] or super-resolution [8]. Moreover deep generative models have started showing significant potential in contributing to natural sciences, though protein design [9], drug development [10] and material synthesis [11]. However, the application of these models to time series data has not seen the same level of advancement [12]. Some notable examples of time series generative models include TimeGAN [13], FourierFlow [14], and RCGAN [15], yet this area remains less explored compared to other applications.

Diffusion Models. In recent years, diffusion models [16, 17, 18, 19] have emerged as one of the most promising research avenues in deep generative modelling, achieving state-of-the-art results across many generative modelling tasks [2, 8]. Diffusion models have been applied to time series modelling, achieving promising results [20]. However, there is substantial room for development and refinement in these early-stage applications.

Fourier analysis. Fourier analysis is a remarkably powerful tool in signal processing, compression and machine learning [21]. It has been shown to significantly improve state-of-the-art the performance of many deep learning based time series analysis techniques [22], with some recent applications in dataset distillation [23]. In the context of deep generative models, this is exemplified in [14], where the application of normalizing flows to Fourier representations

*Denotes equal contribution and corresponding authors.

yielded promising results. More recently, some work by [24] has been done on diffusion on functional spaces, which include Fourier representations of signals although the paper does not specifically focus on the Fourier basis.

Motivation. Despite Fourier analysis’ widespread success, its application to diffusion models for time series remains largely unexplored. This paper seeks to fill this research gap, by examining whether spectral representations can improve diffusion models for time series modelling. Our focus is not on achieving state-of-the-art results, but rather investigating whether representing time series in the frequency domain is a useful inductive bias for diffusion models.

Our contributions. (1) **Formalizing frequency diffusion.** In Section 3, we show theoretically how to translate SDE-based diffusion of time series to the frequency domain. We demonstrate that the denoising score matching recipe can be adapted by replacing standard Brownian motions by what we call mirrored Brownian motions, characterized by mirror symmetries in their components. (2) **Comparing time and frequency diffusion.** In Section 4.1, we compare the ability of the time and frequency score models to generate samples that are faithful to the training sets by leveraging *sliced Wasserstein distances*. Through an extensive analysis on 6 real-world datasets illustrating fields like healthcare, finance, engineering and climate modelling, we demonstrate that frequency score models consistently outperform the time score models. (3) **Understanding why and when frequency diffusion is preferable.** In Section 4.2, we demonstrate that the signals in all 6 datasets concentrate most of their power spectrum on the low frequencies. We hypothesize that this localization in the frequency domain explains the superior performances of frequency diffusion models. In Section 4.3, we confirm this hypothesis by artificially delocalizing the spectral representation of real signals and showing that the gap between time and frequency diffusion closes.

2 Background

Notations. We consider multivariate time series of fixed size² $\mathbf{x} \in \mathbb{R}^{N \times M}$, where $N \in \mathbb{N}$ is the number of time steps and $M \in \mathbb{N}$ is the number of features tracked over time. Often, we will denote by $d_X = N \cdot M$ the total dimension of the time series \mathbf{x} . We shall use *Greek letters* for components of the time series. In this way, $\mathbf{x}_\tau \in \mathbb{R}^M$ denotes the feature vector at time $\tau \in [N]$ and $x_{\tau,\nu}$ denotes the value of feature $\nu \in [M]$ at time τ . We denote by $[K] := \{0, 1, \dots, K-1\}$ the integers between 0 (included) and $K \in \mathbb{N}$ (excluded). To avoid any confusion between time series steps and diffusion steps, we shall use *Latin letters* for the diffusion process. In this way, the diffusion process is described by a family of time series $\{\mathbf{x}(t) \in \mathbb{R}^{d_X}\}_{t=0}^T$ indexed by a continuous diffusion variable $t \in [0, T]$. Thanks to these notations, we unambiguously interpret $\mathbf{x}_\tau(t) \in \mathbb{R}^M$ as the feature vector at time step $\tau \in [N]$ and at diffusion step $t \in [0, T]$. We shall detail below how this diffusion process is defined.

2.1 Score-based generative modeling with SDEs

In continuous-time diffusion modelling, one assumes access to samples drawn from an unknown density p_{data} . The objective of generative modelling is to obtain a tractable approximation of this distribution.

Forward diffusion. Score-based generative modeling with *stochastic differential equation* (SDEs) [19] typically operates by first constructing a forward diffusion process. In the case of time series, forward continuous diffusion is described by the following SDE, with $t \in [0, T]$:

$$d\mathbf{x} = \mathbf{f}(\mathbf{x}, t)dt + \mathbf{G}(t)d\mathbf{w}, \quad (1)$$

where $\mathbf{f} : \mathbb{R}^{d_X} \times [0, T] \rightarrow \mathbb{R}^{d_X}$ is the drift, \mathbf{w} is a standard Brownian motion in \mathbb{R}^{d_X} , and $\mathbf{G} : [0, T] \rightarrow \mathbb{R}^{N \times N}$ is the diffusion matrix. We denote p_t the probability density of the solution $\mathbf{x}(t)$ of Equation (1) at time $t \in [0, T]$. With the slight abuse of notation from [19], we shall abbreviate $p_t(\mathbf{x}(t))$ by $p_t(\mathbf{x})$. Together with the SDE, we impose the initial condition $p_0 = p_{\text{data}}$, which corresponds to samples initially drawn from the data density p_{data} . In practice, we consider \mathbf{f} and \mathbf{G} such that p_{data} is transported to a final density p_T close to an isotropic Gaussian.

Reverse diffusion. The reverse diffusion process performs the inverse transformation by transporting the isotropic Gaussian density p_T to the data density $p_0 = p_{\text{data}}$. Hence, applying reverse diffusion to samples drawn from the isotropic Gaussian permits to sample from the unknown density p_{data} . It was shown by [25] that this reverse diffusion satisfies the following SDE:

$$d\mathbf{x} = \mathbf{b}(\mathbf{x}, t)dt + \mathbf{G}(t)d\hat{\mathbf{w}}, \quad (2)$$

where $\mathbf{b}(\mathbf{x}, t) = \mathbf{f}(\mathbf{x}, t) - \mathbf{G}(t)\mathbf{G}(t)^T \nabla_{\mathbf{x}} \log p_t(\mathbf{x})$, dt is a negative infinitesimal time step, and $\hat{\mathbf{w}}$ is a Brownian time increment with time going backwards from T to 0.

²Padding over time can be used in cases where the datasets contain time series of different lengths.

Denosing score matching. In order to run the reverse diffusion process, one needs access to the score $\mathbf{s}(\mathbf{x}, t) := \nabla_{\mathbf{x}} \log p_t(\mathbf{x})$. In practice, the ground-truth density p_t is unknown. Denosing score matching circumvents this problem by estimating the ground-truth score with a function \mathbf{s}_{θ^*} whose parameters θ^* minimize the following score matching objective computed from the data samples [16, 26]:

$$\theta^* = \arg \min_{\theta \in \Theta} \mathbb{E}_{t, \mathbf{x}(0), \mathbf{x}(t)} [\mathcal{L}_{\text{SM}}(\mathbf{s}_{\theta}, \mathbf{s}_{t|0}, \mathbf{x}, t)] \quad (3)$$

$$\mathcal{L}_{\text{SM}}(\mathbf{s}_{\theta}, \mathbf{s}_{t|0}, \mathbf{x}, t) := \|\mathbf{s}_{\theta}(\mathbf{x}, t) - \mathbf{s}_{t|0}(\mathbf{x}, t)\|^2 \quad (4)$$

where $\|\cdot\|$ denotes the Frobenius norm, $t \sim \mathcal{U}(0, T)$, $\mathbf{x}(0) \sim p_0(\mathbf{x})$, $\mathbf{x}(t) \sim p_{t|0}(\mathbf{x}(t)|\mathbf{x}(0))$ with $p_{t|0}$ denoting the transition kernel from 0 to t , and $\mathbf{s}_{t|0}(\mathbf{x}, t) := \nabla_{\mathbf{x}(t)} \log p_{t|0}(\mathbf{x}(t)|\mathbf{x}(0))$. With sufficient model capacity, the parameters θ^* provide an approximation \mathbf{s}_{θ^*} that is equal to the score \mathbf{s} for almost all \mathbf{x} and t in the large data limit [27]. Equipped with an approximation of the score $\mathbf{s}_{\theta^*} \approx \mathbf{s}$, one can generate data by sampling according to the solution defined by the reverse diffusion process from Equation (2).

2.2 Discrete Fourier Transform

DFT. By considering a time series $\mathbf{x} = (\mathbf{x}_0, \dots, \mathbf{x}_{N-1}) \in \mathbb{R}^{dx}$, the *Discrete Fourier Transform* (DFT), denoted as $\tilde{\mathbf{x}} = \mathcal{F}[\mathbf{x}]$, is defined as

$$\tilde{\mathbf{x}}_{\kappa} := \frac{1}{\sqrt{N}} \sum_{\tau=0}^{N-1} \mathbf{x}_{\tau} \exp\left(-\frac{\kappa 2\pi i}{N} \tau\right) \quad (5)$$

for all $\kappa \in [N]$. In the signal processing literature, each κ corresponds to a harmonic of frequency $\omega_{\kappa} := \frac{\kappa 2\pi}{N}$. For this reason, the DFT $\tilde{\mathbf{x}}$ is said to represent the time series \mathbf{x} in the *frequency domain*, as opposed to the *time domain*. We also note that the DFT is complex-valued ($\tilde{\mathbf{x}} \in \mathbb{C}^{dx}$).

Matrix representation. We note that the DFT operator \mathcal{F} is *linear* with respect to the time components $(\mathbf{x}_0, \dots, \mathbf{x}_{N-1})$. It can therefore be expressed through a left matrix multiplication $\tilde{\mathbf{x}} = \mathcal{F}[\mathbf{x}] = U\mathbf{x}$, where $U \in \mathbb{C}^{N \times N}$ is defined as $[U]_{\kappa\tau} := N^{-1/2} \exp(-i\omega_{\kappa}\tau)$. It can easily be checked (see Appendix A.1) that the matrix U is unitary: $U^*U = UU^* = I_N$, where U^* is the conjugate transpose of U and I_N is the $N \times N$ identity matrix. This implies that the DFT operator is invertible and that the original time series can be reconstructed from its representation in the frequency domain: $\mathbf{x} = \mathcal{F}^{-1}[\tilde{\mathbf{x}}] := U^*\tilde{\mathbf{x}}$.

DFT of a real-valued sequence. While the DFT $\tilde{\mathbf{x}}$ is defined in \mathbb{C}^{dx} , some of its components are made redundant by the fact that \mathbf{x} is a real-valued time-series. One can easily check (see Appendix A.1) that this constraint imposes the following *mirror symmetry* on the DFT for all $\kappa \in [N]$:

$$\tilde{\mathbf{x}}_{\kappa} = \tilde{\mathbf{x}}_{N-\kappa}^* \quad (6)$$

where z^* denotes the complex conjugate of $z \in \mathbb{C}$ and we define $\tilde{\mathbf{x}}_N := \tilde{\mathbf{x}}_0$ for consistency. Through this symmetry, we observe that the components $\tilde{\mathbf{x}}_{\kappa}$ with $\kappa \leq \lfloor N/2 \rfloor$ uniquely define the DFT of a real-valued time series. For this reason, the frequencies beyond the Nyquist frequency $\omega_{\text{Nyq}} := \omega_{\lfloor N/2 \rfloor}$ are redundant with respect to the lower frequencies. In the frequency domain, one then needs only to diffuse N real numbers extracted from the DFT and the rest of $\tilde{\mathbf{x}}$ can be deduced from Equation (6).

Signal Energy. An important quantity related to a time series \mathbf{x} is its *total energy*, which simply corresponds to the squared Frobenius norm $\|\mathbf{x}\|^2 := \sum_{\tau=0}^{N-1} \sum_{\nu=1}^M |x_{\tau, \nu}|^2$, where $|\cdot|$ denotes the modulus of a complex number. Through *Parseval's theorem*, this energy can be evaluated by computing the same norm for the DFT $\tilde{\mathbf{x}}$ of \mathbf{x} : $\|\mathbf{x}\|^2 = \|\tilde{\mathbf{x}}\|^2$. We note that the total energy is obtained by summing over all time steps or frequencies. To characterize how the energy is distributed over the time steps $\tau \in [N]$, we use the *energy density* defined as the squared Euclidean norm $\|\mathbf{x}_{\tau}\|_2^2 := \sum_{\nu=1}^M |x_{\tau, \nu}|^2$. Similarly, the *spectral energy density* defined as $\|\tilde{\mathbf{x}}_{\kappa}\|_2^2$ describes how the signal energy is distributed across the frequencies $\kappa \in [N]$.

Probability density in the complex space. Adapting the diffusion formalism to the frequency domain requires to define a probability density for the complex-valued random variable $\tilde{\mathbf{x}} \in \mathbb{C}^{dx}$. By following [28], this density is written in terms of the real and imaginary parts of the signal $\tilde{p}(\tilde{\mathbf{x}}) := \tilde{p}(\Re[\tilde{\mathbf{x}}], \Im[\tilde{\mathbf{x}}])$. Similarly, the score function follows a similar decomposition in terms of the signal real and imaginary parts $\tilde{\mathbf{s}}(\tilde{\mathbf{x}}) := \nabla_{\Re[\tilde{\mathbf{x}}]} \log \tilde{p}(\tilde{\mathbf{x}}) + i \cdot \nabla_{\Im[\tilde{\mathbf{x}}]} \log \tilde{p}(\tilde{\mathbf{x}})$. We note that the gradient involved in the definition of the scores is non-trivial when the constraint in Equation (6) is enforced. In Appendix A.2, we establish a formal definition in this setting by interpreting the complex signals fulfilling this constraint as a submanifold in \mathbb{C}^{dx} . This constraint implies that the score components follow an analogous mirror symmetry: $\tilde{\mathbf{s}}_{\kappa} = \tilde{\mathbf{s}}_{N-\kappa}^*$ for all $\kappa \in [N]$. In the following, we shall implicitly rely on this definition.

3 Diffusing in the frequency domain

In the previous section, we have described how the typical diffusion formalism applies to time-series. We have also described how the DFT $\tilde{\mathbf{x}} = \mathcal{F}[\mathbf{x}]$ offers a full description of the time series \mathbf{x} in the frequency domain. The first step is to define how time-based diffusion translates in the frequency domain. Note that this is non-trivial as the DFT are complex-valued $\tilde{\mathbf{x}} \in \mathbb{C}^{d_x}$ signals. To solve this, we shall assume that the stochastic process in the time domain $\{\mathbf{x}(t)\}_{t=0}^T$, written compactly as \mathbf{x} , follows the diffusion process described in Equation (1). By leveraging the matrix formulation of the DFT $\tilde{\mathbf{x}} = U\mathbf{x}$, we will now derive diffusion SDEs in the frequency domain.

3.1 Diffusion SDEs

In order to derive the diffusion SDEs in the frequency domain, we shall simply apply the DFT operator to the forward diffusion SDE in the time domain from Equation (1). We note that this equation contains a standard Brownian motion \mathbf{w} . In the below lemma, we describe the DFT of \mathbf{w} and show that it contains two copies of a non-standard Brownian motion related by the constraint from Equation (6). We refer to this as a *mirrored Brownian motion*.

Lemma 3.1. (*DFT of standard Brownian motion*). Let \mathbf{w} be a standard Brownian motion on \mathbb{R}^{d_x} with $d_x = N \cdot M$, where $N \in \mathbb{N}^+$ is the number of time series steps and $M \in \mathbb{N}^+$ is the number of features tracked over time. Then $\mathbf{v} = U\mathbf{w}$ is a continuous stochastic process endowed with:

(1) **Mirror Symmetry.** For all $\kappa \in [N]$, $\mathbf{v}_\kappa = \mathbf{v}_{N-\kappa}^*$.

(2) **Real Brownian Motion.** \mathbf{v}_0 is a (real) standard Brownian motion on \mathbb{R}^M .

(3) **Complex Brownian Motions.** For all κ with $1 \leq \kappa \leq \lfloor N/2 \rfloor$, we can write $\mathbf{v}_\kappa = (\tilde{\mathbf{w}}_\kappa^1 + i\tilde{\mathbf{w}}_\kappa^2)/\sqrt{2}$ where $\tilde{\mathbf{w}}_\kappa^1$ and $\tilde{\mathbf{w}}_\kappa^2$ are independent standard Brownian motions on \mathbb{R}^M , except when N is even and $\kappa = N/2$, where $\mathbf{v}_{N/2}$ is a real standard Brownian motion on \mathbb{R}^M .

(4) **Independence.** The stochastic processes $\{\tilde{\mathbf{v}}_\kappa\}_{\kappa=0}^{\lfloor N/2 \rfloor}$ are mutually independent.

We call any stochastic process satisfying the above constraints a mirrored Brownian motion on \mathbb{C}^{d_x} .

Proof. The proof is given in Appendix A.3. □

Remark 3.2. Note that \mathbf{v} is not strictly speaking a Brownian motion, since it contains duplicate components due to the mirror symmetry. However, our theoretical analysis in Appendix A demonstrates that we can treat it as such by restricting to a subset of non-redundant components.

We now leverage Lemma 3.1 to show that $\tilde{\mathbf{x}}$ can be described by diffusion SDEs in the frequency domain which involve mirrored Brownian motions.

Proposition 3.3. (*Diffusion process in frequency domain*). Let us assume that \mathbf{x} is a diffusion process that is a solution of Equation (1), with $\mathbf{G}(t) = g(t) I_N$. Then $\tilde{\mathbf{x}} = \mathcal{F}[\mathbf{x}]$ is a solution to the forward diffusion process defined by:

$$d\tilde{\mathbf{x}} = \tilde{\mathbf{f}}(\tilde{\mathbf{x}}, t)dt + g(t)d\tilde{\mathbf{v}}, \quad (7)$$

where $\tilde{\mathbf{f}}(\tilde{\mathbf{x}}, t) = U\mathbf{f}(U^*\tilde{\mathbf{x}}, t)$ and $\tilde{\mathbf{v}}$ is a mirrored Brownian motion on \mathbb{C}^{d_x} . The associated reverse diffusion process is defined by:

$$d\tilde{\mathbf{x}} = \tilde{\mathbf{b}}(\tilde{\mathbf{x}}, t)dt + g(t)d\tilde{\mathbf{v}} \quad (8)$$

where $\tilde{\mathbf{b}}(\tilde{\mathbf{x}}, t) = \tilde{\mathbf{f}}(\tilde{\mathbf{x}}, t) - g^2(t)\Lambda^2\tilde{\mathbf{s}}(\tilde{\mathbf{x}}, t)$, $\Lambda \in \mathbb{R}^{N \times N}$ is a diagonal matrix defined in Appendix A.3, dt is a negative infinitesimal time step, and $\tilde{\mathbf{v}}$ is a mirrored Brownian motion on \mathbb{C}^{d_x} with time going from T to 0.

Proof. The proof is given in Appendix A.3. □

Proposition 3.3 gives us a recipe to implement diffusion in the frequency domain. It guarantees that the formalism introduced by [19] extends to this setting with one important difference: the Brownian motion must be replaced by a mirrored Brownian motion. Ignoring this prescription by taking $\tilde{\mathbf{v}}$ to be a Brownian motion on \mathbb{C}^{d_x} could lead to unintended consequences, such as generating complex-valued time series $\mathbf{x} \in \mathbb{C}^{d_x} \setminus \mathbb{R}^{d_x}$.

3.2 Denoising score matching

The reverse diffusion process given in Equation (8) provides an explicit way to samples time-series in the frequency domain provided we can compute $\tilde{\mathbf{b}}(\tilde{\mathbf{x}}, t)$, which involves the unknown score $\tilde{\mathbf{s}}$. Like in the time domain, and motivated by Equation (8), we build an approximation of the score with a function $\tilde{\mathbf{s}}_{\tilde{\theta}^*}$, whose parameters $\tilde{\theta}^*$ minimize the score matching objective:

$$\tilde{\theta}^* = \arg \min_{\tilde{\theta} \in \Theta} \mathbb{E}_{t, \tilde{\mathbf{x}}(0), \tilde{\mathbf{x}}(t)} [\mathcal{L}_{\text{SM}}(\tilde{\mathbf{s}}_{\tilde{\theta}}, \Lambda^2 \tilde{\mathbf{s}}_{t|0}, \tilde{\mathbf{x}}, t)] \quad (9)$$

with $t \sim \mathcal{U}(0, T)$, $\tilde{\mathbf{x}}(0) \sim \tilde{p}_0(\tilde{\mathbf{x}})$, $\tilde{\mathbf{x}}(t) \sim \tilde{p}_{t|0}(\tilde{\mathbf{x}}(t)|\tilde{\mathbf{x}}(0))$ and Λ is the diagonal matrix defined in Proposition 3.3. In practice, this objective is evaluated by first obtaining frequency representations of time-series, and then sampling from $\tilde{p}_{t|0}$ using Equation (7). Having trained $\tilde{\mathbf{s}}_{\tilde{\theta}^*}$, the backward process and $\tilde{\mathbf{s}}_{\tilde{\theta}^*}$ permit to draw samples from \tilde{p}_0 . It then suffices to apply the inverse DFT \mathcal{F}^{-1} to map the resulting complex-valued signals back into the time domain.

One important question remains at this stage. How does training a score in the frequency domain allow to generate DFT of time series sampled from p_{data} ? In other words, how does minimizing the score matching in Equation (9) imply that $\tilde{p}_0 \approx \tilde{p}_{\text{data}}$? To answer this question, a key observation is that we can associate an auxiliary score $\mathbf{s}'_{\tilde{\theta}}$ in the time domain to the score $\tilde{\mathbf{s}}_{\tilde{\theta}}$ by applying an inverse DFT \mathcal{F}^{-1} . Below, we show that minimizing the score matching loss from Equation (9) for the score $\tilde{\mathbf{s}}_{\tilde{\theta}}$ is equivalent to minimizing the score matching loss from Equation (3) for the auxiliary score $\mathbf{s}'_{\tilde{\theta}}$. This important observation connects the reverse diffusion process in the frequency domain described by Equation (8) with a reverse diffusion process in the time domain following Equation (2).

Proposition 3.4. (Score matching equivalence). *Consider a score $\tilde{\mathbf{s}}_{\tilde{\theta}} : \mathbb{C}^{d_x} \times [0, T] \rightarrow \mathbb{C}^{d_x}$ defined in the frequency domain and satisfying the mirror symmetry $[\tilde{\mathbf{s}}_{\tilde{\theta}}]_{\kappa} = [\tilde{\mathbf{s}}_{\tilde{\theta}}^*]_{N-\kappa}$ for all $\kappa \in [N]$. Let us define an auxiliary score $\mathbf{s}'_{\tilde{\theta}} : \mathbb{R}^{d_x} \times [0, T] \rightarrow \mathbb{R}^{d_x}$ as $(\mathbf{x}, t) \mapsto \mathbf{s}'_{\tilde{\theta}}(\mathbf{x}, t) = U^* \tilde{\mathbf{s}}_{\tilde{\theta}}(U\mathbf{x}, t)$ in the time domain. The score matching loss in the frequency domain is equivalent to the score matching loss for the auxiliary score in the time domain:*

$$\mathcal{L}_{\text{SM}}(\tilde{\mathbf{s}}_{\tilde{\theta}}, \Lambda^2 \tilde{\mathbf{s}}_{t|0}, \tilde{\mathbf{x}}, t) = \mathcal{L}_{\text{SM}}(\mathbf{s}'_{\tilde{\theta}}, \mathbf{s}_{t|0}, \mathbf{x}, t) \quad (10)$$

where $\tilde{\mathbf{s}}_{t|0}(\tilde{\mathbf{x}}, t) = \nabla_{\tilde{\mathbf{x}}(t)} \log \tilde{p}_{t|0}(\tilde{\mathbf{x}}(t)|\tilde{\mathbf{x}}(0))$, $\mathbf{s}_{t|0}(\mathbf{x}, t) = \nabla_{\mathbf{x}(t)} \log p_{t|0}(\mathbf{x}(t)|\mathbf{x}(0))$, and Λ is the diagonal matrix in Proposition 3.3.

Proof. The proof is given in Appendix A.4. □

Propositions 3.3 and 3.4 provide an explicit way to translate diffusion in the time domain to diffusion in the frequency domain. We note that attempting to solve Equations (3) and (9) in the finite-sample regime yields a local minimum solution in practice. Hence, there is no guarantee that training a score model in the frequency domain will converge to an auxiliary score $\mathbf{s}'_{\tilde{\theta}^*} = \mathbf{s}_{\theta^*}$ identical to the one obtained by training the score model in the time domain. In particular, having a score function $\tilde{\mathbf{s}}_{\tilde{\theta}}$ defined in the frequency domain is an important inductive bias, which is likely to alter the training dynamic. Through our experiments in the next section, we study the effect of this inductive bias on the resulting diffusion processes.

Take-away 1. Diffusion in the frequency domain can be implemented by replacing the standard Brownian motions with mirrored Brownian motions in the diffusion SDEs. The associated score can be optimized by minimizing a denoising score matching loss.

4 Comparing time and frequency diffusion

Table 1: Various datasets used in our experiments and some of their properties.

Dataset	Reference	Field	# Samples	# Steps N	# Features M
ECG	[29]	Healthcare	87,553	187	1
MIMIC-III	[30]	Healthcare	19,155	24	40
NASDAQ-2019	[31]	Finance	4,827	252	5
NASA-Charge	[32]	Engineering	2,396	251	4
NASA-Discharge			1,755	134	5
US-Droughts	[33]	Climate	2,797	365	13

In this section, we empirically analyze the effect of performing time series diffusion in the frequency domain. In Section 4.1, we show that frequency diffusion models better capture their training distribution than time models. In Section 4.2, we argue that these differences of performance can be attributed the localization of the time series in the frequency domain. Finally, in Section 4.3, we artificially create settings where time models outperform the frequency models in order to confirm this hypothesis. The code necessary to reproduce the results, along with detailed instructions, is provided in the following repository: <https://github.com/JonathanCrabbe/FourierDiffusion>.

Table 2: Sliced Wasserstein distances (\downarrow) evaluated in the time domain ($SW(\mathcal{D}_{\text{train}}, \mathcal{S}_{\text{time}})$, $SW(\mathcal{D}_{\text{train}}, \mathcal{S}_{\text{freq}})$) and in the frequency domain ($SW(\tilde{\mathcal{D}}_{\text{train}}, \tilde{\mathcal{S}}_{\text{time}})$, $SW(\tilde{\mathcal{D}}_{\text{train}}, \tilde{\mathcal{S}}_{\text{freq}})$) on the various datasets. For each distance, we report its mean ± 2 standard errors.

Dataset	Metric Domain	Diffusion Domain	
		Frequency	Time
ECG	Frequency	0.012 \pm 0.000	0.020 \pm 0.000
	Time	0.015 \pm 0.000	0.021 \pm 0.000
MIMIC-III	Frequency	0.144 \pm 0.004	0.206 \pm 0.006
	Time	0.152 \pm 0.004	0.211 \pm 0.006
NASDAQ-2019	Frequency	45.812 \pm 2.096	64.056 \pm 3.040
	Time	43.602 \pm 2.044	60.512 \pm 2.960
NASA-Charge	Frequency	0.211 \pm 0.008	0.27 \pm 0.006
	Time	0.229 \pm 0.008	0.316 \pm 0.008
NASA-Discharge	Frequency	1.999 \pm 0.084	2.974 \pm 0.134
	Time	2.028 \pm 0.082	2.942 \pm 0.134
US-Droughts	Frequency	0.633 \pm 0.018	2.849 \pm 0.090
	Time	0.738 \pm 0.020	2.913 \pm 0.092




Figure 1: Localization of time series in the time and frequency domains. Time series are more localized in the frequency domain.

Data. To illustrate the breadth of time series applications, we work with 6 different datasets described in Table 1. We observe that these datasets cover many use-cases (healthcare, finance, engineering and climate modelling), sample sizes, sequence lengths N and number of features tracked over time M . All the datasets are standardized before being fed to models. We also split the datasets into a training set $\mathcal{D}_{\text{train}}$ and a validation set \mathcal{D}_{val} . We provide more details on the datasets in Appendix B.1.

Models. For each dataset, we parametrize the time score model \mathbf{s}_θ and the frequency score model $\tilde{\mathbf{s}}_{\tilde{\theta}}$ as transformer encoders with 10 attention and MLP layers, each with 12 heads and dimension $d_{\text{model}} = 72$. Both models have learnable positional encoding as well as diffusion time t encoding through random Fourier features composed with a learnable dense layer. This results in models with 3.2M parameters. We use a VP-SDE with linear noise scheduling and $\beta_{\text{min}} = 0.1$ and $\beta_{\text{max}} = 20$, as in [19]. The score models are trained with the denoising score-matching loss, as defined in Section 3. All the models are trained for 200 epochs with batch size 64, AdamW optimizer and cosine learning rate scheduling (20 warmup epochs, $\text{lr}_{\text{max}} = 10^{-3}$). The selected model is the one achieving the lowest validation loss.

Time and frequency. Crucially, the only difference between the time and the frequency diffusion models is the domain in which their input time series are represented. Since all datasets are expressed in the time domain, they can directly be fed to the time diffusion model \mathbf{s}_θ . When it comes to the frequency diffusion model $\tilde{\mathbf{s}}_{\tilde{\theta}}$, the data is first mapped to the frequency domain by applying a DFT \mathcal{F} on each time series. In the time domain, the forward and reverse diffusion obey the SDEs in Equations (1) and (2). In the frequency domain, the forward and reverse diffusion obey the modified SDEs in Equations (7) and (8). The denoised samples $\tilde{\mathbf{x}}(0)$ obtained in the frequency domain can be pulled back to the time domain by applying an inverse DFT $\tilde{\mathbf{x}}(0) \mapsto \mathcal{F}^{-1}[\tilde{\mathbf{x}}(0)]$. In the following, we shall denote by $\mathcal{S}_{\text{time}} \subset \mathbb{R}^{d \times}$ and $\mathcal{S}_{\text{freq}} \subset \mathbb{R}^{d \times}$ the time representation of the samples generated by the time and frequency models. Similarly, we

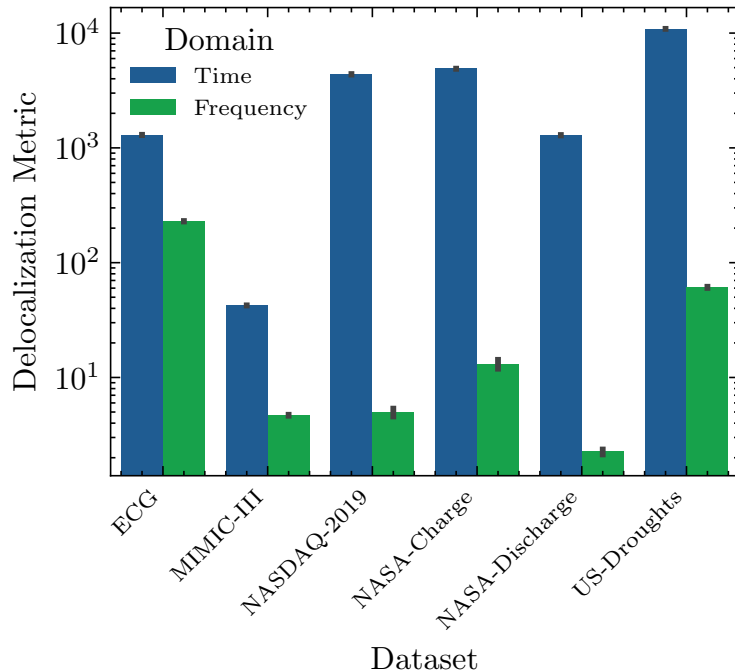


Figure 2: By evaluating our delocalization metrics in the time domain (Δ_{time}) and the frequency domain (Δ_{freq}), we quantitatively confirm that all the datasets are significantly more localized in the frequency domain. All the metrics are averaged over $\mathcal{D}_{\text{train}}$, their mean is reported with a 95% confidence interval.

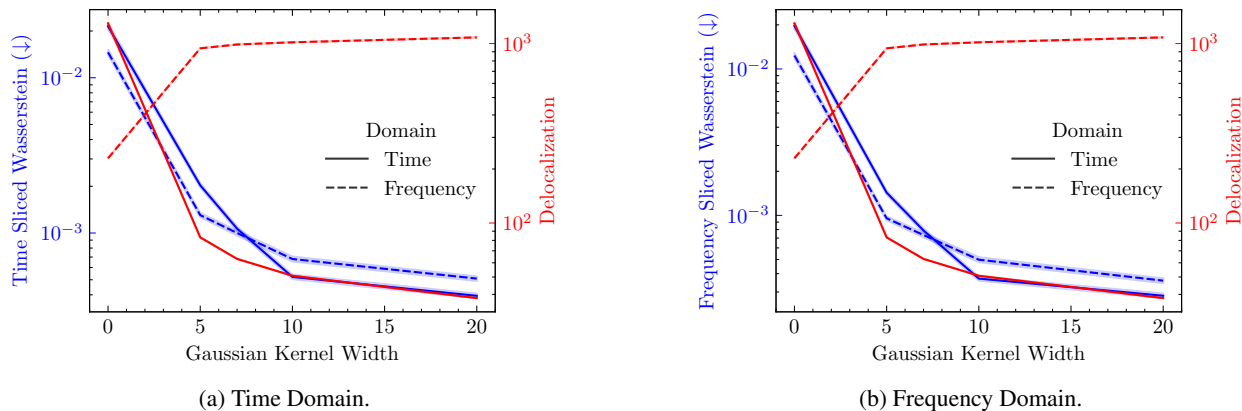


Figure 3: Sliced Wasserstein distances of time and frequency models (blue) and localization metrics in time and frequency domains (red) when smoothing the spectral representations of the time series with Gaussian kernels of variable width. Increasing the kernel width removes the localization in the frequency domain and increases the localization in the time domain. Coincidentally, the time diffusion model becomes better than the frequency diffusion model.

shall denote by $\tilde{\mathcal{S}}_{\text{time}} := \mathcal{F}[\mathcal{S}_{\text{time}}]$ and $\tilde{\mathcal{S}}_{\text{freq}} := \mathcal{F}[\mathcal{S}_{\text{freq}}]$ the frequency representations of these time series. We sample $|\mathcal{S}_{\text{time}}| = |\mathcal{S}_{\text{freq}}| = 10,000$ samples for each model by applying $T = 1,000$ diffusion time steps.

4.1 Which samples better capture the distribution?

Methodology. We are interested in the faithfulness of the samples generated by the time and frequency diffusion models. Ideally, this faithfulness should be evaluated by computing the Wasserstein distance between the true distribution and the distribution spanned by our diffusion models. However, this is impossible since the exact computation of the Wasserstein distance in intractable in input spaces of large dimension $d_X \gg 1$. In the case of images, [34] mitigates these problems

by mapping all the images in a lower dimensional representation space (the activation of the penultimate layer of an Inception-V3 model). This crucially relies on the fact that the Inception-V3 provides high-quality representations of images. Unfortunately, such a general representation of time series does not exist in practice. Hence, our evaluation needs to be performed in the input space \mathbb{R}^{d_x} directly. For this reason, we shall rely on the sliced Wasserstein distance introduced by [35], which has similar properties to the Wasserstein distance and can be efficiently estimated in high dimension spaces. With a slight abuse of notation, we shall denote by $SW(\mathcal{S}_1, \mathcal{S}_2)$ the sliced Wasserstein distances between the empirical distributions corresponding to the samples \mathcal{S}_1 and \mathcal{S}_2 . Its detailed definition is provided in Appendix B.2.

Analysis. The sliced Wasserstein distances are reported in Table 2. Interestingly, we observe that the frequency diffusion models consistently outperform the time diffusion models for all datasets both in the time and the frequency domain. In Appendix B.3, we show that using marginal Wasserstein distances instead of sliced Wasserstein distances essentially leads to the same conclusion. In order to verify that our observations are not specific to transformer backbones, we reproduce the same experiment with LSTM backbones in Appendix C and obtain similar results showing the superior performance of frequency diffusion models. While this observation already suggests the benefits of diffusing in the frequency domain rather than in the time domain, it is important to understand how these performance gains emerge. For this, we need to gain a better understanding of the training distributions, which is the object of next section.

4.2 How to explain the differences?

Signal energy analysis. Before formulating an hypothesis as to why the frequency models are better, it is helpful to gain a better understanding about the training distribution $\mathcal{D}_{\text{train}}$. To that end, we leverage the energy and spectral densities related to the time series, described in Section 2. These densities are represented in Figure 1, where we have averaged the densities over all time series in $\mathcal{D}_{\text{train}}$. By analyzing Figure 1b, we make a key observation: for all datasets, most of the time series energy in the frequency domain is localized on the frequency $\omega_0 = 0$ also known as the *fundamental frequency*. Furthermore, we observe that the energy quickly decays as the frequency increases. This observation suggests that the low frequencies capture most of the time series information. This is in stark contrast with the energy distribution over time in Figure 1a, which is more uniform over time for all the datasets. This asymmetry between the time series spectral localization and their temporal delocalization is a promising candidate to explain the superior performances of frequency diffusion. We now make this observation more quantitative.

Quantitative signal localization. In order to measure how delocalized a time series $\mathbf{x} \in \mathbb{R}^{d_x}$ is in the time domain, we shall use the *delocalization metrics* introduced by [36]:

$$\Delta_{\text{time}}(\mathbf{x}) := \min_{\tau \in [N]} \frac{1}{\|\mathbf{x}\|^2} \sum_{\tau' \in [N]} d_{\text{cyc}}(\tau, \tau') \|\mathbf{x}_{\tau'}\|_2^2, \quad (11)$$

where $d_{\text{cyc}} : [N]^2 \rightarrow [N]$ is the cyclic distance defined as $d_{\text{cyc}}(\tau, \tau') = \min(|\tau - \tau'|, N - |\tau - \tau'|)$, $\|\cdot\|$ denotes the Frobenius norm and $\|\cdot\|_2$ the Euclidean norm. Similarly, we can compute the delocalization in the frequency domain Δ_{freq} by replacing $\mathbf{x} \mapsto \tilde{\mathbf{x}}$ and $\tau, \tau' \mapsto \kappa, \kappa'$ in Equation (11). We report these delocalization metrics for each dataset in Figure 2. This quantitative analysis confirms our previous observations: the time series in all the datasets appear significantly more localized in the frequency domain. Interestingly, we never observe a time series that is localized in both the frequency and time domain simultaneously. This is in agreement with the *uncertainty principle* from [36], which echoes the foundational work of [37]. This verification is made in Appendix B.3.

A localization explanation. Based on the previous observation, we postulate that higher localization of the time series in the frequency domain contributes to the superior performance of frequency diffusion models. While we will test this hypothesis in the next section, it is useful to discuss the intuition behind it first. Due to the frequency localization, the frequency score model is presented with a representation of the time series where most of the relevant information is aligned with few components of the model’s input (i.e. the lower frequencies, especially the fundamental). This is in contrast with the time model, which is presented with an input where all the components matter equally. It follows that the frequency model does not need to learn a good distribution over all frequencies in order to generate samples of high quality, provided the lower frequency distributions are properly learned. The time model, on the other hand, needs to model all the time steps accurately in order to generate high-quality samples. If this intuition is correct, it would imply that delocalizing the signal in the frequency domain should reduce the gap of performance between time and frequency models. This is analyzed in the next section.

4.3 Should we always diffuse in the frequency domain?

Removing spectral localization. We would like to assess whether the localization of these signals in the frequency domain contribute to explain the superior performances of frequency diffusion models over time diffusion models. To

test this hypothesis, we intervene on a given dataset with the objective of varying the localization of frequency and time representations. With this in mind, it is useful to start from a dataset where the imbalance between time and frequency localization is not severe. By looking at Figure 2, it is clear that the ECG dataset is the best candidate. In order to gradually remove the spectral localization from the ECG dataset \mathcal{D}_{ECG} , we convolve the time series \mathbf{x} in the frequency domain with Gaussians of increasing kernel width $\sigma \in \mathbb{R}^+$ and define $\mathbf{x}^\sigma := \mathcal{F}^{-1}[\mathcal{F}[\mathbf{x}] \star \mathbf{g}^\sigma]$, where \star denotes the convolution between two signals and $\mathbf{g}_\kappa^\sigma := Z^{-1} \exp[-\kappa^2/(2\sigma^2)]$ for all $\kappa \in [N]$ is a Gaussian kernel with normalization $Z = \sum_{\kappa=1}^N \exp[-\kappa^2/(2\sigma^2)]$. This results in a family of corrupted datasets $\mathcal{D}_{\text{ECG}}^\sigma$, where the localization in the frequency domain decreases as σ increases. This is indeed what we observe in Figure 3 with the red curves. Coincidentally, the delocalization decreases in the time domain, in agreement with the uncertainty principle. The two curves cross at $\sigma \approx 2$, beyond which the time series are more localized in the time domain. Let us now analyze how the model performances evolve with different values of σ .

Analysis. We train time and frequency diffusion models on the datasets $\mathcal{D}_{\text{ECG}}^\sigma$ for $\sigma \in \{0, 5, 7, 10, 20\}$, where $\sigma = 0$ corresponds to the original ECG dataset: $\mathcal{D}_{\text{ECG}}^{\sigma=0} = \mathcal{D}_{\text{ECG}}$. As in Section 4.1, we measure the quality of 10,000 samples \mathcal{S} produced by these models with the Wasserstein distances $SW(\mathcal{D}_{\text{ECG}}^\sigma, \mathcal{S})$ in Figure 3a and $SW(\tilde{\mathcal{D}}_{\text{ECG}}^\sigma, \tilde{\mathcal{S}})$ in Figure 3b. By inspecting the blue curves from these figures, we notice that decreasing the frequency closes the gap between the time and the frequency models. Moreover, the two curves cross around at $\sigma \approx 7$, beyond which the time model outperforms the frequency model. This confirms our hypothesis that the localization (at least) partially explains the better performance of the frequency diffusion model.

A cautionary remark. Before concluding, we would like to incorporate a bit of nuance in our analysis. While the above results suggest that localization is an important factor to explain the superior performance of frequency diffusion models, we *do not* claim that this is the only explanation. There are essentially two things that suggest that this only partially explains the observations from Section 4.1. (1) In Figure 3, the blue curves and the red curves don't cross for the same value of σ (i.e. the red curves cross at $\sigma \approx 2$ and the blue curves at $\sigma \approx 7$). Hence, the time model requires time series that are substantially more localized in the time domain in order to outperform the frequency model. (2) In the limit $\sigma \rightarrow +\infty$, the time series become constant in the frequency domain. While this corresponds to a minimal localization in the frequency model, this should also be very easy for a frequency diffusion model to learn. Hence, decreasing the localization of the time series in a given domain does not necessarily imply that the resulting time series are more difficult to model in that domain.

Take-away 2. For all datasets in this paper, diffusing in the frequency domain yields better performances than in the time domain. A promising explanation for this is that time series from these datasets are substantially more localized in their spectral representation than in their temporal representation.

5 Discussion

In this work, we have improved the understanding of how diffusion models should be used with time series. We constructed a theoretical framework that extends the score-based SDE formulation of diffusion to complex-valued times series representations in the frequency domain. We have then demonstrated empirically that implementing time series diffusion in the frequency domain consistently outperforms the canonical diffusion in the time domain. Finally, we showed that the spectral localization of the time series plays a significant role to explain this phenomenon. There is a number of interesting ways to extend our work.

Time localized datasets. While all 6 datasets we have studied appear substantially more localized in their spectral representation, we do not claim that is is a universal property of real-world time series. In particular, it would be of interest to survey a large amount of time series datasets to determine the extent to which this phenomenon occurs.

Latent diffusion. Latent diffusion has emerged as a fruitful direction of research in the diffusion literature [5]. A promising direction would be to study how spectral representations of time series can be incorporated to their latent representations and whether this benefits the quality of the generated samples. We leave these insightful research directions for future work.

Acknowledgments

Jonathan Crabbé and Jan Stanczuk are funded by Aviva, Nicolas Huynh by Illumina and Mihaela van der Schaar by the Office of Naval Research (ONR), NSF 172251.

References

- [1] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8110–8119, 2020.
- [2] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021.
- [3] Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. Diffwave: A versatile diffusion model for audio synthesis. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*, 2021.
- [4] Chris Donahue, Julian McAuley, and Miller Puckette. Adversarial audio synthesis. In *International Conference on Learning Representations*, 2018.
- [5] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn 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.
- [6] Sander Dieleman, Laurent Sartran, Arman Roshannai, Nikolay Savinov, Yaroslav Ganin, Pierre H Richemond, Arnaud Doucet, Robin Strudel, Chris Dyer, Conor Durkan, et al. Continuous diffusion for categorical data. *arXiv preprint arXiv:2211.15089*, 2022.
- [7] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc 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.
- [8] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4):4713–4726, 2022.
- [9] Joseph L Watson, David Juergens, Nathaniel R Bennett, Brian L Trippe, Jason Yim, Helen E Eisenach, Woody Ahern, Andrew J Borst, Robert J Ragotte, Lukas F Milles, et al. De novo design of protein structure and function with rdiffusion. *Nature*, 620(7976):1089–1100, 2023.
- [10] Minkai Xu, Lantao Yu, Yang Song, Chence Shi, Stefano Ermon, and Jian Tang. Geodiff: a geometric diffusion model for molecular conformation generation. *ArXiv*, abs/2203.02923, 2022.
- [11] Claudio Zeni, Robert Pinsler, Daniel Zügner, Andrew Fowler, Matthew Horton, Xiang Fu, Sasha Shysheya, Jonathan Crabbé, Lixin Sun, Jake Smith, et al. Mattergen: a generative model for inorganic materials design. *arXiv preprint arXiv:2312.03687*, 2023.
- [12] Federico Gatta, Fabio Giampaolo, Edoardo Prezioso, Gang Mei, Salvatore Cuomo, and Francesco Piccialli. Neural networks generative models for time series. *J. King Saud Univ. Comput. Inf. Sci.*, 34:7920–7939, 2022.
- [13] Jinsung Yoon, Daniel Jarrett, and Mihaela Van der Schaar. Time-series generative adversarial networks. *Advances in neural information processing systems*, 32, 2019.
- [14] Ahmed M. Alaa, Alex J. Chan, and Mihaela van der Schaar. Generative time-series modeling with fourier flows. In *International Conference on Learning Representations*, 2021.
- [15] Cristóbal Esteban, Stephanie L. Hyland, and Gunnar Rätsch. Real-valued (medical) time series generation with recurrent conditional gans, 2017.
- [16] Aapo Hyvärinen and Peter Dayan. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(4), 2005.
- [17] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. PMLR, 2015.
- [18] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

- [19] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2020.
- [20] Lequan Lin, Zhengkun Li, Ruikun Li, Xuliang Li, and Junbin Gao. Diffusion models for time-series applications: a survey. *Frontiers of Information Technology & Electronic Engineering*, pages 1–23, 2023.
- [21] Thomas William Körner. *Fourier analysis*. Cambridge university press, 2022.
- [22] Kun Yi, Qi Zhang, Longbing Cao, Shoujin Wang, Guodong Long, Liang Hu, Hui He, Zhendong Niu, Wei Fan, and Hui Xiong. A survey on deep learning based time series analysis with frequency transformation. *CoRR*, *abs/2302.02173*, 2023.
- [23] Donghyeok Shin, Seungjae Shin, and Il-Chul Moon. Frequency domain-based dataset distillation. *arXiv preprint arXiv:2311.08819*, 2023.
- [24] Angus Phillips, Thomas Seror, Michael Hutchinson, Valentin De Bortoli, Arnaud Doucet, and Emile Mathieu. Spectral diffusion processes. *arXiv preprint arXiv:2209.14125*, 2022.
- [25] Brian D.O. Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982.
- [26] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.
- [27] Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011.
- [28] Peter J. Schreier and Louis L. Scharf. *Statistical Signal Processing of Complex-Valued Data: The Theory of Improper and Noncircular Signals*. Cambridge University Press, 2010.
- [29] Mohammad Kachuee, Shayan Fazeli, and Majid Sarrafzadeh. Ecg heartbeat classification: A deep transferable representation. In *2018 IEEE international conference on healthcare informatics (ICHI)*, pages 443–444. IEEE, 2018.
- [30] Alistair Johnson, Tom Pollard, and Roger Mark. MIMIC-III clinical database (version 1.4). *PhysioNet*, 10(C2XW26):2, 2016.
- [31] Oleh Onyshchak. Stock market dataset, 2020.
- [32] B Saha and T. Goebel. Battery data set, nasa ames prognostics data repository, 2007.
- [33] Christoph Minixhofer. Predict droughts using weather and soil data, 2021.
- [34] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.
- [35] Nicolas Bonneel, Julien Rabin, Gabriel Peyré, and Hanspeter Pfister. Sliced and radon wasserstein barycenters of measures. *Journal of Mathematical Imaging and Vision*, 51:22–45, 2015.
- [36] Sangnam Nam. An Uncertainty Principle for Discrete Signals. In *SampTA*, Bremen, Germany, 2013.
- [37] Werner Heisenberg. Über den anschaulichen inhalt der quantentheoretischen kinematik und mechanik. *Zeitschrift für Physik*, 43(3):172–198, 1927.
- [38] Peter E Kloeden, Eckhard Platen, Peter E Kloeden, and Eckhard Platen. *Stochastic differential equations*. Springer, 1992.
- [39] Shirly Wang, Matthew BA McDermott, Geeticka Chauhan, Marzyeh Ghassemi, Michael C Hughes, and Tristan Naumann. MIMIC-EXTRACT: A data extraction, preprocessing, and representation pipeline for MIMIC-III. In *Proceedings of the ACM conference on health, inference, and learning*, pages 222–235, 2020.

A Mathematical details

A.1 DFT properties

Proposition A.1 (Unitarity of the DFT operator). *The DFT matrix $U \in \mathbb{C}^{N \times N}$ with elements $[U]_{\kappa\tau} := N^{-1/2} \exp(-i\omega_\kappa\tau)$ with $\omega_\kappa := \frac{\kappa 2\pi}{N}$ is unitary.*

Proof. Let U^* denote the conjugate transpose of U . For any κ and τ in $[N]$, we have:

$$\begin{aligned} [UU^*]_{\kappa\tau} &= \sum_{\beta=0}^{N-1} [U]_{\kappa\beta} [U^*]_{\beta\tau} \\ &= \frac{1}{N} \sum_{\beta=0}^{N-1} \exp(-i\omega_\kappa\beta) \exp(i\omega_\beta\tau) \\ &= \frac{1}{N} \sum_{\beta=0}^{N-1} \exp(-i\omega_{\kappa-\tau}\beta) \end{aligned}$$

Hence, if $\kappa = \tau$, we have $[UU^*]_{\kappa\tau} = 1$, otherwise $[UU^*]_{\kappa\tau} = 0$, since $\exp(-i\omega_{\kappa-\tau}N) = 1$. This is equivalent to $UU^* = I_N$, i.e. U is unitary. \square

Proposition A.2. *The DFT $\tilde{\mathbf{x}} = \mathcal{F}[\mathbf{x}] = U\mathbf{x}$ of a real-valued time series $\mathbf{x} \in \mathbb{R}^{d_x}$ verifies the following mirror symmetry for all $\kappa \in [N]$:*

$$\tilde{\mathbf{x}}_\kappa = \tilde{\mathbf{x}}_{N-\kappa}^*$$

Proof. Let κ and τ be in $[N]$. We first note that $\exp(i\omega_{N-\kappa}\tau) = \exp(i(\omega_N - \omega_\kappa)\tau) = \exp(-i\omega_\kappa\tau)$. Hence,

$$\begin{aligned} \tilde{\mathbf{x}}_{N-\kappa}^* &= \sum_{\tau=0}^{N-1} [U]_{N-\kappa,\tau}^* \mathbf{x}_\tau && (\mathbf{x} \text{ is real valued}) \\ &= N^{-1/2} \sum_{\tau=0}^{N-1} \exp(i\omega_{N-\kappa}\tau) \mathbf{x}_\tau \\ &= N^{-1/2} \sum_{\tau=0}^{N-1} \exp(-i\omega_\kappa\tau) \mathbf{x}_\tau \\ &= \sum_{\tau=0}^{N-1} [U]_{\kappa,\tau} \mathbf{x}_\tau \\ &= \tilde{\mathbf{x}}_\kappa \end{aligned}$$

\square

A.2 Densities and scores for constrained signals

As we have mentioned in Section 2, the redundancy between certain components of the DFT $\tilde{\mathbf{x}} \in \mathbb{C}^{d_x}$ of \mathbf{x} expressed by Equation (6) needs to be taken into account if we wish to define a density \tilde{p} for the time series distribution in the frequency domain. In particular, this redundancy implies that the density is really defined on a submanifold $\mathbb{C}_{\text{constr}}^{d_x} := \{\tilde{\mathbf{x}} = (\tilde{\mathbf{x}}_0, \dots, \tilde{\mathbf{x}}_{N-1}) \in \mathbb{C}^{d_x} \mid \tilde{\mathbf{x}}_\kappa = \tilde{\mathbf{x}}_{N-\kappa}^* \forall \kappa \in [N]\}$ of complex signals fulfilling the constraint. We can define a coordinate chart $\varphi : \mathbb{C}_{\text{constr}}^{d_x} \rightarrow \mathbb{R}^{d_x}$ on this submanifold by extracting the unconstrained part of a DFT $\tilde{\mathbf{x}}$. This operator simply concatenates the relevant real and imaginary parts of the DFT and is defined as follows for all $\tilde{\mathbf{x}} \in \mathbb{C}_{\text{constr}}^{d_x}$:

$$\varphi[\tilde{\mathbf{x}}] = \begin{cases} (\Re[\tilde{\mathbf{x}}_\kappa])_{\kappa=0}^{N/2} \oplus (\Im[\tilde{\mathbf{x}}_\kappa])_{\kappa=1}^{N/2-1} & \text{if } N \in 2\mathbb{N} \\ (\Re[\tilde{\mathbf{x}}_\kappa])_{\kappa=0}^{[N/2]} \oplus (\Im[\tilde{\mathbf{x}}_\kappa])_{\kappa=1}^{[N/2]} & \text{else,} \end{cases} \quad (12)$$

where $\mathbf{v}_1 \oplus \mathbf{v}_2$ denotes the concatenation of two vectors $\mathbf{v}_1 \in \mathbb{R}^{d_1}$ and $\mathbf{v}_2 \in \mathbb{R}^{d_2}$, with $d_1, d_2 \in \mathbb{N}$. Due to Equation (6), one can unambiguously reconstruct $\tilde{\mathbf{x}} \in \mathbb{C}_{\text{constr}}^{d_X}$ from $\varphi[\tilde{\mathbf{x}}]$. Hence, the coordinate chart admits an inverse $\varphi^{-1} : \mathbb{R}^{d_X} \rightarrow \mathbb{C}_{\text{constr}}^{d_X}$ defined as follows for all $\mathbf{z} = (\mathbf{z}_0, \dots, \mathbf{z}_{N-1}) \in \mathbb{R}^{d_X}$:

$$\varphi^{-1}[\mathbf{z}] = \begin{cases} (\mathbf{z}_0) \oplus (\mathbf{z}_\kappa + i \cdot \mathbf{z}_{N/2+\kappa})_{\kappa=1}^{N/2-1} \oplus (\mathbf{z}_{N/2}) \oplus (\mathbf{z}_{N/2-\kappa} - i \cdot \mathbf{z}_{N-\kappa})_{\kappa=1}^{N/2-1} & \text{if } N \in 2\mathbb{N} \\ (\mathbf{z}_0) \oplus (\mathbf{z}_\kappa + i \cdot \mathbf{z}_{\lfloor N/2 \rfloor + \kappa})_{\kappa=1}^{\lfloor N/2 \rfloor} \oplus (\mathbf{z}_{\lceil N/2 \rceil - \kappa} - i \cdot \mathbf{z}_{N-\kappa})_{\kappa=1}^{\lfloor N/2 \rfloor} & \text{else.} \end{cases} \quad (13)$$

With this coordinate chart, it becomes possible to rigorously define the density $\tilde{p} : \mathbb{C}_{\text{constr}}^{d_X} \rightarrow \mathbb{R}^+$. One simply defines a probability density $\tilde{p}_\varphi : \mathbb{R}^{d_X} \rightarrow \mathbb{R}^+$ over the real vector space \mathbb{R}^{d_X} on which the coordinate chart is defined and pull it back to the manifold of constrained signals $\tilde{p} := \tilde{p}_\varphi \circ \varphi$. This indeed defines a density that depends on the real and imaginary parts of the frequency representations $\tilde{\mathbf{x}} \in \mathbb{C}_{\text{constr}}^{d_X}$ of time series, as announced in Section 2.

Finally, it remains to rigorously define the score $\tilde{\mathbf{s}} : \mathbb{C}_{\text{constr}}^{d_X} \times [0, T] \rightarrow \mathbb{C}^{d_X}$ in the frequency domain. This can be done by starting from the real score $\tilde{\mathbf{s}}_\varphi : \mathbb{R}^{d_X} \times [0, T] \rightarrow \mathbb{R}^{d_X}$ that is well-defined for all $\mathbf{z} \in \mathbb{R}^{d_X}$ and $t \in [0, T]$ as $\tilde{\mathbf{s}}_\varphi(\mathbf{z}, t) = \nabla_{\mathbf{z}} \log \tilde{p}_{\varphi, t}(\mathbf{z})$. Again, one can expand this vector field to the constrained manifold by defining for all $\tilde{\mathbf{x}} \in \mathbb{C}_{\text{constr}}^{d_X}$ and all $t \in [0, T]$:

$$\tilde{\mathbf{s}}(\tilde{\mathbf{x}}, t) := \varphi^{-1}[\tilde{\mathbf{s}}_\varphi(\varphi(\tilde{\mathbf{x}}), t)]. \quad (14)$$

This indeed defines a vector field involving partial derivatives of the log density with respect to the real and imaginary parts of the frequency representations $\tilde{\mathbf{x}} \in \mathbb{C}_{\text{constr}}^{d_X}$ of time series and respects the mirror symmetry by virtue of Equation (13). Everything is then consistent with the discussion from Section 2.

A.3 Diffusion SDEs in the frequency domain

Lemma 3.1. (*DFT of standard Brownian motion*). Let \mathbf{w} be a standard Brownian motion on \mathbb{R}^{d_X} with $d_X = N \cdot M$, where $N \in \mathbb{N}^+$ is the number of time series steps and $M \in \mathbb{N}^+$ is the number of features tracked over time. Then $\mathbf{v} = U\mathbf{w}$ is a continuous stochastic process endowed with:

(1) **Mirror Symmetry.** For all $\kappa \in [N]$, $\mathbf{v}_\kappa = \mathbf{v}_{N-\kappa}^*$.

(2) **Real Brownian Motion.** \mathbf{v}_0 is a (real) standard Brownian motion on \mathbb{R}^M .

(3) **Complex Brownian Motions.** For all κ with $1 \leq \kappa \leq \lfloor N/2 \rfloor$, we can write $\mathbf{v}_\kappa = (\tilde{\mathbf{w}}_\kappa^1 + i\tilde{\mathbf{w}}_\kappa^2)/\sqrt{2}$ where $\tilde{\mathbf{w}}_\kappa^1$ and $\tilde{\mathbf{w}}_\kappa^2$ are independent standard Brownian motions on \mathbb{R}^M , except when N is even and $\kappa = N/2$, where $\mathbf{v}_{N/2}$ is a real standard Brownian motion on \mathbb{R}^M .

(4) **Independence.** The stochastic processes $\{\tilde{\mathbf{v}}_\kappa\}_{\kappa=0}^{\lfloor N/2 \rfloor}$ are mutually independent.

We call any stochastic process satisfying the above constraints a mirrored Brownian motion on \mathbb{C}^{d_X} .

Proof. (1) **Mirror Symmetry.** This point directly follows from the symmetry of the DFT, proved in Proposition A.2.

In what follows, we consider without loss of generality the case $M = 1$, since the cases $M > 1$ can be handled similarly by flattening matrices into vectors and using the same arguments as below. Let us first decompose U into its real and imaginary parts, i.e. $U = U_{re} + iU_{im}$, where U_{re} and U_{im} are in $\mathbb{R}^{N \times N}$. Note that these two matrices are both symmetric. Computing the distribution of the components of \mathbf{v} will require the knowledge of covariance matrices, which will depend on U_{re}^2 , U_{im}^2 and $U_{re}U_{im}$. For any κ and τ in $[N]$,

$$\begin{aligned} [U_{re}^2]_{\kappa, \tau} &= \frac{1}{N} \sum_{\gamma=0}^{N-1} \cos\left(\frac{2\pi}{N}\kappa\gamma\right) \cos\left(\frac{2\pi}{N}\tau\gamma\right) \\ &= \frac{1}{2N} \sum_{\gamma=0}^{N-1} \left(\cos\left(\frac{2\pi}{N}(\kappa + \tau)\gamma\right) + \cos\left(\frac{2\pi}{N}(\kappa - \tau)\gamma\right) \right) \end{aligned}$$

Similarly,

$$\begin{aligned} [U_{im}^2]_{\kappa, \tau} &= \frac{1}{N} \sum_{\gamma=0}^{N-1} \sin\left(\frac{2\pi}{N}\kappa\gamma\right) \sin\left(\frac{2\pi}{N}\tau\gamma\right) \\ &= \frac{1}{2N} \sum_{\gamma=0}^{N-1} \left(\cos\left(\frac{2\pi}{N}(\kappa - \tau)\gamma\right) - \cos\left(\frac{2\pi}{N}(\kappa + \tau)\gamma\right) \right) \end{aligned}$$

To compute these sums, we consider the situations where:

- $N|\kappa - \tau$: this is equivalent to $\underline{\kappa = \tau}$, since $-(N-1) \leq \kappa - \tau \leq N-1$
- $N|\kappa + \tau$: this is equivalent to $\underline{\kappa = \tau = 0}$ or $\underline{\kappa + \tau = N}$, since $0 \leq \kappa + \tau \leq 2(N-1)$

Hence, if N is even:

$$[U_{re}^2]_{\kappa,\tau} = \begin{cases} 1 & \text{if } \kappa = \tau = 0 \text{ or } \kappa = \tau = \frac{N}{2} \\ \frac{1}{2} & \text{if } \kappa \notin \{0, N/2\} \text{ and } \kappa = \tau \text{ or } \kappa = N - \tau \\ 0 & \text{otherwise} \end{cases} \text{ and } [U_{im}^2]_{\kappa,\tau} = \begin{cases} 0 & \text{if } \kappa = \tau = 0 \text{ or } \kappa = \tau = \frac{N}{2} \\ \frac{1}{2} & \text{if } \kappa \notin \{0, N/2\} \text{ and } \kappa = \tau \\ -\frac{1}{2} & \text{if } \kappa \neq N/2 \text{ and } \kappa = N - \tau \\ 0 & \text{otherwise} \end{cases} \quad (15)$$

If N is odd:

$$[U_{re}^2]_{\kappa,\tau} = \begin{cases} 1 & \text{if } \kappa = \tau = 0 \\ \frac{1}{2} & \text{if } \kappa \neq 0 \text{ and } \kappa = \tau, \text{ or } \kappa = N - \tau \\ 0 & \text{otherwise} \end{cases} \text{ and } [U_{im}^2]_{\kappa,\tau} = \begin{cases} 0 & \text{if } \kappa = \tau = 0 \\ \frac{1}{2} & \text{if } \kappa \neq 0 \text{ and } \kappa = \tau \\ -\frac{1}{2} & \text{if } \kappa = N - \tau \\ 0 & \text{otherwise} \end{cases} \quad (16)$$

Finally, we compute $U_{re}U_{im}$:

$$\begin{aligned} [U_{re}U_{im}]_{\kappa,\tau} &= \frac{1}{N} \sum_{\gamma=0}^{N-1} \cos\left(\frac{2\pi}{N}\kappa\gamma\right) \sin\left(\frac{2\pi}{N}\tau\gamma\right) \\ &= \frac{1}{2N} \sum_{\gamma=0}^{N-1} \left(\sin\left(\frac{2\pi}{N}(\tau - \kappa)\gamma\right) + \sin\left(\frac{2\pi}{N}(\kappa + \tau)\gamma\right) \right) \\ &= 0 \end{aligned}$$

Hence, $U_{re}U_{im} = 0_N$ and similarly $U_{im}U_{re} = 0_N$ by taking the transpose and using the symmetry of U_{re} and U_{im} .

We can now characterize the distribution followed by the stochastic process \mathbf{v} . We first write $U_{col} = \begin{pmatrix} U_{re} \\ U_{im} \end{pmatrix}$, and then notice that \mathbf{v} can be investigated through the lens of its flattened version $\mathbf{v}_{flat} = U_{col}\mathbf{w}$, which is a stochastic process in \mathbb{R}^{2N} .

(2) Real Brownian Motion. First, \mathbf{v}_0 is real-valued, by using **(1) Mirror Symmetry**. We then have the following:

- $\mathbf{v}_0(0) = 0$ almost surely: this stems from $\mathbf{w}(0) = 0$ almost surely, since \mathbf{w} is a multivariate standard Brownian motion and $\mathbf{v} = U\mathbf{w}$.
- Continuity of $t \rightarrow \mathbf{v}_0(t)$ almost surely: \mathbf{w} satisfies the continuity property and the DFT operator (seen as a complex operator) is linear, hence \mathbf{v}_0 is also continuous with respect to t almost surely.
- Stationary and independent increments: this follows from the linearity of the DFT operator and \mathbf{w} being a Brownian motion.
- For any $t, s \geq 0$, $\mathbf{v}_0(t+s) - \mathbf{v}_0(s) \sim \mathcal{N}(0, t)$: to see this, we notice that $\mathbf{v}_0(t+s) - \mathbf{v}_0(s)$ is Gaussian³ since it is a linear transform of $\mathbf{w}(t+s) - \mathbf{w}(s)$. Moreover, its mean and its variance are given by:

$$\begin{aligned} \mathbb{E}[\mathbf{v}_0(t+s) - \mathbf{v}_0(s)] &= [U_{col}\mathbb{E}[\mathbf{w}(t+s) - \mathbf{w}(s)]]_0 \\ &= 0 \end{aligned}$$

$$\begin{aligned} \text{Var}(\mathbf{v}_0(t+s) - \mathbf{v}_0(s)) &= [U_{col}\text{Cov}(\mathbf{w}(t+s) - \mathbf{w}(s), \mathbf{w}(t+s) - \mathbf{w}(s))U_{col}^T]_{0,0} \\ &= t[U_{col}U_{col}^T]_{0,0} \\ &= t[U_{re}^2]_{0,0} \\ &= t \end{aligned}$$

³Note that a random variable almost surely equal to 0 can be seen as a degenerate Gaussian with mean 0 and variance 0.

Hence, we have shown that \mathbf{v}_0 is a real Brownian motion.

(3) Complex Brownian Motions. Let $1 \leq \kappa < \lfloor N/2 \rfloor$. Then, $\Re(\mathbf{v}_\kappa)$ and $\Im(\mathbf{v}_\kappa)$ follow the first three properties of a Brownian motion, using the same arguments as above. For the last point, we first characterize the distribution of $\Re(\mathbf{v}_\kappa)$ and $\Im(\mathbf{v}_\kappa)$, and then show that they are independent.

Distribution of $\Re(\mathbf{v}_\kappa)$. $\sqrt{2}\Re(\mathbf{v}_\kappa)$ is a standard Brownian motion in \mathbb{R} : For any $t, s \geq 0$, $\Re(\mathbf{v}_\kappa)(t+s) - \Re(\mathbf{v}_\kappa)(s)$ is Gaussian since it is a linear transform of $\mathbf{w}(t+s) - \mathbf{w}(s)$ which is a Gaussian vector. We can compute its mean and its variance:

$$\begin{aligned} \mathbb{E}[\Re(\mathbf{v}_\kappa)(t+s) - \Re(\mathbf{v}_\kappa)(s)] &= [U_{col} \mathbb{E}[\mathbf{w}(t+s) - \mathbf{w}(s)]]_{\kappa} \\ &= 0 \end{aligned}$$

$$\begin{aligned} \text{Var}(\Re(\mathbf{v}_\kappa)(t+s) - \Re(\mathbf{v}_\kappa)(s)) &= [U_{col} \text{Cov}(\mathbf{w}(t+s) - \mathbf{w}(s), \mathbf{w}(t+s) - \mathbf{w}(s)) U_{col}^T]_{\kappa, \kappa} \\ &= t [U_{col} U_{col}^T]_{\kappa, \kappa} \\ &= t [U_{re}^2]_{\kappa, \kappa} \\ &= \frac{1}{2} t \end{aligned}$$

Distribution of $\Im(\mathbf{v}_\kappa)$. Similarly, we can prove with the same arguments that $\sqrt{2}\Im(\mathbf{v}_\kappa)$ is a standard Brownian motion in \mathbb{R} .

Independence of $\Re(\mathbf{v}_\kappa)$ and $\Im(\mathbf{v}_\kappa)$. Let k and m be two strictly positive integers, and let $(t_1, \dots, t_k) \in (\mathbb{R}_+^*)^k$ and $(t'_1, \dots, t'_m) \in (\mathbb{R}_+^*)^m$. We need to show that the vectors $\mathbf{v}_\kappa^{re} = (\Re(\mathbf{v}_\kappa)(t_1), \dots, \Re(\mathbf{v}_\kappa)(t_k))$ and $\mathbf{v}_\kappa^{im} = (\Im(\mathbf{v}_\kappa)(t'_1), \dots, \Im(\mathbf{v}_\kappa)(t'_m))$ are independent. First, the concatenation of \mathbf{v}_κ^{re} and \mathbf{v}_κ^{im} can be expressed as a linear transform of $(\mathbf{w}(t_1), \dots, \mathbf{w}(t_k), \mathbf{w}(t'_1), \dots, \mathbf{w}(t'_m))$, which is a Gaussian vector since \mathbf{w} is a Brownian motion. Consequently, $(\mathbf{v}_\kappa^{re}, \mathbf{v}_\kappa^{im})$ is also a Gaussian vector. Now, let $l \in \{1, \dots, k\}$ and $n \in \{1, \dots, m\}$. Then,

$$\begin{aligned} \text{Cov}(\Re(\mathbf{v}_\kappa)(t_l), \Im(\mathbf{v}_\kappa)(t'_n)) &= [U_{re} \text{Cov}(\mathbf{w}(t_l), \mathbf{w}(t'_n)) U_{im}]_{\kappa, \kappa} \\ &= \min(t_l, t'_n) [U_{re} U_{im}]_{\kappa, \kappa} \\ &= 0 \end{aligned}$$

Given this covariance structure and the fact that $(\mathbf{v}_\kappa^{re}, \mathbf{v}_\kappa^{im})$ is a Gaussian vector, we have $\mathbf{v}_\kappa^{re} \perp\!\!\!\perp \mathbf{v}_\kappa^{im}$. Since this holds true for any choice of $(t_1, \dots, t_k) \in \mathbb{R}_+^{*k}$ and $(t'_1, \dots, t'_m) \in \mathbb{R}_+^{*m}$, we conclude that $\Re(\mathbf{v}_\kappa) \perp\!\!\!\perp \Im(\mathbf{v}_\kappa)$. The case $k = \lfloor \frac{N}{2} \rfloor$ can be handled using the same arguments, by distinguishing the cases N odd and N even.

(4) Independence. The mutual independence of the stochastic processes $\{\tilde{\mathbf{v}}_\kappa\}_{\kappa=0}^{\lfloor N/2 \rfloor}$ follows from the structure of U_{re}^2 and U_{im}^2 . Indeed, for any m and n such that $m \neq n$, $0 \leq m \leq \lfloor \frac{N}{2} \rfloor$, and $0 \leq n \leq \lfloor \frac{N}{2} \rfloor$, we have $[U_{re}^2]_{m,n} = [U_{im}^2]_{m,n} = [U_{re} U_{im}]_{m,n} = 0$. We can then apply the same argument as in 3) of **(3) Complex Brownian Motions** to obtain the mutual independence of the stochastic processes $\{\tilde{\mathbf{v}}_\kappa\}_{\kappa=0}^{\lfloor N/2 \rfloor}$. \square

Proposition 3.3. (*Diffusion process in frequency domain*). Let us assume that \mathbf{x} is a diffusion process that is a solution of Equation (1), with $\mathbf{G}(t) = g(t) I_N$. Then $\tilde{\mathbf{x}} = \mathcal{F}[\mathbf{x}]$ is a solution to the forward diffusion process defined by:

$$d\tilde{\mathbf{x}} = \tilde{\mathbf{f}}(\tilde{\mathbf{x}}, t) dt + g(t) d\tilde{\mathbf{v}}, \quad (7)$$

where $\tilde{\mathbf{f}}(\tilde{\mathbf{x}}, t) = U \mathbf{f}(U^* \tilde{\mathbf{x}}, t)$ and $\tilde{\mathbf{v}}$ is a mirrored Brownian motion on \mathbb{C}^{d_x} . The associated reverse diffusion process is defined by:

$$d\tilde{\mathbf{x}} = \tilde{\mathbf{b}}(\tilde{\mathbf{x}}, t) dt + g(t) d\tilde{\mathbf{v}} \quad (8)$$

where $\tilde{\mathbf{b}}(\tilde{\mathbf{x}}, t) = \tilde{\mathbf{f}}(\tilde{\mathbf{x}}, t) - g^2(t) \Lambda^2 \tilde{\mathbf{s}}(\tilde{\mathbf{x}}, t)$, $\Lambda \in \mathbb{R}^{N \times N}$ is a diagonal matrix defined in Appendix A.3, dt is a negative infinitesimal time step, and $\tilde{\mathbf{v}}$ is a mirrored Brownian motion on \mathbb{C}^{d_x} with time going from T to 0.

Proof. Forward SDE. Since $\tilde{\mathbf{x}} = U \mathbf{x}$, we can apply the multivariate Itô's lemma (Eq. 8.3, [38]), and obtain a forward SDE for $\tilde{\mathbf{x}}$:

$$d\tilde{\mathbf{x}} = U \mathbf{f}(\mathbf{x}, t) dt + g(t) U d\mathbf{w}' \quad (17)$$

where we have implicitly used the fact that $\mathbf{G}(t) = g(t) I_N$ and U commute. By Lemma 3.1, $\tilde{\mathbf{v}} = U \mathbf{w}'$ is a mirrored Brownian motion on \mathbb{C}^{d_x} , which gives the result.

Reverse SDE. In order to derive the reverse SDE for $\tilde{\mathbf{x}}$, we follow three steps: (1) we write a forward SDE for which the truncation $\varphi[\tilde{\mathbf{x}}]$ is a solution ; (2) we write its associated reverse-time SDE [25] and (3) we derive the full reverse SDE for the stochastic process $\tilde{\mathbf{x}}$.

Step 1: From Equation (17) and using Lemma 3.1, we can extract the following forward SDE for $\varphi[\tilde{\mathbf{x}}]$, which is defined in Equation (12) :

$$d\varphi[\tilde{\mathbf{x}}] = \varphi[U\mathbf{f}(U^*\varphi^{-1}(\varphi[\tilde{\mathbf{x}}]), t)]dt + g(t)\Lambda d\tilde{\mathbf{w}} \quad (18)$$

where:

- $\varphi^{-1} : \mathbb{R}^{d_x} \rightarrow \mathbb{C}^{d_x}$ satisfies $\varphi^{-1}(\varphi[U\mathbf{y}]) = U\mathbf{y}$ for all $\mathbf{y} \in \mathbb{R}^{d_x}$ as defined in Equation (13).
- $\Lambda \in \mathbb{R}^{N \times N}$ is a diagonal matrix such that $[\Lambda]_{\kappa, \kappa} = \begin{cases} 1 & \text{if } \kappa = 0, \text{ or } N \text{ is even and } \kappa = N/2 \\ \frac{1}{\sqrt{2}} & \text{otherwise} \end{cases}$,
- $\tilde{\mathbf{w}}$ is a stochastic process in \mathbb{R}^{d_x} which satisfies $\Lambda \tilde{\mathbf{w}} = \varphi[U\mathbf{w}']$. The above point, together with Lemma 3.1), then implies that $\tilde{\mathbf{w}}$ is in fact a standard multivariate Brownian motion.

Step 2: The associated reverse-time SDE [25] is given by:

$$d\varphi[\tilde{\mathbf{x}}] = \{\varphi[U\mathbf{f}(U^*\varphi^{-1}(\varphi[\tilde{\mathbf{x}}]), t)] - g(t)^2\Lambda\Lambda^T\nabla_{\varphi[\tilde{\mathbf{x}}]}\log\tilde{p}_t(\varphi[\tilde{\mathbf{x}}])\}dt + g(t)\Lambda d\hat{\mathbf{w}} \quad (19)$$

where $\hat{\mathbf{w}}$ is a standard Brownian motion on \mathbb{R}^{d_x} .

Step 3: Since $\varphi^{-1}(\varphi[\tilde{\mathbf{x}}]) = \tilde{\mathbf{x}}$, we can recover the reverse-time SDE followed by $\tilde{\mathbf{x}}$ by applying the operator φ^{-1} to $\varphi[\tilde{\mathbf{x}}]$, and using the Itô's lemma:

$$d\tilde{\mathbf{x}} = \{\varphi^{-1}(\varphi[U\mathbf{f}(U^*\varphi^{-1}(\varphi[\tilde{\mathbf{x}}]), t)]) - g(t)^2\varphi^{-1}(\Lambda^2\nabla_{\varphi[\tilde{\mathbf{x}}]}\log\tilde{p}_t(\varphi[\tilde{\mathbf{x}}]))\}dt + g(t)\varphi^{-1}(\Lambda)d\hat{\mathbf{w}} \quad (20)$$

where we have exploited the fact that φ^{-1} is linear, and with the slight abuse of notation $\varphi^{-1}(\Lambda) \in \mathbb{C}^{N \times N}$, which is the matrix obtained by applying φ^{-1} to the columns of Λ . Using the definition of φ^{-1} , Equation (20) can further be simplified into:

$$d\tilde{\mathbf{x}} = \{U\mathbf{f}(U^*\tilde{\mathbf{x}}, t) - g(t)^2\varphi^{-1}(\Lambda^2\nabla_{\varphi[\tilde{\mathbf{x}}]}\log\tilde{p}_t(\varphi[\tilde{\mathbf{x}}]))\}dt + g(t)d\check{\mathbf{w}} \quad (21)$$

where $\check{\mathbf{w}} = \varphi^{-1}(\Lambda)\hat{\mathbf{w}} = \varphi^{-1}(\Lambda\hat{\mathbf{w}})$ is a mirrored Brownian motion. Finally, notice that for any $\mathbf{y} \in \mathbb{R}^{d_x}$, we have $\varphi^{-1}(\Lambda^2\mathbf{y}) = \Lambda^2\varphi^{-1}(\mathbf{y})$, which follows from the definition of φ^{-1} . Hence, Equation (21) is equivalent to:

$$d\tilde{\mathbf{x}} = \{U\mathbf{f}(U^*\tilde{\mathbf{x}}, t) - g(t)^2\Lambda^2\tilde{\mathbf{s}}(\tilde{\mathbf{x}}, t)\}dt + g(t)d\check{\mathbf{w}} \quad (22)$$

where $\tilde{\mathbf{s}}(\tilde{\mathbf{x}}, t)$ has been defined in Equation (14). □

A.4 Denoising score matching in the frequency domain

Proposition 3.4. (*Score matching equivalence*). Consider a score $\tilde{\mathbf{s}}_{\tilde{\theta}} : \mathbb{C}^{d_x} \times [0, T] \rightarrow \mathbb{C}^{d_x}$ defined in the frequency domain and satisfying the mirror symmetry $[\tilde{\mathbf{s}}_{\tilde{\theta}}]_{\kappa} = [\tilde{\mathbf{s}}_{\tilde{\theta}}^*]_{N-\kappa}$ for all $\kappa \in [N]$. Let us define an auxiliary score $\mathbf{s}'_{\tilde{\theta}} : \mathbb{R}^{d_x} \times [0, T] \rightarrow \mathbb{R}^{d_x}$ as $(\mathbf{x}, t) \mapsto \mathbf{s}'_{\tilde{\theta}}(\mathbf{x}, t) = U^*\tilde{\mathbf{s}}_{\tilde{\theta}}(U\mathbf{x}, t)$ in the time domain. The score matching loss in the frequency domain is equivalent to the score matching loss for the auxiliary score in the time domain:

$$\mathcal{L}_{\text{SM}}(\tilde{\mathbf{s}}_{\tilde{\theta}}, \Lambda^2\tilde{\mathbf{s}}_{t|0}, \tilde{\mathbf{x}}, t) = \mathcal{L}_{\text{SM}}(\mathbf{s}'_{\tilde{\theta}}, \mathbf{s}_{t|0}, \mathbf{x}, t) \quad (10)$$

where $\tilde{\mathbf{s}}_{t|0}(\tilde{\mathbf{x}}, t) = \nabla_{\tilde{\mathbf{x}}(t)}\log\tilde{p}_{t|0}(\tilde{\mathbf{x}}(t)|\tilde{\mathbf{x}}(0))$, $\mathbf{s}_{t|0}(\mathbf{x}, t) = \nabla_{\mathbf{x}(t)}\log p_{t|0}(\mathbf{x}(t)|\mathbf{x}(0))$, and Λ is the diagonal matrix in Proposition 3.3.

Proof. Let $\Lambda \in \mathbb{R}^{N \times N}$ be the diagonal matrix such that $[\Lambda]_{\kappa, \kappa} = \begin{cases} 1 & \text{if } \kappa = 0, \text{ or } N \text{ is even and } \kappa = N/2 \\ \frac{1}{\sqrt{2}} & \text{otherwise} \end{cases}$

Step 1: We first express the score of \mathbf{x} with respect to the score of the truncation $\varphi[\tilde{\mathbf{x}}]$.

By definition of φ^{-1} in Equation (13), we have $\mathbf{x} = U^*\varphi^{-1}(\varphi[\tilde{\mathbf{x}}])$. Hence, we can write, using the change of variable formula:

$$p_{t|0}(\mathbf{x}(t)|\mathbf{x}(0)) = C \cdot \tilde{p}_{t|0}(\varphi[\tilde{\mathbf{x}}(t)]|\varphi[\tilde{\mathbf{x}}(0)]) \quad (23)$$

where C is a constant which does not depend on \mathbf{x} , since $\mathbf{x} \mapsto \varphi[U\mathbf{x}]$ is linear. Moreover, let us write $\mathbf{x} \mapsto \varphi[U\mathbf{x}]$ in matrix form, i.e. $\forall \mathbf{x} \in \mathbb{R}^{dx}$, $\varphi[U\mathbf{x}] = VU_{col}\mathbf{x} = Q\mathbf{x}$, where $V \in \mathbb{R}^{N \times 2N}$, $U_{col} = \begin{pmatrix} U_{re} \\ U_{im} \end{pmatrix}$ and Q is an invertible matrix in $\mathbb{R}^{N \times N}$. For the rest of the proof, we shall build on the below results:

Result 1. $QQ^T = \Lambda^2$. To see this, write $QQ^T = VU_{col}U_{col}^TV^T$. The matrix $U_{col}U_{col}^T$ is equal to $\begin{pmatrix} U_{re}^2 & 0_N \\ 0_N & U_{im}^2 \end{pmatrix}$ (cf. the proof to Lemma 3.1), while the multiplication by V , on the left and on the right of $U_{col}U_{col}^T$, extracts the submatrix corresponding to the indices represented by the truncature φ . Hence, $VU_{col}U_{col}^TV^T = \Lambda^2$.

Result 2. For any $\mathbf{x} \in \mathbb{R}^{dx}$, we have $Q^T\mathbf{x} = U^*\varphi^{-1}[\Lambda^2\mathbf{x}]$. To see this, notice that Result 1 implies that $Q^T\mathbf{x} = Q^{-1}\Lambda^2\mathbf{x} = U^*\varphi^{-1}[\Lambda^2\mathbf{x}]$ for all $\mathbf{x} \in \mathbb{R}^{dx}$.

Equipped with these results, we can now complete the rest of the proof. First, we have:

$$\nabla_{\mathbf{x}(t)} \log p_{t|0}(\mathbf{x}(t)|\mathbf{x}(0)) = \nabla_{\mathbf{x}(t)} \log \tilde{p}_{t|0}(\varphi[\tilde{\mathbf{x}}(t)]|\varphi[\tilde{\mathbf{x}}(0)]) \quad (24)$$

$$= Q^T \nabla_{\varphi[\tilde{\mathbf{x}}(t)]} \log \tilde{p}_{t|0}(\varphi[\tilde{\mathbf{x}}(t)]|\varphi[\tilde{\mathbf{x}}(0)]) \quad (\text{Chain rule}) \quad (25)$$

Step 2: We then obtain:

$$\begin{aligned} \mathcal{L}_{SM}(\mathbf{s}'_{\tilde{\theta}}, \mathbf{s}_{t|0}, \mathbf{x}, t) &:= \|\mathbf{s}'_{\tilde{\theta}}(\mathbf{x}, t) - \nabla_{\mathbf{x}(t)} \log p_{t|0}(\mathbf{x}(t)|\mathbf{x}(0))\|^2 \\ &= \|U\mathbf{s}'_{\tilde{\theta}}(\mathbf{x}, t) - U\nabla_{\mathbf{x}(t)} \log p_{t|0}(\mathbf{x}(t)|\mathbf{x}(0))\|^2 && (\text{Parseval identity}) \\ &= \|\tilde{\mathbf{s}}_{\tilde{\theta}}(\tilde{\mathbf{x}}, t) - UQ^T \nabla_{\varphi[\tilde{\mathbf{x}}(t)]} \log \tilde{p}_{t|0}(\varphi[\tilde{\mathbf{x}}(t)]|\varphi[\tilde{\mathbf{x}}(0)])\|^2 && (\text{Equation (25)}) \\ &= \|\tilde{\mathbf{s}}_{\tilde{\theta}}(\tilde{\mathbf{x}}, t) - UU^*\varphi^{-1}[\Lambda^2 \nabla_{\varphi[\tilde{\mathbf{x}}(t)]} \log \tilde{p}_{t|0}(\varphi[\tilde{\mathbf{x}}(t)]|\varphi[\tilde{\mathbf{x}}(0)])]\|^2 && (\text{Result 2}) \\ &= \|\tilde{\mathbf{s}}_{\tilde{\theta}}(\tilde{\mathbf{x}}, t) - \Lambda^2 \varphi^{-1}[\nabla_{\varphi[\tilde{\mathbf{x}}(t)]} \log \tilde{p}_{t|0}(\varphi[\tilde{\mathbf{x}}(t)]|\varphi[\tilde{\mathbf{x}}(0)])]\|^2 && (\text{Proposition A.1 \& Definition of } \varphi^{-1}) \\ &= \|\tilde{\mathbf{s}}_{\tilde{\theta}}(\tilde{\mathbf{x}}, t) - \Lambda^2 \tilde{\mathbf{s}}_{t|0}(\tilde{\mathbf{x}}, t)\|^2 && (\text{Equation (14)}) \\ &= \mathcal{L}_{SM}(\tilde{\mathbf{s}}_{\tilde{\theta}}, \Lambda^2 \tilde{\mathbf{s}}_{t|0}, \tilde{\mathbf{x}}, t). \end{aligned}$$

□

B Empirical details

Compute resources. All the models were trained and used for sampling on a single machine equipped with a 18-Core Intel Core i9-10980XE CPU, a NVIDIA RTX A4000 GPU and a NVIDIA GeForce RTX 3080.

B.1 Details on datasets

In this subsection, we give detailed information about the 6 datasets used throughout our experiments and the preprocessing steps for each of them.

ECG. We use two collections of heartbeat signals, from the MIT-BIH Arrhythmia Dataset and the PTB Diagnostic ECG Database [29]. No preprocessing was performed on this dataset.

MIMIC-III. MIMIC-III [30] is a database consisting of deidentified records for patients who were in critical unit care units. *Preprocessing.* We use the "vitals labs" table of the database, which corresponds to time-varying vitals and labs. We extract the rows of the dataset which correspond to the first 24 hours of stay by using *MIMIC-Extract* [39]. The features are then standardized across all times and patients. We also perform imputation to handle missing values in the dataset. To do so, we consider the mean features (average measurement over 1 hour). For each patient, and missing value, we propagate the last observation forward if this is possible. If not, we fill the missing value with the mean value for the patient (which is computed over the whole stay). If no mean value is available, we fill the entry with 0.

NASDAQ-2019. This dataset [31] contains daily prices for tickers trading on NASDAQ, and contains prices for up to 1st of April 2020. *Preprocessing.* We considered one year of daily prices from 1st of January 2019 to 1st of January 2020. Each sample corresponds to one stock, and we remove the stocks which are not active in this whole time interval, or contain missing values.

NASA battery. The NASA battery dataset [32] consists of profiles for Li-on batteries, under charge and discharge. *Preprocessing.* For both the charge and discharge datasets, we bin the time values (bins of size 10 for Charge, 15 for Discharge) and compute the mean of each feature inside each bin.

US-Droughts. This dataset [33] consists of drought levels in different US counties, from 2000 to 2020. *Preprocessing.* We consider one year of history, from 1st of January 2011 to 1st of January 2012, and drop the columns with missing values.

B.2 Details on evaluation

Sliced Wasserstein distances. The sliced Wasserstein distance [35] is a metric which can handle high-dimensional distributions. It is motivated by the fact that the Wasserstein distance is easy to compute when comparing two one-dimensional distributions. The idea of the sliced Wasserstein distance is to map the high-dimensional distributions of interest to one-dimensional distributions, by considering random projections on vectors of the unit sphere. For two distributions μ_1 and μ_2 , it can be written as:

$$SW_p(\mu_1, \mu_2) := \int_{\mathbb{S}^{d-1}} W_p(P_u \# \mu_1, P_u \# \mu_2) du \quad (26)$$

where \mathbb{S}^{d-1} is the unit sphere in dimension d , $P_u(x) = u \cdot x$ denotes the projection of x on u , $P_u \# \mu$ is the push-forward of μ by P_u , and W_p is the Wasserstein distance of order p . To estimate this quantity in practice, we sample $n = 10,000$ random vectors $\{u_i | i \in [n]\}$ which follow a uniform distribution in \mathbb{S}^{d-1} and consider $p = 2$. Hence, we can approximate SW_p by the Monte-carlo estimator:

$$\hat{SW}_p(\mu_1, \mu_2) = \frac{1}{n} \sum_{i=1}^n W_p(P_{u_i} \# \mu_1, P_{u_i} \# \mu_2) \quad (27)$$

Marginal Wasserstein distances. In addition to the sliced Wasserstein distance, we also consider the marginal Wasserstein distance. For any $j \in \{1, \dots, d\}$, the j -th marginal Wasserstein distance is defined as:

$$MW_p^{(j)}(\mu_1, \mu_2) = W_p(P_{e_j} \# \mu_1, P_{e_j} \# \mu_2) \quad (28)$$


where e_j is the j -th vector of the standard basis of \mathbb{R}^d . Throughout our experiments in Section 4, we compute the Wasserstein distances with respect to $\mathcal{D}_{\text{train}}$ and $\tilde{\mathcal{D}}_{\text{train}}$.

B.3 Additional plots

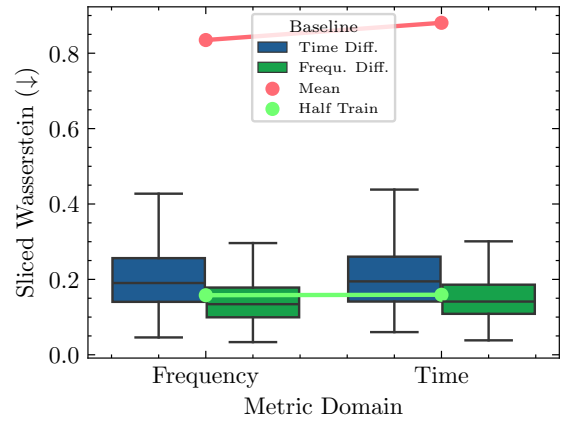
Sliced Wasserstein distances. In Figure 4, we show the distribution of the sliced Wasserstein distances over all slices. In addition, we have included the average sliced Wasserstein distances obtained with 2 baselines. The first baseline is simply the Wasserstein distance between the training set and a set of sample only containing identical copies the average sample $SW(\mathcal{D}_{\text{train}}, \mathcal{S}_{\text{mean}})$, where $\mathcal{S}_{\text{mean}} = \{\mathbb{E}_{X \sim U(\mathcal{D}_{\text{train}})}[X]\}$. It represents the performance of a dummy generator that only generates the average time series and is denoted by *mean* in Figure 4. The second baseline is the Wasserstein distance between two random splits of the training set $SW(\mathcal{D}_{\text{train}}^{1/2}, \mathcal{D}_{\text{train}}^{2/2})$, where $\mathcal{D}_{\text{train}} = \mathcal{D}_{\text{train}}^{1/2} \sqcup \mathcal{D}_{\text{train}}^{2/2}$ is a decomposition of the training set into two disjoint random splits of equal size $|\mathcal{D}_{\text{train}}^{1/2}| = |\mathcal{D}_{\text{train}}^{2/2}|$. It represents the distance between two samples from the ground-truth distribution and is denoted by *self* in Figure 4. As we can observe, both the time and frequency diffusion models substantially outperform the *mean* baseline (as expected) and perform on par with the *self* baseline. This indicates that the models learned a good approximation of the real distribution. Furthermore, we notice that the frequency diffusion models tend to have smaller quantiles than the time diffusion models. This confirms that frequency diffusion models outperform the time diffusion models, as discussed in Section 4.

Marginal Wasserstein distances. In Figure 5, we show the distribution of the marginal Wasserstein distances over all slices. In addition, we have included the average marginal Wasserstein distances obtained with the 2 baselines defined in the previous paragraph. Again, both the time and frequency diffusion models tend to outperform the *mean* baseline (as expected) and perform on par with the *self* baseline. Furthermore, we notice that the frequency diffusion models tend to have smaller quantiles than the time diffusion models. This is consistent with the observations made in the above paragraph.

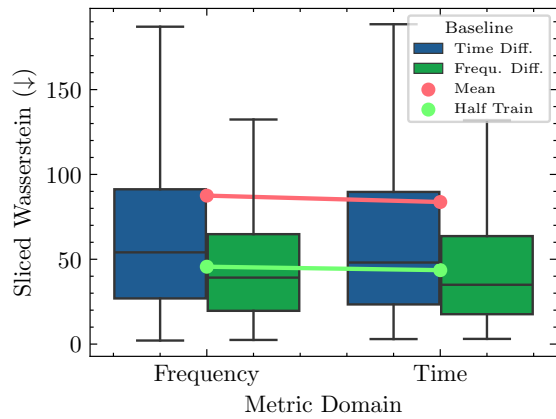
Per-sample localization. In Figure 6, we observe the distribution of our localization metrics Δ_{time} and Δ_{sigma} for each sample and each dataset from Section 4. We notice that most samples are located below the $y = x$ axis, which confirms the fact that most samples are more localized in the frequency domain. Interestingly, we also observe that none of the samples is located close to the origin. This confirms the uncertainty theorem from [36].



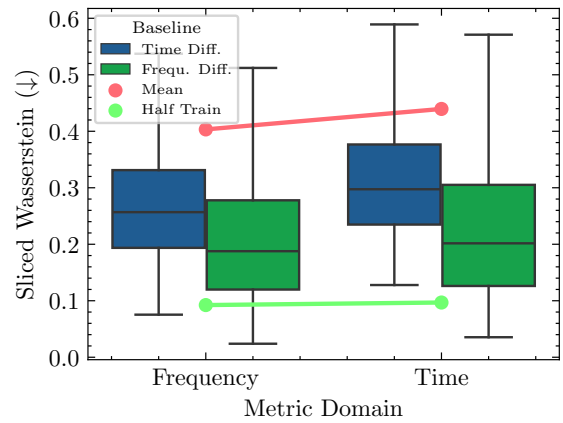
(a) ECG.



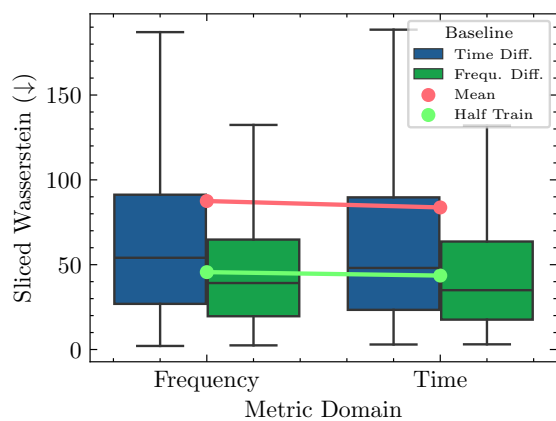
(b) MIMIC-III.




(c) NASDAQ-2019.



(d) NASA-Charge.



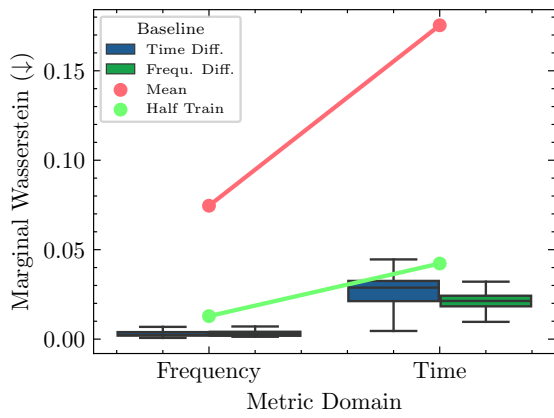
(e) NASA-Discharge.




(f) US-Droughts.

Figure 4: Sliced Wasserstein distances of time and frequency diffusion models.

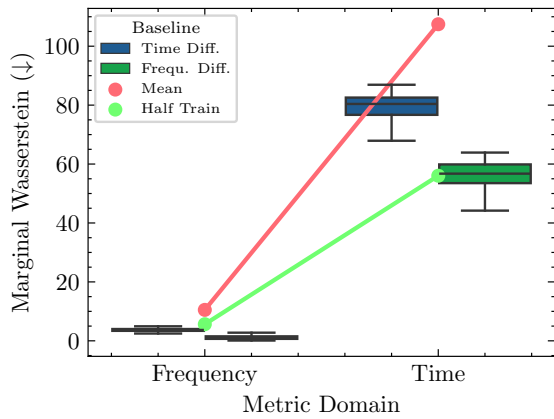
Time Series Diffusion in the Frequency Domain



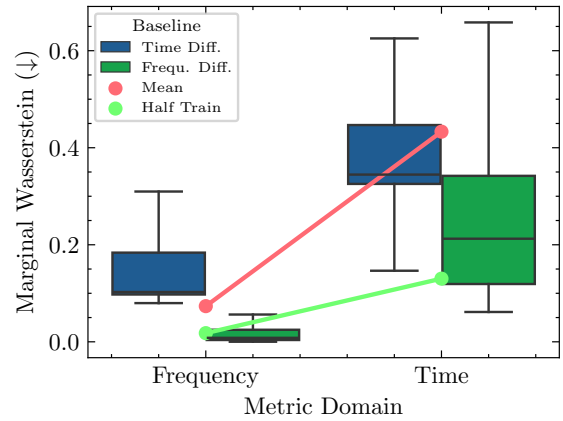
(a) ECG.



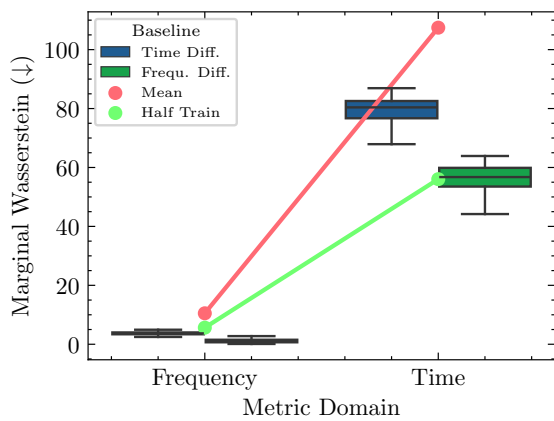
(b) MIMIC-III.




(c) NASDAQ-2019.



(d) NASA-Charge.



(e) NASA-Discharge.



(f) US-Droughts.

Figure 5: Marginal Wasserstein distances of time and frequency diffusion models.




Figure 6: Localization metrics Δ_{time} and Δ_{freq} for all the samples of all datasets. We observe that no sample has a high localization (i.e. low Δ) in the time and frequency domain simultaneously.

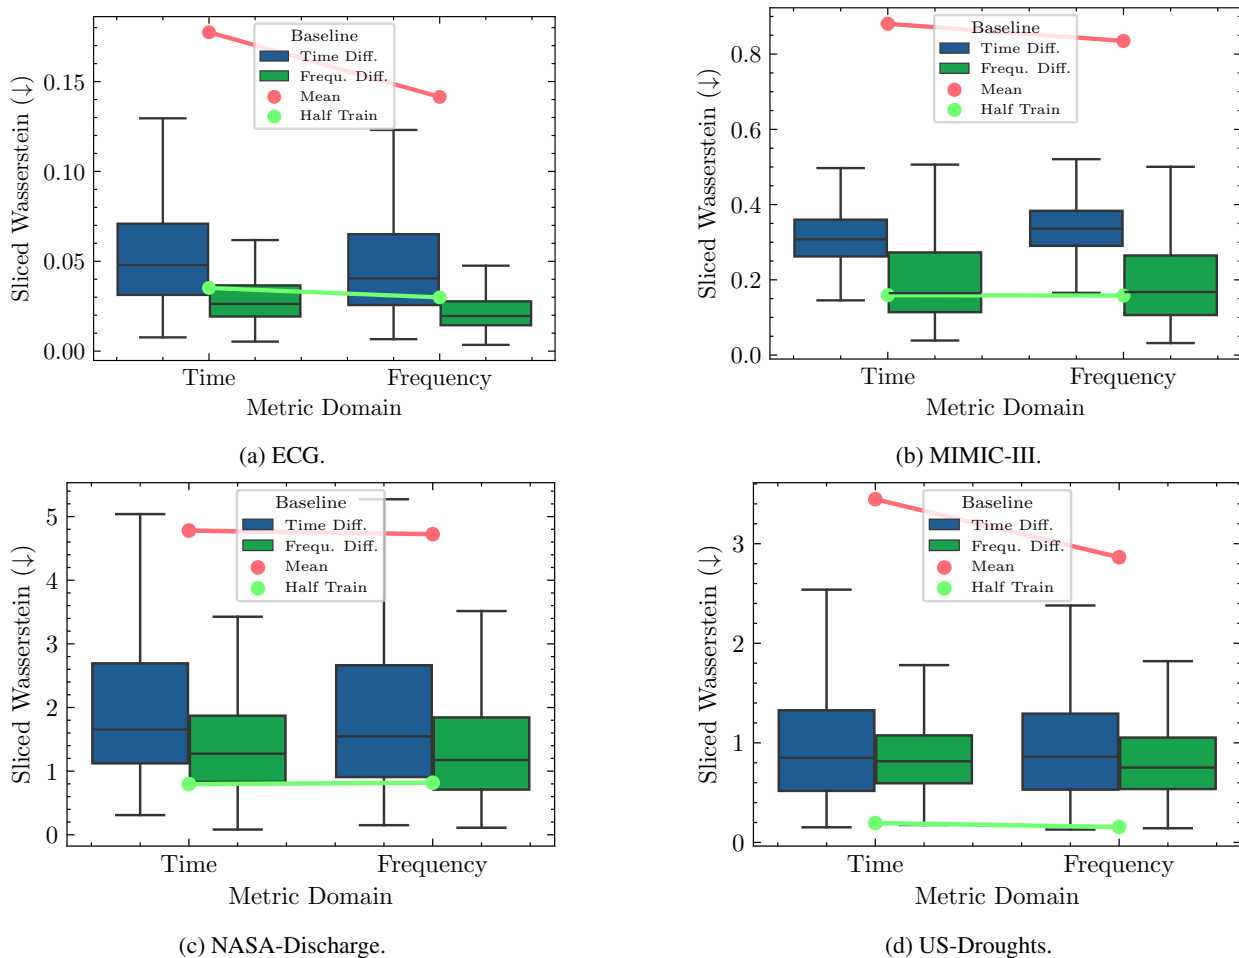


Figure 7: Sliced Wasserstein distances of time and frequency LSTM models.

C Alternative backbone

LSTM models. For each dataset, we try an alternative parametrization of the time score model s_θ and the frequency score model \tilde{s}_θ as LSTM encoders with 10 layers, each with dimension $d_{\text{model}} = 72$. Both models have diffusion time t encoding through random Fourier features composed with a learnable dense layer. This results in models with 427k parameters. The data is noised by using a VP-SDE, as in [19]. The score models are trained with the denoising score-matching loss, as defined in Section 3. All the models are trained for 200 epochs with batch size 64, AdamW optimizer and cosine learning rate scheduling (20 warmup epochs, $\text{lr}_{\text{max}} = 10^{-3}$). The selected model achieves the lowest validation loss.

Sliced Wasserstein distances. In Figure 7, we show the distribution of the sliced Wasserstein distances over all slices for the LSTM models. In addition, we have included the average sliced Wasserstein distances obtained with 2 baselines defined in Appendix B.3. As observed for the transformer models, both the time and frequency diffusion models substantially outperform the *mean* baseline (as expected) and perform on par with the *self* baseline. This indicates that the models learned a good approximation of the real distribution. Furthermore, we notice that the frequency diffusion models tend to have smaller quantiles than the time diffusion models. This confirms that frequency diffusion models outperform the time diffusion models, as observed for the transformer models in Section 4. We note that the Nasa-Charge and the NASDAQ-2019 are absent from Figure 7. This is because we did not manage to obtain diffusion models performing better than the *mean* baseline for these datasets, hence leading to non informative comparisons between models with poor performances.

Other attempts. In order to minimize the inductive bias in our models, we also tried to train diffusion models with simple feed-forward neural networks. Unfortunately, this attempt was unsuccessful and resulted in models performing worse than the *mean* baseline in each case. This emphasizes the value of incorporating inductive biases in time series diffusion models.

D Sample visualization

In Figures 8 to 12, we visualize a few examples generated by each diffusion model, along with ground-truth training examples. We do not include samples from the MIMIC-III dataset in accordance with the dataset licence. We observe that the frequency diffusion models generate samples that are substantially less noisy than the ones generated by time diffusion models. All the generated samples resemble training samples, with the only exception of the NASDAQ-2019 dataset. The models appear to be struggling with the high correlation between the different features.

Time Series Diffusion in the Frequency Domain



Figure 8: Samples for the ECG dataset. The y axis corresponds to the different features, while the x axis corresponds to the different time steps.

Time Series Diffusion in the Frequency Domain

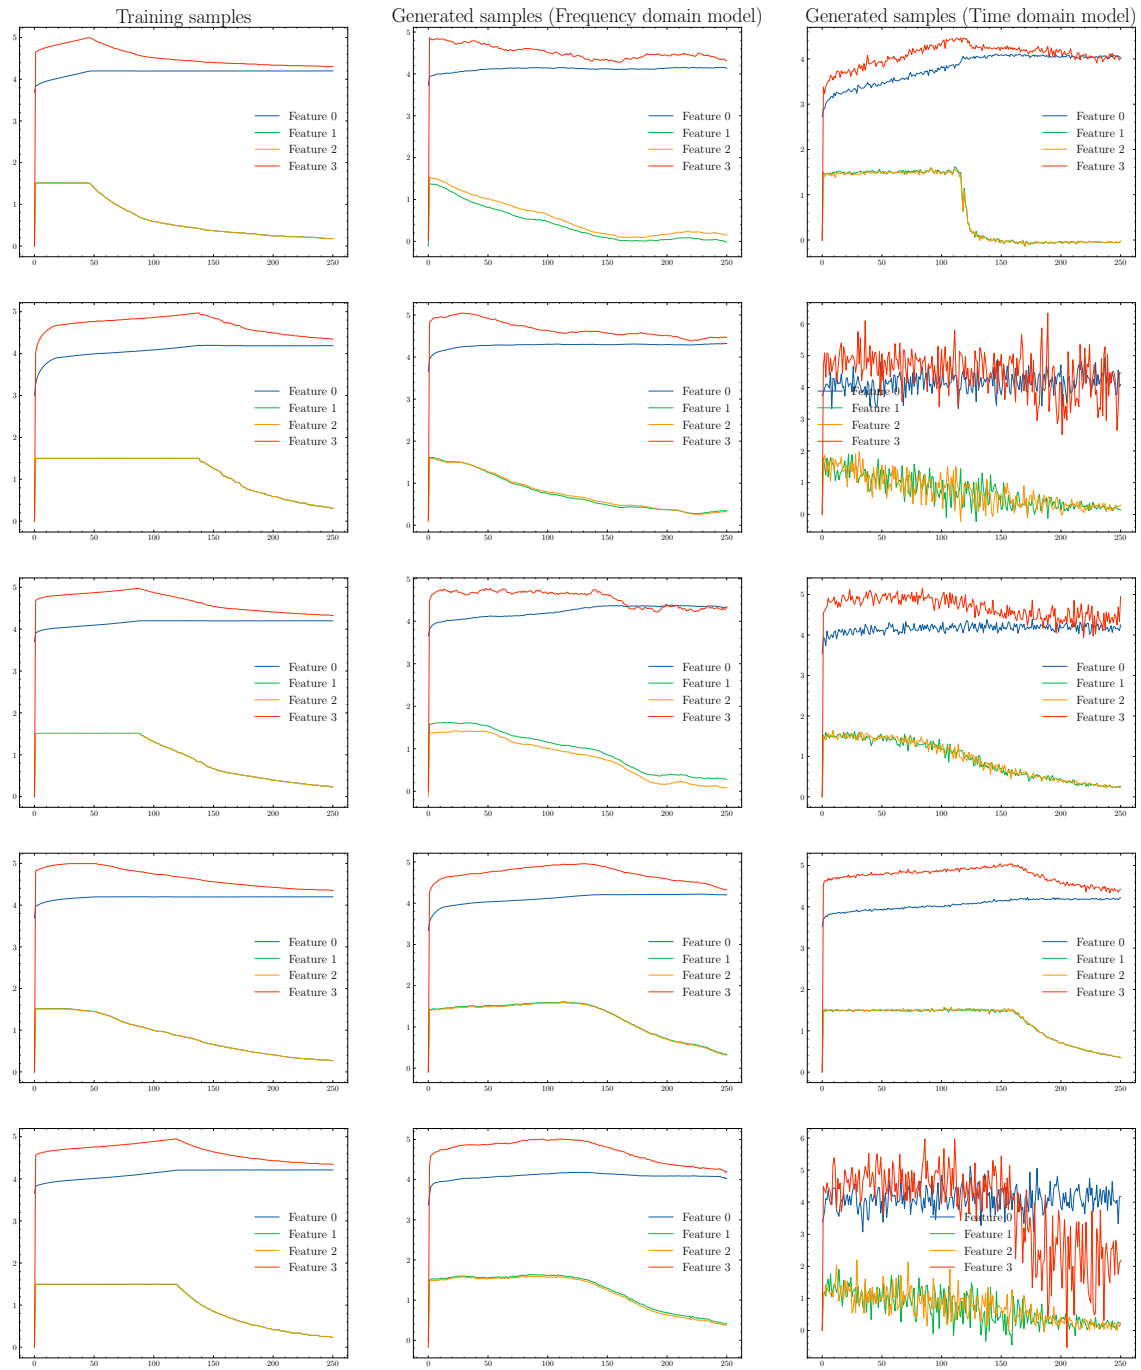


Figure 9: Samples for the NASA-Charge dataset. The y axis corresponds to the different features, while the x axis corresponds to the different time steps.

Time Series Diffusion in the Frequency Domain



Figure 10: Samples for the NASA-Discharge dataset. The y axis corresponds to the different features, while the x axis corresponds to the different time steps.

Time Series Diffusion in the Frequency Domain

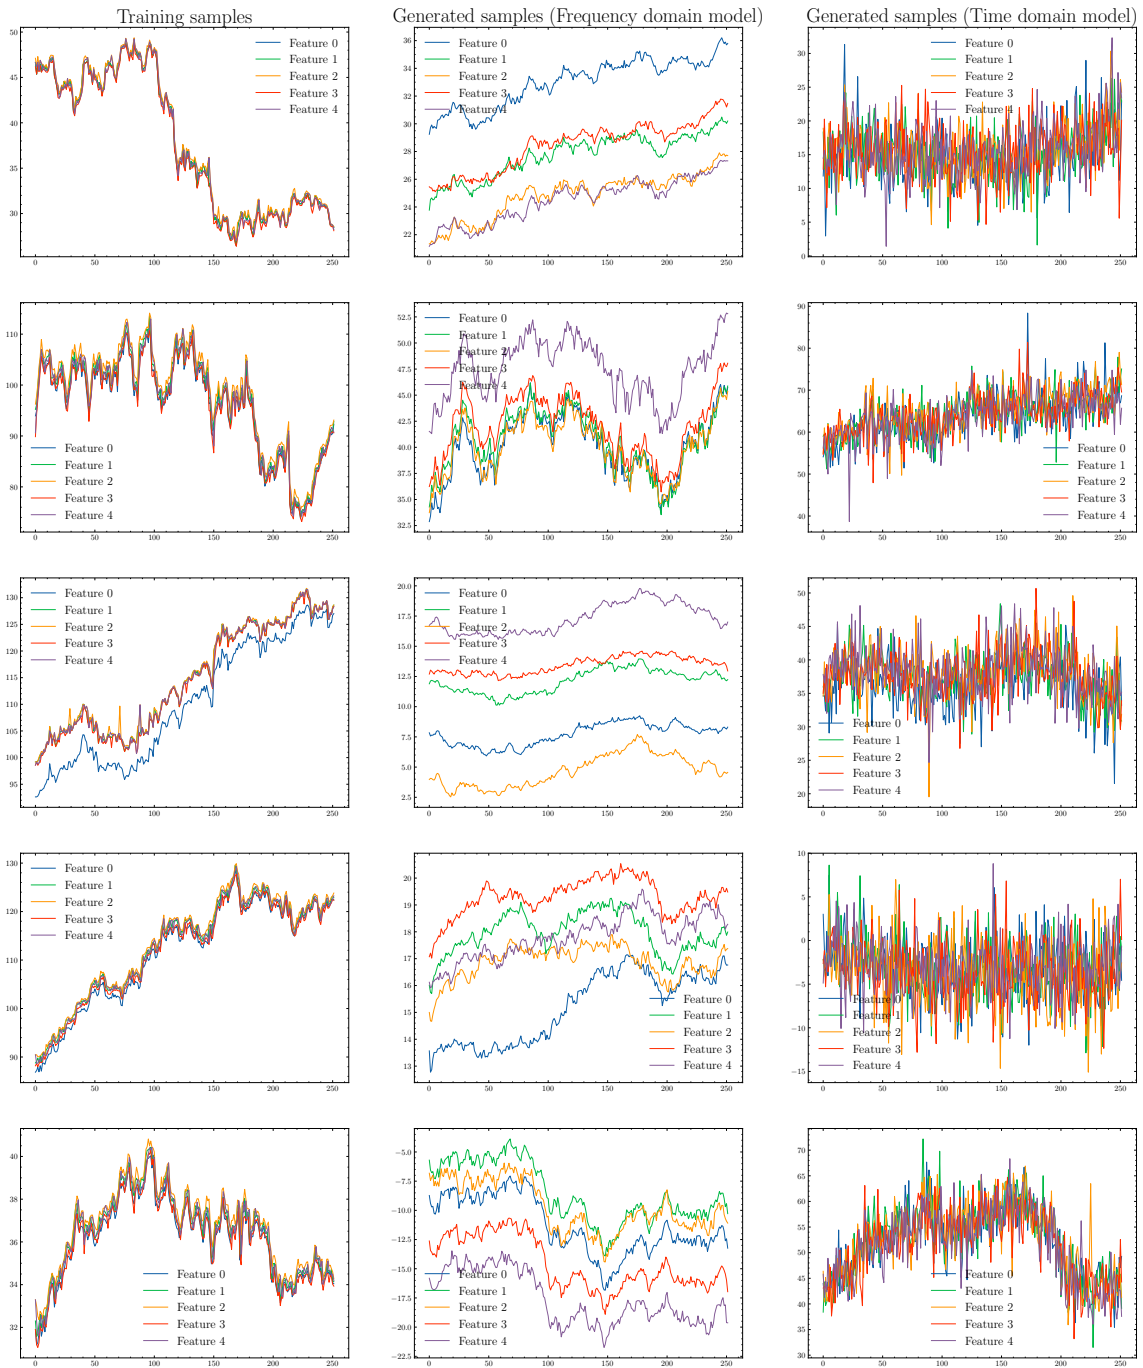


Figure 11: Samples for the NASDAQ-2019 dataset. The y axis corresponds to the different features, while the x axis corresponds to the different time steps.

Time Series Diffusion in the Frequency Domain




Figure 12: Samples for the US-Droughts dataset, represented as heatmaps. The y axis corresponds to the different features, while the x axis corresponds to the different time steps.