DENOISING DIFFUSION SAMPLERS

Anonymous authors
Paper under double-blind review

ABSTRACT

Denoising diffusion models are a popular class of generative models providing state-of-the-art results in many domains. One adds gradually noise to data using a diffusion to transform the data distribution into a Gaussian distribution. Samples from the generative model are then obtained by simulating an approximation of the time-reversal of this diffusion initialized by Gaussian samples. Practically, the intractable score terms appearing in the time-reversed process are approximated using score matching techniques. We explore here a similar idea to sample approximately from unnormalized probability density functions and estimate their normalizing constants. We consider a process where the target density diffuses towards a Gaussian. Denoising Diffusion Samplers (DDS) are obtained by approximating the corresponding time-reversal. While score matching is not applicable in this context, we can leverage many of the ideas introduced in generative modeling for Monte Carlo sampling. Existing theoretical results from denoising diffusion models also provide theoretical guarantees for DDS. We discuss the connections between DDS, optimal control and Schrödinger bridges and finally demonstrate DDS experimentally on a variety of challenging sampling tasks.

1 INTRODUCTION

Let π be a probability density on $\mathbb{R}^d$ of the form

$$\pi(x) = \frac{\gamma(x)}{Z}, \qquad Z = \int_{\mathbb{R}^d} \gamma(x)\mathrm{d}x, \tag{1}$$

where $\gamma : \mathbb{R}^d \to \mathbb{R}^+$ can be evaluated pointwise but the normalizing constant Z is intractable. We are here interested in both estimating Z and obtaining approximate samples from π.

A large variety of Monte Carlo techniques has been developed in the literature to address this problem. In particular Annealed Importance Sampling (AIS) (Neal, 2001) and its Sequential Monte Carlo (SMC) extensions (Del Moral et al., 2006) are often regarded as the "gold standard" approach to compute normalizing constants. These methods rely on the introduction of a sequence of annealed versions of π which are explored using a combination of Markov chain Monte Carlo (MCMC), importance sampling and resampling steps.

Variational techniques are a popular alternative to MCMC and SMC where one considers a flexible family of easy-to-sample distributions q^θ whose parameters are optimized by minimizing a suitable metric w.r.t. π, typically the reverse Kullback–Leibler discrepancy $\mathrm{KL}(q^\theta||\pi)$. Typical choices for q^θ include mean-field approximation (Wainwright & Jordan, 2008) or more sophisticated techniques such as normalizing flows (Papamakarios et al., 2019). To be able to model complex variational distributions, it has often proven useful to model $q^\theta(x)$ as the marginal of an auxiliary extended distribution; i.e. $q^\theta(x) = \int q^\theta(x, u)\mathrm{d}u$. As this marginal is typically intractable, θ is then learned by minimizing a discrepancy measure between $q^\theta(x, u)$ and an extended target distribution $p^\theta(x, u) = \pi(x)p^\theta(u|x)$ where $p^\theta(u|x)$ is an auxiliary conditional distribution (Agakov & Barber, 2004).

Over recent years, variational techniques have also been used fruitfully to optimize the algorithmic parameters of Monte Carlo techniques. For example, AIS can be thought of a procedure where $q^\theta(x, u)$ is the joint distribution of a Markov chain defined by a sequence of MCMC kernels whose final state is x while $p^\theta(x, u)$ is the corresponding AIS extended target distribution (Neal, 2001). The parameters θ of these kernels can then be optimized by minimizing $\mathrm{KL}(q^\theta||p^\theta)$ using stochastic

gradient (Wu et al., 2020; Geffner & Domke, 2021; Thin et al., 2021; Zhang et al., 2021; Doucet et al., 2022; Geffner & Domke, 2022).

In this work, we follow an approach inspired by Denoising Diffusion Probabilistic Models (DDPM), a powerful class of generative models (Sohl-Dickstein et al., 2015; Ho et al., 2020; Song et al., 2021b). In this context, one adds noise progressively to data using a diffusion to transform the complex data distribution into a Gaussian distribution. The time-reversal of this diffusion can then be used to transform a Gaussian sample into a sample from the target. While superficially similar to Langevin dynamics, this process mixes fast even in high dimensions as it inherits the mixing properties of the forward diffusion (De Bortoli et al., 2021). However, as the time-reversal involves the logarithmic derivatives of the intractable marginal densities of the forward diffusion, these so-called scores are practically approximated using score matching techniques. If the score estimation error is small, the approximate time-reversal still enjoys remarkable theoretical properties (De Bortoli, 2022; Chen et al., 2022; Lee et al., 2022).

These results motivate us to introduce Denoising Diffusion Samplers (DDS). Like DDPM, we consider a forward diffusion which progressively transforms the target π into a Gaussian distribution. This defines an extended target distribution $p(x, u) = \pi(x)p(u|x)$. DDS are obtained by approximating the time-reversal of this diffusion using a process of distribution $q^\theta(x, u)$. What distinguishes DDS from DDPM is that we cannot simulate sample paths from the diffusion we want to time-reverse, as its initial state x is distributed according to π. Hence score matching ideas cannot be used to approximate the score terms and we focused here on minimizing $\mathrm{KL}(q^\theta||p)$ as in variational inference. However, we show how one can still exploit the many similarities between DDS and DDPM to leverage some of the ideas developed in generative modeling for Monte Carlo sampling. This includes using the probability flow ordinary differential equation (ODE) (Song et al., 2021b) to derive novel normalizing flows and the use underdamped diffusions as a forward noising diffusion (Dockhorn et al., 2022). Theoretical results for DDPM also provide theoretical guarantees for DDS.

2 Denoising Diffusion Samplers: Continuous Time

We start here by formulating DDS in continuous-time to gain insight on the structure of the time-reversal we want to approximate. Numerical integrators are discussed in the following section.

2.1 Forward diffusion and its time-reversal

Consider the forward noising diffusion given by an Ornstein–Uhlenbeck process[1]

$$\mathrm{d}x_t = -\beta_t x_t \mathrm{d}t + \sigma\sqrt{2\beta_t}\mathrm{d}B_t, \qquad x_0 \sim \pi, \tag{2}$$

where $(B_t)_{t\in[0,T]}$ is a d-dimension Brownian motion and $t \to \beta_t$ is a non-decreasing positive function. This diffusion induces the path measure $\mathcal{P}$ on the time interval $[0, T]$ and the marginal density of x_t is denoted p_t. The transition density of this diffusion is given by $p_{t|0}(x_t|x_0) = \mathcal{N}(x_t; \sqrt{1-\lambda_t}x_0, \sigma^2\lambda_t I)$ where $\lambda_t = 1 - \exp\left(-2\int_0^t \beta_s \mathrm{d}s\right)$. From now on, we will always consider a scenario where $\int_0^T \beta_s \mathrm{d}s \gg 1$ so that $\alpha_T \approx 1$ and $p_T(x_T) \approx \mathcal{N}(x_T; 0, \sigma^2 I)$. We can thus think of (2) as transporting approximately the target density π to this Gaussian density.

From (Haussmann & Pardoux, 1986), its time-reversal $(y_t)_{t\in[0,T]} = (x_{T-t})_{t\in[0,T]}$ satisfies

$$\mathrm{d}y_t = \beta_{T-t}\Big\{y_t + 2\sigma^2\nabla \ln p_{T-t}(y_t)\Big\}\mathrm{d}t + \sigma\sqrt{2\beta_{T-t}}\mathrm{d}W_t, \qquad y_0 \sim p_T, \tag{3}$$

where $(W_t)_{t\in[0,T]}$ is another d-dimensional Brownian motion. By definition this time-reversal starts from $y_0 \sim p_T(y_0) \approx \mathcal{N}(y_0; 0, \sigma^2 I)$ and is such that $y_T \sim \pi$. This suggests that if we could approximately simulate the diffusion (3), then we would obtain approximate samples from π.

However, putting this idea in practice requires being able to approximate the intractable scores $(\nabla \ln p_t(x))_{t\in[0,T]}$. To achieve this, DDPM rely on score matching techniques (Hyvärinen, 2005; Vincent, 2011) as it is easy to sample from (2). This is impossible in our scenario as sampling from (2) requires sampling $x_0 \sim \pi$ which is impossible by assumption.

[1]This is referred to as a Variance Preserving diffusion by Song et al. (2021b).

2.2 REFERENCE DIFFUSION AND VALUE FUNCTION

In our context, it will useful to introduce a reference process defined by the diffusion following (2) but initialized at $p_0^{\text{ref}}(x_0) = \mathcal{N}(x_0; 0, \sigma^2 I)$ rather than $\pi(x_0)$ thus ensuring that the marginals of the resulting path measure $\mathcal{P}^{\text{ref}}$ all satisfy $p_t^{\text{ref}}(x_t) = \mathcal{N}(x_t; 0, \sigma^2 I)$. From the chain rule for KL, one can easily verify that $\mathcal{P}$ can be identified as the path measure minimizing the reverse (or forward) KL discrepancy w.r.t. $\mathcal{P}^{\text{ref}}$ over the set of path measures $\mathcal{Q}$ with marginal $q_0(x_0) = \pi(x_0)$ at time $t = 0$, that is $\mathcal{P} = \arg\min_{\mathcal{Q}} \Big\{ \text{KL}(\mathcal{Q}||\mathcal{P}^{\text{ref}}) : q_0 = \pi \Big\}$.

A time-reversal representation of $\mathcal{P}^{\text{ref}}$ is given by $(y_t)_{t\in[0,T]} = (x_{T-t})_{t\in[0,T]}$ satisfying

$$\mathrm{d}y_t = -\beta_{T-t} y_t \mathrm{d}t + \sigma\sqrt{2\beta_{T-t}}\mathrm{d}W_t, \qquad y_0 \sim p_0^{\text{ref}}, \tag{4}$$

as $y_t + 2\sigma^2 \nabla \ln p_{T-t}^{\text{ref}}(y_t) = -\beta_{T-t} y_t$ and we can rewrite the time-reversal (3) of $\mathcal{P}$ as

$$\mathrm{d}y_t = -\beta_{T-t}\Big\{ y_t - 2\sigma^2 \nabla \ln \phi_{T-t}(y_t) \Big\}\mathrm{d}t + \sigma\sqrt{2\beta_{T-t}}\mathrm{d}W_t, \qquad y_0 \sim p_T, \tag{5}$$

where $\phi_t(x) = p_t(x)/p_t^{\text{ref}}(x)$ is a so-called value function which satisfies a Kolmogorov equation such that $\phi_t(x_t) = \mathbb{E}_{\mathcal{P}^{\text{ref}}}[\phi_0(x_0)|x_t]$, the expectation being w.r.t. the law of the reference process.

2.3 LEARNING THE TIME-REVERSAL

To approximate the time-reversal (3) of $\mathcal{P}$, consider a path measure $\mathcal{Q}^\theta$ whose time-reversal is induced by

$$\mathrm{d}y_t = \beta_{T-t}\Big\{ y_t + 2\sigma^2 s_\theta(T-t, y_t) \Big\}\mathrm{d}t + \sigma\sqrt{2\beta_{T-t}}\mathrm{d}W_t, \qquad y_0 \sim \mathcal{N}(0, \sigma^2 I), \tag{6}$$

so that $y_t \sim q_{T-t}^\theta$. To obtain $s_\theta(t,x) \approx \nabla \ln p_t(x)$, we parameterize $s_\theta(t,x)$ by a neural network whose parameters are obtained by minimizing

$$\begin{aligned} \text{KL}(\mathcal{Q}^\theta||\mathcal{P}) &= \text{KL}(p_T||\mathcal{N}(0,\sigma^2 I)) + \mathbb{E}_{y_0\sim\mathcal{N}(0,\sigma^2 I)}[\text{KL}(\mathcal{Q}^\theta(\cdot|y_0)||\mathcal{P}(\cdot|y_0))] \\ &= \text{KL}(p_T||\mathcal{N}(0,\sigma^2 I)) + \sigma^2 \mathbb{E}_{\mathcal{Q}^\theta}\Big[\int_0^T \beta_{T-t}||s_\theta(T-t,y_t) - \nabla \ln p_{T-t}(y_t)||^2 \mathrm{d}t\Big], \end{aligned} \tag{7}$$

where we have use the chain rule for KL then Girsanov's formula; see e.g. (Klebaner, 2012). This expression of the KL is reminiscent of the expression obtained by (Song et al., 2021a, Theorem 1) in the context of DDPM. However, the expectation appearing in (7) is here w.r.t. $\mathcal{Q}^\theta$ and not w.r.t. $\mathcal{P}$ and we cannot get rid of the intractable scores $(\nabla \ln p_t(x)))_{t\in[0,T]}$ using integration by parts.

Instead, taking inspiration from (5), we reparameterize the time-reversal of $\mathcal{Q}^\theta$ using

$$\mathrm{d}y_t = -\beta_{T-t}\Big\{ y_t - 2\sigma^2 f_\theta(T-t, y_t) \Big\}\mathrm{d}t + \sigma\sqrt{2\beta_{T-t}}\mathrm{d}W_t, \qquad y_0 \sim \mathcal{N}(0, \sigma^2 I), \tag{8}$$

i.e. $f_\theta(t,x) = s_\theta(t,x) - \nabla \ln p_t^{\text{ref}}(x) = s_\theta(t,x) + x/\sigma^2$. This reparameterization allows us to express $\text{KL}(\mathcal{Q}^\theta||\mathcal{P})$ in a compact form.

Proposition 1. *The Radon-Nikodym derivative* $\frac{\mathrm{d}\mathcal{Q}^\theta}{\mathrm{d}\mathcal{P}^{\text{ref}}}(y_{[0,T]})$ *satisfies under* $\mathcal{Q}^\theta$

$$\ln\left(\frac{\mathrm{d}\mathcal{Q}^\theta}{\mathrm{d}\mathcal{P}^{\text{ref}}}\right) = \sigma^2 \int_0^T \beta_{T-t}||f_\theta(T-t,y_t)||^2 \mathrm{d}t + \sigma \int_0^T \sqrt{2\beta_{T-t}} f_\theta(T-t,y_t)^\top \mathrm{d}W_t. \tag{9}$$

From the identity $\text{KL}(\mathcal{Q}^\theta||\mathcal{P}) = \text{KL}(\mathcal{Q}^\theta||\mathcal{P}^{\text{ref}}) + \mathbb{E}_{y_T\sim q_0^\theta}\Big[\ln\Big(\frac{p_0^{\text{ref}}(y_T)}{p_0(y_T)}\Big)\Big]$, *it follows that*

$$\text{KL}(\mathcal{Q}^\theta||\mathcal{P}) = \mathbb{E}_{\mathcal{Q}^\theta}\Big[\sigma^2 \int_0^T \beta_{T-t}||f_\theta(T-t,y_t)||^2 \mathrm{d}t + \ln\Big(\frac{\mathcal{N}(y_T;0,\sigma^2 I)}{\pi(y_T)}\Big)\Big]. \tag{10}$$

For θ minimizing this KL, approximate samples from π can thus be obtained by simulating (8) and returning $y_T \sim q_0^\theta$. We can also obtain an unbiased estimate of Z by using the following importance sampling identity $\hat{Z} = (\gamma(y_T)/\mathcal{N}(y_T;0,\sigma^2 I))\frac{\mathrm{d}\mathcal{P}^{\text{ref}}}{\mathrm{d}\mathcal{Q}^\theta}(y_{[0,T]})$ for $y_{[0,T]} \sim \mathcal{Q}^\theta$ where the second term can be computed directly from (9).

2.4 CONTINUOUS-TIME NORMALIZING FLOW

To approximate the log likelihood of DDPM, it was proposed by Song et al. (2021b) to rely on the probability flow ODE. This is an ODE admitting the same marginal distributions $(p_t)_{t\in[0,T]}$ as the noising diffusion given by

$$\mathrm{d}x_t = -\beta_t\Big\{x_t + \sigma^2\nabla \ln p_t(x_t)\Big\}\mathrm{d}t. \tag{11}$$

We use here this ODE to design a continuous-time normalizing flow to sample from π. For θ such that $f_\theta(t,x) \approx \nabla \ln \phi_t(x)$ (obtained by minimization of the KL in (10)), we have $x_t + \sigma^2\nabla \ln p_t(x) \approx \sigma^2 f_\theta(t,x)$. So it is possible to sample approximately from π by integrating the following ODE over $[0,T]$

$$\mathrm{d}y_t = \sigma^2\beta_{T-t} f_\theta(T-t,y_t)\mathrm{d}t, \qquad y_0 \sim \mathcal{N}(0,\sigma^2 I). \tag{12}$$

If denote by $\bar{q}_{T-t}^\theta$ the distribution of y_t then $y_T \sim \bar{q}_0^\theta$ is an approximate sample from π. We can use this normalizing flow to obtain an unbiased estimate of Z using importance sampling $\hat{Z} = \gamma(y_T)/\bar{q}_0^\theta(y_T)$ for $y_T \sim \bar{q}_0^\theta$. Indeed, contrary to the proposal q_0^θ induced by (8), we can compute pointwise $\bar{q}_0^\theta$ using the instantaneous change of variables formula (Chen et al., 2018)

$$\ln \bar{q}_0^\theta(y_T) = \ln \mathcal{N}(y_0;0,\sigma^2 I) - \sigma^2\int_0^T \beta_{T-t}\nabla\cdot f_\theta(T-t,y_t)\mathrm{d}t \tag{13}$$

where $(y_t)_{t\in[0,T]}$ arises from the ODE (12). The expensive divergence term can be estimated using the Hutchinson estimator (Grathwohl et al., 2018; Song et al., 2021b).

2.5 EXTENSION TO UNDERDAMPED LANGEVIN DYNAMICS

For DDPM, it has been proposed to extend the original state $x \in \mathbb{R}^d$ by a momentum variable $p \in \mathbb{R}^d$. One then diffuses the data distribution augmented by a Gaussian distribution on the momentum using an underdamped Langevin dynamics targeting $\mathcal{N}(x;0,\sigma^2 I)\mathcal{N}(p;0,M)$ where M is a positive definite mass matrix (Dockhorn et al., 2022). It was demonstrated empirically that the resulting scores are smoother, hence easier to estimate and this leads to improved performance. We adapt here this approach to Monte Carlo sampling; see Section B for more details.

We diffuse $\bar{\pi}(x,p) = \pi(x;0,\sigma^2 I)\mathcal{N}(p;0,M)$ using an underdamped Langevin dynamics, i.e.

$$\mathrm{d}x_t = M^{-1}p_t\mathrm{d}t, \quad \mathrm{d}p_t = -\frac{x_t}{\sigma^2}\mathrm{d}t - \beta_t p_t\mathrm{d}t + \sqrt{2\beta_t}M^{1/2}\mathrm{d}B_t, \quad (x_0,p_0)\sim\bar{\pi}. \tag{14}$$

The resulting path measure on $[0,T]$ is denoted $\mathcal{P}$ and the marginal of (x_t,p_t) is denoted η_t.

Here, the reference process $\mathcal{P}^{\text{ref}}$ is defined by the diffusion (14) initialized using $x_0 \sim \mathcal{N}(0,\sigma^2 I), p_0 \sim \mathcal{N}(0,M)$. This ensures that $\eta_t^{\text{ref}}(x_t,p_t) = \mathcal{N}(x_t;0,\sigma^2 I)\mathcal{N}(p_t;0,M)$ for all t and the time-reversal process $(y_t,q_t)_{t\in[0,T]} = (x_{T-t},p_{T-t})_{t\in[0,T]}$ of this stationary diffusion satisfies

$$\mathrm{d}y_t = -M^{-1}q_t\mathrm{d}t, \quad \mathrm{d}q_t = \frac{y_t}{\sigma^2}\mathrm{d}t - \beta_{T-t}q_t\mathrm{d}t + \sqrt{2\beta_{T-t}}M^{1/2}\mathrm{d}W_t. \tag{15}$$

Using manipulations identical to Section 2.2, the time-reversal of $\mathcal{P}$ can be also be written for $\phi_t(x,p) := \eta_t(x,p)/\eta_t^{\text{ref}}(x,p)$ as $\mathrm{d}y_t = -M^{-1}q_t\mathrm{d}t$ and

$$\mathrm{d}q_t = \frac{y_t}{\sigma^2}\mathrm{d}t - \beta_{T-t}q_t\mathrm{d}t + 2\beta_{T-t}\nabla_{q_t}\ln\phi_{T-t}(y_t,q_t)\mathrm{d}t + \sqrt{2\beta_{T-t}}M^{1/2}\mathrm{d}W_t. \tag{16}$$

To approximate $\mathcal{P}$, we will consider a parameterized path measure $\mathcal{Q}^\theta$ whose time reversal is defined for $(y_0,q_0)\sim\mathcal{N}(y_0;0,I)\mathcal{N}(q_0;0,M)$ by $\mathrm{d}y_t = -M^{-1}q_t\mathrm{d}t$ and

$$\mathrm{d}q_t = \frac{y_t}{\sigma^2}\mathrm{d}t - \beta_{T-t}q_t\mathrm{d}t + 2\beta_{T-t}M f_\theta(T-t,y_t,q_t)\mathrm{d}t + \sqrt{2\beta_{T-t}}M^{1/2}\mathrm{d}W_t. \tag{17}$$

We can then express the KL of interest in a compact way.

Proposition 2. *The Radon-Nikodym derivative* $\frac{\mathrm{d}\mathcal{Q}^\theta}{\mathrm{d}\mathcal{P}^{\text{ref}}}(y_{[0,T]},q_{[0,T]})$ *satisfies under* $\mathcal{Q}^\theta$

$$\ln\left(\frac{\mathrm{d}\mathcal{Q}^\theta}{\mathrm{d}\mathcal{P}^{\text{ref}}}\right) = \int_0^T \beta_{T-t}||f_\theta(T-t,y_t,q_t)||_M^2\mathrm{d}t + \int_0^T \sqrt{2\beta_{T-t}}M^{1/2}f_\theta(T-t,y_t,q_t)^\top\mathrm{d}W_t, \tag{18}$$

where $||x||_A^2 := x^\top A x$. *From* $\mathrm{KL}(\mathcal{Q}^\theta||\mathcal{P}) = \mathrm{KL}(\mathcal{Q}^\theta||\mathcal{P}^{\text{ref}}) + \mathbb{E}_{Q^\theta}\Big[\ln\Big(\frac{p_0^{\text{ref}}(y_T,q_T)}{\eta_0(y_T,q_T)}\Big)\Big]$, *it follows that*

$$\mathrm{KL}(\mathcal{Q}^\theta||\mathcal{P}) = \mathbb{E}_{\mathcal{Q}^\theta}\Big[\int_0^T \beta_{T-t}||f_\theta(T-t,y_t,q_t)||_M^2\mathrm{d}t + \ln\Big(\frac{\mathcal{N}(y_T;0,\sigma^2 I)}{\pi(y_T)}\Big)\Big]. \tag{19}$$

Similarly to Section 2.4, after having learned θ by minimizing (19), we can design a continuous-time normalizing flow to approximate π. In practice, this suggests that once we have learned an approximation $f_\theta(t, x, p) \approx \nabla_p \ln \phi_t(x, p)$ by minimization of (19) then we can construct a proposal using the following ODE initialized at $(y_0, q_0) \sim \mathcal{N}(y_0; 0, \sigma^2 I)\mathcal{N}(q_0; 0, M)$

$$\mathrm{d}y_t = -M^{-1} q_t \mathrm{d}t, \quad \mathrm{d}q_t = \frac{y_t}{\sigma^2}\mathrm{d}t + \beta_{T-t} M f_\theta(T-t, y_t, q_t)\mathrm{d}t. \tag{20}$$

The resulting sample $(y_T, q_T) \sim \bar{q}_0$ will have distribution close to $\pi(y_T)\mathcal{N}(q_T; 0, M)$ which can be computed pointwise this distribution via the instantaneous change of variables formula.

3 Denoising Diffusion Samplers: Discrete-Time

We introduce here discrete-time integrators for the SDEs and ODEs introduced in Section 2. Contrary to DDPM, we not only need an integrator for $\mathcal{Q}_\theta$ but also for $\mathcal{P}^{\text{ref}}$ to be able to compute an approximation of the Radon-Nikodym derivative $\mathrm{d}\mathcal{Q}_\theta/\mathrm{d}\mathcal{P}^{\text{ref}}$. For sake of simplicity, we consider here a constant discretization step $\delta > 0$ such that $K = T/\delta$ is an integer. In the indices, we write k for $t_k = k\delta$ to simplify notation; e.g. x_k for $x_{k\delta}$ and p_k for $p_{k\delta}$. Additionally, we want to ensure that the integrator used for $\mathcal{P}^{\text{ref}}$ is such that the resulting marginal distribution at time 0, $p_0^{\text{ref}}(x)$ or $\eta_0^{\text{ref}}(x, p)$, is known analytically.

3.1 Ornstein–Ulhenbeck

We can integrate exactly (4). This is given by $y_0 \sim \mathcal{N}(0, \sigma^2 I)$ and

$$y_{k+1} = \sqrt{1-\alpha_{K-k}}\, y_k + \sigma\sqrt{\alpha_{K-k}}\,\varepsilon_k, \quad \varepsilon_k \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, I), \tag{21}$$

for $\alpha_k = 1 - \exp\left(-2\int_{(k-1)\delta}^{k\delta} \beta_s \mathrm{d}s\right)$ [2]. This defines the discrete-time reference process. We propose the following exponential type integrator (De Bortoli, 2022) for (8) initialized using $y_0 \sim \mathcal{N}(0, \sigma^2 I)$

$$y_{k+1} = \sqrt{1-\alpha_{K-k}}\, y_k + 2\sigma^2(1 - \sqrt{1-\alpha_{K-k}}) f_\theta(K-k, y_k) + \sigma\sqrt{\alpha_{K-k}}\,\varepsilon_k, \tag{22}$$

where $\varepsilon_k \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, I)$. Equations (21) and (22) define the time reversals of the reference process $p^{\text{ref}}(x_{0:K})$ and proposal $q^\theta(x_{0:K})$. For its time reversal, we write $q^\theta(y_{0:K}) = \mathcal{N}(y_0; 0, \sigma^2 I)\prod_{k=1}^K q^\theta_{k-1|k}(y_{K-k+1}|y_{K-k})$ abusing slightly notation, and similarly for $p^{\text{ref}}(y_{0:K})$. While we could adapt the KL decomposition developed in (Sohl-Dickstein et al., 2015; Ho et al., 2020) for DDPM to the reverse KL (see Appendix A.4), we found empirically that it is very unstable in our context. Instead, we will be relying on the discrete-time counterpart of (9).

Proposition 3. *The log density ratio between $q^\theta(y_{0:K})$ and $p^{\text{ref}}(y_{0:K})$ satisfies*

$$\ln\left(\frac{q^\theta(y_{0:K})}{p^{\text{ref}}(y_{0:K})}\right) = 2\sigma^2 \sum_{k=1}^K \frac{\lambda_k^2}{\alpha_k}||f_\theta(k, y_{K-k})||^2 + 2\sigma \sum_{k=1}^K \frac{\lambda_k}{\sqrt{\alpha_k}} f_\theta(k, y_{K-k})^\top \varepsilon_k \tag{23}$$

where $\lambda_k := 1 - \sqrt{1-\alpha_k}$ and ε_k is defined through (22). In particular for $y_{0:K} \sim q^\theta(y_{0:K})$, hence $\varepsilon_k \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, I)$, one obtains from $\mathrm{KL}(q^\theta||p) = \mathrm{KL}(q^\theta||p^{\text{ref}}) + \mathbb{E}_{q_0^\theta}\left[\ln\left(\frac{\pi(y_K)}{\mathcal{N}(y_K; 0, \sigma^2 I)}\right)\right]$ that

$$\mathrm{KL}(q^\theta||p) = \mathbb{E}_{q^\theta}\left[2\sigma^2 \sum_{k=1}^K \frac{\lambda_k^2}{\alpha_k}||f_\theta(k, y_{K-k})||^2 + \ln\left(\frac{\mathcal{N}(y_K; 0, \sigma^2 I)}{\pi(y_K)}\right)\right]. \tag{24}$$

We compute an unbiased gradient of this objective using the reparameterization trick and the JAX software package (Bradbury et al., 2018) for automatic differentiation. The training procedure is summarized in Algorithm 1 in Appendix D Unfortunately, contrary to DDPM, $q^\theta_{k|K}$ is not available in closed form for $k < K-1$ so we can neither mini-batches over the time index k without having to simulate the process until the minimum sampled time nor reparameterize x_k as in Ho et al. (2020). Once we obtain the parameter θ minimizing (24), DDS samples from q^θ using (22). The final sample y_K has a distribution q_0^θ approximating π by design. By using importance sampling, we obtain an unbiased estimate of the normalizing constant $\hat{Z} = (\gamma(y_K)p^{\text{ref}}(y_{0:K}))/(\mathcal{N}(y_K; 0, \sigma^2 q^\theta(y_{0:K})$ for $y_{0:K} \sim q^\theta(y_{0:K})$. Finally we can also discretize (12) using the same type of integrators as the SDEs to obtain $y_{k+1} = y_k + \sigma^2(1 - \sqrt{1-\alpha_{K-k}})f_\theta(K-k, y_k)$ for $y_0 \sim \mathcal{N}(0; \sigma^2 I)$.

[2] Henceforth, it is assumed that α_k can be computed exactly.

3.2 UNDERDAMPED ORNSTEIN–ULHENBECK

Developing suitable integrators for (15) and (17) such that the corresponding Radon-Nikodym derivative is well-defined is complicated by the fact that there is no noise on the position component y_t. A detailed derivation is provided in Section B.3. Let us first propose an integrator for (15). At it is a linear SDE, it could be principally solved exactly but this solution involves integrals one would need to approximate for general $(\beta_t)_{t=0}^T$. Here we will follow a more widely applicable approach where one splits this SDE into two components:

$$a)\ \frac{\mathrm{d}y_t}{\mathrm{d}t} = -M^{-1}q_t,\ \ \frac{\mathrm{d}q_t}{\mathrm{d}t} = \frac{y_t}{\sigma^2}\mathrm{d}t;\ \ b)\ \mathrm{d}y_t = 0,\ \ \mathrm{d}q_t = -\beta_{T-t}q_t\mathrm{d}t + \sqrt{2\beta_{T-t}}M^{1/2}\mathrm{d}W_t. \quad (25)$$

To obtain (y_{k+1}, q_{k+1}) from (y_k, q_k), we first integrate the ODE (a) on the l.h.s. of (25). Consider the Hamiltonian equations $\frac{\mathrm{d}x_t}{\mathrm{d}t} = M^{-1}p_t$, $\frac{\mathrm{d}p_t}{\mathrm{d}t} = -\frac{x_t}{\sigma^2}$ preserving $\mathcal{N}(y; 0, \sigma^2 I)\mathcal{N}(q; 0, M)$ as invariant distribution. They can be integrated exactly over an interval of length δ. Denote the solution $\Phi(x, p)$. We use $(y_{k+1}, q'_k) = \Phi^{-1}(y_k, q_k)$ where this inverse satisfies $\Phi^{-1}(x, p) = \Phi_{\text{flip}}(x, p) \circ \Phi(x, p) \circ \Phi_{\text{flip}}(x, p)$ where $\Phi_{\text{flip}}(x, p) = (x, -p)$. We then integrate exactly the OU process b) appearing on the r.h.s. of (25), i.e. $q_{k+1} \sim \mathcal{N}(\sqrt{1-\alpha_{K-k}}q'_k, \alpha_{k-1}M)$. This ensures that $\eta_0^{\text{ref}}(y_K, q_K) = \mathcal{N}(y_K; 0, \sigma^2 I)\mathcal{N}(q_K; 0, M)$. We have thus designed a transition of the form $p^{\text{ref}}(y_{k+1}, q_{k+1}, q'_k | y_k, q_k) = \delta_{\Phi^{-1}(y_k,q_k)}(y_{k+1}, q'_k)\mathcal{N}(q_{k+1}; \sqrt{1-\alpha_{K-k}}q'_k; \alpha_{K-k}M)$.

Now to integrate (17), we can split it into three parts: a) and b) as in (25) then c) defined by $\mathrm{d}q_t = 2\beta_{T-t}Mf_\theta(T-t, y_t, q_t)\mathrm{d}t$. We first integrate a) using again $(y_{k+1}, q'_k) = \Phi^{-1}(y_k, q_k)$. Then we integrate c) using $q''_k = q'_k + 2(1-\sqrt{1-\alpha_{K-k}})Mf_\theta(K-k, y_{k+1}, q'_k)$. Finally we integrate b) by setting $q_{k+1} = \sqrt{1-\alpha_{K-k}}q''_k + \sqrt{\alpha_{K-k}}M^{1/2}\varepsilon_k$ where $\varepsilon_k \sim \mathcal{N}(0, I)$. Integrating q''_k analytically, we obtain an integrator whose transition density is $q^\theta(y_{k+1}, q_{k+1}, q'_k | y_k, q_k) = \delta_{\Phi^{-1}(y_k,q_k)}(y_{k+1}, q'_k)\mathcal{N}(q_{k+1}; \sqrt{1-\alpha_{K-k}}(q'_k + 2(1-\sqrt{1-\alpha_{K-k}})Mf_\theta(K-k, y_{k+1}, q'_k)); \alpha_{K-k}M)$. Now it is clear that the ratio between these transitions is well-defined as the deterministic parts are identical in both transitions so cancel.

Proposition 4. *The log density ratio* $\mathrm{lr} = \ln(q^\theta(y_{0:K}, q_{0:K}, q'_{0:K})/p^{\text{ref}}(y_{0:K}, q_{0:K}, q'_{0:K}))$ *equals*

$$\mathrm{lr} = \sum_{k=1}^{K} 2\Big[\frac{\kappa_k^2}{\alpha_k}||f_\theta(k, y_{K-k+1}, q'_{K-k}))||_M^2 + M^{1/2}\frac{\kappa_k}{\sqrt{\alpha_k}}f_\theta(k, y_{K-k+1}, q'_{K-k}))^\top \varepsilon_k\Big] \quad (26)$$

where $\kappa_k := \sqrt{1-\alpha_k}(1-\sqrt{1-\alpha_k})$ *and* ε_k *is obtained as a function of* q_{k+1} *and* q''_k *when describing the integrator for* $\mathcal{Q}^\theta$*. In particular for* $y_{0:K}, q_{0:K}, q'_{0:K} \sim q^\theta(\cdot)$*, hence* $\varepsilon_k \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, I)$*, one obtains*

$$\mathrm{KL}(q^\theta||p) = 2\mathbb{E}_{q^\theta}\left[\sum_{k=1}^{K}\frac{\kappa_k^2}{\alpha_k}\Big[||f_\theta(k, y_{K-k+1}, q'_{K-k}))||_M^2 + \ln\Big(\frac{\mathcal{N}(y_K; 0, \sigma^2 I)}{\pi(y_K)}\Big)\right]. \quad (27)$$

3.3 THEORETICAL GUARANTEES

Motivated by DDPM, bounds on the total variation between the target distribution and the distribution of the samples generated by a time-discretization of an approximate time reversal of a forward noising diffusion have been first obtained in (De Bortoli et al., 2021) then refined in (De Bortoli, 2022; Chen et al., 2022; Lee et al., 2022) and extended to the Wasserstein metric. These results are essentially directly applicable to DDS because their proofs rely on assumptions on the score approximation error but not on the way these approximations are learned, i.e. via score matching for DDPM and reverse KL for DDS. For example, the main result in Chen et al. (2022) shows that if the true scores are L-Lipschitz, the $L_2(p_t)$ error on the scores is bounded and other mild integrability assumptions then DDS outputs samples ϵ-close in total variation to π in $O(L^2 d/\epsilon^2)$ time steps up to some logarithmic factors. As pointed out by the authors, this is truly remarkable as it matches state-of-the-art complexity bounds for Langevin Monte Carlo algorithm for sampling targets satisfying a log-Sobolev inequality. However, we emphasize that the assumption on the approximation error for score estimates is less realistic in the DDS setting than in the DDPM setting as we do not observe realizations of the forward diffusion.

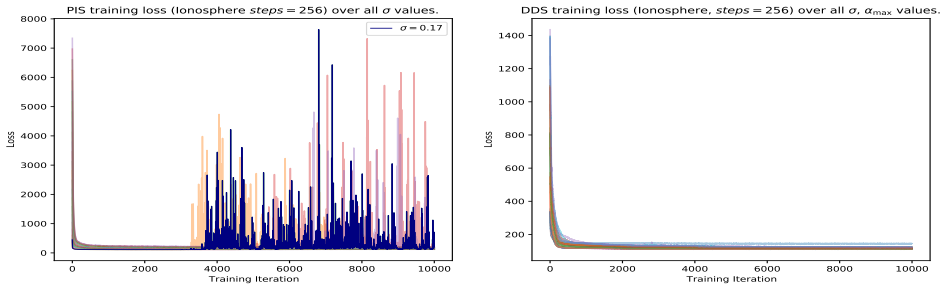

Figure 1: Training loss per hyperparameter: PIS (left) vs DDS (right).

3.4 Intepretation, Related Work and Extensions

DDS as KL and path integral control. The reverse KL we minimize can be expressed as

$$\mathrm{KL}(q^\theta||p) = \mathbb{E}_{q^\theta}\left[\left\{\ln\left(\frac{\mathcal{N}(x_0;0,\sigma^2 I)}{\pi(x_0)}\right) + \sum_{k=1}^{K}\ln\left(\frac{q^\theta_{k-1|k}(x_{k-1}|x_k)}{p^{\mathrm{ref}}_{k-1|k}(x_{k-1}|x_k)}\right)\right\}\right]. \tag{28}$$

This objective is a specific example of KL control problem (Kappen et al., 2012). In continuous-time, (10) corresponds to a path integral control problem; see e.g. (Kappen & Ruiz, 2016).

Heng et al. (2020) use KL control so as to sample from a target π and estimate Z. However, their algorithm relies on $p^{\mathrm{ref}}(x_{0:K})$ being defined by a discretized non-homogeneous Langevin dynamics such that $p_K(x_K)$ is typically not approximating a known distribution. Additionally it approximates the value functions ϕ_k using regression using simple linear/quadratic functions. Finally it relies on a good initial estimate of $p(x_{0:K})$ obtained through SMC. This limits the applicability of this methodology to a restricted class of models.

Connections to Schrödinger Bridges. The Schrödinger Bridge problem (Léonard, 2013) takes the following form in discrete time. Given a reference density $p^{\mathrm{ref}}(x_{0:K})$, we want to find the density $p(x_{0:K})$ such that $p = \arg\min_q\Big\{\mathrm{KL}(q||p^{\mathrm{ref}}) : q_0 = \mu_0,\ \ q_K = \mu_K\Big\}$ where μ_0, μ_K are prescribed distributions. This problem can be solved using iterative proportional fitting (IPF) which is defined by the following recursion with initialization $p^0 = p^{\mathrm{ref}}$

$$p^{2n+1} = \arg\min_q\Big\{\mathrm{KL}(q||p^{2n}) : q_K = \mu_K\Big\},\ \ p^{2n+2} = \arg\min_q\Big\{\mathrm{KL}(q||p^{2n+1}) : q_0 = \mu_0\Big\} \tag{29}$$

For $\mu_0(x_0) = \pi(x_0)$ and $\mu_K(x_K) = \mathcal{N}(x_K;0,\sigma^2 I)$, IPF would converge in one iteration if we had $p^{\mathrm{ref}}_K(x_K) = \mathcal{N}(x_K;0,\sigma^2 I)$. As $p^{\mathrm{ref}}_K(x_K) \approx \mathcal{N}(x_K;0,\sigma^2 I)$, we can thus think of DDS as approximating the solution to this SB as pointed out by De Bortoli et al. (2021).

For $q_0 = \pi$ and $q_K = \delta_0$, IPF also converges in one iteration for $p^{\mathrm{ref}}_K(x_K) = \delta_0(x_K)$. Designing such a reference process is simple, set e.g. a pinned Brownian motion $p^{\mathrm{ref}}_K(x_{0:K}) = \delta_0(x_K)\prod_{k=0}^{K-1}\mathcal{N}(x_k;x_{k+1},\delta\sigma^2 I)$. This was discussed in discrete-time in (Beghi, 1996) and in continuous-time in (Follmer, 1984; Dai Pra, 1991; Tzen & Raginsky, 2019). Algorithmically, Zhang et al. (2021) proposed approximating the solution by some projection on some eigenfunctions. In parallel, Barr et al. (2020), Vargas et al. (2021) and Zhang & Chen (2022) approximated this SB by using a neural network parameterization of the gradient of the logarithm of the corresponding value function trained by minimizing a reverse KL. We will adopt the Path Integral Sampler (PIS) terminology of Zhang & Chen (2022) for this approach. DDS can thus be seen as an alternative to PIS which relies on a reference dynamics corresponding to an overdamped or underdamped Ornstein–Uhlenbeck process instead of relying on a reference dynamics corresponding to a pinned Brownian motion. Empirically this significantly improves numerical stability of the training procedure; see Figure 1. The use of a pinned Brownian motion is also not amenable to the construction of normalizing flows.

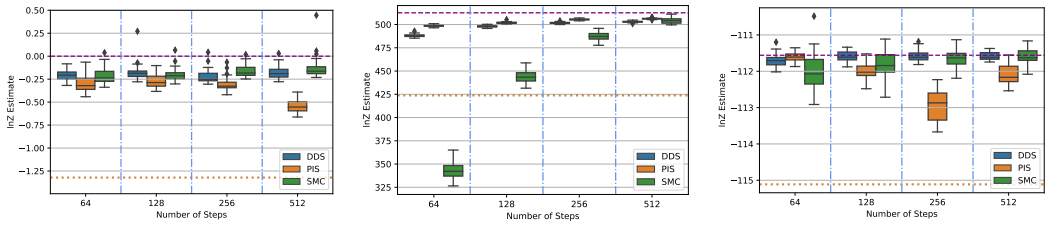

Figure 2: $\ln Z$ estimate as a function of number of steps K - a) Funnel , b) LGCP, c) Logistic Ionosphere dataset. Yellow Dotted line is MF-VI and dashed magenta is the gold standard.

4 EXPERIMENTS

In our implementation, f_θ is implemented using the PIS-GRAD network proposed in (Zhang & Chen, 2022): $f_\theta(k, x) = \text{NN}_1(k, x; \theta) + \text{NN}_2(k; \theta) \circ \nabla \ln \pi(x)$. Across all experiments we use a two layer architecture with 64 hidden units each (for both networks), as done in Zhang & Chen (2022), with the exception of the NICE (Dinh et al., 2014) target where we use 3 layers with 512, 256, and 64 units respectively. The final layers are initialised to 0 in order to make the path regularisation term null. We use $\alpha_k^{1/2} \propto \alpha_{\max}^{1/2} \cos^2\left(\frac{\pi}{2}\frac{1-k/K+s}{1+s}\right)$ with $s = 0.008$ as in (Nichol & Dhariwal, 2021). We found that detaching the target score term (i.e. $\nabla \ln \pi^\perp$) made optimization more stable for both approaches without affecting the final result. We adopt this across experiments, an ablation of this feature can be seen in Appendix C.7.1.

Across all tasks we compare DDS to SMC (Del Moral et al., 2006; Zhou et al., 2016), PIS (Barr et al., 2020; Vargas et al., 2021; Zhang & Chen, 2022), and Mean Field-VI (MF-VI) with a Gaussian variational distribution. We also compare DDS to AIS (Neal, 2001) and optimized variants of AIS using score matching (MCD) (Doucet et al., 2022; Geffner & Domke, 2022) for two standard Bayesian models. Finally we explore a task introduced in (Doucet et al., 2022) that uses a pre-trained normalising flow as a target. Within this setting we propose a benchmarking criterion that allows us to assess mode collapse in high dimensions and explore the benefits of incorporating inductive biases into f_θ. We carefully tuned the hyper-parameters of all algorithms (e.g. step size, diffusion coefficient, and such), details can be found in Appendix C.1. Experiments for the underdamped approach are presented in Appendix C.9. Finally training time can be found in Appendix C.4

4.1 BENCHMARKING TARGETS

We first discuss two standard target distributions which are often us to benchmark methods; see e.g. (Neal, 2003; Arbel et al., 2021; Heng et al., 2020; Zhang & Chen, 2022).

Funnel Distribution: This 10-dimensional challenging distribution is given by $\gamma(x_{0:9}) = \mathcal{N}(x_0; 0, \sigma_f^2)\mathcal{N}(0, \exp(x_0)I)$, where $\sigma_f^2 = 9$ (Neal, 2003). [3].

Log Gaussian Cox process: This model arises in spatial statistics (Møller et al., 1998). We use a $d = M \times M = 1600$ grid, resulting in the unnormalized target density $\gamma(x) = \mathcal{N}(x; \mu, K) \prod_{i \in [1:M]^2} \exp(x_i y_i - a \exp(x_i))$.

4.2 BAYESIAN MODELS

We explore two standard Bayesian models and compare them with standard AIS and MCD benchmarks (See Appendix C.7) in addition to the SMC and VI benchmarks presented so far.

Logistic Regression: We set $x \sim \mathcal{N}(0, \sigma_w^2 I), y_i \sim \text{Bernoulli}(\text{sigmoid}(x^\top u_i))$. This Bayesian logistic model is evaluated on two datasets, Ionosphere ($d = 32$) and Sonar ($d = 61$).

[3]Zhang & Chen (2022) erroneously considered the more favourable $\sigma_f = 1$ scenario for PIS but used $\sigma_f = 3$ for other methods. This explains the significant differences between their results and ours.

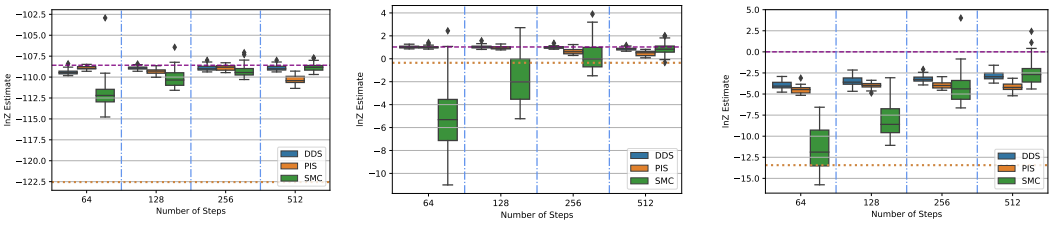

Figure 3: $\ln Z$ estimate as a function of number of steps K - a) Logistic Sonar dataset, b) Brownian motion, c) NICE. Yellow dotted line is MF-VI and dashed magenta is the gold standard.

Brownian Motion: We consider a discretised Brownian motion with a Gaussian observation model and a latent volatility as a time series model, $d = 32$. This model was obtained from the Sountsov et al. (2020) software package and is specified in more detail in Appendix C.2.

4.3 Mode Collapse In High Dimensions

Normalizing Flow Evaluation: Following Doucet et al. (2022) we train NICE (Dinh et al., 2014) on a down-sampled $d = 14 \times 14$ variant of MNIST (LeCun & Cortes, 2010) and use the trained model as our target. As we can generate samples from our target we chose to evaluate the methods samples by measuring the Wasserstein distance between target samples and the respective methods samples. This constitutes of a new evaluation criteria for this benchmark which allows to asses mode collapse for samplers in high dimensional settings.

Method	DDS	PIS	SMC	MCD
$\ln Z$ Estimate	-3.204 ± 0.645	-3.933 ± 0.754	-4.255 ± 2.043	-6.25
$\mathcal{W}^2_{\gamma=0.01}(\pi_{\text{true}}, \hat{\pi})$	658.079	658.778	750.245	NA

Table 1: Results on NICE target, we performed 30 runs with different seeds for each approach. For MCD we used the results from (Doucet et al., 2022).

5 Discussion

We have explored the use of DDPM ideas to develop algorithms to sample unnormalized probability distributions and estimate their normalizing constants. We have compared experimentally the resulting Denoising Diffusion Samplers to state-of-the-art Monte Carlo methods. We found overall that DDS are competitive with SMC and numerically more stable than PIS. However, this comes at the cost of a non-negligible training time compared to SMC. When accounting for this training time, SMC often provide better performance on simple target distributions. However, in the challenging multimodal NICE example, even a carefully tuned SMC sampler using HMC transitions was not competitive to DDS. This is despite the fact that DDS (and PIS) is prone to mode dropping as any method relying on the reverse KL.

Overall, we believe that DDS are a class of algorithms worth considering and further developing. Contrary to SMC methods (Maddison et al., 2017; Le et al., 2018; Naesseth et al., 2018), it provides an objective that can be easily optimized using the reparameterization trick. Additionally, for high-dimensional applications, SMC suffer from high memory requirements as it relies on a large collection of interacting particles. DDS does not suffer from such limitations and, once it has been trained, it allows easy parallel simulation of multiple samples. Moreoever, there has much work devoted to improving successfully DDPM over the past two years including, among many others, modified forward noising mechanisms (Hoogeboom & Salimans, 2022), denoising diffusion implicit models (Song et al., 2020) and sophisticated numerical integrators (Karras et al., 2022). It is highly likely that some of these techniques will lead to more powerful DDS. Advances in KL and path integral control might also be adapted to DDS to provide better training procedures; see e.g. Thalmeier et al. (2020).

REFERENCES

Felix V Agakov and David Barber. An auxiliary variational method. In *Advances in Neural Information Processing Systems*, 2004.

Michael Arbel, Alex Matthews, and Arnaud Doucet. Annealed flow transport Monte Carlo. In *International Conference on Machine Learning*, 2021.

Ariel Barr, Willem Gispen, and Austen Lamacraft. Quantum ground states from reinforcement learning. In *Mathematical and Scientific Machine Learning*, pp. 635–653. PMLR, 2020.

Alessandro Beghi. On the relative entropy of discrete-time Markov processes with given end-point densities. *IEEE Transactions on Information Theory*, 42(5):1529–1535, 1996.

James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL `http://github.com/google/jax`.

Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in Neural Information Processing Systems*, 31, 2018.

Sitan Chen, Sinho Chewi, Jerry Li, Yuanzhi Li, Adil Salim, and Anru R. Zhang. Sampling is as easy as learning the score: theory for diffusion models with minimal data assumptions. *arXiv preprint arXiv:22209.11215*, 2022.

Paolo Dai Pra. A stochastic control approach to reciprocal diffusion processes. *Applied Mathematics and Optimization*, 23(1):313–329, 1991.

Valentin De Bortoli. Convergence of denoising diffusion models under the manifold hypothesis. *arXiv preprint arXiv:2208.05314*, 2022.

Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion Schrödinger bridge with applications to score-based generative modeling. In *Advances in Neural Information Processing Systems*, 2021.

Pierre Del Moral, Arnaud Doucet, and Ajay Jasra. Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(3):411–436, 2006.

Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014.

Tim Dockhorn, Arash Vahdat, and Karsten Kreis. Score-based generative modeling with critically-damped Langevin diffusion. In *International Conference on Learning Representations*, 2022.

Arnaud Doucet, Will Sussman Grathwohl, Alexander G de G Matthews, and Heiko Strathmann. Score-based diffusion meets annealed importance sampling. In *Advances in Neural Information Processing Systems*, 2022.

H. Follmer. An entropy approach to the time reversal of diffusion processes. *Lecture Notes in Control and Information Sciences*, 69:156–163, 1984.

Tomas Geffner and Justin Domke. MCMC variational inference via uncorrected Hamiltonian annealing. In *Advances in Neural Information Processing Systems*, 2021.

Tomas Geffner and Justin Domke. Langevin diffusion variational inference. *arXiv preprint arXiv:2208.07743*, 2022.

Will Grathwohl, Ricky TQ Chen, Jesse Bettencourt, Ilya Sutskever, and David Duvenaud. Ffjord: Free-form continuous dynamics for scalable reversible generative models. In *International Conference on Learning Representations*, 2018.

Ulrich G Haussmann and Etienne Pardoux. Time reversal of diffusions. *The Annals of Probability*, 14(3):1188–1205, 1986.

Jeremy Heng, Adrian N Bishop, George Deligiannidis, and Arnaud Doucet. Controlled sequential Monte Carlo. *The Annals of Statistics*, 48(5):2904–2929, 2020.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, 2020.

Emiel Hoogeboom and Tim Salimans. Blurring diffusion models. *arXiv preprint arXiv:2209.05557*, 2022.

Aapo Hyvärinen. Estimation of non-normalized statistical models by score matching. *The Journal of Machine Learning Research*, 6:695–709, 2005.

Hilbert J Kappen, Vicenç Gómez, and Manfred Opper. Optimal control as a graphical model inference problem. *Machine learning*, 87(2):159–182, 2012.

Hilbert Johan Kappen and Hans Christian Ruiz. Adaptive importance sampling for control and inference. *Journal of Statistical Physics*, 162(5):1244–1266, 2016.

Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *arXiv preprint arXiv:2206.00364*, 2022.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.

Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*, 2013.

Fima C Klebaner. *Introduction to Stochastic Calculus with Applications*. Imperial College Press, 2012.

Tuan Anh Le, Maximilian Igl, Tom Rainforth, Tom Jin, and Frank Wood. Auto-encoding sequential Monte Carlo. In *International Conference on Learning Representations*, 2018.

Yann LeCun and Corinna Cortes. MNIST handwritten digit database. http://yann.lecun.com/exdb/mnist/, 2010. URL `http://yann.lecun.com/exdb/mnist/`.

Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

Holden Lee, Jianfeng Lu, and Yixin Tan. Convergence of score-based generative modeling for general data distributions. *arXiv preprint arXiv:2209.12381*, 2022.

Ben Leimkuhler and Charles Matthews. *Molecular Dynamics with Deterministic and Stochastic Numerical Methods*. Springer, 2016.

Christian Léonard. A survey of the Schrödinger problem and some of its connections with optimal transport. *arXiv preprint arXiv:1308.0215*, 2013.

Chris J Maddison, John Lawson, George Tucker, Nicolas Heess, Mohammad Norouzi, Andriy Mnih, Arnaud Doucet, and Yee Teh. Filtering variational objectives. In *Advances in Neural Information Processing Systems*, 2017.

Jesper Møller, Anne Randi Syversveen, and Rasmus Plenge Waagepetersen. Log Gaussian Cox processes. *Scandinavian Journal of Statistics*, 25(3):451–482, 1998. ISSN 03036898, 14679469.

Christian A Naesseth, Scott W Linderman, Rajesh Ranganath, and David M Blei. Variational sequential Monte Carlo. In *Artificial Intelligence and Statistics*, 2018.

Radford M Neal. Annealed importance sampling. *Statistics and Computing*, 11(2):125–139, 2001.

Radford M. Neal. Slice sampling. *The Annals of Statistics*, 31:705–767, 06 2003.

Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pp. 8162–8171. PMLR, 2021.

George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *arXiv preprint arXiv:1912.02762*, 2019.

Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, 2015.

Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.

Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of score-based diffusion models. In *Advances in Neural Information Processing Systems*, 2021a.

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*, 2021b.

Pavel Sountsov, Alexey Radul, and contributors. Inference gym, 2020. URL `https://pypi.org/project/inference_gym`.

Dominik Thalmeier, Hilbert J Kappen, Simone Totaro, and Vicenç Gómez. Adaptive smoothing for path integral control. 2020.

Achille Thin, Nikita Kotelevskii, Alain Durmus, Eric Moulines, Maxim Panov, and Arnaud Doucet. Monte Carlo variational auto-encoders. In *International Conference on Machine Learning*, 2021.

Belinda Tzen and Maxim Raginsky. Theoretical guarantees for sampling and inference in generative models with latent diffusions. In *Conference on Learning Theory*, pp. 3084–3114. PMLR, 2019.

Francisco Vargas, Andrius Ovsianas, David Fernandes, Mark Girolami, Neil Lawrence, and Nikolas Nüsken. Bayesian learning via neural Schrödinger-Föllmer flows. *arXiv preprint arXiv:2111.10510*, 2021.

Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural Computation*, 23(7):1661–1674, 2011.

Martin J. Wainwright and Michael I. Jordan. Graphical Models, Exponential Families, and Variational Inference. *Foundations and Trends® in Machine Learning*, 1(1–2):1–305, November 2008. URL `https://www.nowpublishers.com/article/Details/MAL-001`.

Hao Wu, Jonas Köhler, and Frank Noé. Stochastic normalizing flows. In *Advances in Neural Information Processing Systems*, 2020.

Benjamin Zhang, Tuhin Sahai, and Youssef Marzouk. Sampling via controlled stochastic dynamical systems. In *I (Still) Can't Believe It's Not Better! NeurIPS 2021 Workshop*, 2021.

Guodong Zhang, Kyle Hsu, Jianing Li, Chelsea Finn, and Roger Grosse. Differentiable annealed importance sampling and the perils of gradient noise. In *Advances in Neural Information Processing Systems*, 2021.

Qinsheng Zhang and Yongxin Chen. Path integral sampler: a stochastic control approach for sampling. In *International Conference on Learning Representations*, 2022.

Yan Zhou, Adam M Johansen, and John AD Aston. Toward automatic model comparison: an adaptive sequential monte carlo approach. *Journal of Computational and Graphical Statistics*, 25 (3):701–726, 2016.