
A Variational Perspective on Diffusion-Based Generative Models and Score Matching

Chin-Wei Huang, Jae Hyun Lim, Aaron Courville
University of Montreal & Mila

{chin-wei.huang, jae.hyun.lim, aaron.courville}@umontreal.ca

Abstract

Discrete-time diffusion-based generative models and score matching methods have shown promising results in modeling high-dimensional image data. Recently, Song et al. (2021) show that diffusion processes that transform data into noise can be reversed via learning the score function, *i.e.* the gradient of the log-density of the perturbed data. They propose to plug the learned score function into an inverse formula to define a generative diffusion process. Despite the empirical success, a theoretical underpinning of this procedure is still lacking. In this work, we approach the (continuous-time) generative diffusion directly and derive a variational framework for likelihood estimation, which includes continuous-time normalizing flows as a special case, and can be seen as an infinitely deep variational autoencoder. Under this framework, we show that minimizing the score-matching loss is equivalent to maximizing a lower bound of the likelihood of the plug-in reverse SDE proposed by Song et al. (2021), bridging the theoretical gap.

1 Introduction

Generative modeling can be thought of as inverting an inference process. If the inference process is invertible, then one can focus on transforming the data into a tractable distribution (Dinh et al., 2016). If the inference process is deterministic yet non-invertible, one could learn to invert it stochastically (Dinh et al., 2019; Nielsen et al., 2020). Most generally, both inference and generation can be stochastic. This is known as the variational autoencoder (Kingma & Welling, 2014; Rezende et al., 2014, VAE).

Under the variational framework, one has a lot of flexibility in choosing the generative and inference models. Recent work on diffusion-based modeling (Sohl-Dickstein et al., 2015; Ho et al., 2020) can be thought of as removing one degree of freedom, by freezing the inference path. The inference model is a fixed discrete-time Markov chain, that slowly transforms the data into a tractable prior, such as the standard normal distribution. The generative model is another Markov chain that is trained to revert this process iteratively. Diffusion-based models have been shown to perform remarkably well on image synthesis (Dhariwal & Nichol, 2021), rivaling the performance of state-of-the-art Generative Adversarial Networks (Brock et al., 2018).

Song et al. (2021) connect diffusion-based model and *score matching* (Hyvärinen & Dayan, 2005), by looking at the stochastic differential equation (SDE) associated with the inference process. They realize that the dynamic of the inference process can be inverted if one has access to the score function of the perturbed data, by solving another SDE reversed in time. They then propose to learn the score function of the inference process and substitute the approximate score into the formula of the reverse SDE to obtain a generative model. We call the resulting generative model the plug-in reverse SDE.




Figure 1: Three special cases of generative SDEs. The stars indicate the initial values, followed by some random sample paths. *Left*: trained with no diffusion $\sigma = 0$ (i.e. neural ODE). *Middle*: trained with some fixed diffusion $\sigma > 0$. *Right*: trained with a fixed inference process, f and g (i.e. the plug-in reverse SDE).

Conceptually simple as this learning procedure may seem, little is known about how the score matching loss relates to the plug-in reverse SDE. In this paper, we propose a variational framework suitable for likelihood estimation for general generative diffusion processes, and use this framework to connect score matching with maximum likelihood. We do so by combining two important theorems in stochastic calculus: the *Feynman-Kac formula* for representing the marginal density of the generative diffusion as an expectation (Section 3), and the *Girsanov theorem* for performing inference in function space (Section 4). We derive a functional evidence lower bound that consistently extends discrete-time diffusion models to have infinite depth, i.e. the number of layers goes to infinity (Section 5). Finally, by reparameterizing our generative and inference SDEs, we obtain a training objective equivalent to minimizing the (implicit) score matching loss (Section 6). Our theory suggests that by matching the score, one actually maximizes a lower bound on the log marginal density of the plug-in reverse SDE, laying a theoretical foundation for this learning procedure. We further generalize our result to a family of *marginal-equivalent* plug-in reverse SDEs, including an equivalent ODE as a limiting case.

Notation: We use (Y_s, s) to denote the inference process (where Y_0 is the data), and (X_t, t) to denote the generative process (where X_0 is a random variable following an unstructured prior). We use s and t to distinguish the two directions, and always integrate the differential equations from 0 to $T > 0$ (different from the literature, where sometimes one might see integration from T to 0). \hat{B}_s and B_t denote the Brownian motions associated with the inference and generative SDEs, respectively. B'_s is a reparameterization of \hat{B}_s (see Section 4). $q(y, s)$ and $p(x, t)$ denote the probability density functions of Y_s and X_t , respectively. We let \mathfrak{s}_θ denote a time-indexed parameterized function that will be used to approximate the score $\nabla \log q(y, s)$. ∇ is the gradient wrt the spatial variable (x or y , which we sometimes call position), ∂_t , ∂_s and ∂_{x_i} are partial derivatives, and H_* denotes Hessian.

2 Background

Assume Y_0 follows the data distribution $q(y, 0)$, and Y_s satisfies the Itô SDE (Øksendal, 2003)

$$dY = f(Y, s) ds + g(Y, s) d\hat{B}_s, \quad (1)$$

where f and g are chosen such that the density $q(y, s)$ will converge to some tractable prior p_0 as $s \rightarrow T$. Following Song et al. (2021), we assume g is position-independent. It is possible to find a “reverse” SDE, whose marginal density evolves according to $q(y, s)$, reversed in time, for example¹

$$dX = (gg^\top \nabla \log q(X, T - t) - f) dt + g dB_t. \quad (2)$$

If $X_0 \sim p_0$, then the density $p(x, t)$ of X_t is equal to $q(x, T - t)$. This means that if we have access to the score function $\nabla \log q$, we can solve the above SDE to obtain $X_T \stackrel{d}{=} Y_0$. Song et al. (2021)

¹See Appendix G for a family of equivalent (reverse) SDEs indexed by some parameter λ , of which equation (2) is a special case with $\lambda = 0$.

Method	Loss
\mathcal{L}_{ESM}	$\frac{1}{2}\mathbb{E}[\ \mathbf{s}_\theta(Y_s, s) - \nabla \log q(Y_s)\ _\Lambda^2]$
\mathcal{L}_{ISM}	$\mathbb{E}[\frac{1}{2}\ \mathbf{s}_\theta(Y_s, s)\ _\Lambda^2 + \nabla \cdot (\Lambda^\top \mathbf{s}_\theta)]$
\mathcal{L}_{SSM}	$\mathbb{E}[\frac{1}{2}\ \mathbf{s}_\theta(Y_s, s)\ _\Lambda^2 + v^\top \nabla (\Lambda^\top \mathbf{s}_\theta)v]$
\mathcal{L}_{DSM}	$\frac{1}{2}\mathbb{E}[\ \mathbf{s}_\theta(Y_s, s) - \nabla \log q(Y_s Y_0)\ _\Lambda^2]$

Table 1: Score matching losses. v follows the Rademacher distribution.

F-K	F-P
$v(y, \varsigma)$	$p(y, T - \varsigma)$
$c(y, \varsigma)$	$-\nabla \cdot \mu(y, T - \varsigma)$
$b(y, \varsigma)$	$-\mu(y, T - \varsigma)$
$\eta(y, \varsigma)$	$\sigma(T - \varsigma)$
$g(y)$	$p_0(y)$

Table 2: Feynman-Kac coefficients.

propose to approximate the score via a parameterized score function \mathbf{s}_θ by minimizing

$$\int_0^T \mathbb{E}_{Y_s} \left[\frac{1}{2} \|\mathbf{s}_\theta(Y_s, s) - \nabla \log q(Y_s, s)\|_{\Lambda(s)}^2 \right] ds$$

where the expectation in the integral is known as the explicit score matching (ESM) loss \mathcal{L}_{ESM} , and $\Lambda(s)$ is a positive definite matrix² that serves as a weighting function for the overall loss. \mathcal{L}_{ESM} is not immediately useful, since we do not have access to the ground truth score $\nabla \log q$. A few alternative losses can be used, which are all equal to one another up to a constant, including implicit score matching (Hyvärinen & Dayan, 2005, ISM), sliced score matching (Song et al., 2020, SSM), and denoising score matching (Vincent, 2011, DSM). The losses are summarized in Table 1, and are related through the following identity (see Appendix A for the derivation):

$$\mathcal{L}_{\text{ESM}} - \frac{1}{2}\mathcal{I}(q(y_s, s)) = \mathcal{L}_{\text{ISM}} = \mathcal{L}_{\text{SSM}} = \mathcal{L}_{\text{DSM}} - \frac{1}{2}\mathbb{E}_{Y_0}[\mathcal{I}(q(y_s|y_0))], \quad (3)$$

where $\mathcal{I}(q) = \mathbb{E}[\|\nabla \log q\|_\Lambda^2]$ is a constant. After training, Song et al. (2021) plug \mathbf{s}_θ into (2) to define a generative model. We refer to this SDE as the *plug-in reverse SDE*. The plug-in reverse SDE has been demonstrated to have impressive empirical results, but a theoretical underpinning of this learning framework is still lacking. For example, it is unclear how the training objective (minimizing the score matching loss) relates to the sampling procedure, *e.g.* whether the probability distribution induced by the plug-in reverse SDE gets closer to the data distribution in the sense of any statistical divergence or metric. We seek to answer the following question in this paper: *How will minimizing the score-matching loss impact the plug-in reverse SDE?* We first provide a framework to estimate the likelihood of generative SDEs, and then get back to this question in Section 6.

3 Marginal density and stochastic instantaneous change of variable

Let X_t be a diffusion process solving the following Itô SDE³:

$$dX = \mu(X, t) dt + \sigma(X, t) dB_t \quad (4)$$

with the initial condition $X_0 \sim p_0$, which induces a family of densities $X_t \sim p(\cdot, t)$. We use this SDE as the generative SDE, and we are interested in $\log p(x, T)$ for maximum likelihood. The density $p(x, t)$ follows the *Kolmogorov forward* (or the *Fokker Planck*) equation:

$$\partial_t p(x, t) = - \sum_j \partial_{x_j} [\mu_j(x, t) p(x, t)] + \sum_{i,j} \partial_{x_i, x_j}^2 [D_{ij}(x, t) p(x, t)] \quad (5)$$

with the initial value $p(\cdot, 0) = p_0(\cdot)$, where $D = \frac{1}{2}\sigma\sigma^\top$ is the diffusion matrix. We can expand the Fokker Planck and rearrange the terms to obtain

$$\begin{aligned} \partial_t p(x, t) = & \left[-\nabla \cdot \mu(x, t) + \sum_{i,j} \partial_{x_i, x_j}^2 D_{ij}(x, t) \right] p(x, t) + \\ & \sum_i \left[-\mu_i(x, t) + 2 \sum_j \partial_{x_j} D_{ij}(x, t) \right] \partial_{x_i} p(x, t) + \sum_{i,j} D_{ij}(x, t) \partial_{x_i, x_j}^2 p(x, t) \end{aligned} \quad (6)$$

²We use this matrix to induce a Mahalanobis norm $\|x\|_\Lambda^2 := x^\top \Lambda x$, which will be used in Section 6.

³For generality, we use the notation μ and σ to describe a generative SDE, which will be set to $g^2 \mathbf{s}_\theta - f$ and g when we come back to the discussion of the plug-in reverse SDE in Section 6.

so that all coefficients of the same order are grouped together. For simplicity, we assume the diffusion term σ is independent of x throughout the paper. Then (6) reduces to

$$\partial_t p(x, t) = -(\nabla \cdot \mu(x, t)) p(x, t) - \mu(x, t)^\top \nabla p(x, t) + D(t) : H_p(x, t) \quad (7)$$

where $:$ denotes the Frobenius inner product between matrices. Even with this simplification, solving (7) is not trivial. Fortunately, we can estimate this quantity using the *Feynman-Kac formula*, which tells us that the solution of certain second-order linear partial differential equations have a probabilistic representation.

Theorem 1 (Feynman-Kac representation, Chapter 5.7 of Karatzas & Shreve (2014)). *Let $T > 0$. Let y and ς be the spatial and temporal arguments to the function $v \in C^{2,1}(\mathbb{R}^d \times [0, T])$ solving*

$$\partial_\varsigma v + cv + b^\top \nabla v + A : H_v = 0 \quad (8)$$

with the terminal condition $v(y, T) = h(y)$, where $A = \frac{1}{2}\eta\eta^\top$ for some matrix-valued function $\eta(y, \varsigma)$. Under the assumption stated in Appendix B, if B'_s is a Brownian motion and Y_s solves

$$dY = b(Y, s) ds + \eta(Y, s) dB'_s, \quad (9)$$

with the initial datum $Y_\varsigma = y$, then

$$v(y, \varsigma) = \mathbb{E} \left[h(Y_T) \exp \left(\int_\varsigma^T c(Y_s, s) ds \right) \middle| Y_\varsigma = y \right]. \quad (10)$$

To estimate the density $p(\cdot, T)$ of (7), we can apply the change of variable $p(x, t) := v(x, T - t)$ by letting the Feynman-Kac (F-K) coefficients correspond to their Fokker-Planck (F-P) counterparts according to Table 2. This way, solving (8) backward is equivalent to solving (7) forward, and we have the following representation of the marginal density at T :

$$p(x, T) = \mathbb{E} \left[p_0(Y_T) \exp \left(\int_0^T -\nabla \cdot \mu(Y_s, T - s) ds \right) \middle| Y_0 = x \right], \quad (11)$$

where Y_s is a diffusion process solving

$$dY = -\mu(Y, T - s) ds + \sigma(T - s) dB'_s. \quad (12)$$

Remark 1 (Marginalization). This representation can be interpreted as a mixture of continuous time flows. Assume a sample path of the Brownian motion is given, and we are interested in how the density evolves following the dynamic (4). In the infinitesimal setting, it can be seen as applying the invertible map $x \mapsto x + \mu(x, t)\Delta t + \sigma(t)\Delta B_i$, where $\Delta B_i := B_{(i+1)\Delta t} - B_{i\Delta t}$ is the Brownian increment. Since the diffusion term is independent of the spatial variable, it can be seen as a constant additive transformation, which is volume preserving, so it will not be taken into account when computing the change of density. The only contribution to the change of density will be from the log-determinant of the Jacobian of $\text{id} + \mu\Delta t$, which means we can simply apply the instantaneous change of variable formula (Chen et al., 2018). This will be the conditional density given the entire $\{B_t : t \geq 0\}$, and marginalizing it out results in the expectation in (11). See Appendix C for details.

Our framework also works with the general case where σ depends on x , but the formulae need to be adapted to account for the spatial partial derivatives. See Appendix D for the derivation.

4 Inferring latent Brownian motion

As our goal is to estimate likelihood, we would like to compute the log density value using (11). However, this involves integrating out all possible Brownian paths, which is intractable. To resolve this, we view the Brownian motion as a latent variable, and perform inference by assigning higher probability to sample paths that are more likely to generate the observation. One can view this as a VAE, except we have an infinite dimensional latent variable.

Formally, let $(\Omega, \mathcal{F}, \mathbb{P})$ be the underlying probability space for which B'_s is a Brownian motion. Suppose \mathbb{Q} is another probability measure on (Ω, \mathcal{F}) equivalent to \mathbb{P} ; that is, \mathbb{P} and \mathbb{Q} are similar in

the sense that they have the same measure zero sets. This allows us to apply the change-of-measure trick and lower bound the log-likelihood with a finite quantity using Jensen’s inequality:

$$\log p(x, T) \geq \mathbb{E}_{\mathbb{Q}} \left[\log \frac{d\mathbb{P}}{d\mathbb{Q}} + \log p_0(Y_T) - \int_0^T \nabla \cdot \mu \, ds \mid Y_0 = x \right]. \quad (13)$$

Note that $\frac{d\mathbb{P}}{d\mathbb{Q}}$ is the *Radon-Nikodym* derivative of \mathbb{P} wrt \mathbb{Q} . When both measures are absolutely continuous wrt a third measure, say Lebesgue, then the derivative can be expressed as the ratio of the two densities. However, since we are dealing with an infinite dimensional space, we are immediately faced with the following problems:

1. Is there a measure \mathbb{Q} (equiv. to \mathbb{P}) for which $\frac{d\mathbb{P}}{d\mathbb{Q}}$ can be easily computed, or at least numerically approximated?
2. Can we find a reparameterization (similar to the Gaussian reparameterization) of B'_s under the new law \mathbb{Q} to estimate the gradient needed for training?

We resort to the *Girsanov theorem*, which describes a general framework for dealing with the change of measure of Gaussian random variables under additive perturbation. It allows us to consider the law of a diffusion process as \mathbb{Q} . See Appendix E for an explanation using the more familiar notion of probability densities.

Theorem 2 (Girsanov theorem, Theorem 8.6.3 of Øksendal (2003)). *Let \hat{B}_s be an Itô process solving*

$$d\hat{B}_s = -a(\omega, s) \, ds + dB'_s, \quad (14)$$

for $\omega \in \Omega$, $0 \leq s \leq T$ and $\hat{B}_0 = 0$, where $a(\omega, s)$ satisfies the Novikov’s condition $\mathbb{E} \left[\exp \left(\frac{1}{2} \int_0^T a^2 \, ds \right) \right] < \infty$. Then \hat{B}_s is a Brownian motion wrt \mathbb{Q} where

$$\frac{d\mathbb{Q}}{d\mathbb{P}}(\omega) := \exp \left(\int_0^T a(\omega, s) \cdot dB'_s - \frac{1}{2} \int_0^T \|a(\omega, s)\|_2^2 \, ds \right). \quad (15)$$

Equation (14) provides a standarization formula of B'_s under \mathbb{Q} , which means we can “invert” it to reparameterize B'_s . This leads to the following lower bound.

Theorem 3 (Continuous-time ELBO). *Let \mathbb{Q} be defined via the density (15). Then the RHS of (13) can be rewritten as*

$$\mathbb{E} \left[-\frac{1}{2} \int_0^T \|a(\omega, s)\|_2^2 \, ds + \log p_0(Y_T) - \int_0^T \nabla \cdot \mu \, ds \mid Y_0 = x \right] =: \mathcal{E}^\infty, \quad (16)$$

where the expectation is taken wrt the Brownian motion \hat{B}_s , and Y_s solves⁴

$$dY = (-\mu + \sigma a) \, ds + \sigma d\hat{B}_s. \quad (17)$$

We call Y_s solving (17) the inference SDE, and \mathcal{E}^∞ the continuous-time ELBO (CT-ELBO).

Remark 2 (Computation). This lower bound can be numerically estimated by using any black box SDE solver, by augmenting the dynamic of y with the accumulation of $\|a\|^2$ and $\nabla \cdot \mu$. Computing the divergence term $\nabla \cdot \mu$ directly can be expensive, but it can be efficiently estimated using the Hutchinson trace estimator (Hutchinson, 1989) along with reverse-mode automatic differentiation, similar to Grathwohl et al. (2018). As the parameters of both the generative and inference models are decoupled from the random variable \hat{B}_s , their gradients can be estimated via the reparameterization trick (Kingma & Welling, 2014; Rezende et al., 2014). Furthermore, backpropagation can be computed using an adjoint method with a constant memory cost (Li et al., 2020).

Remark 3 (Drift a). (i) In general, the drift term of the approximate posterior to the latent Brownian motion can be amortized, so that it will encode the information of individual datum x . (ii) The regularization $\|a\|^2$ ensures that a is kept close to 0, since it represents the deviation of the measure

⁴Note that μ and σ run backward in time from T , whereas a runs forward.

it induces (i.e. \mathbb{Q}) from the classical Wiener measure (which is a centered Gaussian measure). (iii) When the diffusion coefficient σ is 0, the inference SDE reduces to the reverse dynamic of the generative ODE, and if $a \equiv 0$ in this case, the lower bound is tight. (iv) There is generally no constraint on the form of $a(\omega, s)$, so one can potentially augment it with additional dimensions to have a non-Markovian inference SDE. For simplicity, we let the inference SDE be a Markovian model, i.e. $a = a(y, s)$. This is justified by the following theorem.

Theorem 4 (Variational gap and optimal inference SDE). *The variational gap can be written as*

$$\log p(x, T) - \mathcal{E}^\infty = \int_0^T \mathbb{E} [\|a(\omega, s) - \sigma^\top \nabla \log p(Y_s, T - s)\|^2] ds. \quad (18)$$

In particular, $\mathcal{E}^\infty = \log p(x, T)$ if and only if $a(\omega, s)$ can be written as $a(\omega, s) = a(Y_s(\omega), s)$ for almost every $s \in [0, T]$ and $\omega \in \Omega$, and $a(y, s) = \sigma^\top \nabla \log p(y, T - s)$ almost everywhere.

Remark 4 (Variational gap). Even though the inference SDE seemingly takes a simple form, it is sufficiently flexible in that this type of variational problem can be generally solved by taking the supremum over all progressively measurable processes $a(\omega, s)$ (Boué et al., 1998). In fact, the above theorem shows that $\mathcal{E}^\infty = \log p(x, T)$ if and only if $a(y, s) = \sigma^\top \nabla \log p(y, T - s)$. This means a non-amortized Markovian inference process is powerful enough.

5 Infinitely deep hierarchical VAE

Before we make the connection to score matching, we formally address the common belief that “diffusion models can be viewed as the continuous limit of hierarchical VAEs” (Tzen & Raginsky, 2019), and show that the CT-ELBO consistently extends their discrete-time counterpart. We do so by inspecting the ELBO of a hierarchical VAE defined as discretized⁵ generative and inference SDEs. We assume the generative model (i.e. the decoder) follows the transitional distributions

$$p(x_{i+1}|x_i) = \mathcal{N}(x_{i+1}; \tilde{\mu}_i(x_i), \tilde{\sigma}_i^2) \quad (19)$$

$$\tilde{\sigma}_i^2 = \Delta t \sigma^2 (i \Delta t), \quad (20)$$

where $\Delta t = T/L$ is the step size and L is the number of layers. For the inference model (i.e. the encoder), we assume

$$q(x_i|x_{i+1}) = \mathcal{N}(x_i; \hat{\mu}_{i+1}(x_{i+1}), \hat{\sigma}_{i+1}^2) \quad (21)$$

$$\hat{\mu}_i(x) = x + \Delta t(-\mu(x, i \Delta t) + \sigma(i \Delta t)a(x, T - i \Delta t)) \quad \hat{\sigma}_i^2 = \Delta t \sigma^2 (i \Delta t). \quad (22)$$

These transition kernels constitute a hierarchical variational autoencoder of L stochastic layers, whose marginal likelihood can be lower bounded by

$$\log p(x_L) \geq \mathbb{E}_q \left[\log p(x_0) + \sum_{i=0}^{L-1} \log \frac{p(x_{i+1}|x_i)}{q(x_i|x_{i+1})} \right] =: \mathcal{E}^L, \quad (23)$$

which we refer to as the discrete-time ELBO (DT-ELBO). The reconstruction error of the stochastic layer can be seen as some form of finite difference approximation to differentiation, which gives rise to $\nabla \cdot \mu$ in the CT-ELBO in the infinitesimal limit (as Δt approaches 0). The regularization of $\|a\|^2$ pops up when we compare the difference between $\tilde{\mu}_i$ and $\hat{\mu}_i$ using the Gaussian reparameterization to compute the reconstruction error. We formalize this idea in the following theorem.

Theorem 5 (Consistency). *Assume $\mu, \sigma, \sigma^{-2}, a, \|a\|^2$ and their derivatives up to the fourth order are all bounded and continuous, and that σ is non-singular. Then $\mathcal{E}^L \rightarrow \mathcal{E}^\infty$ as $L \rightarrow \infty$.*

This theorem tells us that the CT-ELBO we derive for continuous-time diffusion models is not that different from the traditional ELBO, and that maximizing the CT-ELBO can be seen as training an infinitely deep hierarchical VAE. We present the proof in Appendix F, which formalizes the above intuition, using Taylor’s theorem to control the polynomial approximation error, which will go to 0 as the step size Δt vanishes when the number of layers L increases to infinity.




Figure 2: *Left*: Samples from plug-in reverse SDEs with different λ values (rows). We use the same score function \mathbf{s}_θ trained on the Swiss roll dataset, and plug it into (27). For generation, we use the Euler Maruyama method with a step size of $\Delta t = 1/1000$. We visualize the samples for the i -th iterates (columns), which approximately represent the same marginal distribution when the score function is well trained. *Right*: Lower bound on the marginal likelihood of a continuum of plug-in reverse SDEs. The lower bound is optimized when the score matching loss is minimized, which will push up the entire dark blue curve.

6 Score-based generative modeling

Recall that our goal is to analyze the plug-in reverse SDE and draw connection to score matching. To this end, we reparameterize the generative (4) and inference (17) SDEs as

$$dX = (gg^\top \mathbf{s}_\theta - f) dt + g dB_t \quad \text{and} \quad dY = f ds + g d\hat{B}_s, \quad (24)$$

by letting $a = g^\top \mathbf{s}_\theta$, where the time variable is reversed ($T - t$) for the generative process, and forward in time (s) for inference. The ELBO (16) can be rewritten as

$$\mathcal{E}^\infty = \mathbb{E}_{Y_T} [\log p_0(Y_T) | Y_0 = x] - \int_0^T \mathbb{E}_{Y_s} \left[\frac{1}{2} \|\mathbf{s}_\theta\|_{gg^\top}^2 + \nabla \cdot (gg^\top \mathbf{s}_\theta - f) \Big| Y_0 = x \right] ds. \quad (25)$$

Comparing the integrand to the implicit score matching loss in Table 1, we immediately see that the network \mathbf{s}_θ approximates $\nabla \log q(y, s)$, the score function of the marginal density of Y_s . That is, *matching the score of $q(y, t)$ amounts to maximizing the lower bound on the marginal likelihood of the plug-in reverse SDE.*

Recently, Durkan & Song (2021)⁶ also attempt to establish the equivalency between maximum likelihood and score matching, by showing the following relationship between the forward KL divergence and a weighted sum of score matching loss (aka the Fisher divergence):

$$D_{\text{KL}}(q(y, 0) || r(y, 0)) = \frac{1}{2} \int_0^T \mathbb{E}_{q(\cdot, s)} \left[\|\nabla \log r(Y_s, s) - \nabla \log q(Y_s, s)\|_{gg^\top}^2 \right] ds, \quad (26)$$

where $r(y, s)$ is the density of Y_s solving the same inference SDE with the initial condition $y_0 \sim r(\cdot, 0)$, assuming $q(y, T) = r(y, T)$. However, it is inaccurate to claim that score matching is equivalent to maximum likelihood. This is because if we simply let $r(y, 0) = p(y, T)$, *i.e.* the density of the generative SDE evaluated at y , $r(y, s)$ will not necessarily be the same as either $p(y, T - s)$ or $\mathbf{s}_\theta(y, s)$. This means the KL divergence is not equal to the integral of the weighted score matching loss $\mathbb{E}[\frac{1}{2} \|\mathbf{s}_\theta - \nabla \log q\|_{gg^\top}^2]$. In fact, the latter corresponds to a lower bound on the likelihood (the cross-entropy term of the KL) up to some constant, as equation (25) suggests.

More generally, we can apply our analysis to a family of plug-in reverse SDEs indexed by some parameter $\lambda \leq 1$:

$$dX = \left((1 - \frac{\lambda}{2}) g^2 \mathbf{s}_\theta - f \right) dt + \sqrt{1 - \lambda} g dB_t \quad \text{and} \quad dY = \left(f - \frac{\lambda}{2} g^2 \nabla \log q \right) ds + \sqrt{1 - \lambda} g d\hat{B}_s, \quad (27)$$

⁵We follow the Euler-Maruyama (EM) scheme. Other discretization scheme may also work; we leave that for future work.

⁶This refers to the v1 of the paper on arXiv. This version was later on replaced with a new version where they derived a similar bound as ours.

where we assume g is diagonal for simplicity. We defer the formal discussion to Appendix G, but the essence is that this inference SDE induces the same marginal distribution as (1), and the generative SDE is its corresponding plug-in reverse. Equation (27) includes the original plug-in reverse SDE (24) and an equivalent ODE as special cases with $\lambda = 0$ and $\lambda = 1$. Denote its corresponding CT-ELBO by $\mathcal{E}_\lambda^\infty$. Specifically, (25) becomes \mathcal{E}_0^∞ . Then we have the following relationship.

Theorem 6 (Plug-in reverse SDE ELBO, abridged). *For $\lambda < 1$,*

$$\mathbb{E}_{Y_0}[\mathcal{E}_\lambda^\infty] = \mathbb{E}_{Y_0}[\mathcal{E}_0^\infty] - \left(\frac{\lambda^2}{4(1-\lambda)}\right) \int_0^T \mathbb{E}_{Y_s} \left[\frac{1}{2} \|\mathbf{s}_\theta(Y_s, s) - \nabla \log q(Y_s, s)\|_{g^2}^2 \right] ds \quad (28)$$

We state the full theorem in Appendix H, where we rearrange the terms to show that the average CT-ELBO of the λ -plug-in reverse SDE is also equivalent to the ISM loss, similarly to (25) but up to some multiplying and additive constants. The implication is that *while minimizing the score matching loss, we implicitly maximize the likelihood of a continuum of plug-in reverse SDEs which include the ODE as a limiting case ($\lambda \rightarrow 1$)*. See Figure 2 (right) for illustration. This suggests the likelihood of the equivalent ODE can be improved by minimizing the score matching loss, as the ODE’s likelihood will be close to plug-in reverse SDEs with $\lambda \approx 1$, which explains the good likelihood of the equivalent ODE reported in Song et al. (2021). In practice, we can only estimate the ELBO of the case $\lambda = 0$ since otherwise there will be some constant we do not have access to, but their gradients can all be estimated via score matching.

6.1 Computational trade-off

Having a general framework for estimating the likelihood of diffusion processes allows us to compare a wide family of models, including continuous-time flows and plug-in reverse SDEs trained by score matching. We compare the two by measuring the negative ELBO throughout training to highlight their computation-estimation trade-off. We train the models on the Swiss roll toy data. For continuous-time flow, we set $\sigma = 0$, using the Hutchinson trace estimator following Grathwohl et al. (2018). The ELBO in this case is tight since a will be penalized to be 0. We use the `torchdiffeq` library (Chen et al., 2018) for numerical integration for fairer comparison⁷. For plug-in reverse SDE, we train the drift network a using SSM and DSM (for DSM the loss is weighted to reduce variance, which introduces some bias; see the next subsection). We use the variance-preserving inference SDE from Song et al. (2021), which allows us to sample Y_s using a closed form formula, for s sampled uniformly between $[0, T]$. The trained models are visualized in Figure 1, the learning curves presented in Figure 3.

From the learning curve figures, we see that neg-likelihood decreases rapidly for the continuous-time flow in terms of the number of parameter updates. But once the x-axis is normalized by runtime, the convergence speed becomes almost indistinguishable. This is because for continuous-time flows, numerical integration takes time, whereas for plug-in reverse SDEs, we train on a random time step s ; that is, within a fixed amount of time the latter can make more parameter updates at the cost of noisier gradients. Note that both models have constant memory cost (wrt T or L , the number of integration steps), so a large batch size can be used to reduce variance for training.

6.2 Bias and variance trade-off

The integral in equation (25) can be estimated by sampling (Y_s, s) , and using the Hutchinson trace estimator to estimate the divergence, which corresponds to implicit score matching. However, in practice the variance of this estimator is very high when the norm of the Jacobian $\nabla \mathbf{s}_\theta$ is large. Another popular approach is to use the denoising estimator (recall the identity from (3)),

$$\mathbb{E}_{Y_s} \left[\frac{1}{2} \|\mathbf{s}_\theta(Y_s, s) - \nabla \log q(Y_s | Y_0)\|_{gg^\top}^2 \mid Y_0 = x \right]. \quad (29)$$

The inference SDE is typically chosen so that Y_s can be easily sampled, e.g. following $\mathcal{N}(\mu_s, \sigma_s^2)$, where μ_s and σ_s are functions of Y_0 and s . In this case, if we reparameterize $Y_s = \mu_s + \sigma_s \epsilon$ where

⁷Black-box SDE solvers such as `torchsde` (Li et al., 2020) might not be optimized for the deterministic case, since their stochastic adjoint method scales $\mathcal{O}(L \log L)$ in time whereas deterministic numerical solvers are usually faster. This matters for our runtime comparison.




Figure 3: Neural ODE vs plug-in reverse SDE (denoising or slice score matching). The learning curves are presented as a function of iterations (left) and runtime (right) to emphasize the computational distinction between the two families of models.

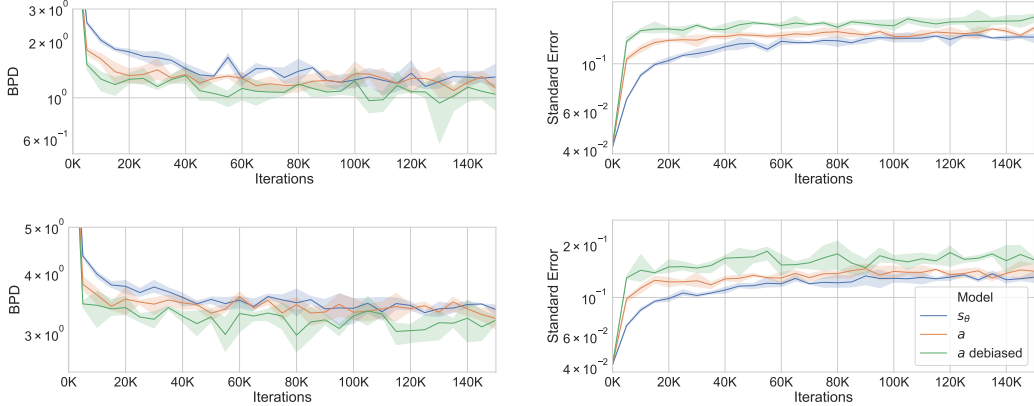


Figure 4: Likelihood estimation on MNIST (first row) and CIFAR10 (second row). S_θ and a denote which model we parameterize. Y-axes are bits-per-dim and the standard error of BPD of the test set. The debiased curves improve upon the original biased gradient estimator (Song et al., 2021) since it maximizes a proper ELBO. Shaded area reflects the uncertainty estimated by 3 random seeds.

$\epsilon \sim \mathcal{N}(0, \mathbf{I})$, then the score becomes $\nabla \log q = -\frac{\epsilon}{\sigma_s}$. Since $\sigma_s \rightarrow 0$ as $s \rightarrow 0$, this estimator normally has unbounded variance. Song & Ermon (2019); Song et al. (2021) propose to remedy this by multiplying the DSM loss by σ_s^2/g^2 (assuming g is a scalar for simplicity), so that the target has constant magnitude on average $\mathbb{E}[\frac{1}{2}|\sigma_s \mathbf{s}_\theta + \epsilon|^2]$, which would result in a biased gradient estimate with much smaller variance. We can debias this estimator by sampling $s \sim q(s) \propto g^2/\sigma_s^2$. This ratio, however, is usually not normalizable in practice (as it integrates to ∞). As an alternative, we consider the following unnormalized density $\tilde{q}_\epsilon(s) = g^2(s_\epsilon)/\sigma_{s_\epsilon}^2$ for $s \in [0, s_\epsilon]$, and $\tilde{q}_\epsilon(s) = g^2(s)/\sigma_s^2$ for $s \in [s_\epsilon, T]$. We experiment with this debiased procedure by sampling $s \sim q_\epsilon \propto \tilde{q}_\epsilon$, for f and g chosen to be the variance-preserving SDE. s_ϵ is small so that the bias is negligible.

We train the model on MNIST (LeCun et al., 1998) and CIFAR10 (Krizhevsky et al., 2009). We present the learning curves and the standard error of the estimate of the ELBO in Figure 4. The lower bound is estimated using the Hutchinson trace estimator with s sampled uniformly from $[0, T]$, with the same batch size, so the only thing that will affect the dispersion is the magnitude of $\nabla \mathbf{s}_\theta$. Since smaller values of s are more likely to be sampled under q_ϵ , the debiased model will see samples with less perturbation more often. On the contrary, sampling s uniformly will bias the model to learn from noisier data, causing the learned score to be smoother. We also experiment with parameterizing S_θ vs parameterizing a . We find the latter parameterization to be helpful since the relationship $\mathbf{s}_\theta = g^{-1}a$ has the effect of negating the multiplier σ_s in the reweighted loss, *i.e.* $\mathbb{E}[\frac{1}{2}|\frac{\sigma_s}{g}a + \epsilon|^2]$. This is similar to the noise conditioning technique introduced in Song & Ermon (2020).

7 Related work

Diffusion-based generative models Our work lays a theoretical foundation for Song et al. (2021), which recognizes that conditional denoising score matching (Song & Ermon, 2019, 2020) and

discrete-time diffusion-based generative models (Sohl-Dickstein et al., 2015; Goyal et al., 2017; Ho et al., 2020) can be viewed as learning to revert an inference process (using the plug-in reverse SDE). Different from Ho et al. (2020), which shows the ELBO of discrete time diffusion process can be likened to DSM (Section 3.2 of the paper), we show that ISM loss naturally arises from the Fokker-Planck equation of the marginal density, via the Feynman-Kac representation and the Girsanov change of measure. This line of work has been successfully applied to modeling high dimensional natural images (Dhariwal & Nichol, 2021; Saharia et al., 2021), audio (Kong et al., 2020), 3D point cloud (Cai et al., 2020; Zhou et al., 2021), and discrete data (Hoogeboom et al., 2021).

Time-reversal of diffusion processes Plenty of works have studied the reverse-time diffusion processes (2), including Anderson (1982); Föllmer (1985); Elliott & Anderson (1985); Haussmann & Pardoux (1986). These are different from our marginal-equivalent (reverse) processes (27) when $s_\theta = \nabla \log q$, since the latter is related by the marginals only.

Score matching for energy-based models Besides the connection to diffusion models, score matching is also often used as a method for learning energy based models (EBM)— see Song & Kingma (2021) for a comprehensive review on useful techniques—. When used as an EBM, sampling from the conditional score model can be achieved by running the annealed Langevin diffusion (Neal, 2001), which is connected to free-energy estimator in physics (Jarzynski, 1997), wherein the *path integral* is essentially a Feynman Kac representation.

De Bruijn’s identity To connect maximum likelihood and score matching, Durkan & Song (2021) shows that KL divergence can be represented as an integral of weighted Fisher divergence, generalizing the case of Lyu (2009) where the inference perturbation is a simple Brownian motion. This type of formulas fall into the category of de Bruijn’s identity (Cover, 1999) for relative entropy. A similar differential form result can be found in Wibisono et al. (2017).

Learning SDEs Tzen & Raginsky (2019); Li et al. (2020) also propose to learn a neural SDE by applying Girsanov’s theorem. The key difference is that they treat the SDE entirely as a latent variable, with an additional emission probability, whereas we use the Feynman-Kac formula to directly express the marginal density as an expectation, side-stepping the need to smooth out the density using the emission probability (which will be a Dirac point mass in our case). In their case, the inference direction is the same as the generative direction, since they infer the latent SDE directly, whereas we apply Girsanov to the Feynman-Kac diffusion (opposite the generative direction). Xu et al. (2021) further apply neural SDE as an infinitely deep Bayesian neural network.

8 Conclusion and Discussion

In this work, we derive a general variational framework for estimating the marginal likelihood of continuous-time diffusion models. This framework allows us to study a wide spectrum of models, including continuous-time normalizing flows and score-based generative models. Using our framework, we show that performing score matching with a particular choice of mixture weighting is equivalent to maximizing a lower bound on the marginal likelihood of a family of plug-in reverse SDEs, of which the one used in Song et al. (2021) and the equivalent ODE are special cases. Empirically, we validate our theory by monitoring the ELBO while performing score matching, and discuss the implication of the choice of mixture weighting and the potential of debiasing via non-uniform sampling. We emphasize that our theory does not explain the impressive sample quality of this family of models, which is still an open research problem and we leave it for future work.

This work introduces a general framework to estimate the likelihood of diffusion-based models, which allows the parameters of both the generative and inference SDEs to be learned, using a numerical solver with constant memory cost (as per Remark 2). The training time can be reduced via the connection to score matching and the reverse-time parameterization (24) as long as f and g take a simple form (so that Y_t can be sampled without numerical integration). For example, one can generalize the Ornstein-Uhlenbeck process to have non-linear (in time) f and g , similar to the variance-preserving SDE, by parameterizing the integral of f using a monotone network (Sill, 1998; Kay & Ungar, 2000; Daniels & Velikova, 2010; Huang et al., 2018). This has been explored in a concurrent work by Kingma et al. (2021) in a different framework.

Acknowledgements

We would like to thank David Kanaa, Ricky Chen, Simon Verret, Rémi Piché-Taillefer, Alexia Jolicoeur-Martineau, and Faruk Ahmed for giving their feedback on this manuscript. We would also like to thank the INNF+ 2021 reviewers and NeurIPS 2021 reviewers for their constructive suggestions, which help us improve the clarity of the paper. Chin-Wei is supported by the Google PhD fellowship.

We also acknowledge the Python community (Van Rossum & Drake Jr, 1995; Oliphant, 2007) for developing the tools that enabled this work, including numpy (Oliphant, 2006; Van Der Walt et al., 2011; Walt et al., 2011; Harris et al., 2020), PyTorch (Paszke et al., 2019), Matplotlib (Hunter, 2007), seaborn (Waskom et al., 2018), and SciPy (Jones et al., 2014).

References

- Anderson, B. D. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982.
- Boué, M., Dupuis, P., et al. A variational representation for certain functionals of brownian motion. *The Annals of Probability*, 26(4):1641–1659, 1998.
- Brock, A., Donahue, J., and Simonyan, K. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018.
- Cai, R., Yang, G., Averbuch-Elor, H., Hao, Z., Belongie, S., Snavely, N., and Hariharan, B. Learning gradient fields for shape generation. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part III*, 2020.
- Chen, R. T., Rubanova, Y., Bettencourt, J., and Duvenaud, D. Neural ordinary differential equations. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pp. 6572–6583, 2018.
- Cover, T. M. *Elements of information theory*. John Wiley & Sons, 1999.
- Daniels, H. and Velikova, M. Monotone and partially monotone neural networks. *IEEE Transactions on Neural Networks*, 21(6):906–917, 2010.
- Dhariwal, P. and Nichol, A. Diffusion models beat gans on image synthesis. *arXiv preprint arXiv:2105.05233*, 2021.
- Dinh, L., Sohl-Dickstein, J., and Bengio, S. Density estimation using real nvp. In *International Conference on Learning Representations*, 2016.
- Dinh, L., Sohl-Dickstein, J., Larochelle, H., and Pascanu, R. A rad approach to deep mixture models. *arXiv preprint arXiv:1903.07714*, 2019.
- Durkan, C. and Song, Y. On maximum likelihood training of score-based generative models. *arXiv preprint arXiv:2101.09258*, 2021.
- Elliott, R. J. and Anderson, B. D. Reverse time diffusions. *Stochastic processes and their applications*, 19(2):327–339, 1985.
- Föllmer, H. An entropy approach to the time reversal of diffusion processes. In *Stochastic Differential Systems Filtering and Control*, pp. 156–163. Springer, 1985.
- Goyal, A., Ke, N. R., Ganguli, S., and Bengio, Y. Variational walkback: Learning a transition operator as a stochastic recurrent net. *Advances in Neural Information Processing Systems*, 2017.
- Grathwohl, W., Chen, R. T., Bettencourt, J., Sutskever, I., and Duvenaud, D. Ffjord: Free-form continuous dynamics for scalable reversible generative models. In *International Conference on Learning Representations*, 2018.
- Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., et al. Array programming with numpy. *Nature*, 585(7825): 357–362, 2020.

- Hausmann, U. G. and Pardoux, E. Time reversal of diffusions. *The Annals of Probability*, pp. 1188–1205, 1986.
- Ho, J., Jain, A., and Abbeel, P. Denoising diffusion probabilistic models. In *Advances in neural information processing systems*, 2020.
- Hoogeboom, E., Nielsen, D., Jaini, P., Forré, P., and Welling, M. Argmax flows and multinomial diffusion: Towards non-autoregressive language models. *arXiv preprint arXiv:2102.05379*, 2021.
- Huang, C.-W., Krueger, D., Lacoste, A., and Courville, A. Neural autoregressive flows. In *International Conference on Machine Learning*, pp. 2078–2087, 2018.
- Hunter, J. D. Matplotlib: A 2d graphics environment. *Computing in science & engineering*, 9(3): 90, 2007.
- Hutchinson, M. F. A stochastic estimator of the trace of the influence matrix for laplacian smoothing splines. *Communications in Statistics-Simulation and Computation*, 18(3):1059–1076, 1989.
- Hyvärinen, A. and Dayan, P. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(4), 2005.
- Jarzynski, C. Equilibrium free-energy differences from nonequilibrium measurements: A master-equation approach. *Physical Review E*, 56(5):5018, 1997.
- Jones, E., Oliphant, T., and Peterson, P. {SciPy}: Open source scientific tools for {Python}. 2014.
- Karatzas, I. and Shreve, S. *Brownian motion and stochastic calculus*, volume 113. springer, 2014.
- Kay, H. and Ungar, L. H. Estimating monotonic functions and their bounds. *AICHE Journal*, 46 (12):2426–2434, 2000.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. In *International Conference on Learning Representations*, 2014.
- Kingma, D. P., Salimans, T., Poole, B., and Ho, J. Variational diffusion models. *arXiv preprint arXiv:2107.00630*, 2021.
- Kong, Z., Ping, W., Huang, J., Zhao, K., and Catanzaro, B. Diffwave: A versatile diffusion model for audio synthesis. *arXiv preprint arXiv:2009.09761*, 2020.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. 2009.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Li, X., Wong, T.-K. L., Chen, R. T., and Duvenaud, D. Scalable gradients for stochastic differential equations. In *International Conference on Artificial Intelligence and Statistics*, pp. 3870–3882. PMLR, 2020.
- Lyu, S. Interpretation and generalization of score matching. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, pp. 359–366, 2009.
- Milshtein, G. Approximate integration of stochastic differential equations. *Theory of Probability & Its Applications*, 19(3):557–562, 1975.
- Neal, R. M. Annealed importance sampling. *Statistics and computing*, 11(2):125–139, 2001.
- Nielsen, D., Jaini, P., Hoogeboom, E., Winther, O., and Welling, M. Survae flows: Surjections to bridge the gap between vaes and flows. *Advances in Neural Information Processing Systems*, 33, 2020.
- Øksendal, B. Stochastic differential equations. In *Stochastic differential equations*, pp. 65–84. Springer, 2003.
- Oliphant, T. E. *A guide to NumPy*, volume 1. Trelgol Publishing USA, 2006.

- Oliphant, T. E. Python for scientific computing. *Computing in Science & Engineering*, 9(3):10–20, 2007.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*, pp. 8026–8037, 2019.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. In *International conference on machine learning*, pp. 1278–1286. PMLR, 2014.
- Ronneberger, O., Fischer, P., and Brox, T. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pp. 234–241. Springer, 2015.
- Saharia, C., Ho, J., Chan, W., Salimans, T., Fleet, D. J., and Norouzi, M. Image super-resolution via iterative refinement. *arXiv preprint arXiv:2104.07636*, 2021.
- Sill, J. Monotonic networks. In *Advances in neural information processing systems*, pp. 661–667, 1998.
- Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., and Ganguli, S. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pp. 2256–2265. PMLR, 2015.
- Song, Y. and Ermon, S. Generative modeling by estimating gradients of the data distribution. In *Advances in neural information processing systems*, 2019.
- Song, Y. and Ermon, S. Improved techniques for training score-based generative models. *arXiv preprint arXiv:2006.09011*, 2020.
- Song, Y. and Kingma, D. P. How to train your energy-based models. *arXiv preprint arXiv:2101.03288*, 2021.
- Song, Y., Garg, S., Shi, J., and Ermon, S. Sliced score matching: A scalable approach to density and score estimation. In *Uncertainty in Artificial Intelligence*, pp. 574–584. PMLR, 2020.
- Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021.
- Tzen, B. and Raginsky, M. Neural stochastic differential equations: Deep latent gaussian models in the diffusion limit. *arXiv preprint arXiv:1905.09883*, 2019.
- Van Der Walt, S., Colbert, S. C., and Varoquaux, G. The numpy array: a structure for efficient numerical computation. *Computing in Science & Engineering*, 13(2):22, 2011.
- Van Rossum, G. and Drake Jr, F. L. *Python reference manual*. Centrum voor Wiskunde en Informatica Amsterdam, 1995.
- Vincent, P. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011.
- Walt, S. v. d., Colbert, S. C., and Varoquaux, G. The numpy array: a structure for efficient numerical computation. *Computing in science & engineering*, 13(2):22–30, 2011.
- Waskom, M., Botvinnik, O., O’Kane, D., Hobson, P., Ostblom, J., Lukauskas, S., Gemperline, D. C., Augspurger, T., Halchenko, Y., Cole, J. B., Warmenhoven, J., de Ruiter, J., Pye, C., Hoyer, S., Vanderplas, J., Villalba, S., Kunter, G., Quintero, E., Bachant, P., Martin, M., Meyer, K., Miles, A., Ram, Y., Brunner, T., Yarkoni, T., Williams, M. L., Evans, C., Fitzgerald, C., Brian, and Qalieh, A. mwaskom/seaborn: v0.9.0, July 2018.
- Wibisono, A., Jog, V., and Loh, P.-L. Information and estimation in fokker-planck channels. In *2017 IEEE International Symposium on Information Theory (ISIT)*, pp. 2673–2677. IEEE, 2017.

Xu, W., Chen, R. T., Li, X., and Duvenaud, D. Infinitely deep bayesian neural networks with stochastic differential equations. *arXiv preprint arXiv:2102.06559*, 2021.

Zhou, L., Du, Y., and Wu, J. 3d shape generation and completion through point-voxel diffusion. *arXiv preprint arXiv:2104.03670*, 2021.

A Score matching losses

In this section, we prove the score matching loss identity for completeness. These proofs are adapted from [Hyvärinen & Dayan \(2005\)](#); [Song et al. \(2020\)](#); [Vincent \(2011\)](#) with slight modifications since we project the score onto the eigen-basis of $\Lambda(s)$. Recall the definition of the ESM loss

$$\mathcal{L}_{\text{ESM}} = \mathbb{E}_{Y_s} \left[\frac{1}{2} \|\mathbf{s}_\theta(Y_s, s) - \nabla \log q(Y_s, s)\|_{\Lambda(s)}^2 \right] \quad (30)$$

where $Y_s \sim q(Y_s, s)$. Expanding the quadratic equation, we have

$$\mathcal{L}_{\text{ESM}} = \mathbb{E}_{Y_s} \left[\frac{1}{2} \|\mathbf{s}_\theta(Y_s, s)\|_{\Lambda(s)}^2 - \mathbf{s}_\theta(Y_s, s)^\top \Lambda(s) \nabla \log q(Y_s, s) + \frac{1}{2} \|\nabla \log q(Y_s, s)\|_{\Lambda(s)}^2 \right] \quad (31)$$

Moving $\mathcal{I}(q(Y_s, s))$ from the RHS to the LHS gives us

$$\mathcal{L}_{\text{ESM}} - \frac{1}{2} \mathcal{I}(q(Y_s, s)) = \mathbb{E}_{Y_s} \left[\frac{1}{2} \|\mathbf{s}_\theta(Y_s, s)\|_{\Lambda(s)}^2 - \mathbf{s}_\theta(Y_s, s)^\top \Lambda(s) \nabla \log q(Y_s, s) \right] \quad (32)$$

A.1 Implicit score matching

Now to draw connection to ISM, we apply integration by parts and the general Stokes' theorem (with mild regularity condition on \mathbf{s}_θ) to the inner product term to obtain

$$\begin{aligned} \int q(y, s) \mathbf{s}_\theta(y, s)^\top \Lambda(s) \nabla \log q(y, s) dy &= \int \mathbf{s}_\theta(y, s)^\top \Lambda(s) \nabla q(y, s) dy \\ &= \int \nabla \cdot (q \Lambda^\top \mathbf{s}_\theta) dy - \int q \nabla \cdot (\Lambda^\top \mathbf{s}_\theta) dy \\ &= \mathbb{E}_{Y_s} [\nabla \cdot (\Lambda^\top \mathbf{s}_\theta)] \end{aligned}$$

A.2 Sliced score matching

For the SSM loss, use the Hutchinson trace estimator ([Hutchinson, 1989](#)) to replace the divergence operator, which is simply the trace of the Jacobian matrix.

A.3 Denoising score matching

For DSM, similarly we first look at the inner product term

$$\begin{aligned} \int q(y, s) \mathbf{s}_\theta(y, s)^\top \Lambda(s) \nabla \log q(y, s) dy &= \int \mathbf{s}_\theta(y, s)^\top \Lambda(s) \nabla q(y, s) dy \\ &= \int \mathbf{s}_\theta(y_s, s)^\top \Lambda(y_s) \nabla \int q(y_s | y_0) q(y_0, 0) dy_0 dy_s \\ &= \int \int q(y_0, 0) \mathbf{s}_\theta^\top \Lambda \nabla q(y_s | y_0) dy_s dy_0 \\ &= \int \int q(y_0, 0) q(y_s | y_0) \mathbf{s}_\theta^\top \Lambda \nabla \log q(y_s | y_0) dy_s dy_0 \\ &= \mathbb{E}_{Y_0, Y_s} [\mathbf{s}_\theta^\top \Lambda \nabla \log q(Y_s | Y_0)] \end{aligned}$$

where $q(y_s | y_0)$ denotes the conditional density of Y_s given Y_0 . Combining this with $\mathbb{E}_{Y_s} [\|\mathbf{s}_\theta\|_{\Lambda}^2]$, we have

$$\mathbb{E}_{Y_0, Y_s} \left[\frac{1}{2} \|\mathbf{s}_\theta\|_{\Lambda}^2 - \mathbf{s}_\theta^\top \Lambda \nabla \log q(Y_s | Y_0) \right] = \mathbb{E}_{Y_0, Y_s} \left[\frac{1}{2} \|\mathbf{s}_\theta - \nabla \log q(Y_s | Y_0)\|_{\Lambda}^2 \right] - \frac{1}{2} \mathbb{E}_{Y_0} [\mathcal{I}(q(y_s | Y_0))]$$

B Assumption of Feynman-Kac

Assumption 1 (Feynman-Kac). *We assume the following: There exist some constants $B_h, B_v > 0$ and $p_h, p_v \geq 1$ such that $h \in C^0(\mathbb{R}^d)$ and $v \in C^{2,1}(\mathbb{R}^d \times [0, T])$ satisfy*

$$|h(y)| \leq B_h (1 + \|y\|^{2p_h}) \quad \text{or} \quad h(y) \geq 0 \quad (33)$$

$$\max_{0 \leq \varsigma \leq T} |v(y, \varsigma)| \leq B_v (1 + \|y\|^{2p_v}) \quad (34)$$

C Mixture of continuous-time flows

Continuing the discussion in Remark 1, we analyze the limit of the determinant of the Jacobian of the finite approximation given the Brownian path: $x \leftarrow x + \mu(x, t)\Delta t + \sigma(t)\Delta B_i$. When the step size decreases to 0, this should converge to the Itô integral. When Δt is small enough, under the assumption that μ is uniformly Lipschitz, the finite approximation will be invertible for all steps. Then the determinant of the Jacobian of the overall transformation is just the product of determinant of each step:

$$\begin{aligned} \prod_i \det(\nabla(x + \mu(x, t)\Delta t + \sigma(t)\Delta B_i)) &= \prod_i \det(\mathbf{I} + \Delta t \nabla \mu) \\ &= \prod_i (1 + \Delta t \text{Tr}(\nabla \mu) + \mathcal{O}(\Delta t^2)) \\ &= \exp\left(\sum_i \log(1 + \Delta t \nabla \cdot \mu + \mathcal{O}(\Delta t^2))\right) \\ &= \exp\left(\sum_i \Delta t \nabla \cdot \mu + \mathcal{O}(\Delta t^2)\right) \\ &\rightarrow \exp\left(\int \nabla \cdot \mu\right) \quad \text{as } \Delta t \rightarrow 0 \end{aligned}$$

This leads to the same derivation for the instantaneous change of variable formula for continuous time flow (Chen et al., 2018), but the argument of μ will be the solution of the Itô integral, instead of the solution of the deterministic dynamics only.

D Marginal density of diffusion models (general case)

In Section 3, we assume σ is position-independent for simplicity. The general case of the Fokker-Planck equation can also be represented using the Feynmann-Kac formula. Following a similar conversion as in Table 2, we have

$$p(x, T) = \mathbb{E} \left[p_0(Y_T) \exp \left(\int_0^T -\nabla \cdot \mu(Y_s, T-s) + \sum_{i,j} \partial_{x_i, x_j}^2 D_{ij}(Y_s, T-s) ds \right) \middle| Y_0 = x \right]$$

where Y_s solves

$$dY = -\tilde{\mu}(Y, T-s) ds + \sigma(Y, T-s) dB'_s$$

where $\tilde{\mu}(y, s)_i := \mu_i(y, s) - 2 \sum_j \partial_{x_j} D_{ij}(y, s)$.

E One-dimensional explanation of Girsanov and variational inference

The Girsanov theorem (aka the Cameron–Martin–Girsanov theorem) describes how translation affects Wiener (Gaussian) measures. In Section 4, we deal with the infinite dimensional case, therefore demanding a formal measure-theoretic treatment. In this section, we use a one-dimensional case to illustrate how to interpret the Girsanov theorem and how we use it to derive the CT-ELBO, using the more familiar notion of probability density functions. Now imagine we do not have an

infinite-dimensional latent variable (*i.e.* the Brownian motion B'). Instead, imagine we have a one-dimensional latent variable ϵ' following a standard normal distribution. One can think about it as a VAE. This way, instead of having a classical Wiener measure (*i.e.* the distribution of Brownian motion) we would only need to deal with the standard Gaussian distribution, then \mathbb{P} in (13) has density $p = \mathcal{N}(0, 1)$. Suppose \mathbb{Q} also has density q . Then we can rewrite (13) using the more familiar density ratio

$$\mathbb{E}_q \left[\log \frac{p}{q} + \dots \mid \dots \right]$$

Recall $p(\epsilon') = \mathcal{N}(\epsilon'; 0, 1) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}\epsilon'^2}$. If we translate this density by a and let it be q , we have

$$q(\epsilon') = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}(\epsilon' - a)^2}$$

This definition of q gives us the density ratio

$$\frac{q(\epsilon')}{p(\epsilon')} = e^{a\epsilon' - \frac{1}{2}a^2}$$

Also, under the density q ,

$$\hat{\epsilon} := \epsilon' - a$$

is again a standard normal random variable (which means ϵ' is a Gaussian random variable with mean a). Note the striking resemblance between the last two formulas and (14,15).

Now if we want to use this q to perform inference as well as reparameterization, we simply just invert the standardization formula, by first sampling $\hat{\epsilon}$ from the standard normal distribution, and letting $\epsilon' = \hat{\epsilon} + a$. Under this reparameterization, the log-likelihood ratio $\log p/q$ in the ELBO becomes

$$-a\epsilon' + \frac{1}{2}a^2 = -a(\hat{\epsilon} + a) + \frac{1}{2}a^2 = -a\hat{\epsilon} - \frac{1}{2}a^2$$

Note that since $\hat{\epsilon}$ is the standard normal (under q), the first term is equal to 0 in expectation. This derivation leads to the CT-ELBO in (16). See Section F for the formal proof.

F Proofs

Theorem 3 (Continuous-time ELBO). *Let \mathbb{Q} be defined via the density (15). Then the RHS of (13) can be rewritten as*

$$\mathbb{E} \left[-\frac{1}{2} \int_0^T \|a(\omega, s)\|_2^2 ds + \log p_0(Y_T) - \int_0^T \nabla \cdot \mu ds \mid Y_0 = x \right] =: \mathcal{E}^\infty, \quad (16)$$

where the expectation is taken wrt the Brownian motion \hat{B}_s , and Y_s solves⁸

$$dY = (-\mu + \sigma a) ds + \sigma d\hat{B}_s. \quad (17)$$

Proof. By inverting the relationship (14), we have

$$dB'_s = d\hat{B}_s + a(\omega, s) ds$$

This allows us to reparameterize Y_s as

$$dY = -\mu ds + \sigma dB'_s = -\mu ds + \sigma(d\hat{B}_s + a ds) = (-\mu + \sigma a) ds + \sigma d\hat{B}_s$$

The log density can be written as

$$\begin{aligned} \log \frac{d\mathbb{P}}{d\mathbb{Q}} &= - \int_0^T a \cdot dB'_s + \frac{1}{2} \int_0^T \|a\|^2 ds \\ &= - \int_0^T a \cdot (d\hat{B}_s + a ds) + \frac{1}{2} \int_0^T \|a\|^2 ds \\ &= - \int_0^T a \cdot d\hat{B}_s - \frac{1}{2} \int_0^T \|a\|^2 ds \end{aligned}$$

⁸Note that μ and σ run backward in time from T , whereas a runs forward.

Finally, since the first term is in expectation equal to zero (Øksendal, 2003, Theorem 3.2.1), we conclude the proof. \square

Theorem 4 (Variational gap and optimal inference SDE). *The variational gap can be written as*

$$\log p(x, T) - \mathcal{E}^\infty = \int_0^T \mathbb{E} [\|a(\omega, s) - \sigma^\top \nabla \log p(Y_s, T - s)\|^2] ds. \quad (18)$$

In particular, $\mathcal{E}^\infty = \log p(x, T)$ if and only if $a(\omega, s)$ can be written as $a(\omega, s) = a(Y_s(\omega), s)$ for almost every $s \in [0, T]$ and $\omega \in \Omega$, and $a(y, s) = \sigma^\top \nabla \log p(y, T - s)$ almost everywhere.

Proof. To characterize the variational gap, we directly subtract the lower bound from the marginal likelihood:

$$\log p(x, T) - \mathcal{E}^\infty = \mathbb{E} \left[\log p(Y_0, T) - \log p(Y_T, 0) + \frac{1}{2} \int_0^T \|a(\omega, s)\|_2^2 ds + \int_0^T \nabla \cdot \mu ds \middle| Y_0 = x \right]$$

The first two terms can be written as an integral

$$\log p(Y_0, T) - \log p(Y_T, 0) = - \int_0^T d \log p(Y_s, T - s) \quad (35)$$

Using Itô's formula, we can rewrite the differential as

$$d \log p(Y_s, T - s) = - \frac{\partial_s p(Y_s, T - s)}{p(Y_s, T - s)} ds + \nabla \log p \cdot dY_s + \frac{1}{2} H_{\log p} : dY_s dY_s^\top$$

where $dY_s dY_s^\top = \sigma \sigma^\top ds$.

After rearrangement, we have

$$\int_0^T \left[\frac{\partial_s p}{p} - \nabla \log p^\top (-\mu + \sigma a) - \frac{1}{2} H_{\log p} : \sigma \sigma^\top + \frac{1}{2} \|a\|^2 + \nabla \cdot \mu \right] ds - \int_0^T \nabla \log p^\top \sigma d\hat{B}_s \quad (36)$$

where the second term is equal to 0 in expectation.

Now using the Fokker Planck equation to expand $\partial_s p$, with further rearrangement and cancellation and by a final application of the conditional Fubini's theorem, we end us with the desired characterization of the gap. \square

Theorem 5 (Consistency). *Assume $\mu, \sigma, \sigma^{-2}, a, \|a\|^2$ and their derivatives up to the fourth order are all bounded and continuous, and that σ is non-singular. Then $\mathcal{E}^L \rightarrow \mathcal{E}^\infty$ as $L \rightarrow \infty$.*

Proof. By definition of the log transitional distributions

$$\log p(x_{i+1} | x_i) = - \frac{d}{2} \log 2\pi - \log \det(\tilde{\sigma}_i) - \frac{1}{2} \|x_{i+1} - \tilde{\mu}_i(x_i)\|_{\tilde{\sigma}_i}^2 \quad (37)$$

Using the definition of $\tilde{\mu}_i$, the quadratic term becomes

$$\|x_{i+1} - x_i - \Delta t \mu(x_i, i \Delta t)\|_{\tilde{\sigma}_i}^2$$

Due the the Gaussian reparameterization (under q), we can write

$$\begin{aligned} x_i &= \hat{\mu}_{i+1}(x_{i+1}) + \hat{\sigma}_{i+1} \epsilon \\ &= x_{i+1} + \Delta t (-\mu(x_{i+1}, (i+1)\Delta t) \\ &\quad + \sigma((i+1)\Delta t) a(x_{i+1}, T - (i+1)\Delta t)) + \sqrt{\Delta t} \sigma((i+1)\Delta t) \epsilon \end{aligned} \quad (38)$$

Plugging this into the quadratic term yields

$$\begin{aligned} \|\cdots\|^2 &= \|\Delta t(\mu(x_{i+1}, (i+1)\Delta t) - \mu(x_i, i\Delta t)) \\ &\quad - \Delta t\sigma((i+1)\Delta t)a(x_{i+1}, T - (i+1)\Delta t) - \sqrt{\Delta t}\sigma((i+1)\Delta t)\epsilon\|^2 \end{aligned} \quad (39)$$

We take care of the deviation in μ first, by taking the Taylor expansion around $(x_i, i\Delta t)$:

$$\mu(x_{i+1}, (i+1)\Delta t) = \mu(x_i, i\Delta t) + \nabla\mu(x_i, i\Delta t)^\top (x_{i+1} - x_i) + \mathcal{O}(\Delta t) \quad (40)$$

Note that the first order term wrt the time variable is also $\mathcal{O}(\Delta t)$, so it's absorbed into the remainder. Combining the last three identities, we have

$$\begin{aligned} &\frac{1}{2}\|x_{i+1} - \tilde{\mu}_i(x_i)\|_{\hat{\sigma}_i}^2 \\ &= \frac{1}{2}\epsilon^\top \sigma^\top (\sigma\sigma^\top)^{-1} \sigma\epsilon + \Delta t\epsilon^\top \sigma^\top \nabla\mu^\top (\sigma\sigma^\top)^{-1} \sigma\epsilon + \frac{1}{2}\Delta t a^\top \sigma^\top (\sigma\sigma^\top)^{-1} \sigma a \\ &\quad + o(\Delta t) + (\Delta t)^{1/2}\epsilon^\top \sigma^\top (\sigma\sigma^\top)^{-1} \sigma a \end{aligned} \quad (41)$$

Note that we've dropped the arguments of the functions for notational convenience. All the σ s in the denominator are $\sigma(i\Delta t)$. The $o(\Delta t)$ term can be neglected since it decays fast enough even though there are $L = 1/\Delta t$ of them. The last term is 0 in expectation since ϵ is Gaussian distributed. To take care of the first term (*), we turn to the log density of the inference model.

$$\log q(x_i|x_{i+1}) = -\frac{d}{2}\log 2\pi - \log \det(\hat{\sigma}_{i+1}) - \frac{1}{2}\|x_i - \hat{\mu}_{i+1}(x_{i+1})\|_{\hat{\sigma}_{i+1}}^2 \quad (42)$$

$$= -\frac{d}{2}\log 2\pi - \log \det(\hat{\sigma}_{i+1}) - \frac{1}{2}\|\hat{\sigma}_{i+1}\epsilon\|_{\hat{\sigma}_{i+1}}^2 \quad (43)$$

Comparing the third term with (*), we have

$$\frac{1}{2}\epsilon^\top \sigma_{i+1}^\top ((\sigma_{i+1}\sigma_{i+1}^\top)^{-1} - (\sigma_i\sigma_i^\top)^{-1}) \sigma_{i+1}\epsilon, \quad (44)$$

where $\sigma_i := \sigma(i\Delta)$. Using the differential notation, in expectation, the above can be rewritten as

$$\mathbb{E}\left[\frac{1}{2}\epsilon^\top \sigma^\top (\partial_t(\sigma\sigma^\top)^{-1}) \sigma\epsilon\right] dt = -\text{tr}(\sigma^{-1}\partial_t\sigma) dt = -\partial_t \log \det(\sigma) dt, \quad (45)$$

where we used Hutchinson's trace identity and Jacobi's formula. Therefore, the summation of the differences will converge to $\log \det(\sigma(0)) - \log \det(\sigma(T))$. This quantity will be negated by summing up the differences between the normalizing constants for all L terms, which gives us $\log \det(\sigma(T)) - \log \det(\sigma(0))$, by the telescoping cancellation.

Now we only have two terms from the quadratic function, which will converge to

$$\epsilon^\top \sigma^\top \nabla\mu^\top \sigma^{-\top} \epsilon dt + \frac{1}{2}\|a\|^2 dt$$

Using the trace identity again, and the fact that trace is similarity-invariant, we see that the above quantity is equal to

$$\left(\nabla \cdot \mu + \frac{1}{2}\|a\|^2\right) dt$$

in expectation. Now summing up all the layers, we can decompose the approximate error as

$$\begin{aligned} &\left|\mathbb{E}\left[\sum \log \frac{p}{q}\right] - \mathbb{E}\left[-\int \left(\nabla \cdot \mu + \frac{1}{2}\|a\|^2\right)\right]\right| \leq \left|\mathbb{E}\left[\sum \log \frac{p}{q} + \sum \left(\nabla \cdot \mu + \frac{1}{2}\|a\|^2\right) \Delta t\right]\right| \\ &\quad + \mathbb{E}\left[\left|\sum \left(\nabla \cdot \mu + \frac{1}{2}\|a\|^2\right) \Delta t - \int \left(\nabla \cdot \mu + \frac{1}{2}\|a\|^2\right)\right|\right] \end{aligned}$$

As all the approximation errors are bounded and converge to 0 as $L \rightarrow \infty$, the first term goes to 0 by the *Dominated Convergence Theorem*. The assumption on the coefficients also guarantees the convergence in mean square error (Milshtein, 1975) of the Euler Maruyama scheme, which implies the second term goes to 0. The same applies to the last step for the prior term: $x_0 \rightarrow y(T)$ in L^2 .

□

G Equivalent SDEs

We use the following definition to formalize what we mean by equivalent SDEs

Definition 1 (Equivalent processes / SDEs). Let Y_s, \tilde{Y}_s and X_t be stochastic processes for $0 \leq s, t \leq T$. If Y_s and \tilde{Y}_s have the same distribution for all s , then they are said to be equivalent. If X_t and \tilde{Y}_{T-t} have the same distribution for all t , then we say X_t is an equivalent reverse process. Two SDEs are equivalent if the processes they induce are equivalent. Two SDEs are equivalent reverse of each other if the processes they induce are equivalent reverse of one another.

Note that when talking about the equivalency between SDEs, the dependency on an initial condition is implied.

In this section, we show how to construct a family of equivalent (reverse) SDEs. Let Y_s be a diffusion process solving

$$dY = f ds + g d\hat{B}_s$$

We assume g is position-independent and diagonal for simplicity. Let $\lambda \leq 1$, We can rearrange the Fokker-Planck equation to get

$$\partial_s q = -\nabla \cdot (fq) + \frac{1}{2}g^2 : H_q = -\nabla \cdot \left(\left(f - \frac{\lambda}{2}g^2 \nabla \log q \right) q \right) + \frac{1-\lambda}{2}g^2 : H_q \quad (46)$$

Now let $f_\lambda := f - \frac{\lambda}{2}g^2 \nabla \log q$, and $g_\lambda := \sqrt{1-\lambda}g$. Then the SDE $dY = f_\lambda ds + g_\lambda d\hat{B}_s$ has the same Fokker Planck equation as (46), which means the SDEs defined this way form a family of equivalent SDEs⁹.

Now to construct an equivalent reverse SDE, we rearrange the Fokker Planck of this new SDE,

$$\partial_s q = -\nabla \cdot (f_\lambda q) + \frac{1}{2}g_\lambda^2 : H_q = -\nabla \cdot \left((f_\lambda - g_\lambda^2 \nabla \log q) q \right) - \frac{1}{2}g_\lambda^2 : H_q \quad (47)$$

Now let $\mu_\lambda(x, t) := g_\lambda^2(x, T-t) \nabla \log q(x, T-t) - f_\lambda(x, T-t)$ and $\sigma_\lambda = g_\lambda(x, T-t)$. Then the SDE $dX = \mu_\lambda dt + \sigma_\lambda dB_t$ with the initial condition $X_0 \sim q(\cdot, T)$ is an equivalent reverse SDE, since

$$\partial_t p = -\nabla \cdot (\mu_\lambda p) + \frac{1}{2}\sigma_\lambda^2 : H_p = \nabla \cdot \left((f_\lambda - g_\lambda^2 \nabla \log q) p \right) + \frac{1}{2}g_\lambda^2 : H_p \quad (48)$$

is the time-reversal of (47). This also means there is a family of plug-in reverse SDEs parameterized by λ and \mathfrak{s}_θ :

$$dX = (g_\lambda^2 \mathfrak{s}_\theta - f_\lambda) dt + \sigma_\lambda dB_t \quad (49)$$

$$= \left(\left(1 - \frac{\lambda}{2} \right) g^2 \mathfrak{s}_\theta - f \right) dt + \sqrt{1-\lambda}g dB_t \quad (50)$$

The plug-in reverse SDE used by Song et al. (2021) corresponds to $\lambda = 0$, and the equivalent (plug-in) reverse ODE corresponds to $\lambda = 1$. See Figure 5 for the simulation.

H Score matching and plug-in reverse SDEs

In Section 6 we establish the connection between the score matching loss and the CT-ELBO of the plug-in reverse SDE for $\lambda = 0$. If we want to do the same for different values of λ , we need to make sure the generative and inference SDEs have the same diffusion coefficient (this is to make sure the Radon-Nikodym derivative is finite). In light of this, we define the following generative and inference pair

$$dX = \left(\left(1 - \frac{\lambda}{2} \right) g^2 \mathfrak{s}_\theta - f \right) dt + \sqrt{1-\lambda}g dB_t \quad \text{and} \quad dY = \left(f - \frac{\lambda}{2}g^2 \nabla \log q \right) ds + \sqrt{1-\lambda}g d\hat{B}_s \quad (27)$$

Note that this is just the same equivalent SDE and equivalent (plug-in) reverse SDE from the Appendix G. We show that maximizing the ELBO of this family of plug-in reverse SDEs is also equivalent to performing score matching.

⁹Note that more generally the same would also hold if we let λ be a time-dependent function.




Figure 5: Samples from plug-in reverse SDEs with different λ values (rows).

Theorem 7 (Plug-in reverse SDE ELBO). Assume the generative and inference SDEs follow (27). For $\lambda < 1$, then the CT-ELBO (denoted by $\mathcal{E}_\lambda^\infty$) can be written as

$$\begin{aligned} \mathcal{E}_\lambda^\infty = \mathbb{E}_{Y_T}[\log p_0(Y_T) | Y_0 = x] &- \int_0^T \left(1 - \frac{\lambda}{2}\right) \mathbb{E}_{Y_s} \left[\frac{1}{2} \|\mathbf{s}_\theta\|_{g^2}^2 + \nabla \cdot \left(g^2 \mathbf{s}_\theta - \left(\frac{2}{2-\lambda} \right) f \right) \middle| Y_0 = x \right] \\ &+ \frac{\lambda}{2} \mathbb{E}_{Y_s} \left[\frac{1}{2} \|\mathbf{s}_\theta\|_{g^2}^2 - g^2 \mathbf{s}_\theta^\top \nabla \log q(Y_s, s) \middle| Y_0 = x \right] \\ &+ \frac{\lambda^2}{4(1-\lambda)} \mathbb{E}_{Y_s} \left[\frac{1}{2} \|\mathbf{s}_\theta - \nabla \log q(Y_s, s)\|_{g^2}^2 \middle| Y_0 = x \right] ds \end{aligned}$$

As a result, averaging the ELBO over the data distribution and applying the identity (3) yield

$$\mathbb{E}_{Y_0}[\mathcal{E}_\lambda^\infty] = \mathbb{E}_{Y_T}[\log p_0(Y_T)] - \int_0^T \left(1 + \frac{\lambda^2}{4(1-\lambda)}\right) \mathbb{E}_{Y_s} \left[\frac{1}{2} \|\mathbf{s}_\theta\|_{g^2}^2 + \nabla \cdot (g^2 \mathbf{s}_\theta) \right] ds + \text{Const.} \quad (51)$$

$$= \mathbb{E}_{Y_0}[\mathcal{E}_0^\infty] - \left(\frac{\lambda^2}{4(1-\lambda)} \right) \int_0^T \mathbb{E}_{Y_s} \left[\frac{1}{2} \|\mathbf{s}_\theta(Y_s, s) - \nabla \log q(Y_s, s)\|_{g^2}^2 \right] ds \quad (52)$$

Before proving this theorem, we first make a few remarks. First, setting $\lambda = 0$, this ELBO will reduce to (25). Second, (51) tells us that while matching the score, we implicitly maximize the likelihood of the entire family of plug-in reverse SDEs. Third, (52) tells us that the average CT-ELBO is maximized when $\lambda = 0$ (recall Figure 2). Lastly, the theorem excludes the case where $\lambda = 1$, *i.e.* the equivalent ODE, since otherwise there will be a division-by-zero problem. But an ODE can be seen as having λ very close to 1, which will make the SDE effectively deterministic in practice. This explains the low BPD of the equivalent plug-in ODE reported in Song et al. (2021).

Proof. Plugging (27) in (4) and (17), we get

$$\begin{aligned}\mu &= \left(1 - \frac{\lambda}{2}\right) g^2 \mathbf{s}_\theta - f \\ \sigma &= \sqrt{1 - \lambda} g \\ a &= \frac{1}{\sqrt{1 - \lambda}} \left[(1 - \lambda) g \mathbf{s}_\theta + \frac{\lambda}{2} g (\mathbf{s}_\theta - \nabla \log q) \right]\end{aligned}$$

Then we have

$$\begin{aligned}\frac{1}{2} \|a\|_2^2 &= \frac{1}{2(1 - \lambda)} \left[(1 - \lambda)^2 \|\mathbf{s}_\theta\|_{g^2}^2 + (1 - \lambda) \lambda g^2 \mathbf{s}_\theta^\top (\mathbf{s}_\theta - \nabla \log q) + \frac{\lambda^2}{4} \|\mathbf{s}_\theta - \nabla \log q\|_{g^2}^2 \right] \\ &= \left(1 - \frac{\lambda}{2}\right) \frac{1}{2} \|\mathbf{s}_\theta\|_{g^2}^2 + \frac{\lambda}{2} \left(\frac{1}{2} \|\mathbf{s}_\theta\|_{g^2}^2 - g^2 \mathbf{s}_\theta^\top \nabla \log q \right) + \frac{\lambda^2}{4(1 - \lambda)} \frac{1}{2} \|\mathbf{s}_\theta - \nabla \log q\|_{g^2}^2\end{aligned}$$

$$\nabla \cdot \mu = \left(1 - \frac{\lambda}{2}\right) \nabla \cdot \left(g^2 \mathbf{s}_\theta - \left(\frac{2}{2 - \lambda} \right) f \right)$$

Summing up these two parts gives us $\mathcal{E}_\lambda^\infty$. Under the expectation, we can rewrite $\mathbb{E}_{Y_s} [g^2 \mathbf{s}_\theta^\top \nabla \log q] = -\mathbb{E}_{Y_s} [\nabla \cdot (g^2 \mathbf{s}_\theta)]$ using the score matching loss identity (see Appendix A), to obtain the second part of the statement. \square

I Non-uniform sampling for debiasing

We perform non-uniform sampling to debias the denoising score matching loss weighted by σ_s^2/g^2 , as discussed in subsection 6.2. We experiment with the variance-preserving SDE from Song et al. (2021) (originally from Ho et al. (2020)), whose drift and diffusion coefficients are

$$f(y, s) = -\frac{1}{2} \beta(s) y \quad (53)$$

$$g(y, s) = g(s) = \sqrt{\beta(s)} \quad (54)$$

where $\beta(s) = (\beta_{\max} - \beta_{\min})s + \beta_{\min}$, for some constants β_{\max} and β_{\min} .

Solving the Fokker Planck of this SDE with a Dirac point mass as initial condition gives us a conditional Gaussian, whose variance is

$$\sigma_s^2 := \int_0^s g^2(s') ds' = \frac{1}{2} s^2 (\beta_{\max} - \beta_{\min}) + s \beta_{\min} \quad (55)$$

Our goal is to sample from a density function proposal to g^2/σ_s^2 for most of the part. So for some small $s_\epsilon > 0$, we define the following unnormalized density

$$\tilde{q}_\epsilon(s) = \begin{cases} \frac{g^2(s_\epsilon)}{\sigma_{s_\epsilon}^2} & s \in [0, s_\epsilon] \\ \frac{g^2(s)}{\sigma_s^2} & s \in [s_\epsilon, T] \end{cases} \quad (56)$$

To simplify our notation, we let

$$\phi(s) := \log \left(\exp \left(\frac{1}{2} s^2 (\beta_{\max} - \beta_{\min}) s \beta_{\min} \right) - 1 \right) \quad (57)$$

$$\varphi(u) := \log \left(1 + \exp \left(Z u + \phi(s_\epsilon) - \frac{g^2(s_\epsilon)}{\sigma_{s_\epsilon}^2} s_\epsilon \right) \right) \quad (58)$$

$$\tilde{\Phi}_\epsilon(s) := \begin{cases} \frac{g^2(s_\epsilon)}{\sigma_{s_\epsilon}^2} s & s \in [0, s_\epsilon] \\ \frac{g^2(s_\epsilon)}{\sigma_{s_\epsilon}^2} s_\epsilon + \phi(s) - \phi(s_\epsilon) & s \in [s_\epsilon, T] \end{cases} \quad (59)$$

where $\tilde{\Phi}_\epsilon$ is the cumulative function of the unnormalized density. Evaluating it at T gives us the normalizing constant $Z = \tilde{\Phi}_\epsilon(T)$, from which we obtain the CDF, $\Phi_\epsilon(s) = \frac{\tilde{\Phi}_\epsilon(s)}{Z}$, the pdf $q_\epsilon(s) = \frac{\tilde{q}_\epsilon(s)}{Z}$, and the inverse CDF that we need for sampling (using the inverse CDF transform):

$$\Phi_\epsilon^{-1}(u) = \begin{cases} Z \frac{\sigma^2(s_\epsilon)}{g^2(s_\epsilon)} u & u \in \left[0, s_\epsilon \frac{g^2(s_\epsilon)}{Z\sigma_{s_\epsilon}^2}\right) \\ \frac{1}{\beta_{\max} - \beta_{\min}} \left(-\beta_{\min} + \sqrt{\beta_{\min}^2 + 2(\beta_{\max} - \beta_{\min})\varphi(u)}\right) & u \in \left[s_\epsilon \frac{g^2(s_\epsilon)}{Z\sigma_{s_\epsilon}^2}, 1\right] \end{cases} \quad (60)$$

J Experiments

J.1 MNIST and CIFAR 10

We use the variance preserving SDE described in Appendix I, with $\beta_{\min} = 0.1$, $\beta_{\max} = 20$, and $T = 1$. We use the same architecture following Ho et al. (2020) for the CIFAR10 experiment (which is a modified U-Net (Ronneberger et al., 2015)). For MNIST, we use 3 feature map resolutions (instead of 4) and reduce the number of channels from 128 to 32. Also we did not apply dropout.

For optimization, we use the Adam optimizer with a learning rate of 0.0001. We use minibatch size 128 for all experiments. We apply the standard uniform dequantization, and map the data to the real space using the logit transform (with a squeeze coefficient $\alpha = 0.05$ to avoid numerical instability). For CIFAR10, we additionally apply random horizontal flipping for regularization.

More details can be found in <https://github.com/CW-Huang/sdeflow-light>.