
Maximum Likelihood Training of Score-Based Diffusion Models

Yang Song*

Computer Science Department
Stanford University
yangsong@cs.stanford.edu

Conor Durkan*

School of Informatics
University of Edinburgh
conor.durkan@ed.ac.uk

Iain Murray

School of Informatics
University of Edinburgh
i.murray@ed.ac.uk

Stefano Ermon

Computer Science Department
Stanford University
ermon@cs.stanford.edu

Abstract

Score-based diffusion models synthesize samples by reversing a stochastic process that diffuses data to noise, and are trained by minimizing a weighted combination of score matching losses. The log-likelihood of score-based diffusion models can be tractably computed through a connection to continuous normalizing flows, but log-likelihood is not directly optimized by the weighted combination of score matching losses. We show that for a specific weighting scheme, the objective upper bounds the negative log-likelihood, thus enabling approximate maximum likelihood training of score-based diffusion models. We empirically observe that maximum likelihood training consistently improves the likelihood of score-based diffusion models across multiple datasets, stochastic processes, and model architectures. Our best models achieve negative log-likelihoods of 2.83 and 3.76 bits/dim on CIFAR-10 and ImageNet 32×32 without any data augmentation, on a par with state-of-the-art autoregressive models on these tasks.

1 Introduction

Score-based generative models [44, 45, 48] and diffusion probabilistic models [43, 19] have recently achieved state-of-the-art sample quality in a number of tasks, including image generation [48, 11], audio synthesis [5, 27, 37], and shape generation [3]. Both families of models perturb data with a sequence of noise distributions, and generate samples by learning to reverse this path from noise to data. Through stochastic calculus, these approaches can be unified into a single framework [48] which we refer to as score-based diffusion models in this paper.

The framework of score-based diffusion models [48] involves gradually diffusing the data distribution towards a given noise distribution using a stochastic differential equation (SDE), and learning the time reversal of this SDE for sample generation. Crucially, the reverse-time SDE has a closed-form expression which depends solely on the time-dependent gradient field (a.k.a., score) of the perturbed data distribution. This gradient field can be efficiently estimated by training a neural network (called a score-based model [44, 45]) with a weighted combination of score matching losses [23, 56, 46] as the objective. A key advantage of score-based diffusion models is that they can be transformed into continuous normalizing flows (CNFs) [6, 15], thus allowing tractable likelihood computation with numerical ODE solvers.

*Equal contribution.

Compared to vanilla CNFs, score-based diffusion models are much more efficient to train. This is because the maximum likelihood objective for training CNFs requires running an expensive ODE solver for every optimization step, while the weighted combination of score matching losses for training score-based models does not. However, unlike maximum likelihood training, minimizing a combination of score matching losses does not necessarily lead to better likelihood values. Since better likelihoods are useful for applications including compression [21, 20, 51], semi-supervised learning [10], adversarial purification [47], and comparing against likelihood-based generative models, we seek a training objective for score-based diffusion models that is as efficient as score matching but also promotes higher likelihoods.

We show that such an objective can be readily obtained through slight modification of the weighted combination of score matching losses. Our theory reveals that with a specific choice of weighting, which we term the *likelihood weighting*, the combination of score matching losses actually upper bounds the negative log-likelihood. We further prove that this upper bound becomes tight when our score-based model corresponds to the true time-dependent gradient field of a certain reverse-time SDE. Using likelihood weighting increases the variance of our objective, which we counteract by introducing a variance reduction technique based on importance sampling. Our bound is analogous to the classic evidence lower bound used for training latent-variable models in the variational autoencoding framework [26, 39], and can be viewed as a continuous-time generalization of [43].

With our likelihood weighting, we can minimize the weighted combination of score matching losses for approximate maximum likelihood training of score-based diffusion models. Compared to weightings in previous work [48], we consistently improve likelihood values across multiple datasets, model architectures, and SDEs, with only slight degradation of Fréchet Inception distances [17]. Moreover, our upper bound on negative log-likelihood allows training with variational dequantization [18], with which we reach negative log-likelihood of **2.83** bits/dim on CIFAR-10 [28] and **3.76** bits/dim on ImageNet 32×32 [55] with no data augmentation. Our models present the first instances of normalizing flows which achieve comparable likelihood to cutting-edge autoregressive models.

2 Score-based diffusion models

Score-based diffusion models are deep generative models that smoothly transform data to noise using a diffusion process, and synthesize samples by learning and simulating the time-reversal of this diffusion. The overall approach is illustrated in Fig. 1.

2.1 Diffusing data to noise with an SDE

Let $p(\mathbf{x})$ denote the unknown distribution of a dataset consisting of D -dimensional i.i.d. samples. Score-based diffusion models [48] employ a stochastic differential equation (SDE) to diffuse $p(\mathbf{x})$ towards a noise distribution. The SDEs are of the form

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

where $\mathbf{f}(\cdot, t) : \mathbb{R}^D \rightarrow \mathbb{R}^D$ is the drift coefficient, $g(t) \in \mathbb{R}$ is the diffusion coefficient, and $\mathbf{w} \in \mathbb{R}^D$ denotes a standard Wiener process (a.k.a., Brownian motion). Intuitively, we can interpret $d\mathbf{w}$ as infinitesimal Gaussian noise. The solution of an SDE is a diffusion process $\{\mathbf{x}(t)\}_{t \in [0, T]}$, where $[0, T]$ is a fixed time horizon. We let $p_t(\mathbf{x})$ denote the marginal distribution of $\mathbf{x}(t)$, and $p_{0t}(\mathbf{x}' | \mathbf{x})$ denote the transition distribution from $\mathbf{x}(0)$ to $\mathbf{x}(t)$. Note that by definition we always have $p_0 = p$ when using an SDE to perturb the data distribution.

The role of the SDE is to smooth the data distribution by adding noise, gradually removing structure until little of the original signal remains. In the framework of score-based diffusion models, we choose $\mathbf{f}(\mathbf{x}, t)$, $g(t)$, and T such that the diffusion process $\{\mathbf{x}(t)\}_{t \in [0, T]}$ approaches some analytically tractable prior distribution $\pi(\mathbf{x})$ at $t = T$, meaning $p_T(\mathbf{x}) \approx \pi(\mathbf{x})$. Three families of SDEs suitable for this task are outlined in [48], namely Variance Exploding (VE) SDEs, Variance Preserving (VP) SDEs, and subVP SDEs.

2.2 Generating samples with the reverse SDE

Sample generation in score-based diffusion models relies on time-reversal of the diffusion process. For well-behaved drift and diffusion coefficients, the forward diffusion described in Eq. (1) has an




Figure 1: We can use an SDE to diffuse data to a simple noise distribution. This SDE can be reversed once we know the score of the marginal distribution at each intermediate time step, $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$.

associated reverse-time diffusion process [1, 16] given by the following SDE

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

where $\bar{\mathbf{w}}$ is now a standard Wiener process in the reverse-time direction. Here dt represents an infinitesimal negative time step, meaning that the above SDE must be solved from $t = T$ to $t = 0$. This reverse-time SDE results in exactly the same diffusion process $\{\mathbf{x}(t)\}_{t \in [0, T]}$ as Eq. (1), assuming it is initialized with $\mathbf{x}(T) \sim p_T(\mathbf{x})$. This result allows for the construction of diffusion-based generative models, and its functional form reveals the key target for learning: the time-dependent score function $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$. Again, see Fig. 1 for a helpful visualization of this two-part formulation.

In order to estimate $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ from a given dataset, we fit the parameters of a neural network $\mathbf{s}_{\theta}(\mathbf{x}, t)$, termed a *score-based model*, such that $\mathbf{s}_{\theta}(\mathbf{x}, t) \approx \nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ for almost all $\mathbf{x} \in \mathbb{R}^D$ and $t \in [0, T]$. Unlike many likelihood-based generative models, a score-based model does not need to satisfy the integral constraints of a density function, and is therefore much easier to parameterize. Good score-based models should keep the following least squares loss small

$$\mathcal{J}_{\text{SM}}(\boldsymbol{\theta}; \lambda(\cdot)) := \frac{1}{2} \int_0^T \mathbb{E}_{p_t(\mathbf{x})} [\lambda(t) \|\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) - \mathbf{s}_{\theta}(\mathbf{x}, t)\|_2^2] dt, \quad (3)$$

where $\lambda: [0, T] \rightarrow \mathbb{R}_{>0}$ is a positive weighting function. The integrand features the well-known score matching [23] objective $\mathbb{E}_{p_t(\mathbf{x})} [\|\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) - \mathbf{s}_{\theta}(\mathbf{x}, t)\|_2^2]$. We therefore refer to Eq. (3) as a weighted combination of score matching losses.

With score matching techniques [56, 46], we can compute Eq. (3) up to an additive constant and minimize it for training score-based models. For example, we can use denoising score matching [56] to transform $\mathcal{J}_{\text{SM}}(\boldsymbol{\theta}; \lambda(\cdot))$ into the following, which is equivalent up to a constant independent of $\boldsymbol{\theta}$:

$$\mathcal{J}_{\text{DSM}}(\boldsymbol{\theta}; \lambda(\cdot)) := \frac{1}{2} \int_0^T \mathbb{E}_{p(\mathbf{x}) p_{0t}(\mathbf{x}' | \mathbf{x})} [\lambda(t) \|\nabla_{\mathbf{x}'} \log p_{0t}(\mathbf{x}' | \mathbf{x}) - \mathbf{s}_{\theta}(\mathbf{x}', t)\|_2^2] dt. \quad (4)$$

Whenever the drift coefficient $\mathbf{f}_{\theta}(\mathbf{x}, t)$ is linear in \mathbf{x} (which is true for all SDEs in [48]), the transition density $p_{0t}(\mathbf{x}' | \mathbf{x})$ is a tractable Gaussian distribution. We can form a Monte Carlo estimate of both the time integral and expectation in $\mathcal{J}_{\text{DSM}}(\boldsymbol{\theta}; \lambda(\cdot))$ with a sample $(t, \mathbf{x}, \mathbf{x}')$, where t is uniformly drawn from $[0, T]$, $\mathbf{x} \sim p(\mathbf{x})$ is a sample from the dataset, and $\mathbf{x}' \sim p_{0t}(\mathbf{x}' | \mathbf{x})$. The gradient $\nabla_{\mathbf{x}'} \log p_{0t}(\mathbf{x}' | \mathbf{x})$ can also be computed in closed form since $p_{0t}(\mathbf{x}' | \mathbf{x})$ is Gaussian.

After training a score-based model $\mathbf{s}_{\theta}(\mathbf{x}, t)$ with $\mathcal{J}_{\text{DSM}}(\boldsymbol{\theta}; \lambda(\cdot))$, we can plug it into the reverse-time SDE in Eq. (2). Samples are then generated by solving this reverse-time SDE with numerical SDE solvers, given an initial sample from $\pi(\mathbf{x})$ at $t = T$. Since the forward SDE Eq. (1) is designed such that $p_T(\mathbf{x}) \approx \pi(\mathbf{x})$, the reverse-time SDE will closely trace the diffusion process given by Eq. (1) in the reverse time direction, and yield an approximate data sample at $t = 0$ (as visualized in Fig. 1).

3 Likelihood of score-based diffusion models

The forward and backward diffusion processes in score-based diffusion models induce two probabilistic models for which we can define a likelihood. The first probabilistic model, denoted as $p_{\theta}^{\text{SDE}}(\mathbf{x})$, is given by the approximate reverse-time SDE constructed from our score-based model $\mathbf{s}_{\theta}(\mathbf{x}, t)$. In particular, suppose $\{\hat{\mathbf{x}}_{\theta}(t)\}_{t \in [0, T]}$ is a stochastic process given by

$$d\hat{\mathbf{x}} = [\mathbf{f}(\hat{\mathbf{x}}, t) - g(t)^2 \mathbf{s}_{\theta}(\hat{\mathbf{x}}, t)] dt + g(t) d\bar{\mathbf{w}}, \quad \hat{\mathbf{x}}_{\theta}(T) \sim \pi. \quad (5)$$

We define p_{θ}^{SDE} as the marginal distribution of $\hat{\mathbf{x}}_{\theta}(0)$. The probabilistic model p_{θ}^{SDE} is jointly defined by the score-based model $\mathbf{s}_{\theta}(\mathbf{x}, t)$, the prior π , plus the drift and diffusion coefficients of the forward SDE in Eq. (1). We can obtain a sample $\hat{\mathbf{x}}_{\theta}(0) \sim p_{\theta}^{\text{SDE}}$ by numerically solving the reverse-time SDE in Eq. (5) with an initial noise vector $\hat{\mathbf{x}}_{\theta}(T) \sim \pi$.

The other probabilistic model, denoted $p_{\theta}^{\text{ODE}}(\mathbf{x})$, is derived from the SDE’s associated *probability flow ODE* [32, 48]. Every SDE has a corresponding probability flow ODE whose marginal distribution at each time t matches that of the SDE, so that they share the same $p_t(\mathbf{x})$ for all time. In particular, the ODE corresponding to the SDE in Eq. (1) is given by

$$\frac{d\mathbf{x}}{dt} = \mathbf{f}(\mathbf{x}, t) - \frac{1}{2}g(t)^2 \nabla_{\mathbf{x}} \log p_t(\mathbf{x}). \quad (6)$$

Unlike the SDEs in Eq. (1) and Eq. (2), this ODE describes fully deterministic dynamics for the process. Notably, it still features the same time-dependent score function $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$. By approximating this score function with our model $\mathbf{s}_{\theta}(\mathbf{x}, t)$, the probability flow ODE becomes

$$\frac{d\tilde{\mathbf{x}}}{dt} = \mathbf{f}(\tilde{\mathbf{x}}, t) - \frac{1}{2}g(t)^2 \mathbf{s}_{\theta}(\tilde{\mathbf{x}}, t). \quad (7)$$

In fact, this ODE is an instance of a continuous normalizing flow (CNF) [15], and we can quantify how the ODE dynamics transform volumes across time in exactly the same way as these traditional flow-based models [6]. Given a prior distribution $\pi(\mathbf{x})$, and a trajectory function $\tilde{\mathbf{x}}_{\theta}: [0, T] \rightarrow \mathbb{R}^D$ satisfying the ODE in Eq. (7), we define p_{θ}^{ODE} as the marginal distribution of $\tilde{\mathbf{x}}_{\theta}(0)$ when $\tilde{\mathbf{x}}_{\theta}(T) \sim \pi$. Similarly to p_{θ}^{SDE} , the model p_{θ}^{ODE} is jointly defined by the score-based model $\mathbf{s}_{\theta}(\mathbf{x}, t)$, the prior π , and the forward SDE in Eq. (1). Leveraging the instantaneous change-of-variables formula [6], we can evaluate $\log p_{\theta}^{\text{ODE}}(\mathbf{x})$ exactly with numerical ODE solvers. Since p_{θ}^{ODE} is a CNF, we can generate a sample $\tilde{\mathbf{x}}_{\theta}(0) \sim p_{\theta}^{\text{ODE}}$ by numerically solving the ODE in Eq. (7) with an initial value $\tilde{\mathbf{x}}_{\theta}(T) \sim \pi$.

Although computing $\log p_{\theta}^{\text{ODE}}(\mathbf{x})$ is tractable, training p_{θ}^{ODE} with maximum likelihood will require calling an ODE solver for every optimization step [6, 15], which can be prohibitively expensive for large-scale score-based models. Unlike p_{θ}^{ODE} , we cannot evaluate $\log p_{\theta}^{\text{SDE}}(\mathbf{x})$ exactly for an arbitrary data point \mathbf{x} . However, we have a lower bound on $\log p_{\theta}^{\text{SDE}}(\mathbf{x})$ which allows both efficient evaluation and optimization, as will be shown in Section 4.2.

4 Bounding the likelihood of score-based diffusion models

Many applications benefit from models which achieve high likelihood. One example is lossless compression, where log-likelihood directly corresponds to the minimum expected number of bits needed to encode a message. Indeed, popular likelihood-based models such as variational autoencoders and normalizing flows have already found success in image compression [51, 20, 21]. Despite some known drawbacks [50], likelihood is still one of the most popular metrics for evaluating and comparing generative models.

Maximizing the likelihood of score-based diffusion models can be accomplished by either maximizing the likelihood of p_{θ}^{SDE} or p_{θ}^{ODE} . Although p_{θ}^{ODE} is a continuous normalizing flow (CNF) and its log-likelihood is tractable, training with maximum likelihood is expensive. As mentioned already, it requires solving an ODE at every optimization step in order to evaluate the log-likelihood on a batch of training data. In contrast, training with the weighted combination of score matching losses is much more efficient, yet in general it does not directly promote high likelihood of either p_{θ}^{SDE} or p_{θ}^{ODE} .

In what follows, we show that with a specific choice of the weighting function $\lambda(t)$, the combination of score matching losses $\mathcal{J}_{\text{SM}}(\theta; \lambda(\cdot))$ actually becomes an upper bound on $D_{\text{KL}}(p \parallel p_{\theta}^{\text{SDE}})$, and can therefore serve as an efficient proxy for maximum likelihood training. In addition, we provide a related lower bound on $\log p_{\theta}^{\text{SDE}}(\mathbf{x})$ that can be evaluated efficiently on any individual datapoint \mathbf{x} .

4.1 Bounding the KL divergence with likelihood weighting

It is well-known that maximizing the log-likelihood of a probabilistic model is equivalent to minimizing the KL divergence from the data distribution to the model distribution. We show in the following theorem that for the model p_{θ}^{SDE} , this KL divergence can be upper bounded by $\mathcal{J}_{\text{SM}}(\theta; \lambda(\cdot))$ when using the weighting function $\lambda(t) = g(t)^2$, where $g(t)$ is the diffusion coefficient of SDE in Eq. (1).

Table 1: SDEs and their corresponding weightings for score matching losses.

SDE	Formula	$\lambda(t)$ in [48]	likelihood weighting
VE	$d\mathbf{x} = \sigma(t) d\mathbf{w}$	$\sigma^2(t)$	$\sigma^2(t)$
VP	$d\mathbf{x} = -\frac{1}{2}\beta(t)\mathbf{x} dt + \sqrt{\beta(t)} d\mathbf{w}$	$1 - e^{-\int_0^t \beta(s) ds}$	$\beta(t)$
subVP	$d\mathbf{x} = -\frac{1}{2}\beta(t)\mathbf{x} dt + \sqrt{\beta(t)(1 - e^{-2\int_0^t \beta(s) ds})} d\mathbf{w}$	$(1 - e^{-\int_0^t \beta(s) ds})^2$	$\beta(t)(1 - e^{-2\int_0^t \beta(s) ds})$

Theorem 1. Let $p(\mathbf{x})$ be the data distribution, $\pi(\mathbf{x})$ be a known prior distribution, and p_{θ}^{SDE} be defined as in Section 3. Suppose $\{\mathbf{x}(t)\}_{t \in [0, T]}$ is a stochastic process defined by the SDE in Eq. (1) with $\mathbf{x}(0) \sim p$, where the marginal distribution of $\mathbf{x}(t)$ is denoted as p_t . Under some regularity conditions detailed in Appendix A, we have

$$D_{\text{KL}}(p \parallel p_{\theta}^{\text{SDE}}) \leq \mathcal{J}_{\text{SM}}(\boldsymbol{\theta}; g(\cdot)^2) + D_{\text{KL}}(p_T \parallel \pi). \quad (8)$$

Sketch of proof. Let $\boldsymbol{\mu}$ and $\boldsymbol{\nu}$ denote the path measures of SDEs in Eq. (1) and Eq. (5) respectively. Intuitively, $\boldsymbol{\mu}$ is the joint distribution of the diffusion process $\{\mathbf{x}(t)\}_{t \in [0, T]}$ given in Section 2.1, and $\boldsymbol{\nu}$ represents the joint distribution of the process $\{\hat{\mathbf{x}}_{\theta}(t)\}_{t \in [0, T]}$ defined in Section 3. Since we can marginalize $\boldsymbol{\mu}$ and $\boldsymbol{\nu}$ to obtain distributions p and p_{θ}^{SDE} , the data processing inequality gives $D_{\text{KL}}(p \parallel p_{\theta}^{\text{SDE}}) \leq D_{\text{KL}}(\boldsymbol{\mu} \parallel \boldsymbol{\nu})$. From the chain rule for the KL divergence, we also have $D_{\text{KL}}(\boldsymbol{\mu} \parallel \boldsymbol{\nu}) = D_{\text{KL}}(p_T \parallel \pi) + \mathbb{E}_{p_T(\mathbf{z})}[D_{\text{KL}}(\boldsymbol{\mu}(\cdot \mid \mathbf{x}(T) = \mathbf{z}) \parallel \boldsymbol{\nu}(\cdot \mid \hat{\mathbf{x}}_{\theta}(T) = \mathbf{z}))]$, where the KL divergence in the final term can be computed by applying the Girsanov theorem [34] to Eq. (5) and the reverse-time SDE of Eq. (1). \square

When the prior distribution π is fixed, Theorem 1 guarantees that optimizing the weighted combination of score matching losses $\mathcal{J}_{\text{SM}}(\boldsymbol{\theta}; g(\cdot)^2)$ is equivalent to minimizing an upper bound on the KL divergence from the data distribution p to the model distribution p_{θ}^{SDE} . Due to well-known equivalence between minimizing KL divergence and maximizing likelihood, we have the following corollary.

Corollary 1. Consider the same conditions and notations in Theorem 1. When π is a fixed prior distribution that does not depend on $\boldsymbol{\theta}$, we have

$$-\mathbb{E}_{p(\mathbf{x})}[\log p_{\theta}^{\text{SDE}}(\mathbf{x})] \leq \mathcal{J}_{\text{SM}}(\boldsymbol{\theta}; g(\cdot)^2) + C_1 = \mathcal{J}_{\text{DSM}}(\boldsymbol{\theta}; g(\cdot)^2) + C_2,$$

where C_1 and C_2 are constants independent of $\boldsymbol{\theta}$.

In light of the result in Corollary 1, we henceforth term $\lambda(t) = g(t)^2$ the *likelihood weighting*. The original weighting functions in [48] are inspired from earlier work such as [44, 45] and [19], which are motivated by balancing different score matching losses in the combination, and justified by empirical performance. In contrast, likelihood weighting is motivated from maximizing the likelihood of a probabilistic model induced by the diffusion process, and derived by theoretical analysis. There are three types of SDEs considered in [48]: the Variance Exploding (VE) SDE, the Variance Preserving (VP) SDE, and the subVP SDE. In Table 1, we summarize all these SDEs and contrast their original weighting functions with our likelihood weighting. For VE SDE, our likelihood weighting incidentally coincides with the original weighting used in [48], whereas for VP and subVP SDEs they differ from one another.

Theorem 1 leaves two questions unanswered. First, what are the conditions for the bound to be tight (become an equality)? Second, is there any connection between p_{θ}^{SDE} and p_{θ}^{ODE} under some conditions? We provide both answers in the following theorem.

Theorem 2. Suppose $p(\mathbf{x})$ and $q(\mathbf{x})$ have continuous second-order derivatives and finite second moments. Let $\{\mathbf{x}(t)\}_{t \in [0, T]}$ be the diffusion process defined by the SDE in Eq. (1). We use p_t and q_t to denote the distributions of $\mathbf{x}(t)$ when $\mathbf{x}(0) \sim p$ and $\mathbf{x}(0) \sim q$, and assume they satisfy the same assumptions in Appendix A. Under the conditions $q_T = \pi$ and $\mathbf{s}_{\theta}(\mathbf{x}, t) \equiv \nabla_{\mathbf{x}} \log q_t(\mathbf{x})$ for all $t \in [0, T]$, we have the following equivalence in distributions

$$p_{\theta}^{\text{SDE}} = p_{\theta}^{\text{ODE}} = q. \quad (9)$$

Moreover, we have

$$D_{\text{KL}}(p \parallel p_{\theta}^{\text{SDE}}) = \mathcal{J}_{\text{SM}}(\boldsymbol{\theta}; g(\cdot)^2) + D_{\text{KL}}(p_T \parallel \pi). \quad (10)$$

Sketch of proof. When $s_\theta(\mathbf{x}, t)$ matches $\nabla_{\mathbf{x}} \log q_t(\mathbf{x})$, they both represent the time-dependent score of the same stochastic process so we immediately have $p_\theta^{\text{SDE}} = q$. According to the theory of probability flow ODEs, we also have $p_\theta^{\text{ODE}} = q = p_\theta^{\text{SDE}}$. To prove Eq. (10), we note that $D_{\text{KL}}(p \parallel p_\theta^{\text{SDE}}) = D_{\text{KL}}(p \parallel q) = D_{\text{KL}}(p_T \parallel q_T) - \int_0^T \frac{d}{dt} D_{\text{KL}}(p_t \parallel q_t) dt = D_{\text{KL}}(p_T \parallel \pi) - \int_0^T \frac{d}{dt} D_{\text{KL}}(p_t \parallel q_t) dt$. We can now complete the proof by simplifying the integrand using the Fokker-Planck equation of p_t and q_t followed by integration by parts. \square

In practice, the conditions of Theorem 2 are hard to satisfy since our score-based model $s_\theta(\mathbf{x}, t)$ will not exactly match the score function $\nabla_{\mathbf{x}} \log q_t(\mathbf{x})$ of some reverse-time diffusion process with the initial distribution $q_T = \pi$. In other words, our score model may not be a valid time-dependent score function of a stochastic process with an appropriate initial distribution. Therefore, although score matching with likelihood weighting performs approximate maximum likelihood training for p_θ^{SDE} , we emphasize that it is not theoretically guaranteed to make the likelihood of p_θ^{ODE} better. That said, p_θ^{ODE} will closely match p_θ^{SDE} if our score-based model well-approximates the true score such that $s_\theta(\mathbf{x}, t) \approx \nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ for all \mathbf{x} and $t \in [0, T]$. Moreover, we empirically observe in our experiments (see Table 2) that training with the likelihood weighting is actually able to consistently improve the likelihood of p_θ^{ODE} across multiple datasets, SDEs, and model architectures.

4.2 Bounding the log-likelihood on individual datapoints

The bound in Theorem 1 is for the entire distributions of p and p_θ^{SDE} , but we often seek to bound the log-likelihood for an individual data point \mathbf{x} . In addition, $\mathcal{J}_{\text{SM}}(\theta; \lambda(\cdot))$ in the bound is not directly computable due to the unknown quantity $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$, and can only be evaluated up to an additive constant through $\mathcal{J}_{\text{DSM}}(\theta; \lambda(\cdot))$ (as we already discussed in Section 2.2). Therefore, the bound in Theorem 1 is only suitable for training purposes. To address these issues, we provide the following bounds for individual data points.

Theorem 3. *Let $p_{0t}(\mathbf{x}' \mid \mathbf{x})$ denote the transition distribution from $p_0(\mathbf{x})$ to $p_t(\mathbf{x})$ for the SDE in Eq. (1). With the same notations and conditions in Theorem 1, we have*

$$-\log p_\theta^{\text{SDE}}(\mathbf{x}) \leq \mathcal{L}_\theta^{\text{SM}}(\mathbf{x}) = \mathcal{L}_\theta^{\text{DSM}}(\mathbf{x}), \quad (11)$$

where $\mathcal{L}_\theta^{\text{SM}}(\mathbf{x})$ is defined as

$$-\mathbb{E}_{p_{0T}(\mathbf{x}' \mid \mathbf{x})}[\log \pi(\mathbf{x}')] + \frac{1}{2} \int_0^T \mathbb{E}_{p_{0t}(\mathbf{x}' \mid \mathbf{x})} \left[2g(t)^2 \nabla_{\mathbf{x}'} \cdot \mathbf{s}_\theta(\mathbf{x}', t) + g(t)^2 \|\mathbf{s}_\theta(\mathbf{x}', t)\|_2^2 - 2\nabla_{\mathbf{x}'} \cdot \mathbf{f}(\mathbf{x}', t) \right] dt,$$

and $\mathcal{L}_\theta^{\text{DSM}}(\mathbf{x})$ is given by

$$\begin{aligned} &-\mathbb{E}_{p_{0T}(\mathbf{x}' \mid \mathbf{x})}[\log \pi(\mathbf{x}')] + \frac{1}{2} \int_0^T \mathbb{E}_{p_{0t}(\mathbf{x}' \mid \mathbf{x})} \left[g(t)^2 \|\mathbf{s}_\theta(\mathbf{x}', t) - \nabla_{\mathbf{x}'} \log p_{0t}(\mathbf{x}' \mid \mathbf{x})\|_2^2 \right] dt \\ &\quad - \frac{1}{2} \int_0^T \mathbb{E}_{p_{0t}(\mathbf{x}' \mid \mathbf{x})} \left[g(t)^2 \|\nabla_{\mathbf{x}'} \log p_{0t}(\mathbf{x}' \mid \mathbf{x})\|_2^2 + 2\nabla_{\mathbf{x}'} \cdot \mathbf{f}(\mathbf{x}', t) \right] dt. \end{aligned}$$

Sketch of proof. For any continuous data distribution p , we have $-\mathbb{E}_{p(\mathbf{x})}[\log p_\theta^{\text{SDE}}(\mathbf{x})] = D_{\text{KL}}(p \parallel p_\theta^{\text{SDE}}) + \mathcal{H}(p)$, where $\mathcal{H}(p)$ denotes the differential entropy of p . The KL term can be bounded according to Theorem 1, while the differential entropy has an identity similar to Theorem 2 (see Theorem 4 in Appendix A). Combining the bound of $D_{\text{KL}}(p \parallel p_\theta^{\text{SDE}})$ and the identity of $\mathcal{H}(p)$, we obtain a bound on $-\mathbb{E}_{p(\mathbf{x})}[\log p_\theta^{\text{SDE}}(\mathbf{x})]$ that holds for all continuous distribution p . Removing the expectation over p on both sides then gives us a bound on $-\log p_\theta^{\text{SDE}}(\mathbf{x})$ for an individual datapoint \mathbf{x} . We can simplify this bound to $\mathcal{L}_\theta^{\text{SM}}(\mathbf{x})$ and $\mathcal{L}_\theta^{\text{DSM}}(\mathbf{x})$ with similar techniques to [23] and [56]. \square

We provide two equivalent bounds $\mathcal{L}_\theta^{\text{SM}}(\mathbf{x})$ and $\mathcal{L}_\theta^{\text{DSM}}(\mathbf{x})$. The former bears resemblance to score matching while the second resembles denoising score matching. Both admit efficient unbiased estimators when $\mathbf{f}(\cdot, t)$ is linear, as the time integrals and expectations in $\mathcal{L}_\theta^{\text{SM}}(\mathbf{x})$ and $\mathcal{L}_\theta^{\text{DSM}}(\mathbf{x})$ can be estimated by samples of the form (t, \mathbf{x}') , where t is uniformly sampled over $[0, T]$, and $\mathbf{x}' \sim p_{0t}(\mathbf{x}' \mid \mathbf{x})$. Since the transition distribution $p_{0t}(\mathbf{x}' \mid \mathbf{x})$ is a tractable Gaussian when $\mathbf{f}(\cdot, t)$ is linear, we can easily sample from it as well as evaluating $\nabla_{\mathbf{x}'} \log p_{0t}(\mathbf{x}' \mid \mathbf{x})$ for computing $\mathcal{L}_\theta^{\text{DSM}}(\mathbf{x})$.

Moreover, the divergences $\nabla_{\mathbf{x}} \cdot \mathbf{s}_{\theta}(\mathbf{x}, t)$ and $\nabla_{\mathbf{x}} \cdot \mathbf{f}(\mathbf{x}, t)$ in $\mathcal{L}_{\theta}^{\text{SM}}(\mathbf{x})$ and $\mathcal{L}_{\theta}^{\text{DSM}}(\mathbf{x})$ have efficient unbiased estimators via the Skilling–Hutchinson trick [42, 22].

We can view $\mathcal{L}_{\theta}^{\text{DSM}}(\mathbf{x})$ as a continuous-time generalization of the evidence lower bound (ELBO) in diffusion probabilistic models [43, 19]. Our bounds in Theorem 3 are not only useful for optimizing and estimating $\log p_{\theta}^{\text{SDE}}(\mathbf{x})$, but also for training the drift and diffusion coefficients $\mathbf{f}(\mathbf{x}, t)$ and $g(t)$ jointly with the score-based model $\mathbf{s}_{\theta}(\mathbf{x}, t)$; we leave this avenue of research for future work. In addition, we can plug the bounds in Theorem 3 into any objective that involves minimizing $-\log p_{\theta}^{\text{SDE}}(\mathbf{x})$ to obtain an efficient surrogate. Section 5.2 provides an example, where we perform variational dequantization to further improve the likelihood of score-based diffusion models.

Similar to the observation in Section 4.1, $\mathcal{L}_{\theta}^{\text{SM}}(\mathbf{x})$ and $\mathcal{L}_{\theta}^{\text{DSM}}(\mathbf{x})$ are not guaranteed to upper bound $-\log p_{\theta}^{\text{ODE}}(\mathbf{x})$. However, they should become approximate upper bounds when $\mathbf{s}_{\theta}(\mathbf{x}, t)$ is trained sufficiently close to the ground truth. In fact, we empirically observe that $-\log p_{\theta}^{\text{ODE}}(\mathbf{x}) \leq \mathcal{L}_{\theta}^{\text{SM}}(\mathbf{x}) = \mathcal{L}_{\theta}^{\text{DSM}}(\mathbf{x})$ holds true for \mathbf{x} sampled from the dataset in all experiments.

4.3 Numerical stability

So far we have assumed that the SDEs are defined in the time horizon $[0, T]$ in all theoretical analysis. In practice, however, we often face numerical instabilities when $t \rightarrow 0$. To avoid them, we choose a small non-zero starting time $\epsilon > 0$, and train/evaluate score-based diffusion models in the time horizon $[\epsilon, T]$ instead of $[0, T]$. Since ϵ is small, training score-based diffusion models with likelihood weighting still approximately maximizes their model likelihood. Yet at test time, the likelihood bound as computed in Theorem 3 is slightly biased, rendering the values not directly comparable to results reported in other works. We use Jensen’s inequality to correct for this bias in our experiments, for which we provide a detailed explanation in Appendix B.

4.4 Related work

Our result in Theorem 2 can be viewed as a generalization of De Bruijn’s identity ([49], Eq. 2.12) from its original differential form to an integral form. De Bruijn’s identity relates the rate of change of the Shannon entropy under an additive Gaussian noise channel to the Fisher information, a result which can be interpreted geometrically as relating the rate of change of the volume of a distribution’s typical set to its surface area. Ref. [2] (Lemma 1) builds on this result and presents an integral and relative form of de Bruijn’s identity which relates the KL divergence to the integral of the relative Fisher information for a distribution of interest and a reference standard normal. More generally, various identities and inequalities involving the (relative) Shannon entropy and (relative) Fisher information have found use in proofs of the central limit theorem [24]. Ref. [31] (Theorem 1) covers similar ground to the relative form of de Bruijn’s identity, but is perhaps the first to consider its implications for learning in probabilistic models by framing the discussion in terms of the score matching objective ([23], Eq. 2).

5 Improving the likelihood of score-based diffusion models

Our theoretical analysis implies that training with the likelihood weighting should improve the likelihood of score-based diffusion models. To verify this empirically, we test likelihood weighting with different model architectures, SDEs, and datasets. We observe that switching to likelihood weighting increases the variance of the training objective and propose to counteract it with importance sampling. We additionally combine our bound with variational dequantization [18] which narrows the gap between the likelihood of continuous and discrete probability models. All combined, we observe consistent improvement of likelihoods for both p_{θ}^{SDE} and p_{θ}^{ODE} across all settings. We term the model p_{θ}^{ODE} trained in this way *ScoreFlow*, and show that it achieves excellent likelihoods on CIFAR-10 [28] and ImageNet 32×32 [55], on a par with cutting-edge autoregressive models.

5.1 Variance reduction via importance sampling

As mentioned in Section 2.2, we typically use Monte Carlo sampling to approximate the time integral in $\mathcal{J}_{\text{DSM}}(\theta; \lambda(\cdot))$ during training. In particular, we first uniformly sample a time step $t \sim \mathcal{U}[0, T]$, and then use the denoising score matching loss at t as an estimate for the whole time integral. This

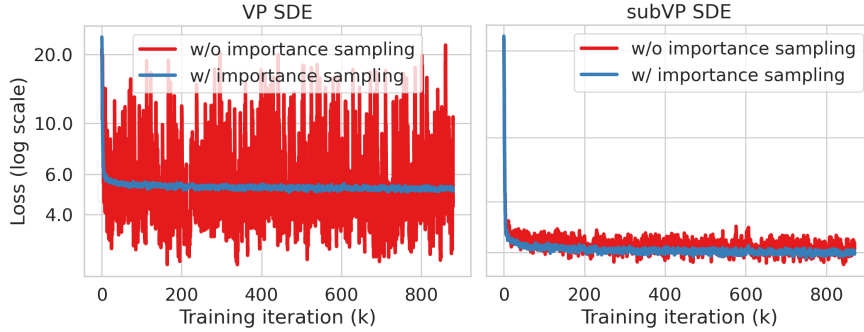


Figure 2: Learning curves with the likelihood weighting on the CIFAR-10 dataset (smoothed with exponential moving average). Importance sampling significantly reduces the loss variance.

Monte Carlo approximation is much faster than computing the time integral accurately, but introduces additional variance to the training loss.

We empirically observe that this Monte Carlo approximation suffers from a larger variance when using our likelihood weighting instead of the original weightings in [48]. Leveraging importance sampling, we propose a new Monte Carlo approximation that significantly reduces the variance of learning curves under likelihood weighting, as demonstrated in Fig. 2. In fact, with importance sampling, the loss variance (after convergence) decreases from 98.48 to 0.068 on CIFAR-10, and decreases from 0.51 to 0.043 on ImageNet.

Let $\lambda(t) = \alpha(t)^2$ denote the weightings in [48] (reproduced in Table 1), and recall that our likelihood weighting is $\lambda(t) = g(t)^2$. Since $\alpha(t)^2$ empirically leads to lower variance, we can use a proposal distribution $p(t) := g(t)^2 / \alpha(t)^2 Z$ to change the weighting in $\mathcal{J}_{\text{DSM}}(\theta; g(\cdot)^2)$ from $g(t)^2$ to $\alpha(t)^2$ with importance sampling, where Z is a normalizing constant that ensures $\int p(t) dt = 1$. Specifically, for any function $h(t)$, we estimate the time integral $\int_0^T g(t)^2 h(t) dt$ with

$$\int_0^T g(t)^2 h(t) dt = Z \int_0^T p(t) \alpha(t)^2 h(t) dt \approx TZ \alpha(\tilde{t})^2 h(\tilde{t}), \quad (12)$$

where \tilde{t} is a sample from $p(t)$. When training score-based models with likelihood weighting, $h(t)$ corresponds to the denoising score matching loss at time t .

Ref. [33] also observes that optimizing the ELBO for diffusion probabilistic models has large variance, and proposes to reduce it with importance sampling. They build their proposal distribution based on historical loss values stored at thousands of discrete time steps. Despite this similarity, our method is easier to implement without needing to maintain history, can be used for evaluation, and is particularly suited to the continuous-time setting.

5.2 Variational dequantization

Digital images are discrete data, and must be dequantized when training continuous density models like normalizing flows [12, 13] and score-based diffusion models. One popular approach to this is uniform dequantization [53, 50], where we add small uniform noise over $[0, 1)$ to images taking values in $\{0, 1, \dots, 255\}$. As shown in [50], training a continuous model $p_\theta(\mathbf{x})$ on uniformly dequantized data implicitly maximizes a lower bound on the log-likelihood of a certain discrete model $P_\theta(\mathbf{x})$. Due to the gap between $p_\theta(\mathbf{x})$ and $P_\theta(\mathbf{x})$, comparing the likelihood of continuous density models to models which fit discrete data directly, such as autoregressive models [55] or variational autoencoders, naturally puts the former at a disadvantage.

To minimize the gap between $p_\theta(\mathbf{x})$ and $P_\theta(\mathbf{x})$, ref. [18] proposes variational dequantization, where a separate normalizing flow model $q_\phi(\mathbf{u} | \mathbf{x})$ is trained to produce the dequantization noise by optimizing the following objective

$$\max_{\phi} \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} \mathbb{E}_{\mathbf{u} \sim q_\phi(\cdot | \mathbf{x})} [\log p_\theta(\mathbf{x} + \mathbf{u}) - \log q_\phi(\mathbf{u} | \mathbf{x})]. \quad (13)$$

Plugging in the lower bound on $\log p_\theta(\mathbf{x})$ from Theorem 3, we can optimize Eq. (13) to improve the likelihood of score-based diffusion models.

Table 2: Negative log-likelihood (bits/dim) and sample quality (FID scores) on CIFAR-10 and ImageNet 32×32 . Abbreviations: “NLL” for “negative log-likelihood”; “Uni. deq.” for “Uniform dequantization”; “Var. deq.” for “Variational dequantization”; “LW” for “likelihood weighting”; and “IS” for “importance sampling”. Bold indicates best result in the corresponding column. Shaded rows represent models trained with both likelihood weighting and importance sampling.

Model	SDE	CIFAR-10					ImageNet 32×32				
		Uni. deq.		Var. deq.		FID↓	Uni. deq.		Var. deq.		FID↓
		NLL↓	Bound↓	NLL↓	Bound↓		NLL↓	Bound↓	NLL↓	Bound↓	
Baseline	VP	3.16	3.28	3.04	3.14	3.98	3.90	3.96	3.84	3.91	8.34
Baseline + LW	VP	3.06	3.18	2.94	3.03	5.18	3.91	3.96	3.86	3.92	17.75
Baseline + LW + IS	VP	2.95	3.08	2.83	2.94	6.03	3.86	3.92	3.80	3.88	11.15
Deep	VP	3.13	3.25	3.01	3.10	3.09	3.89	3.95	3.84	3.90	8.40
Deep + LW	VP	3.06	3.17	2.93	3.02	7.88	3.91	3.96	3.86	3.92	17.73
Deep + LW + IS	VP	2.93	3.06	2.80	2.92	5.34	3.85	3.92	3.79	3.88	11.20
Baseline	subVP	2.99	3.09	2.88	2.98	3.20	3.87	3.92	3.82	3.88	8.71
Baseline + LW	subVP	2.97	3.07	2.86	2.96	7.33	3.87	3.92	3.82	3.88	12.99
Baseline + LW + IS	subVP	2.94	3.05	2.84	2.94	5.58	3.84	3.91	3.79	3.87	10.57
Deep	subVP	2.96	3.06	2.85	2.95	2.86	3.86	3.91	3.81	3.87	8.87
Deep + LW	subVP	2.95	3.05	2.85	2.94	6.57	3.88	3.93	3.83	3.88	16.55
Deep + LW + IS	subVP	2.90	3.02	2.81	2.90	5.40	3.82	3.90	3.76	3.86	10.18

5.3 Experiments

We empirically test the performance of likelihood weighting, importance sampling and variational dequantization across multiple architectures of score-based models, SDEs, and datasets. In particular, we consider DDPM++ (“Baseline” in Table 2) and DDPM++ (deep) (“Deep” in Table 2) models with VP and subVP SDEs [48] on CIFAR-10 [28] and ImageNet 32×32 [55] datasets. We omit experiments on the VE SDE since (i) under this SDE our likelihood weighting is the same as the original weighting in [48]; (ii) we empirically observe that the best VE SDE model achieves around 3.4 bits/dim on CIFAR-10 in our experiments, which is significantly worse than other SDEs. For each experiment, we report $-\mathbb{E}[\log p_{\theta}^{\text{ODE}}(\mathbf{x})]$ (“Negative log-likelihood” in Table 2), and the upper bound $\mathbb{E}[\mathcal{L}_{\theta}^{\text{DSM}}(\mathbf{x})]$ on $-\mathbb{E}[\log p_{\theta}^{\text{SDE}}(\mathbf{x})]$ (“Bound” in Table 2). In addition, we report FID scores [17] for samples from p_{θ}^{ODE} , produced by solving the corresponding ODE with the Dormand–Prince RK45 [14] solver. Unless otherwise noted, we apply horizontal flipping as data augmentation for training models on CIFAR-10, so as to match the settings in [48, 19]. Detailed description of all our experiments can be found in Appendices B and C.

We summarize all results in Table 2. Our key observations are as follows:

1. Although Theorem 3 only guarantees $\mathbb{E}[\mathcal{L}_{\theta}^{\text{DSM}}(\mathbf{x})] \geq -\mathbb{E}[\log p_{\theta}^{\text{SDE}}(\mathbf{x})]$, and in general we have $p_{\theta}^{\text{SDE}} \neq p_{\theta}^{\text{ODE}}$, we still find that $\mathbb{E}[\mathcal{L}_{\theta}^{\text{DSM}}(\mathbf{x})]$ (“Bound” in Table 2) $\geq -\mathbb{E}[\log p_{\theta}^{\text{ODE}}(\mathbf{x})]$ (“NLL” in Table 2) in all our settings.
2. When all conditions are fixed except for the weighting in the training objective, having a lower value of the bound for p_{θ}^{SDE} always leads to a lower negative log-likelihood for p_{θ}^{ODE} .
3. With only likelihood weighting, we can uniformly improve the likelihood of p_{θ}^{ODE} and the bound of p_{θ}^{SDE} on CIFAR-10 across model architectures and SDEs, but it is not sufficient to guarantee likelihood improvement on ImageNet 32×32 .
4. By combining importance sampling and likelihood weighting, we are able to achieve uniformly better likelihood for p_{θ}^{ODE} and bounds for p_{θ}^{SDE} across all model architectures, SDEs, and datasets, with only slight degradation of sample quality as measured by FID [17].
5. Variational dequantization uniformly improves both the bound for p_{θ}^{SDE} and the negative log-likelihood (NLL) of p_{θ}^{ODE} in all settings, regardless of likelihood weighting.

Our experiments confirm that with importance sampling, likelihood weighting is not only effective for maximizing the lower bound for the log-likelihood of p_{θ}^{SDE} , but also improving the log-likelihood of p_{θ}^{ODE} . In agreement with [19, 33], we observe that models achieving better likelihood tend to have worse FIDs. However, we emphasize that this degradation of FID is small, and samples actually

have no obvious difference in visual quality (see Figs. 3 and 4). To trade likelihood for FID, we can use weighting functions that interpolate between likelihood weighting and the original weighting functions in [48]. Our FID scores are still much better than most other likelihood-based models.

We term p_{θ}^{ODE} a *ScoreFlow* when its corresponding score-based model $s_{\theta}(x, t)$ is trained with likelihood weighting, importance sampling, and variational dequantization combined. It can be viewed as a continuous normalizing flow, but is parameterized by a score-based model and trained in a more efficient way. With variational dequantization, we show ScoreFlows obtain competitive negative log-likelihoods (NLLs) of 2.83 bits/dim on CIFAR-10 and 3.76 bits/dim on ImageNet 32×32 . Here the ScoreFlow on CIFAR-10 is trained without horizontal flipping (different from the setting in Table 2). As shown in Table 3, our results are on a par with the state-of-the-art autoregressive models on these tasks, and outperform all existing normalizing flow models. The likelihood for CIFAR-10 can be significantly improved by incorporating advanced data augmentation, as demonstrated in [25, 41]. While we do not compare against these approaches, we believe that incorporating the same data augmentation techniques could also improve the likelihood of ScoreFlows.

Table 3: NLLs on CIFAR-10 and ImageNet 32×32 .

Model	CIFAR-10	ImageNet
FFJORD [15]	3.40	-
Flow++ [18]	3.08	3.86
Gated PixelCNN [35]	3.03	3.83
VFlow [4]	2.98	3.83
PixelCNN++ [40]	2.92	-
NVAE [54]	2.91	3.92
Image Transformer [36]	2.90	3.77
Very Deep VAE [8]	2.87	3.80
PixelSNAIL [7]	2.85	3.80
δ -VAE [38]	2.83	3.77
Sparse Transformer [9]	2.80	-
ScoreFlow (Ours)	2.83	3.76

6 Conclusion

We propose an efficient training objective for approximate maximum likelihood training of score-based diffusion models. Our theoretical analysis shows that the weighted combination of score matching losses upper bounds the negative log-likelihood when using a particular weighting function which we term the likelihood weighting. By minimizing this upper bound, we consistently improve the likelihood of score-based diffusion models across multiple model architectures, SDEs, and datasets. When combined with variational dequantization, we achieve competitive likelihoods on CIFAR-10 and ImageNet 32×32 , matching the performance of best-in-class autoregressive models.

Our upper bound is analogous to the evidence lower bound commonly used for training variational autoencoders. Aside from promoting higher likelihood, the bound can be combined with other objectives that depend on the negative log-likelihood, and also enables joint training of the forward and backward SDEs, which we leave as a future research direction. Our results suggest that score-based diffusion models are competitive alternatives to continuous normalizing flows which enjoy the same tractable likelihood computation but with more efficient maximum likelihood training.

Limitations and broader impact Despite promising experimental results, we would like to emphasize that there is no theoretical guarantee that improving the SDE likelihood will improve the ODE likelihood, and this is explicitly a limitation of our work. Score-based diffusion models also suffer from slow sampling. In our experiments, the ODE solver typically need around 550 and 450 evaluations of the score-based model for generation and likelihood computation on CIFAR-10 and ImageNet respectively, which is considerably slower than alternative generative models like VAEs and GANs. In addition, the current formulation of score-based diffusion models only supports continuous data, and cannot be naturally adapted to discrete data without resorting to dequantization. Similarly to other deep generative models, score-based diffusion models can potentially be used to generate harmful media contents such as ‘deepfakes’, and might reflect and amplify undesirable social bias that could exist in the training dataset.

Author Contributions

Yang Song wrote the code, ran the experiments, proposed and proved Theorems 1 and 3, and wrote most of the paper. Conor Durkan proposed and proved a first version of Theorem 2, and wrote the paper. Iain Murray and Stefano Ermon co-advised the project and provided helpful edits to the draft.

Acknowledgments and Disclosure of Funding

The authors would like to thank Sam Power, George Papamakarios, Adji Dieng for helpful feedback, and Duoduo for providing her photos in Fig. 1. This research was supported by NSF (#1651565, #1522054, #1733686), ONR (N000141912145), AFOSR (FA95501910024), ARO (W911NF-21-1-0125), Sloan Fellowship, and Google TPU Research Cloud. This research was also supported by the EPSRC Centre for Doctoral Training in Data Science, funded by the UK Engineering and Physical Sciences Research Council (grant EP/L016427/1), and the University of Edinburgh. Yang Song was supported by the Apple PhD Fellowship in AI/ML.

References

- [1] B. D. Anderson. Reverse-Time Diffusion Equation Models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982.
- [2] A. R. Barron. Entropy and the Central Limit Theorem. *Annals of Probability*, 14(1):336–342, 1986.
- [3] R. Cai, G. Yang, H. Averbuch-Elor, Z. Hao, S. Belongie, N. Snavely, and B. Hariharan. Learning Gradient Fields for Shape Generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020.
- [4] J. Chen, C. Lu, B. Chenli, J. Zhu, and T. Tian. Vflow: More expressive generative flows with variational data augmentation. In *International Conference on Machine Learning*, pages 1660–1669. PMLR, 2020.
- [5] N. Chen, Y. Zhang, H. Zen, R. J. Weiss, M. Norouzi, and W. Chan. WaveGrad: Estimating Gradients for Waveform Generation. *arXiv preprint arXiv:2009.00713*, 2020.
- [6] R. T. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud. Neural Ordinary Differential Equations. In *Advances in neural information processing systems*, pages 6571–6583, 2018.
- [7] X. Chen, N. Mishra, M. Rohaninejad, and P. Abbeel. Pixelsnail: An improved autoregressive generative model. In *International Conference on Machine Learning*, pages 864–872. PMLR, 2018.
- [8] R. Child. Very deep VAEs generalize autoregressive models and can outperform them on images. In *International Conference on Learning Representations*, 2021.
- [9] R. Child, S. Gray, A. Radford, and I. Sutskever. Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509*, 2019.
- [10] Z. Dai, Z. Yang, F. Yang, W. W. Cohen, and R. Salakhutdinov. Good semi-supervised learning that requires a bad GAN. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 6513–6523, 2017.
- [11] P. Dhariwal and A. Nichol. Diffusion models beat GANs on image synthesis. *arXiv preprint arXiv:2105.05233*, 2021.
- [12] L. Dinh, D. Krueger, and Y. Bengio. NICE: Non-Linear Independent Components Estimation. *arXiv preprint arXiv:1410.8516*, 2014.
- [13] L. Dinh, J. Sohl-Dickstein, and S. Bengio. Density estimation using real NVP. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017.
- [14] J. R. Dormand and P. J. Prince. A family of embedded runge-kutta formulae. *Journal of computational and applied mathematics*, 6(1):19–26, 1980.
- [15] W. Grathwohl, R. T. Q. Chen, J. Bettencourt, I. Sutskever, and D. Duvenaud. FFFJORD: Free-form Continuous Dynamics for Scalable Reversible Generative Models. In *International Conference on Learning Representations*, 2019.
- [16] U. G. Haussmann and E. Pardoux. Time reversal of diffusions. *The Annals of Probability*, pages 1188–1205, 1986.
- [17] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. In I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, and R. Garnett, editors, *Advances*

- in *Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 6626–6637, 2017.
- [18] J. Ho, X. Chen, A. Srinivas, Y. Duan, and P. Abbeel. Flow++: Improving flow-based generative models with variational dequantization and architecture design. In *International Conference on Machine Learning*, pages 2722–2730. PMLR, 2019.
- [19] J. Ho, A. Jain, and P. Abbeel. Denoising Diffusion Probabilistic Models. *Advances in Neural Information Processing Systems*, 33, 2020.
- [20] J. Ho, E. Lohn, and P. Abbeel. Compression with flows via local bits-back coding. In H. M. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. B. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 3874–3883, 2019.
- [21] E. Hoogeboom, J. W. T. Peters, R. van den Berg, and M. Welling. Integer discrete flows and lossless compression. In H. M. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. B. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 12134–12144, 2019.
- [22] M. F. Hutchinson. 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.
- [23] A. Hyvärinen. Estimation of Non-Normalized Statistical Models by Score Matching. *Journal of Machine Learning Research*, 6(Apr):695–709, 2005.
- [24] O. Johnson and A. Barron. Fisher Information inequalities and the Central Limit Theorem. *Probability Theory and Related Fields*, 129(3):391–409, 2004.
- [25] H. Jun, R. Child, M. Chen, J. Schulman, A. Ramesh, A. Radford, and I. Sutskever. Distribution augmentation for generative modeling. In *International Conference on Machine Learning*, pages 5006–5019. PMLR, 2020.
- [26] D. P. Kingma and M. Welling. Auto-Encoding Variational Bayes. In *International Conference on Learning Representations*, 2014.
- [27] Z. Kong, W. Ping, J. Huang, K. Zhao, and B. Catanzaro. DiffWave: A Versatile Diffusion Model for Audio Synthesis. *arXiv preprint arXiv:2009.09761*, 2020.
- [28] A. Krizhevsky, V. Nair, and G. Hinton. The CIFAR-10 Dataset. *online: <http://www.cs.toronto.edu/kriz/cifar.html>*, 55, 2014.
- [29] C. Léonard. Some properties of path measures. In *Séminaire de Probabilités XLVI*, pages 207–230. Springer, 2014.
- [30] X. Li, T.-K. L. Wong, R. T. Q. Chen, and D. Duvenaud. Scalable Gradients for Stochastic Differential Equations. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, 2020.
- [31] S. Lyu. Interpretation and Generalization of Score Matching. In *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence*, pages 359–366, 2009.
- [32] D. Maoutsa, S. Reich, and M. Opper. Interacting particle solutions of Fokker–Planck equations through gradient-log-density estimation. *arXiv preprint arXiv:2006.00702*, 2020.
- [33] A. Nichol and P. Dhariwal. Improved denoising diffusion probabilistic models. *arXiv preprint arXiv:2102.09672*, 2021.
- [34] B. Oksendal. *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media, 2013.
- [35] A. v. d. Oord, N. Kalchbrenner, O. Vinyals, L. Espeholt, A. Graves, and K. Kavukcuoglu. Conditional image generation with pixelcnn decoders. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pages 4797–4805, 2016.
- [36] N. Parmar, A. Vaswani, J. Uszkoreit, L. Kaiser, N. Shazeer, A. Ku, and D. Tran. Image transformer. In *International Conference on Machine Learning*, pages 4055–4064. PMLR, 2018.

- [37] V. Popov, I. Vovk, V. Gogoryan, T. Sadekova, and M. Kudinov. Grad-TTS: A diffusion probabilistic model for text-to-speech. *arXiv preprint arXiv:2105.06337*, 2021.
- [38] A. Razavi, A. van den Oord, B. Poole, and O. Vinyals. Preventing posterior collapse with delta-vaes. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.
- [39] D. J. Rezende, S. Mohamed, and D. Wierstra. Stochastic Backpropagation and Approximate Inference in Deep Generative Models. In *Proceedings of the 31st International Conference on Machine Learning*, volume 32, 2014.
- [40] T. Salimans, A. Karpathy, X. Chen, and D. P. Kingma. Pixelcnn++: Improving the pixelcnn with discretized logistic mixture likelihood and other modifications. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017.
- [41] S. Sinha and A. B. Dieng. Consistency regularization for variational auto-encoders. *arXiv preprint arXiv:2105.14859*, 2021.
- [42] J. Skilling. The eigenvalues of mega-dimensional matrices. In *Maximum Entropy and Bayesian Methods*, pages 455–466. Springer, 1989.
- [43] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli. Deep Unsupervised Learning Using Nonequilibrium Thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265, 2015.
- [44] Y. Song and S. Ermon. Generative Modeling by Estimating Gradients of the Data Distribution. In *Advances in Neural Information Processing Systems*, pages 11918–11930, 2019.
- [45] Y. Song and S. Ermon. Improved Techniques for Training Score-Based Generative Models. *Advances in Neural Information Processing Systems*, 33, 2020.
- [46] Y. Song, S. Garg, J. Shi, and S. Ermon. Sliced Score Matching: A Scalable Approach to Density and Score Estimation. In *Proceedings of the Thirty-Fifth Conference on Uncertainty in Artificial Intelligence*, page 204, 2019.
- [47] Y. Song, T. Kim, S. Nowozin, S. Ermon, and N. Kushman. Pixeldefend: Leveraging generative models to understand and defend against adversarial examples. In *International Conference on Learning Representations*, 2018.
- [48] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole. Score-Based Generative Modeling Through Stochastic Differential Equations. In *International Conference on Learning Representations*, 2021.
- [49] A. Stam. Some Inequalities Satisfied by the Quantities of Information of Fisher and Shannon. *Information and Control*, 2(2):101–112, June 1959.
- [50] L. Theis, A. van den Oord, and M. Bethge. A note on the evaluation of generative models. In Y. Bengio and Y. LeCun, editors, *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016.
- [51] J. Townsend, T. Bird, and D. Barber. Practical lossless compression with latent variables using bits back coding. In *International Conference on Learning Representations*, 2019.
- [52] B. Tzen and M. Raginsky. Neural Stochastic Differential Equations: Deep Latent Gaussian Models in the Diffusion Limit. *arXiv:1905.09883*, 2019.
- [53] B. Uria, I. Murray, and H. Larochelle. Rnade: the real-valued neural autoregressive density-estimator. In *Proceedings of the 26th International Conference on Neural Information Processing Systems-Volume 2*, pages 2175–2183, 2013.
- [54] A. Vahdat and J. Kautz. NVAE: A deep hierarchical variational autoencoder. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [55] A. Van den Oord, N. Kalchbrenner, and K. Kavukcuoglu. Pixel recurrent neural networks. In *International Conference on Machine Learning*, pages 1747–1756. PMLR, 2016.
- [56] P. Vincent. A Connection Between Score Matching and Denoising Autoencoders. *Neural Computation*, 23(7):1661–1674, 2011.

[57] K. Yang, J. Yau, L. Fei-Fei, J. Deng, and O. Russakovsky. A study of face obfuscation in imagenet. *arXiv preprint arXiv:2103.06191*, 2021.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes] The limitations of our theory are discussed inline in Section 4.
 - (c) Did you discuss any potential negative societal impacts of your work? [Yes] Discussed in conclusion.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes] The full set of assumptions are in Appendix A.
 - (b) Did you include complete proofs of all theoretical results? [Yes] The full proofs are all provided in Appendix A.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] Code is released at https://github.com/yang-song/score_flow.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] All experimental details are in Appendix C.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No] Training our models is expensive and we do not have enough resource/time for multiple repetitions of our experiments.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] Compute and resource types are given in Appendix C.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [Yes] We cited the creators of CIFAR-10 and down-sampled ImageNet.
 - (b) Did you mention the license of the assets? [No] Licenses are standard and can be found online.
 - (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] Code and checkpoint are released at https://github.com/yang-song/score_flow.
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [No] All datasets used in our work are publicly available.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [Yes] We mentioned privacy issues of the ImageNet dataset inline in Appendix C.
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Proofs

We first summarize the notations and assumptions used in our theorems.

Notations The drift and diffusion coefficients of the SDE in Eq. (1) are denoted as $\mathbf{f} : \mathbb{R}^D \times [0, T] \rightarrow \mathbb{R}^D$ and $g : [0, T] \rightarrow \mathbb{R}$ respectively, where $[0, T]$ represents a fixed time horizon, and \times denotes the Cartesian product. The solution to Eq. (1) is a stochastic process $\{\mathbf{x}(t)\}_{t \in [0, T]}$. We use p_t to represent the marginal distribution of $\mathbf{x}(t)$, and $p_{0t}(\mathbf{x}' | \mathbf{x})$ to denote the transition distribution from $\mathbf{x}(0)$ to $\mathbf{x}(t)$. The data distribution and prior distribution are given by p and π . We use \mathcal{C} to denote all continuous functions, and let \mathcal{C}^k denote the family of functions with continuous k -th order derivatives. For any vector-valued function $\mathbf{h} : \mathbb{R}^D \times [0, T] \rightarrow \mathbb{R}^D$, we use $\nabla \cdot \mathbf{h}(\mathbf{x}, t)$ to represent its divergence with respect to the first input variable.

Assumptions We make the following assumptions throughout the paper:

- (i) $p(\mathbf{x}) \in \mathcal{C}^2$ and $\mathbb{E}_{\mathbf{x} \sim p}[\|\mathbf{x}\|_2^2] < \infty$.
- (ii) $\pi(\mathbf{x}) \in \mathcal{C}^2$ and $\mathbb{E}_{\mathbf{x} \sim \pi}[\|\mathbf{x}\|_2^2] < \infty$.
- (iii) $\forall t \in [0, T] : \mathbf{f}(\cdot, t) \in \mathcal{C}^1, \exists C > 0 \forall \mathbf{x} \in \mathbb{R}^D, t \in [0, T] : \|\mathbf{f}(\mathbf{x}, t)\|_2 \leq C(1 + \|\mathbf{x}\|_2)$.
- (iv) $\exists C > 0, \forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^D : \|\mathbf{f}(\mathbf{x}, t) - \mathbf{f}(\mathbf{y}, t)\|_2 \leq C \|\mathbf{x} - \mathbf{y}\|_2$.
- (v) $g \in \mathcal{C}$ and $\forall t \in [0, T], |g(t)| > 0$.
- (vi) For any open bounded set $\mathcal{O}, \int_0^T \int_{\mathcal{O}} \|p_t(\mathbf{x})\|_2^2 + Dg(t)^2 \|\nabla_{\mathbf{x}} p_t(\mathbf{x})\|_2^2 d\mathbf{x} dt < \infty$.
- (vii) $\exists C > 0 \forall \mathbf{x} \in \mathbb{R}^D, t \in [0, T] : \|\nabla_{\mathbf{x}} \log p_t(\mathbf{x})\|_2 \leq C(1 + \|\mathbf{x}\|_2)$.
- (viii) $\exists C > 0, \forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^D : \|\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) - \nabla_{\mathbf{y}} \log p_t(\mathbf{y})\|_2 \leq C \|\mathbf{x} - \mathbf{y}\|_2$.
- (ix) $\exists C > 0 \forall \mathbf{x} \in \mathbb{R}^D, t \in [0, T] : \|\mathbf{s}_{\theta}(\mathbf{x}, t)\|_2 \leq C(1 + \|\mathbf{x}\|_2)$.
- (x) $\exists C > 0, \forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^D : \|\mathbf{s}_{\theta}(\mathbf{x}, t) - \mathbf{s}_{\theta}(\mathbf{y}, t)\|_2 \leq C \|\mathbf{x} - \mathbf{y}\|_2$.
- (xi) Novikov's condition: $\mathbb{E} \left[\exp \left(\frac{1}{2} \int_0^T \|\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) - \mathbf{s}_{\theta}(\mathbf{x}, t)\|_2^2 dt \right) \right] < \infty$.
- (xii) $\forall t \in [0, T] \exists k > 0 : p_t(\mathbf{x}) = O(e^{-\|\mathbf{x}\|_2^k})$ as $\|\mathbf{x}\|_2 \rightarrow \infty$.

Below we provide all proofs for our theorems.

Theorem 1. Let $p(\mathbf{x})$ be the data distribution, $\pi(\mathbf{x})$ be a known prior distribution, and p_{θ}^{SDE} be defined as in Section 3. Suppose $\{\mathbf{x}(t)\}_{t \in [0, T]}$ is a stochastic process defined by the SDE in Eq. (1) with $\mathbf{x}(0) \sim p$, where the marginal distribution of $\mathbf{x}(t)$ is denoted as p_t . Under some regularity conditions detailed in Appendix A, we have

$$D_{\text{KL}}(p \parallel p_{\theta}^{\text{SDE}}) \leq \mathcal{J}_{\text{SM}}(\theta; g(\cdot)^2) + D_{\text{KL}}(p_T \parallel \pi). \quad (8)$$

Proof. We denote the path measure of $\{\mathbf{x}(t)\}_{t \in [0, T]}$ and $\{\hat{\mathbf{x}}_{\theta}(t)\}_{t \in [0, T]}$ as μ and ν respectively. Due to assumptions (i) (ii) (iii) (iv) (v) (ix) and (x), both μ and ν are uniquely given by the corresponding SDEs. Consider a Markov kernel $K(\{\mathbf{z}(t)\}_{t \in [0, t]}, \mathbf{y}) := \delta(\mathbf{z}(0) = \mathbf{y})$. Since $\mathbf{x}(0) \sim p_0$ and $\hat{\mathbf{x}}_{\theta}(0) \sim p_{\theta}$, we have the following result

$$\begin{aligned} \int K(\{\mathbf{x}(t)\}_{t \in [0, T]}, \mathbf{x}) d\mu(\{\mathbf{x}(t)\}_{t \in [0, T]}) &= p_0(\mathbf{x}) \\ \int K(\{\hat{\mathbf{x}}_{\theta}(t)\}_{t \in [0, T]}, \mathbf{x}) d\nu(\{\hat{\mathbf{x}}_{\theta}(t)\}_{t \in [0, T]}) &= p_{\theta}(\mathbf{x}). \end{aligned}$$

Here the Markov kernel K essentially performs marginalization of path measures to obtain ‘‘sliced’’ distributions at $t = 0$. We can use the data processing inequality with this Markov kernel to obtain

$$\begin{aligned} D_{\text{KL}}(p \parallel p_{\theta}) &= D_{\text{KL}}(p_0 \parallel p_{\theta}) \\ &= D_{\text{KL}} \left(\int K(\{\mathbf{x}(t)\}_{t \in [0, T]}, \mathbf{x}) d\mu(\{\mathbf{x}(t)\}_{t \in [0, T]}) \parallel \int K(\{\hat{\mathbf{x}}_{\theta}(t)\}_{t \in [0, T]}, \mathbf{x}) d\nu(\{\hat{\mathbf{x}}_{\theta}(t)\}_{t \in [0, T]}) \right) \\ &\leq D_{\text{KL}}(\mu \parallel \nu). \end{aligned} \quad (14)$$

Recall that by definition $\mathbf{x}(T) \sim p_T$ and $\hat{\mathbf{x}}_\theta(T) \sim \pi$. Leveraging the chain rule of KL divergences (see, for example, Theorem 2.4 in [29]), we have

$$D_{\text{KL}}(\boldsymbol{\mu} \parallel \boldsymbol{\nu}) = D_{\text{KL}}(p_T \parallel \pi) + \mathbb{E}_{\mathbf{z} \sim p_T} [D_{\text{KL}}(\boldsymbol{\mu}(\cdot \mid \mathbf{x}(T) = \mathbf{z}) \parallel \boldsymbol{\nu}(\cdot \mid \hat{\mathbf{x}}_\theta(T) = \mathbf{z}))]. \quad (15)$$

Under assumptions (i) (iii) (iv) (v) (vi) (vii) (viii), the SDE in Eq. (1) has a corresponding reverse-time SDE given by

$$d\mathbf{x} = [\mathbf{f}(\mathbf{x}, t) - g(t)^2 \nabla_{\mathbf{x}} \log p_t(\mathbf{x})] dt + g(t) d\bar{\mathbf{w}}. \quad (16)$$

Since Eq. (16) is the time reversal of Eq. (1), it induces the same path measure $\boldsymbol{\mu}$. As a result, $D_{\text{KL}}(\boldsymbol{\mu}(\cdot \mid \mathbf{x}(T) = \mathbf{z}) \parallel \boldsymbol{\nu}(\cdot \mid \hat{\mathbf{x}}_\theta(T) = \mathbf{z}))$ can be viewed as the KL divergence between the path measures induced by the following two (reverse-time) SDEs:

$$\begin{aligned} d\mathbf{x} &= [\mathbf{f}(\mathbf{x}, t) - g(t)^2 \nabla_{\mathbf{x}} \log p_t(\mathbf{x})] dt + g(t) d\bar{\mathbf{w}}, & \mathbf{x}(T) &= \mathbf{x} \\ d\hat{\mathbf{x}} &= [\mathbf{f}(\hat{\mathbf{x}}, t) - g(t)^2 \mathbf{s}_\theta(\hat{\mathbf{x}}, t)] dt + g(t) d\bar{\mathbf{w}}, & \hat{\mathbf{x}}_\theta(T) &= \mathbf{x}. \end{aligned}$$

The KL divergence between two SDEs with shared diffusion coefficients and starting points exists under assumptions (vii) (viii) (ix) (x) (xi) (see, e.g., [52, 30]), and can be computed via the Girsanov theorem [34]:

$$\begin{aligned} & D_{\text{KL}}(\boldsymbol{\mu}(\cdot \mid \mathbf{x}(T) = \mathbf{z}) \parallel \boldsymbol{\nu}(\cdot \mid \hat{\mathbf{x}}_\theta(T) = \mathbf{z})) \\ &= -\mathbb{E}_{\boldsymbol{\mu}} \left[\log \frac{d\boldsymbol{\nu}}{d\boldsymbol{\mu}} \right] \quad (17) \\ &\stackrel{(i)}{=} \mathbb{E}_{\boldsymbol{\mu}} \left[\int_0^T g(t) (\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) - \mathbf{s}_\theta(\mathbf{x}, t)) d\bar{\mathbf{w}}_t + \frac{1}{2} \int_0^T g(t)^2 \|\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) - \mathbf{s}_\theta(\mathbf{x}, t)\|_2^2 dt \right] \\ &\stackrel{(ii)}{=} \mathbb{E}_{\boldsymbol{\mu}} \left[\frac{1}{2} \int_0^T g(t)^2 \|\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) - \mathbf{s}_\theta(\mathbf{x}, t)\|_2^2 dt \right] \\ &= \frac{1}{2} \int_0^T \mathbb{E}_{p_t(\mathbf{x})} [g(t)^2 \|\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) - \mathbf{s}_\theta(\mathbf{x}, t)\|_2^2] dt \\ &= \mathcal{J}_{\text{SM}}(\boldsymbol{\theta}; g(\cdot)^2), \quad (18) \end{aligned}$$

where (i) is due to Girsanov Theorem II [34, Theorem 8.6.6], and (ii) is due to the martingale property of Itô integrals. Combining Eqs. (14), (15) and (18) completes the proof. \square

Theorem 2. *Suppose $p(\mathbf{x})$ and $q(\mathbf{x})$ have continuous second-order derivatives and finite second moments. Let $\{\mathbf{x}(t)\}_{t \in [0, T]}$ be the diffusion process defined by the SDE in Eq. (1). We use p_t and q_t to denote the distributions of $\mathbf{x}(t)$ when $\mathbf{x}(0) \sim p$ and $\mathbf{x}(0) \sim q$, and assume they satisfy the same assumptions in Appendix A. Under the conditions $q_T = \pi$ and $\mathbf{s}_\theta(\mathbf{x}, t) \equiv \nabla_{\mathbf{x}} \log q_t(\mathbf{x})$ for all $t \in [0, T]$, we have the following equivalence in distributions*

$$p_\theta^{\text{SDE}} = p_\theta^{\text{ODE}} = q. \quad (9)$$

Moreover, we have

$$D_{\text{KL}}(p \parallel p_\theta^{\text{SDE}}) = \mathcal{J}_{\text{SM}}(\boldsymbol{\theta}; g(\cdot)^2) + D_{\text{KL}}(p_T \parallel \pi). \quad (10)$$

Proof. When $\pi = q_T$ and $\mathbf{s}_\theta(\mathbf{x}, t) \equiv \nabla_{\mathbf{x}} \log q_t(\mathbf{x})$, the reverse-time SDE that defines p_θ^{SDE} , i.e.,

$$d\hat{\mathbf{x}} = [\mathbf{f}(\hat{\mathbf{x}}, t) - g(t)^2 \mathbf{s}_\theta(\hat{\mathbf{x}}, t)] dt + g(t) d\bar{\mathbf{w}}, \quad \hat{\mathbf{x}}_\theta(T) \sim \pi, \quad (19)$$

becomes equivalent to

$$d\hat{\mathbf{x}} = [\mathbf{f}(\hat{\mathbf{x}}, t) - g(t)^2 \nabla_{\hat{\mathbf{x}}} \log q_t(\hat{\mathbf{x}})] dt + g(t) d\bar{\mathbf{w}}, \quad \hat{\mathbf{x}}_\theta(T) \sim q_T, \quad (20)$$

which yields the same stochastic process as the following forward-time SDE

$$d\hat{\mathbf{x}} = \mathbf{f}(\hat{\mathbf{x}}, t) dt + g(t) d\mathbf{w}, \quad \hat{\mathbf{x}}_\theta(0) \sim q. \quad (21)$$

Since $\hat{\mathbf{x}}_\theta(0) \sim p_\theta^{\text{SDE}}$ by definition, we immediately have $p_\theta^{\text{SDE}} = q$. Similarly, the ODE that defines p_θ^{ODE} is

$$\frac{d\tilde{\mathbf{x}}}{dt} = \mathbf{f}_\theta(\tilde{\mathbf{x}}, t) - \frac{1}{2} g(t)^2 \mathbf{s}_\theta(\tilde{\mathbf{x}}, t), \quad \tilde{\mathbf{x}}_\theta(T) \sim \pi, \quad (22)$$

which is equivalent to the following when $q_T = \pi$ and $\mathbf{s}_\theta(\mathbf{x}, t) \equiv \nabla_{\mathbf{x}} \log q_t(\mathbf{x})$,

$$\frac{d\tilde{\mathbf{x}}}{dt} = \mathbf{f}_\theta(\tilde{\mathbf{x}}, t) - \frac{1}{2}g(t)^2 \nabla_{\tilde{\mathbf{x}}} \log q_t(\tilde{\mathbf{x}}, t), \quad \tilde{\mathbf{x}}_\theta(T) \sim q_T. \quad (23)$$

The theory of probability flow ODEs [48] guarantees that Eq. (21) and Eq. (23) share the same set of marginal distributions, $\{q_t\}_{t \in [0, T]}$, which implies that $\tilde{\mathbf{x}}_\theta(0) \sim q$. Since by definition $\tilde{\mathbf{x}}_\theta(0) \sim p_\theta^{\text{ODE}}$, we have $p_\theta^{\text{ODE}} = q$.

The next part of the theorem can be proved by first rewriting the KL divergence from p to q in an integral form:

$$\begin{aligned} D_{\text{KL}}(p(\mathbf{x}) \parallel q(\mathbf{x})) &\stackrel{(i)}{=} D_{\text{KL}}(p_0(\mathbf{x}) \parallel q_0(\mathbf{x})) - D_{\text{KL}}(p_T(\mathbf{x}) \parallel q_T(\mathbf{x})) + D_{\text{KL}}(p_T(\mathbf{x}) \parallel q_T(\mathbf{x})) \\ &\stackrel{(ii)}{=} \int_T^0 \frac{\partial D_{\text{KL}}(p_t(\mathbf{x}) \parallel q_t(\mathbf{x}))}{\partial t} dt + D_{\text{KL}}(p_T(\mathbf{x}) \parallel q_T(\mathbf{x})), \end{aligned} \quad (24)$$

where (i) holds due to our definition $p_0(\mathbf{x}) \equiv p(\mathbf{x})$ and $q_0(\mathbf{x}) \equiv q(\mathbf{x})$; (ii) is due to the fundamental theorem of calculus.

Next, we show how to rewrite Eq. (24) as a mixture of score matching losses. The Fokker–Planck equation for the SDE in Eq. (1) describes the time-evolution of the stochastic process’s associated probability density function, and is given by

$$\frac{\partial p_t(\mathbf{x})}{\partial t} = \nabla_{\mathbf{x}} \cdot \left(\frac{1}{2}g^2(t)p_t(\mathbf{x})\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) - \mathbf{f}(\mathbf{x}, t)p_t(\mathbf{x}) \right) = \nabla_{\mathbf{x}} \cdot (\mathbf{h}_p(\mathbf{x}, t)p_t(\mathbf{x})),$$

where for simplified notations we define $\mathbf{h}_p(\mathbf{x}, t) := \frac{1}{2}g^2(t)\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) - \mathbf{f}(\mathbf{x}, t)$. Similarly, $\frac{\partial q_t(\mathbf{x})}{\partial t} = \nabla_{\mathbf{x}} \cdot (\mathbf{h}_q(\mathbf{x}, t)q_t(\mathbf{x}))$. Since we assume $\log p_t(\mathbf{x})$ and $\log q_t(\mathbf{x})$ are smooth functions with at most polynomial growth at infinity (assumption (xii)), we have $\lim_{\mathbf{x} \rightarrow \infty} \mathbf{h}_p(\mathbf{x}, t)p_t(\mathbf{x}) = \mathbf{0}$ and $\lim_{\mathbf{x} \rightarrow \infty} \mathbf{h}_q(\mathbf{x}, t)q_t(\mathbf{x}) = \mathbf{0}$ for all t . Then, the time-derivative of $D_{\text{KL}}(p_t \parallel q_t)$ can be rewritten in the following way:

$$\begin{aligned} \frac{\partial D_{\text{KL}}(p_t(\mathbf{x}) \parallel q_t(\mathbf{x}))}{\partial t} &= \frac{\partial}{\partial t} \int p_t(\mathbf{x}) \log \frac{p_t(\mathbf{x})}{q_t(\mathbf{x})} d\mathbf{x} \\ &= \int \frac{\partial p_t(\mathbf{x})}{\partial t} \log \frac{p_t(\mathbf{x})}{q_t(\mathbf{x})} d\mathbf{x} + \underbrace{\int \frac{\partial p_t(\mathbf{x})}{\partial t} d\mathbf{x}}_{=0} - \int \frac{p_t(\mathbf{x})}{q_t(\mathbf{x})} \frac{\partial q_t(\mathbf{x})}{\partial t} d\mathbf{x} \\ &= \int \nabla_{\mathbf{x}} \cdot (\mathbf{h}_p(\mathbf{x}, t)p_t(\mathbf{x})) \log \frac{p_t(\mathbf{x})}{q_t(\mathbf{x})} d\mathbf{x} - \int \frac{p_t(\mathbf{x})}{q_t(\mathbf{x})} \nabla_{\mathbf{x}} \cdot (\mathbf{h}_q(\mathbf{x}, t)q_t(\mathbf{x})) d\mathbf{x} \\ &\stackrel{(i)}{=} - \int p_t(\mathbf{x}) [\mathbf{h}_p^\top(\mathbf{x}, t) - \mathbf{h}_q^\top(\mathbf{x}, t)] [\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) - \nabla_{\mathbf{x}} \log q_t(\mathbf{x})] d\mathbf{x} \\ &= - \frac{1}{2} \int p_t(\mathbf{x}) g(t)^2 \|\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) - \nabla_{\mathbf{x}} \log q_t(\mathbf{x})\|_2^2 d\mathbf{x}, \end{aligned}$$

where (i) is due to integration by parts. Combining with Eq. (24), we can conclude that

$$D_{\text{KL}}(p \parallel q) = \frac{1}{2} \int_0^T \mathbb{E}_{\mathbf{x} \sim p_t(\mathbf{x})} [g(t)^2 \|\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) - \nabla_{\mathbf{x}} \log q_t(\mathbf{x})\|_2^2] dt + D_{\text{KL}}(p_T \parallel q_T). \quad (25)$$

Since $p_\theta^{\text{SDE}} = q$ and $q_T = \pi$, we also have

$$\begin{aligned} D_{\text{KL}}(p \parallel p_\theta^{\text{SDE}}) &= \frac{1}{2} \int_0^T \mathbb{E}_{\mathbf{x} \sim p_t(\mathbf{x})} [g(t)^2 \|\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) - \nabla_{\mathbf{x}} \log q_t(\mathbf{x})\|_2^2] dt + D_{\text{KL}}(p_T \parallel q_T) \\ &= \mathcal{J}_{\text{SM}}(\boldsymbol{\theta}; g(\cdot)^2) + D_{\text{KL}}(p_T \parallel q_T), \end{aligned} \quad (26)$$

which completes the proof. \square

Using a similar technique to Theorem 2, we can express the entropy of a distribution in terms of a time-dependent score function, as detailed in the following theorem.

Theorem 4. Let $\mathcal{H}(p(\mathbf{x}))$ be the differential entropy of the initial probability density $p(\mathbf{x})$. Under the same conditions in Theorem 2, we have

$$\mathcal{H}(p(\mathbf{x})) = \mathcal{H}(p_T(\mathbf{x})) + \frac{1}{2} \int_0^T \mathbb{E}_{\mathbf{x} \sim p_t(\mathbf{x})} \left[2\mathbf{f}(\mathbf{x}, t)^\top \nabla_{\mathbf{x}} \log p_t(\mathbf{x}) - g(t)^2 \|\nabla_{\mathbf{x}} \log p_t(\mathbf{x})\|_2^2 \right] dt. \quad (27)$$

$$= \mathcal{H}(p_T(\mathbf{x})) - \frac{1}{2} \int_0^T \mathbb{E}_{\mathbf{x} \sim p_t(\mathbf{x})} \left[2\nabla \cdot \mathbf{f}(\mathbf{x}, t) + g(t)^2 \|\nabla_{\mathbf{x}} \log p_t(\mathbf{x})\|_2^2 \right] dt. \quad (28)$$

Proof. Once more we proceed analogously to the proofs of Theorem 2. We have

$$\mathcal{H}(p(\mathbf{x})) - \mathcal{H}(p_T(\mathbf{x})) = \int_T^0 \frac{\partial}{\partial t} \mathcal{H}(p_t(\mathbf{x})) dt. \quad (29)$$

Expanding the integrand, we have

$$\begin{aligned} \frac{\partial}{\partial t} \mathcal{H}(p_t(\mathbf{x})) &= -\frac{\partial}{\partial t} \int p_t(\mathbf{x}) \log p_t(\mathbf{x}) d\mathbf{x} \\ &= -\int \frac{\partial p_t(\mathbf{x})}{\partial t} \log p_t(\mathbf{x}) + \frac{\partial p_t(\mathbf{x})}{\partial t} d\mathbf{x} \\ &= -\int \frac{\partial p_t(\mathbf{x})}{\partial t} \log p_t(\mathbf{x}) d\mathbf{x} - \underbrace{\frac{\partial}{\partial t} \int p_t(\mathbf{x}) d\mathbf{x}}_{=1} \\ &= -\int \nabla_{\mathbf{x}} \cdot (\mathbf{h}_p(\mathbf{x}, t) p_t(\mathbf{x})) \log p_t(\mathbf{x}) d\mathbf{x} \\ &\stackrel{(i)}{=} \int p_t(\mathbf{x}) \mathbf{h}_p^\top(\mathbf{x}, t) \nabla_{\mathbf{x}} \log p_t(\mathbf{x}) d\mathbf{x} \\ &= \frac{1}{2} \mathbb{E}_{\mathbf{x} \sim p_t(\mathbf{x})} [g(t)^2 \|\nabla_{\mathbf{x}} \log p_t(\mathbf{x})\|_2^2 - 2\mathbf{f}(\mathbf{x}, t)^\top \nabla_{\mathbf{x}} \log p_t(\mathbf{x})], \end{aligned}$$

where again (i) follows from integration by parts and the limiting behaviour of \mathbf{h}_p given by assumption (xii). Plugging this expression in for the integrand in Eq. (29) then completes the proof for Eq. (27). For Eq. (28), we can once again perform integration by parts and leverage the limiting behavior of $p_t(\mathbf{x})$ in assumption (xii) to get

$$\mathbb{E}_{p_t(\mathbf{x})} [\mathbf{f}(\mathbf{x}, t)^\top \nabla_{\mathbf{x}} \log p_t(\mathbf{x})] = \int \mathbf{f}(\mathbf{x}, t)^\top \nabla_{\mathbf{x}} p_t(\mathbf{x}) d\mathbf{x} = -\int p_t(\mathbf{x}) \nabla \cdot \mathbf{f}(\mathbf{x}, t) d\mathbf{x},$$

which establishes the equivalence between Eq. (28) and Eq. (27). \square

Remark The formula in Theorem 4 provides a new way to estimate the entropy of a data distribution from i.i.d. samples. Specifically, given $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\} \stackrel{\text{i.i.d.}}{\sim} p(\mathbf{x})$ and an SDE like Eq. (1), we can first apply score matching to train a time-dependent score-based model such that $\mathbf{s}_\theta(\mathbf{x}, t) \approx \nabla_{\mathbf{x}} \log p_t(\mathbf{x})$, and then plug $\mathbf{s}_\theta(\mathbf{x}, t)$ into Eq. (27) to obtain the following estimator of $\mathcal{H}(p(\mathbf{x}))$:

$$\mathcal{H}(p_T(\mathbf{x})) + \frac{1}{2N} \sum_{i=1}^N \int_0^T \left[2\mathbf{f}(\mathbf{x}_i, t)^\top \mathbf{s}_\theta(\mathbf{x}_i, t) - g(t)^2 \|\mathbf{s}_\theta(\mathbf{x}_i, t)\|_2^2 \right] dt,$$

or plug it into Eq. (28) to obtain the following alternative estimator

$$\mathcal{H}(p_T(\mathbf{x})) - \frac{1}{2N} \sum_{i=1}^N \int_0^T \left[2\nabla \cdot \mathbf{f}(\mathbf{x}_i, t) + g(t)^2 \|\mathbf{s}_\theta(\mathbf{x}_i, t)\|_2^2 \right] dt.$$

Both estimators can be computed from a score-based model alone, and do not require training a density model.

Theorem 5. Let $p_{0t}(\mathbf{x}' | \mathbf{x})$ denote the transition kernel from $p_0(\mathbf{x})$ to $p_t(\mathbf{x})$ for any $t \in (0, T]$. With the same conditions and notations in Theorem 1, we have

$$-\mathbb{E}_{p(\mathbf{x})}[\log p_{\theta}^{\text{SDE}}(\mathbf{x})] \leq -\mathbb{E}_{p_T(\mathbf{x})}[\log \pi(\mathbf{x})] + \frac{1}{2} \int_0^T \mathbb{E}_{\mathbf{x} \sim p_t(\mathbf{x})} [2g(t)^2 \nabla \cdot \mathbf{s}_{\theta}(\mathbf{x}, t) + g(t)^2 \|\mathbf{s}_{\theta}(\mathbf{x}, t)\|_2^2 - 2\nabla \cdot \mathbf{f}(\mathbf{x}, t)] dt. \quad (30)$$

$$\begin{aligned} &= -\mathbb{E}_{p_T(\mathbf{x})}[\log \pi(\mathbf{x})] \\ &+ \frac{1}{2} \int_0^T \mathbb{E}_{p_{0t}(\mathbf{x}' | \mathbf{x}) p(\mathbf{x})} [g(t)^2 \|\mathbf{s}_{\theta}(\mathbf{x}', t) - \nabla_{\mathbf{x}'} \log p_{0t}(\mathbf{x}' | \mathbf{x})\|_2^2 \\ &- g(t)^2 \|\nabla_{\mathbf{x}'} \log p_{0t}(\mathbf{x}' | \mathbf{x})\|_2^2 - 2\nabla \cdot \mathbf{f}(\mathbf{x}', t)] dt. \end{aligned} \quad (31)$$

Proof. Since $-\mathbb{E}_{p(\mathbf{x})}[\log p_{\theta}^{\text{SDE}}(\mathbf{x})] = D_{\text{KL}}(p \parallel p_{\theta}^{\text{SDE}}) + \mathcal{H}(p)$, we can combine Theorem 1 and Theorem 4 to obtain

$$\begin{aligned} -\mathbb{E}_{p(\mathbf{x})}[\log p_{\theta}^{\text{SDE}}(\mathbf{x})] &\leq \frac{1}{2} \int_0^T \mathbb{E}_{p_t(\mathbf{x})} [g(t)^2 \|\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) - \mathbf{s}_{\theta}(\mathbf{x}, t)\|_2^2] dt + D_{\text{KL}}(p_T \parallel \pi) \\ &+ \mathcal{H}(p_T(\mathbf{x})) - \frac{1}{2} \int_0^T \mathbb{E}_{p_t(\mathbf{x})} [2\nabla \cdot \mathbf{f}(\mathbf{x}, t) + g(t)^2 \|\nabla_{\mathbf{x}} \log p_t(\mathbf{x})\|_2^2] dt \\ &= -\mathbb{E}_{p_T(\mathbf{x})}[\log \pi(\mathbf{x})] \\ &+ \frac{1}{2} \int_0^T \mathbb{E}_{p_t(\mathbf{x})} [g(t)^2 \|\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) - \mathbf{s}_{\theta}(\mathbf{x}, t)\|_2^2 - g(t)^2 \|\nabla_{\mathbf{x}} \log p_t(\mathbf{x})\|_2^2] dt \\ &- \int_0^T \mathbb{E}_{p_t(\mathbf{x})} [\nabla \cdot \mathbf{f}(\mathbf{x}, t)] dt. \end{aligned} \quad (32)$$

The second term of Eq. (32) can be simplified via integration by parts

$$\begin{aligned} &\frac{1}{2} \int_0^T \mathbb{E}_{p_t(\mathbf{x})} [g(t)^2 \|\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) - \mathbf{s}_{\theta}(\mathbf{x}, t)\|_2^2 - g(t)^2 \|\nabla_{\mathbf{x}} \log p_t(\mathbf{x})\|_2^2] dt \\ &= \frac{1}{2} \int_0^T \mathbb{E}_{p_t(\mathbf{x})} [g(t)^2 \|\mathbf{s}_{\theta}(\mathbf{x}, t)\|_2^2 - 2g(t)^2 \mathbf{s}_{\theta}(\mathbf{x}, t)^{\top} \nabla_{\mathbf{x}} \log p_t(\mathbf{x})] dt \\ &= \frac{1}{2} \int_0^T \mathbb{E}_{p_t(\mathbf{x})} [g(t)^2 \|\mathbf{s}_{\theta}(\mathbf{x}, t)\|_2^2] dt - \int_0^T \mathbb{E}_{p_t(\mathbf{x})} [g(t)^2 \mathbf{s}_{\theta}(\mathbf{x}, t)^{\top} \nabla_{\mathbf{x}} \log p_t(\mathbf{x})] dt \\ &= \frac{1}{2} \int_0^T \mathbb{E}_{p_t(\mathbf{x})} [g(t)^2 \|\mathbf{s}_{\theta}(\mathbf{x}, t)\|_2^2] dt - \int_0^T g(t)^2 \int p_t(\mathbf{x}) \mathbf{s}_{\theta}(\mathbf{x}, t)^{\top} \nabla_{\mathbf{x}} \log p_t(\mathbf{x}) d\mathbf{x} dt \\ &= \frac{1}{2} \int_0^T \mathbb{E}_{p_t(\mathbf{x})} [g(t)^2 \|\mathbf{s}_{\theta}(\mathbf{x}, t)\|_2^2] dt - \int_0^T g(t)^2 \int \mathbf{s}_{\theta}(\mathbf{x}, t)^{\top} \nabla_{\mathbf{x}} p_t(\mathbf{x}) d\mathbf{x} dt \\ &\stackrel{(i)}{=} \frac{1}{2} \int_0^T \mathbb{E}_{p_t(\mathbf{x})} [g(t)^2 \|\mathbf{s}_{\theta}(\mathbf{x}, t)\|_2^2] dt + \int_0^T g(t)^2 \int p_t(\mathbf{x}) \nabla \cdot \mathbf{s}_{\theta}(\mathbf{x}, t) d\mathbf{x} dt \\ &= \frac{1}{2} \int_0^T \mathbb{E}_{p_t(\mathbf{x})} [g(t)^2 \|\mathbf{s}_{\theta}(\mathbf{x}, t)\|_2^2 + 2g(t)^2 \nabla \cdot \mathbf{s}_{\theta}(\mathbf{x}, t)] dt, \end{aligned} \quad (33)$$

where (i) is due to integration by parts and the limiting behavior of $p_t(\mathbf{x})$ given by assumption (xii). Combining Eq. (33) and Eq. (32) completes the proof for Eq. (30).

The proof for Eq. (31) parallels that of denoising score matching [56]. Observe that $p_t(\mathbf{x}) = \int p(\mathbf{x}')p_{0t}(\mathbf{x} | \mathbf{x}') d\mathbf{x}'$. As a result,

$$\begin{aligned}
& \int_0^T \mathbb{E}_{p_t(\mathbf{x})} [g(t)^2 \mathbf{s}_\theta(\mathbf{x}, t)^\top \nabla_{\mathbf{x}} \log p_t(\mathbf{x})] dt = \int_0^T g(t)^2 \int \mathbf{s}_\theta(\mathbf{x}, t)^\top \nabla_{\mathbf{x}} p_t(\mathbf{x}) d\mathbf{x} dt \\
&= \int_0^T g(t)^2 \int \mathbf{s}_\theta(\mathbf{x}, t)^\top \nabla_{\mathbf{x}} \int p(\mathbf{x}') p_{0t}(\mathbf{x} | \mathbf{x}') d\mathbf{x}' d\mathbf{x} dt \\
&= \int_0^T g(t)^2 \int \mathbf{s}_\theta(\mathbf{x}, t)^\top \int p(\mathbf{x}') \nabla_{\mathbf{x}} p_{0t}(\mathbf{x} | \mathbf{x}') d\mathbf{x}' d\mathbf{x} dt \\
&= \int_0^T g(t)^2 \int \mathbf{s}_\theta(\mathbf{x}, t)^\top \int p(\mathbf{x}') p_{0t}(\mathbf{x} | \mathbf{x}') \nabla_{\mathbf{x}} \log p_{0t}(\mathbf{x} | \mathbf{x}') d\mathbf{x}' d\mathbf{x} dt \\
&= \int_0^T \mathbb{E}_{p(\mathbf{x})p_{0t}(\mathbf{x}'|\mathbf{x})} [g(t)^2 \mathbf{s}_\theta(\mathbf{x}', t)^\top \nabla_{\mathbf{x}'} \log p_{0t}(\mathbf{x}' | \mathbf{x})] dt. \tag{34}
\end{aligned}$$

Substituting Eq. (34) into the second term of Eq. (32), we have

$$\begin{aligned}
& \frac{1}{2} \int_0^T \mathbb{E}_{p_t(\mathbf{x})} [g(t)^2 \|\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) - \mathbf{s}_\theta(\mathbf{x}, t)\|_2^2 - g(t)^2 \|\nabla_{\mathbf{x}} \log p_t(\mathbf{x})\|_2^2] dt \\
&= \frac{1}{2} \int_0^T \mathbb{E}_{p_t(\mathbf{x})} [g(t)^2 \|\mathbf{s}_\theta(\mathbf{x}, t)\|_2^2 - 2g(t)^2 \mathbf{s}_\theta(\mathbf{x}, t)^\top \nabla_{\mathbf{x}} \log p_t(\mathbf{x})] dt \\
&= \frac{1}{2} \int_0^T \mathbb{E}_{p_t(\mathbf{x})} [g(t)^2 \|\mathbf{s}_\theta(\mathbf{x}, t)\|_2^2 - 2g(t)^2 \mathbf{s}_\theta(\mathbf{x}, t)^\top \nabla_{\mathbf{x}} \log p_t(\mathbf{x})] dt \\
&= \frac{1}{2} \int_0^T \mathbb{E}_{p(\mathbf{x})p_{0t}(\mathbf{x}'|\mathbf{x})} [g(t)^2 \|\mathbf{s}_\theta(\mathbf{x}', t)\|_2^2 - 2g(t)^2 \mathbf{s}_\theta(\mathbf{x}', t)^\top \nabla_{\mathbf{x}'} \log p_{0t}(\mathbf{x}' | \mathbf{x})] dt \\
&= \frac{1}{2} \int_0^T \mathbb{E}_{p(\mathbf{x})p_{0t}(\mathbf{x}'|\mathbf{x})} [g(t)^2 \|\mathbf{s}_\theta(\mathbf{x}', t) - \nabla_{\mathbf{x}'} \log p_{0t}(\mathbf{x}' | \mathbf{x})\|_2^2 - g(t)^2 \|\nabla_{\mathbf{x}'} \log p_{0t}(\mathbf{x}' | \mathbf{x})\|_2^2] dt. \tag{35}
\end{aligned}$$

We can now complete the proof for Eq. (31) by combining Eq. (35) and Eq. (32). \square

Theorem 3. Let $p_{0t}(\mathbf{x}' | \mathbf{x})$ denote the transition distribution from $p_0(\mathbf{x})$ to $p_t(\mathbf{x})$ for the SDE in Eq. (1). With the same notations and conditions in Theorem 1, we have

$$-\log p_\theta^{\text{SDE}}(\mathbf{x}) \leq \mathcal{L}_\theta^{\text{SM}}(\mathbf{x}) = \mathcal{L}_\theta^{\text{DSM}}(\mathbf{x}), \tag{11}$$

where $\mathcal{L}_\theta^{\text{SM}}(\mathbf{x})$ is defined as

$$-\mathbb{E}_{p_{0T}(\mathbf{x}'|\mathbf{x})} [\log \pi(\mathbf{x}')] + \frac{1}{2} \int_0^T \mathbb{E}_{p_{0t}(\mathbf{x}'|\mathbf{x})} \left[2g(t)^2 \nabla_{\mathbf{x}'} \cdot \mathbf{s}_\theta(\mathbf{x}', t) + g(t)^2 \|\mathbf{s}_\theta(\mathbf{x}', t)\|_2^2 - 2\nabla_{\mathbf{x}'} \cdot \mathbf{f}(\mathbf{x}', t) \right] dt,$$

and $\mathcal{L}_\theta^{\text{DSM}}(\mathbf{x})$ is given by

$$\begin{aligned}
& -\mathbb{E}_{p_{0T}(\mathbf{x}'|\mathbf{x})} [\log \pi(\mathbf{x}')] + \frac{1}{2} \int_0^T \mathbb{E}_{p_{0t}(\mathbf{x}'|\mathbf{x})} \left[g(t)^2 \|\mathbf{s}_\theta(\mathbf{x}', t) - \nabla_{\mathbf{x}'} \log p_{0t}(\mathbf{x}' | \mathbf{x})\|_2^2 \right] dt \\
& \quad - \frac{1}{2} \int_0^T \mathbb{E}_{p_{0t}(\mathbf{x}'|\mathbf{x})} \left[g(t)^2 \|\nabla_{\mathbf{x}'} \log p_{0t}(\mathbf{x}' | \mathbf{x})\|_2^2 + 2\nabla_{\mathbf{x}'} \cdot \mathbf{f}(\mathbf{x}', t) \right] dt.
\end{aligned}$$

Proof. The result in Theorem 5 can be re-written as

$$\begin{aligned}
-\mathbb{E}_{p(\mathbf{x})} [\log p_\theta^{\text{SDE}}(\mathbf{x})] &\leq -\mathbb{E}_{p(\mathbf{x})p_{0T}(\mathbf{x}'|\mathbf{x})} [\log \pi(\mathbf{x}')] + \frac{1}{2} \int_0^T \mathbb{E}_{p(\mathbf{x})p_{0t}(\mathbf{x}'|\mathbf{x})} [2g(t)^2 \nabla \cdot \mathbf{s}_\theta(\mathbf{x}', t) \\
& \quad + g(t)^2 \|\mathbf{s}_\theta(\mathbf{x}', t)\|_2^2 - 2\nabla \cdot \mathbf{f}(\mathbf{x}', t)] dt. \\
&= -\mathbb{E}_{p(\mathbf{x})p_{0T}(\mathbf{x}'|\mathbf{x})} [\log \pi(\mathbf{x}')] \\
& \quad + \frac{1}{2} \int_0^T \mathbb{E}_{p(\mathbf{x})p_{0t}(\mathbf{x}'|\mathbf{x})} [g(t)^2 \|\mathbf{s}_\theta(\mathbf{x}', t) - \nabla_{\mathbf{x}'} \log p_{0t}(\mathbf{x}' | \mathbf{x})\|_2^2 \\
& \quad - g(t)^2 \|\nabla_{\mathbf{x}'} \log p_{0t}(\mathbf{x}' | \mathbf{x})\|_2^2 - 2\nabla \cdot \mathbf{f}(\mathbf{x}', t)] dt.
\end{aligned}$$

Given a fixed SDE (and its transition kernel $p_{0t}(\mathbf{x}' | \mathbf{x})$), Theorem 5 holds for any data distribution p that satisfies our assumptions. Leveraging proof by contradiction, we can easily see that $\mathbb{E}_{p(\mathbf{x})}$ in both sides of Eqs. (30) and (31) can be cancelled to get

$$-\log p_{\theta}^{\text{SDE}}(\mathbf{x}) \leq \mathcal{L}_{\theta}^{\text{SM}}(\mathbf{x}) = \mathcal{L}_{\theta}^{\text{DSM}}(\mathbf{x}),$$

which finishes the proof. \square

B Numerical stability

In our previous theoretical discussion, we always assume that data are perturbed with an SDE starting from $t = 0$. However, in practical implementations, $t = 0$ often leads to numerical instability. As a pragmatic solution, we choose a small non-zero starting time $\epsilon > 0$, and consider the SDE in the time horizon $[\epsilon, T]$. Using the same proof techniques, we can easily see that when the time horizon is $[\epsilon, T]$ instead of $[0, T]$, the original bound in Theorem 1,

$$\begin{aligned} D_{\text{KL}}(p \parallel p_{\theta}^{\text{SDE}}) &\leq \mathcal{J}_{\text{SM}}(\theta; g(\cdot)^2) + D_{\text{KL}}(p_T \parallel \pi) \\ &= \frac{1}{2} \int_0^T \mathbb{E}_{p_t(\mathbf{x})} [g(t)^2 \|\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) - \mathbf{s}_{\theta}(\mathbf{x}, t)\|_2^2] dt + D_{\text{KL}}(p_T \parallel \pi) \end{aligned}$$

shall be replaced with

$$D_{\text{KL}}(\tilde{p} \parallel \tilde{p}_{\theta}^{\text{SDE}}) \leq \frac{1}{2} \int_{\epsilon}^T \mathbb{E}_{p_t(\mathbf{x})} [g(t)^2 \|\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) - \mathbf{s}_{\theta}(\mathbf{x}, t)\|_2^2] dt + D_{\text{KL}}(p_T \parallel \pi) \quad (36)$$

where $\tilde{p}(\mathbf{x}) := \int p(\tilde{\mathbf{x}}) p_{0\epsilon}(\mathbf{x} | \tilde{\mathbf{x}}) d\mathbf{x}$, and $\tilde{p}_{\theta}^{\text{SDE}}$ denotes the marginal distribution of $\hat{\mathbf{x}}_{\theta}(\epsilon)$. Here the stochastic process $\{\hat{\mathbf{x}}_{\theta}(t)\}_{t \in [0, T]}$ is defined according to Eq. (5). When ϵ is sufficiently small, we always have

$$D_{\text{KL}}(\tilde{p} \parallel \tilde{p}_{\theta}^{\text{SDE}}) \approx D_{\text{KL}}(p \parallel p_{\theta}^{\text{SDE}}),$$

so we train with Eq. (36) to approximately maximize the model likelihood for p_{θ}^{SDE} . However, at test time, we should report the likelihood bound for p_{θ}^{SDE} for mathematical rigor, not $\tilde{p}_{\theta}^{\text{SDE}}$. To this end, we first derive an analogous result to Theorem 3 with the time horizon $[\epsilon, T]$, given as below.

Theorem 6. *Let $p_{0t}(\mathbf{x}' | \mathbf{x})$ denote the transition distribution from $p_0(\mathbf{x})$ to $p_t(\mathbf{x})$ for the SDE in Eq. (1). With the same notations and conditions in Theorem 3, as well as the definitions of \tilde{p} and $\tilde{p}_{\theta}^{\text{SDE}}$ given above, we have*

$$-\mathbb{E}_{p_{0\epsilon}(\mathbf{x}' | \mathbf{x})} [\log \tilde{p}_{\theta}^{\text{SDE}}(\mathbf{x}')] \leq \mathcal{L}_{\theta}^{\text{SM}}(\mathbf{x}, \epsilon) = \mathcal{L}_{\theta}^{\text{DSM}}(\mathbf{x}, \epsilon), \quad (37)$$

where $\mathcal{L}_{\theta}^{\text{SM}}(\mathbf{x}, \epsilon)$ is defined as

$$-\mathbb{E}_{p_{0T}(\mathbf{x}' | \mathbf{x})} [\log \pi(\mathbf{x}')] + \frac{1}{2} \int_{\epsilon}^T \mathbb{E}_{p_{0t}(\mathbf{x}' | \mathbf{x})} \left[2g(t)^2 \nabla_{\mathbf{x}'} \cdot \mathbf{s}_{\theta}(\mathbf{x}', t) + g(t)^2 \|\mathbf{s}_{\theta}(\mathbf{x}', t)\|_2^2 - 2\nabla_{\mathbf{x}'} \cdot \mathbf{f}(\mathbf{x}', t) \right] dt,$$

and $\mathcal{L}_{\theta}^{\text{DSM}}(\mathbf{x}, \epsilon)$ is given by

$$\begin{aligned} -\mathbb{E}_{p_{0T}(\mathbf{x}' | \mathbf{x})} [\log \pi(\mathbf{x}')] + \frac{1}{2} \int_{\epsilon}^T \mathbb{E}_{p_{0t}(\mathbf{x}' | \mathbf{x})} \left[g(t)^2 \|\mathbf{s}_{\theta}(\mathbf{x}', t) - \nabla_{\mathbf{x}'} \log p_{0t}(\mathbf{x}' | \mathbf{x})\|_2^2 \right] dt \\ - \frac{1}{2} \int_{\epsilon}^T \mathbb{E}_{p_{0t}(\mathbf{x}' | \mathbf{x})} \left[g(t)^2 \|\nabla_{\mathbf{x}'} \log p_{0t}(\mathbf{x}' | \mathbf{x})\|_2^2 + 2\nabla_{\mathbf{x}'} \cdot \mathbf{f}(\mathbf{x}', t) \right] dt. \end{aligned}$$

Proof. The proof closely parallels that of Theorem 3, by noting that $\tilde{p}(\mathbf{x}) = \int p(\mathbf{x}') p_{0\epsilon}(\mathbf{x} | \mathbf{x}') d\mathbf{x}'$. \square

Although $\tilde{p}_{\theta}^{\text{SDE}}$ is a probabilistic model for \tilde{p} , we can transform it into a probabilistic model for p leveraging a denoising distribution $q_{\theta}(\mathbf{x} | \mathbf{x}')$ that approximately converts \tilde{p} to p . Suppose $p_{0\epsilon}(\mathbf{x}' | \mathbf{x}) = \mathcal{N}(\mathbf{x}' | \alpha \mathbf{x}, \beta^2 \mathbf{I})$. Inspired by Tweedie's formula, we choose

$$q_{\theta}(\mathbf{x} | \mathbf{x}') := \mathcal{N}\left(\mathbf{x} \mid \frac{\mathbf{x}'}{\alpha} + \frac{\beta^2}{\alpha} \mathbf{s}_{\theta}(\mathbf{x}', \epsilon), \frac{\beta^2}{\alpha^2} \mathbf{I}\right),$$

and define $p_\theta(\mathbf{x}) := \int q_\theta(\mathbf{x} | \mathbf{x}') \tilde{p}_\theta^{\text{SDE}}(\mathbf{x}') d\mathbf{x}'$, which is a probabilistic model for p . With slight abuse of notation, we identify p_θ with p_θ^{SDE} in Table 2. With Jensen’s inequality, we have

$$-\log p_\theta(\mathbf{x}) \leq -\mathbb{E}_{p_{0\epsilon}(\mathbf{x}'|\mathbf{x})} \left[\log \frac{q_\theta(\mathbf{x} | \mathbf{x}') \tilde{p}_\theta^{\text{SDE}}(\mathbf{x}')}{p_{0\epsilon}(\mathbf{x}' | \mathbf{x})} \right].$$

Combined with Theorem 6, we have

$$-\log p_\theta(\mathbf{x}) \leq -\mathbb{E}_{p_{0\epsilon}(\mathbf{x}'|\mathbf{x})} [\log q_\theta(\mathbf{x} | \mathbf{x}') - \log p_{0\epsilon}(\mathbf{x}' | \mathbf{x})] + \mathcal{L}_\theta^{\text{SM}}(\mathbf{x}, \epsilon) \quad (38)$$

$$= -\mathbb{E}_{p_{0\epsilon}(\mathbf{x}'|\mathbf{x})} [\log q_\theta(\mathbf{x} | \mathbf{x}') - \log p_{0\epsilon}(\mathbf{x}' | \mathbf{x})] + \mathcal{L}_\theta^{\text{DSM}}(\mathbf{x}, \epsilon) \quad (39)$$

The above bound Eq. (39) was applied to both computing the test-time likelihood bounds in Table 2, and training the flow model used in variational dequantization. Note that it was not used to train the time-dependent score-based model.

In practice, we choose $\epsilon = 10^{-5}$ for VP SDEs and $\epsilon = 10^{-2}$ for subVP SDEs, except that on ImageNet we use $\epsilon = 5 \times 10^{-5}$ for VP SDE models trained with likelihood weighting and importance sampling. Note that [48] chooses $\epsilon = 10^{-5}$ for all cases. We found that when using likelihood weighting and optionally importance sampling, $\epsilon = 10^{-5}$ for subVP SDEs can cause stiffness for numerical ODE solvers. In contrast, using $\epsilon = 10^{-2}$ for subVP SDEs sidesteps numerical issues without hurting the performance for score-based models trained with original weightings in [48]. For the bound values in Table 2, we draw 1000 time values uniformly in $[\epsilon, T]$ and use them to estimate $\mathcal{L}_\theta^{\text{DSM}}$ for each datapoint, with the same importance sampling technique in Eq. (12). We use the correction in Eq. (39) and report upper bounds for $-\log p_\theta(\mathbf{x})$. For computing the likelihood of p_θ^{ODE} , we use the Dormand-Prince RK45 ODE solver [14] with absolute and relevant tolerances set to 10^{-5} . We do not use the correction in Eq. (39) for p_θ^{ODE} , because it is still a valid likelihood for the data distribution even in the time horizon $[\epsilon, T]$.

Below is a related result to bound $\log \tilde{p}_\theta^{\text{SDE}}(\mathbf{x})$ directly. We include it here for completeness, though we do not use it for either training or inference in our experiments.

Theorem 7. *Let $p_{0t}(\mathbf{x}' | \mathbf{x})$ denote the transition distribution from $p_0(\mathbf{x})$ to $p_t(\mathbf{x})$ for the SDE in Eq. (1). With the same notations and conditions in Theorem 3, as well as the definitions of \tilde{p} and $\tilde{p}_\theta^{\text{SDE}}$ in Theorem 6, we have*

$$-\log \tilde{p}_\theta^{\text{SDE}}(\mathbf{x}) \leq \mathcal{L}_{\theta,\epsilon}^{\text{SM}}(\mathbf{x}) = \mathcal{L}_{\theta,\epsilon}^{\text{DSM}}(\mathbf{x}), \quad (40)$$

where $\mathcal{L}_{\theta,\epsilon}^{\text{SM}}(\mathbf{x})$ is defined as

$$-\mathbb{E}_{p_{\epsilon T}(\mathbf{x}'|\mathbf{x})} [\log \pi(\mathbf{x}')] + \frac{1}{2} \int_\epsilon^T \mathbb{E}_{p_{\epsilon t}(\mathbf{x}'|\mathbf{x})} \left[2g(t)^2 \nabla_{\mathbf{x}'} \cdot \mathbf{s}_\theta(\mathbf{x}', t) + g(t)^2 \|\mathbf{s}_\theta(\mathbf{x}', t)\|_2^2 - 2\nabla_{\mathbf{x}'} \cdot \mathbf{f}(\mathbf{x}', t) \right] dt,$$

and $\mathcal{L}_{\theta,\epsilon}^{\text{DSM}}(\mathbf{x})$ is given by

$$\begin{aligned} & -\mathbb{E}_{p_{\epsilon T}(\mathbf{x}'|\mathbf{x})} [\log \pi(\mathbf{x}')] + \frac{1}{2} \int_\epsilon^T \mathbb{E}_{p_{\epsilon t}(\mathbf{x}'|\mathbf{x})} \left[g(t)^2 \|\mathbf{s}_\theta(\mathbf{x}', t) - \nabla_{\mathbf{x}'} \log p_{\epsilon t}(\mathbf{x}' | \mathbf{x})\|_2^2 \right] dt \\ & - \frac{1}{2} \int_\epsilon^T \mathbb{E}_{p_{\epsilon t}(\mathbf{x}'|\mathbf{x})} \left[g(t)^2 \|\nabla_{\mathbf{x}'} \log p_{\epsilon t}(\mathbf{x}' | \mathbf{x})\|_2^2 + 2\nabla_{\mathbf{x}'} \cdot \mathbf{f}(\mathbf{x}', t) \right] dt. \end{aligned}$$

Proof. Proof closely parallels those of Theorems 3 and 6. □

C Experimental details

Datasets All our experiments are performed on two image datasets: CIFAR-10 [28] and down-sampled ImageNet [55]. Both contain images of resolution 32×32 . CIFAR-10 has 50000 images as the training set and 10000 images as the test set. Down-sampled ImageNet has 1281149 training images and 49999 test images. It is well-known that ImageNet contains some personal sensitive information and may cause privacy concern [57]. We minimize this risk by using the dataset with a small resolution (32×32).

Model architectures Our variational dequantization model, $q_\phi(\mathbf{u} | \mathbf{x})$, follows the same architecture of Flow++ [18]. We do not use dropout for score-based models trained on ImageNet. We did not tune model architectures or training hyper-parameters specifically for maximizing likelihoods. All likelihood values were reported using the last checkpoint of each setting.



(a) DDPM++ (deep, subVP) [48], FID = 2.86



(b) ScoreFlow, FID = 5.34

Figure 3: Samples on CIFAR-10. (a) Model with the best FID. (b) ScoreFlow trained with likelihood weighting + importance sampling + VP SDE. Samples of both models are generated with the same random seed.

Training We follow the same training procedure for score-based models in [48]. We also use the same hyperparameters for training the variational dequantization model, except that we train it for only 300000 iterations while fixing the score-based model. All models are trained on Cloud TPU v3-8 (roughly equivalent to 4 Tesla V100 GPUs). The baseline DDPM++ model requires around 33 hours to finish training, while the deep DDPM++ model requires around 44 hours. The variational dequantization model for the former requires around 7 hours to train, and for the latter it requires around 9.5 hours.

Confidence intervals All likelihood values are obtained by averaging the results on around 50000 datapoints, sampled with replacement from the test dataset. We can compute the confidence intervals with Student’s t-test. On CIFAR-10, the radius of 95% confidence intervals is typically around 0.006 bits/dim, while on ImageNet it is around 0.008 bits/dim.

Sample quality All FID values are computed on 50000 samples from p_{θ}^{ODE} , generated with numerical ODE solvers as in [48]. We compute FIDs between samples and training/test data for CIFAR-10/ImageNet. Although likelihood weighting + importance sampling slightly increases FID scores, their samples have comparable visual quality, as demonstrated in Figs. 3 and 4.



(a) DDPM++ (VP) [48], FID = 8.34



(b) ScoreFlow, FID = 10.18

Figure 4: Samples on ImageNet 32×32 . (a) Model with the best FID. (b) ScoreFlow trained with likelihood weighting + importance sampling + VP SDE. Samples of both models are generated with the same random seed.