

# Consistent Diffusion Models: Mitigating Sampling Drift by Learning to be Consistent

Giannis Daras\*

Department of Computer Science  
University of Texas at Austin

Alexandros G. Dimakis

Department of ECE  
University of Texas at Austin

Yuval Dagan\*

Electrical Engineering and Computer Science  
University of California, Berkeley

Constantinos Daskalakis

Electrical Engineering and Computer Science  
Massachusetts Institute of Technology

February 20, 2023

## Abstract

Imperfect score-matching leads to a shift between the training and the sampling distribution of diffusion models. Due to the recursive nature of the generation process, errors in previous steps yield sampling iterates that drift away from the training distribution. Yet, the standard training objective via Denoising Score Matching (DSM) is only designed to optimize over non-drifted data. To train on drifted data, we propose to enforce a *consistency* property which states that predictions of the model on its own generated data are consistent across time. Theoretically, we show that if the score is learned perfectly on some non-drifted points (via DSM) and if the consistency property is enforced everywhere, then the score is learned accurately everywhere. Empirically we show that our novel training objective yields state-of-the-art results for conditional and unconditional generation in CIFAR-10 and baseline improvements in AFHQ and FFHQ. We open-source our code and models: <https://github.com/giannisdaras/cdm>.

## 1 Introduction

The diffusion-based (Sohl-Dickstein et al., 2015; Song and Ermon, 2019; Ho et al., 2020) approach to generative models has been successful across various modalities, including images (Ramesh et al., 2022; Saharia et al., 2022; Dhariwal and Nichol, 2021; Nichol and Dhariwal, 2021; Kim et al., 2022; Song et al., 2021b; Ruiz et al., 2022; Gal et al., 2022; Daras and Dimakis, 2022; Daras et al., 2022a), videos (Ho et al., 2022a,b; Hong et al., 2022), audio (Kong et al., 2021), 3D structures (Poole et al., 2022), proteins (Anand and Achim, 2022; Trippe et al., 2022; Schneuing et al., 2022; Corso et al., 2022), and medical applications (Jalal et al., 2021; Arvinte et al., 2022).

Diffusion models generate data by first drawing a sample from a noisy distribution and slowly *denoising* this sample to ultimately obtain a sample from the target distribution. This is achieved by sampling, in reverse from time  $t = 1$  down to  $t = 0$ , a stochastic process  $\{x_t\}_{t \in [0,1]}$  wherein  $x_0$  is distributed according to the target distribution  $p_0$  and, for all  $t$ ,

$$x_t \sim p_t \quad \text{where} \quad p_t := p_0 \oplus N(0, \sigma_t^2 I_d). \quad (1)$$

That is,  $p_t$  is the distribution resulting from corrupting a sample from  $p_0$  with noise sampled from  $N(0, \sigma_t^2 I_d)$ , where  $\sigma_t$  is an increasing function such that  $\sigma_0 = 0$  and  $\sigma_1$  is sufficiently large so that  $p_1$  is nearly indistinguishable from pure noise. We note that diffusion models have been generalized to other types of

---

\*These authors contributed equally to this work

corruptions by the recent works of Daras et al. (2022b); Bansal et al. (2022); Hoogeboom and Salimans (2022); Deasy et al. (2021); Nachmani et al. (2021).

In order to sample from a diffusion model, i.e. sample the afore-described process in reverse time, it suffices to know the *score function*  $s(x, t) = \nabla_x \log p(x, t)$ , where  $p(x, t)$  is the density of  $x_t \sim p_t$ . Indeed, given a sample  $x_t \sim p_t$ , one can use the score function at  $x_t$ , i.e.  $s(x_t, t)$ , to generate a sample from  $p_{t-dt}$  by taking an infinitesimal step of a stochastic or an ordinary differential equation (Song et al., 2021b,a), or by using Langevin dynamics (Grenander and Miller, 1994; Song and Ermon, 2020).<sup>1</sup> Hence, in order to train a diffusion model to sample from a target distribution of interest  $p_0^*$  it suffices to learn the score function  $s^*(x, t)$  using samples from the corrupted distributions  $p_t^*$  resulting from  $p_0^*$  and a particular noise schedule  $\sigma_t$ . Notice that those samples can be easily drawn given samples from  $p_0^*$ .

**The Sampling Drift Challenge:** Unfortunately the true score function  $s^*(x, t)$  is not perfectly learned during training. Thus, at generation time, the samples  $x_t$  drawn using the learned score function,  $s(x, t)$ , in the ways discussed above, drift astray in distribution from the true corrupted distributions  $p_t^*$ . This drift becomes larger for smaller  $t$  due to compounding of errors and is accentuated by the fact that the further away a sample  $x_t$  is from the likely support of the true  $p_t^*$  the larger is also the error  $\|s(x_t, t) - s^*(x_t, t)\|$  between the learned and the true score function at  $x_t$ , which feeds into an even larger drift between the distribution of  $x_{t'}$  from  $p_{t'}^*$ , for  $t' < t$ ; see e.g. (Sehwag et al., 2022; Ho et al., 2020; Nichol and Dhariwal, 2021; Chen et al., 2022a). These challenges motivate the question:

*Question 1.* How can one train diffusion models to improve the error  $\|s(x, t) - s^*(x, t)\|$  between the learned and true score function on inputs  $(x, t)$  where  $x$  is unlikely under the target noisy distribution  $p_t^*$ ?

A direct approach to this challenge is to train our model to minimize the afore-described error on pairs  $(x, t)$  where  $x$  is sampled from distributions other than  $p_t^*$ . However, there is no straightforward way to do so, because we do not have direct access to the values of the true score function  $s^*(x, t)$ .

This motivates us to propose a novel training method to mitigate sampling drift by enforcing that the learned score function satisfies an invariant, that we call “consistency property.” This property relates multiple inputs to  $s(\cdot, \cdot)$  and can be optimized without using any samples from the target distribution  $p_0^*$ . As we will show theoretically, enforcing this consistency in conjunction with minimizing a very weakened form of the standard score matching objective (for a single  $t$  and an open set of  $x$ ’s) suffices to learn the correct score everywhere. We also provide experiments illustrating that regularizing the standard score matching objective using our consistency property leads to state-of-the-art models.

**Our Approach:** The true score function  $s^*(x, t)$  is closely related to another function, called *optimal denoiser*, which predicts a clean sample  $x_0 \sim p_0^*$  from a noisy observation  $x_t = x_0 + \sigma_t \eta$  where the noise is  $\eta \sim \mathcal{N}(0, I_d)$ . The optimal denoiser (under the  $\ell_2$  loss) is the conditional expectation:

$$h^*(x, t) := \mathbb{E}[x_0 | x_t = x],$$

and the true score function can be obtained from the optimal denoiser as follows:  $s^*(x, t) = (h^*(x, t) - x)/\sigma_t^2$ . Indeed, the standard training technique, via *score-matching*, explicitly trains for the score through the denoiser  $h^*$  (Vincent, 2011; Efron, 2011; Meng et al., 2021; Kim and Ye, 2021; Luo, 2022).

We are now ready to state our consistency property. We will say that a (denoising) function  $h(x, t)$  is *consistent* iff

$$\forall t, \forall x : \mathbb{E}[x_0 | x_t = x] = h(x, t),$$

where the *expectation is with respect to a sample from the learned reverse process*, defined in terms of the implied score function  $s(x, t) = (h(x, t) - x)/\sigma_t^2$ , when this is initialized at  $x_t = x$  and run backwards in time to sample  $x_0$ . See Eq. (3) for the precise stochastic differential equation and its justification. In particular,  $h$  is called consistent if the prediction  $h(x, t)$  of the conditional expectation of the clean image  $x_0$  given  $x_t = x$  equals the expected value of an image that is generated by the learned reversed process, starting from  $x_t = x$ .

<sup>1</sup>Some of these methods, such as Langevin dynamics, require also to know the score function in the neighborhood of  $x_t$ .

While there are several other properties that the score function of a diffusion process must satisfy, e.g. the Fokker-Planck equation (Lai et al., 2022), our first theoretical result is that the consistency of  $h(x, t)$  suffices (in conjunction with the conservativeness of its score function  $s(x, t) = (h(x, t) - x)/\sigma_t^2$ ) to guarantee that  $s$  must be the score function of a diffusion process (and must thus satisfy any other property that a diffusion process must satisfy). If additionally  $s(x, t)$  equals the score function  $s^*(x, t)$  of a target diffusion process at a single time  $t = t_0$  and an open subset of  $x \in \mathbb{R}^d$ , then it equals  $s^*$  everywhere. Intuitively, this suggests that learning the score in-sample for a single  $t = t_0$ , and satisfying the consistency and conservativeness properties off-sample, also yields a correct estimate off-sample. This can be summarized as follows below:

**Theorem 1.1** (informal). *If some denoiser  $h(x, t)$  is consistent and its corresponding score function  $s(x, t) = (h(x, t) - x)/\sigma_t^2$  is a conservative field, then  $s(x, t)$  is the score function of a diffusion process, i.e. the generation process using score function  $s$  is the inverse of a diffusion process. If additionally  $s(x, t) = s^*(x, t)$  for a single  $t = t_0$  and all  $x$  in an open subset of  $\mathbb{R}^d$ , where  $s^*$  is the score function of a target diffusion process, then  $s(x, t) = s^*(x, t)$  everywhere, i.e. to learn the score function everywhere it suffices to learn it for a single  $t_0$  and an open subset of  $x$ 's.*

We propose a loss function to train for the consistency property and we show experimentally that regularizing the standard score matching objective using our consistency property leads to better models.

### Summary of Contributions:

1. We identify an invariant property, consistency of the denoiser  $h$ , that any perfectly trained model should satisfy.
2. We prove that if the denoiser  $h(x, t)$  is consistent and its implied score function  $s(x, t) = (h(x, t) - x)/\sigma_t^2$  is a conservative field, then  $s(x, t)$  is the score function of *some* diffusion process, even if there are learning errors with respect to the score of the target process, which generates the training data.
3. We prove that if these two properties are satisfied, then optimizing perfectly the score for a single  $t = t_0$  and an open subset  $S \subseteq \mathbb{R}^d$ , guarantees that the score is learned perfectly everywhere.
4. We propose a novel training objective that enforces the consistency property. Our new objective optimizes the network to have consistent predictions on data points from the *learned* distribution.
5. We show experimentally that, paired with the original Denoising Score Matching (DSM) loss, our objective achieves a new state-of-the-art on conditional and unconditional generation in CIFAR10 and baseline improvements in AFHQ and FFHQ.
6. We open-source our code and models: <https://github.com/giannisdaras/cdm>.

## 2 Background

**Diffusion processes, score functions and denoising.** Diffusion models are trained by solving a supervised regression problem (Song and Ermon, 2019; Ho et al., 2020). The function that one aims to learn, called the score function, defined below, is equivalent (up to a linear transformation) to a denoising function (Efron, 2011; Vincent, 2011), whose goal is to denoise an image that was injected with noise. In particular, for some target distribution  $p_0$ , one's goal is to learn the following function  $h: \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$ :

$$h(x, t) = \mathbb{E}[x_0 \mid x_t = x]; \quad x_0 \sim p_0, \quad x_t \sim N(x_0, \sigma_t^2 I_d). \quad (2)$$

In other words, the goal is to predict the expected “clean” image  $x_0$  given a corrupted version of it, assuming that the image was sampled from  $p_0$  and its corruption was done by adding to it noise from  $N(0, \sigma_t^2 I_d)$ , where  $\sigma_t^2$  is a non-negative and increasing function of  $t$ . Given such a function  $h$ , we can generate samples from  $p_0$

by solving a Stochastic Differential Equation (SDE) that depends on  $h$  (Song et al., 2021b). Specifically, one starts by sampling  $x_1$  from some fixed distribution and then runs the following SDE backwards in time:

$$dx_t = -g(t)^2 \frac{h(x_t, t) - x_t}{\sigma_t^2} dt + g(t) d\bar{B}_t, \quad (3)$$

where  $\bar{B}_t$  is a reverse-time Brownian motion and  $g(t)^2 = \frac{d\sigma_t^2}{dt}$ . To explain how Eq. (3) was derived, consider the *forward* SDE that starts with a clean image  $x_0$  and slowly injects noise:

$$dx_t = g(t) dB_t, \quad x_0 \sim p_0. \quad (4)$$

We notice here that the  $x_t$  under Eq. (4) is  $N(x_0, \sigma_t^2 I_d)$ , where  $x_0 \sim p_0$ , so it has the same distribution that it has in Eq. (2). Remarkably, such SDEs are reversible in time (Anderson, 1982). Hence, the diffusion process of Eq. (4) can be viewed as a reversed-time diffusion:

$$dx_t = -g(t)^2 \nabla_x \log p(x_t, t) dt + g(t) d\bar{B}_t, \quad (5)$$

where  $p(x_t, t)$  is the density of  $x_t$  at time  $t$ . We note that  $s(x, t) := \nabla_x \log p(x, t)$  is called the *score function* of  $x_t$  at time  $t$ . Using Tweedie’s lemma (Efron, 2011), one obtains the following relationship between the denoising function  $h$  and the score function:

$$\nabla_x \log p(x, t) = \frac{h(x, t) - x}{\sigma_t^2}. \quad (6)$$

Substituting Eq. (6) in Eq. (5), one obtains Eq. (3).

**Training via denoising score matching.** The standard way to train for  $h$  is via *denoising score matching*. This is performed by obtaining samples of  $x_0 \sim p_0$  and  $x_t \sim N(x_0, \sigma_t^2 I_d)$  and training to minimize

$$\mathbb{E}_{x_0 \sim p_0, x_t \sim N(x_0, \sigma_t^2 I_d)} L_{t, x_t, x_0}^1(\theta) = \mathbb{E}_{x_0 \sim p_0, x_t \sim N(x_0, \sigma_t^2 I_d)} \|h_\theta(x_t, t) - x_0\|^2,$$

where the optimization is over some family of functions,  $\{h_\theta\}_{\theta \in \Theta}$ . It was shown by Vincent (2011) that optimizing Eq. (2) is equivalent to optimizing  $h$  in mean-squared-error on a random point  $x_t$  that is a noisy image,  $x_t \sim N(x_0, \sigma_t^2 I_d)$  where  $x_0 \sim p_0$ :

$$\mathbb{E}_{x_t} \|h_\theta(x_t, t) - h^*(x_t, t)\|^2,$$

where  $h^*$  is the true denoising function from Eq. (2).

### 3 Theory

We define below the consistency property that a function  $h$  should satisfy. This states that the output of  $h(x, t)$  (which is meant to approximate the conditional expectation of  $x_0$  conditioned on  $x_t = x$ ) is indeed consistent with the average point  $x_0$  generated using  $h$  and conditioning on  $x_t = x$ . Recall from the previous section that generation according to  $h$  conditioning on  $x_t = x$  is done by running the following SDE backwards in time conditioning on  $x_t = x$ :

$$dx_t = -g(t)^2 \frac{h(x_t, t) - x_t}{\sigma_t^2} dt + g(t)^2 d\bar{B}_t, \quad (7)$$

The consistency property is therefore defined as follows:

*Property 1 (Consistency).* A function  $h: \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$  is said to be *consistent* iff for all  $t \in (0, 1]$  and all  $x \in \mathbb{R}^d$ ,

$$h(x, t) = \mathbb{E}_h[x_0 \mid x_t = x], \quad (8)$$

where  $\mathbb{E}_h[x_0 \mid x_t = x]$  corresponds to the conditional expectation of  $x_0$  in the process that starts with  $x_t = x$  and samples  $x_0$  by running the SDE of Eq. (7) backwards in time (where note that the SDE uses  $h$ ).

The following Lemma states that Property 1 holds if and only if the model prediction,  $h(x, t)$ , is consistent with the average output of  $h$  on samples that are generated using  $h$  and conditioning on  $x_t = x$ , i.e. that  $h(x_t, t)$  is a reverse-Martingale under the same process of Eq. (7).

**Lemma 3.1.** *Property 1 holds if and only if the following two properties hold:*

- *The function  $h$  is a reverse-Martingale, namely: for all  $t > t'$  and for any  $x$ :*

$$h(x, t) = \mathbb{E}_h[h(x_{t'}, t') \mid x_t = x],$$

*where the expectation is over  $x_{t'}$  that is sampled according to Eq. (7) with the same function  $h$ , given the initial condition  $x_t = x$ .*

- *For all  $x \in \mathbb{R}^d$ ,  $h(x, 0) = x$ .*

The proof of this Lemma is included in the Appendix. Further, we introduce one more property that will be required for our theoretical results: the learned vector-field should be conservative.

**Property 2 (Conservative vector field / Score Property).** Let  $h: \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$ . We say that  $h$  induces a *conservative vector field* (or that it satisfies the score property) if for any  $t \in (0, 1]$  there exists some probability density  $p(\cdot, t)$  such that

$$\frac{h(x, t) - x}{\sigma_t^2} = \nabla \log p(x, t).$$

We note that the optimal denoiser, i.e.  $h$  defined as in Eq. (2) satisfies both of the properties we introduced. In the paper, we will focus on enforcing the consistency property and we are going to assume conservativeness for our theoretical results. This assumption can be relaxed to hold only at a *single*  $t \in (0, 1]$  using results of Lai et al. (2022).

Next, we show the theoretical consequences of enforcing Properties 1 and 2. First, we show that this enforces  $h$  to indeed correspond to a denoising function, namely,  $h$  satisfies Eq. (2) for some distribution  $p'_0$  over  $x_0$ . Yet, this does not imply that  $p_0$  is the *correct* underlying distribution that we are trying to learn. Indeed, these properties can apply to any distribution  $p_0$ . Yet, we can show that if we learn  $h$  correctly for some inputs and if these properties apply everywhere then  $h$  is learned correctly everywhere.

**Theorem 3.2.** *Let  $h: \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$  be a continuous function. Then:*

1. *The function  $h$  satisfies both Properties 1 and 2 if and only if  $h$  is defined by Eq. (2) for some distribution  $p_0$ .*
2. *Assume that  $h$  satisfies Properties 1 and 2. Further, let  $h^*$  be another function that corresponds to Eq. (2) with some initial distribution  $p_0^*$ . Assume that  $h = h^*$  on some open set  $U \subseteq \mathbb{R}^d$  and some fixed  $t_0 \in (0, 1]$ , namely,  $h(x, t_0) = h^*(x, t_0)$  for all  $x \in U$ . Then,  $h^*(x, t) = h(x, t)$  for all  $x$  and all  $t$ .*

*Proof overview.* We start with the first part of the theorem. We assume that  $h$  satisfies Properties 1 and 2 and we will show that  $h$  is defined by Eq. (2) for some distribution  $p_0$  (while the other direction in the equivalence follows trivially from the definitions of these properties). Motivated by Eq. (6), define the function  $s: \mathbb{R}^d \times (0, 1]$  according to

$$s(x, t) = \frac{h(x, t) - x}{\sigma_t^2}. \tag{9}$$

We will first show that  $s$  satisfies the partial differential equation

$$\frac{\partial s}{\partial t} = g(t)^2 \left( J_s s + \frac{1}{2} \Delta s \right), \tag{10}$$

where  $J_s \in \mathbb{R}^{d \times d}$  is the Jacobian of  $s$ ,  $(J_s)_{ij} = \frac{\partial s_i}{\partial x_j}$  and each coordinate  $i$  of  $\Delta s \in \mathbb{R}^d$  is the Laplacian of coordinate  $i$  of  $s$ ,  $(\Delta s)_i = \sum_{j=1}^n \frac{\partial^2 s_i}{\partial x_j^2}$ . In order to obtain Eq. (10), first, we use a generalization of Ito's lemma, which states that for an SDE

$$dx_t = \mu(x_t, t)dt + g(t)d\bar{B}_t \quad (11)$$

and for  $f: \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$ ,  $f(x_t, t)$  satisfies the SDE

$$df(x_t, t) = \left( \frac{\partial f}{\partial t} + J_f \mu - \frac{g(t)^2}{2} \Delta f \right) dt + g(t) J_f d\bar{B}_t.$$

If  $f$  is a reverse-Martingale then the term that multiplies  $dt$  has to equal zero, namely,

$$\frac{\partial f}{\partial t} + J_f \mu - \frac{g(t)^2}{2} \Delta f = 0.$$

By Lemma 3.1,  $h(x_t, t)$  is a reverse Martingale, therefore we can substitute  $f = h$  and substitute  $\mu = -g(t)^2 s$  according to Eq. (7), to deduce that

$$\frac{\partial h}{\partial t} - g(t)^2 J_h s - \frac{g(t)^2}{2} \Delta h = 0.$$

Substituting  $h(x, t) = \sigma_t^2 s(x, t) + x$  according to Eq. (6) yields Eq. (10) as required.

Next, we show that any  $s'$  that is the score-function (i.e. gradient of log probability) of some diffusion process that follows the SDE Eq. (4), also satisfies Eq. (10). To obtain this, one can use the Fokker-Planck equation, whose special case states that the density function  $p(x, t)$  of any stochastic process that satisfies the SDE Eq. (4) satisfies the PDE

$$\frac{\partial p}{\partial t} = \frac{g(t)^2}{2} \Delta p$$

where  $\Delta$  corresponds to the Laplacian operator. Using this one can obtain a PDE for  $\nabla_x \log p$  which happens to be exactly Eq. (10) if the process is defined by Eq. (4).

Next, we use Property 2 to deduce that there exists some densities  $p(\cdot, t)$  for  $t \in [0, 1]$  such that

$$s(x, t) = \frac{h(x, t) - x}{\sigma_t^2} = \nabla_x \log p(x, t).$$

Denote by  $p'(x, t)$  the score function of the diffusion process that is defined by the SDE of Eq. (4) with the initial condition that  $p(x, 0) = p'(x, 0)$  for all  $x$ . Denote by  $s'(x, t) = \nabla_x \log p'(x, t)$  the score function of  $p'$ . As we proved above, both  $s$  and  $s'$  satisfy the PDE Eq. (10) and the same initial condition at  $t = 0$ . By the uniqueness of the PDE, it holds that  $s(x, t) = s'(x, t)$  for all  $t \geq t_0$ . Denote by  $h^*$  the function that satisfies Eq. (2) with the initial condition  $x_0 \sim p_0$ . By Eq. (6),

$$s'(x, t) = \frac{h^*(x, t) - x}{\sigma_t^2}.$$

By Eq. (9) and since  $s = s'$ , it follows that  $h = h^*$  and this is what we wanted to prove.

We proceed with proving part 2 of the theorem. We use the notion of an *analytic function* on  $\mathbb{R}^d$ : that is a function  $f: \mathbb{R}^d \rightarrow \mathbb{R}$  such that at any  $x_0 \in \mathbb{R}^d$ , the Taylor series of  $f$  centered at  $x_0$  converges for all  $x \in \mathbb{R}^d$  to  $f(x)$ . We use the property that an analytic function is uniquely determined by its value on any open subset: *If  $f$  and  $g$  are analytic functions that identify in some open subset  $U \subset \mathbb{R}^d$  then  $f = g$  everywhere.* We prove this statement in the remainder of this paragraph, as follows: Represent  $f$  and  $g$  as Taylor series around some  $x_0 \in U$ . The Taylor series of  $f$  and  $g$  identify: indeed, these series are functions of the derivatives of  $f$  and  $g$  which are functions of only the values in  $U$ . Since  $f$  and  $g$  equal their Taylor series, they are equal.

Next, we will show that for any diffusion process that is defined by Eq. (4), the probability density of  $p(x, t_0)$  at any time  $t_0 > 0$  is analytic as a function of  $x$ . Recall that the distribution of  $x_0$  is defined in Eq. (4) as  $p_0$  and it holds that the distribution of  $x_{t_0}$  is obtained from  $p_0$  by adding a Gaussian noise  $N(0, \sigma_t^2 I)$  and its density at any  $x$  equals

$$p(x, t_0) = \int_{a \in \mathbb{R}^d} \frac{1}{\sqrt{2\pi}\sigma_{t_0}} \exp\left(-\frac{(x-a)^2}{2\sigma_{t_0}^2}\right) dp_0(a).$$

Since the function  $\exp(-(x-a)^2/(2\sigma_t^2))$  is analytic, one could deduce that  $p(x, t_0)$  is also analytic. Further,  $p(x, t_0) > 0$  for all  $x$  which implies that there is no singularity for  $\log p(x, t_0)$  which can be used to deduce that  $\log p(x, t_0)$  is also analytic and further that  $\nabla_x \log p(x, t_0)$  is analytic as well.

We use the first part of the theorem to deduce that  $s$  is the score function of some diffusion process hence it is analytic. By assumption,  $s$  identifies with some target score function  $s^*$  in some open subset  $U \subseteq \mathbb{R}^d$  at some  $t_0$ , which, by the fact that  $s(x, t_0)$  and  $s^*(x, t_0)$  are analytic, implies that  $s(x, t_0) = s^*(x, t_0)$  for all  $x$ . Finally, since  $s$  and  $s^*$  both satisfy the PDE Eq. (10) and they satisfy the same initial condition at  $t_0$ , it holds that by uniqueness of the PDE  $s(x, t) = s^*(x, t)$  for all  $x$  and  $t$ .  $\square$

## 4 Method

Theorem 3.2 motivates enforcing the consistency property on the learned model. We notice that the consistency equation Eq. (8) may be expensive to train for, because it requires one to generate whole trajectories. Rather, we use the equivalent Martingale assumption of Lemma 3.1, which can be observed locally with only partial trajectories:<sup>2</sup> We suggest the following loss function, for some fixed  $t, t'$  and  $x$ :

$$L_{t,t',x}^2(\theta) = (\mathbb{E}_\theta[h_\theta(x_{t'}, t') \mid x_t = x] - h_\theta(x, t))^2 / 2,$$

where the expectation  $\mathbb{E}_\theta[\cdot \mid x_t = x]$  is taken according to process Eq. (7) parameterized by  $h_\theta$  with the initial condition  $x_t = x$ . Differentiating this expectation, one gets the following (see Section B.1 for full derivation):

$$\begin{aligned} \nabla L_{t,t',x}^2(\theta) = \mathbb{E}_\theta [h_\theta(x_{t'}, t') - h_\theta(x_t, t) \mid x_t = x]^\top \mathbb{E}_\theta \left[ h_\theta(x_{t'}, t') \nabla_\theta \log(p_\theta(x_{t'} \mid x_t = x)) + \right. \\ \left. \nabla_\theta h_\theta(x_{t'}, t') - \nabla_\theta h_\theta(x_t, t) \mid x_t = x \right], \end{aligned}$$

where  $p_\theta$  corresponds to the same probability measure where the expectation  $\mathbb{E}_\theta$  is taken from and  $\nabla_\theta h_\theta$  corresponds to the Jacobian matrix of  $h_\theta$  where the derivatives are taken with respect to  $\theta$ . Notice, however, that computing the expectation accurately might require a large number of samples. Instead, it is possible to obtain a stochastic gradient of this target by taking two samples,  $x_{t'}$  and  $x_{t'}$ , independently, from the conditional distribution of  $x_{t'}$  conditioned on  $x_t = x$  and replace each of the two expectations in the formula above with one of these two samples.

We further notice the gradient of the consistency loss can be written as

$$\nabla_\theta L_{t,t',x}^2(\theta) = \frac{1}{2} \nabla_\theta \|\mathbb{E}_\theta[h_\theta(x_{t'}, t')] - h_\theta(x, t)\|^2 + \mathbb{E}_\theta [h_\theta(x_{t'}, t') - h_\theta(x, t)]^\top \mathbb{E}_\theta [\nabla_\theta \log(p(x_{t'})) h_\theta(x_{t'}, t')]$$

In order to save on computation time, we trained by taking gradient steps with respect to only the first summand in this decomposition and notice that if the consistency property is preserved then this term becomes zero, which implies that no update is made, as desired.

It remains to determine how to select  $t, t'$  and  $x_{t'}$ . Notice that  $t$  has to vary throughout the whole range of  $[0, 1]$  whereas  $t'$  can either vary over  $[0, t]$ , however, it sufficient to take  $t' \in [t - \epsilon, t]$ . However, the further

<sup>2</sup>According to Lemma 3.1, in order to completely train for Property 1, one has to also enforce  $h(x, 0) = x$ , however, this is taken care from the denoising score matching objective Eq. (2).

away  $t$  and  $t'$  are, we need to run more steps of the reverse SDE to avoid large discretization errors. Instead, we enforce the property only on small time windows using that consistency over small intervals implies global consistency. We notice that  $x_t$  can be chosen arbitrarily and two possible choices are to sample it from the target noisy distribution  $p_t$  or from the model.

*Remark 4.1.* It is important to sample  $x_{t'}$  conditioned on  $x_t$  according to the specific SDE Eq. (7). While a variety of alternative SDEs exist which preserve the same marginal distribution at any  $t$ , they might not preserve the conditionals.

## 5 Experiments

For all our experiments, we rely on the official open-sourced code and the training and evaluation hyperparameters from the paper “*Elucidating the Design Space of Diffusion-Based Generative Models*” (Karras et al., 2022) that, to the best of our knowledge, holds the current state-of-the-art on conditional generation on CIFAR-10 and unconditional generation on CIFAR-10, AFHQ (64x64 resolution), FFHQ (64x64 resolution). We refer to the models trained with our regularization as “CDM (Ours)” and to models trained with vanilla Denoising Score Matching (DSM) as “EDM” models. “CDM” models are trained with the weighted objective:

$$L_{\lambda}^{\text{ours}}(\theta) = \mathbb{E}_t \left[ \mathbb{E}_{x_0 \sim p_0, x_t \sim \mathcal{N}(x_0, \sigma_t^2 I_d)} L_{t, x_t, x_0}^1(\theta) + \lambda \mathbb{E}_{x_t \sim p_t} \mathbb{E}_{t' \sim \mathcal{U}[t-\epsilon, t]} L_{t, t', x_t}^2(\theta) \right],$$

while the “EDM” models are trained only with the first term of the outer expectation. We also denote in the name whether the models have been trained with the Variance Preserving (VP) Song et al. (2021b); Ho et al. (2020) or the Variance Exploding Song et al. (2021b); Song and Ermon (2020, 2019), e.g. we write EDM-VP. Finally, for completeness, we also report scores from the models of Song et al. (2021b), following the practice of the EDM paper. We refer to the latter baselines as “NCSNv3” baselines.

We train diffusion models, with and without our regularization, for conditional generation on CIFAR-10 and unconditional generation on CIFAR-10 and AFHQ (64x64 resolution). For the re-trained models on CIFAR-10, we use exactly the same training hyperparameters as in Karras et al. (2022) and we verify that our re-trained models match (within 1%) the FID numbers mentioned in the paper. For AFHQ, we had to drop the batch size from the suggested value of 512 to 256 to fit in memory, which increased the FID from 1.96 (reported value) to 2.29. All models were trained for 200k iterations, as in Karras et al. (2022). Finally, we retrain a baseline model on FFHQ for 150k iterations and we finetune it for 5k steps using our proposed objective.

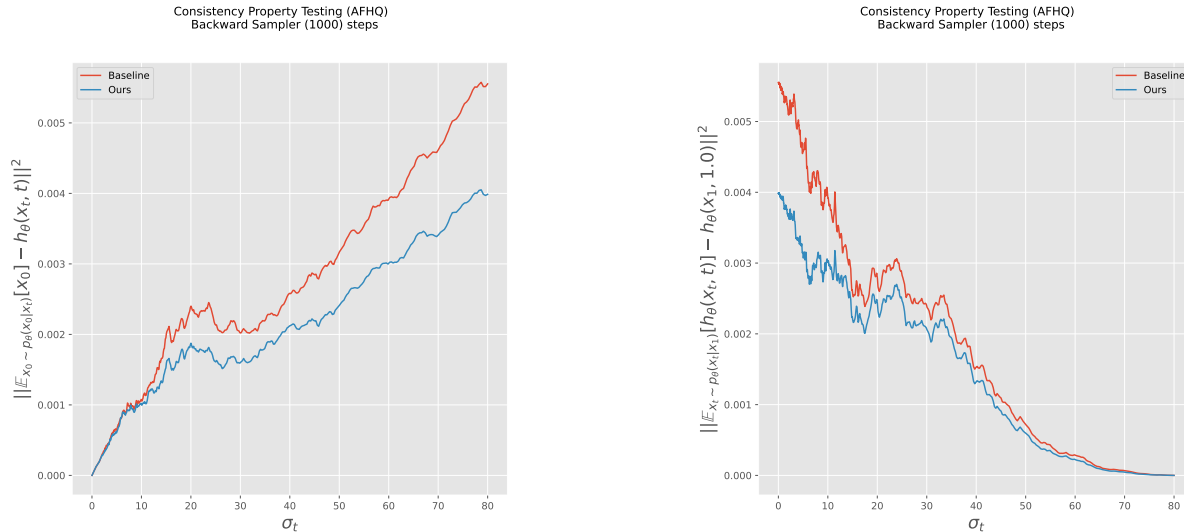
**Implementation Choices and Computational Requirements.** As mentioned, when enforcing the Consistency Property, we are free to choose  $t'$  anywhere in the interval  $[0, t]$ . When  $t, t'$  are far away, sampling  $x'_t$  from the distribution  $p_{t'}^{\theta}(x'_t | x_t)$  requires many sampling steps (to reduce discretization errors). Since this needs to be done for every Gradient Descent update, the training time increases significantly. Instead, we notice that local consistency implies global consistency. Hence, we first fix the number of sampling steps to run in every training iteration and then we sample  $t'$  uniformly in the interval  $[t - \epsilon, t]$  for some specified  $\epsilon$ . For all our experiments, we fix the number of sampling steps to 6 which roughly increases the training time needed by 1.5x. We train all our models on a DGX server with 8 A100 GPUs with 80GBs of memory each.

### 5.1 Consistency Property Testing

We are now ready to present our results. The first thing that we check is whether regularizing for the Consistency Property actually leads to models that are more consistent. Specifically, we want to check that the model trained with  $L_{\lambda}^{\text{ours}}$  achieves lower consistency error, i.e. lower  $L_{t, t', x_t}^2$ . To check this, we do the following two tests: i) we fix  $t = 1$  and we show how  $L_{t, t', x_t}^2$  changes as  $t'$  changes in  $[0, 1]$ , ii) we fix  $t' = 0$  and we show how the loss is changing as you change  $t$  in  $[0, 1]$ . Intuitively, the first test shows how the violation of the consistency property splits across the sampling process and the second test shows how much



you finally ( $t' = 0$ ) violate the property if the violation started at time  $t$ . The results are shown in Figures 1a, 1b, respectively, for the models trained on AFHQ. We include additional results for CIFAR-10, FFHQ in Figures 4, 5, 6, 7 of the Appendix. As shown, indeed regularizing for the Consistency Loss drops the  $L_{t,t',x_t}^2$  as expected.



(a) Consistency Property Testing on AFHQ. The plot illustrates how the Consistency Loss,  $L_{t,t',x_t}^2$ , behaves for  $t' = 0$ , as  $t$  changes.

(b) Consistency Property Testing on AFHQ. The plot illustrates how the Consistency Loss,  $L_{t,t',x_t}^2$ , behaves for  $t = 0$ , as  $t'$  changes.

Figure 1: Consistency Property Testing on AFHQ.


**Performance.** We evaluate performance of the models trained from scratch. Following the methodology of Karras et al. (2022), we generate 150k images from each model and we report the minimum FID computed on three sets of 50k images each. We keep checkpoints during training and we report FID for 30k, 70k, 100k, 150k, 180k and 200k iterations in Table 1. We also report the best FID found for each model, after evaluating checkpoints every 5k iterations (i.e. we evaluate 40 models spanning 200k steps of training). As shown in the Table, the proposed consistency regularization yields improvements throughout the training. In the case of CIFAR-10 (conditional and unconditional) where the re-trained baseline was trained with exactly the same hyperparameters as the models in the EDM Karras et al. (2022) paper, our CDM models achieve a new state-of-the-art.

We further show that our consistency regularization can be applied on top of a pre-trained model. Specifically, we train a baseline EDM-VP model on FFHQ  $64 \times 64$  for 150k using vanilla Denoising Score Matching. We then do 5k steps of finetuning, with and without our consistency regularization and we measure the FID score of both models. The baseline model achieves FID 2.68 while the model finetuned with consistency regularization achieves 2.61. This experiment shows the potential of applying our consistency regularization to pre-trained models, potentially even at large scale, e.g. we could apply this idea with text-to-image models such as Stable Diffusion Rombach et al. (2022). We leave this direction for future work.

Uncurated samples from our best models on AFHQ, CIFAR-10 and FFHQ are given in Figures 2a, 2b, 8. One benefit of the deterministic samplers is the unique identifiability property (Song et al., 2021b). Intuitively, this means that by using the same noise and the same deterministic sampler, we can directly compare visually models that might have been trained in completely different ways. We select a couple of images from Figure 2a (AFHQ generations) and we compare the generated images from our model with the ones from the EDM baseline for the same noises. The results are shown in Figure 3. As shown, the consistency regularization fixes several geometric inconsistencies for the picked images. We underline that the shown images are examples for



(a) Uncurated images by our model trained on AFHQ. FID: 2.21, NFEs: 79.



(b) Uncurated images by our conditional CIFAR-10 model. FID: 1.77, NFEs: 35.

Figure 2: Comparison of uncurated images generated by two different models.

Model		30k	70k	100k	150k	180k	200k	Best
<b>CDM-VP (Ours)</b>	AFHQ	<b>3.00</b>	2.44	<b>2.30</b>	<b>2.31</b>	<b>2.25</b>	<b>2.44</b>	2.21
EDM-VP (retrained)		3.27	<b>2.41</b>	2.61	2.43	2.29	2.61	2.26
EDM-VP (reported)* <sup>3</sup>								<b>1.96</b>
EDM-VE (reported)*								2.16
NCSNv3-VP (reported)*								2.58
NCSNv3-VE (reported)*								18.52
<b>CDM-VP (Ours)</b>	CIFAR10 (cond.)	<b>2.44</b>	<b>1.94</b>	<b>1.88</b>	1.88	<b>1.80</b>	<b>1.82</b>	<b>1.77</b>
EDM-VP (retrained)		2.50	1.99	1.94	<b>1.85</b>	1.86	1.90	1.82
EDM-VP (reported)								1.79
EDM-VE (reported)								1.79
NCSNv3-VP (reported)								2.48
NCSNv3-VE (reported)								3.11
<b>CDM-VP (Ours)</b>	CIFAR10 (uncond.)	<b>2.83</b>	<b>2.21</b>	<b>2.14</b>	<b>2.08</b>	<b>1.99</b>	<b>2.03</b>	<b>1.95</b>
EDM-VP (retrained)		2.90	2.32	2.15	2.09	2.01	2.13	2.01
EDM-VP (reported)								1.97
EDM-VE (reported)								1.98
NCSNv3-VP (reported)								3.01
NCSNv3-VE (reported)								3.77

Table 1: FID results for deterministic sampling, using the Karras et al. (2022) second-order samplers. For the CIFAR-10 models, we do 35 function evaluations and for AFHQ 79.

which consistency regularization helped and that potentially there are images for which the baseline models give more realistic results.

**Ablation Study for Theoretical Predictions.** One interesting implication of Theorem 3.2 is that it suggests that we only need to learn the score perfectly on some fixed  $t_0$  and then the consistency property implies that the score is learned everywhere (for all  $t$  and in the whole space). This motivates the following



Figure 3: Visual comparison of EDM model (top) and CDM model (Ours, bottom) using deterministic sampling initiated with the same noise. As seen, the consistency regularization fixes several geometric inconsistencies and artifacts in the generated images.

Model	FID
EDM (baseline)	5.81
CDM, all times $t$	5.45
CDM, for some $t$	6.59
CDM, for some $t$ , early stopped sampling	14.52

Table 2: Ablation study on removing the DSM loss for some  $t$ . Table reports FID results after 10k steps of training in CIFAR-10.

experiment: instead of using as our loss the weighted sum of DSM and our consistency regularization for all  $t$ , we will not use DSM for  $t \leq t_{\text{threshold}}$ , for some  $t_{\text{threshold}}$  that we test our theory for.

We pick  $t_{\text{threshold}}$  such that for 20% of the diffusion (on the side of clean images), we do not train with DSM. For the rest 80% we train with both DSM and our consistency regularization. Since this is only an ablation study, we train for only 10k steps on (conditional) CIFAR-10. We report FID numbers for three models: i) training with only DSM, ii) training with DSM and consistency regularization everywhere, iii) training with DSM for 80% of times  $t$  and consistency regularization everywhere. In our reported models, we also include FID of an early stopped sampling of the latter model, i.e. we do not run the sampling for  $t < t_{\text{threshold}}$  and we just output  $h_{\theta}(x_{t_{\text{threshold}}}, t_{\text{threshold}})$ . The numbers are summarized in Table 2. As shown, the theory is predictive since early stopping the generation at time  $t$  gives significantly worse results than continuing the sampling through the times that were never explicitly trained for approximating the score (i.e. we did not use DSM for those times). That said, the best results are obtained by combining DSM and our consistency regularization everywhere, which is what we did for all the other experiments in the paper.

## 6 Related Work

The fact that imperfect learning of the score function introduces a shift between the training and the sampling distribution has been well known. Chen et al. (2022a) analyze how the  $l_2$  error in the approximation of the score function propagates to Total Variation distance error bounds between the true and the learned distribution. Several methods for mitigating this issue have been proposed, but the majority of the attempts focus on changing the sampling process Song et al. (2021b); Karras et al. (2022); Jolicoeur-Martineau et al. (2021); Sehwag et al. (2022). A related work is the Analog-Bits paper Chen et al. (2022b) that conditions the model during training with past model predictions.

Karras et al. (2022) discusses potential violations of invariances, such as the non-conservativity of the induced vector field, due to imperfect score matching. However, they do not formally test or enforce this

property. Lai et al. (2022) study the problem of regularizing diffusion models to satisfy the Fokker-Planck equation. While we show in Theorem 3.2 that perfect conservative training enforces the Fokker-Planck equation, we notice that their training method is different: they suggest to enforce the equation locally by using the finite differences method to approximate the derivatives. Further, they do not train on drifted data. Instead, we notice that our consistency loss is well suited to handle drifted data since it operates across trajectories generated by the model. Finally, they show benchmark improvements on MNIST whereas we achieve state-of-the-art performance and benchmark improvements in more challenging datasets such as CIFAR-10 and AFHQ.

## 7 Conclusions and Future Work

We proposed a novel objective that enforces the trained network to have self-consistent predictions over time. We optimize this objective with points from the sampling distribution, effectively reducing the sampling drift observed in prior empirical works. Theoretically, we show that the consistency property implies that we are sampling from the reverse of some diffusion process. Together with the assumption that the network has learned perfectly the score for some time  $t_0$  and some open set  $U$ , we can prove that the consistency property implies that we learn the score perfectly everywhere. Empirically, we use our objective to obtain state-of-the-art for CIFAR-10 and baseline improvements on AFHQ and FFHQ.

There are limitations of our method and several directions for future work. The proposed regularization increases the training time by approximately 1.5x. It would be interesting to explore how to enforce consistency in more effective ways in future work. Further, our method does not test nor enforce that the induced vector-field is conservative, which is a key theoretical assumption. Our method guarantees only indirectly improve the performance in the samples from the learned distribution by enforcing some invariant. Finally, our theoretical result assumes perfect learning of the score in some subset of  $\mathbb{R}^d$ . An important next step would be to understand how errors propagate if the score-function is only approximately learned.

## 8 Acknowledgments

This research has been supported by NSF Grants CCF 1763702, AF 1901292, CNS 2148141, Tripods CCF 1934932, IFML CCF 2019844, the Texas Advanced Computing Center (TACC) and research gifts by Western Digital, WNCG IAP, UT Austin Machine Learning Lab (MLL), Cisco and the Archie Straiton Endowed Faculty Fellowship. Giannis Daras has been supported by the Onassis Fellowship, the Bodossaki Fellowship and the Leventis Fellowship. Constantinos Daskalakis has been supported by NSF Awards CCF-1901292, DMS-2022448 and DMS2134108, a Simons Investigator Award, the Simons Collaboration on the Theory of Algorithmic Fairness and a DSTA grant.

## References

- Anand, N. and Achim, T. (2022). Protein structure and sequence generation with equivariant denoising diffusion probabilistic models. *arXiv preprint arXiv:2205.15019*.
- Anderson, B. D. (1982). Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326.
- Arvinte, M., Jalal, A., Daras, G., Price, E., Dimakis, A., and Tamir, J. I. (2022). Single-shot adaptation using score-based models for mri reconstruction. In *International Society for Magnetic Resonance in Medicine, Annual Meeting*.
- Bansal, A., Borgnia, E., Chu, H.-M., Li, J. S., Kazemi, H., Huang, F., Goldblum, M., Geiping, J., and Goldstein, T. (2022). Cold Diffusion: Inverting arbitrary image transforms without noise. *arXiv preprint arXiv:2208.09392*.

- Chen, S., Chewi, S., Li, J., Li, Y., Salim, A., and Zhang, A. R. (2022a). Sampling is as easy as learning the score: theory for diffusion models with minimal data assumptions. *arXiv preprint arXiv:2209.11215*.
- Chen, T., Zhang, R., and Hinton, G. (2022b). Analog bits: Generating discrete data using diffusion models with self-conditioning. *arXiv preprint arXiv:2208.04202*.
- Corso, G., Stärk, H., Jing, B., Barzilay, R., and Jaakkola, T. (2022). Diffdock: Diffusion steps, twists, and turns for molecular docking. *arXiv preprint arXiv:2210.01776*.
- Daras, G., Dagan, Y., Dimakis, A. G., and Daskalakis, C. (2022a). Score-guided intermediate layer optimization: Fast langevin mixing for inverse problem. *arXiv preprint arXiv:2206.09104*.
- Daras, G., Delbraccio, M., Talebi, H., Dimakis, A. G., and Milanfar, P. (2022b). Soft diffusion: Score matching for general corruptions. *arXiv preprint arXiv:2209.05442*.
- Daras, G. and Dimakis, A. G. (2022). Multiresolution textual inversion. *arXiv preprint arXiv:2211.17115*.
- Deasy, J., Simidjievski, N., and Liò, P. (2021). Heavy-tailed denoising score matching. *arXiv preprint arXiv:2112.09788*.
- Dhariwal, P. and Nichol, A. (2021). Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794.
- Efron, B. (2011). Tweedie’s formula and selection bias. *Journal of the American Statistical Association*, 106(496):1602–1614.
- Gal, R., Alaluf, Y., Atzmon, Y., Patashnik, O., Bermano, A. H., Chechik, G., and Cohen-Or, D. (2022). An image is worth one word: Personalizing text-to-image generation using textual inversion.
- Grenander, U. and Miller, M. I. (1994). Representations of knowledge in complex systems. *Journal of the Royal Statistical Society: Series B (Methodological)*, 56(4):549–581.
- Ho, J., Chan, W., Saharia, C., Whang, J., Gao, R., Gritsenko, A., Kingma, D. P., Poole, B., Norouzi, M., Fleet, D. J., et al. (2022a). Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*.
- Ho, J., Jain, A., and Abbeel, P. (2020). Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851.
- Ho, J., Salimans, T., Gritsenko, A., Chan, W., Norouzi, M., and Fleet, D. J. (2022b). Video diffusion models. *arXiv:2204.03458*.
- Hong, W., Ding, M., Zheng, W., Liu, X., and Tang, J. (2022). Cogvideo: Large-scale pretraining for text-to-video generation via transformers. *arXiv preprint arXiv:2205.15868*.
- Hoogeboom, E. and Salimans, T. (2022). Blurring diffusion models. *arXiv preprint arXiv:2209.05557*.
- Jalal, A., Arvinte, M., Daras, G., Price, E., Dimakis, A. G., and Tamir, J. (2021). Robust compressed sensing mri with deep generative priors. *Advances in Neural Information Processing Systems*, 34:14938–14954.
- Jolicœur-Martineau, A., Li, K., Piché-Taillefer, R., Kachman, T., and Mitliagkas, I. (2021). Gotta go fast when generating data with score-based models. *arXiv preprint arXiv:2105.14080*.
- Karras, T., Aittala, M., Aila, T., and Laine, S. (2022). Elucidating the design space of diffusion-based generative models. *arXiv preprint arXiv:2206.00364*.
- Kim, D., Shin, S., Song, K., Kang, W., and Moon, I.-C. (2022). Soft truncation: A universal training technique of score-based diffusion model for high precision score estimation. In *International Conference on Machine Learning*, pages 11201–11228. PMLR.

- Kim, K. and Ye, J. C. (2021). Noise2score: tweedie’s approach to self-supervised image denoising without clean images. *Advances in Neural Information Processing Systems*, 34:864–874.
- Kong, Z., Ping, W., Huang, J., Zhao, K., and Catanzaro, B. (2021). Diffwave: A versatile diffusion model for audio synthesis. In *International Conference on Learning Representations*.
- Lai, C.-H., Takida, Y., Murata, N., Uesaka, T., Mitsufuji, Y., and Ermon, S. (2022). Regularizing score-based models with score fokker-planck equations. In *NeurIPS 2022 Workshop on Score-Based Methods*.
- Luo, C. (2022). Understanding diffusion models: A unified perspective. *arXiv preprint arXiv:2208.11970*.
- Meng, C., Song, Y., Li, W., and Ermon, S. (2021). Estimating high order gradients of the data distribution by denoising. *Advances in Neural Information Processing Systems*, 34:25359–25369.
- Nachmani, E., Roman, R. S., and Wolf, L. (2021). Denoising diffusion gamma models. *arXiv preprint arXiv:2110.05948*.
- Nichol, A. Q. and Dhariwal, P. (2021). Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pages 8162–8171. PMLR.
- Poole, B., Jain, A., Barron, J. T., and Mildenhall, B. (2022). Dreamfusion: Text-to-3d using 2d diffusion. *arXiv*.
- Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., and Chen, M. (2022). Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Ommer, B. (2022). High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695.
- Ruiz, N., Li, Y., Jampani, V., Pritch, Y., Rubinstein, M., and Aberman, K. (2022). Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation.
- Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E., Ghasemipour, S. K. S., Ayan, B. K., Mahdavi, S. S., Lopes, R. G., et al. (2022). Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*.
- Schneuing, A., Du, Y., Harris, C., Jamasb, A., Igashov, I., Du, W., Blundell, T., Lió, P., Gomes, C., Welling, M., et al. (2022). Structure-based drug design with equivariant diffusion models. *arXiv preprint arXiv:2210.13695*.
- Sehwag, V., Hazirbas, C., Gordo, A., Ozgenel, F., and Canton, C. (2022). Generating high fidelity data from low-density regions using diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11492–11501.
- Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., and Ganguli, S. (2015). Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR.
- Song, J., Meng, C., and Ermon, S. (2021a). Denoising diffusion implicit models. In *International Conference on Learning Representations*.
- Song, Y. and Ermon, S. (2019). Generative modeling by estimating gradients of the data distribution. *Advances in Neural Information Processing Systems*, 32.
- Song, Y. and Ermon, S. (2020). Improved techniques for training score-based generative models. *Advances in neural information processing systems*, 33:12438–12448.

- Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. (2021b). Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*.
- Trippe, B. L., Yim, J., Tischer, D., Broderick, T., Baker, D., Barzilay, R., and Jaakkola, T. (2022). Diffusion probabilistic modeling of protein backbones in 3d for the motif-scaffolding problem. *arXiv preprint arXiv:2206.04119*.
- Vincent, P. (2011). A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674.

## A Proof of Theorem 3.2

In Section A.1 we present some preliminaries to the proof, in Section A.2 we include the proof, with proofs of some lemmas omitted and in the remaining sections we prove these lemmas.

### A.1 Preliminaries

**Preliminaries on diffusion processes** In the next definition we define for a function  $F: \mathbb{R}^d \rightarrow \mathbb{R}^d$  its Jacobian  $J_F$ , its divergence  $\nabla \cdot F$  and its Laplacian  $\Delta F$  that is computed separately on each coordinate of  $F$ :

**Definition A.1.** Given a function  $F = (f_1, \dots, f_n): \mathbb{R}^d \rightarrow \mathbb{R}^d$ , denote by  $J_F: \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$  its Jacobian:

$$(J_F)_{ij} = \frac{\partial f_i(x)}{\partial x_j}.$$

The *divergence* of  $F$  is defined as

$$\nabla \cdot F(x) := \sum_{i=1}^n \frac{\partial f_i(x)}{\partial x_i}.$$

Denote by  $\Delta F: \mathbb{R}^d \rightarrow \mathbb{R}^d$  the function whose  $i$ th entry is the Laplacian of  $f_i$ :

$$(\Delta F(x))_i = \sum_{j=1}^n \frac{\partial^2 f_i(x)}{\partial x_j^2}.$$

If  $F$  is a function of both  $x \in \mathbb{R}^d$  and  $t \in \mathbb{R}$ , then  $J_F$ ,  $\Delta f$  and  $\nabla \cdot F$  correspond to  $F$  as a function of  $x$ , whereas  $t$  is kept fixed. In particular,

$$(J_F(x, t))_{ij} = \frac{\partial f_i(x, t)}{\partial x_j}, \quad (\Delta F(x, t))_i = \sum_{j=1}^n \frac{\partial^2 f_i(x, t)}{\partial x_j^2}, \quad \nabla \cdot F = \sum_{i=1}^n \frac{\partial f_i(x, t)}{\partial x_i}.$$

We use the celebrated Ito's lemma and some of its immediate generalizations:

**Lemma A.2** (Ito's Lemma). *Let  $x_t$  be a stochastic process  $x_t \in \mathbb{R}^d$ , that is defined by the following SDE:*

$$dx_t = \mu(x_t, t)dt + g(t)dB_t,$$

where  $B_t$  is a standard Brownian motion. Let  $f: \mathbb{R}^d \times \mathbb{R} \rightarrow \mathbb{R}$ . Then,

$$df(x_t, t) = \left( \frac{df}{dt} + \nabla_x f^\top \mu(x_t, t) + \frac{g(t)^2}{2} \Delta f \right) dt + g(t) \nabla_x f^\top dB_t.$$

Further, if  $F: \mathbb{R}^d \times \mathbb{R} \rightarrow \mathbb{R}^d$  is a multi-valued function, then

$$dF(x_t, t) = \left( \frac{dF}{dt} + J_F \mu + \frac{g(t)^2}{2} \Delta F \right) dt + g(t) J_F dB_t.$$

Lastly, if  $x_t$  is instead defined with a reverse noise,

$$dx_t = \mu(x_t, t)dt + g(t)d\bar{B}_t,$$

then the multi-valued Ito's lemma is modified as follows:

$$dF(x_t, t) = \left( \frac{dF}{dt} + J_F \mu - \frac{g(t)^2}{2} \Delta F \right) dt + g(t) J_F d\bar{B}_t. \quad (12)$$



Lastly, we present the Fokker-Planck equation which states that the probability distribution that corresponds to diffusion processes satisfy a certain partial differential equation:

**Lemma A.3** (Fokker-Planck equation). *Let  $x_t$  be defined by*

$$dx_t = \mu(x_t, t)dt + g(t)dB_t,$$

where  $x_t, \mu(x, t) \in \mathbb{R}^d$  and  $B_t$  is a Brownian motion in  $\mathbb{R}^d$ . Denote by  $p(x, t)$  the density at point  $x$  on time  $t$ . Then,

$$\frac{\partial}{\partial t}p(x, t) = -\nabla \cdot (\mu(x, t)p(x, t)) + \frac{g(t)^2}{2}\Delta p(x, t) = -p\nabla \cdot \mu - \mu\nabla \cdot p + \frac{g(t)^2}{2}\Delta p.$$

### Preliminaries on analytic functions

**Definition A.4.** A function  $f: \mathbb{R}^d \rightarrow \mathbb{R}$  is analytic on  $\mathbb{R}^d$  if for any  $x_0, x \in \mathbb{R}^d$ , the Taylor series of  $f$  around  $x_0$ , evaluated at  $x$ , converges to  $f(x)$ . We say that  $F = (f_1, \dots, f_n): \mathbb{R}^d \rightarrow \mathbb{R}^d$  is an analytic function if  $f_i$  is analytic for all  $i \in \{1, \dots, n\}$ .

The following holds:

**Lemma A.5.** *If  $F, G: \mathbb{R}^d \rightarrow \mathbb{R}^d$  are two analytic functions and if  $F = G$  for all  $x \in U$  where  $U \subseteq \mathbb{R}^d$ ,  $U \neq \emptyset$ , is an open set, then  $F = G$  on all  $\mathbb{R}^d$ .*

This is a well known result and a proof sketch was given in Section 3.

**The heat equation.** The following is a Folklore lemma on the uniqueness of the solutions to the heat equation:

**Lemma A.6.** *Let  $p$  and  $p'$  be two continuous functions on  $\mathbb{R}^d \times [t_0, 1]$  that satisfy the heat equation*

$$\frac{\partial p}{\partial t} = \frac{g(t)^2}{2}\Delta p. \tag{13}$$

Further, assume that  $p(\cdot, t_0) = p'(\cdot, t_0)$ . Then,  $p = p'$  for all  $t \in [t_0, 1]$ .

## A.2 Main proof

In what appears below we denote

$$s(x, t) := \frac{h(x, t) - x}{\sigma_t^2}. \tag{14}$$

We start by claiming that if  $h$  satisfies Property 1, then  $s$  satisfies the PDE Eq. (10): (proof in Section A.3)

**Lemma A.7.** *Let  $h$  satisfy Property 1 and define  $s$  according to Eq. (14). Then,  $s$  satisfies Eq. (10).*

Next, we claim that the score function of any diffusion process satisfies the PDE Eq. (10): (proof in Section A.4)

**Lemma A.8.** *Let  $s$  be the score function of some diffusion process that is defined by Eq. (4). Then,  $s$  satisfies the PDE Eq. (10).*

To complete the first part of the proof, denote by  $p(\cdot, t)$  the probability distribution such that  $s(x, t) = \nabla \log p(x, t)$ , whose existence follows from Property 2. We would like to argue that  $\{p(\cdot, t)\}_{t \in (0, 1]}$  corresponds the probability density of the diffusion

$$dx_t = g(t)dB_t. \tag{15}$$

It suffices to show that for any  $t_0 > 0$ ,  $\{p(\cdot, t)\}_{t \in (t_0, 1]}$  corresponds to the same diffusion. To show the latter, let  $t_0 \in (0, 1)$  and consider the diffusion process according to Eq. (15) with the initial condition that

$x_{t_0} \sim p(\cdot, t_0)$ . Denote its score function by  $s'$  and notice that it satisfies the PDE Eq. (10) and the initial condition  $s'(x, t_0) = \nabla_x \log p(x, t_0) = s(x, t_0)$ , where the first equality follows from the definition of a score function and the second from the construction of  $p(x, t_0)$ . Further, recall that  $s(x, t)$  satisfies the same PDE Eq. (10) by Lemma A.3. Next we will show that  $s = s'$  for all  $t \in [t_0, 1]$ , and this will follow from the following lemma: (proof in Section A.5)

**Lemma A.9.** *Let  $s$  and  $s'$  be two solutions for the PDE (10) on the domain  $\mathbb{R}^d \times [t_0, 1]$  that satisfy the same initial condition at  $t_0$ :  $s(x, t_0) = s'(x, t_0)$  for all  $x$ . Further, assume that for all  $t \in [t_0, 1]$  there exist probability densities  $p(\cdot, t)$  and  $p'(\cdot, t)$  such that  $s(x, t) = \nabla_x \log p(x, t)$  and  $s'(x, t) = \nabla_x \log p'(x, t)$  for all  $x$ . Then,  $s = s'$  on all of  $\mathbb{R}^d \times [t_0, 1]$ .*

Then, by uniqueness of the PDE one obtains that  $s = s'$  for all  $t \in [t_0, 1]$ . Hence,  $s$  is the score of a diffusion for all  $t \geq t_0$  and this holds for any  $t_0 > 0$ , hence this holds for any  $t > 0$ . This concludes the proof of the first part of the theorem.

For the second part, let  $s^*$  denote some score function of a diffusion process that satisfies Eq. (4). Assume that for some  $t_0 > 0$  and some open subset  $U \subseteq \mathbb{R}^d$ ,  $s = s^*$ , namely  $s(x, t_0) = s^*(x, t_0)$  for all  $t_0 > 0$  and all  $x \in U$ . First, we would like to argue that if  $s(x, t)$  is the score function of some diffusion process that satisfies Eq. (4), then for any  $t_0 > 0$  it holds that  $s(x, t_0)$  is an analytic function (proof in Section A.6)

**Lemma A.10.** *Let  $x_t$  obey the SDE Eq. (4) with the initial condition  $x_0 \sim \mu_0$ . Let  $t > 0$  and let  $s(x, t)$  denote the score function of  $x_t$ , namely,  $s(x, t) = \nabla_x \log p(x, t)$  where  $p(x, t)$  is the density of  $x_t$ . Assume that  $\mu_0$  is a bounded-support distribution. Then,  $s(x, t)$  is an analytic function.*

Since both  $s$  and  $s^*$  are scores of diffusion processes, then  $s(x, t_0)$  and  $s^*(x, t_0)$  are analytic functions. Using the fact that  $s = s^*$  on  $U \times \{t_0\}$  and using Lemma A.5 we derive that  $s(x, t_0) = s^*(x, t_0)$  for all  $x$ . Let  $p$  and  $p^*$  denote the densities that correspond to the score functions  $s$  and  $s^*$  and by definition of a score function, we obtain that for all  $x$ ,

$$\nabla \log p(x, t_0) = s(x, t_0) = s^*(x, t_0) = \nabla \log p^*(x, t_0),$$

which implies, by integration, that

$$\log p(x, t_0) = \log p^*(x, t_0) + c$$

for some constant  $c \in \mathbb{R}$ . However,  $c = 0$ . Indeed,

$$1 = \int p(x, t_0) dx = \int e^{\log p(x, t_0)} dx = \int e^{\log p^*(x, t_0) + c} dx = \int p^*(x, t_0) e^c dx = e^c,$$

which implies that  $c = 0$  as required. As a consequence, the following lemma implies that  $p(x, 0) = p^*(x, 0)$  for all  $x$  (proof in Section A.7):

**Lemma A.11.** *Let  $x_t$  and  $y_t$  be stochastic processes that follow Eq. (4) with initial conditions  $x_0 \sim \mu_0$  and  $y_0 \sim \mu'_0$  and assume that  $\mu_0$  and  $\mu'_0$  are bounded-support. Assume that for some  $t_0 > 0$ ,  $x_{t_0}$  and  $y_{t_0}$  have the same distribution. Then,  $\mu_0 = \mu'_0$ .*

Without loss of generality, one can replace 0 with any  $\tilde{t} \in (0, t_0)$ , to obtain that  $p(x, \tilde{t}) = p^*(x, \tilde{t})$  for any  $\tilde{t} \in [0, t_0]$ . Now,  $p(x, t_0)$  is analytic, from Lemma A.5, hence it is continuous. Consequently, Lemma A.6 implies that  $p = p^*$  in  $\mathbb{R}^d \times [t_0, 1]$ . This concludes that  $p = p^*$  in all the domain, which implies that  $s = \nabla \log p = \nabla \log p^* = s^*$ , as required.

### A.3 Proof of Lemma A.7

We use Ito's lemma, and in particular Eq. (12), to get a PDE for the function  $h(x_t, t)$  where  $x_t$  satisfies the stochastic process

$$dx_t = -g(t)^2 s(x_t, t) dt + g(t) d\bar{B}_t.$$

Ito's formula yields that

$$dh(x_t, t) = \left( \frac{\partial h}{\partial t} - g(t)^2 J_F s - \frac{g(t)^2}{2} \Delta h \right) dt + \sigma J_h d\bar{B}_t.$$

Since  $(h, s)$  satisfies Property 1 and using Lemma 3.1,  $h$  is a reverse martingale which implies that the term that multiplies  $dt$  has to equal zero. In particular, we have that

$$\frac{\partial h}{\partial t} - g(t)^2 J_h s - \frac{g(t)^2}{2} \Delta h = 0. \quad (16)$$

By Eq. (14),

$$s = \frac{h - x}{\sigma_t^2}.$$

Therefore,

$$h = x + \sigma_t^2 s.$$

Substituting this in Eq. (16) and using the relation  $d\sigma_t^2/dt = g(t)^2$  that follows from Eq. (14), one obtains that

$$\begin{aligned} 0 &= \frac{\partial}{\partial t} (x + \sigma_t^2 s) - g(t)^2 J_{x + \sigma_t^2 s} s - \frac{g(t)^2}{2} \Delta (x + \sigma_t^2 s) \\ &= g(t)^2 s + \sigma_t^2 \frac{\partial s}{\partial t} - g(t)^2 (I + \sigma_t^2 J_s) s - \frac{g(t)^2 \sigma_t^2}{2} \Delta s \\ &= \sigma_t^2 \frac{\partial s}{\partial t} - g(t)^2 \sigma_t^2 J_s s - \frac{g(t)^2 \sigma_t^2}{2} \Delta s. \end{aligned}$$

Dividing by  $\sigma_t^2$ , we get that

$$\frac{\partial s}{\partial t} - g(t)^2 J_s s - \frac{g(t)^2}{2} \Delta s = 0,$$

which is what we wanted to prove.

#### A.4 Proof of Lemma A.8

We present as a consequence of the Fokker-Plank equation (Lemma A.3) a PDE for the log density  $\log p$ :

**Lemma A.12.** *Let  $x_t$  be defined by*

$$dx_t = \mu(x_t, t)dt + g(t)dB_t.$$

*Then,*

$$\frac{\partial \log p}{\partial t} = -\nabla \cdot \mu - \mu \nabla \cdot \log p + \frac{g(t)^2 \|\nabla \log p\|^2}{2} + \frac{g(t)^2 \Delta \log p}{2}$$

*Proof.* We would like to replace the partial derivatives of  $p$  that appears in Lemma A.3 with the partial derivatives of  $\log p$ . Using the formula

$$\frac{\partial \log p}{\partial t} = \frac{1}{p} \frac{\partial p}{\partial t},$$

one obtains that

$$\frac{\partial p}{\partial t} = p \frac{\partial \log p}{\partial t}.$$

Similarly,

$$\frac{\partial p}{\partial x_i} = p \frac{\partial \log p}{\partial x_i} \quad (17)$$

which also implies that

$$\nabla p = p \nabla \log p, \quad \nabla \cdot p = p \nabla \cdot \log p.$$

Differentiating Eq. (17) again with respect to  $x_i$  and applying Eq. (17) once more, one obtains that

$$\frac{\partial^2 p}{\partial x_i^2} = \frac{\partial}{\partial x_i} \left( p \frac{\partial \log p}{\partial x_i} \right) = \frac{\partial p}{\partial x_i} \frac{\partial \log p}{\partial x_i} + p \frac{\partial^2 \log p}{\partial x_i^2} = p \left( \left( \frac{\partial \log p}{\partial x_i} \right)^2 + \frac{\partial^2 \log p}{\partial x_i^2} \right).$$

Summing over  $i$ , one obtains that

$$\Delta p = p \sum_{i=1}^n \left( \left( \frac{\partial \log p}{\partial x_i} \right)^2 + \frac{\partial^2 \log p}{\partial x_i^2} \right) = p \|\nabla \log p\|^2 + p \Delta \log p. \quad (18)$$

Substituting the partials derivatives of  $p$  inside the Fokker-Planck equation in Lemma A.3, one obtains that

$$p \frac{\partial \log p}{\partial t} = -p \nabla \cdot \mu - \mu (p \nabla \cdot \log p) + \frac{g(t)^2}{2} (p \|\nabla \log p\|^2 + p \Delta \log p).$$

Dividing by  $p$ , one gets that

$$\frac{\partial \log p}{\partial t} = -\nabla \cdot \mu - \mu \nabla \cdot \log p + \frac{g(t)^2 \|\nabla \log p\|^2}{2} + \frac{g(t)^2 \Delta \log p}{2}.$$

as required.  $\square$

We are ready to prove Lemma A.8: Substituting  $\mu = 0$  in Lemma A.12, one obtains that

$$\frac{\partial \log p}{\partial t} = \frac{g(t)^2 \|\nabla \log p\|^2}{2} + \frac{g(t)^2 \Delta \log p}{2}.$$

Taking the gradient with respect to  $x$ , one obtains that

$$\nabla \frac{\partial \log p}{\partial t} = \frac{g(t)^2 \nabla \|\nabla \log p\|^2}{2} + \frac{g(t)^2 \nabla \Delta \log p}{2}. \quad (19)$$

Since  $\partial/\partial x_i$  commutes with  $\partial/\partial t$ , it holds that

$$\nabla \frac{\partial \log p}{\partial t} = \frac{\partial}{\partial t} \nabla \log p = \frac{\partial s}{\partial t}, \quad (20)$$

recalling that by definition  $s = \nabla \log p$ . Further,

$$\frac{\partial}{\partial x_i} \|\nabla \log p\|^2 = \sum_{j=1}^n \frac{\partial}{\partial x_i} \left( \frac{\partial \log p}{\partial x_j} \right)^2 = 2 \sum_{j=1}^n \frac{\partial^2 \log p}{\partial x_i \partial x_j} \frac{\partial \log p}{\partial x_j} = 2(H_{\log p} \nabla \log p)_i,$$

where for any function  $f: \mathbb{R}^d \rightarrow \mathbb{R}$ ,  $H_f$  is the Hessian function of  $f$  that is defined by

$$(H_f)_{ij} = \frac{\partial^2 f}{\partial x_i \partial x_j}$$

This implies that

$$\nabla \|\nabla \log p\|^2 = 2H_{\log p} \nabla \log p.$$

Further, notice that

$$H_f = J_{\nabla f},$$

which implies that

$$\nabla \|\nabla \log p\|^2 = 2J_{\nabla \log p} \nabla \log p = 2J_s s. \quad (21)$$

Lastly, we get that by the commutative property of partial derivatives,

$$\nabla \Delta \log p = \Delta \nabla \log p = \Delta s. \quad (22)$$

Substituting Eq. (20), Eq. (21) and Eq. (22) in Eq. (19), one obtains that

$$\frac{\partial s}{\partial t} = g(t)^2 J_s s + \frac{g(t)^2 \Delta s}{2},$$

as required.

## A.5 Proof of Lemma A.9

We will prove that  $p$  and  $p'$  satisfy the same PDE (which is the heat equation). Recall that  $s$  and  $s'$  satisfy

$$\frac{\partial s}{\partial t} = g(t)^2 \left( J_s s + \frac{1}{2} \Delta s \right) = \frac{g(t)^2}{2} (\nabla \|s\|^2 + \Delta s)$$

By substituting  $s = \nabla \log p$ ,

$$\frac{\partial \nabla \log p}{\partial t} = \frac{g(t)^2}{2} (\nabla \|\nabla \log p\|^2 + \Delta \nabla \log p).$$

By exchanging the order of derivatives, we obtain that

$$\nabla \frac{\partial \log p}{\partial t} = \nabla \frac{g(t)^2}{2} (\|\nabla \log p\|^2 + \Delta \log p).$$

By integrating, this implies that

$$\frac{\partial \log p}{\partial t} = \frac{g(t)^2}{2} (\|\nabla \log p\|^2 + \Delta \log p) + c(t),$$

where  $c(t)$  depends only on  $t$ . Eq. (18) shows that

$$\Delta \log p = \frac{\Delta p}{p} - \|\nabla \log p\|^2.$$

By substituting this in the equation above, we obtain that

$$\frac{\partial \log p}{\partial t} = \frac{g(t)^2}{2} \frac{\Delta p}{p} + c(t).$$

By multiplying both sides with  $p$ , we get that

$$\frac{\partial p}{\partial t} = p \frac{\partial \log p}{\partial t} = \frac{g(t)^2}{2} \Delta p + c(t). \quad (23)$$

Since  $p$  is a probability distribution,

$$\int_{\mathbb{R}^d} p(x, t) dx = 1,$$

therefore,

$$\int \frac{\partial p(x, t)}{\partial t} dx = \frac{\partial}{\partial t} \int_{\mathbb{R}^d} p(x, t) dx = \frac{\partial 1}{\partial t} = 0.$$

Integrating over Eq. (23) we obtain that

$$0 = \int \frac{g(t)^2}{2} \Delta p + c(t) dx = 0 + \int c(t) dx,$$

where the last equation holds since the integral of a Laplacian of probability density integrates to 0. It follows that  $c(t) = 0$  which implies that

$$\frac{\partial p}{\partial t} = \frac{g(t)^2}{2} \Delta p, \quad (24)$$

and the same PDE holds where  $p'$  replaces  $p$ , and this follows without loss of generality. Further, since  $\log p$  and  $\log p'$  are differentiable, it holds that  $p(\cdot, t)$  and  $p'(\cdot, t)$  are continuous for all fixed  $t$ . This implies that  $p$  and  $p'$  are continuous as functions of  $x$  and  $t$  since they both satisfy the heat equation Eq. (13). Consequently, Lemma A.6 implies that  $p = p'$  on  $\mathbb{R}^d \times [t_0, 1]$ . Finally,  $s = \nabla \log p = \nabla \log p' = s'$ , as required.

## A.6 Proof of Lemma A.10

First, recall that since  $x_t$  satisfies Eq. (4) with the initial condition  $x_0 \sim \mu_0$ , then  $x_t \sim \mu_0 + N(0, \sigma_t^2 I)$ , namely,  $x_t$  is the addition of a random variable drawn from  $\mu_0$  and an independent Gaussian  $N(0, \sigma_t^2 I)$ . Therefore, the density of  $x_t$ , which we denote by  $p(x, t)$ , equals

$$p(x, a) = \mathbb{E}_{a \sim \mu_0} \left[ \frac{1}{\sqrt{2\pi}\sigma_t} \exp\left(-\frac{\|x-a\|^2}{2\sigma_t^2}\right) \right].$$

Using the equation

$$\nabla_x \log f(x) = \frac{\nabla_x f(x)}{f(x)},$$

we get that

$$s(x, a) = \nabla_x \log p(x, a) = \frac{\nabla_x p(x, a)}{p(x, a)} = \frac{\mathbb{E}_{a \sim \mu_0} \left[ \frac{1}{\sqrt{2\pi}\sigma_t} \frac{x-a}{\sigma_t^2} \exp\left(-\frac{\|x-a\|^2}{2\sigma_t^2}\right) \right]}{\mathbb{E}_{a \sim \mu_0} \left[ \frac{1}{\sqrt{2\pi}\sigma_t} \exp\left(-\frac{\|x-a\|^2}{2\sigma_t^2}\right) \right]} \quad (25)$$

By using the fact that the Taylor formula for  $e^x$  equals

$$e^x = \sum_{i=0}^{\infty} \frac{e^i}{i!},$$

we obtain that the right hand side of Eq. (25) equals

$$\frac{\mathbb{E}_{a \sim \mu_0} \left[ \frac{1}{\sqrt{2\pi}\sigma_t} \frac{x-a}{\sigma_t^2} \sum_{i=0}^{\infty} \frac{(-1)^i}{i!} \left( \frac{\|x-a\|^2}{2\sigma_t^2} \right)^i \right]}{\mathbb{E}_{a \sim \mu_0} \left[ \frac{1}{\sqrt{2\pi}\sigma_t} \sum_{i=0}^{\infty} \frac{(-1)^i}{i!} \left( \frac{\|x-a\|^2}{2\sigma_t^2} \right)^i \right]} = \frac{\mathbb{E}_{a \sim \mu_0} \left[ \frac{x-a}{\sigma_t^2} \sum_{i=0}^{\infty} \frac{(-1)^i}{i!} \left( \frac{\|x-a\|^2}{2\sigma_t^2} \right)^i \right]}{\mathbb{E}_{a \sim \mu_0} \left[ \sum_{i=0}^{\infty} \frac{(-1)^i}{i!} \left( \frac{\|x-a\|^2}{2\sigma_t^2} \right)^i \right]} \quad (26)$$

We will use the following property of analytic functions: if  $f$  and  $g$  are analytic functions over  $\mathbb{R}^d$  and  $g(x) \neq 0$  for all  $x$  then  $f/g$  is analytic over  $\mathbb{R}^d$ . Since the denominator at the right hand side of Eq. (26) is nonzero, it suffices to prove that the numerator and the denominator are analytic. We will prove for the denominator and the proof for the numerator is nearly identical. By assumption of this lemma, the support of  $\mu_0$  is bounded, hence there is some  $M > 0$  such that  $\|x\| \leq M$  for any  $x$  in the support. Then,

$$\left| \frac{(-1)^i}{i!} \left( \frac{\|x-a\|^2}{2\sigma_t^2} \right)^i \right| \leq \frac{1}{i!} \left( \frac{x^2 + a^2}{\sigma_t^2} \right)^i = \frac{M^{2i}}{\sigma_t^{2i} i!}.$$

This bound is independent on  $a$ , and summing these absolute values of coefficients for  $i \in \mathbb{N}$ , one obtains a convergent series. Hence we can replace the summation and the expectation in the denominator at the right hand side of Eq. (26) to get that it equals

$$\sum_{i=0}^{\infty} \frac{(-1)^i}{i!} \mathbb{E}_{a \sim \mu_0} \left[ \left( \frac{\|x-a\|^2}{2\sigma_t^2} \right)^i \right]. \quad (27)$$

This is the Taylor series around 0 of the above-described denominator it converges to the value of the denominator at any  $x$ . While this Taylor series is taken around 0, we note the Taylor series around any other point  $x_0 \in \mathbb{R}^n$  converges as well. This can be shown by shifting the coordinate system by a constant vector such that  $x_0$  shifts to 0 and applying the same proof. One deduces that the Taylor series for the denominator around any point  $x_0$  converges on all  $\mathbb{R}^d$ , which implies that the denominator in the right hand side of Eq. (26) is analytic. The numerator is analytic as well by the same argument. Therefore the ratio, which equals  $s(x, t)$ , is analytic as well as required.

## A.7 Proof of Lemma A.11

Let  $t > 0$ , denote by  $\mu_t$  and  $\mu'_t$  the distributions of  $x_t$  and  $x'_t$ , respectively, and by  $p(x, t)$  and  $p'(x, t)$  the densities of these variables. Then,  $\mu_t = \mu_0 + N(0, \sigma_t^2 I)$ , namely,  $\mu_t$  is obtained by adding an independent sample from  $\mu_0$  with an independent  $N(0, \sigma_t^2 I)$  variables, and similarly for  $\mu'_t$ . Hence, the density  $p(x, t)$  is the convolution of the densities  $p(x, 0)$  with the density of a Gaussian  $N(0, \sigma_t^2 I)$ . Denote by  $\hat{p}(y, t)$  the Fourier transform of the density  $p(x, t)$  with respect to  $x$  (while keeping  $t$  fixed) and similarly define  $\hat{p}'$  as the Fourier transform of  $p'$ . Denote by  $g$  and by  $\hat{g}$  the density of  $N(0, \sigma_t^2 I)$  and its Fourier transform, respectively. Denote the convolution of two functions by the operator  $*$ . Then,

$$p(x, t) = p(x, 0) * g(x), \quad p'(x, t) = p'(x, 0) * g(x).$$

Since the Fourier transform turns convolutions into multiplications, one obtains that

$$\hat{p}(y, t) = \hat{p}(y, 0)\hat{g}(y), \quad \hat{p}'(y, t) = \hat{p}'(y, 0)\hat{g}(y).$$

Since  $p(x, t) = p'(x, t)$  we obtain that  $\hat{p}(y, t) = \hat{p}'(y, t)$ . Consequently,

$$\hat{p}(y, 0)\hat{g}(y) = \hat{p}'(y, 0)\hat{g}(y)$$

Since the Fourier transform of a Gaussian is nonzero, we can divide by  $\hat{g}(y)$  to get that

$$\hat{p}(y, 0) = \hat{p}'(y, 0).$$

This implies that the Fourier transform of  $p(x, 0)$  equals that of  $p'(x, 0)$  hence  $p(x, 0) = p'(x, 0)$  for all  $x$ , as required.

## B Other proofs

### B.1 Differentiating the loss function

Denote our parameter space as  $\Theta \subseteq \mathbb{R}^m$ . In order to differentiate  $L_{t, t', x}^1(\theta)$  with respect to  $\theta \in \Theta$ , we make the following calculations below, and we notice that  $\mathbb{E}_\theta$  is used to denote an expectation with respect to the distribution of  $x_{[t', t]}$  according to Eq. (7) with  $s = s_\theta$  and the initial condition  $x_t = x$ . In other words, the expectation is over  $x_{[t', t]}$  that is taken with respect to the sampler that is parameterized by  $\theta$  with the initial condition  $x_t = x$ . We denote by  $p_\theta(x_{[t', t]} | x_t = x)$  the corresponding density of  $x_{[t', t]}$ . For any function  $f = (f_1, \dots, f_n): \Theta \rightarrow \mathbb{R}^n$ , denote by  $\nabla_\theta f$  the Jacobian matrix of  $f$ , where

$$(\nabla_\theta f)_{i,j} = \frac{\partial f_i}{\partial \theta_j}.$$

For notational consistency, if  $f$  is a single-valued function, namely, if  $n = 1$ , then  $\nabla_\theta f$  is a column vector. We begin with the following:

$$\begin{aligned}
\nabla_\theta \mathbb{E}_\theta [h_\theta(x_{t'}, t')] &= \nabla_\theta \int_{\mathbb{R}^d} h_\theta(x_{t'}, t') p_\theta(x_{[t', t]} | x_t = x) dx_{t'} \\
&= \int_{\mathbb{R}^d} \nabla_\theta h_\theta(x_{t'}, t') p_\theta(x_{[t', t]} | x_t = x) dx_{t'} + \int_{\mathbb{R}^d} h_\theta(x_{t'}, t') \nabla_\theta p_\theta(x_{[t', t]} | x_t = x) dx_{t'} \\
&= \mathbb{E}_\theta [\nabla_\theta h_\theta(x_{t'}, t')] + \mathbb{E}_\theta \left[ h_\theta(x_{t'}, t') \frac{\nabla_\theta p_\theta(x_{[t', t]} | x_t = x)}{p_\theta(x_{[t', t]} | x_t = x)} \right] \\
&= \mathbb{E}_\theta [\nabla_\theta h_\theta(x_{t'}, t')] + \mathbb{E}_\theta [h_\theta(x_{t'}, t') \nabla_\theta \log(p_\theta(x_{[t', t]} | x_t = x))]
\end{aligned}$$

Differentiating the whole loss, we get the following:

$$\begin{aligned}
\nabla_\theta L_{t, t', x}^1(\theta) &= \frac{1}{2} \nabla_\theta (\mathbb{E}_\theta [h_\theta(x_{t'}, t')] - h_\theta(x, t))^2 \\
&= (\mathbb{E}_\theta [h_\theta(x_{t'}, t')] - h_\theta(x, t))^\top (\nabla_\theta \mathbb{E}_\theta [h_\theta(x_{t'}, t')] - \nabla_\theta h_\theta(x, t)) \\
&= \mathbb{E}_\theta [h_\theta(x_{t'}, t') - h_\theta(x, t)]^\top \mathbb{E}_\theta [\nabla_\theta h_\theta(x_{t'}, t') - \nabla_\theta h_\theta(x, t)] \\
&\quad + \mathbb{E}_\theta [h_\theta(x_{t'}, t') - h_\theta(x, t)]^\top \mathbb{E}_\theta [h_\theta(x_{t'}, t') \nabla_\theta \log(p_\theta(x_{[t', t]} | x_t = x))]
\end{aligned}$$

Let us compute the gradient of the log density. We use the discrete process, and let us assume that  $t = t_0 > t_1 > \dots > t_k = t'$  are the sampling times. Then,

$$p_\theta(x_{[t', t]} | x_t = x) = \prod_{i=1}^k p_\theta(x_{t_i} | x_{t_{i-1}}).$$

We assume that

$$p_\theta(x_{t_i} | t_{i-1}) = \mathcal{N}(\mu_{\theta, i}, g_i I_d).$$

Then,

$$p_\theta(x_{[t', t]} | x_t = x) \propto \prod_{i=1}^k \exp\left(-\frac{\|\mu_{\theta, i} - (x_{t_i} - x_{t_{i-1}})\|^2}{2g_i^2}\right)$$

Therefore

$$\log p_\theta(x_{[t', t]} | x_t = x) = C + \sum_{i=1}^k \frac{\|\mu_{\theta, i} - (x_{t_i} - x_{t_{i-1}})\|^2}{2g_i^2}$$

where  $C$  corresponds to the normalizing factor that is independent of  $\theta$ . Differentiating, we get that

$$\nabla_\theta \log p_\theta(x_{[t', t]} | x_t = x) = \sum_{i=1}^k \frac{(\mu_{\theta, i} - (x_{t_i} - x_{t_{i-1}}))^\top \nabla_\theta \mu_{\theta, i}}{g_i^2}$$

## B.2 Proof of Lemma 3.1

In what appears below, the expectation  $\mathbb{E}[\cdot | x_t = x]$  is taken with respect to the distribution obtained by Eq. (7), namely, the backward SDE that corresponds to the function  $s$ , with the initial condition  $x_t = x$ . Similarly,  $\mathbb{E}[\cdot | x_{t'}]$  is taken with the initial condition at  $x_{t'}$ . To prove the first direction in the equivalence, assume that Property 1 holds and our goal is to prove the two consequences as described in the lemma. To prove the first consequence, by the law of total expectation and by the fact that  $x_t - x_{t'} - x_0$  is a Markov chain, namely,  $x_0$  and  $x_t$  are independent conditioned on  $x_{t'}$ , we obtain that

$$h(x, t) = \mathbb{E}[x_0 | x_t = x] = \mathbb{E}[\mathbb{E}[x_0 | x_{t'}] | x_t = x] = \mathbb{E}[h(x_{t'}, t') | x_t = x].$$



To prove the second consequence, by Property 1

$$h(x, 0) = \mathbb{E}[x_0 | x_0 = x] = x_0.$$

This concludes the first direction in the equivalence.

To prove the second direction, assume that  $h(x, t) = \mathbb{E}[h(x_{t'}, t') | x_t = x]$  and that  $h(x, 0) = x$  and notice that by substituting  $t' = 0$  we derive the following:

$$h(x, t) = \mathbb{E}[h(x_0, 0) | x_t = x] = \mathbb{E}[x_0 | x_t = x],$$

as required.

## C Additional Results

### C.1 Property Testing




Figure 4: Consistency Property Testing on CIFAR10. The plot illustrates how the Consistency Loss,  $L_{t,t',x_t}^2$ , behaves for  $t' = 0$ , as  $t$  changes.

Consistency Property Testing (CIFAR10)  
Backward Sampler (300) steps



Figure 5: Consistency Property Testing on CIFAR10. The plot illustrates how the Consistency Loss,  $L_{t,t',x_t}^2$ , behaves for  $t = 0$ , as  $t'$  changes.




Figure 6: Consistency Property Testing on FFHQ. The plot illustrates how the Consistency Loss,  $L_{t,t',x_t}^2$ , behaves for  $t' = 0$ , as  $t$  changes.




Figure 7: Consistency Property Testing on FFHQ. The plot illustrates how the Consistency Loss,  $L_{t,t',x_t}^2$ , behaves for  $t = 0$ , as  $t'$  changes.

## C.2 Uncurated Samples



Figure 8: Uncurated generated images by our fine-tuned model on FFHQ. FID: 2.61, NFEs: 79.