



Article How Much Is Enough? A Study on Diffusion Times in Score-Based Generative Models

Giulio Franzese ^{1,*}, Simone Rossi ¹, Lixuan Yang ², Alessandro Finamore ², Dario Rossi ², Maurizio Filippone ¹ and Pietro Michiardi ¹

- ¹ EURECOM Data Science Department, 06410 Biot, France
- ² Huawei Technologies Paris, 92100 Boulogne-Billancourt, France

* Correspondence: giulio.franzese@eurecom.fr

Abstract: Score-based diffusion models are a class of generative models whose dynamics is described by stochastic differential equations that map noise into data. While recent works have started to lay down a theoretical foundation for these models, a detailed understanding of the role of the diffusion time *T* is still lacking. Current best practice advocates for a large *T* to ensure that the forward dynamics brings the diffusion sufficiently close to a known and simple noise distribution; however, a smaller value of *T* should be preferred for a better approximation of the score-matching objective and higher computational efficiency. Starting from a variational interpretation of diffusion models, in this work we quantify this trade-off and suggest a new method to improve quality and efficiency of both training and sampling, by adopting smaller diffusion times. Indeed, we show how an auxiliary model can be used to bridge the gap between the ideal and the simulated forward dynamics, followed by a standard reverse diffusion process. Empirical results support our analysis; for image data, our method is competitive with regard to the state of the art, according to standard sample quality metrics and log-likelihood.

Keywords: generative modelling; diffusion models; variational inference; efficiency



Citation: Franzese, G.; Rossi, S.; Yang, L.; Finamore, A.; Rossi, D.; Filippone, M.; Michiardi, P. How Much Is Enough? A Study on Diffusion Times in Score-Based Generative Models. *Entropy* **2023**, *25*, 633. https://doi.org/10.3390/ e25040633

Academic Editor: Jakub Tomczak

Received: 10 March 2023 Revised: 28 March 2023 Accepted: 29 March 2023 Published: 7 April 2023



Copyright: © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https:// creativecommons.org/licenses/by/ 4.0/). 1. Introduction

Diffusion-based generative models [1–7] have recently gained popularity due to their ability to synthesize high-quality audio [8,9], image [10,11] and other data modalities [12], outperforming known methods based on Generative Adversarial Networks (GANS) [13], normalizing flows (NFs) [14] or Variational Autoencoders (VAEs) and Bayesian autoencoders (BAEs) [15,16].

Diffusion models learn to generate samples from an unknown density p_{data} by reversing a *diffusion process* which transforms the distribution of interest into noise. The forward dynamics injects noise into the data following a diffusion process that can be described by a Stochastic Differential Equation (SDE) of the form

$$dx_t = f(x_t, t)dt + g(t)dw_t \quad \text{with} \quad x_0 \sim p_{\text{data}} , \qquad (1)$$

where x_t is a random variable at time t, $f(\cdot, t)$ is the *drift term*, $g(\cdot)$ is the *diffusion term* and w_t is a *Wiener process* (or Brownian motion). We also consider a special class of linear SDEs, for which the drift term is decomposed as $f(x_t, t) = \alpha(t)x_t$, where the function $\alpha(t) \le 0$ for all t, and the diffusion term is independent of x_t . This class of parameterizations of SDEs is known as *affine* and it admits analytic solutions. We denote the time-varying probability density by p(x, t), where, by definition $p(x, 0) = p_{data}(x)$, and the conditional on the initial condition x_0 by $p(x, t | x_0)$. The forward SDE is usually considered for a "sufficiently long" *diffusion time* T, leading to the density p(x, T). In principle, when $T \to \infty$, p(x, T) converges to Gaussian noise, regardless of initial conditions.

For generative modeling purposes, we are interested in the inverse dynamics of such process, i.e., transforming samples of the noisy distribution p(x, T) into $p_{data}(x)$. Such dynamics can be obtained by considering the solutions of the inverse diffusion process [17],

$$\mathbf{d}\mathbf{x}_t = \left[-f(\mathbf{x}_t, t') + g^2(t') \nabla \log p(\mathbf{x}_t, t')\right] \mathbf{d}t + g(t') \mathbf{d}\mathbf{w}_t,$$
(2)

where $t' \stackrel{\text{def}}{=} T - t$, with the inverse dynamics involving a new Wiener process. Given p(x, T) as the initial condition, the solution of Equation (2) after a *reverse diffusion time* T, will be distributed as $p_{data}(x)$. We refer to the density associated to the backward process as q(x, t'). The simulation of the backward process is referred to as *sampling* and, differently from the forward process, this process is not *affine* and a closed form solution is out of reach.

Practical considerations on diffusion times. In practice, diffusion models are challenging to work with [3]. Indeed, a direct access to the true *score* function $\nabla \log p(x_t, t)$ required in the dynamics of the reverse diffusion is unavailable. This can be solved by approximating it with a parametric function $s_{\theta}(x_t, t)$, e.g., a neural network, which is trained using the following loss function:

$$\mathcal{L}(\boldsymbol{\theta}) = \int_0^T \mathbb{E}_{\sim(1)} \lambda(t) \|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t) - \boldsymbol{\nabla} \log p(\boldsymbol{x}_t, t \,|\, \boldsymbol{x}_0)\|^2$$
(3)

where $\lambda(t)$ is a positive weighting factor and the notation $\mathbb{E}_{\sim(1)}$ means that the expectation is taken with respect to the random process x_t in Equation (1): for a generic function h, $\mathbb{E}_{\sim(1)}[h(x_t, x_0, t)] = \int h(x, z, t)p(x, t | z)p_{data}(z)dxdz$. The loss in Equation (3), usually referred to as *score matching loss*, is the cost function considered in [18] (Equation (4)). The condition $\lambda(t) = g(t)^2$, which we use in this work, is referred to a *likelihood reweighting*. Due to the affine property of the drift, the term $p(x_t, t | x_0)$ is analytically known and normally distributed for all t (expression available in Table 1, and in Särkkä and Solin [19]). Intuitively, the estimation of the *score* is akin to a denoising objective, which operates in a challenging regime. Later, we will quantify the difficulty of learning the *score*, as a function of T.

While the forward and reverse diffusion processes are valid for all *T*, the noise distribution p(x, T) is analytically known only when the diffusion time is $T \rightarrow \infty$. Then, the common solution is to replace p(x, T) with a simple (i.e., easy to sample) distribution $p_{\text{noise}}(x)$, which, for the classes of SDEs that we consider in this work, is a Gaussian distribution.

In the literature, the discrepancy between p(x, T) and $p_{noise}(x)$ has been neglected, under the informal assumption of a sufficiently large diffusion time. Unfortunately, while this approximation seems a valid approach to simulate and generate samples, the reverse diffusion process starts from an initial condition q(x, 0) which is different from p(x, T) and, as a consequence, it will converge to a solution q(x, T) that is different from the true $p_{data}(x)$. Later, we will expand on the error introduced by this approximation, but for illustration purposes, Figure 1 shows quantitatively this behavior for a simple 1D toy example where we set the data distribution equal to a mixture of normal (\mathcal{N}) distributions as $p_{data}(x) = \pi \mathcal{N}(1, 0.1^2) + (1 - \pi) \mathcal{N}(3, 0.5^2)$, with $\pi = 0.3$. When *T* is small, the distribution $p_{noise}(x)$ is very different from p(x, T) and samples from q(x, T) exhibit very low likelihood of being generated from $p_{data}(x)$.

Crucially, Figure 1 (zoomed region) illustrates an unknown behavior of diffusion models, which we unveil in our analysis. The right balance between efficient *score* estimation and sampling quality can be achieved by diffusion times that are smaller than common best practices. Moreover, even excessively large diffusion times can be detrimental. This is a key observation that we explore in our work.



Effect of diffusion time

Figure 1. Effect of T on a toy model: low diffusion times are detrimental for sample quality (likelihood of 1024 samples as median and 95 quantile, on 8 random seeds).

Contributions. An appropriate choice of the diffusion time T is a key factor that impacts training convergence, sampling time and quality. On the one hand, the approximation error introduced by considering initial conditions for the reverse diffusion process drawn from a simple distribution $p_{\text{roise}}(\mathbf{x}) \neq p(\mathbf{x}, T)$ increases when T is small. This is why the current best practice is to choose a sufficiently long diffusion time. On the other hand, training convergence of the score model $s_{\theta}(x_t, t)$ becomes more challenging to achieve with a large T, which also imposes extremely high computational costs **both** for training and for sampling. This would suggest to choose a smaller diffusion time. Given the importance of this problem, in this work, we set off to study the existence of suitable operating regimes to strike the right balance between computational efficiency and model quality. The main contributions of this work are the following:

Contribution 1: We use an evidence lower bound (ELBO) decomposition which allows us to study the impact of the diffusion time T. This ELBO decomposition emphasizes the roles of (i) the discrepancy between the "ending" distribution of the diffusion and the "starting" distribution of the reverse diffusion processes, and (ii) of the score matching objective. Crucially, our analysis does not rely on assumptions on the quality of the score models. We explicitly study the existence of a trade-off and explore experimentally, for the first time, current approaches for selecting the diffusion time T. Contribution 2: In Section 3, we propose a novel method to improve *both* the training and sampling efficiency of diffusion-based models, while maintaining high sample quality. Our method introduces an auxiliary distribution, allowing us to transform the simple "starting" distribution of the reverse process used in the literature so as to minimize the discrepancy to the "ending" distribution of the forward process. Then, a standard reverse diffusion can be used to closely match the data distribution. Intuitively, our method allows to build "bridges" across multiple distributions, and to set T toward the advantageous regime of small diffusion times.

In addition to our methodological contributions, in Section 4, we provide experimental evidence of the benefits of our method, in terms of sample quality and log likelihood. Finally, we conclude this work in Section 5.

Related Work. A concurrent work by Zheng et al. [20] presents an empirical study of a truncated diffusion process but lacks a rigorous analysis and a clear justification for the proposed approach. Recent attempts by Lee et al. [9] to optimize p_{noise} , or the proposal to do so [21], have been studied in different contexts. Related work focus primarily on improving sampling efficiency (but not training efficiency), using a wide array of techniques. Sample generation times can be drastically reduced considering adaptive step-size integrators [22]. Such methods are complementary to our approach, and can be used in combination with the techniques we propose in this work. Other popular choices are based on merging multiple steps of a pretrained model through distillation techniques [23] or by taking larger sampling steps with GANs [24]. Approaches closer to ours *modify* the SDE, or the discrete time processes, to obtain inference efficiency gains. In particular, Song et al. [7] considers

implicit non-Markovian diffusion processes, while Watson et al. [25] changes the diffusion processes by optimal scheduling selection, and Dockhorn et al. [26] considers overdamped SDEs. Finally, hybrid techniques combining VAEs and diffusion models [4] or simple auto encoders and diffusion models [27] have positive effects on training and sampling times.

Moreover, we remark that a simple modification of the noise schedule to steer the diffusion process toward a small diffusion time [5,28] is not a viable solution. As we discuss in Section 2.4, the optimal value of the ELBO, in the case of affine SDEs, is invariant to the choice of the noise schedule. Naively selecting a faster noise schedule does not provide any practical benefit in terms of computational complexity, as it requires smaller step sizes to keep the same accuracy of the original noise schedule simulation. However, the optimization of the noise schedule can have important practical effects on the stability of training and variance of estimations [5]. Finally, few other works in the literature attempt to study the convergence properties of diffusion models. For instance, De Bortoli et al. [29] obtain a total variation bound between the generated and data distribution under maximum error assumptions between true and approximated score. De Bortoli [30] relaxes this requirement obtaining a bound in terms of Wasserstein distance. Lee et al. [31] show how the total variation bound can be expressed as a function of the maximum score error and find that the bound is optimized for a diffusion time that depends on this error. Our work, on the other hand, does not make any assumption and aims at selecting the smallest possible diffusion time to maximize training and sampling efficiency.

2. A Tradeoff on Diffusion Time

The dynamics of a diffusion model can be studied through the lens of variational inference, which allows us to bound the (log-)likelihood using an evidence lower bound (ELBO) [32]. The interpretation we consider in this work (see also [18], Theorem 1) emphasizes the two main factors affecting the quality of sample generation: an imperfect *score* and a mismatch, measured in terms of KL[log $p(x, T) \parallel p_{noise}(x)$], the Kullback-Leibler (KL) divergence between the noise distribution p(x, T) of the forward process and the distribution p_{noise} used to initialize the backward process.

2.1. Preliminaries: The ELBO Decomposition

Our goal is to study the quality of the generated data distribution as a function of the diffusion time *T*. Instead of focusing on the log-likelihood bounds for single datapoints $\log q(x, T)$, we consider the average over the data distribution, i.e., the *cross-entropy* $\mathbb{E}_{p_{data}(x)} \log q(x, T)$. By rewriting the $\mathcal{L}_{\text{ELBO}}$ derived in Huang et al. [32] [Equation (25)] (details of the steps in Appendix B), we have that

$$\mathbb{E}_{p_{\textit{data}}(\boldsymbol{x})}\log q(\boldsymbol{x},T) \ge \mathcal{L}_{\text{ELBO}}(\boldsymbol{s}_{\boldsymbol{\theta}},T) = \mathbb{E}_{\sim(1)}\log p_{\textit{noise}}(\boldsymbol{x}_{T}) - I(\boldsymbol{s}_{\boldsymbol{\theta}},T) + R(T)$$
(4)

where $R(T) = \frac{1}{2} \int_{t=0}^{T} \mathbb{E}_{\sim(1)} \Big[g^2(t) \| \nabla \log p(\mathbf{x}_t, t \mid \mathbf{x}_0) \|^2 - 2f^{\top}(\mathbf{x}_t, t) \nabla \log p(\mathbf{x}_t, t \mid \mathbf{x}_0) \Big] dt$, and $I(\mathbf{s}_{\theta}, T) = \frac{1}{2} \int_{t=0}^{T} g^2(t) \mathbb{E}_{\sim(1)} \Big[\| \mathbf{s}_{\theta}(\mathbf{x}_t, t) - \nabla \log p(\mathbf{x}_t, t \mid \mathbf{x}_0) \|^2 \Big] dt$ is equal to the loss term

Equation (3) when $\lambda(t) = g^2(t)$.

Note that R(T) depends neither on s_{θ} nor on p_{noise} , while $I(s_{\theta}, T)$, or an equivalent reparameterization [18,32] [Equation (1)], is used to learn the approximated *score*, by optimization of the parameters θ . It is then possible to show that

$$I(s_{\theta}, T) \ge \underbrace{I(\nabla \log p, T)}_{\stackrel{\text{def}}{=} K(T)} = \frac{1}{2} \int_{t=0}^{1} g^{2}(t) \mathbb{E}_{\sim(1)} [\|\nabla \log p(x_{t}, t) - \nabla \log p(x_{t}, t \mid x_{0})\|]^{2} dt.$$
(5)

Note that the term $K(T) = I(\nabla \log p, T)$ does not depend on θ . Consequently, we can define $\mathcal{G}(s_{\theta}, T) = I(s_{\theta}, T) - K(T)$ (see Appendix C for details), where $\mathcal{G}(s_{\theta}, T)$ is a positive term that we call the *gap* term, accounting for the practical case of an imperfect *score*, i.e., $s_{\theta}(x_t, t) \neq \nabla \log p(x_t, t)$. It also holds that

$$\mathbb{E}_{\sim(1)}\log p_{\text{noise}}(\mathbf{x}_T) = \int [\log p_{\text{noise}}(\mathbf{x}) - \log p(\mathbf{x}, T) + \log p(\mathbf{x}, T)]p(\mathbf{x}, T)d\mathbf{x} =$$
$$= \mathbb{E}_{\sim(1)}\log p(\mathbf{x}_T, T) - \mathrm{KL}[\log p(\mathbf{x}, T) \parallel p_{\text{noise}}(\mathbf{x})]. \tag{6}$$

Therefore, we can substitute the cross-entropy term $\mathbb{E}_{\sim(1)} \log p_{noise}(x_T)$ of the ELBO in Equation (4) to obtain

$$\mathbb{E}_{p_{data}(x)} \log q(x,T) \geq -\operatorname{KL}[p(x,T) \parallel p_{noise}(x)] + \mathbb{E}_{\sim(1)} \log p(x_T,T) - K(T) + R(T) - \mathcal{G}(s_{\theta},T).$$
(7)

Before concluding our derivation, we show how to combine different terms of Equation (7) into the negative entropy term $\mathbb{E}_{p_{data}(x)} \log p_{data}(x)$. Given the stochastic dynamics defined in Equation (1), it holds that (see derivation and details in Appendix D)

$$\mathbb{E}_{\sim(1)}\log p(\boldsymbol{x}_T, T) - K(T) + R(T) = \mathbb{E}_{p_{data}(\boldsymbol{x})}\log p_{data}(\boldsymbol{x}).$$
(8)

Finally, we can now bound the value of $\mathbb{E}_{p_{data}(x)} \log q(x, T)$ as

$$\mathbb{E}_{p_{data}(x)} \log q(x,T) \ge \underbrace{\mathbb{E}_{p_{data}(x)} \log p_{data}(x) - \mathcal{G}(s_{\theta},T) - \mathrm{KL}[p(x,T) \parallel p_{noise}(x)]}_{\mathcal{L}_{\mathrm{ELBO}}(s_{\theta},T)}.$$
(9)

Equation (9) clearly emphasizes the roles of an approximate score function, through the gap term $\mathcal{G}(\cdot)$, and the discrepancy between the noise distribution of the forward process and the initial distribution of the reverse process through the KL term. The (negative) entropy term $\mathbb{E}_{p_{data}}(x) \log p_{data}(x)$, which is constant with regard to T and θ , is the best value achievable by the ELBO. Indeed, by rearranging Equation (9), $\text{KL}[p_{data}(x) \parallel q(x, T)] \leq \mathcal{G}(s_{\theta}, T) + \text{KL}[p(x, T) \parallel p_{noise}(x)]$. Optimality is achieved when (i) we have perfect *score* matching and (ii) the initial conditions for the reverse process are ideal, i.e., q(x, 0) = p(x, T).

Next, we show the existence of a tradeoff: the KL decreases with T, while the gap increases with T.

2.2. The Tradeoff on Diffusion Time

We begin by showing that the KL term in Equation (9) decreases with the diffusion time *T*, which induces to select large *T* to maximize the ELBO.

We consider the two main classes of SDEs for the forward diffusion process defined in Equation (1): SDEs whose steady state distribution is the standard multivariate Gaussian, referred to as *Variance Preserving* (VP), and SDEs without a stationary distribution, referred to as *Variance Exploding* (VE), which we summarize in Table 1. The standard approach to generate new samples relies on the backward process defined in Equation (2), and consists in setting p_{noise} in agreement with the form of the forward process SDE. The following result bounds the discrepancy between the noise distribution p(x, T) and p_{noise} .

Table 1. Two main families of diffusion processes, where $\sigma^2(t) = \left(\frac{\sigma_{max}^2}{\sigma_{min}^2}\right)^t$ and $\beta(t) = \beta_0 + (\beta_1 - \beta_0)t$.

	Diffusion Process	$p(x_t,t x_0) = \mathcal{N}(m,sI)$	$p_{noise}\left(x ight)$
Variance Exploding	$\alpha(t) = 0, g(t) = \sqrt{\frac{\mathrm{d}\sigma^2(t)}{\mathrm{d}t}}$	$\boldsymbol{m} = \boldsymbol{x}_0, \boldsymbol{s} = \sigma^2(t) - \sigma^2(0)$	$\mathcal{N}(0, (\sigma^2(T) - \sigma^2(0))\mathbf{I})$
Variance Preserving	$\alpha(t) = -\frac{1}{2}\beta(t), g(t) = \sqrt{\beta(t)}$	$m = e^{-\frac{1}{2}\int_0^t \beta(\mathrm{d} au)} x_0, s = 1 - e^{-\int_0^t \beta(\mathrm{d} au)}$	$\mathcal{N}(0, I)$

Lemma 1. For the classes of SDEs considered (Table 1), the discrepancy between p(x, T) and the $p_{\text{noise}}(x)$ can be bounded as follows.

For Variance Preserving SDEs, it holds that: $\operatorname{KL}[p(\mathbf{x},T) \parallel p_{\operatorname{noise}}(\mathbf{x})] \leq C_1 \exp\left(-\int_0^T \beta(t) dt\right)$. For Variance Exploding SDEs, it holds that: $\operatorname{KL}[p(\mathbf{x},T) \parallel p_{\operatorname{noise}}(\mathbf{x})] \leq C_2 \frac{1}{\sigma^2(T) - \sigma^2(0)}$.

Our proof uses results from Villani [33], the logarithmic Sobolev Inequality and Gronwall inequality (see Appendix E for details). The consequence of Lemma 1 is that to maximize the ELBO, the diffusion time *T* should be as large as possible (ideally, $T \rightarrow \infty$), such that the KL term vanishes. This result is in line with current practices for training score-based diffusion processes, which argue for sufficiently long diffusion times [29]. Our analysis, on the other hand, highlights how this term is only one of the two contributions to the ELBO.

Now, we focus our attention on studying the behavior of the second component, $\mathcal{G}(\cdot)$. Before that, we define a few quantities that allow us to write the next important result.

Definition 1. We define the **optimal score** \hat{s}_{θ} for any diffusion time *T*, as the score obtained using parameters that minimize $I(s_{\theta}, T)$. Similarly, we define the **optimal score gap** $\mathcal{G}(\hat{s}_{\theta}, T)$ for any diffusion time *T*, as the gap attained when using the optimal score.

The optimal score gap term $\mathcal{G}(\hat{s}_{\theta}, T)$ is a non-decreasing function in T. That is, given $T_2 > T_1$, and $\theta_1 = \arg \min_{\theta} I(s_{\theta}, T_1), \theta_2 = \arg \min_{\theta} I(s_{\theta}, T_2)$, then $\mathcal{G}(s_{\theta_2}, T_2) \ge \mathcal{G}(s_{\theta_1}, T_1)$. The proof (see Appendix F) is a direct consequence of the definition of \mathcal{G} and the optimality of the score.

Note that Section 2.2 does not imply that $\mathcal{G}(s_{\theta_a}, T_2) \geq \mathcal{G}(s_{\theta_b}, T_1)$ holds for generic parameters θ_a, θ_b .

2.3. Is There an Optimal Diffusion Time?

While diffusion processes are generally studied for $T \rightarrow \infty$, diffusion times in scorebased models have been arbitrarily set to be "sufficiently large" in the literature. Here we formally argue about the existence of an optimal diffusion time, which strikes the right balance between the gap $\mathcal{G}(\cdot)$ and the KL terms of the ELBO in Equation (9).

Before proceeding any further, we clarify that **our final objective in this work is not to find and use an optimal diffusion time**. Instead, our result on the existence of optimal diffusion times (which can be smaller than the ones set by than popular heuristics) serves the purpose of motivating the choice of small diffusion times, which, however, calls for a method to overcome approximation errors. For completeness, in Appendix H, we show that optimizing the ELBO to obtain an optimal diffusion time T^* is technically feasible, without resorting to exhaustive grid search.

Consider the ELBO decomposition in Equation (9). We study it as a function of time T, and seek its optimal argument $T^* = \arg \max_T \mathcal{L}_{\text{ELBO}}(\widehat{s}_{\theta}, T)$. Then, the optimal diffusion time $T^* \in \mathbb{R}^+$, and thus not necessarily $T^* = \infty$. Additional assumptions on the gap term $\mathcal{G}(\cdot)$ can be used to guarantee strict finiteness of T^* .

It is trivial to verify that, since the optimal gap term $\mathcal{G}(\hat{s}_{\theta}, T)$ is a non decreasing function in *T* (Section 2.2), we have $\frac{\partial \mathcal{G}}{\partial T} \geq 0$. Then, we study the sign of the KL derivative, which is always negative as shown in Appendix G. Moreover, we know that that $\lim_{T\to\infty} \frac{\partial KL}{\partial T} = 0$. Consequently, the function $\frac{\partial \mathcal{L}_{\text{ELBO}}}{\partial T} = \frac{\partial \mathcal{G}}{\partial T} + \frac{\partial KL}{\partial T}$ has at least one zero in its

domain \mathbb{R}^+ . To guarantee a stricter bounding of T^* , we could study asymptotically the growth rates of \mathcal{G} and the KL terms for large T. The investigation is technically involved and outside the scope of this paper. Nevertheless, as discussed hereafter, the numerical investigation carried out in this work suggests finiteness of T^* .

Empirically, we use Figure 2 to illustrate the tradeoff and the optimality arguments through the lens of the same toy example we use in Section 1. On the first and third columns, we show the ELBO decomposition. We can verify that $\mathcal{G}(s_{\theta}, T)$ is an increasing function of T, whereas the KL term is a decreasing function of T. Even in the simple case of a toy example, the tension between small and large values of T is clear. On the second and fourth columns, we show the values of the ELBO and of the likelihood as a function of T. We then verify the validity of our claims: the ELBO is neither maximized by an infinite diffusion time, nor by a "sufficiently large" value. Instead, there exists an optimal diffusion time which, for this example, is smaller than T = 1.0, which is typically used in practice.



Figure 2. ELBO decomposition, ELBO and likelihood for a 1D toy model, as a function of diffusion time *T*. Tradeoff and optimality numerical results confirm our theory.

In Section 3, we present a new method that admits much smaller diffusion times and we show that the ELBO of our approach is at least as good as the one of a standard diffusion model, configured to use its optimal diffusion time T^* .

2.4. Relation with Diffusion Process Noise Schedule

We remark that a simple modification of the noise schedule to steer the the diffusion process toward a small diffusion time [5,28] is not a viable solution. In Appendix J, we discuss how the optimal value of the ELBO, in the case of affine SDEs, is *invariant* to the choice of the noise schedule. Indeed, its value depends uniquely on the relative level of corruption of the initial data at the considered final diffusion time T, that is, the Signalto-Noise Ratio. Naively, we could think that, by selecting a twice as fast noise schedule, we would be able to obtain the same ELBO of the original schedule by diffusing only for half the time. While true, this does not provide any practical benefit in terms of computational complexity. If the noise schedule is faster, the drift terms involved in the reverse process changes more rapidly. Consequently, to *simulate* the reverse SDE with a numerical integration scheme, smaller step sizes are required to keep the same accuracy of the original noise schedule simulation. The effect is that, while the diffusion time for the continuous time dynamics is smaller, the number of integration steps is larger, inducing no computational gains. The optimization of the noise schedule can, however, have important practical effects in terms of stability of the training and variance of the estimations, which we do not tackle in this work [5].

2.5. Relation with Literature on Bounds and Goodness of Score Assumptions

Few other works in the literature attempt to study the convergence properties of Diffusion models. In the work of De Bortoli et al. [29] (Theorem 1), a total variation (TV) bound between the generated and data distribution is obtained in the form $C_1 \exp(a_1 T) +$ $C_2 \exp(-a_2 T)$, where the constant C_1 depends on the maximum error over [0, T] between the true and approximated score, i.e., $\max_{t \in [0,T]} \|s_{\theta}(x,t) - \nabla \log p(x,t)\|$. In the work of De Bortoli [30], the requirement is relaxed by setting $\max_{t \in [0,T]} \frac{\sigma^2(t)}{1+||x||} ||s_{\theta}(x,t) - \nabla \log p(x,t)||$, where the 1-Wasserstein distance between generated and true data is bounded as C_1 + $C_2 \exp(-a_2T) + C_3$ (Theorem 1). Other works consider the more realistic average square norm instead of the infinity norm, which is consistent with standard training of diffusion models. Moreover, Lee et al. [31] show how the TV bound can be expressed as a function of $\max_{t \in [0,T]} \mathbb{E} \left| \| \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t) - \boldsymbol{\nabla} \log p(\boldsymbol{x}_t, t) \|^2 \right|$ (Theorems 2.2, 3.1 and 3.2). Related to our work, Lee et al. [31] find that the TV bound is optimized for a diffusion time that depends, among others, on the maximum score error. Finally, the work by Chen et al. [34] (Theorem 2), which is concurrent to ours, shows that if $\max_{t \in [0,T]} \mathbb{E} \left| \| s_{\theta}(x_t, t) - \nabla \log p(x_t, t) \|^2 \right|$ is bounded, then the TV distance between true and generated data can be bound as $C_1 \exp(-a_1 T) + \sqrt{\epsilon T}$, plus a discretization error.

All prior approaches require assumptions on the maximum score error, which *implicitly* depends on: (i) the maximum diffusion time *T* and (ii) the class of parametric score networks considered. Hence, such methods allow for the study of convergence properties, but with the following limitations. It is not clear how the score error behaves as the fitting domain ([0, T]) is increased, for generic class of parametric functions and generic p_{data} . Moreover, it is difficult to link the error assumptions with the actual training loss of diffusion models. In this work, instead, we follow a more agnostic path, as we make no assumptions about the error behavior. We notice that the optimal gap term is **always** a non decreasing function of *T*. First, we question whether the current best practice for setting diffusion times is adequate: we find that, in realistic implementations, diffusion times are larger than necessary. Second, we introduce a new approach, with provably the same performance of standard diffusion models but lower computational complexity, as highlighted in Section 3.

3. A New, Practical Method for Decreasing Diffusion Times

The ELBO decomposition in Equation (9) and the bounds in Lemma 1 and Section 2.2 highlight a dilemma. We thus propose a simple method that allows us to achieve **both** a small gap $\mathcal{G}(s_{\theta}, T)$ and a small discrepancy $KL[p(x, T) \parallel p_{noise}(x)]$. Before that, let us use Figure 3 to summarize all densities involved and the effects of the various approximations, which will be useful to visualize our proposal.

The data distribution $p_{data}(\mathbf{x})$ is transformed into the noise distribution $p(\mathbf{x}, T)$ through the forward diffusion process. Ideally, starting from $p(\mathbf{x}, T)$, we can recover the data distribution by simulating using the exact score $\nabla \log p$. Using the approximated score s_{θ} and the same initial conditions, the backward process ends up in $q^{(1)}(\mathbf{x}, T)$, whose discrepancy (1) to $p_{data}(\mathbf{x})$ is $\mathcal{G}(s_{\theta}, T)$. However, the distribution $p(\mathbf{x}, T)$ is unknown and replaced with an easy distribution $p_{noise}(\mathbf{x})$, accounting for an error (1) measured as $KL[p(\mathbf{x}, T) \parallel p_{noise}(\mathbf{x})]$. With the score and initial distribution approximated, the backward process ends up in $q^{(3)}(\mathbf{x}, T)$, where the discrepancy (1) from p_{data} is the sum of the terms $\mathcal{G}(s_{\theta}, T) + KL[p(\mathbf{x}, T) \parallel p_{noise}]$.

Multiple bridges across densities. In a nutshell, our method allows us to reduce the gap term by selecting smaller diffusion times and by using a learned auxiliary model to transform the initial density $p_{noise}(x)$ into a density $v_{\phi}(x)$, which is as close as possible to p(x, T), thus avoiding the penalty of a large KL term. To implement this, we first *transform* the simple distribution p_{noise} into the distribution $v_{\phi}(x)$, whose discrepancy (b) KL $[p(x, T) \parallel v_{\phi}(x)]$ is smaller than (a). Then, starting from from the auxiliary model $v_{\phi}(x)$, we use the approximate score s_{θ} to simulate the backward process reaching $q^{(2)}(x, T)$. This solution has a discrepancy (c) from the data distribution of $\mathcal{G}(s_{\theta}, T) + \text{KL}[p(x, T) \parallel v_{\phi}(x)]$.

which we will quantify later in the section. Intuitively, we introduce two bridges. The first bridge connects the noise distribution p_{noise} to an auxiliary distribution $v_{\phi}(x)$ that is as close as possible to that obtained by the forward diffusion process. The second bridge—a standard reverse diffusion process—connects the smooth distribution $v_{\phi}(x)$ to the data distribution. Notably, our approach has important guarantees, which we discuss next.



Figure 3. Intuitive illustration of the forward and backward diffusion processes. Discrepancies between distributions are illustrated as distances. Color coding is discussed in the text.

3.1. Auxiliary Model Fitting and Guarantees

We begin by stating the requirements we consider for the density $v_{\phi}(x)$. First, as it is the case for p_{nvise} , it should be easy to generate samples from $v_{\phi}(x)$ in order to initialize the reverse diffusion process. Second, the auxiliary model should allow us to compute the likelihood of the samples generated through the overall generative process, which begins in p_{nvise} , passes through $v_{\phi}(x)$, and arrives in q(x, T).

The fitting procedure of the auxiliary model is straightforward. First, we recognize that minimizing $\operatorname{KL}[p(\mathbf{x},T) \parallel v_{\phi}(\mathbf{x})]$ with respect to ϕ also minimizes $\mathbb{E}_{p(\mathbf{x},T)}[\log v_{\phi}(\mathbf{x})]$, which we can use as loss function. To obtain the set of optimal parameters ϕ^* , we require samples from $p(\mathbf{x},T)$, which can be easily obtained even if the density $p(\mathbf{x},T)$ is not available. Indeed, by sampling from p_{data} , and $p(\mathbf{x},T \mid \mathbf{x}_0)$, we obtain an unbiased Monte Carlo estimate of $\mathbb{E}_{p(\mathbf{x},T)}[\log v_{\phi}(\mathbf{x})]$, and optimization of the loss can be performed. Note that, due to the affine nature of the drift, the conditional distribution $p(\mathbf{x},T \mid \mathbf{x}_0)$ is easy to sample from, as shown in Table 1. From a practical point of view, it is important to notice that the fitting of v_{ϕ} is independent from the training of the score-matching objective, i.e., the result of $I(\mathbf{s}_{\theta})$ does not depend on the shape of the auxiliary distribution v_{ϕ} . This implies that the two training procedures can be run in parallel, thus enabling considerable time savings.

Next, we show that the first bridge in our model reduces the KL term, even for small diffusion times.

Proposition 1. Let us assume that $p_{noise}(x)$ is in the family spanned by v_{ϕ} , i.e., there exists $\tilde{\phi}$ such that $v_{\tilde{\phi}} = p_{noise}$. Then we have that

$$\mathrm{KL}\left[p(\boldsymbol{x},T) \parallel v_{\boldsymbol{\phi}^*}(\boldsymbol{x})\right] \leq \mathrm{KL}\left[p(\boldsymbol{x},T) \parallel v_{\widetilde{\boldsymbol{\phi}}}(\boldsymbol{x})\right] = \mathrm{KL}\left[p(\boldsymbol{x},T) \parallel p_{\mathrm{noise}}(\boldsymbol{x})\right].$$
(10)

Since we introduce the auxiliary distribution ν , we shall define a new ELBO for our method:

$$\mathcal{L}^{\boldsymbol{\phi}}_{\text{ELBO}}(\boldsymbol{s}_{\boldsymbol{\theta}}, T) = \mathbb{E}_{p_{\text{data}}(\boldsymbol{x})} \log p_{\text{data}}(\boldsymbol{x}) - \mathcal{G}(\boldsymbol{s}_{\boldsymbol{\theta}}, T) - \text{KL}[p(\boldsymbol{x}, T) \parallel \nu_{\boldsymbol{\phi}}(\boldsymbol{x})]$$
(11)

Recalling that \hat{s}_{θ} is the optimal score for a generic time *T*, Proposition 1 allows us to claim that $\mathcal{L}_{\text{ELBO}}^{\phi^*}(\hat{s}_{\theta}, T) \geq \mathcal{L}_{\text{ELBO}}(\hat{s}_{\theta}, T)$.

Then, we can state the following important result:

Proposition 2. Given the existence of T^* , defined as the diffusion time such that the ELBO is maximized (Section 2.3), there exists at least one diffusion time $\tau \leq T^*$, such that $\mathcal{L}_{ELBO}^{\phi^*}(\widehat{s}_{\theta}, \tau) \geq \mathcal{L}_{ELBO}(\widehat{s}_{\theta}, T^*)$.

Proposition 2, which we prove in Appendix I, has two interpretations. On the one hand, given two score models optimally trained for their respective diffusion times, our approach guarantees an ELBO that is at least as good as that of a standard diffusion model configured with its optimal time T^* . Our method achieves this with a smaller diffusion time τ , which offers sampling efficiency and generation quality. On the other hand, if we settle for an equivalent ELBO for the standard diffusion model and our approach, with our method we can afford a sub-optimal score model, which requires a smaller computational budget to be trained, while guaranteeing shorter sampling times. We elaborate on this interpretation in Section 4, where our approach obtains substantial savings in terms of training iterations.

A final note is in order. The choice of the auxiliary model depends on the selected diffusion time. The larger the *T*, the "simpler" the auxiliary model can be. Indeed, the noise distribution $p(\mathbf{x}, T)$ approaches p_{noise} , so that a simple auxiliary model is sufficient to transform p_{noise} into a distribution v_{ϕ} . Instead, for a small *T*, the distribution $p(\mathbf{x}, T)$ is closer to the data distribution. Then, the auxiliary model requires high flexibility and capacity. In Section 4, we substantiate this discussion empirically on synthetic and real data.

3.2. Comparison with Schrödinger Bridges

In this section, we briefly compare our method with the Schrödinger bridges approach [29,35,36], which allows one to move from an arbitrary p_{noise} to p_{data} in any finite amount of time *T*. This is achieved by simulating the SDE

$$\mathrm{d}\boldsymbol{x}_t = \left[-f(\boldsymbol{x}_t, t') + g^2(t')\boldsymbol{\nabla}\log\hat{\psi}(\boldsymbol{x}_t, t')\right]\mathrm{d}t + g(t')\mathrm{d}\boldsymbol{w}_t, \quad \boldsymbol{x}_0 \sim p_{\text{noise}}, \tag{12}$$

where $\hat{\psi}, \psi$ solve the Partial Differential Equation (PDE) system

$$\begin{cases} \frac{\partial \psi(\mathbf{x},t)}{\partial t} = -\boldsymbol{\nabla}^{\top} (f(\mathbf{x},t)) \psi(\mathbf{x},t) - \frac{g^2(t)}{2} \Delta(\psi(\mathbf{x},t)), \\ \frac{\partial \hat{\psi}}{\partial t} = -\boldsymbol{\nabla}^{\top} (\hat{\psi}(\mathbf{x},t) f(\mathbf{x},t)) + \frac{g^2(t)}{2} \Delta(\hat{\psi}(\mathbf{x},t)), \end{cases}$$
(13)

with boundary conditions $\psi(\mathbf{x}, 0)\hat{\psi}(\mathbf{x}, 0) = p_{data}(\mathbf{x}), \psi(\mathbf{x}, T)\hat{\psi}(\mathbf{x}, T) = p_{noise}(\mathbf{x})$. In the above equation, $\nabla^{\top}(f(\mathbf{x}, t)) = \sum_{i=1}^{N} \frac{\partial f^{i}(\mathbf{x}, t)}{\partial x^{i}}$, being *N* the dimension of the vectors \mathbf{x} , f and the notation f^{i} , x^{i} indicating their i_{th} component. This approach presents drawbacks compared to classical Diffusion models. First, the functions ψ , $\hat{\psi}$ are not known, and their parametric approximation is costly and complex. Second, it is much harder to obtain quantitative bounds between true and generated data as a function of the quality of such approximations.

The $\hat{\psi}$, ψ estimation procedure simplifies considerably in the particular case where $p_{\text{noise}}(x) = p(x, T)$, for arbitrary *T*. The solution of Equation (13) is indeed $\psi(x, t) = 1$, $\hat{\psi}(x, t) = p(x, t)$. The first PDE of the system is satisfied when ψ is a constant. The second PDE is the Fokker–Planck equation, satisfied by $\hat{\psi}(x, t) = p(x, t)$. Boundary conditions are

also satisfied. In this scenario, a sensible objective is the score-matching, as getting $\nabla \log \hat{\psi}$ equal to the true score $\nabla \log p$ allows perfect generation.

Unfortunately, it is difficult to generate samples from $p(\mathbf{x}, T)$, the starting conditions of Equation (12). A trivial solution is to select $T \to \infty$ in order to have p_{noise} as the simple and analytically known steady state distribution of Equation (1). This corresponds to the classical diffusion models approach, which we discussed in the previous sections. An alternative solution is to keep *T* finite and *cover* the first part of the bridge from p_{noise} to $p(\mathbf{x}, T)$ with an auxiliary model. This provides a different interpretation of our method, which allows for smaller diffusion times while keeping good generative quality.

3.3. An Extension for Density Estimation

Diffusion models can be also used for density estimation by transforming the diffusion SDE into an equivalent Ordinary Differential Equation (ODE) whose marginal distribution $p(\mathbf{x}, t)$ at each time instant coincide to that of the corresponding SDE [3]. The exact equivalent ODE requires the score $\nabla \log p(\mathbf{x}_t, t)$, which in practice is replaced by the score model s_{θ} , leading to the following ODE

$$d\mathbf{x}_t = \left(f(\mathbf{x}_t, t) - \frac{1}{2}g(t)^2 \mathbf{s}_{\boldsymbol{\theta}}(\mathbf{x}_t, t)\right) dt \quad \text{with} \quad \mathbf{x}_0 \sim p_{data} , \qquad (14)$$

whose time varying probability density is indicated with $\tilde{p}(\mathbf{x}, t)$. Note that the density $\tilde{p}(\mathbf{x}, t)$, is in general **not** equal to the density $p(\mathbf{x}, t)$ associated to Equation (1), with the exception of perfect score matching [18]. The reverse time process is modeled as a Continuous Normalizing Flow (CNF) [37,38] initialized with distribution $p_{noise}(\mathbf{x})$; then, the likelihood of a given value \mathbf{x}_0 is

$$\log \widetilde{p}(\mathbf{x}_0) = \log p_{\text{noise}}(\mathbf{x}_T) + \int_{t=0}^{T} \nabla \cdot \left(f(\mathbf{x}_t, t) - \frac{1}{2}g(t)^2 s_{\theta}(\mathbf{x}_t, t) \right) \mathrm{d}t.$$
(15)

To use our proposed model for density estimation, we also need to take into account the ODE dynamics. We focus again on the term $\log p_{noise}(x_T)$ to improve the expected log likelihood. For consistency, our auxiliary density ν_{ϕ} should now maximize $\mathbb{E}_{\sim(14)} \log \nu_{\phi}(x_T)$ instead of $\mathbb{E}_{\sim(1)} \log \nu_{\phi}(x_T)$. However, the simulation of Equation (14) requires access to s_{θ} which, in the endeavor of density estimation, is available only once the score model has been trained. Consequently, optimization with respect to ϕ can only be performed sequentially, whereas, for generative purposes, it could be done concurrently. While the sequential version is expected to perform better, experimental evidence indicates that improvements are marginal, justifying the adoption of the more efficient concurrent version.

4. Experiments

We now present numerical results on the MNIST and CIFAR10 datasets, to support our claims in Sections 2 and 3. We follow a standard experimental setup [5,7,18,32]: we use a standard U-Net architecture with time embeddings [6] and we report the log-likelihood in terms of bit per dimension (BPD) and the Fréchet Inception Distance (FID) scores (uniquely for CIFAR10). Although the FID score is a standard metric for ranking generative models, caution should be used against over-interpreting FID improvements [39]. Similarly, while the theoretical properties of the models we consider are obtained through the lens of ELBO maximization, the log-likelihood measured in terms of BPD should be considered with care [40]. Finally, we also report the number of neural function evaluations (NFE) for computing the relevant metrics. We compare our method to the standard score-based model [3]. The full description on the experimental setup is presented in Appendix K.

On the existence of T^* . We look for further empirical evidence of the existence of a $T^* < \infty$, as stated in Section 2.3. For the moment, we shall focus on the baseline model [3], where no auxiliary models are introduced. Results are reported in Table 2. For

MNIST, we observe how times T = 0.6 and T = 1.0 have comparable performance in terms of BPD, implying that any $T \ge 1.0$ is at best unnecessary and generally detrimental. Similarly, for CIFAR10, it is possible to notice that the best value of BPD is achieved for T = 0.6, outperforming all other values.

Table 2. Optimal *T* in [3].

Dataset	Time T	BPD (↓)
	1.0	1.16
MUCT	0.6	1.16
MINIST	0.4	1.25
	0.2	1.75
	1.0	3.09
	0.6	3.07
CIFARIO	0.4	3.09
	0.2	3.38

Our auxiliary models. In Section 3, we introduced an auxiliary model to minimize the mismatch between initial distributions of the backward process. We now specify the family of parametric distributions we have considered. Clearly, the choice of an auxiliary model also depends on the data distribution, in addition to the choice of diffusion time *T*.

For our experiments, we consider two auxiliary models: (i) a Dirichlet process Gaussian mixture model (DPGMM) [41,42] for MNIST and (ii) Glow [43], a flexible normalizing flow for CIFAR10. Both of them satisfy our requirements: they allow exact likelihood computation and they are equipped with a simple sampling procedure. As discussed in Section 3, auxiliary model complexity should be adjusted as a function of *T*.

This is confirmed experimentally in Figure 4, where we use the number of mixture components of the DPGMM as a proxy to measure the complexity of the auxiliary model.



Figure 4. Complexity of the auxiliary model as function of diffusion time (reported median and 95 quantiles on 4 random seeds).

Reducing T with auxiliary models. We now show how it is possible to obtain a comparable (or better) performance than the baseline model for a wide range of diffusion times T. 0.4 produces good performance both in terms For MNIST, setting τ = of BPD (Table 3) and visual sample quality (Figure 5). We also consider the sequential extension (S) to compute the likelihood, but remark marginal improvements compared to a concurrent implementation. Similarly for the CIFAR10 dataset, in Table 4 we observe how our method achieves better BPD than the baseline diffusion for T = 1. Moreover, our approach outperforms the baselines for the corresponding diffusion time in terms of FID score (Figure 6 and additional non-curated samples in the Appendix K). In Figure A3 we provide a non curated subset of qualitative results, showing that our method for a diffusion time equal to 0.4 still produces appealing images, while the vanilla approach fails. We finally notice how the proposed method has comparable performance with regard to several other competitors, while stressing that many orthogonal to our solutions (like diffusion in latent space [4], or the selection of higher order schemes [22]) can actually be combined with our methodology.

	NFE(↓)	BPD ((↓)
Model	(ODE)		
ScoreSDE	300	1.16	
ScoreSDE ($T = 0.6$)	258	1.16	
Our $(T = 0.6)$	258	1.16	1.14 (S)
ScoreSDE ($T = 0.4$)	235	1.25	
Our $(T = 0.4)$	235	1.17	1.16 (S)
ScoreSDE ($T = 0.2$)	191	1.75	
Our $(T = 0.2)$	191	1.33	1.31 (S)

Table 3. Experiment results on MNIST. For our method, (S) is for the extension in Section 3.3.

Table 4. Experimental results on CIFAR10, including other relevant baselines and sampling efficiency enhancements from the literature.

	FID(↓)	BPD (\downarrow)	NFE (\downarrow)	NFE (\downarrow)
Model			(SDE)	(ODE)
ScoreSDE [3]	3.64	3.09	1000	221
ScoreSDE ($T = 0.6$)	5.74	3.07	600	200
ScoreSDE ($T = 0.4$)	24.91	3.09	400	187
ScoreSDE ($T = 0.2$)	339.72	3.38	200	176
Our $(T = 0.6)$	3.72	3.07	600	200
Our $(T = 0.4)$	5.44	3.06	400	187
Our $(T = 0.2)$	14.38	3.06	200	176
ARDM [44]	-	2.69	3072	
VDM [5]	4.0	2.49	1000	
D3PMs [21]	7.34	3.43	1000	
DDPM [6]	3.21	3.75	1000	
Gotta Go Fast [22]	2.44	-	180	
LSGM [4]	2.10	2.87	120/138	
ARDM-P [44]		2.68/2.74	200/50	

Training and sampling efficiency. In Figure 7, the horizontal line corresponds to the best performance of a fully trained baseline model for T = 1.0 [3]. To achieve the same performance of the baseline, variants of our method require fewer iterations, which translate in training efficiency. For the sake of fairness, the total training cost of our method should account for the auxiliary model training, which, however, can be done concurrently to the diffusion process. As an illustration for CIFAR10, using four GPUs, the baseline model requires ~6.4 days of training. With our method we trained the auxiliary and diffusion models for ~2.3 and 2 days, respectively, leading to a total training time of max{2.3,2} = 2.3 days. Similar training curves can be obtained for the MNIST dataset, where the training time for DPGMMs is negligible.

Sampling speed benefits are evident from Tables 3 and 4. When considering the SDE version of the methods the number of sampling steps can decrease linearly with T, in accordance with theory [45], while retaining good BPD and FID scores. Similarly, although not in a linear fashion, the number of steps of the ODE samplers can be reduced by using a smaller diffusion time T.

Finally, we test the proposed methodology on the more challenging CELEBA 64x64 dataset. In this case, we use a variance exploding diffusion and we consider again Glow as the auxiliary model. The results, presented in Table 5, report the log-likelihood performance of different methods (qualitative results are reported in Appendix K). On the two extremes of the complexity we have the original diffusion (VE, T = 1.0) with the best BPD and the highest complexity, and Glow which provides a much simpler scheme with worse performance. In the table we report the BPD and the NFE metrics for smaller diffusion times, in three different configurations: naively neglecting the mismatch (ScoreSDE) or using the auxiliary model (Our). Interestingly, we found that the best results are obtained by using a combination of diffusion models pretrained for T = 1.0. The summary of the

content of this table is the following: by accepting a small degradation in terms of BPD, we can reduce the computational cost by almost one order of magnitude. We think it would be interesting to study more performing auxiliary models to improve performance of our method on challenging datasets.

Table 5. Experimental results on CELEBA 64.

Model	BPD (↓) (ODE)	NFE (\downarrow)
ScoreSDE [3]	2.13	68
ScoreSDE ($T = 0.5$)	8.06	15
ScoreSDE ($T = 0.2$)	12.1	9
Our $(T = 0.5)$	2.48	16
Our $(T = 0.2)$	2.58	9
Our with pretrain diffusion ($T = 0.5$)	2.36	16
Our with pretrain diffusion ($T = 0.2$)	2.32	9
Glow [43]	3.74	1



Figure 5. Visualization of some samples. Top to Bottom: ScoreSDE [3] (T = 1, BPD = 1.16), ScoreSDE (T = 0.4, BPD = 1.25), **Our** (T = 0.4, BPD = 1.17)



Figure 6. Visualization of some samples on CIFAR10.



Figure 7. Training curves of score models for different diffusion time *T*, recorded during the span of 1.3 million iterations.

5. Conclusions

Diffusion-based generative models emerged as an extremely competitive approach for a wide range of application domains. In practice, however, these models are resourcehungry, both for their training and for sampling new data points. In this work, we have introduced the key idea of considering diffusion times *T* as a free variable which should be chosen appropriately. We have shown that the choice of *T* introduces a trade-off, for which a "sweet spot" exists. In standard diffusion-based models, smaller values of *T* are preferable for efficiency reasons, but sufficiently large *T* are required to reduce approximation errors of the forward dynamics. Thus, we devised a novel method that allows for an arbitrary selection of diffusion times, where even small values are allowed. Our method closes the gap between practical and ideal diffusion dynamics, using an auxiliary model. Our empirical validation indicated that the performance of our approach was comparable and often superior to standard diffusion models, while being efficient both in training and in sampling.

Author Contributions: Conceptualization, G.F., S.R., M.F. and P.M.; Methodology, G.F., S.R., M.F. and P.M.; Software, G.F., S.R., L.Y., A.F., D.R., M.F. and P.M.; Validation, G.F., S.R., M.F. and P.M.; Formal analysis, G.F., S.R., M.F. and P.M.; Investigation, G.F., S.R., M.F. and P.M.; Resources, L.Y., A.F. and D.R.; Writing—original draft, G.F., S.R., M.F. and P.M.; Writing—review & editing, G.F., S.R., L.Y., A.F., D.R., M.F. and agreed to the published version of the manuscript.

Funding: This work has been partially supported by the Netception project, funded by Huawei Paris.

Data Availability Statement: All used datasets are publicly available

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A. Generic Definitions and Assumptions

Our work builds upon the work in [18], which should be considered as a basis for the developments hereafter. In this supplementary material, we use the following shortened notation for a generic $\omega > 0$:

$$\mathcal{N}_{\omega}(\mathbf{x}) \stackrel{\text{\tiny def}}{=} \mathcal{N}(\mathbf{x}; \mathbf{0}, \omega \mathbf{I}). \tag{A1}$$

It is useful to notice that $\nabla \log(\mathcal{N}_{\omega}(\mathbf{x})) = -\frac{1}{\omega}\mathbf{x}$.

For an arbitrary probability density p(x), we define the convolution (* operator) with N_{ω} using notation

$$p_{\omega}(\mathbf{x}) = p(\mathbf{x}) * \mathcal{N}_{\omega}(\mathbf{x}). \tag{A2}$$

Equivalently, $p_{\omega}(\mathbf{x}) = \exp(\frac{\omega}{2}\Delta)p(\mathbf{x})$, and consequently $\frac{dp_{\omega}(\mathbf{x})}{d\omega} = \frac{1}{2}\Delta p(\mathbf{x})$, where $\Delta = \nabla^{\top}\nabla$. Notice that, by considering the Dirac delta function $\delta(\mathbf{x})$, we have the equality $\delta_{\omega}(\mathbf{x}) = \mathcal{N}_{\omega}(\mathbf{x})$.

In the following derivations, we make use of the Stam–Gross logarithmic Sobolev inequality result in [33] (p. 562 Example 21.3):

$$\operatorname{KL}[p(\mathbf{x}) \parallel \mathcal{N}_{\omega}(\mathbf{x})] = \int p(\mathbf{x}) \log\left(\frac{p(\mathbf{x})}{\mathcal{N}_{\omega}(\mathbf{x})}\right) d\mathbf{x} \leq \frac{\omega}{2} \int \left\| \nabla\left(\log\frac{p(\mathbf{x})}{\mathcal{N}_{\omega}(\mathbf{x})}\right) \right\|^2 p(\mathbf{x}) d\mathbf{x}.$$
(A3)

Appendix B. Deriving Equation (4) from [32]

We start with Equation (25) of [32], which, in our notation, reads

$$\log q(\mathbf{x}, T) \geq \mathbb{E}[\log p_{\text{noise}}(\mathbf{x}_T) | \mathbf{x}_0 = \mathbf{x}] - \int_0^T \mathbb{E}\left[\frac{1}{2}g^2(t) \|\mathbf{s}_{\theta}(\mathbf{x}_t)\|^2 + \boldsymbol{\nabla}^\top \left(g^2(t)\mathbf{s}_{\theta}(\mathbf{x}_t) - f(\mathbf{x}_t, t)\right) | \mathbf{x}_0 = \mathbf{x}\right] \mathrm{d}t.$$

The first step is to take the expected value with respect to $x_0 \sim p_{data}$ on both sides of the above inequality

$$\mathbb{E}_{p_{\textit{data}}}\left[\log q(\boldsymbol{x}, T)\right] \ge \mathbb{E}\left[\log p_{\textit{noise}}\left(\boldsymbol{x}_{T}\right)\right] - \int_{0}^{T} \mathbb{E}\left[\frac{1}{2}g^{2}(t)\|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_{t})\|^{2} + \boldsymbol{\nabla}^{\top}\left(g^{2}(t)\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_{t}) - \boldsymbol{f}(\boldsymbol{x}_{t}, t)\right)\right] \mathrm{d}t.$$
(A4)

We focus on rewriting the term

$$\begin{split} &\int_0^T \mathbb{E}\Big[\boldsymbol{\nabla}^\top \Big(g^2(t) \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t) - f(\boldsymbol{x}_t, t)\Big)\Big] \mathrm{d}t = \\ &\int_0^T p(\boldsymbol{x}, t | \boldsymbol{x}_0) p_{\text{data}}(\boldsymbol{x}_0) \boldsymbol{\nabla}^\top \Big(g^2(t) \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}) - f(\boldsymbol{x}, t)\Big) \mathrm{d}\boldsymbol{x} \mathrm{d}\boldsymbol{x}_0 \mathrm{d}t = \\ &- \int_0^T \boldsymbol{\nabla}^\top (p(\boldsymbol{x}, t | \boldsymbol{x}_0) p_{\text{data}}(\boldsymbol{x}_0)) \Big(g^2(t) \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}) - f(\boldsymbol{x}, t)\Big) \mathrm{d}\boldsymbol{x} \mathrm{d}\boldsymbol{x}_0 \mathrm{d}t = \\ &- \int_0^T (p(\boldsymbol{x}, t | \boldsymbol{x}_0) p_{\text{data}}(\boldsymbol{x}_0))^{-1} \boldsymbol{\nabla}^\top (p(\boldsymbol{x}, t | \boldsymbol{x}_0) p_{\text{data}}(\boldsymbol{x}_0)) \Big(g^2(t) \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}) - f(\boldsymbol{x}, t)\Big) \\ &(p(\boldsymbol{x}, t | \boldsymbol{x}_0) p_{\text{data}}(\boldsymbol{x}_0)) \mathrm{d}\boldsymbol{x} \mathrm{d}\boldsymbol{x}_0 \mathrm{d}t = \\ &- \int_0^T \boldsymbol{\nabla}^\top (\log(p(\boldsymbol{x}, t | \boldsymbol{x}_0)) + \log(p_{\text{data}}(\boldsymbol{x}_0))) \Big(g^2(t) \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}) - f(\boldsymbol{x}, t)\Big) \\ &(p(\boldsymbol{x}, t | \boldsymbol{x}_0) p_{\text{data}}(\boldsymbol{x}_0)) \mathrm{d}\boldsymbol{x} \mathrm{d}\boldsymbol{x}_0 \mathrm{d}t = \\ &- \int_0^T \boldsymbol{\nabla}^\top (\log(p(\boldsymbol{x}, t | \boldsymbol{x}_0))) \Big(g^2(t) \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}) - f(\boldsymbol{x}, t)\Big) (p(\boldsymbol{x}, t | \boldsymbol{x}_0) p_{\text{data}}(\boldsymbol{x}_0)) \mathrm{d}\boldsymbol{x} \mathrm{d}\boldsymbol{x}_0 \mathrm{d}t = \\ &- \int_0^T \boldsymbol{\nabla}^\top (\log(p(\boldsymbol{x}, t | \boldsymbol{x}_0))) \Big(g^2(t) \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}) - f(\boldsymbol{x}, t)\Big) (p(\boldsymbol{x}, t | \boldsymbol{x}_0) p_{\text{data}}(\boldsymbol{x}_0)) \mathrm{d}\boldsymbol{x} \mathrm{d}\boldsymbol{x}_0 \mathrm{d}t = \\ &- \int_0^T \boldsymbol{\nabla}^\top (\log(p(\boldsymbol{x}, t | \boldsymbol{x}_0))) \Big(g^2(t) \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}) - f(\boldsymbol{x}, t)\Big) (p(\boldsymbol{x}, t | \boldsymbol{x}_0) p_{\text{data}}(\boldsymbol{x}_0)) \mathrm{d}\boldsymbol{x} \mathrm{d}\boldsymbol{x}_0 \mathrm{d}t = \\ &- \int_0^T \boldsymbol{\nabla}^\top (\log(p(\boldsymbol{x}, t | \boldsymbol{x}_0))) \Big(g^2(t) \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}) - f(\boldsymbol{x}, t)\Big) (p(\boldsymbol{x}, t | \boldsymbol{x}_0) p_{\text{data}}(\boldsymbol{x}_0)) \mathrm{d}\boldsymbol{x} \mathrm{d}\boldsymbol{x}_0 \mathrm{d}t = \\ &- \int_0^T \boldsymbol{\nabla}^\top (\log(p(\boldsymbol{x}, t | \boldsymbol{x}_0))) \Big(g^2(t) \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}) - f(\boldsymbol{x}, t)\Big) \Big] \mathrm{d}t. \end{aligned}$$

Consequently, we can rewrite the r.h.s of Equation (A4) as

$$\begin{split} & \mathbb{E}[\log p_{\textit{noise}}\left(\mathbf{x}_{T}\right)] - \int_{0}^{T} \mathbb{E}\left[\frac{1}{2}g^{2}(t)\|\mathbf{s}_{\theta}(\mathbf{x}_{t})\|^{2} - g^{2}(t)\boldsymbol{\nabla}^{\top}(\log(p(\mathbf{x}_{t},t|\mathbf{x}_{0})))\mathbf{s}_{\theta}(\mathbf{x})\right] + \\ & \mathbb{E}\left[+g^{2}(t)\boldsymbol{\nabla}^{\top}(\log(p(\mathbf{x}_{t},t|\mathbf{x}_{0})))f(\mathbf{x},t)\right] \mathrm{d}t = \\ & \mathbb{E}[\log p_{\textit{noise}}\left(\mathbf{x}_{T}\right)] - \int_{0}^{T} \mathbb{E}\left[\frac{1}{2}g^{2}(t)\|\mathbf{s}_{\theta}(\mathbf{x}_{t}) - \boldsymbol{\nabla}(\log(p(\mathbf{x}_{t},t|\mathbf{x}_{0})))\|^{2}\right] \mathrm{d}t - \\ & \frac{1}{2}\int_{0}^{T} \mathbb{E}\left[g^{2}(t)\|\boldsymbol{\nabla}(\log(p(\mathbf{x}_{t},t|\mathbf{x}_{0})))\| + 2\boldsymbol{\nabla}^{\top}(\log(p(\mathbf{x}_{t},t|\mathbf{x}_{0})))f(\mathbf{x},t)\right] \mathrm{d}t, \end{split}$$

which is exactly Equation (4).

Appendix C. Proof of Equation (5)

We prove the following result

$$I(\boldsymbol{s}_{\boldsymbol{\theta}}, T) \geq \underbrace{I(\boldsymbol{\nabla}\log p, T)}_{\overset{\text{def}}{=} K(T)} = \frac{1}{2} \int_{t=0}^{T} g^{2}(t) \mathbb{E}_{\sim(1)} [\|\boldsymbol{\nabla}\log p(\boldsymbol{x}_{t}, t) - \boldsymbol{\nabla}\log p(\boldsymbol{x}_{t}, t \mid \boldsymbol{x}_{0})\|]^{2} \mathrm{d}t.$$

Proof. We prove that for generic positive $\lambda(\cdot)$, and $T_2 > T_1$ the following holds:

$$\int_{t=T_1}^{T_2} \lambda(t) \mathbf{E}_{\sim(1)} \Big[|| \boldsymbol{s}(\boldsymbol{x}_t, t) - \boldsymbol{\nabla} \log p(\boldsymbol{x}_t, t | \boldsymbol{x}_0) ||^2 \Big] dt \ge \int_{t=T_1}^{T_2} \lambda(t) \mathbf{E}_{\sim(1)} \Big[|| \boldsymbol{\nabla} \log p(\boldsymbol{x}_t, t) - \boldsymbol{\nabla} \log p(\boldsymbol{x}_t, t | \boldsymbol{x}_0) ||^2 \Big] dt.$$
(A5)

First, we compute the functional derivative (with regard to *s*)

$$\begin{split} &\frac{\delta}{\delta s} \int_{t=T_1}^{T_2} \lambda(t) \mathbf{E}_{\sim(1)} \Big[|| \boldsymbol{s}(\boldsymbol{x}_t, t) - \boldsymbol{\nabla} \log p(\boldsymbol{x}_t, t | \boldsymbol{x}_0) ||^2 \Big] \mathrm{d}t = \\ &2 \int_{t=T_1}^{T_2} \lambda(t) \mathbf{E}_{\sim(1)} [(\boldsymbol{s}(\boldsymbol{x}_t, t) - \boldsymbol{\nabla} \log p(\boldsymbol{x}_t, t | \boldsymbol{x}_0))] \mathrm{d}t = \\ &2 \int_{t=T_1}^{T_2} \lambda(t) \mathbf{E}_{\sim(1)} [(\boldsymbol{s}(\boldsymbol{x}_t, t) - \boldsymbol{\nabla} \log p(\boldsymbol{x}_t, t | \boldsymbol{x}_0))] \mathrm{d}t, \end{split}$$

where we used

Consequently, we can obtain the optimal *s* through

$$\frac{\delta}{\delta s} \int_{t=T_1}^{T_2} \lambda(t) \mathbb{E}_{\sim(1)} \Big[|| \boldsymbol{s}(\boldsymbol{x}_t, t) - \boldsymbol{\nabla} \log p(\boldsymbol{x}_t, t | \boldsymbol{x}_0) ||^2 \Big] \mathrm{d}t = 0 \to \boldsymbol{s}(\boldsymbol{x}, t) = \boldsymbol{\nabla} \log p(\boldsymbol{x}, t).$$
(A6)

Substitution of this result into Equation (A5) directly proves the desired inequality.

As a byproduct, we prove the correctness of Equation (5), since it is a particular case of Equation (A5), with $\lambda = g^2$, $T_1 = 0$, $T_2 = T$. Since K(T) is a minimum, the decomposition $I(s_{\theta}, T) = K(T) + \mathcal{G}(s_{\theta}, T)$ implies $K(T) + \mathcal{G}(s_{\theta}, T) \ge K(T) \rightarrow \mathcal{G}(s_{\theta}, T) \ge 0$. \Box

Appendix D. Proof of Equation (8)

Proof. We consider the pair of equations

$$d\mathbf{x}_{t} = \left[-f(\mathbf{x}_{t}, t') + g^{2}(t')\nabla \log q(\mathbf{x}_{t}, t)\right]dt + g(t')d\mathbf{w}(t),$$

$$d\mathbf{x}_{t} = f(\mathbf{x}_{t}, t)dt + g(t)d\mathbf{w}(t),$$
(A7)

where t' = T - t, q is the density of the backward process and p is the density of the forward process. These equations can be interpreted as a particular case of the following pair of SDEs (corresponding to [32] Equations (4) and (17) (Notice that our notation for the roles of p, q is swapped with respect to [32])).

$$d\mathbf{x}_{t} = \underbrace{\left[-f(\mathbf{x}_{t}, t') + g^{2}(t') \nabla \log q(\mathbf{x}_{t}, t)\right]}_{\mu(\mathbf{x}_{t}, t)} dt + \underbrace{g(t')}_{\sigma(t)} dw(t),$$

$$d\mathbf{x}_{t} = \left[\underbrace{f(\mathbf{x}_{t}, t) - g^{2}(t) \nabla \log q(\mathbf{x}_{t}, t')}_{-\mu(\mathbf{x}_{t}, t')} + \underbrace{g(t)}_{\sigma(t')} a(\mathbf{x}_{t}, t)\right] dt + g(t) dw(t), \quad (A8)$$

where Equation (A7) is recovered considering $a(x,t) = \sigma(t')\nabla \log q(x,t') = g(t)\nabla \log q(x,t')$. Equation (A8) is associated to an ELBO ([32], Theorem 3) that is attained

with equality if and only if $a(x, t) = \sigma(t')\nabla \log q(x, t')$. Consequently, we can write the following equality associated to the backward process of Equation (A7)

$$\log q(\mathbf{x}, T) = \mathbb{E}\left[-\frac{1}{2}\int_{0}^{T} ||\mathbf{a}(\mathbf{x}_{t}, t)||^{2} + 2\nabla^{\top} \mu(\mathbf{x}_{t}, t') ds + \log q(\mathbf{x}_{T}, 0) \quad \Big|\mathbf{x}_{0} = \mathbf{x}\right], \quad (A9)$$

where the expected value is taken with respect to the dynamics of the associated forward process.

By careful inspection of the couple of equations, we notice that, in the process x_t , the drift includes the $\nabla \log q(x_t, t)$ term, while in our main (1) we have $\nabla \log p(x_t, t')$. In general the two vector fields do not agree. However, if we select as starting distribution of the generating process p(x, T), i.e., q(x, 0) = p(x, T), then $\forall t, q(x, t) = p(x, t')$.

Given the initial conditions, the time evolution of the density p is fully described by the Fokker–Planck equation

$$\frac{d}{dt}p(\mathbf{x},t) = -\boldsymbol{\nabla}^{\top}(f(\mathbf{x},t)p(\mathbf{x},t)) + \frac{g^2(t)}{2}\Delta(p(\mathbf{x},t)), \quad p(\mathbf{x},0) = p_{data}(\mathbf{x}).$$
(A10)

Similarly, for the density *q*,

$$\frac{d}{dt}q(\mathbf{x},t) = -\boldsymbol{\nabla}^{\top} \left(-f(\mathbf{x},t')q(\mathbf{x},t) + g^2(t')\boldsymbol{\nabla}\log q(\mathbf{x},t)q(\mathbf{x},t) \right) + \frac{g^2(t')}{2}\Delta(q(\mathbf{x},t)),$$
(A11)

with q(x, 0) = p(x, T). By Taylor expansion we have

$$\begin{aligned} q(\mathbf{x},\delta t) &= q(\mathbf{x},0) + \delta t \left(\frac{d}{dt} q(\mathbf{x},t) \right)_{t=0} + \mathcal{O}(\delta t^2) = \\ q(\mathbf{x},0) + \delta t \left(-\boldsymbol{\nabla}^\top \left(-f(\mathbf{x},T)q(\mathbf{x},0) + g^2(T)\boldsymbol{\nabla}\log q(\mathbf{x},0)q(\mathbf{x},0) \right) + \frac{g^2(T)}{2}\Delta(q(\mathbf{x},0)) \right) + \mathcal{O}(\delta t^2) = \\ q(\mathbf{x},0) + \delta t \left(\boldsymbol{\nabla}^\top (f(\mathbf{x},T)q(\mathbf{x},0)) - \frac{g^2(T)}{2}\Delta(q(\mathbf{x},0)) \right) + \mathcal{O}(\delta t^2), \end{aligned}$$

and

$$p(\mathbf{x}, T - \delta t) = p(\mathbf{x}, T) - \delta t \left(\frac{d}{dt} p(\mathbf{x}, t) \right)_{t=T} + \mathcal{O}(\delta t^2) =$$

$$p(\mathbf{x}, T) - \delta t \left(-\nabla^\top (f(\mathbf{x}, T) p(\mathbf{x}, T)) + \frac{g^2(T)}{2} \Delta(p(\mathbf{x}, T)) \right) + \mathcal{O}(\delta t^2) =$$

$$p(\mathbf{x}, T) + \delta t \left(\nabla^\top (f(\mathbf{x}, T) p(\mathbf{x}, T)) - \frac{g^2(T)}{2} \Delta(p(\mathbf{x}, T)) \right) + \mathcal{O}(\delta t^2)$$

Since $q(\mathbf{x}, 0) = p(\mathbf{x}, T)$, we finally have $q(\mathbf{x}, \delta t) - p(\mathbf{x}, T - \delta t) = \mathcal{O}(\delta t^2)$. This holds for arbitrarily small δt . By induction, with similar reasoning, we claim that $q(\mathbf{x}, t) = p(\mathbf{x}, t')$. This last result allows us to rewrite Equation (A7) as the pair of SDEs

$$d\mathbf{x}_{t} = \left[-f(\mathbf{x}_{t}, t') + g^{2}(t')\nabla \log p(\mathbf{x}_{t}, t')\right]dt + g(t')dw(t),$$

$$d\mathbf{x}_{t} = f(\mathbf{x}_{t}, t)dt + g(t)dw(t).$$
(A12)

Moreover, since $q(x, T) = p(x, 0) = p_{data}(x)$, together with the result Equation (A9), we have the following equality

$$\log p_{data}(\mathbf{x}) = \mathbb{E}\left[-\frac{1}{2}\int_{0}^{T} ||\mathbf{a}(\mathbf{x}_{t},t)||^{2} + 2\nabla^{\top} \boldsymbol{\mu}(\mathbf{x}_{t},t') dt + \log p(\mathbf{x}_{T},T) \quad |\mathbf{x}_{0} = \mathbf{x}\right].$$
 (A13)

Consequently,

$$\mathbb{E}_{\mathbf{x}\sim p_{\text{data}}}\left[\log p_{\text{data}}(\mathbf{x})\right] = \mathbb{E}\left[\log p(\mathbf{x}_{T},T)\right] + \mathbb{E}\left[-\frac{1}{2}\int_{0}^{T}||\mathbf{a}(\mathbf{x}_{t},t)||^{2} + 2\nabla^{\top}\mu(\mathbf{x}_{t},t')dt\right] = \\\mathbb{E}\left[\log p(\mathbf{x}_{T},T)\right] + \mathbb{E}\left[-\frac{1}{2}\int_{0}^{T}g(t)^{2}||\nabla\log p(\mathbf{x}_{t},t)||^{2} + 2\nabla^{\top}\left(-f(\mathbf{x}_{t},t) + g^{2}(t)\nabla\log p(\mathbf{x}_{t},t)\right)dt\right] = \\\mathbb{E}\left[\log p(\mathbf{x}_{T},T)\right] + \mathbb{E}\left[-\frac{1}{2}\int_{0}^{T}g(t)^{2}||\nabla\log p(\mathbf{x}_{t},t)||^{2} - 2g^{2}(t)\nabla_{\mathbf{x}}^{\top}\log p(\mathbf{x}_{t},t)\nabla\log p(\mathbf{x}_{t},t|\mathbf{x}_{0})dt\right] \\ + \mathbb{E}\left[-\frac{1}{2}\int_{0}^{T}2f^{\top}(\mathbf{x}_{t},t)\nabla\log p(\mathbf{x}_{t},t|\mathbf{x}_{0})dt\right] =$$

$$\mathbb{E}[\log p(\mathbf{x}_T, T)] + \mathbb{E}\left[-\frac{1}{2}\int_0^T g(t)^2 ||\nabla \log p(\mathbf{x}_t, t) - \nabla \log p(\mathbf{x}_t, t|\mathbf{x}_0)||^2 dt\right] + \\\mathbb{E}\left[-\frac{1}{2}\int_0^T -g(t)^2 ||\nabla \log p(\mathbf{x}_t, t|\mathbf{x}_0)||^2 + 2f^\top(\mathbf{x}_t, t)\nabla \log p(\mathbf{x}_t, t|\mathbf{x}_0) dt\right].$$

Remembering the definitions

$$K(T) = \frac{1}{2} \int_{t=0}^{T} g^{2}(t) \mathbb{E}_{\sim(1)} [||\nabla \log p(\mathbf{x}_{t}, t) - \nabla \log p(\mathbf{x}_{t}, t | \mathbf{x}_{0})||]^{2} dt$$

$$R(T) = \frac{1}{2} \int_{t=0}^{T} \mathbb{E}_{\sim(1)} \Big[g^{2}(t) ||\nabla \log p(\mathbf{x}, t | \mathbf{x}_{0})|| \Big]^{2} - 2f^{\top}(\mathbf{x}, t) \nabla \log p(\mathbf{x}, t | \mathbf{x}_{0}) dt,$$

we finally conclude the proof that

$$\mathbb{E}_{\sim(1)}[\log p(\boldsymbol{x}_T, T)] - K(T) + R(T) = \mathbb{E}_{\boldsymbol{x} \sim p_{data}}[\log p_{data}(\boldsymbol{x})].$$
(A14)

Appendix E. Proof of Lemma 1

In this section, we prove the validity of Lemma 1 for the case of Variance Preserving (VP) and Variance Exploding (VE) SDEs. Remember that, as reported also in Table 1, the above mentioned classes correspond to $\alpha(t) = -\frac{1}{2}\beta(t), g(t) = \sqrt{\beta(t)}, \beta(t) = \beta_0 + (\beta_1 - \beta_0)t$ and $\alpha(t) = 0, g(t) = \sqrt{\frac{d\sigma^2(t)}{dt}}, \sigma^2(t) = \left(\frac{\sigma_{max}}{\sigma_{min}}\right)^t$, respectively.

Appendix E.1. The Variance Preserving (VP) Convergence

We associate this class of SDEs to the Fokker Planck operator

$$\mathcal{L}^{\dagger}(t) = \frac{1}{2}\beta(t)\boldsymbol{\nabla}^{\top}(\boldsymbol{x}\cdot+\boldsymbol{\nabla}(\cdot)), \qquad (A15)$$

and consequently $\frac{dp(\mathbf{x},t)}{dt} = \mathcal{L}^{\dagger}(t)p(\mathbf{x},t)$. Simple calculations show that $\lim_{T\to\infty} p(\mathbf{x},T) = \mathcal{N}_1(\mathbf{x})$.

We compute bound the time derivative of the KL term as

$$\frac{d}{dt} \operatorname{KL}[p(\mathbf{x},T) \parallel \mathcal{N}_{1}(\mathbf{x})] = \int \frac{dp(\mathbf{x},t)}{dt} \log\left(\frac{p(\mathbf{x},t)}{\mathcal{N}_{1}(\mathbf{x})}\right) d\mathbf{x} + \int \frac{p(\mathbf{x},t)}{p(\mathbf{x},t)} \frac{dp(\mathbf{x},t)}{dt} d\mathbf{x} =
\frac{1}{2}\beta(t) \int \nabla^{\top}(-\nabla \log(\mathcal{N}_{1}(\mathbf{x}))p(\mathbf{x},t)) + \nabla p(\mathbf{x},t))) \log\left(\frac{p(\mathbf{x},t)}{\mathcal{N}_{1}(\mathbf{x})}\right) d\mathbf{x} =
- \frac{1}{2}\beta(t) \int p(\mathbf{x},t)(-\nabla \log(\mathcal{N}_{1}(\mathbf{x})) + \nabla \log p(\mathbf{x},t)))^{\top} \nabla\left(\log\left(\frac{p(\mathbf{x},t)}{\mathcal{N}_{1}(\mathbf{x})}\right)\right) d\mathbf{x} =
- \frac{1}{2}\beta(t) \int p(\mathbf{x},t) \nabla\left(\log\left(\frac{p(\mathbf{x},t)}{\mathcal{N}_{1}(\mathbf{x})}\right)\right)^{\top} \nabla\left(\log\left(\frac{p(\mathbf{x},t)}{\mathcal{N}_{1}(\mathbf{x})}\right)\right) d\mathbf{x} =
- \frac{1}{2}\beta(t) \int p(\mathbf{x},t) ||\nabla\left(\log\left(\frac{p(\mathbf{x},t)}{\mathcal{N}_{1}(\mathbf{x})}\right)\right)||^{2} d\mathbf{x} =
- \frac{1}{2}\beta(t) \operatorname{KL}[p(\mathbf{x},T) \parallel \mathcal{N}_{1}(\mathbf{x})].$$
(A16)

We then apply Gronwall's inequality [33] to $\frac{d}{dt} \operatorname{KL}[p(\mathbf{x}, T) \parallel \mathcal{N}_1(\mathbf{x})] \le -\beta(t) \operatorname{KL}[p(\mathbf{x}, T) \parallel \mathcal{N}_1(\mathbf{x})]$ to claim

$$\operatorname{KL}[p(\boldsymbol{x},T) \parallel \mathcal{N}_{1}(\boldsymbol{x})] \leq \operatorname{KL}[p(\boldsymbol{x},0) \parallel \mathcal{N}_{1}(\boldsymbol{x})] \exp\left(-\int_{0}^{T} \beta(s) ds\right).$$
(A17)

To claim validity of the result, we need to assume that p(x, t) has finite first and second order derivatives, and that $KL[p(x, 0) \parallel N_1(x)] < \infty$.

Appendix E.2. The Variance Exploding (VE) Convergence

The first step is to bound the derivative with respect to to ω of the divergence $KL[p_{\omega}(\mathbf{x}) \parallel \mathcal{N}_{\omega}(\mathbf{x})]$, i.e.,

$$\frac{d}{d\omega} \operatorname{KL}[p_{\omega}(\mathbf{x}) \parallel \mathcal{N}_{\omega}(\mathbf{x})] = \int \frac{dp_{\omega}(\mathbf{x})}{d\omega} \log\left(\frac{p_{\omega}(\mathbf{x})}{\mathcal{N}_{\omega}(\mathbf{x})}\right) d\mathbf{x} + \int \frac{p_{\omega}(\mathbf{x})}{p_{\omega}(\mathbf{x})} \frac{dp_{\omega}(\mathbf{x})}{d\omega} d\mathbf{x} - \int \frac{p_{\omega}(\mathbf{x})}{\mathcal{N}_{\omega}(\mathbf{x})} \frac{d\mathcal{N}_{\omega}(\mathbf{x})}{d\omega} d\mathbf{x} = \frac{1}{2} \int (\Delta p_{\omega}(\mathbf{x})) \log\left(\frac{p_{\omega}(\mathbf{x})}{\mathcal{N}_{\omega}(\mathbf{x})}\right) - (\Delta \mathcal{N}_{\omega}(\mathbf{x})) \frac{p_{\omega}(\mathbf{x})}{\mathcal{N}_{\omega}(\mathbf{x})} d\mathbf{x} = \frac{1}{2} \int \nabla^{\top}(p_{\omega}(\mathbf{x})\nabla\log p_{\omega}(\mathbf{x})) \log\left(\frac{p_{\omega}(\mathbf{x})}{\mathcal{N}_{\omega}(\mathbf{x})}\right) - \nabla^{\top}(\mathcal{N}_{\omega}(\mathbf{x})\nabla\log \mathcal{N}_{\omega}(\mathbf{x})) \frac{p_{\omega}(\mathbf{x})}{\mathcal{N}_{\omega}(\mathbf{x})} d\mathbf{x} = -\frac{1}{2} \int (p_{\omega}(\mathbf{x})\nabla\log p_{\omega}(\mathbf{x}))^{\top}\nabla\left(\log\left(\frac{p_{\omega}(\mathbf{x})}{\mathcal{N}_{\omega}(\mathbf{x})}\right)\right) - (\mathcal{N}_{\omega}(\mathbf{x})\nabla\log \mathcal{N}_{\omega}(\mathbf{x}))^{\top}\nabla\left(\frac{p_{\omega}(\mathbf{x})}{\mathcal{N}_{\omega}(\mathbf{x})}\right) d\mathbf{x} = -\frac{1}{2} \int (p_{\omega}(\mathbf{x})\nabla\log p_{\omega}(\mathbf{x}))^{\top}\nabla\left(\log\left(\frac{p_{\omega}(\mathbf{x})}{\mathcal{N}_{\omega}(\mathbf{x})}\right)\right) - (p_{\omega}(\mathbf{x})\nabla\log \mathcal{N}_{\omega}(\mathbf{x}))^{\top}\nabla\left(\log\left(\frac{p_{\omega}(\mathbf{x})}{\mathcal{N}_{\omega}(\mathbf{x})}\right)\right) d\mathbf{x} = -\frac{1}{2} \int p_{\omega}(\mathbf{x})||\nabla\left(\log\left(\frac{p_{\omega}(\mathbf{x})}{\mathcal{N}_{\omega}(\mathbf{x})}\right)\right)||^{2}d\mathbf{x} \le -\frac{1}{\omega}\operatorname{KL}[p_{\omega}(\mathbf{x}) \parallel \mathcal{N}_{\omega}(\mathbf{x})]. \tag{A18}$$

Consequently, using again Gronwall inequality, for all $\omega_1 > \omega_0 > 0$ we have

$$\begin{split} & \operatorname{KL}[p_{\omega_1}(\boldsymbol{x}) \parallel \mathcal{N}_{\omega_1}(\boldsymbol{x})] \leq \operatorname{KL}[p_{\omega_0}(\boldsymbol{x}) \parallel \mathcal{N}_{\omega_0}(\boldsymbol{x})] \exp(-(\log(\omega_1) - \log(\omega_0))) = \\ & \operatorname{KL}[p_{\omega_0}(\boldsymbol{x}) \parallel \mathcal{N}_{\omega_0}(\boldsymbol{x})] \omega_0 \frac{1}{\omega_1}. \end{split}$$

This can be directly applied to obtain the bound for VE SDE. Consider $\omega_1 = \sigma^2(T) - \sigma^2(0)$ and $\omega_0 = \sigma^2(\tau) - \sigma^2(0)$ for an arbitrarily small $\tau < T$. Then, since for the considered class of variance exploding SDE we have $p(\mathbf{x}, T) = p_{\sigma^2(T) - \sigma^2(0)}(\mathbf{x})$

$$\operatorname{KL}\left[p(\boldsymbol{x},T) \parallel \mathcal{N}_{\sigma^{2}(T)-\sigma^{2}(0)}(\boldsymbol{x})\right] \leq C \frac{1}{\sigma^{2}(T)-\sigma^{2}(0)}$$
(A19)
where $C = \operatorname{KL}\left[p(\boldsymbol{x},\tau) \parallel \mathcal{N}_{\sigma^{2}(\tau)-\sigma^{2}(0)}(\boldsymbol{x})\right](\sigma^{2}(\tau)-\sigma^{2}(0)).$

Similarly to the previous case, we assume that p(x, t) has finite first and second order derivatives, and that $C < \infty$.

Appendix F. Proof for the Optimal Score Gap Term, Section 2.2

The optimal score gap term $\mathcal{G}(\widehat{s}_{\theta}, T)$ is a non-decreasing function in T. That is, given $T_2 > T_1$, and $\theta_1 = \arg \min_{\theta} I(s_{\theta}, T_1), \theta_2 = \arg \min_{\theta} I(s_{\theta}, T_2)$, then $\mathcal{G}(s_{\theta_2}, T_2) \ge \mathcal{G}(s_{\theta_1}, T_1)$.

Proof. For θ_1 defined as in the lemma, $I(s_{\theta_1}, T_1) = K(T_1) + \mathcal{G}(s_{\theta_1}, T_1)$. Next, select $T_2 > T_1$. Then, for a generic θ , including θ_2 ,

$$I(s_{\theta}, T_{2}) = \int_{t=0}^{T_{1}} g^{2}(t) \mathbb{E}_{\sim(1)} \Big[||s_{\theta}(x_{t}, t) - \nabla \log p(x_{t}, t|x_{0})||^{2} \Big] dt + \int_{=I(s_{\theta}, T_{1}) \ge K(T_{1}) + \mathcal{G}(s_{\theta_{1}}, T_{1}) = I(s_{\theta_{1}}, T_{1})} \int_{t=T_{1}}^{T_{2}} g^{2}(t) \mathbb{E}_{\sim(1)} \Big[||s_{\theta}(x_{t}, t) - \nabla \log p(x_{t}, t|x_{0})||^{2} \Big] dt \ge \mathcal{G}(s_{\theta_{1}}, T_{1}) + K(T_{2}),$$
$$\geq \int_{t=T_{1}}^{T_{2}} g^{2}(t) \mathbb{E}_{\sim(1)} [||\nabla \log p(x_{t}, t) - \nabla \log p(x_{t}, t|x_{0})||^{2}] dt = K(T_{2}) - K(T_{1})$$

from which $\mathcal{G}(s_{\theta}, T_2) = I(s_{\theta}, T_2) - K(T_2) \ge \mathcal{G}(s_{\theta_1}, T_1).$ \Box

Appendix G. Proof of Section 2.3

Consider the ELBO decomposition in Equation (9). We study it as a function of time T, and seek its optimal argument $T^* = \arg \max_T \mathcal{L}_{\text{ELBO}}(\widehat{s}_{\theta}, T)$. Then, the optimal diffusion time $T^* \in \mathbb{R}^+$, and thus not necessarily $T^* = \infty$. Additional assumptions on the gap term $\mathcal{G}(\cdot)$ can be used to guarantee strict finiteness of T^* .

Proof. It is trivial to verify that since the optimal gap term $\mathcal{G}(\hat{s}_{\theta}, T)$ is an increasing function in *T* Section 2.2, then $\frac{\partial \mathcal{G}}{\partial T} \geq 0$. Then, we study the sign of the KL derivative, which is always negative as shown by Equation (A16) and Equation (A18) (where we also notice $\frac{d}{dt} = \frac{d\omega}{dt} \frac{d}{dt} \frac{d}{d\omega}$ keep the sign). Moreover, we know that that $\lim_{T\to\infty} \frac{\partial KL}{\partial T} = 0$. Then, the function $\frac{\partial \mathcal{L}_{\text{ELBO}}}{\partial T} = \frac{\partial \mathcal{G}}{\partial T} + \frac{\partial KL}{\partial T}$ has at least one zero in $[0, \infty]$. \Box

Appendix H. Optimization of T^{\star}

It is possible to treat the diffusion time *T* as a hyper-parameter and perform gradientbased optimization jointly with the score model parameters θ . Indeed, simple calculations show that

$$\frac{\partial \mathcal{L}_{\text{ELBO}}(\boldsymbol{s}_{\boldsymbol{\theta}}, T)}{\partial T} = \mathbb{E}\Big[\Big(\boldsymbol{f}^{\top}(\boldsymbol{x}_{T}, T)\boldsymbol{\nabla} + \boldsymbol{g}^{2}(T)\Delta\Big)\log p_{\text{noise}}(\boldsymbol{x}_{T})\Big] +$$
(A20)

$$-\frac{1}{2}\mathbb{E}\Big[\|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_{T},T)-\boldsymbol{\nabla}\log p(\boldsymbol{x}_{T},T\,|\,\boldsymbol{x}_{0})\|^{2}\Big]+$$
(A21)

$$\frac{1}{2}\mathbb{E}\left[g^{2}(T)\|\boldsymbol{\nabla}\log p(\boldsymbol{x}_{T},T\mid\boldsymbol{x}_{0})\|^{2}-2\boldsymbol{f}^{\top}(\boldsymbol{x}_{T},T)\boldsymbol{\nabla}\log p(\boldsymbol{x}_{T},T\mid\boldsymbol{x}_{0})\right]$$
(A22)

Appendix I. Proof of Proposition 2

Proof. Since $\forall T$ we have $\mathcal{L}_{ELBO}^{\phi}(s_{\theta}, T) \geq \mathcal{L}_{ELBO}(s_{\theta}, T)$, there exists a countable set of intervals \mathcal{I} contained in $[0, T^*]$ of variable supports, where $\mathcal{L}_{ELBO}^{\phi}$ is greater than $\mathcal{L}_{ELBO}(s_{\theta}, T)$. Assuming continuity of $\mathcal{L}_{ELBO}^{\phi}$, in these intervals it is possible to find at least one $\tau \leq T^*$ where $\mathcal{L}_{ELBO}^{\phi^*}(\widehat{s}_{\theta}, \tau) \geq \mathcal{L}_{ELBO}(\widehat{s}_{\theta}, T^*)$. \Box We notice that the degenerate case $\mathcal{I} = T^*$ is obtained only when $\forall T \leq T^*$, $\operatorname{KL}[p(\mathbf{x}, T) \parallel v_{\phi^*}(\mathbf{x})] = \operatorname{KL}[p(\mathbf{x}, T) \parallel p_{\operatorname{noise}}(\mathbf{x})]$. We expect this condition to never occur in practice.

Appendix J. Invariance to Noise Schedule

We here discuss about the claims made in Section 2.4 about the invariance of the ELBO to the particular choice of noise schedule. First, in Appendix J.1, we explain how different SDEs corresponding to different noise schedules can be translated one into the other. We introduce the concept of the signal-to-noise ratio (SNR). We clarify the unified score parametrization used in practice in the literature [5,46]. Then, in Appendix J.2, we prove how the single elements of the ELBO depend only on the value of the SNR at the final diffusion time T, as claimed in the main paper.

Appendix J.1. Preliminaries

We consider as reference SDE a pure Wiener process diffusion,

$$dx_t = dw_t$$
 with $x_0 \sim p_{data}$, (A23)

It is easily seen that the solution of the random process admits representation

$$\mathbf{x}_t = \mathbf{x}_0 + \sqrt{t}\boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I})$$
 (A24)

In this case, the time varying probability density, which we indicate with ψ , satisfies

$$\psi(\mathbf{x},t) = \exp\left(\frac{t}{2}\Delta\right) p_{\text{data}}(\mathbf{x}), \quad \psi(\mathbf{x},t \mid \mathbf{x}_0) = \exp\left(\frac{t}{2}\Delta\right) \delta(\mathbf{x}-\mathbf{x}_0)$$
(A25)

Simple calculations show that

$$\nabla \log \psi(\mathbf{x}, \sigma^2) = \frac{\mathbb{E}[\mathbf{x}_0 \mid \mathbf{x}_0 + \sigma \boldsymbol{\epsilon} = \mathbf{x}] - \mathbf{x}}{\sigma^2} \stackrel{\text{def}}{=} \frac{d(\mathbf{x}; \sigma^2) - \mathbf{x}}{\sigma^2}, \quad (A26)$$

where again $x_0 \sim p_{data}$ and the function *d* can be interpreted as a *denoiser*.

Our goal is to show the relationship between equations like Equation (1), and Equation (A23). In particular, we focus on *affine* SDEs, as classically done with Diffusion models. The class of considered affine SDEs is the following:

$$d\mathbf{x}_t = \alpha(t)\mathbf{x}_t dt + g(t)d\mathbf{w}_t$$
 with $\mathbf{x}_0 \sim p_{data}$, (A27)

In this simple linear case the process admits representation

$$\mathbf{x}_t = k(t)\mathbf{x}_0 + \sigma(t)\boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I})$$
 (A28)

where $k(t) = \exp\left(\int_{0}^{t} \alpha(s)ds\right), \sigma^{2}(t) = k^{2}(t)\int_{0}^{t} \frac{g^{2}(s)}{k^{2}(s)}ds$. We can rewrite Equation (A28)

as $\mathbf{x}_t = k(t)(\mathbf{x}_0 + \tilde{\sigma}(t)\boldsymbol{\epsilon})$, and define the SNR as $\tilde{\sigma}(t) = \frac{\sigma(t)}{k(t)}$. The density associated to Equation (A27) can be expressed as a function of ψ as follows

$$p(\mathbf{x},t) = k(t)^{-D} \left[\exp\left(\frac{\tilde{\sigma}^2(t)}{2}\Delta\right) p_{\text{data}}(\mathbf{x}) \right]_{\frac{\mathbf{x}}{k(t)}} = k(t)^{-D} \psi(\frac{\mathbf{x}}{k(t)}, \tilde{\sigma}^2(t)).$$
(A29)

The score function associated to Equation (A28) has consequently expression

$$\boldsymbol{\nabla}_{\boldsymbol{x}} \log p(\boldsymbol{x}, t) = \boldsymbol{\nabla}_{\boldsymbol{x}} \log \psi(\frac{\boldsymbol{x}}{k(t)}, \tilde{\sigma}^2(t)) = \frac{1}{k(t)} \boldsymbol{\nabla}_{\frac{\boldsymbol{x}}{k(t)}} \log \psi(\frac{\boldsymbol{x}}{k(t)}, \tilde{\sigma}^2(t)) = \frac{k(t)\boldsymbol{d}(\frac{\boldsymbol{x}}{k(t)}; \tilde{\sigma}^2(t)) - \boldsymbol{x}}{\sigma^2(t)}.$$
 (A30)

Appendix J.2. Different Noise Schedules

Consider a diffusion of the form of Equation (A23) and a score network \bar{s}_{θ} that approximate the true score. Inspecting Equation (A30), we parametrize the score network associated to a generic diffusion Equation (A27) as a function of the score of the reference diffusion. The score parametrization considered in [5] can be generalized to arbitrary SDEs [46]. In particular, as suggested by Equation (A26), we select

$$\bar{s}_{\theta}(x,t) = \frac{k(t)d_{\theta}(\frac{x}{k(t)};\tilde{\sigma}^2(t)) - x}{\sigma^2(t)}$$
(A31)

We proceed by showing that the different components of the ELBO depends on the diffusion time *T* only through $\tilde{\sigma}(T)$, but not on k(t), $\sigma(t)$ singularly for any time t < T.

Theorem A1. Consider a generic diffusion Equation (A27) and parametrize the score network as $\bar{s}_{\theta}(\frac{x}{k(t)}, \tilde{\sigma}(t))$. Then, the gap term $\mathcal{G}(\bar{s}_{\theta}, T)$ associated to Equation (A27) for a diffusion time T depends only on $\tilde{\sigma}(T)$ but not on $k(t), \sigma(t)$ singularly for any time t < T.

Proof. We first rearrange the gap term

$$2\mathcal{G}(\bar{s}_{\theta}, T) = \int_{t=0}^{T} g^{2}(t) \mathbb{E}_{\sim(A27)} \Big[\|\bar{s}_{\theta}(\boldsymbol{x}_{t}, t) - \boldsymbol{\nabla} \log p(\boldsymbol{x}_{t}, t \mid \boldsymbol{x}_{0})\|^{2} \Big] dt - \int_{t=0}^{T} g^{2}(t) \mathbb{E}_{\sim(A27)} \Big[\|\boldsymbol{\nabla} \log p(\boldsymbol{x}_{t}, t) - \boldsymbol{\nabla} \log p(\boldsymbol{x}_{t}, t \mid \boldsymbol{x}_{0})\|^{2} \Big] dt = \int_{t=0}^{T} g^{2}(t) \mathbb{E}_{\sim(A27)} \Big[\|\bar{s}_{\theta}(\boldsymbol{x}_{t}, t) - \boldsymbol{\nabla} \log p(\boldsymbol{x}_{t}, t \mid \boldsymbol{x}_{0})\|^{2} \Big] dt$$

Then

Т

$$\begin{split} &\int_{t=0}^{T} g^{2}(t) \| \bar{\mathbf{s}}_{\theta}(\mathbf{x},t) - \nabla \log p(\mathbf{x},t) \|^{2} p(\mathbf{x},t \mid \mathbf{x}_{0}) p_{\text{data}}(\mathbf{x}_{0}) d\mathbf{x} d\mathbf{x}_{0} dt = \\ &\int_{t=0}^{T} g^{2}(t) \left\| \frac{k(t) d_{\theta}(\frac{\mathbf{x}}{k(t)}; \bar{\sigma}^{2}(t)) - \mathbf{x}}{\sigma^{2}(t)} - \frac{k(t) d_{\theta}(\frac{\mathbf{x}}{k(t)}; \bar{\sigma}^{2}(t)) - \mathbf{x}}{\sigma^{2}(t)} \right\|^{2} p(\mathbf{x},t \mid \mathbf{x}_{0}) p_{\text{data}}(\mathbf{x}_{0}) d\mathbf{x} d\mathbf{x}_{0} dt = \\ &\int_{t=0}^{T} g^{2}(t) \left\| \frac{k(t) d_{\theta}(\frac{\mathbf{x}}{k(t)}; \bar{\sigma}^{2}(t)) - k(t) d(\frac{\mathbf{x}}{k(t)}; \bar{\sigma}^{2}(t))}{\sigma^{2}(t)} \right\|^{2} p(\mathbf{x},t \mid \mathbf{x}_{0}) p_{\text{data}}(\mathbf{x}_{0}) d\mathbf{x} d\mathbf{x}_{0} dt = \\ &\int_{t=0}^{T} \frac{g^{2}(t)}{k^{2}(t)} \left\| \frac{d_{\theta}(\frac{\mathbf{x}}{k(t)}; \bar{\sigma}^{2}(t)) - d(\frac{\mathbf{x}}{k(t)}; \bar{\sigma}^{2}(t))}{\bar{\sigma}^{2}(t)} \right\|^{2} p(\mathbf{x},t \mid \mathbf{x}_{0}) p_{\text{data}}(\mathbf{x}_{0}) d\mathbf{x} d\mathbf{x}_{0} dt = \\ &\int_{t=0}^{T} \frac{g^{2}(t)}{k^{2}(t)} \left\| \frac{d_{\theta}(\frac{\mathbf{x}}{k(t)}; \bar{\sigma}^{2}(t)) - d(\frac{\mathbf{x}}{k(t)}; \bar{\sigma}^{2}(t))}{\bar{\sigma}^{2}(t)} \right\|^{2} \psi(\frac{\mathbf{x}}{k(t)}, \bar{\sigma}^{2}(t) \mid \mathbf{x}_{0}) p_{\text{data}}(\mathbf{x}_{0}) k(t)^{-D} d\mathbf{x} d\mathbf{x}_{0} dt = \\ &\text{subst.} \quad \tilde{\mathbf{x}} = \frac{\mathbf{x}}{k(t)}, \quad d\tilde{\mathbf{x}} = d\mathbf{x} k^{-D}(t) \\ &\int_{t=0}^{T} \frac{g^{2}(t)}{k^{2}(t)} \left\| \frac{d_{\theta}(\tilde{\mathbf{x}}; \bar{\sigma}^{2}(t)) - d(\tilde{\mathbf{x}}; \bar{\sigma}^{2}(t))}{\bar{\sigma}^{2}(t)} \right\|^{2} \psi(\tilde{\mathbf{x}}, \bar{\sigma}^{2}(t) \mid \mathbf{x}_{0}) p_{\text{data}}(\mathbf{x}_{0}) d\mathbf{x} d\mathbf{x}_{0} dt = \\ &\text{subst.} \quad r = \tilde{\sigma}^{2}(t), \quad dr = \frac{g^{2}(t)}{k^{2}(t)} dt \end{aligned}$$

$$\int_{t=0}^{\tilde{\sigma}^{2}(T)} \|\bar{\boldsymbol{s}}_{\boldsymbol{\theta}}(\tilde{\boldsymbol{x}}, r) - \boldsymbol{\nabla} \log \psi(\tilde{\boldsymbol{x}}, r \,|\, \boldsymbol{x}_{0})\|^{2} \psi(\tilde{\boldsymbol{x}}, r) p_{\text{data}}(\boldsymbol{x}_{0}) \mathrm{d}\tilde{\boldsymbol{x}} \mathrm{d}\boldsymbol{x}_{0} \mathrm{d}r$$

For any k(t), $\sigma(t)$ such that $\tilde{\sigma}(T)$ is the same, the score matching loss is the same. \Box

Theorem A2. Suppose that for any ϕ of the auxiliary model $v_{\phi}(\mathbf{x})$ it exists one ϕ' such that $v_{\phi'}(\mathbf{x}) = k^{-D}v_{\phi}(\frac{\mathbf{x}}{k})$, for any k > 0. Notice that this condition is trivially satisfied if the considered parametric model has the expressiveness to multiply its output by the scalar k. Then the minimum Kullback–Leibler divergence between $p(\mathbf{x}, T)$ associated to a generic diffusion Equation (A27) and the density of an auxiliary model $v_{\phi}(\mathbf{x})$ depends only on $\tilde{\sigma}(T)$ and not on $\sigma(T)$ alone.

Proof. We start with the equality

$$\begin{split} & \operatorname{KL}\left[p(\mathbf{x},T) \parallel \nu_{\boldsymbol{\phi}}(\mathbf{x})\right] = \operatorname{KL}\left[k(T)^{-D}\psi(\frac{\mathbf{x}}{k(T)},\tilde{\sigma}(T)) \parallel \nu_{\boldsymbol{\phi}}(\mathbf{x})\right] = \\ & \operatorname{KL}\left[k(T)^{-D}\psi(\frac{\mathbf{x}}{k(T)},\tilde{\sigma}(T)) \parallel k(T)^{-D}\nu_{\boldsymbol{\phi}'}(\frac{\mathbf{x}}{k(T)})\right] = \\ & \int k(T)^{-D}\psi(\frac{\mathbf{x}}{k(T)},\tilde{\sigma}(T)) \log\left(\frac{\psi(\frac{\mathbf{x}}{k(T)},\tilde{\sigma}(T))}{\nu_{\boldsymbol{\phi}'}(\frac{\mathbf{x}}{k(T)})}\right) \mathrm{d}\mathbf{x} = \int \psi(\tilde{\mathbf{x}},\tilde{\sigma}(T)) \log\left(\frac{\psi(\tilde{\mathbf{x}},\tilde{\sigma}(T))}{\nu_{\boldsymbol{\phi}'}(\tilde{\mathbf{x}})}\right) \mathrm{d}\tilde{\mathbf{x}} = \\ & \operatorname{KL}\left[\psi(\mathbf{x},\tilde{\sigma}(T)) \parallel \nu_{\boldsymbol{\phi}'}(\mathbf{x})\right] \end{split}$$

Then the minimimum only depends on $\tilde{\sigma}(T)$, as it is always possible to achieve the same value independently on the SDE by rescaling the auxiliary model output. \Box

Appendix K. Experimental Details

We here give some additional details concerning the experimental (Section 4) settings.

Appendix K.1. Toy Example Details

In the toy example, we use 8192 samples from a simple Gaussian mixture with two components as target $p_{data}(\mathbf{x})$. In detail, we have $p_{data}(\mathbf{x}) = \pi \mathcal{N}(1, 0.1^2) + (1 - \pi)\mathcal{N}(3, 0.5^2)$, with $\pi = 0.3$. The choice of Gaussian mixture allows to write down explicitly the time-varying density

$$p(\mathbf{x}_t, t) = \pi \mathcal{N}(1, s^2(t) + 0.1^2) + (1 - \pi) \mathcal{N}(3, s^2(t) + 0.5^2),$$
(A32)

where $s^2(t)$ is the marginal variance of the process at time *t*. We consider a variance exploding SDE of the type $dx_t = \sigma^t dw_t$, which corresponds to $s^2(t) = \frac{\sigma^{2t}-1}{2\log\sigma}$.



Figure A1. Visualization of few samples at different diffusion times *T*.

Appendix K.2. Section 4 Details

We considered Variance Preserving SDE with default β_0 , β_1 parameter settings. When experimenting on CIFAR10 we considered the NCSN++ architecture as implemented

in [3]. Training of the score matching network has been carried out with the default set of optimizers and schedulers of [3], independently of the selected *T*.

For the MNIST dataset we reduced the architecture by considering 64 features, $ch_mult = (1, 2)$ and attention resolutions equal to 8. The optimizer has been selected as the one in the CIFAR10 experiment but the warmup has been reduced to 1000 and the total number of iterations to 65,000.

Appendix K.3. Varying T

We clarify about the *T* truncation procedure during both training and testing. The SDE parameters are kept unchanged irrespective of *T*. During training, as evident from Equation (3), it is sufficient to sample randomly the diffusion time from distribution U(0, T), where *T* can take any positive value. For testing (sampling), we simply modified the algorithmic routines to begin the reverse diffusion processes from a generic *T* instead of the default 1.0.

Appendix L. Non-Curated Samples

We provide for completeness collection of non-curated samples for the CIFAR10 (Figures A2–A5), MNIST dataset (Figures A6–A9) and CELEBA dataset (Figure A10 and Table 5).



Figure A2. CIFAR10: Our method (left) and the Vanilla method (right) at T = 0.2.



Figure A3. CIFAR10: Our method (left) and the Vanilla method (right) at T = 0.4.



Figure A4. CIFAR10: Our method (**left**) and the Vanilla method (**right**) at T = 0.6.



Figure A5. Vanilla method at T = 1.0.

Š Ë Ş 38 Z S. 30 P 3 8 G mit. Z Ĩ Ŗ Ľ Ğ Z ₹¥ 57 3 0 12 14 ろう 7 ま લુગ あ B Ŕ Ş Ť. 320 100 ゆ 64 Ċ ij 19 Q Ħ ¢\$ 5 440 5. S.S. 1 3 Z B Ŋ R Ż S 2 劣 P Ŷ Ċ, R g S ÷ S. 4 K Ŋ, Z. ery. C a Ŷ Y 3 **5**2; Ş ŝ Ľ 44 ŝ Ì 3 夸 Q, HO ୍ଯ ちゃ 19 3 È 認 50 R Ŋ Y'A Ş 3 野 R 4 Ř 5 ويلا 45 γ_{π} 4 89 ଟ୍ଟ 8 ¥ G. 3 ష Ŷ ₹¥ Ľ な ¥ Ð B 老 窃 Ĥ Ĉ, 9.H s' S ð دمل Ð b ₽¶• 3)

Figure A6. MNIST: Our method (**left**) and the Vanilla method (**right**) at T = 0.2.

G ન ጓ テグ \$. 2.5 ら Q \$ а ĥ F q ઝ ij в 5) Ŷ Я シ ¥ Ø Q Э P Ο.-~ ζ P e + ¥ в Э q A જ (P > С X Ð Э S ð J З ç -3 **Figure A7.** MNIST: Our method (left) and the Vanilla method (right) at T = 0.4. Ð L а ч わ C З z Ь B Ъ රී ລ Ø \mathcal{L} Н \mathcal{O} گ ア I ð て Э Б Ц О J æ H a З 7.2 Ъ Š (a)S Г З \mathcal{O} ¥ Ĉ. u Ç \boldsymbol{z} с, Ч \circ Ð З B З Г ß Ł Ď G -

Figure A8. MNIST: Our method (**left**) and the Vanilla method (**right**) at T = 0.6.

Figure A9. MNIST: Vanilla method at T = 1.0.



Figure A10. CELEBA images. **Top Left**: our method with pretrained score model and Glow (T = 0.2), **Top Right**: our method with pretrained score model and Glow (T = 0.5) and **Bottom Left**: baseline diffusion (T = 1.0). **Bottom Right**: FID scores for our method and baseline (T = 1.0).

References

- 1. Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; Ganguli, S. Deep unsupervised learning using nonequilibrium thermodynamics. In Proceedings of the International Conference on Machine Learning, Lille, France, 6–11 July 2015.
- Song, Y.; Ermon, S. Generative Modeling by Estimating Gradients of the Data Distribution. In Proceedings of the Advances in Neural Information Processing Systems, Vancouver, BC, Canada, 8–14 December 2019.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D.P.; Kumar, A.; Ermon, S.; Poole, B. Score-Based Generative Modeling through Stochastic Differential Equations. In Proceedings of the International Conference on Learning Representations, Virtual, 30 April–3 May 2021.
- 4. Vahdat, A.; Kreis, K.; Kautz, J. Score-based Generative Modeling in Latent Space. In Proceedings of the Advances in Neural Information Processing Systems, Virtual, 6–14 December 2021.
- Kingma, D.; Salimans, T.; Poole, B.; Ho, J. Variational Diffusion Models. In Proceedings of the Advances in Neural Information Processing Systems, Virtual, 6–14 December 2021.
- 6. Ho, J.; Jain, A.; Abbeel, P. Denoising Diffusion Probabilistic Models. In Proceedings of the Advances in Neural Information Processing Systems, Online, 6–12 December 2020.
- 7. Song, J.; Meng, C.; Ermon, S. Denoising Diffusion Implicit Models. In Proceedings of the International Conference on Learning Representations, Virtual, 30 April–3 May 2021.

- Kong, Z.; Ping, W.; Huang, J.; Zhao, K.; Catanzaro, B. DiffWave: A Versatile Diffusion Model for Audio Synthesis. In Proceedings of the International Conference on Learning Representations, Virtual, 30 April–3 May 2021.
- Lee, S.g.; Kim, H.; Shin, C.; Tan, X.; Liu, C.; Meng, Q.; Qin, T.; Chen, W.; Yoon, S.; Liu, T.Y. PriorGrad: Improving Conditional Denoising Diffusion Models with Data-Dependent Adaptive Prior. In Proceedings of the International Conference on Learning Representations, Virtual, 25–29 April 2022.
- 10. Dhariwal, P.; Nichol, A. Diffusion Models Beat GANs on Image Synthesis. In Proceedings of the Advances in Neural Information Processing Systems, Virtual, 6–14 December 2021.
- 11. Nichol, A.Q.; Dhariwal, P. Improved Denoising Diffusion Probabilistic Models. In Proceedings of the International Conference on Machine Learning, Virtual, 18–24 July 2021.
- 12. Tashiro, Y.; Song, J.; Song, Y.; Ermon, S. CSDI: Conditional Score-based Diffusion Models for Probabilistic Time Series Imputation. In Proceedings of the Advances in Neural Information Processing Systems, Virtual, 6–14 December 2021.
- 13. Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; Bengio, Y. Generative Adversarial Nets. In Proceedings of the Advances in Neural Information Processing Systems, Montreal, QC, Canada, 8–13 December 2014.
- Kingma, D.P.; Salimans, T.; Jozefowicz, R.; Chen, X.; Sutskever, I.; Welling, M. Improved Variational Inference with Inverse Autoregressive Flow. In Proceedings of the Advances in Neural Information Processing Systems 29, Barcelona, Spain, 5–10 December 2016.
- 15. Kingma, D.P.; Welling, M. Auto-Encoding Variational Bayes. In Proceedings of the International Conference on Learning Representations, Banff, AB, Canada, 14–16 April 2014.
- Tran, B.H.; Rossi, S.; Milios, D.; Michiardi, P.; Bonilla, E.V.; Filippone, M. Model Selection for Bayesian Autoencoders. In Proceedings of the Advances in Neural Information Processing Systems, Virtual, 6–14 December 2021.
- 17. Anderson, B.D. Reverse-Time Diffusion Equation Models. Stoch. Process. Their Appl. 1982, 12, 313–326. [CrossRef]
- Song, Y.; Durkan, C.; Murray, I.; Ermon, S. Maximum Likelihood Training of Score-Based Diffusion Models. In Proceedings of the Advances in Neural Information Processing Systems, Virtual, 6–14 December 2021.
- 19. Särkkä, S.; Solin, A. *Applied Stochastic Differential Equations*; Institute of Mathematical Statistics Textbooks, Cambridge University Press: Cambridge, UK, 2019. [CrossRef]
- Zheng, H.; He, P.; Chen, W.; Zhou, M. Truncated Diffusion Probabilistic Models. CoRR 2022. abs/2202.09671. Available online: http://xxx.lanl.gov/abs/2202.09671 (accessed on 31 March 2023).
- Austin, J.; Johnson, D.D.; Ho, J.; Tarlow, D.; van den Berg, R. Structured Denoising Diffusion Models in Discrete State-Spaces. In Proceedings of the Advances in Neural Information Processing Systems, Virtual, 6–14 December 2021.
- Jolicoeur-Martineau, A.; Li, K.; Piché-Taillefer, R.; Kachman, T.; Mitliagkas, I. Gotta Go Fast When Generating Data with Score-Based Models. CoRR 2021. abs/2105.14080. Available online: http://xxx.lanl.gov/abs/2105.14080 (accessed on 31 March 2023).
- Salimans, T.; Ho, J. Progressive Distillation for Fast Sampling of Diffusion Models. In Proceedings of the International Conference on Learning Representations, Virtual, 25–29 April 2022.
- 24. Xiao, Z.; Kreis, K.; Vahdat, A. Tackling the Generative Learning Trilemma with Denoising Diffusion GANs. In Proceedings of the International Conference on Learning Representations, Virtual, 25–29 April 2022.
- Watson, D.; Ho, J.; Norouzi, M.; Chan, W. Learning to Efficiently Sample from Diffusion Probabilistic Models. CoRR 2021. abs/2106.03802. Available online: http://xxx.lanl.gov/abs/2106.03802 (accessed on 31 March 2023).
- Dockhorn, T.; Vahdat, A.; Kreis, K. Score-Based Generative Modeling with Critically-Damped Langevin Diffusion. In Proceedings
 of the International Conference on Learning Representations, Virtual, 25–29 April 2022.
- 27. Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; Ommer, B. High-Resolution Image Synthesis with Latent Diffusion Models. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, New Orleans, LA, USA, 19–20 June 2022.
- Bao, F.; Li, C.; Zhu, J.; Zhang, B. Analytic-DPM: an Analytic Estimate of the Optimal Reverse Variance in Diffusion Probabilistic Models. In Proceedings of the International Conference on Learning Representations, Virtual, 25–29 April 2022.
- 29. De Bortoli, V.; Thornton, J.; Heng, J.; Doucet, A. Diffusion Schrödinger Bridge with Applications to Score-Based Generative Modeling. In Proceedings of the Advances in Neural Information Processing Systems, Virtual, 6–14 December 2021.
- 30. De Bortoli, V. Convergence of denoising diffusion models under the manifold hypothesis. *arXiv* **2022**, arXiv:2208.05314.
- 31. Lee, H.; Lu, J.; Tan, Y. Convergence for score-based generative modeling with polynomial complexity. *arXiv* 2022, arXiv:2206.06227.
- Huang, C.W.; Lim, J.H.; Courville, A.C. A Variational Perspective on Diffusion-Based Generative Models and Score Matching. In Proceedings of the Advances in Neural Information Processing Systems, Virtual, 6–14 December 2021.
- 33. Villani, C. Optimal Transport: Old and New; Springer: Berlin/Heidelberg, Germany, 2009; Volume 338.
- 34. Chen, S.; Chewi, S.; Li, J.; Li, Y.; Salim, A.; Zhang, A.R. Sampling is as easy as learning the score: theory for diffusion models with minimal data assumptions. *arXiv* **2022**, arXiv:2209.11215.
- 35. Chen, Y.; Georgiou, T.T.; Pavon, M. Stochastic control liaisons: Richard sinkhorn meets gaspard monge on a schrodinger bridge. *SIAM Rev.* 2021, *63*, 249–313. [CrossRef]
- Chen, T.; Liu, G.H.; Theodorou, E. Likelihood Training of Schrödinger Bridge using Forward-Backward SDEs Theory. In Proceedings of the International Conference on Learning Representations, Virtual, 25–29 April 2022.
- 37. Chen, R.T.Q.; Rubanova, Y.; Bettencourt, J.; Duvenaud, D.K. Neural Ordinary Differential Equations. In Proceedings of the Advances in Neural Information Processing Systems, Montreal, QC, Canada, 3–8 December 2018.

- Grathwohl, W.; Chen, R.T.Q.; Bettencourt, J.; Duvenaud, D. Scalable Reversible Generative Models with Free-form Continuous Dynamics. In Proceedings of the International Conference on Learning Representations, New Orleans, LA, USA, 6–9 May 2019.
 Kynkäänniemi, T.; Karras, T.; Aittala, M.; Aila, T.; Lehtinen, J. The Role of ImageNet Classes in Fréchet Inception Distance. CoRR
- 2022. abs/2203.06026. Available online: http://xxx.lanl.gov/abs/2203.06026 (accessed on 31 March 2023).
- 40. Theis, L.; van den Oord, A.; Bethge, M. A Note on the Evaluation of Generative Models. In Proceedings of the International Conference on Learning Representations, San Juan, Puerto Rico, 2–4 May 2016.
- 41. Rasmussen, C. The Infinite Gaussian Mixture Model. In Proceedings of the Advances in Neural Information Processing Systems, Denver, CO, USA, 29 November–4 December 1999.
- 42. Görür, D.; Edward Rasmussen, C. Dirichlet Process Gaussian Mixture Models: Choice of the Base Distribution. *J. Comput. Sci. Technol.* **2010**, *25*, 653–664. [CrossRef]
- 43. Kingma, D.P.; Dhariwal, P. Glow: Generative Flow with Invertible 1 × 1 Convolutions. In Proceedings of the Advances in Neural Information Processing Systems, Montreal, QC, Canada, 3–8 December 2018.
- Hoogeboom, E.; Gritsenko, A.A.; Bastings, J.; Poole, B.; van den Berg, R.; Salimans, T. Autoregressive Diffusion Models. In Proceedings of the International Conference on Learning Representations, Virtual, 25–29 April 2022.
- 45. Kloeden, P.E.; Platen, E. Numerical Solution of Stochastic Differential Equations; Springer: Berlin/Heidelberg, Germany, 1995.
- 46. Karras, T.; Aittala, M.; Aila, T.; Laine, S. Elucidating the Design Space of Diffusion-Based Generative Models. *arXiv* 2022, arXiv:2206.00364.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.