
On Calibrating Diffusion Probabilistic Models

Tianyu Pang^{†1}, Cheng Lu², Chao Du¹, Min Lin¹, Shuicheng Yan¹, Zhijie Deng^{†3}

¹Sea AI Lab, Singapore

²Department of Computer Science, Tsinghua University

³Qing Yuan Research Institute, Shanghai Jiao Tong University

{tianyupang, duchao, linmin, yansc}@sea.com;

lucheng.lc15@gmail.com; zhijied@sjtu.edu.cn

Abstract

Recently, diffusion probabilistic models (DPMs) have achieved promising results in diverse generative tasks. A typical DPM framework includes a forward process that gradually diffuses the data distribution and a reverse process that recovers the data distribution from time-dependent data scores. In this work, we observe that the stochastic reverse process of data scores is a martingale, from which concentration bounds and the optional stopping theorem for data scores can be derived. Then, we discover a simple way for *calibrating* an arbitrary pretrained DPM, with which the score matching loss can be reduced and the lower bounds of model likelihood can consequently be increased. We provide general calibration guidelines under various model parametrizations. Our calibration method is performed only once and the resulting models can be used repeatedly for sampling. We conduct experiments on multiple datasets to empirically validate our proposal. Our code is available at <https://github.com/thudzj/Calibrated-DPMs>.

1 Introduction

In the past few years, denoising diffusion probabilistic modeling [17, 40] and score-based Langevin dynamics [42, 43] have demonstrated appealing results on generating images. Later, Song et al. [46] unify these two generative learning mechanisms through stochastic/ordinary differential equations (SDEs/ODEs). In the following we refer to this unified model family as diffusion probabilistic models (DPMs). The emerging success of DPMs has attracted broad interest in downstream applications, including image generation [10, 22, 48], shape generation [4], video generation [18, 19], super-resolution [35], speech synthesis [5], graph generation [51], textual inversion [13, 34], improving adversarial robustness [50], and text-to-image large models [32, 33], just to name a few.

A typical framework of DPMs involves a *forward* process gradually diffusing the data distribution $q_0(x_0)$ towards a noise distribution $q_T(x_T)$. The transition probability for $t \in [0, T]$ is a conditional Gaussian distribution $q_{0t}(x_t|x_0) = \mathcal{N}(x_t|\alpha_t x_0, \sigma_t^2 \mathbf{I})$, where $\alpha_t, \sigma_t \in \mathbb{R}^+$. Song et al. [46] show that there exist *reverse* SDE/ODE processes starting from $q_T(x_T)$ and sharing the same marginal distributions $q_t(x_t)$ as the forward process. The only unknown term in the reverse processes is the data score $\nabla_{x_t} \log q_t(x_t)$, which can be approximated by a time-dependent score model $s_\theta^t(x_t)$ (or with other model parametrizations). $s_\theta^t(x_t)$ is typically learned via score matching (SM) [20].

In this work, we observe that the stochastic process of the scaled data score $\alpha_t \nabla_{x_t} \log q_t(x_t)$ is a *martingale* w.r.t. the reverse-time process of x_t from T to 0 , where the timestep t can be either continuous or discrete. Along the reverse-time sampling path, this martingale property leads to concentration bounds for scaled data scores. Moreover, a martingale satisfies the optional stopping theorem that the expected value at a stopping time is equal to its initial expected value.

[†]Corresponding authors.

Based on the martingale property of data scores, for any $t \in [0, T]$ and any pretrained score models $s^t(x_t)$ (or with other model parametrizations), we calibrate the model by subtracting its expectation over $q_t(x_t)$, i.e., $E_{q_t(x_t)}[s^t(x_t)]$. We formally demonstrate that the calibrated score models $s^t(x_t) - E_{q_t(x_t)}[s^t(x_t)]$ achieves lower values of SM objectives. By the connections between SM objectives and model likelihood of the SDE process [45] or the ODE process [28], the calibrated score model has higher evidence lower bounds. Similar conclusions also hold for the conditional case, in which we calibrate a conditional score model $s^t(x_t; y)$ by subtracting its conditional expectation $E_{q_t(x_t; y)}[s^t(x_t; y)]$.

In practice, $E_{q_t(x_t)}[s^t(x_t)]$ or $E_{q_t(x_t; y)}[s^t(x_t; y)]$ can be approximated using noisy training data when the score model has been pretrained. We can also utilize an auxiliary shallow model to estimate these expectations dynamically during pretraining. When we do not have access to training data, we could calculate the expectations using data generated from q_t or $s^t(x_t; y)$. In experiments, we evaluate our calibration tricks on the CIFAR-100 [105] and CelebA64 [64] [27] datasets, reporting the FID scores [6]. We also provide insightful visualization results on the AFHQ [7], FFHQ [21] and ImageNet [9] at 64 resolution.

2 Diffusion probabilistic models

In this section, we briefly review the notations and training paradigms used in diffusion probabilistic models (DPMs). While recent works develop DPMs based on general corruption [106], we mainly focus on conventional Gaussian-based DPMs.

2.1 Forward and reverse processes

We consider a k -dimensional random variable $x_0 \in \mathbb{R}^k$ and define a forward diffusion process $\{x_t\}_{t \in [0, T]}$ with $T > 0$, which satisfies $t \in [0, T]$,

$$x_0 \sim q_0(x_0), \quad q_t(x_t | x_0) = N(x_t | \mu_t(x_0); \Sigma_t). \quad (1)$$

Here $q_0(x_0)$ is the data distribution; μ_t and Σ_t are two positive real-valued functions that are differentiable w.r.t. t with bounded derivatives. Let $p_t(x_t) = \int q_t(x_t | x_0) q_0(x_0) dx_0$ be the marginal distribution of x_t . The schedules of μ_t , Σ_t need to ensure that $p_T(x_T) = N(x_T | \mu_T; \Sigma_T)$ for some μ_T, Σ_T . Kingma et al. [23] prove that there exists a stochastic differential equation (SDE) satisfying the forward transition distribution in Eq. (1), and this SDE can be written as

$$dx_t = f(t)x_t dt + g(t)dw_t, \quad (2)$$

where $w_t \in \mathbb{R}^k$ is the standard Wiener process $dw_t = \frac{d \log \Sigma_t}{dt} dt$, and $g(t)^2 = \frac{d \Sigma_t}{dt} - 2 \frac{d \mu_t}{dt} \frac{d \log \Sigma_t}{dt}$. Song et al. [46] demonstrate that the forward SDE in Eq. (2) corresponds to a reverse SDE constructed as

$$dx_t = f(t)x_t dt - g(t)^2 r_{x_t} \log q_t(x_t) dt + g(t)dw_t^-, \quad (3)$$

where $w_t^- \in \mathbb{R}^k$ is the standard Wiener process in reverse time. Starting from $q_T(x_T)$, the marginal distribution of the reverse SDE process is $q_t(x_t)$ for $t \in [0, T]$. There also exists a deterministic process described by an ordinary differential equation (ODE) as

$$\frac{dx_t}{dt} = f(t)x_t - \frac{1}{2}g(t)^2 r_{x_t} \log q_t(x_t), \quad (4)$$

which starts from $q_T(x_T)$ and shares the same marginal distribution $q_t(x_t)$ as the reverse SDE in Eq. (3). Moreover, let $q_t(x_0 | x_t) = \frac{q_t(x_t | x_0) q_0(x_0)}{q_t(x_t)}$ and by Tweedie's formula [12], we know that $E_{q_t(x_0 | x_t)}[x_0] = x_t + \frac{1}{2} r_{x_t} \log q_t(x_t)$.

2.2 Training paradigm of DPMs

To estimate the data score $r_{x_t} \log q_t(x_t)$ at timestep t , a score-based model $s^t(x_t)$ [46] with shared parameters is trained to minimize the score matching (SM) objective [20] as

$$J_{SM}^t(\theta), \quad \frac{1}{2} E_{q_t(x_t)} \| s^t(x_t) - r_{x_t} \log q_t(x_t) \|_2^2. \quad (5)$$

To eliminate the intractable computation of $r_{x_t} \log q_t(x_t)$, denoising score matching (DSM) [9] transforms $J_{SM}^t(\theta)$ into $J_{DSM}^t(\theta)$, $\frac{1}{2} E_{q_0(x_0); q(t)} \| s^t(x_t) + \frac{1}{t} x_t \|^2$, where $x_t = \mu_t(x_0) + \Sigma_t$ and

$q(\cdot) = N(\cdot; \mu, \Sigma)$ is a standard Gaussian distribution. Under mild boundary conditions, we know $J_{SM}^t(\cdot)$ and $J_{DSM}^t(\cdot)$ is equivalent up to a constant, i.e. $J_{SM}^t(\cdot) = J_{DSM}^t(\cdot) + C^t$ and C^t is a constant independent of the model parameter. Other SM variants [31, 44] are also applicable here. The total SM objective for training is a weighted sum of $J_{SM}^t(\cdot)$ across $t \in [0, T]$, defined as $J_{SM}(\cdot; \omega(t)) = \int_0^T \omega(t) J_{SM}^t(\cdot) dt$, where $\omega(t)$ is a positive weighting function. Similarly, the total DSM objective is $J_{DSM}(\cdot; \omega(t)) = \int_0^T \omega(t) J_{DSM}^t(\cdot) dt$. The training objectives under other model parametrizations such as noise prediction $\mathcal{L}_{\text{noise}}^t(x_t)$ [17, 33], data prediction $\mathcal{L}_{\text{data}}^t(x_t)$ [23, 32], and velocity prediction $\mathcal{L}_{\text{vel}}^t(x_t)$ [18, 38] are recapped in Appendix B.1.

2.3 Likelihood of DPMs

Suppose that the reverse processes start from a tractable prior $p_T(x_T) = N(x_T; \mu_T, \Sigma_T)$. We can approximate the reverse-time SDE process by substituting $\log q(x_t)$ with $s^t(x_t)$ in Eq. (3) as $dx_t = f(t)x_t - g(t)^2 s^t(x_t) dt + g(t) dW_t$, which induces the marginal distribution $p_t^{SDE}(x_t; \cdot)$ for $t \in [0, T]$. In particular, at $t = 0$, the KL divergence between $p_0(x_0)$ and $p_0^{SDE}(x_0; \cdot)$ can be bounded by the total SM objective $J_{SM}(\cdot; g(t)^2)$ with the weighing function $\omega(t)^2$, as stated below:

Lemma 1. (Proof in Song et al. [45]) Let $q(x_t)$ be constructed from the forward process in Eq. (2). Then under regularity conditions, we have $D_{KL}(p_0 \| p_0^{SDE}(\cdot)) \leq J_{SM}(\cdot; g(t)^2) + D_{KL}(q_T \| p_T)$.

Here $D_{KL}(q_T \| p_T)$ is the prior loss independent of θ . Similarly, we approximate the reverse-time ODE process by substituting $\log q(x_t)$ with $s^t(x_t)$ in Eq. (4) as $\frac{dx_t}{dt} = f(t)x_t - \frac{1}{2}g(t)^2 s^t(x_t)$, which induces the marginal distribution $p_t^{ODE}(x_t; \cdot)$ for $t \in [0, T]$. By the instantaneous change of variables formula [4], we have $\frac{d \log p_t^{ODE}(x_t; \cdot)}{dt} = \text{tr} \left(r_{x_t} \left(f(t)x_t - \frac{1}{2}g(t)^2 s^t(x_t) \right) \right)$, where $\text{tr}(\cdot)$ denotes the trace of a matrix. Integrating change of $\log p_t^{ODE}(x_t; \cdot)$ from $t = 0$ to T can give the value of $\log p_T(x_T) - \log p_0^{ODE}(x_0; \cdot)$, but requires tracking the path from x_0 to x_T . On the other hand, at $t = 0$, the KL divergence between $p_0(x_0)$ and $p_0^{ODE}(x_0; \cdot)$ can be decomposed:

Lemma 2. (Proof in Lu et al. [28]) Let $q(x_t)$ be constructed from the forward process in Eq. (2). Then under regularity conditions, we have $D_{KL}(p_0 \| p_0^{ODE}(\cdot)) = J_{SM}(\cdot; g(t)^2) + D_{KL}(q_T \| p_T) + J_{\text{Diff}}(\cdot)$, where the term $J_{\text{Diff}}(\cdot)$ measures the difference between $s^t(x_t)$ and $r_{x_t} \log p_t^{ODE}(x_t; \cdot)$.

Directly computing $J_{\text{Diff}}(\cdot)$ is intractable due to the term $r_{x_t} \log p_t^{ODE}(x_t; \cdot)$, nevertheless, we could bound $J_{\text{Diff}}(\cdot)$ via bounding high-order SM objectives [28].

3 Calibrating pretrained DPMs

In this section we begin with deriving the relationship between data scores at different timesteps, which leads us to a straightforward method for calibrating any pretrained DPMs. We investigate further how the dataset bias of finite samples prevents empirical learning from achieving calibration.

3.1 The stochastic process of data score

According to Kingma et al. [23], the form of the forward process in Eq. (1) can be generalized to any two timesteps $0 \leq s < t \leq T$. Then, the transition probability from x_s to x_t is written as

$q_{st}(x_t | x_s) = N(x_t | \mu_{st}(x_s), \Sigma_{st}(x_s))$, where $\mu_{st} = \frac{t}{s} x_s$ and $\Sigma_{st} = \frac{t^2}{s^2} \Sigma_s$. Here the marginal distribution satisfies $q(x_t) = \int q_{st}(x_t | x_s) q_s(x_s) dx_s$. We can generally derive the connection between data scores $r_{x_t} \log q(x_t)$ and $r_{x_s} \log q_s(x_s)$ as stated below:

Theorem 1. (Proof in Appendix A.1) Let $q(x_t)$ be constructed from the forward process in Eq. (2). Then under some regularity conditions, we have $s < t \leq T$,

$$r_{x_t} \log q(x_t) = E_{q_{st}(x_s | x_t)} [r_{x_s} \log q_s(x_s)], \quad (6)$$

where $q_{st}(x_s | x_t) = \frac{q_{st}(x_t | x_s) q_s(x_s)}{q_t(x_t)}$ is the transition probability from x_t to x_s .

Theorem 1 indicates that the stochastic process of $r_{x_t} \log q(x_t)$ is a martingale w.r.t. the reverse-time process of x_t from timestep T to 0. From the optional stopping theorem [4], the expected

value of a martingale at a stopping time is equal to its initial expected value $E_{q_0(x_0)} [r_{x_0} \log q_0(x_0)]$. It is known that, under a mild boundary condition $q(x_0)$, there is $E_{q_0(x_0)} [r_{x_0} \log q_0(x_0)] = 0$ (proof is recapped in Appendix A.2). Consequently, as to the stochastic process, the martingale property results in $E_{q_t(x_t)} [r_{x_t} \log q_t(x_t)] = 0$ for $8t \in [0; T]$. Moreover, the martingale property of the (scaled) data score $r_{x_t} \log q_t(x_t)$ leads to concentration bounds using Azuma's inequality and Doob's martingale inequality as derived in Appendix A.3. Although we do not use these concentration bounds further in this paper, there are other concurrent works that use roughly similar concentration bounds in diffusion models, such as proving consistency [47] or justifying trajectory retrieval [52].

3.2 A simple calibration trick

Given a pretrained model $f(x_t)$ in practice, there is usually $E_{q_t(x_t)} [s^t(x_t)] \neq 0$, despite the fact that the expected data score is zero $E_{q_t(x_t)} [r_{x_t} \log q_t(x_t)] = 0$. This motivates us to calibrate $f(x_t)$ to $s^t(x_t) = \tau$, where τ is a time-dependent calibration term that is independent of any particular input x_t . The calibrated SM objective is written as follows:

$$\begin{aligned} J_{SM}^t(\tau; t) &= \frac{1}{2} E_{q_t(x_t)} \|ks^t(x_t) - \tau - r_{x_t} \log q_t(x_t)\|_2^2 \\ &= J_{SM}^t(\tau) + E_{q_t(x_t)} [s^t(x_t) - \tau]^2 + \frac{1}{2} k_t^2 k_2^2, \end{aligned} \quad (7)$$

where the second equation holds after the result of $E_{q_t(x_t)} [r_{x_t} \log q_t(x_t)] = 0$, and there is $J_{SM}^t(\tau; 0) = J_{SM}^t(\tau)$ specially when $t = 0$. Note that the orange part in Eq. (7) is a quadratic function w.r.t. τ . We look for the optimal $\tau = \arg \min_{\tau} J_{SM}^t(\tau; t)$ that minimizes the calibrated SM objective, from which we can derive

$$\tau = E_{q_t(x_t)} [s^t(x_t)]. \quad (8)$$

After taking τ into $J_{SM}^t(\tau; t)$, we have

$$J_{SM}^t(\tau; t) = J_{SM}^t(\tau) + \frac{1}{2} E_{q_t(x_t)} [s^t(x_t) - \tau]^2. \quad (9)$$

Since there is $J_{SM}^t(\tau) = J_{DSM}^t(\tau) + C^t$, we have $J_{DSM}^t(\tau; t) = J_{DSM}^t(\tau) + \frac{1}{2} E_{q_t(x_t)} [s^t(x_t) - \tau]^2$ for the DSM objective. Similar calibration tricks are also valid under other model parametrizations and SM variants, as formally described in Appendix B.2.

Remark. For any pretrained score model $f(x_t)$, we can calibrate it into $s^t(x_t) = E_{q_t(x_t)} [s^t(x_t)]$, which reduces the SM/DSM objectives at time t to $\frac{1}{2} E_{q_t(x_t)} [s^t(x_t) - \tau]^2$. The expectation of the calibrated score model is always zero, i.e., $E_{q_t(x_t)} [s^t(x_t) - \tau] = E_{q_t(x_t)} [s^t(x_t)] - \tau = 0$ holds for any τ , which is consistent with $E_{q_t(x_t)} [r_{x_t} \log q_t(x_t)] = 0$ satisfied by data scores.

Calibration preserves conservativeness. A theoretical law of score-based modeling is that $f(x_t)$ may not correspond to a probability distribution. To solve this issue, Salimans et al. [17] develop an energy-based model design, which utilizes the power of score-based modeling and simultaneously makes sure that $f(x_t)$ is conservative, i.e., there exists a probability distribution $p(x_t)$ such that $\forall x_t \in \mathbb{R}^k$, we have $s^t(x_t) = -r_{x_t} \log p^t(x_t)$. In this case, after we calibrate $f(x_t)$ by subtracting τ , there is $s^t(x_t) - \tau = -r_{x_t} \log \frac{p^t(x_t)}{\exp(\tau r_{x_t}) Z_t(\tau)}$, where $Z_t(\tau) = \int_{\mathbb{R}^k} p^t(x_t) \exp(\tau r_{x_t}) dx_t$ represents the normalization factor. Intuitively, subtracting τ corresponds to a shift in the vector space, so if $f(x_t)$ is conservative, its calibrated version $f(x_t) - \tau$ is also conservative.

Conditional cases. As to the conditional DPMS, we usually employ a conditional model $f(x_t; y)$, where $y \in Y$ is the conditional context (e.g., class label or text prompt). To learn the conditional data score $r_{x_t} \log q_t(x_t; y) = r_{x_t} \log q_t(x_t; y)$, we minimize the SM objective defined as $J_{SM}^t(\tau; t(y)) = \frac{1}{2} E_{q_t(x_t; y)} \|ks^t(x_t; y) - \tau(y) - r_{x_t} \log q_t(x_t; y)\|_2^2$. Similar to the conclusion of $E_{q_t(x_t)} [r_{x_t} \log q_t(x_t)] = 0$, there is $E_{q_t(x_t; y)} [r_{x_t} \log q_t(x_t; y)] = 0$. To calibrate $f(x_t; y)$, we use the conditional term $\tau(y)$ and the calibrated SM objective is formulated as

$$\begin{aligned} J_{SM}^t(\tau; t(y)) &= \frac{1}{2} E_{q_t(x_t; y)} \|ks^t(x_t; y) - \tau(y) - r_{x_t} \log q_t(x_t; y)\|_2^2 \\ &= J_{SM}^t(\tau; t(y)) + E_{q_t(x_t; y)} [s^t(x_t; y) - \tau(y)]^2 + \frac{1}{2} k_t^2 k_2^2, \end{aligned} \quad (10)$$

and for any $2 \leq Y$, the optimal $\hat{y}_t(y)$ is given by $\hat{y}_t(y) = E_{q_t(x_t, y)} [s^t(x_t; y)]$. We highlight the conditional context in contrast to the unconditional form in Eq. (7). After taking $\hat{y}_t(y)$ into $J_{SM}^t(\cdot; \hat{y}_t(y))$, we have $J_{SM}^t(\cdot; \hat{y}_t(y)) = J_{SM}^t(\cdot) + \frac{1}{2} E_{q_t(y)} E_{q_t(x_t, y)} [s^t(x_t; y)]^2$. This conditional calibration form can naturally generalize to other model parametrizations and SM variants.

3.3 Likelihood of calibrated DPMS

Now we discuss the effects of calibration on model likelihood. Following the notations in Section 2.3, we use $p_0^{SDE}(\cdot; t)$ and $p_0^{ODE}(\cdot; t)$ to denote the distributions induced by the reverse-time SDE and ODE processes, respectively, where $\log q_t(x_t)$ is substituted with $s^t(x_t)$.

Likelihood of $p_0^{SDE}(\cdot; t)$. Let $J_{SM}(\cdot; t; g(t)^2)$, $\int_0^{R_T} g(t)^2 J_{SM}^t(\cdot; t) dt$ be the total SM objective after the score model is calibrated by then according to Lemma 1, we have $D_{KL}(q_T \| p_T) + J_{SM}(\cdot; t; g(t)^2) + D_{KL}(q_T \| p_T)$. From the result in Eq. (9), there is

$$J_{SM}(\cdot; t; g(t)^2) = J_{SM}(\cdot; g(t)^2) + \frac{1}{2} \int_0^T g(t)^2 E_{q_t(x_t)} [s^t(x_t)]^2 dt. \quad (11)$$

Therefore, the likelihood $D_{KL}(q_0 \| p_0^{SDE}(\cdot; t))$ after calibration has a lower upper bound of $J_{SM}(\cdot; t; g(t)^2) + D_{KL}(q_T \| p_T)$, compared to the bound of $J_{SM}(\cdot; g(t)^2) + D_{KL}(q_T \| p_T)$ for the original $D_{KL}(q_0 \| p_0^{SDE}(\cdot; t))$. However, we need to clarify that $D_{KL}(q_0 \| p_0^{SDE}(\cdot; t))$ may not necessarily smaller than $D_{KL}(q_0 \| p_0^{SDE}(\cdot; t))$, since we can only compare their upper bounds.

Likelihood of $p_0^{ODE}(\cdot; t)$. Note that in Lemma 2, there is a term $\text{Diff}(\cdot)$, which is usually small in practice since $s^t(x_t)$ and $r_{x_t} \log p_t^{ODE}(x_t; \cdot)$ are close. Thus, we have

$$D_{KL}(q_0 \| p_0^{ODE}(\cdot; t)) \approx J_{SM}(\cdot; t; g(t)^2) + D_{KL}(q_T \| p_T),$$

and $D_{KL}(q_0 \| p_0^{ODE}(\cdot; t))$ approximately achieves its lowest value. Lu et al [23] show that $D_{KL}(q_0 \| p_0^{ODE}(\cdot; t))$ can be further bounded by high-order SM objectives (as detailed in Appendix A.4), which depend on $r_{x_t} s^t(x_t)$ and $r_{x_t} \text{tr}(r_{x_t} s^t(x_t))$. Since the calibration term is independent of x_t , i.e., $r_{x_t} s^t(x_t) = 0$, it does not affect the values of high-order SM objectives, and achieves a lower upper bound due to the lower value of the first-order SM objective.

3.4 Empirical learning fails to achieve $E_{q_t(x_t)} [s^t(x_t)] = 0$

A question that naturally arises is whether better architectures or learning algorithms for DPMS (e.g., EDMs [22]) could empirically achieve $E_{q_t(x_t)} [s^t(x_t)] = 0$ without calibration? The answer may be negative, since in practice we only have access to a dataset sampled from $p_0(x_0)$.

More specifically, assuming that we have a training dataset $\{x_0^n\}_{n=1}^N$ where $x_0^n \sim p_0(x_0)$, and defining the kernel density distribution induced by $\hat{q}_t(x_t; D) = \frac{1}{N} \sum_{n=1}^N \frac{x_t - x_0^n}{t} \delta(\cdot)$.

When the quantity of training data approaches infinity, we have $\hat{q}_t(x_t; D) = q_t(x_t)$ holds for $\forall t \in [0; T]$. Then the empirical DSM objective trained on D is written as

$$J_{DSM}^t(\cdot; D) = \frac{1}{2N} \sum_{n=1}^N E_{q_t(\cdot)} [s^t(x_0^n + t)]^2, \quad (12)$$

and it is easy to show that the optimal solution for minimizing $J_{DSM}^t(\cdot; D)$ satisfies (assuming q^t has universal model capacity) $r_{x_t} \log q_t(x_t; D) = 0$. Given a finite dataset D , there is

$$E_{q_t(x_t; D)} [r_{x_t} \log q_t(x_t; D)] = 0, \text{ but typically } E_{q_t(x_t)} [r_{x_t} \log q_t(x_t; D)] \neq 0, \quad (13)$$

indicating that even if the score model is learned to be optimal, there is still $E_{q_t(x_t)} [s^t(x_t)] \neq 0$. Thus, the mis-calibration of DPMS is partially due to the dataset bias, i.e., during training we can only access a finite dataset sampled from $p_0(x_0)$.

Furthermore, when trained on a finite dataset in practice, the learned model will not converge to the optimal solution [15], so there is typically $r_{x_t} \log q_t(x_t; D) \neq 0$ and $E_{q_t(x_t; D)} [s^t(x_t)] \neq 0$.

After calibration, we can at least guarantee that $E_{q_t(x_t; D)} [s^t(x_t) - E_{q_t(x_t; D)} [s^t(x_t)]] = 0$ always holds on any finite dataset D . In Figure 3, we demonstrate that even state-of-the-art EDMs still have non-zero and semantic $E_{q_t(x_t)} [s^t(x_t)]$, which emphasises the significance of calibrating DPMS.

Figure 1: Time-dependent values $\frac{1}{2} \mathbb{E}_{q_t(x_t)} [\|s^t(x_t)\|_2^2]$ (the first row) and $\frac{g(t)^2}{2} \mathbb{E}_{q_t(x_t)} [\|s^t(x_t)\|_2^2]$ (the second row) calculated on different datasets. The models on CIFAR-10 and CelebA is trained on discrete timesteps $t \in [0; 1; \dots; 1000]$, while those on AFHQv2, FFHQ, and ImageNet are trained on continuous timesteps $t \in [0; 1]$. We convert data prediction $\hat{x}_t(x_t)$ into noise prediction $s^t(x_t)$ based on $s^t(x_t) = (x_t - \hat{x}_t(x_t)) / \sigma_t$. The y-axis is clamped into $[0; 500]$.

3.5 Amortized computation of $\mathbb{E}_{q_t(x_t)} [s^t(x_t)]$

By default, we are able to calculate and restore the value of $\mathbb{E}_{q_t(x_t)} [s^t(x_t)]$ for a pretrained model $s^t(x_t)$, where the selection of timesteps is determined by the inference algorithm, and the expectation over $q_t(x_t)$ can be approximated by Monte Carlo sampling from a noisy training set. When we do not have access to training data, we can approximate the expectation using data generated from $p_t^{\text{ODE}}(x_t; \cdot)$ or $p_t^{\text{SDE}}(x_t; \cdot)$. Since we only need to calculate $\mathbb{E}_{q_t(x_t)} [s^t(x_t)]$ once, the raised computational overhead is amortized as the number of generated samples increases.

Dynamically recording. In the preceding context, we focus primarily on post-training computing of $\mathbb{E}_{q_t(x_t)} [s^t(x_t)]$. An alternative strategy would be to dynamically record $\mathbb{E}_{q_t(x_t)} [s^t(x_t)]$ during the pretraining phase of $s^t(x_t)$. Specifically, we could construct an auxiliary shallow network $\text{rec}(t)$ parameterized by θ , whose input is the timestep t . We define the expected mean squared error as

$$J_{\text{Cal}}^t(\theta), \mathbb{E}_{q_t(x_t)} \|\text{rec}(t) - s^t(x_t)\|_2^2, \quad (14)$$

where the superscript t denotes the stopping gradient and is the optimal solution of minimizing $J_{\text{Cal}}^t(\theta)$ w.r.t. θ , satisfying $\text{rec}(t) = \mathbb{E}_{q_t(x_t)} [s^t(x_t)]$ (assuming sufficient model capacity).

The total training objective can therefore be expressed as $\int_0^T \alpha(t) J_{\text{Cal}}^t(\theta) dt$, where $\alpha(t)$ is a time-dependent trade-off coefficient for $t \in [0; T]$.

4 Experiments

In this section, we demonstrate that sample quality and model likelihood can be both improved by calibrating DPMs. Instead of establishing a new state-of-the-art, the purpose of our empirical studies is to testify the efficacy of our calibration technique as a simple way to repair DPMs.

4.1 Sample quality

Setup. We apply post-training calibration to discrete-time models trained on CIFAR-10 and CelebA [27], which apply parametrization of noise prediction $\hat{x}_t(x_t)$. In the sampling phase, we employ DPM-Solver [29], an ODE-based sampler that achieves a promising balance between sample efficiency and image quality. Because our calibration directly acts on model scores, it is also compatible with other ODE/SDE-based samplers [26], while we only focus on DPM-Solver cases in this paper. In accordance with the recommendation, we set the end time of DPM-Solver to 10^{-4} when the number of sampling steps is less than 1000 and to 10^{-4} otherwise. Additional details can be

Table 1: Comparison on sample quality measured by FID with varying NFE on CIFAR-10. Experiments are conducted using a linear noise schedule on the discrete-time model. We consider three variants of DPM-Solver with different orders. The results within the actual NFE is order $b \frac{NFE}{order} c$ which is smaller than the given NFE, following the setting in [29].

Noise prediction	DPM-Solver	Number of evaluations (NFE)						
		10	15	20	25	30	35	40
$t(x_t)$	1-order	20.49	12.47	9.72	7.89	6.84	6.22	5.75
	2-order	7.35	\checkmark 4.52	4.14	\checkmark 3.92	3.74	\checkmark 3.71	3.68
	3-order	\checkmark 23.96	4.61	\checkmark 3.89	\checkmark 3.73	3.65	\checkmark 3.65	\checkmark 3.60
$t(x_t) \quad E_{q_t(x_t)} [t(x_t)]$	1-order	19.31	11.77	8.86	7.35	6.28	5.76	5.36
	2-order	6.76	\checkmark 4.36	4.03	\checkmark 3.66	3.54	\checkmark 3.44	3.48
	3-order	\checkmark 53.50	4.22	\checkmark 3.32	\checkmark 3.33	3.35	\checkmark 3.32	\checkmark 3.31

Table 2: Comparison on sample quality measured by FID with varying NFE on CelebA 64x64. Experiments are conducted using a linear noise schedule on the discrete-time model. The settings of DPM-Solver are the same as on CIFAR-10.

Noise prediction	DPM-Solver	Number of evaluations (NFE)						
		10	15	20	25	30	35	40
$t(x_t)$	1-order	16.74	11.85	7.93	6.67	5.90	5.38	5.01
	2-order	4.32	\checkmark 3.98	2.94	\checkmark 2.88	2.88	\checkmark 2.88	2.84
	3-order	\checkmark 11.92	3.91	\checkmark 2.84	\checkmark 2.76	2.82	\checkmark 2.81	\checkmark 2.85
$t(x_t) \quad E_{q_t(x_t)} [t(x_t)]$	1-order	16.13	11.29	7.09	6.06	5.28	4.87	4.39
	2-order	4.42	\checkmark 3.94	2.61	\checkmark 2.66	2.54	\checkmark 2.52	2.49
	3-order	\checkmark 35.47	3.62	\checkmark 2.33	\checkmark 2.43	2.40	\checkmark 2.43	\checkmark 2.49

found in Lu et al. [29]. By default, we employ the FID score [6] to quantify the sample quality using 50,000 samples. Typically, a lower FID indicates a higher sample quality. In addition, in Table 3, we evaluate using other metrics such as sFID [30], IS [39], and Precision/Recall [36].

Computing $E_{q_t(x_t)} [t(x_t)]$. To estimate the expectation over $q_t(x_t)$, we construct $x_t = t x_0 + \sqrt{1-t} \epsilon$, where $x_0 = q_0(x_0)$ is sampled from the training set and $N(\epsilon; 0, I)$ is sampled from a standard Gaussian distribution. The selection of timestep depends on the sampling schedule of DPM-Solver. The computed values $E_{q_t(x_t)} [t(x_t)]$ are stored in a dictionary and warped into the output layers of DPMs, allowing existing inference pipelines to be reused.

We first calibrate the model trained by Ho et al. [17] on the CIFAR-10 dataset and compare it to the original one for sampling with DPM-Solvers. We conduct a systematical study with varying NFE (i.e., number of function evaluations) and solver order. The results are presented in Tables 1 and 3. After calibrating the model, the sample quality is consistently enhanced, which demonstrates the significance of doing so and the efficacy of our method. We highlight the significant improvement in sample quality (4.61 \rightarrow 4.22 when using 15 NFE and 3-order DPM-Solver; 3.89 \rightarrow 3.32 when using 20 NFE and 3-order DPM-Solver). After model calibration, the number of steps required to achieve convergence for 3-order DPM-Solver is reduced from 30 to 20, making our method a new option for expediting the sampling of DPMs. In addition, as a point of comparison, 3-order DPM-Solver with 1,000 NFE can only yield an FID score of 3.45 when using the original model, which, along with the results in Table 1, indicates that model calibration helps to improve the convergence of sampling.

Then, we conduct experiments with the discrete-time model trained on the CelebA 64x64 dataset by Song et al. [41]. The corresponding sample quality comparison is shown in Table 2. Clearly, model calibration brings significant gains (3.91 \rightarrow 3.62 when using 15 NFE and 3-order DPM-Solver; 2.84 \rightarrow 2.33 when using 20 NFE and 3-order DPM-Solver) that are consistent with those on the CIFAR-10 dataset. This demonstrates the prevalence of the mis-calibration issue in existing DPMs and the efficacy of our correction. We still observe that model calibration improves convergence of sampling, and as shown in Figure 2, our calibration could help to reduce ambiguous generations. More generated images are displayed in Appendix C.

Table 3: Comparison on sample quality measured by different metrics, including FID#, inception score (IS), precision and recall with varying NFE on CIFAR-10. We use Base to denote the baseline \hat{x}_t and Ours to denote calibrated score $\hat{x}_t - E_{q_t(x_t)}[\hat{x}_t]$. The sampler is DPM-Solver with different orders. Note that FID is computed by the PyTorch checkpoint of Inception-v3, while sFID/IS/Precision/Recall are computed by the TensorFlow checkpoint of Inception-v3 following github.com/kynkaat/improved-precision-and-recall-metric

Method	Number of evaluations (NFE)															
	20					25					30					
	FID	sFID	IS	Pre.	Rec.	FID	sFID	IS	Pre.	Rec.	FID	sFID	IS	Pre.	Rec.	
Base	1-ord.	9.72	6.03	8.49	0.641	0.547	7.89	5.45	8.68	0.644	0.556	6.84	5.12	8.76	0.650	0.565
	2-ord.	4.14	4.36	9.15	0.654	0.590	9.92	4.22	9.17	0.657	0.591	7.74	4.18	9.20	0.658	0.591
	3-ord.	3.89	4.18	9.29	0.652	0.597	7.73	4.15	9.21	0.657	0.595	5.65	4.12	9.22	0.658	0.593
Ours	1-ord.	8.86	6.01	8.56	0.649	0.547	7.35	5.42	8.76	0.653	0.560	6.28	5.09	8.84	0.653	0.568
	2-ord.	4.03	4.31	9.17	0.661	0.592	3.66	4.20	9.20	0.664	0.594	4.54	4.14	9.23	0.662	0.599
	3-ord.	3.32	4.14	9.38	0.657	0.603	3.33	4.11	9.28	0.665	0.597	3.35	4.08	9.27	0.662	0.600

Figure 2: Selected images on CIFAR-10 (generated with NFE = 20 using 3-order DPM-Solver) demonstrating that our calibration could reduce ambiguous generations, such as generations that resemble both horse and dog. However, we must emphasize that not all generated images have a visually discernible difference before and after calibration.

4.2 Model likelihood

As described in Section 3.3, calibration contributes to reducing the SM objective, thereby decreasing the upper bound of the KL divergence between model distribution at timestep t (either $p_0^{\text{SDE}}(\cdot; t)$ or $p_0^{\text{ODE}}(\cdot; t)$) and data distribution q . Consequently, it aids in raising the lower bound of model likelihood. In this subsection, we examine such effects by evaluating the aforementioned DPMs on the CIFAR-10 and CelebA datasets. We also conduct experiments with continuous-time models trained by Karras et al [22] on AFHQv2 64 64 [7], FFHQ 64 64 [21], and ImageNet 64 64 [9] datasets considering their top performance. These models apply parametrization of data prediction $x^t(x_t)$, and for consistency, we convert it to align with \hat{x}_t based on the relationship $\hat{x}_t = (x_t - \hat{x}_t) + \hat{x}_t$, as detailed in Kingma et al. [23] and Appendix B.2.

Given that we employ noise prediction models in practice, we first estimate $\frac{1}{2} E_{q_t(x_t)}[\hat{x}_t - x_t]^2$ at timestep $t \in [0; T]$, which reflects the decrement on the SM objective according to Eq. (9) (up to a scaling factor of σ_t^2). We approximate the expectation using Monte Carlo (MC) estimation with training data points. The results are displayed in the first row of Figure 1. Notably, the value of $\frac{1}{2} E_{q_t(x_t)}[\hat{x}_t - x_t]^2$ varies significantly along with timestep t ; it decreases relative to for CelebA but increases in all other cases (except for $[0; 4; 1; 0]$ on ImageNet 64 64). Ideally, there should be $\frac{1}{2} E_{q_t(x_t)}[\hat{x}_t - x_t]^2 = 0$ at any t . Such inconsistency reveals that mis-calibration issues exist in general, although the phenomenon may vary across datasets and training mechanisms.

Then, we quantify the gain of model calibration on increasing the lower bound of model likelihood, which is $\frac{1}{2} \int_0^T g(t)^2 E_{q_t(x_t)}[s^t(x_t)]^2 dt$ according to Eq. (11). We first rewrite it with the model parametrization of noise prediction $h(x_t)$, and it can be straightforwardly demonstrated that it equals $\int_0^T \frac{g(t)^2}{2} E_{q_t(x_t)}[\hat{x}_t - x_t]^2 dt$. Therefore, we calculate the value of $\frac{g(t)^2}{2} E_{q_t(x_t)}[\hat{x}_t - x_t]^2$ using MC

Figure 3: Visualization of the expected predicted noises with increasing t for each dataset, the first row displays $E_{q_t(x_t)}[\epsilon_t(x_t)]$ (after normalization) and the second row highlights the 10% pixels that $E_{q_t(x_t)}[\epsilon_t(x_t)]$ has high values. The DPM on CelebA is a discrete-time model with 1000 timesteps [41] and that on FFHQ is a continuous-time one [22].

estimation and report the results in the second row of Figure 1. The integral is represented by the area under the curve (i.e., the gain of model calibration on the lower bound of model likelihood). Various datasets and model architectures exhibit non-trivial gains, as observed. In addition, we notice that the DPMs trained by Karras et al. [22] show patterns distinct from those of DDPM [7] and DDIM [41], indicating that different DPM training mechanisms may result in different mis-calibration effects.

Visualizing $E_{q_t(x_t)}[\epsilon_t(x_t)]$. To better understand the inductive bias learned by DPMs, we visualize the expected predicted noises $E_{q_t(x_t)}[\epsilon_t(x_t)]$ for timestep from 0 to T , as seen in Figure 3. For each dataset, the first row normalizes the values $E_{q_t(x_t)}[\epsilon_t(x_t)]$ into $[0; 255]$; the second row calculates pixel-wise norm (across RGB channels) and highlights the 10% locations with the highest norm. As we can observe, on facial datasets like CelebA and FFHQ, there are obvious facial patterns inside $E_{q_t(x_t)}[\epsilon_t(x_t)]$, while on other datasets like CIFAR-10, ImageNet, as well as the animal face dataset AFHQv2, the patterns inside $E_{q_t(x_t)}[\epsilon_t(x_t)]$ are more like random noises. Besides, the facial patterns in Figure 3 are more significant when t is smaller, and become blurry when t is close to T . This phenomenon may be attributed to the bias of finite training data, which is detrimental to generalization during sampling and justifies the importance of calibration as described in Section 3.4.

4.3 Ablation studies

We conduct ablation studies focusing on the estimation method of $E_{q_t(x_t)}[\epsilon_t(x_t)]$.

Estimating $E_{q_t(x_t)}[\epsilon_t(x_t)]$ with partial training data. In the post-training calibration setting, our primary algorithmic change is to subtract the calibration term $E_{q_t(x_t)}[\epsilon_t(x_t)]$ from the pretrained DPMs' output. In the aforementioned studies, the expectation $E_{q_t(x_t)}[\epsilon_t(x_t)]$ (or its variant of other model parametrizations) is approximated with MC estimation using all training images. However, there may be situations where training data are (partially) inaccessible. To evaluate the effectiveness of our method under these cases, we examine the number of training images used to estimate the calibration term on CIFAR-10. To determine the quality of the estimated calibration term, we sample from the calibrated models using our DPM-Solver running for 20 steps and evaluate the corresponding FID score. The results are listed in the left part of Table 4. As observed, we need to use the majority of training images (at least 20,000) to estimate the calibration term. We deduce that this is because the CIFAR-10 images are rich in diversity, necessitating a non-trivial number of training images to cover the various modes and produce a nearly unbiased calibration term.

Estimating $E_{q_t(x_t)}[\epsilon_t(x_t)]$ with generated data. In the most extreme case where we do not have access to any training data (e.g., due to privacy concerns), we could still estimate the expectation over $q_t(x_t)$ with data generated from $p_0^{\text{ODE}}(x_0; \cdot)$ or $p_0^{\text{SDE}}(x_0; \cdot)$. Specifically, under the hypothesis that $p_0^{\text{ODE}}(x_0; \cdot) = q_0(x_0)$ (DPM-Solver is an ODE-based sampler), we first generate $p_0^{\text{ODE}}(x_0; \cdot)$

Table 4: Sample quality varies w.r.t. the number of training images (left part) and generated images (right part) used to estimate the calibration term on CIFAR-10. In the generated data case, the images used to estimate the calibration term rapidly decrease, across different timesteps.

Training data		Generated data	
# of samples	FID#	# of samples	FID#
500	55.38	2,000	8.80
1,000	18.72	5,000	4.53
2,000	8.05	10,000	3.78
5,000	4.31	20,000	3.31
10,000	3.47	50,000	3.46
20,000	3.25	100,000	3.47
50,000	3.32	200,000	3.46

and construct $\mathbf{x}_t = \mathbf{x}_0 + \mathbf{p}_t$, where $\mathbf{p}_t = \mathbb{E}_{\mathbf{q}_t(x_t)}[\mathbf{s}^t(x_t)]$. Then, the expectation over $\mathbf{q}_t(x_t)$ could be approximated by the expectation over \mathbf{p}_t .

Empirically, on the CIFAR-10 dataset, we adopt a 3-order DPM-Solver to generate a set of samples from the pretrained model of Ho et al. [17], using a relatively large number of sampling steps (e.g., 50 steps). This set of generated data is used to calculate the calibration term $\mathbb{E}_{\mathbf{q}_t(x_t)}[\mathbf{s}^t(x_t)]$. Then, we obtain the calibrated model $\mathbb{E}_{\mathbf{q}_t(x_t)}[\mathbf{s}^t(x_t)]$ and craft new images based on a 20-step DPM-Solver. In the right part of Table 4, we present the results of an empirical investigation into how the number of generated images influences the quality of model calibration.

Using the same sampling setting, we also provide two reference points: 1) the originally mis-calibrated model can reach the FID score of 3.89, and 2) the model calibrated with training data can reach the FID score of 3.32. Comparing these results reveals that the DPM calibrated with a large number of high-quality generations can achieve comparable FID scores to those calibrated with training samples (see the result of using 20,000 generated images). Additionally, it appears that using more generations is not advantageous. This may be because the generations from DPMs, despite being known to cover diverse modes, still exhibit semantic redundancy and deviate slightly from the data distribution.

Dynamical recording. We simulate the proposed dynamical recording technique. Specifically, we use a 3-layer MLP of width 512 to parameterize the aforementioned network and train it with an Adam optimizer [14] to approximate the expected predicted noises $\mathbb{E}_{\mathbf{q}_t(x_t)}[\mathbf{s}^t(x_t)]$, where $\mathbf{s}^t(x_t)$ comes from the pretrained noise prediction model on CIFAR-10. The training of $\mathbf{h}(\mathbf{t})$ runs for 1,000 epochs. Meanwhile, using the training data, we compute the expected predicted noises with MC estimation and treat them as the ground truth. In Figure 4, we compare them to the outputs of $\mathbf{h}(\mathbf{t})$ and visualize the disparity measured by mean square error. As demonstrated, as the number of training epochs increases, the network $\mathbf{h}(\mathbf{t})$ quickly converges and can form a relatively reliable approximation to the ground truth. Dynamic recording has a distinct advantage of being able to be performed during the training of DPMs to enable immediate generation. We clarify that better timestep embedding techniques and NN architectures can improve approximation quality even further.

5 Discussion

We propose a straightforward method for calibrating any pretrained DPM that can provably reduce the values of SM objectives and, as a result, induce higher values of lower bounds for model likelihood. We demonstrate that the mis-calibration of DPMs may be inherent due to the dataset bias and/or sub-optimally learned model scores. Our findings also provide a potentially new metric for assessing a diffusion model by its degree of “uncalibration”, namely, how far the learned scores deviate from the essential properties (e.g., the expected data scores should be zero).

Limitations. While our calibration method provably improves the model’s likelihood, it does not necessarily yield a lower FID score, as previously discussed. Besides, for text-to-image generation, post-training computation $\mathbb{E}_{\mathbf{q}_t(x_t, y)}[\mathbf{s}^t(x_t; y)]$ becomes infeasible due to the exponentially large number of conditions, necessitating dynamic recording with multimodal modules.

Acknowledgements

Zhijie Deng was supported by Natural Science Foundation of Shanghai (No. 23ZR1428700) and the Key Research and Development Program of Shandong Province, China (No. 2023CXGC010112).

References

- [1] Kazuoki Azuma. Weighted sums of certain dependent random variables. *Tohoku Mathematical Journal, Second Series* 9(3):357–367, 1967.
- [2] Arpit Bansal, Eitan Borgnia, Hong-Min Chu, Jie S Li, Hamid Kazemi, Furong Huang, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Cold diffusion: Inverting arbitrary image transforms without noise. *arXiv preprint arXiv:2208.09392*, 2022.
- [3] Fan Bao, Chongxuan Li, Jun Zhu, and Bo Zhang. Analytic-dpm: an analytic estimate of the optimal reverse variance in diffusion probabilistic models. *International Conference on Learning Representations (ICLR)*, 2022.
- [4] Ruojin Cai, Guandao Yang, Hadar Averbuch-Elor, Zekun Hao, Serge Belongie, Noah Snaveley, and Bharath Hariharan. Learning gradient fields for shape generation. *European Conference on Computer Vision (ECCV)*, 2020.
- [5] Nanxin Chen, Yu Zhang, Heiga Zen, Ron J Weiss, Mohammad Norouzi, and William Chan. Wavegrad: Estimating gradients for waveform generation. *International Conference on Learning Representations (ICLR)*, 2021.
- [6] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [7] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. *IEEE International Conference on Computer Vision (CVPR)* 2020.
- [8] Giannis Daras, Mauricio Delbracio, Hossein Talebi, Alexandros G Dimakis, and Peyman Milanfar. Soft diffusion: Score matching for general corruptions. *arXiv preprint arXiv:2209.05442*, 2022.
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* 2009.
- [10] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*.
- [11] Joseph L Doob and Joseph L Doob. *Stochastic processes*, volume 7. Wiley New York, 1953.
- [12] Bradley Efron. Tweedie's formula and selection bias. *Journal of the American Statistical Association*, 106(496):1602–1614, 2011.
- [13] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint arXiv:2208.01618*, 2022.
- [14] Geoffrey Grimmett and David Stirzaker. *Probability and random processes*. Oxford university press, 2001.
- [15] Xiangming Gu, Chao Du, Tianyu Pang, Chongxuan Li, Min Lin, and Ye Wang. On memorization in diffusion models. *arXiv preprint arXiv:2310.02664*, 2023.
- [16] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in Neural Information Processing Systems (NeurIPS)*, pages 6626–6637, 2017.

- [17] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In Advances in Neural Information Processing Systems (NeurIPS).
- [18] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. arXiv preprint arXiv:2210.02303, 2022.
- [19] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. arXiv preprint arXiv:2204.03458, 2022.
- [20] Aapo Hyvärinen. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research (JMLR)* (Apr):695–709, 2005.
- [21] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. *IEEE International Conference on Computer Vision (CVPR)* 19.
- [22] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *Advances in Neural Information Processing Systems (NeurIPS)* 2022.
- [23] Diederik Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. In Advances in Neural Information Processing Systems (NeurIPS).
- [24] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014.
- [25] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- [26] Luping Liu, Yi Ren, Zhijie Lin, and Zhou Zhao. Pseudo numerical methods for diffusion models on manifolds. In International Conference on Learning Representations (ICLR) 22.
- [27] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In International Conference on Computer Vision (ICCV) 15.
- [28] Cheng Lu, Kaiwen Zheng, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Maximum likelihood training for score-based diffusion odes by high order denoising score matching. In International Conference on Machine Learning (ICML) 22.
- [29] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural Information Processing Systems (NeurIPS)* 22.
- [30] Charlie Nash, Jacob Menick, Sander Dieleman, and Peter W Battaglia. Generating images with sparse representations. arXiv preprint arXiv:2103.03841, 2021.
- [31] Tianyu Pang, Kun Xu, Chongxuan Li, Yang Song, Stefano Ermon, and Jun Zhu. Efficient learning of generative models via finite-difference score matching. *Advances in Neural Information Processing Systems (NeurIPS)* 20.
- [32] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. arXiv preprint arXiv:2204.06125, 2022.
- [33] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* 22.
- [34] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. arXiv preprint arXiv:2208.12242, 2022.
- [35] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* 2022.

- [36] Mehdi SM Sajjadi, Olivier Bachem, Mario Lucic, Olivier Bousquet, and Sylvain Gelly. Assessing generative models via precision and recall. *Advances in Neural Information Processing Systems (NeurIPS)* 2018.
- [37] Tim Salimans and Jonathan Ho. Should ebms model the energy or the score? *Energy Based Models Workshop-ICLR* 2021.
- [38] Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. In *International Conference on Learning Representations (ICLR)* 2022.
- [39] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. *Advances in Neural Information Processing Systems (NeurIPS)* 2016.
- [40] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning (ICML)*, pages 2256–2265. PMLR, 2015.
- [41] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations (ICLR)* 2021.
- [42] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 11895–11907, 2019.
- [43] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. In *Advances in Neural Information Processing Systems (NeurIPS)* 2020.
- [44] Yang Song, Sahaj Garg, Jiaxin Shi, and Stefano Ermon. Sliced score matching: A scalable approach to density and score estimation. *Conference on Uncertainty in Artificial Intelligence (UAI)*, 2019.
- [45] Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of score-based diffusion models. *Advances in Neural Information Processing Systems (NeurIPS)* 2021.
- [46] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations (ICLR)* 2021.
- [47] Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. preprint arXiv:2303.01469 2023.
- [48] Arash Vahdat, Karsten Kreis, and Jan Kautz. Score-based generative modeling in latent space. In *Advances in Neural Information Processing Systems (NeurIPS)* 2021.
- [49] Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011.
- [50] Zekai Wang, Tianyu Pang, Chao Du, Min Lin, Weiwei Liu, and Shuicheng Yan. Better diffusion models further improve adversarial training. *International Conference on Machine Learning (ICML)*, 2023.
- [51] Minkai Xu, Lantao Yu, Yang Song, Chence Shi, Stefano Ermon, and Jian Tang. Geodiff: A geometric diffusion model for molecular conformation generation. *International Conference on Learning Representations (ICLR)* 2022.
- [52] Kexun Zhang, Xianjun Yang, William Yang Wang, and Lei Li. Redi: Efficient learning-free diffusion inference via trajectory retrieval. arXiv preprint arXiv:2302.02285 2023.

A Detailed derivations

In this section, we provide detailed derivations for the Theorem and equations shown in the main text. We follow the regularization assumptions listed in Song et al. [45].

A.1 Proof of Theorem 1

Proof. For any two timesteps $0 \leq s < t \leq T$, i.e., the transition probability from x_s to x_t is written as $q_{st}(x_t|x_s) = \mathcal{N}_{\mathbb{R}^k} \left(x_t | \alpha_{t|s} x_s, \sigma_{t|s}^2 \mathbf{I} \right)$, where $\alpha_{t|s} = \frac{\alpha_t}{\alpha_s}$ and $\sigma_{t|s}^2 = \sigma_t^2 - \alpha_{t|s}^2 \sigma_s^2$. The marginal distribution $q_t(x_t) = \int_{\mathbb{R}^k} q_{st}(x_t|x_s) q_s(x_s) dx_s$ and we have

$$\begin{aligned}
\nabla_{x_t} \log q_t(x_t) &= \frac{1}{\alpha_{t|s}} \nabla_{\alpha_{t|s}^{-1} x_t} \log \frac{1}{\alpha_{t|s}^k} \mathbb{E}_{\mathcal{N}_{\mathbb{R}^k} \left(x_s | \alpha_{t|s}^{-1} x_t, \alpha_{t|s}^2 \sigma_{t|s}^2 \mathbf{I} \right)} [q_s(x_s)] \\
&= \frac{1}{\alpha_{t|s}} \nabla_{\alpha_{t|s}^{-1} x_t} \log \mathbb{E}_{\mathcal{N}_{\mathbb{R}^k} \left(\eta | 0, \alpha_{t|s}^2 \sigma_{t|s}^2 \mathbf{I} \right)} [q_s(\alpha_{t|s}^{-1} x_t + \eta)] \\
&= \frac{\mathbb{E}_{\mathcal{N}_{\mathbb{R}^k} \left(\eta | 0, \alpha_{t|s}^2 \sigma_{t|s}^2 \mathbf{I} \right)} \left[\nabla_{\alpha_{t|s}^{-1} x_t} q_s(\alpha_{t|s}^{-1} x_t + \eta) \right]}{\alpha_{t|s} \mathbb{E}_{\mathcal{N}_{\mathbb{R}^k} \left(\eta | 0, \alpha_{t|s}^2 \sigma_{t|s}^2 \mathbf{I} \right)} [q_s(\alpha_{t|s}^{-1} x_t + \eta)]} \\
&= \frac{\mathbb{E}_{\mathcal{N}_{\mathbb{R}^k} \left(\eta | 0, \alpha_{t|s}^2 \sigma_{t|s}^2 \mathbf{I} \right)} [q_s(\alpha_{t|s}^{-1} x_t + \eta) \nabla_{\alpha_{t|s}^{-1} x_t + \eta} \log q_s(\alpha_{t|s}^{-1} x_t + \eta)]}{\alpha_{t|s} \mathbb{E}_{\mathcal{N}_{\mathbb{R}^k} \left(\eta | 0, \alpha_{t|s}^2 \sigma_{t|s}^2 \mathbf{I} \right)} [q_s(\alpha_{t|s}^{-1} x_t + \eta)]} \quad (15) \\
&= \frac{\mathbb{E}_{\mathcal{N}_{\mathbb{R}^k} \left(x_s | \alpha_{t|s}^{-1} x_t, \alpha_{t|s}^2 \sigma_{t|s}^2 \mathbf{I} \right)} [q_s(x_s) \nabla_{x_s} \log q_s(x_s)]}{\alpha_{t|s} \mathbb{E}_{\mathcal{N}_{\mathbb{R}^k} \left(x_s | \alpha_{t|s}^{-1} x_t, \alpha_{t|s}^2 \sigma_{t|s}^2 \mathbf{I} \right)} [q_s(x_s)]} \\
&= \frac{\int_{\mathbb{R}^k} \mathcal{N}_{\mathbb{R}^k} \left(x_t | \alpha_{t|s} x_s, \sigma_{t|s}^2 \mathbf{I} \right) q_s(x_s) \nabla_{x_s} \log q_s(x_s) dx_s}{\int_{\mathbb{R}^k} \mathcal{N}_{\mathbb{R}^k} \left(x_t | \alpha_{t|s} x_s, \sigma_{t|s}^2 \mathbf{I} \right) q_s(x_s) dx_s} \\
&= \frac{1}{\alpha_{t|s}} \mathbb{E}_{q_{st}(x_s|x_t)} [\nabla_{x_s} \log q_s(x_s)].
\end{aligned}$$

Note that when the transition probability $q_{st}(x_t|x_s)$ corresponds to a well-defined forward process, there is $\alpha_t > 0$ for $\forall t \in [0, T]$, and thus we achieve $\alpha_t \nabla_{x_t} \log q_t(x_t) = \mathbb{E}_{q_{st}(x_s|x_t)} [\alpha_s \nabla_{x_s} \log q_s(x_s)]$. \square

A.2 Proof of $\mathbb{E}_{q_0(x_0)} [\nabla_{x_0} \log q_0(x_0)] = 0$

Proof. The input variable $x \in \mathbb{R}^k$ and $q_0(x_0) \in \mathcal{C}^2$, where \mathcal{C}^2 denotes the family of functions with continuous second-order derivatives.¹ We use x^i denote the i -th element of x , then we can derive the expectation

$$\begin{aligned}
\mathbb{E}_{q_0(x_0)} \frac{\partial}{\partial x_0^i} \log q_0(x_0) &= \int_{\mathbb{R}^k} \cdots \int_{\mathbb{R}^k} q_0(x_0) \frac{\partial}{\partial x_0^i} \log q_0(x_0) dx_0^1 dx_0^2 \cdots dx_0^k \\
&= \int_{\mathbb{R}^k} \cdots \int_{\mathbb{R}^k} \frac{\partial}{\partial x_0^i} q_0(x_0) dx_0^1 dx_0^2 \cdots dx_0^k \\
&= \int_{\mathbb{R}^k} \frac{\partial}{\partial x_0^i} q_0(x_0^i, x_0^{\setminus i}) dx_0^{\setminus i} dx_0^i \\
&= \int_{\mathbb{R}^k} \frac{d}{dx_0^i} q_0(x_0^i) dx_0^i = 0,
\end{aligned} \quad (16)$$

¹This continuously differentiable assumption can be satisfied by adding a small Gaussian noise (e.g., with variance of 0.0001) on the original data distribution, as done in Song and Ermon [42].

where $x_0^{\setminus i}$ denotes all the $k - 1$ elements in x_0 except for the i -th one. The last equation holds under the boundary condition that $\lim_{x_0^i \rightarrow \infty} q_0(x_0^i) = 0$ hold for any $i \in [K]$. Thus, we achieve the conclusion that $\mathbb{E}_{q_0(x_0)} [\nabla_{x_0} \log q_0(x_0)] = 0$. \square

A.3 Concentration bounds

We describe concentration bounds [11, 1] of the martingale $\alpha_t \nabla_{x_t} \log q_t(x_t)$.

Azuma's inequality. For discrete reverse timestep $t = T, T - 1, \dots, 0$, Assuming that there exist constants $0 < c_1, c_2, \dots, < \infty$ such that for the i -th element of x ,

$$A_t \leq \frac{\partial}{\partial x_{t-1}^i} \alpha_{t-1} \log q_{t-1}(x_{t-1}) - \frac{\partial}{\partial x_t^i} \alpha_t \log q_t(x_t) \leq B_t \text{ and } B_t - A_t \leq c_t \quad (17)$$

almost surely. Then $\forall \epsilon > 0$, the probability (note that $\alpha_0 = 1$)

$$P \left[\frac{\partial}{\partial x_0^i} \log q_0(x_0) - \frac{\partial}{\partial x_T^i} \alpha_T \log q_T(x_T) \geq \epsilon \right] \leq 2 \exp \left[-\frac{2\epsilon^2}{\sum_{t=1}^T c_t^2} \right]. \quad (18)$$

Specially, considering that $q_T(x_T) \approx \mathcal{N}(x_T | 0, \Theta^2 \mathbf{I})$, there is $\frac{\partial}{\partial x_T^i} \log q_T(x_T) \approx -\frac{x_T^i}{\Theta^2}$. Thus, we can approximately obtain

$$P \left[\frac{\partial}{\partial x_0^i} \log q_0(x_0) + \frac{\alpha_T x_T^i}{\Theta^2} \geq \epsilon \right] \leq 2 \exp \left[-\frac{2\epsilon^2}{\sum_{t=1}^T c_t^2} \right]. \quad (19)$$

Doob's inequality. For continuous reverse timestep t from T to 0 , if the sample paths of the martingale are almost surely right-continuous, then for the i -th element of x we have (note that $\alpha_0 = 1$)

$$P \left[\sup_{0 \leq t \leq T} \frac{\partial}{\partial x_t^i} \alpha_t \log q_t(x_t) \geq C \right] \leq \frac{\mathbb{E}_{q_0(x_0)} \max \left[\frac{\partial}{\partial x_0^i} \log q_0(x_0), 0 \right]}{C}. \quad (20)$$

A.4 High-order SM objectives

Lu et al. [28] show that the KL divergence $\mathcal{D}_{\text{KL}}(q_0 \| p_0^{\text{ODE}}(\theta))$ can be bounded as

$$\mathcal{D}_{\text{KL}}(q_0 \| p_0^{\text{ODE}}(\theta)) \leq \mathcal{D}_{\text{KL}}(q_T \| p_T) + \frac{1}{\mathcal{J}_{\text{SM}}(\theta; g(t)^2)} \cdot \frac{1}{\mathcal{J}_{\text{Fisher}}(\theta)}, \quad (21)$$

where $\mathcal{J}_{\text{Fisher}}(\theta)$ is a weighted sum of Fisher divergence between $q_t(x_t)$ and $p_t^{\text{ODE}}(\theta)$ as

$$\mathcal{J}_{\text{Fisher}}(\theta) = \frac{1}{2} \int_0^T g(t)^2 D_F(q_t \| p_t^{\text{ODE}}(\theta)) dt. \quad (22)$$

Moreover, Lu et al. [28] prove that if $\forall t \in [0, T]$ and $\forall x_t \in \mathbb{R}^k$, there exist a constant C_F such that the spectral norm of Hessian matrix $\|\nabla_{x_t}^2 \log p_t^{\text{ODE}}(x_t; \theta)\|_2 \leq C_F$, and there exist $\delta_1, \delta_2, \delta_3 > 0$ such that

$$\begin{aligned} \|\mathbf{s}_\theta^t(x_t) - \nabla_{x_t} \log q_t(x_t)\|_2 &\leq \delta_1, \\ \|\nabla_{x_t} \mathbf{s}_\theta^t(x_t) - \nabla_{x_t}^2 \log q_t(x_t)\|_F &\leq \delta_2, \\ \|\nabla_{x_t} \mathbf{tr} \nabla_{x_t} \mathbf{s}_\theta^t(x_t) - \nabla_{x_t} \mathbf{tr} \nabla_{x_t}^2 \log q_t(x_t)\|_2 &\leq \delta_3, \end{aligned} \quad (23)$$

where $\|\cdot\|_F$ is the Frobenius norm of matrix. Then there exist a function $U(t; \delta_1, \delta_2, \delta_3, q)$ that independent of θ and strictly increasing (if $g(t) \neq 0$) w.r.t. δ_1, δ_2 , and δ_3 , respectively, such that the Fisher divergence can be bounded as $D_F(q_t \| p_t^{\text{ODE}}(\theta)) \leq U(t; \delta_1, \delta_2, \delta_3, q)$.

The case after calibration. When we impose the calibration term $\eta_t^* = \mathbb{E}_{q_t(x_t)} [\mathbf{s}_\theta^t(x_t)]$ to get the score model $\mathbf{s}_\theta^t(x_t) - \eta_t^*$, there is $\nabla_{x_t} \eta_t^* = 0$ and thus $\nabla_{x_t} (\mathbf{s}_\theta^t(x_t) - \eta_t^*) = \nabla_{x_t} \mathbf{s}_\theta^t(x_t)$. Then we have

$$\begin{aligned} \|\mathbf{s}_\theta^t(x_t) - \eta_t^* - \nabla_{x_t} \log q_t(x_t)\|_2 &\leq \delta'_1 \leq \delta_1, \\ \|\nabla_{x_t} \mathbf{s}_\theta^t(x_t) - \eta_t^* - \nabla_{x_t}^2 \log q_t(x_t)\|_F &\leq \delta_2, \\ \|\nabla_{x_t} \mathbf{tr} \nabla_{x_t} \mathbf{s}_\theta^t(x_t) - \eta_t^* - \nabla_{x_t} \mathbf{tr} \nabla_{x_t}^2 \log q_t(x_t)\|_2 &\leq \delta_3. \end{aligned} \quad (24)$$

From these, we know that the Fisher divergence $D_F q_t \| p_t^{\text{ODE}}(\theta, \eta_t^*) \leq U(t; \delta'_1, \delta_2, \delta_3, q) \leq U(t; \delta_1, \delta_2, \delta_3, q)$, namely, $D_F q_t \| p_t^{\text{ODE}}(\theta, \eta_t^*)$ has a lower upper bound compared to $D_F q_t \| p_t^{\text{ODE}}(\theta)$. Consequently, we can get lower upper bounds for both $\mathcal{J}_{\text{Fisher}}(\theta, \eta_t^*)$ and $\mathcal{D}_{\text{KL}} q_0 \| p_0^{\text{ODE}}(\theta, \eta_t^*)$, compared to $\mathcal{J}_{\text{Fisher}}(\theta)$ and $\mathcal{D}_{\text{KL}} q_0 \| p_0^{\text{ODE}}(\theta)$, respectively.

B Model parametrization

This section introduces different parametrizations used in diffusion models and provides their calibrated instantiations.

B.1 Preliminary

Along the research routine of diffusion models, different model parametrizations have been used, including score prediction $\mathbf{s}_\theta^t(x_t)$ [42, 46], noise prediction $\epsilon_\theta^t(x_t)$ [17, 33], data prediction $\mathbf{x}_\theta^t(x_t)$ [23, 32], and velocity prediction $\mathbf{v}_\theta^t(x_t)$ [38, 18]. Taking the DSM objective as the training loss, its instantiation at timestep $t \in [0, T]$ is written as

$$\mathcal{J}_{\text{DSM}}^t(\theta) = \begin{cases} \frac{1}{2} \mathbb{E}_{q_0(x_0), q(\epsilon)} \|\mathbf{s}_\theta^t(x_t) + \frac{\epsilon}{\sigma_t}\|_2^2, & \text{score prediction;} \\ \frac{\alpha_t^2}{2\sigma_t^4} \mathbb{E}_{q_0(x_0), q(\epsilon)} \|\mathbf{x}_\theta^t(x_t) - x_0\|_2^2, & \text{data prediction;} \\ \frac{1}{2\sigma_t^2} \mathbb{E}_{q_0(x_0), q(\epsilon)} \|\epsilon_\theta^t(x_t) - \epsilon\|_2^2, & \text{noise prediction;} \\ \frac{\alpha_t^2}{2\sigma_t^2} \mathbb{E}_{q_0(x_0), q(\epsilon)} \|\mathbf{v}_\theta^t(x_t) - (\alpha_t \epsilon - \sigma_t x_0)\|_2^2, & \text{velocity prediction.} \end{cases} \quad (25)$$

B.2 Calibrated instantiation

Under different model parametrizations, we can derive the optimal calibration terms η_t^* that minimizing $\mathcal{J}_{\text{DSM}}^t(\theta, \eta_t)$ as

$$\eta_t^* = \begin{cases} \mathbb{E}_{q_t(x_t)} [\mathbf{s}_\theta^t(x_t)], & \text{score prediction;} \\ \mathbb{E}_{q_t(x_t)} [\mathbf{x}_\theta^t(x_t)] - \mathbb{E}_{q_0(x_0)} [x_0], & \text{data prediction;} \\ \mathbb{E}_{q_t(x_t)} [\epsilon_\theta^t(x_t)], & \text{noise prediction;} \\ \mathbb{E}_{q_t(x_t)} [\mathbf{v}_\theta^t(x_t)] + \sigma_t \mathbb{E}_{q_0(x_0)} [x_0], & \text{velocity prediction.} \end{cases} \quad (26)$$

Taking η_t^* into $\mathcal{J}_{\text{DSM}}^t(\theta, \eta_t)$ we can obtain the gap

$$\mathcal{J}_{\text{DSM}}^t(\theta) - \mathcal{J}_{\text{DSM}}^t(\theta, \eta_t^*) = \begin{cases} \frac{1}{2} \|\mathbb{E}_{q_t(x_t)} [\mathbf{s}_\theta^t(x_t)]\|_2^2, & \text{score prediction;} \\ \frac{\alpha_t^2}{2\sigma_t^4} \|\mathbb{E}_{q_t(x_t)} [\mathbf{x}_\theta^t(x_t)] - \mathbb{E}_{q_0(x_0)} [x_0]\|_2^2, & \text{data prediction;} \\ \frac{1}{2\sigma_t^2} \|\mathbb{E}_{q_t(x_t)} [\epsilon_\theta^t(x_t)]\|_2^2, & \text{noise prediction;} \\ \frac{\alpha_t^2}{2\sigma_t^2} \|\mathbb{E}_{q_t(x_t)} [\mathbf{v}_\theta^t(x_t)] + \sigma_t \mathbb{E}_{q_0(x_0)} [x_0]\|_2^2, & \text{velocity prediction.} \end{cases} \quad (27)$$

C Visualization of the generations

We further show generated images in Figure 5 to double confirm the efficacy of our calibration method. Our calibration could help to reduce ambiguous generations on both CIFAR-10 and CelebA.



(a) CIFAR-10, w/ calibration (b) CIFAR-10, w/o calibration (c) CelebA, w/ calibration (d) CelebA, w/o calibration

Figure 5: Unconditional generation results on CIFAR-10 and CelebA using models from [17] and [41] respectively. The number of sampling steps is 20 based on the results in Tables 1 and 2.