

DENOISING LIKELIHOOD SCORE MATCHING FOR CONDITIONAL SCORE-BASED DATA GENERATION

Chen-Hao Chao*
National Tsing Hua University

Wei-Fang Sun
National Tsing Hua University

Bo-Wun Cheng
National Tsing Hua University

Yi-Chen Lo
Mediatek Inc.

Chia-Che Chang
Mediatek Inc.

Yu-Lun Liu
Mediatek Inc.

Yu-Lin Chang
Mediatek Inc.

Chia-Ping Chen
Mediatek Inc.

Chun-Yi Lee[†]
National Tsing Hua University

ABSTRACT

Many existing conditional score-based data generation methods utilize Bayes' theorem to decompose the gradients of a log posterior density into a mixture of scores. These methods facilitate the training procedure of conditional score models, as a mixture of scores can be separately estimated using a score model and a classifier. However, our analysis indicates that the training objectives for the classifier in these methods may lead to a serious *score mismatch issue*, which corresponds to the situation that the estimated scores deviate from the true ones. Such an issue causes the samples to be misled by the deviated scores during the diffusion process, resulting in a degraded sampling quality. To resolve it, we formulate a novel training objective, called *Denoising Likelihood Score Matching* (DLSM) loss, for the classifier to match the gradients of the true log likelihood density. Our experimental evidences show that the proposed method outperforms the previous methods on both Cifar-10 and Cifar-100 benchmarks noticeably in terms of several key evaluation metrics. We thus conclude that, by adopting DLSM, the conditional scores can be accurately modeled, and the effect of the score mismatch issue is alleviated.

1 INTRODUCTION

Score-based generative models are probabilistic generative models that estimate score functions, i.e., the gradients of the log density for some given data distribution. As described in the pioneering work (Hyvärinen, 2005), the process of training score-based generative models is called *Score Matching* (SM), in which a score-based generative model is iteratively updated to approximate the true score function. Such a process often incurs heavy computational burdens, since the calculation of the score-matching objective involves the explicit computation of the partial derivatives of the score model during training. Therefore, a branch of study in this research domain (Vincent, 2011; Song et al., 2019) resorts to reformulating the score-matching objective to reduce the training cost. Among these works, the author in (Vincent, 2011) introduced the *Denoising Score-Matching* (DSM) method. This method facilitates the training process of score-based generative models, and thus lays the foundation for a number of subsequent researches. Recently, the authors in (Song & Ermon, 2019) proposed a unified framework based on DSM, and achieved remarkable performance on several real-world datasets. Their success inspired several succeeding works (Song & Ermon, 2020; Ho et al., 2020; Song et al., 2021a;b; Dhariwal & Nichol, 2021), which together contribute to making score-based generative models an attractive choice for contemporary image generation tasks.

A favorable aspect of score-based generative models is their flexibility to be easily extended to conditional variants. This characteristic comes from a research direction that utilizes Bayes' theorem

*Work done during an internship at Mediatek Inc. E-mail: lance_chao@gapp.nthu.edu.tw

[†]Corresponding author. E-mail: cylee@gapp.nthu.edu.tw

to decompose a conditional score into a mixture of scores (Nguyen et al., 2017). Recent endeavors followed this approach and further extended the concept of conditional score-based models to a number of application domains, including colorization (Song et al., 2021b), inpainting (Song et al., 2021b), and source separation (Jayaram & Thickstun, 2020). In particular, some recent researchers (Song et al., 2021b; Dhariwal & Nichol, 2021) applied this method to the field of class-conditional image generation tasks, and proposed the *classifier-guidance* method. Different from the *classifier-guidance-free* method adopted by (Ho et al., 2021), they utilized a score model and a classifier to generate the posterior scores (i.e., the gradients of the log posterior density), with which the data samples of certain classes can be generated through the diffusion process. The authors in (Dhariwal & Nichol, 2021) showed that the classifier guidance method is able to achieve improved performance on large image generation benchmarks. In spite of their success, our analysis indicates that the conditional generation methods utilizing a score model and a classifier may suffer from a *score mismatch issue*, which is the situation that the estimated posterior scores deviate from the true ones. This issue causes the samples to be guided by inaccurate scores during the diffusion process, and may result in a degraded sampling quality consequently.

To resolve this problem, we first analyze the potential causes for the score mismatch issue through a motivational low-dimensional example. Then, we formulate a new loss function called *Denoising Likelihood Score-Matching* (DLSM) loss, and explain how it can be integrated into the current training method. Finally, we evaluate the proposed method under various configurations, and demonstrate its advantages in improving the sampling quality over the previous methods in terms of several evaluation metrics.

2 BACKGROUND

In this section, we introduce the essential background material for understanding the contents of this paper. We first introduce Langevin diffusion (Roberts & Tweedie, 1996; Roberts & Rosenthal, 1998) for generating data samples $\tilde{\mathbf{x}} \in \mathbb{R}^d$ of a certain dimension d from an unknown probability density function (pdf) $p(\tilde{\mathbf{x}})$ using the score function $\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{x}})$. Next, we describe Parzen density estimation (Parzen, 1962) and denoising score matching (DSM) (Vincent, 2011) for approximating $\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{x}})$ with limited data samples. Finally, we elaborate on the conditional variant of the score function, i.e., $\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{x}}|\tilde{\mathbf{y}})$, and explains how it can be decomposed into $\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{x}})$ and $\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}})$ for some conditional variable $\tilde{\mathbf{y}} \in \mathbb{R}^c$ of dimension c .

2.1 LANGEVIN DIFFUSION

Langevin diffusion (Roberts & Tweedie, 1996; Roberts & Rosenthal, 1998) can be used to generate data samples from an unknown data distribution $p(\tilde{\mathbf{x}})$ using only the score function $\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{x}})$, which is said to be well-defined if $p(\tilde{\mathbf{x}})$ is everywhere non-zero and differentiable. Under the condition that $\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{x}})$ is well-defined, Langevin diffusion enables $p(\tilde{\mathbf{x}})$ to be approximated iteratively based on the following equation:

$$\tilde{\mathbf{x}}_t = \tilde{\mathbf{x}}_{t-1} + \frac{\epsilon^2}{2} \nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{x}}_{t-1}) + \epsilon \mathbf{z}_t, \quad (1)$$

where $\tilde{\mathbf{x}}_0$ is sampled from an arbitrary distribution, ϵ is a fixed positive step size, and \mathbf{z}_t is a noise vector sampled from a normal distribution $\mathcal{N}(\mathbf{0}, I_{d \times d})$ for simulating a d -dimensional standard Brownian motion. Under suitable regularity conditions, when $\epsilon \rightarrow 0$ and $T \rightarrow \infty$, $\tilde{\mathbf{x}}_T$ is generated as if it is directly sampled from $p(\tilde{\mathbf{x}})$ (Roberts & Tweedie, 1996; Welling & Teh, 2011). In practice, however, the data samples are generated with $\epsilon > 0$ and $T < \infty$, which violates the convergence guarantee. Although it is possible to use Metropolized algorithms (Roberts & Tweedie, 1996; Roberts & Rosenthal, 1998) to recover the convergence guarantee, we follow the assumption of the prior work (Song & Ermon, 2019) and presume that the errors are sufficiently small to be negligible when ϵ is small and T is large.

The sampling process introduced in Eq. (1) can be extended to a time-inhomogeneous variant by making $p(\tilde{\mathbf{x}}_t)$ and ϵ dependent on t (i.e., $p_t(\tilde{\mathbf{x}}_t)$ and ϵ_t). Such a time-inhomogeneous variant is commonly adopted by recent works on score-based generative models (Song & Ermon, 2019; Song et al., 2021b), as it provides flexibility in controlling $p_t(\tilde{\mathbf{x}}_t)$ and ϵ_t . The experimental results in (Song & Ermon, 2019; Song et al., 2021b) demonstrated that such a time-inhomogeneous sampling process

can improve the sampling quality on real-world datasets. In Appendix A.6.1, we offer a discussion on the detailed implementation of such a time-inhomogeneous sampling process in this work.

2.2 PARZEN DENSITY ESTIMATION

Given a true data distribution p_{data} , the empirical data distribution $p_0(\mathbf{x})$ is constructed by sampling m independent and identically distributed data points $\{\mathbf{x}^{(i)}\}_{i=1}^m$, and can be represented as a sum of Dirac functions $\frac{1}{m} \sum_{i=1}^m \delta(\|\mathbf{x} - \mathbf{x}^{(i)}\|)$. Such a discrete data distribution $p_0(\mathbf{x})$ constructed from the dataset often violates the previous assumptions that everywhere is non-zero and is differentiable. Therefore, it is necessary to somehow adjust the empirical data distribution $p_0(\mathbf{x})$ before applying Langevin diffusion in such cases.

To deal with the above issue, a previous literature (Vincent, 2011) utilized Parzen density estimation to replace the Dirac functions with isotropic Gaussian smoothing kernels $p_\sigma(\tilde{\mathbf{x}}|\mathbf{x}) = \frac{1}{(2\pi)^{d/2}\sigma^d} e^{-\frac{1}{2\sigma^2}\|\tilde{\mathbf{x}}-\mathbf{x}\|^2}$ with variance σ^2 . Specifically, Parzen density estimation enables the calculation of $p_\sigma(\tilde{\mathbf{x}}) = \frac{1}{m} \sum_{i=1}^m p_\sigma(\tilde{\mathbf{x}}|\mathbf{x}^{(i)})$. When $\sigma > 0$, the score function becomes well-defined and can thus be represented as the following:

$$\nabla_{\tilde{\mathbf{x}}} \log p_\sigma(\tilde{\mathbf{x}}) = \frac{\sum_{i=1}^m \frac{1}{\sigma^2} (\mathbf{x}^{(i)} - \tilde{\mathbf{x}}) p_\sigma(\tilde{\mathbf{x}}|\mathbf{x}^{(i)})}{\sum_{i=1}^m p_\sigma(\tilde{\mathbf{x}}|\mathbf{x}^{(i)})}. \quad (2)$$

The proof for Eq. (2) is provided in Appendix A.2. This equation can be directly applied to Eq. (1) to generate samples with Langevin diffusion. Unfortunately, this requires summation over all m data points during every iteration, preventing it from scaling to large datasets due to the rapid growth in computational complexity.

2.3 DENOISING SCORE MATCHING

Score matching (SM) (Hyvärinen, 2005) was proposed to estimate the score function with a model $s(\tilde{\mathbf{x}}; \phi)$, parameterized by ϕ . Given a trained score model $s(\tilde{\mathbf{x}}; \phi)$, the scores can be generated by a single forward pass, which reduces the computational complexity of Eq. (2) by a factor of m . To train such a score model, a straightforward approach is to use the Explicit Score-Matching (ESM) loss L_{ESM} , represented as:

$$L_{\text{ESM}}(\phi) = \mathbb{E}_{p_\sigma(\tilde{\mathbf{x}})} \left[\frac{1}{2} \|s(\tilde{\mathbf{x}}; \phi) - \nabla_{\tilde{\mathbf{x}}} \log p_\sigma(\tilde{\mathbf{x}})\|^2 \right]. \quad (3)$$

This objective requires evaluating Eq. (2) for each training step, which also fails to scale well to large datasets. Based on Parzen density estimation, an efficient alternative, called Denoising Score-Matching (DSM) loss (Vincent, 2011), is proposed to efficiently calculate the equivalent loss L_{DSM} , expressed as:

$$L_{\text{DSM}}(\phi) = \mathbb{E}_{p_\sigma(\tilde{\mathbf{x}}, \mathbf{x})} \left[\frac{1}{2} \|s(\tilde{\mathbf{x}}; \phi) - \nabla_{\tilde{\mathbf{x}}} \log p_\sigma(\tilde{\mathbf{x}}|\mathbf{x})\|^2 \right]. \quad (4)$$

where $\nabla_{\tilde{\mathbf{x}}} \log p_\sigma(\tilde{\mathbf{x}}|\mathbf{x})$ is simply $\frac{1}{\sigma^2}(\mathbf{x} - \tilde{\mathbf{x}})$. Since the computational cost of denoising score matching is relatively lower in comparison to other reformulation techniques (Hyvärinen, 2005; Song et al., 2019), it is extensively adopted in recent score-based generative models (Song & Ermon, 2019; 2020; Song et al., 2021b).

2.4 CONDITIONAL SCORE DECOMPOSITION VIA BAYES' THEOREM

Score models can be extended to conditional models when conditioned on a certain label $\tilde{\mathbf{y}}$. Similar to $\tilde{\mathbf{x}}$, the smoothing kernels with variance τ^2 can be applied on \mathbf{y} to meet the requirement that the pdf is everywhere non-zero. Typically, τ is assumed to be sufficiently small so that $p_\tau(\tilde{\mathbf{y}}) \approx p(\mathbf{y})$. A popular approach adopted by researchers utilizes Bayes' theorem $p_{\sigma, \tau}(\tilde{\mathbf{x}}|\tilde{\mathbf{y}}) = p_{\sigma, \tau}(\tilde{\mathbf{y}}|\tilde{\mathbf{x}})p_\sigma(\tilde{\mathbf{x}})/p_\tau(\tilde{\mathbf{y}})$ to decompose the conditional score $\nabla_{\tilde{\mathbf{x}}} \log p_{\sigma, \tau}(\tilde{\mathbf{x}}|\tilde{\mathbf{y}})$ into a mixture of scores (Nguyen et al., 2017), which enables conditional data generation. Following the assumptions in the previous study (Song et al., 2021b), the decomposition can be achieved by taking the log-gradient on both sides of the equation, expressed as follows:

$$\nabla_{\tilde{\mathbf{x}}} \log p_{\sigma, \tau}(\tilde{\mathbf{x}}|\tilde{\mathbf{y}}) = \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma, \tau}(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}) + \nabla_{\tilde{\mathbf{x}}} \log p_\sigma(\tilde{\mathbf{x}}) - \underbrace{\nabla_{\tilde{\mathbf{x}}} \log p_\tau(\tilde{\mathbf{y}})}_{=0}, \quad (5)$$

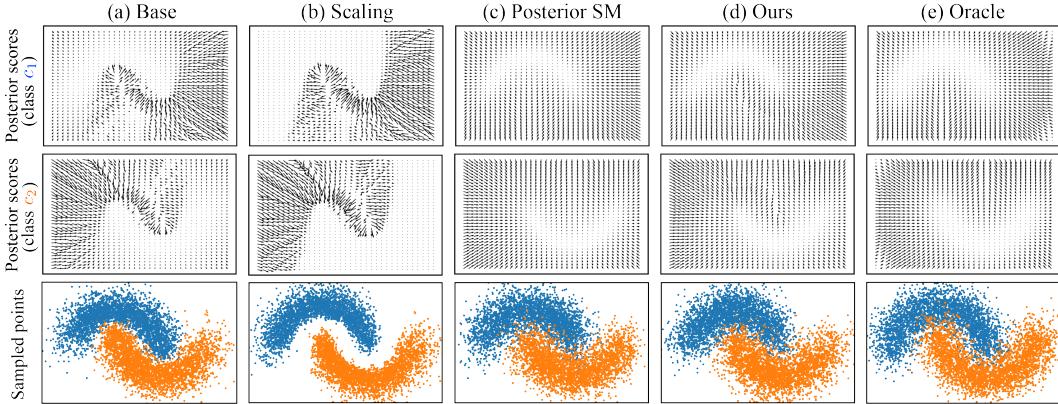


Figure 1: The visualized results on the *inter-twining moon* dataset. The plots presented in the first two rows correspond to the visualized vector fields for the posterior scores of the class c_1 (upper crescent) and c_2 (lower crescent), respectively. The plots in the third row are the sampled points. Different columns correspond to different experimental settings in Section 3. The detailed configurations are presented in Appendix A.6.2.

where $\nabla_{\tilde{\mathbf{x}}} \log p_{\sigma, \tau}(\tilde{\mathbf{x}}|\tilde{\mathbf{y}})$ is the posterior score, $\nabla_{\tilde{\mathbf{x}}} \log p_{\sigma, \tau}(\tilde{\mathbf{y}}|\tilde{\mathbf{x}})$ is the likelihood score, and $\nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}})$ is the prior score. Base on Eq. (5), a conditional score model $\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{x}}|\tilde{\mathbf{y}}; \phi, \theta)$ can be represented as the combination of the log-gradient of a differentiable classifier $p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta)$ and a prior score model $s(\tilde{\mathbf{x}}; \phi)$:

$$\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{x}}|\tilde{\mathbf{y}}; \theta, \phi) = \nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta) + s(\tilde{\mathbf{x}}; \phi). \quad (6)$$

This formulation enables conditional data generation using a classifier $p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta)$ trained with cross-entropy loss $L_{CE}(\theta) \triangleq \mathbb{E}_{p_{\sigma, \tau}(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})}[-\log p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta)]$, and a score model $s(\tilde{\mathbf{x}}; \phi)$ trained with denoising score matching loss $L_{DSM}(\phi)$. Unfortunately, a few previous studies (Nguyen et al., 2017; Dhariwal & Nichol, 2021) have noticed that the approximation of $\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{x}}|\tilde{\mathbf{y}}; \theta, \phi)$ is empirically inaccurate (referred to as the score mismatch issue in this work), and leveraged a scaling factor $\alpha > 0$ to adjust the likelihood score $\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta)$. Such a scaling factor is a hyperparameter that empirically controls the amount of conditional information incorporated during the sampling process (see Appendix A.3). However, this usually causes the diversity of generated data samples to degrade noticeably (Dhariwal & Nichol, 2021), which is later discussed in Section 5.2.

In this work, we examine and investigate the score mismatch issue from a different perspective. We first offer a motivational example to show that the issue remains even with the scaling technique. Then, we theoretically formulate a new loss term that enables a classifier $p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta)$ to produce better likelihood score estimation.

3 ANALYSIS ON THE SCORE MISMATCH ISSUE

To further investigate the score mismatch issue, we first leverage a motivational experiment on the *inter-twining moon* dataset to examine the extent of the discrepancy between the estimated and true posterior scores. In this experiment, we consider five different methods to calculate the posterior scores denoted as (a)~(e):

- **(a) Base method.** The posterior scores are estimated using a score model $s(\tilde{\mathbf{x}}; \phi)$ trained with L_{DSM} and a classifier $p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta)$ trained with L_{CE} , as described in Eq. (6).
- **(b) Scaling method.** The posterior scores are estimated in a similar fashion as method (a) except that the scaling technique (Dhariwal & Nichol, 2021) is adopted to scale the likelihood scores, i.e., $\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{x}}|\tilde{\mathbf{y}}; \theta, \phi) = \alpha \nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta) + s(\tilde{\mathbf{x}}; \phi)$, where $\alpha > 0$.
- **(c) Posterior SM method.** The posterior scores for different class conditions, in this case, c_1 and c_2 , are separately estimated using different score models $s(\tilde{\mathbf{x}}; \phi_1)$ and $s(\tilde{\mathbf{x}}; \phi_2)$ trained with L_{DSM} . Note that, since this method requires network architectures different from

Table 1: The experimental results on the inter-twining moon dataset. The quality of the sampled data for different methods are measured in terms of the precision and recall metrics. The errors of the score functions for different methods are measured using $\mathbb{E}_{U(\tilde{\mathbf{x}})}[D_P(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})]$ and $\mathbb{E}_{U(\tilde{\mathbf{x}})}[D_L(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})]$, where $U(\tilde{\mathbf{x}})$ is a uniform distribution over a two-dimensional subspace (i.e., the space of each subplot in Fig. 1), $\tilde{\mathbf{y}}$ represents the classes specified in the second row of the table, and D_P and D_L are defined in Eq. (7). The configurations and the hyperparameters for different experimental settings of this experiment are presented in Appendix A.6.2.

	(a) Base		(b) Scaling		(c) Posterior SM		(d) Ours		(e) Oracle	
	c_1	c_2	c_1	c_2	c_1	c_2	c_1	c_2	c_1	c_2
$S_P(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})$	$\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}} \tilde{\mathbf{x}}; \theta) + s(\tilde{\mathbf{x}}; \phi)$		$\alpha \nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}} \tilde{\mathbf{x}}; \theta) + s(\tilde{\mathbf{x}}; \phi)$		$s(\tilde{\mathbf{x}}; \phi_1)$	$s(\tilde{\mathbf{x}}; \phi_2)$	$\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}} \tilde{\mathbf{x}}; \theta) + s(\tilde{\mathbf{x}}; \phi)$		$\nabla_{\tilde{\mathbf{x}}} \log p_{\sigma, \tau}(\tilde{\mathbf{x}} \tilde{\mathbf{y}})$	
$S_L(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})$	$\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}} \tilde{\mathbf{x}}; \theta)$		$\alpha \nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}} \tilde{\mathbf{x}}; \theta)$		-		$\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}} \tilde{\mathbf{x}}; \theta)$		$\nabla_{\tilde{\mathbf{x}}} \log p_{\sigma, \tau}(\tilde{\mathbf{y}} \tilde{\mathbf{x}})$	
$\mathbb{E}_{U(\tilde{\mathbf{x}})}[D_P(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})]$	0.198	0.197	2.734	2.627	0.063	0.060	0.066	0.064	0.000	0.000
$\mathbb{E}_{U(\tilde{\mathbf{x}})}[D_L(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})]$	0.190	0.181	2.730	2.618	-	-	0.052	0.052	0.000	0.000
Precision	0.94	0.93	0.95	0.94	0.97	0.97	0.96	0.95	1.00	1.00
Recall	1.00	1.00	0.97	0.97	1.00	1.00	1.00	1.00	1.00	1.00

those used in methods (a), (b), and (d), it is only adopted for analytical purposes. For more detailed discussions about this method, please refer to Appendix A.4.

- **(d) Ours.** The posterior scores are estimated in a similar fashion as method (a) except that the classifier is trained with the proposed loss function L_{Total} , which is described in detail in Section 4.
- **(e) Oracle.** The oracle posterior scores are directly computed using the conditional variant of Eq. (2), which is detailed in Appendix A.6.2.

Fig. 1 visualizes the posterior scores and the sampled points based on the five methods. It is observed that the posterior scores estimated using methods (a) and (b) are significantly different from the true posterior scores measured by method (e). This causes the sampled points in methods (a) and (b) to deviate from those sampled based on method (e). On the other hand, the estimated posterior scores and the sampled points in method (c) are relatively similar to those in method (e). The above results therefore suggest that the score mismatch issue is severe under the cases of methods (a) and (b), but is alleviated when method (c) is used.

In order to inspect the potential causes for the differences between the results produced by methods (a), (b), and (c), we incorporate metrics for evaluating the sampling quality and the errors between the scores in an quantitative manner. The sampling quality is evaluated using the precision and recall (Kynkäänniemi et al., 2019) metrics. On the other hand, the estimation errors of the score functions are measured by the expected values of D_P and D_L , which are formulated according to the Euclidean distances between the estimated scores and the oracle scores. The expressions of D_P and D_L are represented as the following:

$$\begin{aligned} D_P(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) &= \|S_P(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) - \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma, \tau}(\tilde{\mathbf{x}}|\tilde{\mathbf{y}})\|, \\ D_L(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) &= \|S_L(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) - \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma, \tau}(\tilde{\mathbf{y}}|\tilde{\mathbf{x}})\|, \end{aligned} \quad (7)$$

where the subscripts P and L denote ‘Posterior’ and ‘Likelihood’, while the terms $S_P(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})$ and $S_L(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})$ correspond to the estimated posterior score and the likelihood score for a certain pair $(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})$, respectively.

Table 1 presents $S_P(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})$ and $S_L(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})$, the expectations of D_P and D_L , and the precision and recall for different methods. It can be seen that the numerical results in Table 1 are consistent with the observations revealed in Fig. 1, since the expectations of D_P are greater in methods (a) and (b) as compared to method (c), and the evaluated recall values in methods (a) and (b) are lower than those in method (c). Furthermore, the results in Table 1 suggest that the possible reasons for the disparity of the posterior scores between methods (a)~(c) are twofold. First, adding the scaling factor α to the likelihood scores increases the expected distance between the estimated score and oracle score, which may exacerbate the score mismatch issue. Second, since the main difference between methods (a)~(c) lies in their score-matching objectives, i.e., the parameters θ in methods (a) and (b) are optimized using L_{CE} only while the parameters ϕ_1 and ϕ_2 of the score models in method (c) are

optimized through L_{DSM} , the adoption of the score-matching objective may potentially be the key factor to the success of method (c).

The above experimental clues therefore shed light on two essential issues to be further explored and dealt with. First, although employing a classifier trained with L_{CE} to assist estimating the oracle posterior score is theoretically feasible, this method may potentially lead to considerable discrepancies in practice. This implies that the score mismatch issue stated in Section 2.4 may be the result of the inaccurate likelihood scores produced by a classifier. Second, the comparisons between methods (a), (b), and (c) suggest that score matching may potentially be the solution to the score mismatch issue. Based on these hypotheses, this paper explores an alternative approach to the previous works, called denoising likelihood score matching, which incorporates the score-matching technique when training the classifier to enhance its capability to capture the true likelihood scores. In the next section, we detail the theoretical derivation and the property of our method.

4 DENOISING LIKELIHOOD SCORE MATCHING

In this section, we introduce the proposed denoising likelihood score-matching (DLSM) loss, a new training objective that encourages the classifier to capture the true likelihood score. In Section 4.1, we derive this loss function and discuss the intuitions behind it. In Section 4.2, we elaborate on the training procedure of DLSM, and highlight the key property and implications of this training objective.

4.1 THE DERIVATION OF THE DENOISING LIKELIHOOD SCORE MATCHING LOSS

As discussed in Section 3, a score model trained with the score-matching objective can potentially be beneficial in producing a better posterior score estimation. In light of this, a classifier may be enhanced if the score-matching process is involved during its training procedure. An intuitive way to accomplish this aim is through minimizing the explicit likelihood score-matching loss L_{ELSM} , which is defined as the following:

$$L_{\text{ELSM}}(\theta) = \mathbb{E}_{p_{\sigma,\tau}(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})} \left[\frac{1}{2} \|\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta) - \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma,\tau}(\tilde{\mathbf{y}}|\tilde{\mathbf{x}})\|^2 \right]. \quad (8)$$

This loss term, however, involves the calculation of the true likelihood score $\nabla_{\tilde{\mathbf{x}}} \log p_{\sigma,\tau}(\tilde{\mathbf{y}}|\tilde{\mathbf{x}})$, whose computational cost grows with respect to the dataset size. In order to reduce the computational cost, we follow the derivation of DSM as well as Bayes’ theorem, and formulate an alternative objective called DLSM loss (L_{DLSM}):

$$L_{\text{DLSM}}(\theta) = \mathbb{E}_{p_{\sigma,\tau}(\tilde{\mathbf{x}}, \mathbf{x}, \tilde{\mathbf{y}}, \mathbf{y})} \left[\frac{1}{2} \|\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta) + \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}}) - \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}}|\mathbf{x})\|^2 \right]. \quad (9)$$

Theorem 1. $L_{\text{DLSM}}(\theta) = L_{\text{ELSM}}(\theta) + C$, where C is a constant with respect to θ .

The proof of this theorem is provided in Appendix A.5. Theorem 1 suggests that optimizing θ with $L_{\text{ELSM}}(\theta)$ is equivalent to optimizing θ with $L_{\text{DLSM}}(\theta)$. In contrast to L_{ELSM} , L_{DLSM} can be approximated in a computationally feasible fashion. This is because $\nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}})$ can be estimated using a score model $s(\tilde{\mathbf{x}}; \phi)$ trained with L_{DSM} , and $\nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}}|\mathbf{x}) = \frac{1}{\sigma^2}(\mathbf{x} - \tilde{\mathbf{x}})$ as described in Section 2.3. As $\nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}})$ and $\nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}}|\mathbf{x})$ can be computed in a tractable manner, the classifier can be updated by minimizing the approximated variant of L_{DLSM} , defined as:

$$L_{\text{DLSM}'}(\theta) = \mathbb{E}_{p_{\sigma,\tau}(\tilde{\mathbf{x}}, \mathbf{x}, \tilde{\mathbf{y}}, \mathbf{y})} \left[\frac{1}{2} \|\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta) + s(\tilde{\mathbf{x}}; \phi) - \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}}|\mathbf{x})\|^2 \right]. \quad (10)$$

The underlying intuition of L_{DLSM} and $L_{\text{DLSM}'}$ is to match the likelihood score via matching the posterior score. More specifically, the sum of the first two terms $\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta)$ and $\nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}})$ in L_{DLSM} should ideally construct the posterior score (i.e., Eq. (5)). By following the posterior score, the perturbed data sample $\tilde{\mathbf{x}}$ should move towards the clean sample \mathbf{x} , since $\nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}}|\mathbf{x})$ is equal to the weighted difference $\frac{1}{\sigma^2}(\mathbf{x} - \tilde{\mathbf{x}})$ (i.e., Eq. (4)). In the next section, we discuss the training procedure with $L_{\text{DLSM}'}$ as well as the property of it.




Figure 2: The training procedure of the proposed methodology.

4.2 THE APPROXIMATED DLSM OBJECTIVE IN THE CLASSIFIER TRAINING PROCESS

Following the theoretical derivation in Section 4.1, we next discuss the practical aspects during training, and propose to train the classifier by jointly minimizing the approximated denoising likelihood score-matching loss $L_{\text{DLSM}'}$ and the cross-entropy loss L_{CE} . In practice, the total training objective of the classifier can be written as follows:

$$L_{\text{Total}}(\theta) = L_{\text{DLSM}' }(\theta) + \lambda L_{\text{CE}}(\theta), \quad (11)$$

where $\lambda > 0$ is a balancing coefficient. According to Theorem 1, employing $L_{\text{DLSM}'}$ allows the gradients of the classifier to estimate the likelihood scores. However, since the computation of the term $L_{\text{DLSM}'}$ requires the approximation of $\nabla_{\tilde{x}} \log p_{\sigma}(\tilde{x})$ from a score model $s(\tilde{x}; \phi)$, the training process that utilizes $L_{\text{DLSM}'}$ alone is undesirable due to its instability on real-world datasets. To reinforce the stability of it, an additional cross-entropy loss L_{CE} is adopted during the training process of the classifier. The advantage of L_{CE} is that it leverages the ground truth labels to assist the classifier to learn to match the true likelihood density $p_{\sigma, \tau}(\tilde{y}|\tilde{x})$, which in turn helps the score estimation to be improved. To validate the above advantage, an ablation analysis are offered in Section 5.3 to demonstrate the differences between the classifiers trained with L_{Total} , L_{CE} , and $L_{\text{DLSM}'}$. The analysis reveals that L_{Total} does provide the best score-matching results while maintaining the stability.

Fig. 2 depicts a two-stage training procedure adopted in this work. In stage 1, a score model $s(\tilde{x}; \phi)$ is updated using $L_{\text{DLSM}'}$ to match $\nabla_{\tilde{x}} \log p_{\sigma}(\tilde{x})$ (i.e., Eq. (4)). In stage 2, the weights of the trained score model are fixed, and a classifier $p(\tilde{y}|\tilde{x}; \theta)$ is updated using L_{Total} . After these two training stages, $s(\tilde{x}; \phi)$ and $\nabla_{\tilde{x}} \log p(\tilde{y}|\tilde{x}; \theta)$ can then be added together based on Eq. (6) to perform conditional sampling.

5 EXPERIMENTS

In this section, we present experiments to validate the effectiveness of the proposed method. First, we provide the experimental configurations and describe the baseline methods. Next, we evaluate the proposed method quantitatively against the baselines, and demonstrate its improved performance on the Cifar-10 and Cifar-100 datasets. Lastly, we offer an ablation analysis to inspect the impacts of L_{CE} and $L_{\text{DLSM}'}$ in L_{Total} .

5.1 EXPERIMENTAL SETUPS

In this work, we quantitatively compare the base method, scaling method, and our method (i.e., (a), (b), and (d) in Section 3) on the Cifar-10 and Cifar-100 datasets, which contain real-world RGB images with 10 and 100 categories of objects, respectively. For both of the datasets, the image size is $32 \times 32 \times 3$, and the pixel values in an image are first rescaled to $[-1, 1]$ before training. For the sampling algorithm, we adopt the predictor-corrector (PC) sampler described in (Song et al., 2021b) with the sampling steps set to $T=1,000$. The score model architecture is exactly the same as the one used in (Song et al., 2021b), while the architecture of the classifier is based on ResNet (He et al., 2016) with a conditional branch for encoding the information of the standard deviation σ (Song et al., 2021b). When training the classifier, τ is assumed to be sufficiently small, thus the unperturbed labels \mathbf{y} are adopted. Furthermore, the balancing coefficient λ is set to 1 for the Cifar-10 and Cifar-100 datasets, and 0.125 for the motivational example. For more details about our implementation and additional experimental results, please refer to Appendix A.6 and Appendix A.7.

Table 2: The evaluation results on the Cifar-10 and Cifar-100 datasets. The P/R/D/C metrics with ‘(CW)’ in the last four rows represents the average class-wise metrics described in Section 5.2. The arrow symbols \uparrow / \downarrow represent that a higher / lower evaluation result correspond to a better performance.

		Cifar-10			Cifar-100		
		Base	Scaling	Ours	Base	Scaling	Ours
FID	\downarrow	4.10	12.48	2.25	4.52	12.58	3.86
IS	\uparrow	9.08	9.37	9.90	11.53	11.59	11.62
Precision	\uparrow	0.67	0.75	0.65	0.62	0.65	0.61
Recall	\uparrow	0.61	0.49	0.62	0.62	0.52	0.63
Density	\uparrow	1.05	1.36	0.96	0.84	0.93	0.82
Coverage	\uparrow	0.80	0.75	0.81	0.71	0.61	0.71
CAS	\uparrow	0.38	0.46	0.58	0.15	0.24	0.28
(CW) Precision	\uparrow	0.51	0.70	0.56	0.33	0.59	0.42
(CW) Recall	\uparrow	0.59	0.42	0.61	0.68	0.41	0.68
(CW) Density	\uparrow	0.63	1.23	0.76	0.30	0.78	0.43
(CW) Coverage	\uparrow	0.60	0.66	0.71	0.38	0.51	0.52



Figure 3: A comparison of the curated samples generated via the base method and the scaling method. As for the uncurated visualized samples, please refer to Appendix A.7.4.

5.2 RESULTS ON CIFAR-10 AND CIFAR-100

In this section, we examine the effectiveness of the base method, the scaling method, and our proposed method on the Cifar-10 and Cifar-100 benchmarks with several key evaluation metrics. We adopt the Inception Score (IS) (Barratt & Sharma, 2018) and the Fréchet Inception Distance (FID) (Heusel et al., 2017) as the metrics for evaluating the overall sampling quality by comparing the similarity between the distributions of the generated images and the real images. We also evaluate the methods using the Precision (P), Recall (R) (Kynkäänniemi et al., 2019), Density (D), and Coverage (C) (Naeem et al., 2020) metrics to further examine the fidelity and diversity of the generated images. In addition, we report the Classification Accuracy Score (CAS) (Ravuri & Vinyals, 2019) to measure if the generated samples bear representative class information. Given a dataset containing m_i images for each class i , we first conditionally generate the same number of images (i.e., m_i) for each class. The FID, IS, P/R/D/C, and CAS metrics are then evaluated based on all the generated $m = \sum_{i=1}^c m_i$ images, where c is the total number of classes. In other words, the above metrics are evaluated in an unconditional manner.

Table 2 reports the quantitative results of the above methods. It is observed that the proposed method outperforms the other two methods with substantial margins in terms of FID and IS, indicating that the generated samples bear closer resemblance to the real data. Meanwhile, for the P/R/D/C metrics, the scaling method is superior to the other two methods in terms of the fidelity metrics (i.e., precision and density). However, this method may cause the diversity of the generated images to degrade, as depicted in Fig. 3, resulting in significant performance drops for the diversity metrics (i.e., the recall and the coverage metrics).

Another insight is that the base method achieves relatively better performance on the precision and density metrics in comparison to our method. However, it fails to deliver analogous tendency on the CAS metric. This behavior indicates that the base method may be susceptible to generating false positive samples, since the evaluation of the P/R/D/C metrics does not involve the class information, and thus may fail to consider samples with wrong classes. Such a phenomenon motivates us to further introduce a set of class-wise (CW) metrics, which takes the class information into account by evaluating the P/R/D/C metrics on a per-class basis. Specifically, the class-wise metrics are evaluated separately for each class i with m_i images. The evaluation results of these metrics shown in Fig. 4 reveal that our method outperforms the base method for a majority of classes. Moreover, the results of the average class-wise metrics presented in the last four rows of Table 2 also show that our method yields a better performance as compared to the base method.

Base on these evidences, it can be concluded that the proposed method outperforms both baseline methods in terms of FID and IS, implying that our method does possess a better ability to capture the true data distribution. Additionally, the evaluation results on the CAS and the class-wise metrics suggest that our method does offer a superior ability for a classifier to learn accurate class information as compared to the base method.

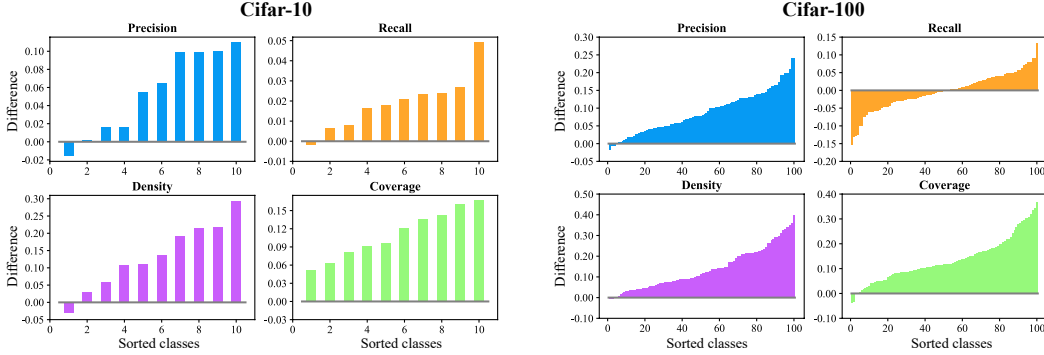


Figure 4: The sorted differences between the proposed method and the base method evaluated on the Cifar-10 and Cifar-100 datasets for the class-wise P / R / D / C metrics. Each colored bar in the plots represents the differences between our method and the base method evaluated using one of the P / R / D / C metrics for a certain class. A positive difference represents that our method outperforms the base method for that class.

5.3 ABLATION ANALYSIS

To further investigate the characteristic of L_{Total} , we perform an ablation analysis on the two components L_{CE} and $L_{DLSM'}$ in L_{Total} using the inter-twinning moon dataset, and observe the trends of (a) the score errors $\frac{1}{2} \mathbb{E}_{U(\tilde{x})} [D_L(\tilde{x}, c_1)] + \frac{1}{2} \mathbb{E}_{U(\tilde{x})} [D_L(\tilde{x}, c_2)]$, and (b) the cross entropy $\mathbb{E}_{p_{\sigma, \tau}(\tilde{x}, \tilde{y})} [-\log p(\tilde{y}|\tilde{x}; \theta)]$ during the training iterations. Metrics (a) and (b) measure the accuracies of the estimated likelihood scores $\nabla_{\tilde{x}} \log p(\tilde{y}|\tilde{x}; \theta)$ and the estimated likelihood density $p(\tilde{y}|\tilde{x}; \theta)$, respectively. As depicted in Fig 5, an increasing trend is observed for metric (a) when the classifier is trained with L_{CE} alone. On the contrary, the opposite tendency is observed for those trained with L_{Total} and $L_{DLSM'}$. The results thus suggest that matching the likelihood scores implicitly through training a classifier with L_{CE} alone leads to larger approximation errors. In contrast, $L_{DLSM'}$ and L_{Total} explicitly encourage the classifier to capture accurate likelihood scores, since they involve the score-matching objective. On the other hand, for metric (b), the classifiers trained with L_{Total} and L_{CE} yield stable decreasing trend in comparison to that trained with $L_{DLSM'}$ alone. These results suggest that in addition to minimizing $L_{DLSM'}$ between $\nabla_{\tilde{x}} \log p_{\sigma, \tau}(\tilde{y}|\tilde{x})$ and $\nabla_{\tilde{x}} \log p(\tilde{y}|\tilde{x}; \theta)$, the utilization of L_{CE} as an auxiliary objective enhances the stability during training. Based on the above observations, the interplay of L_{CE} and $L_{DLSM'}$ synergistically achieves the best results in terms of metric (a) while ensuring the stability of the training process. These clues thus validate the adoption of L_{Total} during the training of the classifier.

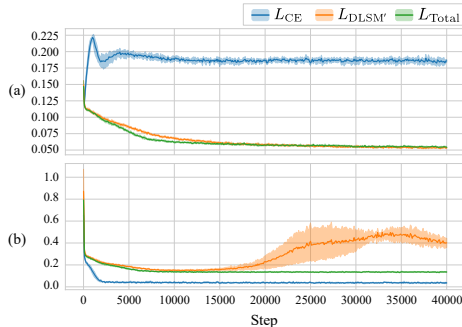


Figure 5: The evaluation curves of (a) the score errors and (b) the cross entropy during the training iterations for L_{CE} , $L_{DLSM'}$, and L_{Total} . The curves depict the mean and 95% confidence interval of five times of training.

6 CONCLUSION

In this paper, we highlighted the score mismatch issue in the existing conditional score-based data generation methods, and theoretically derived a new denoising likelihood score-matching (DLSM) loss, which is a training objective for the classifier to match the true likelihood score. We have demonstrated that, by adopting the proposed DLSM loss, the likelihood scores can be better estimated, and the negative impact of the score mismatch issue can be alleviated. Our experimental results validated that the proposed method does offer benefits in producing higher-quality conditional sampling results on both Cifar-10 and Cifar-100 datasets.

ACKNOWLEDGMENTS

This work was supported by the Ministry of Science and Technology (MOST) in Taiwan under grant number MOST 111-2628-E-007-010. The authors acknowledge the financial support from MediaTek Inc., Taiwan. The authors would also like to acknowledge the donation of the GPUs from NVIDIA Corporation and NVIDIA AI Technology Center (NVAITC) used in this research work. The authors thank National Center for High-Performance Computing (NCHC) for providing computational and storage resources. Finally, the authors would also like to thank the time and effort of the anonymous reviewers for reviewing this paper.

REFERENCES

- S. Barratt and R. Sharma. A note on the inception score. 2018.
- P. Dhariwal and A. Nichol. Diffusion models beat GANs on image synthesis. *arXiv preprint arXiv:2105.05233*, 2021.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. of the Conf. on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016.
- M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter. GANs trained by a two time-scale update rule converge to a local nash equilibrium. In *Proc. of Conf. on Neural Information Processing Systems (NeurIPS)*, volume 30, 2017.
- J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. In *Proc. of Conf. on Neural Information Processing Systems (NeurIPS)*, 2020.
- J. Ho, Chitwan S., W. Chan, D. Fleet, M. Norouzi, and T. Salimans. Cascaded diffusion models for high fidelity image generation. *ArXiv*, abs/2106.15282, 2021.
- A. Hyvärinen. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research (JMLR)*, 6(24):695–709, 2005. URL <http://jmlr.org/papers/v6/hyvarinen05a.html>.
- V. Jayaram and J. Thickstun. Source separation with deep generative priors. In *Proc. of the Int. Conf. on Machine Learning (ICML)*, 2020.
- A. Krizhevsky, G. Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- T. Kynkäänniemi, T. Karras, S. Laine, J. Lehtinen, and T. Aila. Improved precision and recall metric for assessing generative models. In *Proc. of Conf. on Neural Information Processing Systems (NeurIPS)*, 2019.
- M. F. Naeem, S. J. Oh, Y. Uh, Y. Choi, and J. Yoo. Reliable fidelity and diversity metrics for generative models. In *Proc. of the Int. Conf. on Machine Learning (ICML)*, 2020.
- A. Nguyen, J. Clune, Y. Bengio, A. Dosovitskiy, and J. Yosinski. Plug & play generative networks: Conditional iterative generation of images in latent space. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pp. 4467–4477, 2017.
- E. Parzen. On estimation of a probability density function and mode. In *The annals of mathematical statistics*, volume 33, pp. 1065–1076. JSTOR, 1962.
- S. Ravuri and O. Vinyals. Classification accuracy score for conditional generative models. In *Proc. of Conf. on Neural Information Processing Systems (NeurIPS)*, 2019.
- G. O Roberts and J. S Rosenthal. Optimal scaling of discrete approximations to langevin diffusions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 60(1):255–268, 1998.
- G. O Roberts and R. L. Tweedie. Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli*, 2(4):341 – 363, 1996. doi: [bj/1178291835](https://doi.org/10.1080/10236199608839533). URL [https://doi.org/](https://doi.org/10.1080/10236199608839533).

- J. Song, C. Meng, and S. Ermon. Denoising diffusion implicit models. In *Proc. of the Int. Conf. on Machine Learning (ICML)*, 2021a.
- Y. Song and S. Ermon. Generative modeling by estimating gradients of the data distribution. In *Proc. of Conf. on Neural Information Processing Systems (NeurIPS)*, 2019.
- Y. Song and S. Ermon. Improved techniques for training score-based generative models. In *Proc. of Conf. on Neural Information Processing Systems (NeurIPS)*, 2020.
- Y. Song, S. Garg, J. Shi, and S. Ermon. Sliced score matching: A scalable approach to density and score estimation. In *Proc. of the Conf. on Uncertainty in Artificial Intelligence (UAI)*, pp. 204, 2019. URL <http://auai.org/uai2019/proceedings/papers/204.pdf>.
- Y. Song, Jascha S.-D., Diederik P K., Abhishek K., Stefano E., and Ben P. Score-based generative modeling through stochastic differential equations. In *Int. Conf. on Learning Representations (ICLR)*, 2021b. URL <https://openreview.net/forum?id=PxtTIG12RRHS>.
- P. Vincent. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011.
- M. Welling and Y. W Teh. Bayesian learning via stochastic gradient langevin dynamics. In *Proc. of the Int. Conf. on Machine Learning (ICML)*, pp. 681–688. Citeseer, 2011.

A APPENDIX

In this Appendix, we first provide the definitions for the symbols used in the main manuscript and the Appendix in Section A.1. Next, we derive the theoretical proof for the closed form of $\nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}})$ in Section A.2. Then, we offer a number of discussions on the scaling technique in Section A.3, and different conditional score-based data generation methods in Section A.4. Subsequently, we offer the theoretical proof for Theorem 1 in Section A.5. Finally, in Sections A.6 and A.7, we provide the detailed experimental setups and additional experimental results.

A.1 THE LIST OF SYMBOLS USED IN THE PAPER

In this section, we provide the list of symbols used throughout the main manuscript and the Appendix. These symbols are summarized in Tables A1 and A2, containing their notations as well as their detailed descriptions.

Symbol	Description
$\langle \mathbf{u}, \mathbf{v} \rangle = \mathbf{u}^T \mathbf{v} = \sum_i \mathbf{u}_i \mathbf{v}_i$ $\ \mathbf{u}\ = \sqrt{\langle \mathbf{u}, \mathbf{u} \rangle}$ $\mathbf{I}_{d \times d}$	dot (inner) product between two vectors. Euclidean norm of vector. identity matrix with dimension $d \times d$.
ϵ $t \in \{0, \dots, T\}$ $\mathbf{x}, \mathbf{x}_t \in \mathbb{R}^d$ $\tilde{\mathbf{x}} \in \mathbb{R}^d$ $\mathbf{y} \in \mathbb{R}^c$ $\tilde{\mathbf{y}} \in \mathbb{R}^c$ $\mathbf{z} : \mathbb{R}^d$	the step size used in Langevin Diffusion. the discrete timestep used in Langevin Diffusion. data sample with dimension d , where \mathbf{x}_0 is sampled from $\mathcal{N}(\mathbf{0}, \mathbf{I}_{d \times d})$. perturbed data sample used in Denoising Score Matching. conditional class label with dimension c . perturbed class label used in Denoising Score Matching. the noise vector used in Langevin Diffusion. (sampled from $\mathcal{N}(\mathbf{0}, \mathbf{I}_{d \times d})$)
$\mathbf{x}^{(i)} \in \mathbb{R}^d$ $\mathbf{y}^{(i)} \in \{0, 1\}^c$ $\mathbb{D} = \{(\mathbf{x}^{(1)}, \mathbf{y}^{(1)}), \dots, (\mathbf{x}^{(m)}, \mathbf{y}^{(m)})\}$	the i -th example (input) from a dataset \mathbb{D} . the class label associated with $\mathbf{x}^{(i)}$ for supervised learning (one-hot encoded). Training set: m data-label pairs from p_{data} . (i.i.d. samples from p_{data})
$p_{\text{data}}(\mathbf{x})$ $p_0(\mathbf{x}) = \frac{1}{m} \sum_{i=1}^m \delta(\ \mathbf{x} - \mathbf{x}^{(i)}\)$ $p_{\sigma}(\tilde{\mathbf{x}} \mathbf{x}) = \frac{1}{(2\pi)^{d/2} \sigma^d} e^{-\frac{1}{2\sigma^2} \ \tilde{\mathbf{x}} - \mathbf{x}\ ^2}$ $p_{\sigma}(\tilde{\mathbf{x}}, \mathbf{x}) = p_{\sigma}(\tilde{\mathbf{x}} \mathbf{x})p_0(\mathbf{x})$ $p_{\sigma}(\tilde{\mathbf{x}}) = \frac{1}{m} \sum_{i=1}^m p_{\sigma}(\tilde{\mathbf{x}} \mathbf{x}^{(i)})$	unknown true probability density function (pdf) of data samples. empirical pdf of data samples associated with the training set \mathbb{D} . smoothing kernel or noise model: sampled from isotropic Gaussian $\mathcal{N}(\mathbf{x}, \sigma^2 \mathbf{I}_{d \times d})$ with mean \mathbf{x} and variance σ^2 . joint pdf of perturbed data and clean data samples. Parzen density estimate based on \mathbb{D} , obtainable by marginalizing $p_{\sigma}(\tilde{\mathbf{x}}, \mathbf{x})$.
$p_{\text{data}}(\mathbf{y})$ $p_0(\mathbf{y}) = \frac{1}{m} \sum_{i=1}^m \delta(\ \mathbf{y} - \mathbf{y}^{(i)}\)$ $p_{\tau}(\tilde{\mathbf{y}} \mathbf{y}) = \frac{1}{(2\pi)^{c/2} \tau^c} e^{-\frac{1}{2\tau^2} \ \tilde{\mathbf{y}} - \mathbf{y}\ ^2}$ $p_{\tau}(\tilde{\mathbf{y}}, \mathbf{y}) = p_{\tau}(\tilde{\mathbf{y}} \mathbf{y})p_0(\mathbf{y})$ $p_{\tau}(\tilde{\mathbf{y}}) = \frac{1}{m} \sum_{i=1}^m p_{\tau}(\tilde{\mathbf{y}} \mathbf{y}^{(i)})$	unknown true pdf of data labels. empirical pdf of data labels associated with the training set \mathbb{D} . smoothing kernel or noise model: sampled from isotropic Gaussian $\mathcal{N}(\mathbf{y}, \tau^2 \mathbf{I}_{c \times c})$ with mean \mathbf{y} and variance τ^2 . joint pdf of perturbed labels and clean labels. Parzen density estimate based on \mathbb{D} , obtainable by marginalizing $p_{\tau}(\tilde{\mathbf{y}}, \mathbf{y})$.
$p_{0,0}(\mathbf{x}, \mathbf{y})$ $p_{\sigma,0}(\tilde{\mathbf{x}}, \mathbf{x}, \mathbf{y})$ $p_{0,\tau}(\mathbf{x}, \tilde{\mathbf{y}}, \mathbf{y})$ $p_{\sigma,\tau}(\tilde{\mathbf{x}}, \mathbf{x}, \tilde{\mathbf{y}}, \mathbf{y})$ $s(\tilde{\mathbf{x}}; \phi)$ $p(\tilde{\mathbf{y}} \tilde{\mathbf{x}}; \theta)$	defined as $\frac{1}{m} \sum_{i=1}^m \delta(\ \mathbf{x} - \mathbf{x}^{(i)}\) \delta(\ \mathbf{y} - \mathbf{y}^{(i)}\)$. equals to $p_{\sigma}(\tilde{\mathbf{x}} \mathbf{x})p_{0,0}(\mathbf{x}, \mathbf{y})$. equals to $p_{\tau}(\tilde{\mathbf{y}} \mathbf{y})p_{0,0}(\mathbf{x}, \mathbf{y})$. defined as $p_{\sigma}(\tilde{\mathbf{x}} \mathbf{x})p_{\tau}(\tilde{\mathbf{y}} \mathbf{y})p_{0,0}(\mathbf{x}, \mathbf{y})$. Score Model: approximated score function with parameters ϕ . Classifier: approximated probability function parameterized by θ .

Table A1: The list of symbols used in this paper.

α	the scaling factor.
$\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{x}} \mathbf{y})$	posterior score.
$\nabla_{\tilde{\mathbf{x}}} \log p(\mathbf{y} \tilde{\mathbf{x}})$	likelihood score.
$\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{x}})$	prior score.
$L_{\text{ESM}}(\phi)$	Explicit Score Matching (ESM) loss for optimizing $s(\tilde{\mathbf{x}}; \phi)$.
$L_{\text{DSM}}(\phi)$	Denosing Score Matching loss (DSM) for optimizing $s(\tilde{\mathbf{x}}; \phi)$.
$L_{\text{ELSM}}(\theta)$	Explicit Likelihood Score Matching (ELSM) loss for optimizing $p(\mathbf{y} \tilde{\mathbf{x}}; \theta)$.
$L_{\text{DLSM}}(\theta)$	Denosing Likelihood Score Matching (DLSM) loss for optimizing $p(\mathbf{y} \tilde{\mathbf{x}}; \theta)$.
$L_{\text{DLSM}'}(\theta)$	approximated Denosing Likelihood Score Matching loss for optimizing $p(\mathbf{y} \tilde{\mathbf{x}}; \theta)$.
$L_{\text{CE}}(\theta)$	Cross Entropy (CE) loss for optimizing $p(\mathbf{y} \tilde{\mathbf{x}}; \theta)$.
$L_{\text{Total}}(\theta)$	full loss for optimizing $p(\mathbf{y} \tilde{\mathbf{x}}; \theta)$.

Table A2: The list of symbols used in this paper.

A.2 A PROOF OF THE CLOSED FORM FOR $\nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}})$

In Sections 2.2 and 3, we utilized the closed form formula to calculate the true scores $\nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}})$. To show the equivalence holds for Eq. (2), we provide a formal derivation as shown in the following proposition.

Proposition 1. *Given $p_{\sigma}(\tilde{\mathbf{x}}) = \frac{1}{m} \sum_{i=1}^m p_{\sigma}(\tilde{\mathbf{x}}|\mathbf{x}^{(i)})$, $p_{\sigma}(\tilde{\mathbf{x}}|\mathbf{x}) = \frac{1}{(2\pi)^{d/2}\sigma^d} e^{-\frac{1}{2\sigma^2}\|\tilde{\mathbf{x}}-\mathbf{x}\|^2}$, and $\sigma > 0$, the closed form of $\nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}})$ can be represented as follows:*

$$\nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}}) = \frac{\sum_{i=1}^m \frac{1}{\sigma^2} (\mathbf{x}^{(i)} - \tilde{\mathbf{x}}) p_{\sigma}(\tilde{\mathbf{x}}|\mathbf{x}^{(i)})}{\sum_{i=1}^m p_{\sigma}(\tilde{\mathbf{x}}|\mathbf{x}^{(i)})}.$$

Proof.

$$\begin{aligned} \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}}) &= \nabla_{\tilde{\mathbf{x}}} \log \frac{1}{m} \sum_{i=1}^m p_{\sigma}(\tilde{\mathbf{x}}|\mathbf{x}^{(i)}) \\ &= \nabla_{\tilde{\mathbf{x}}} \log \sum_{i=1}^m p_{\sigma}(\tilde{\mathbf{x}}|\mathbf{x}^{(i)}) + \underbrace{\nabla_{\tilde{\mathbf{x}}} \log \frac{1}{m}}_{=0} \\ &= \frac{\nabla_{\tilde{\mathbf{x}}} \sum_{i=1}^m p_{\sigma}(\tilde{\mathbf{x}}|\mathbf{x}^{(i)})}{\sum_{i=1}^m p_{\sigma}(\tilde{\mathbf{x}}|\mathbf{x}^{(i)})} \\ &= \frac{\sum_{i=1}^m \nabla_{\tilde{\mathbf{x}}} p_{\sigma}(\tilde{\mathbf{x}}|\mathbf{x}^{(i)})}{\sum_{i=1}^m p_{\sigma}(\tilde{\mathbf{x}}|\mathbf{x}^{(i)})} \\ &= \frac{\sum_{i=1}^m \nabla_{\tilde{\mathbf{x}}} \left[\frac{1}{(2\pi)^{d/2}\sigma^d} e^{-\frac{1}{2\sigma^2}\|\tilde{\mathbf{x}}-\mathbf{x}^{(i)}\|^2} \right]}{\sum_{i=1}^m p_{\sigma}(\tilde{\mathbf{x}}|\mathbf{x}^{(i)})} \\ &= \frac{\sum_{i=1}^m \left[\frac{(\mathbf{x}^{(i)} - \tilde{\mathbf{x}})/\sigma^2}{(2\pi)^{d/2}\sigma^d} e^{-\frac{1}{2\sigma^2}\|\tilde{\mathbf{x}}-\mathbf{x}^{(i)}\|^2} \right]}{\sum_{i=1}^m p_{\sigma}(\tilde{\mathbf{x}}|\mathbf{x}^{(i)})} \\ &= \frac{\sum_{i=1}^m \frac{1}{\sigma^2} (\mathbf{x}^{(i)} - \tilde{\mathbf{x}}) p_{\sigma}(\tilde{\mathbf{x}}|\mathbf{x}^{(i)})}{\sum_{i=1}^m p_{\sigma}(\tilde{\mathbf{x}}|\mathbf{x}^{(i)})} \end{aligned}$$

□

A.3 A DISCUSSION ON THE IMPACTS OF THE SCALING TECHNIQUE

In this section, we first provide a proof to demonstrate the reason behind the non-equivalence between the scaling technique and the re-normalization of a likelihood density. Then, we present the

derivation of the posterior scores based on a re-normalized likelihood density. Finally, we adopt a low-dimensional example to explain the effect of re-normalization on a posterior density.

The authors in (Dhariwal & Nichol, 2021) leveraged a scaling factor $\alpha > 0$ to adjust the likelihood score $\nabla_{\mathbf{x}} \log p(\mathbf{y}|\mathbf{x})$. To demonstrate the effect of such a scaling factor α , the authors claimed that there always exists a normalizing constant Z such that $\alpha \nabla_{\mathbf{x}} \log p(\mathbf{y}|\mathbf{x}) = \nabla_{\mathbf{x}} \log \frac{1}{Z} p(\mathbf{y}|\mathbf{x})^\alpha$, where $\frac{1}{Z} p(\mathbf{y}|\mathbf{x})^\alpha$ is a re-normalized likelihood density. However, such an equality only holds for specific conditions, and is not applicable to general cases, as explained in **Proposition 2**.

Proposition 2. *Given a constant $\alpha > 0$, and a probability density function (pdf) $p(\mathbf{y}|\mathbf{x})$ that is everywhere non-zero and differentiable, there does not exist a constant Z such that $\frac{1}{Z} p(\mathbf{y}|\mathbf{x})^\alpha$ is a pdf for all \mathbf{x} in general.*

Proof. The proposition can be proved by contradiction. Assume that such a constant Z exists and consider the following case:

$$\begin{aligned} \alpha &= 2, \\ x, y &\in \mathbb{R}, \\ f(y; \mu, \sigma^2) &= (2\pi\sigma^2)^{-\frac{1}{2}} e^{-\frac{(y-\mu)^2}{2\sigma^2}}, \\ p(y|x) &= \frac{1}{2(x^2+1)} \left((x+1)^2 f(y; 0, 1) + (x-1)^2 f(y; 0, \frac{1}{2}) \right). \end{aligned}$$

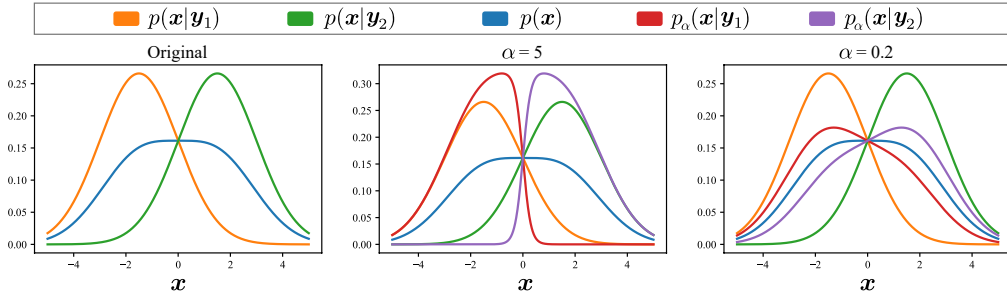
Note that $p(y|x)$ is a pdf and is everywhere non-zero and differentiable.

Since a pdf sums to 1 (i.e., $\int_{\mathbf{y}} \frac{1}{Z} p(\mathbf{y}|\mathbf{x})^\alpha d\mathbf{y} = 1, \forall \mathbf{x}$), by utilizing the Gaussian integral and the above assumption, the following derivation holds:

$$\begin{aligned} 1 &= \int_{\mathbf{y}} \frac{1}{Z} p(\mathbf{y}|\mathbf{x})^\alpha d\mathbf{y} = \int_{\mathbf{y}} \frac{1}{Z} p(\mathbf{y}|1)^\alpha d\mathbf{y} \\ \Leftrightarrow Z &= \int_{\mathbf{y}} p(\mathbf{y}|\mathbf{x})^2 d\mathbf{y} = \int_{\mathbf{y}} p(\mathbf{y}|1)^2 d\mathbf{y} \\ \Leftrightarrow Z &= \int_{\mathbf{y}} f(y; 0, \frac{1}{2})^2 d\mathbf{y} = \int_{\mathbf{y}} f(y; 0, 1)^2 d\mathbf{y} \\ \Leftrightarrow Z &= \int_{\mathbf{y}} \left((\pi)^{-\frac{1}{2}} e^{-y^2} \right)^2 d\mathbf{y} = \int_{\mathbf{y}} \left((2\pi)^{-\frac{1}{2}} e^{-\frac{y^2}{2}} \right)^2 d\mathbf{y} \\ \Leftrightarrow Z &= \frac{1}{\pi} \int_{\mathbf{y}} \left(e^{-y^2} \right)^2 d\mathbf{y} = \frac{1}{2\pi} \int_{\mathbf{y}} \left(e^{-\frac{y^2}{2}} \right)^2 d\mathbf{y} \\ \Leftrightarrow Z &= \frac{1}{\pi} \int_{\mathbf{y}} e^{-2y^2} d\mathbf{y} = \frac{1}{2\pi} \int_{\mathbf{y}} e^{-y^2} d\mathbf{y} \\ \Leftrightarrow Z &= \frac{1}{\sqrt{2\pi}} \int_{\mathbf{u}} e^{-u^2} d\mathbf{u} = \frac{1}{2\pi} \int_{\mathbf{y}} e^{-y^2} d\mathbf{y} \\ \Leftrightarrow Z &= \frac{1}{\sqrt{2\pi}} \sqrt{\pi} = \frac{1}{2\pi} \sqrt{\pi} \\ \Leftrightarrow Z &= \frac{1}{\sqrt{2\pi}} = \frac{1}{2\sqrt{\pi}} \end{aligned}$$

The final equation leads to a contradiction for the value of Z , which completes the proof of **Proposition 2**. Therefore, we concluded that there does not exist a constant Z such that $\frac{1}{Z} p(\mathbf{y}|\mathbf{x})^\alpha$ is a pdf for all \mathbf{x} in general. \square

To the best of our knowledge, the scaling technique still lacks a theoretical justification. However, if Z is allowed to be a function of \mathbf{x} rather than a constant, re-normalization of the likelihood density becomes feasible. In other words, setting $Z(\mathbf{x}) = \int_{\mathbf{y}} p(\mathbf{y}|\mathbf{x})^\alpha d\mathbf{y}$ enables $\int_{\mathbf{y}} \frac{1}{Z(\mathbf{x})} p(\mathbf{y}|\mathbf{x})^\alpha d\mathbf{y} = 1$,

Figure A1: The original and the renormalized densities under different choices of α .

where $\frac{1}{Z(\mathbf{x})}p(\mathbf{y}|\mathbf{x})^\alpha \triangleq p_\alpha(\mathbf{y}|\mathbf{x})$ is a re-normalized likelihood density. In this case, the score of the re-normalized likelihood density can be written as,

$$\begin{aligned}
 \nabla_{\mathbf{x}} \log p_\alpha(\mathbf{y}|\mathbf{x}) &= \nabla_{\mathbf{x}} \log \left(\frac{1}{Z(\mathbf{x})} p(\mathbf{y}|\mathbf{x})^\alpha \right) \\
 &= \nabla_{\mathbf{x}} \left(\log p(\mathbf{y}|\mathbf{x})^\alpha - \log Z(\mathbf{x}) \right) \\
 &= \nabla_{\mathbf{x}} \log p(\mathbf{y}|\mathbf{x})^\alpha - \nabla_{\mathbf{x}} \log Z(\mathbf{x}) \\
 &= \alpha \nabla_{\mathbf{x}} \log p(\mathbf{y}|\mathbf{x}) - \nabla_{\mathbf{x}} \log Z(\mathbf{x}).
 \end{aligned} \tag{A1}$$

Given the joint density $p_\alpha(\mathbf{x}, \mathbf{y}) \triangleq p_\alpha(\mathbf{y}|\mathbf{x})p(\mathbf{x})$, the posterior of the re-normalized density can be derived by Bayes' Theorem, i.e., $p_\alpha(\mathbf{x}|\mathbf{y}) = p_\alpha(\mathbf{y}|\mathbf{x})p(\mathbf{x}) / \int_{\mathbf{x}} p_\alpha(\mathbf{x}, \mathbf{y}) d\mathbf{x}$. Based on Eq. (5), the posterior score of the re-normalized density can be written as:

$$\begin{aligned}
 \nabla_{\mathbf{x}} \log p_\alpha(\mathbf{x}|\mathbf{y}) &= \nabla_{\mathbf{x}} \log p_\alpha(\mathbf{y}|\mathbf{x}) + \nabla_{\mathbf{x}} \log p(\mathbf{x}) \\
 &= \alpha \nabla_{\mathbf{x}} \log p(\mathbf{y}|\mathbf{x}) - \nabla_{\mathbf{x}} \log Z(\mathbf{x}) + \nabla_{\mathbf{x}} \log p(\mathbf{x}).
 \end{aligned} \tag{A2}$$

Through the above equation, conditional sampling based on the re-normalized likelihood density becomes feasible.

To further investigate the effect of re-normalization on $p_\alpha(\mathbf{x}|\mathbf{y})$, we provide an illustrative example in Fig. A1. In this example, the condition \mathbf{y} is assumed to be sampled from $\{\mathbf{y}_1, \mathbf{y}_2\}$ with equal probability 1/2. The original posterior densities $p(\mathbf{x}|\mathbf{y}_1)$ and $p(\mathbf{x}|\mathbf{y}_2)$ are normal distributions with variance 2.25 centered at $\mathbf{x} = -1.5$ and $\mathbf{x} = 1.5$, which are plotted in orange and green, respectively. On the other hand, the prior density $p(\mathbf{x}) \triangleq \frac{1}{2}p(\mathbf{x}|\mathbf{y}_1) + \frac{1}{2}p(\mathbf{x}|\mathbf{y}_2)$ is plotted in blue, and the likelihood density $p(\mathbf{y}|\mathbf{x}) = p(\mathbf{x}|\mathbf{y})p(\mathbf{y})/p(\mathbf{x})$ is defined according to Bayes' theorem.

Subsequently, the re-normalized posterior densities $p_\alpha(\mathbf{x}|\mathbf{y}_1)$ and $p_\alpha(\mathbf{x}|\mathbf{y}_2)$ can be constructed based on the equation $\frac{1}{Z(\mathbf{x})}p(\mathbf{y}_1|\mathbf{x})^\alpha$, $\frac{1}{Z(\mathbf{x})}p(\mathbf{y}_2|\mathbf{x})^\alpha$, and Bayes' theorem. In Fig. A1, it is observed that under the case that $\alpha = 5$, the re-normalized posterior densities are more concentrated. Conversely, when $\alpha = 0.2$, the re-normalized posterior densities are much more dispersed. These results therefore suggest that the re-normalized posterior density can be adjusted by controlling the value of α .

A.4 CONDITIONAL SCORE-BASED DATA GENERATION METHODS

As described in Section 3, methods (a), (b), and (d) utilize a classifier $p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta)$ and a score model $s(\tilde{\mathbf{x}}; \phi)$ to estimate the posterior scores. In contrast, method (c) separately trains the posterior score models $s(\tilde{\mathbf{x}}; \phi_1)$ and $s(\tilde{\mathbf{x}}; \phi_2)$ for different class conditions. Although method (c) achieves better evaluation results, it requires training score models for each possible $\tilde{\mathbf{y}}$, which is usually considered impractical in real-world applications. A potential solution to this problem is to train a score model $s(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}; \phi)$ that can estimate the scores for the joint probability of $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{y}}$ (denoted as 'Method (c*)' in Fig. A2). This method, however, is inflexible in comparison to methods (a), (b), and (d), since any change on the condition $\tilde{\mathbf{y}}$ could lead to re-training the joint score model $s(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}; \phi)$. Therefore, in this work, we only compare the performance of methods (a), (b), and (d), and consider method (c) to be another orthogonal research direction.

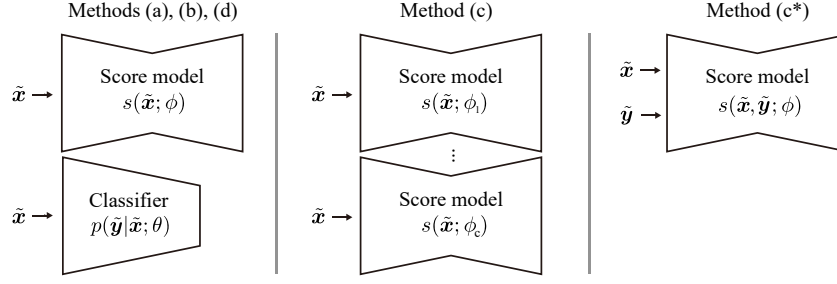


Figure A2: An illustration for different conditional score-based data generation methods in Section A.4.

A.5 A PROOF FOR DENOISING LIKELIHOOD SCORE MATCHING

In Section 4, we stated the equivalence between $L_{\text{ELSM}}(\theta)$ and $L_{\text{DLSTM}}(\theta)$ when optimizing θ . To demonstrate how such an equivalence stands, we provide the theoretical proof as follows.

Theorem 1. $L_{\text{DLSTM}}(\theta) = L_{\text{ELSM}}(\theta) + C$, where C is a constant with respect to θ .

Proof. $L_{\text{ELSM}}(\theta)$ and $L_{\text{DLSTM}}(\theta)$ are defined as follows:

$$L_{\text{ELSM}}(\theta) = \mathbb{E}_{p_{\sigma, \tau}(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})} \left[\frac{1}{2} \|\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta) - \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma, \tau}(\tilde{\mathbf{y}}|\tilde{\mathbf{x}})\|^2 \right].$$

$$L_{\text{DLSTM}}(\theta) = \mathbb{E}_{p_{\sigma, \tau}(\tilde{\mathbf{x}}, \mathbf{x}, \tilde{\mathbf{y}}, \mathbf{y})} \left[\frac{1}{2} \|\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta) + \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}}) - \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}}|\mathbf{x})\|^2 \right].$$

The proof expands $L_{\text{ELSM}}(\theta)$ and $L_{\text{DLSTM}}(\theta)$ individually, leading to constant difference.

$$\begin{aligned} L_{\text{ELSM}}(\theta) &= \mathbb{E}_{p_{\sigma, \tau}(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})} \left[\frac{1}{2} \|\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta) - \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma, \tau}(\tilde{\mathbf{y}}|\tilde{\mathbf{x}})\|^2 \right] \\ &= \mathbb{E}_{p_{\sigma, \tau}(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})} \left[\frac{1}{2} \|\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta)\|^2 \right] - F(\theta) + C_1, \text{ where} \\ C_1 &= \mathbb{E}_{p_{\sigma, \tau}(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})} \left[\frac{1}{2} \|\nabla_{\tilde{\mathbf{x}}} \log p_{\sigma, \tau}(\tilde{\mathbf{y}}|\tilde{\mathbf{x}})\|^2 \right] \text{ is a constant, and} \\ F(\theta) &= \mathbb{E}_{p_{\sigma, \tau}(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})} [\langle \nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta), \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma, \tau}(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}) \rangle] \\ &= \mathbb{E}_{p_{\sigma, \tau}(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})} [\langle \nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta), \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma, \tau}(\tilde{\mathbf{x}}|\tilde{\mathbf{y}}) - \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}}) \rangle] \\ &= G(\theta) - \mathbb{E}_{p_{\sigma, \tau}(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})} [\langle \nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta), \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}}) \rangle], \text{ where} \\ G(\theta) &= \mathbb{E}_{p_{\sigma, \tau}(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})} [\langle \nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta), \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma, \tau}(\tilde{\mathbf{x}}|\tilde{\mathbf{y}}) \rangle] \\ &= \int_{\tilde{\mathbf{x}}} \int_{\tilde{\mathbf{y}}} p_{\sigma, \tau}(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) \langle \nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta), \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma, \tau}(\tilde{\mathbf{x}}|\tilde{\mathbf{y}}) \rangle d\tilde{\mathbf{y}} d\tilde{\mathbf{x}} \\ &= \int_{\tilde{\mathbf{x}}} \int_{\tilde{\mathbf{y}}} p_{\tau}(\tilde{\mathbf{y}}) p_{\sigma, \tau}(\tilde{\mathbf{x}}|\tilde{\mathbf{y}}) \langle \nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta), \frac{\nabla_{\tilde{\mathbf{x}}} p_{\sigma, \tau}(\tilde{\mathbf{x}}|\tilde{\mathbf{y}})}{p_{\sigma, \tau}(\tilde{\mathbf{x}}|\tilde{\mathbf{y}})} \rangle d\tilde{\mathbf{y}} d\tilde{\mathbf{x}} \\ &= \int_{\tilde{\mathbf{x}}} \int_{\tilde{\mathbf{y}}} p_{\tau}(\tilde{\mathbf{y}}) \langle \nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta), \nabla_{\tilde{\mathbf{x}}} p_{\sigma, \tau}(\tilde{\mathbf{x}}|\tilde{\mathbf{y}}) \rangle d\tilde{\mathbf{y}} d\tilde{\mathbf{x}} \\ &= \int_{\tilde{\mathbf{x}}} \int_{\tilde{\mathbf{y}}} p_{\tau}(\tilde{\mathbf{y}}) \langle \nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta), \nabla_{\tilde{\mathbf{x}}} \int_{\mathbf{x}} p_{0, \tau}(\mathbf{x}|\tilde{\mathbf{y}}) p_{\sigma, \tau}(\tilde{\mathbf{x}}|\mathbf{x}, \tilde{\mathbf{y}}) d\mathbf{x} \rangle d\tilde{\mathbf{y}} d\tilde{\mathbf{x}} \\ &= \int_{\tilde{\mathbf{x}}} \int_{\tilde{\mathbf{y}}} p_{\tau}(\tilde{\mathbf{y}}) \langle \nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta), \nabla_{\tilde{\mathbf{x}}} \int_{\mathbf{x}} \int_{\mathbf{y}} p_{0, \tau}(\mathbf{x}|\tilde{\mathbf{y}}) p_{0, \tau}(\mathbf{y}|\mathbf{x}, \tilde{\mathbf{y}}) p_{\sigma, \tau}(\tilde{\mathbf{x}}|\mathbf{x}, \tilde{\mathbf{y}}, \mathbf{y}) d\mathbf{y} d\mathbf{x} \rangle d\tilde{\mathbf{y}} d\tilde{\mathbf{x}} \\ &= \int_{\tilde{\mathbf{x}}} \int_{\tilde{\mathbf{y}}} p_{\tau}(\tilde{\mathbf{y}}) \langle \nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta), \int_{\mathbf{x}} \int_{\mathbf{y}} p_{0, \tau}(\mathbf{x}|\tilde{\mathbf{y}}) p_{0, \tau}(\mathbf{y}|\mathbf{x}, \tilde{\mathbf{y}}) \nabla_{\tilde{\mathbf{x}}} p_{\sigma, \tau}(\tilde{\mathbf{x}}|\mathbf{x}, \tilde{\mathbf{y}}, \mathbf{y}) d\mathbf{y} d\mathbf{x} \rangle d\tilde{\mathbf{y}} d\tilde{\mathbf{x}} \end{aligned}$$

$$\begin{aligned}
&= \int_{\tilde{\mathbf{x}}} \int_{\tilde{\mathbf{y}}} \int_{\mathbf{x}} \int_{\mathbf{y}} p_{\sigma, \tau}(\tilde{\mathbf{y}}) p_{0, \tau}(\mathbf{x} | \tilde{\mathbf{y}}) p_{0, \tau}(\mathbf{y} | \mathbf{x}, \tilde{\mathbf{y}}) p_{\sigma, \tau}(\tilde{\mathbf{x}} | \mathbf{x}, \tilde{\mathbf{y}}, \mathbf{y}) \langle \nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}} | \tilde{\mathbf{x}}; \theta), \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma, \tau}(\tilde{\mathbf{x}} | \mathbf{x}, \tilde{\mathbf{y}}, \mathbf{y}) \rangle d\mathbf{y} d\mathbf{x} d\tilde{\mathbf{y}} d\tilde{\mathbf{x}} \\
&= \int_{\tilde{\mathbf{x}}} \int_{\tilde{\mathbf{y}}} \int_{\mathbf{x}} \int_{\mathbf{y}} p_{\sigma, \tau}(\tilde{\mathbf{x}}, \mathbf{x}, \tilde{\mathbf{y}}, \mathbf{y}) \langle \nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}} | \tilde{\mathbf{x}}; \theta), \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma, \tau}(\tilde{\mathbf{x}} | \mathbf{x}, \tilde{\mathbf{y}}, \mathbf{y}) \rangle d\mathbf{y} d\mathbf{x} d\tilde{\mathbf{y}} d\tilde{\mathbf{x}} \\
&= \int_{\tilde{\mathbf{x}}} \int_{\tilde{\mathbf{y}}} \int_{\mathbf{x}} \int_{\mathbf{y}} p_{\sigma, \tau}(\tilde{\mathbf{x}}, \mathbf{x}, \tilde{\mathbf{y}}, \mathbf{y}) \langle \nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}} | \tilde{\mathbf{x}}; \theta), \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma, \tau}(\tilde{\mathbf{x}}, \mathbf{x}, \tilde{\mathbf{y}}, \mathbf{y}) \rangle d\mathbf{y} d\mathbf{x} d\tilde{\mathbf{y}} d\tilde{\mathbf{x}} \\
&= \int_{\tilde{\mathbf{x}}} \int_{\tilde{\mathbf{y}}} \int_{\mathbf{x}} \int_{\mathbf{y}} p_{\sigma, \tau}(\tilde{\mathbf{x}}, \mathbf{x}, \tilde{\mathbf{y}}, \mathbf{y}) \langle \nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}} | \tilde{\mathbf{x}}; \theta), \nabla_{\tilde{\mathbf{x}}} \log(p_{\sigma}(\tilde{\mathbf{x}} | \mathbf{x}) p_{\tau}(\tilde{\mathbf{y}} | \mathbf{y}) p_{0,0}(\mathbf{x}, \mathbf{y})) \rangle d\mathbf{y} d\mathbf{x} d\tilde{\mathbf{y}} d\tilde{\mathbf{x}} \\
&= \int_{\tilde{\mathbf{x}}} \int_{\tilde{\mathbf{y}}} \int_{\mathbf{x}} \int_{\mathbf{y}} p_{\sigma, \tau}(\tilde{\mathbf{x}}, \mathbf{x}, \tilde{\mathbf{y}}, \mathbf{y}) \langle \nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}} | \tilde{\mathbf{x}}; \theta), \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}} | \mathbf{x}) \rangle d\mathbf{y} d\mathbf{x} d\tilde{\mathbf{y}} d\tilde{\mathbf{x}} \\
&= \mathbb{E}_{p_{\sigma, \tau}(\tilde{\mathbf{x}}, \mathbf{x}, \tilde{\mathbf{y}}, \mathbf{y})} [\langle \nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}} | \tilde{\mathbf{x}}; \theta), \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}} | \mathbf{x}) \rangle], \text{ and now we expand } L_{\text{DLISM}}(\theta) :
\end{aligned}$$

$$\begin{aligned}
L_{\text{DLISM}}(\theta) &= \mathbb{E}_{p_{\sigma, \tau}(\tilde{\mathbf{x}}, \mathbf{x}, \tilde{\mathbf{y}}, \mathbf{y})} \left[\frac{1}{2} \|\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}} | \tilde{\mathbf{x}}; \theta) + \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}}) - \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}} | \mathbf{x})\|^2 \right] \\
&= \mathbb{E}_{p_{\sigma, \tau}(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})} \left[\frac{1}{2} \|\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}} | \tilde{\mathbf{x}}; \theta)\|^2 \right] - \mathbb{E}_{p_{\sigma}(\tilde{\mathbf{x}}, \mathbf{x}, \tilde{\mathbf{y}}, \mathbf{y})} [\langle \nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}} | \tilde{\mathbf{x}}; \theta), \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}} | \mathbf{x}) \rangle] + \\
&\quad \mathbb{E}_{p_{\sigma, \tau}(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})} [\langle \nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}} | \tilde{\mathbf{x}}; \theta), \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}}) \rangle] + C_2, \text{ where} \\
C_2 &= \mathbb{E}_{p_{\sigma}(\tilde{\mathbf{x}}, \mathbf{x})} \left[\frac{1}{2} \|\nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}})\|^2 \right] + \mathbb{E}_{p_{\sigma}(\mathbf{x}, \tilde{\mathbf{x}})} \left[\frac{1}{2} \|\nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}} | \mathbf{x})\|^2 \right] - \\
&\quad \mathbb{E}_{p_{\sigma}(\tilde{\mathbf{x}}, \mathbf{x})} [\langle \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}}), \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}} | \mathbf{x}) \rangle], \text{ is a constant, leading to the conclusion} \\
L_{\text{DLISM}}(\theta) &= L_{\text{ELISM}}(\theta) - C_1 + C_2
\end{aligned}$$

Note that the domain of the integrals is $\{(\tilde{\mathbf{x}}, \mathbf{x}, \tilde{\mathbf{y}}, \mathbf{y}) : p_{\sigma, \tau}(\tilde{\mathbf{x}}, \mathbf{x}, \tilde{\mathbf{y}}, \mathbf{y}) \neq 0\}$ by the definition of expected value. \square

Remark. Theorem 1 also holds for the clean label \mathbf{y} , which can be shown as follows:

$$\begin{aligned}
&\mathbb{E}_{p_{\sigma}(\tilde{\mathbf{x}}, \mathbf{y})} \left[\frac{1}{2} \|\nabla_{\tilde{\mathbf{x}}} \log p(\mathbf{y} | \tilde{\mathbf{x}}; \theta) - \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\mathbf{y} | \tilde{\mathbf{x}})\|^2 \right] \\
&= \mathbb{E}_{p_{\sigma}(\tilde{\mathbf{x}}, \mathbf{x}, \mathbf{y})} \left[\frac{1}{2} \|\nabla_{\tilde{\mathbf{x}}} \log p(\mathbf{y} | \tilde{\mathbf{x}}; \theta) + \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}}) - \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}} | \mathbf{x})\|^2 \right] + C.
\end{aligned}$$

A.6 THE DETAILED EXPERIMENTAL CONFIGURATIONS

In this section, we elaborate on the implementation details and the experimental setups. In Section A.6.1, we describe the noise perturbation method and the sampling algorithm adopted in this work. In Sections A.6.2 and A.6.3, we present the training configurations for the experiments on the inter-twinning moon, Cifar-10, and Cifar-100 datasets. The code implementation for these experiments is provided in the Github repository (<https://github.com/chen-hao-chao/dlsm>).

A.6.1 NOISE PERTURBATION AND SAMPLING ALGORITHM

Sampling Algorithm. In this work, we adopt the PC sampler (VE-SDE) identical to that in (Song et al., 2021b), which is a time-inhomogeneous sampling algorithm. The PC sampler (VE-SDE) is described in **Algorithm 1**. In **Algorithm 1**, a noise-conditioned score model $s(\tilde{\mathbf{x}}; \phi, \sigma_t)$ is adopted, where $\tilde{\mathbf{x}}$ is perturbed by noises sampled from Gaussian distributions using different standard deviations σ_t . The standard deviations of the noises, denoted as $\sigma_t = \sigma_{\min} \left(\frac{\sigma_{\max}}{\sigma_{\min}} \right)^{\frac{t}{T}}$, are determined by the time step t in **Algorithm 1**. We set $\sigma_{\min} = 0.01$ for all of our experiments. On the other hand, we set $\sigma_{\max} = 50$ for the Cifar-10 and Cifar-100 datasets, and $\sigma_{\max} = 10$ for the inter-twinning moon dataset. As for the total time steps, the value of T is set to 1,000.

Algorithm 1 The PC Sampler (VE-SDE)

```

1:  $\tilde{\mathbf{x}} = \mathcal{N}(\mathbf{0}, \sigma_{\max}^2 \mathbf{I})$ 
2: for  $t = T - 1$  to 0 do
3:   /* Predictor */
4:    $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
5:    $\tilde{\mathbf{x}}_t \leftarrow \tilde{\mathbf{x}}_{t+1} + (\sigma_{t+1}^2 - \sigma_t^2) s(\tilde{\mathbf{x}}_{t+1}; \phi, \sigma_{t+1}) + \sqrt{\sigma_{t+1}^2 - \sigma_t^2} \mathbf{z}$ 
6:   /* Corrector */
7:    $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ 
8:    $\tilde{\mathbf{x}}_t \leftarrow \tilde{\mathbf{x}}_t + \epsilon_t s(\tilde{\mathbf{x}}_t; \phi, \sigma_t) - \sqrt{2\epsilon_t} \mathbf{z}$ 
9: return  $\tilde{\mathbf{x}}_0$ 

```

Training with Different Standard Deviations. Since the PC sampler (VE-SDE) described in **Algorithm 1** requires the score model $s(\tilde{\mathbf{x}}; \phi, \sigma_t)$ conditioned on different values of standard deviation σ_t , the training objectives described in Eqs. (4) and (11) are modified accordingly. For the DSM loss described in Eq. (4), the modified loss $L_{\text{DSM}}(\phi)$ is represented as:

$$\mathbb{E}_{\mathcal{U}_{[0, T]}(t)} \left[\lambda_{\text{DSM}}(t) \mathbb{E}_{p_{\sigma_t}(\tilde{\mathbf{x}}, \mathbf{x})} \left[\|s(\tilde{\mathbf{x}}; \phi, \sigma_t) - \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma_t}(\tilde{\mathbf{x}}|\mathbf{x})\|^2 \right] \right], \quad (\text{A3})$$

where $\mathcal{U}_{[0, T]}(t)$ is a uniform distribution on the interval $[0, T]$. In Eq. (A3), The additional coefficient $\lambda_{\text{DSM}}(t)$ is multiplied with the original DSM loss for better training results (Song & Ermon, 2019). For the approximated DLSM loss described in Eq. (9), the modified loss $L_{\text{DLSM}'}(\theta)$ is represented as:

$$\mathbb{E}_{\mathcal{U}_{[0, T]}(t)} \left[\lambda_{\text{DLSM}'}(t) \mathbb{E}_{p_{\sigma_t, \tau}(\tilde{\mathbf{x}}, \mathbf{x}, \tilde{\mathbf{y}}, \mathbf{y})} \left[\|\nabla_{\tilde{\mathbf{x}}} \log p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta, \sigma_t) + s(\tilde{\mathbf{x}}; \phi, \sigma_t) - \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma_t}(\tilde{\mathbf{x}}|\mathbf{x})\|^2 \right] \right], \quad (\text{A4})$$

where $p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta, \sigma_t)$ is a noise-conditioned classifier. $\lambda_{\text{DLSM}'}(t)$ is the balancing coefficient for different σ_t . For the cross-entropy loss, the modified loss $L_{\text{CE}}(\theta)$ is represented as:

$$\mathbb{E}_{\mathcal{U}_{[0, T]}(t)} \left[\lambda_{\text{CE}}(t) \mathbb{E}_{p_{\sigma_t, \tau}(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})} \left[-\log p(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}; \theta, \sigma_t) \right] \right], \quad (\text{A5})$$

where $\lambda_{\text{CE}}(t)$ is the balancing coefficient for different σ_t . In all of our experiments, $\lambda_{\text{CE}}(t)$ is set to 1. The values of $\lambda_{\text{DSM}}(t)$ and $\lambda_{\text{DLSM}'}(t)$ are customized according to different settings, and specified in Sections A.6.2 and A.6.3. Please note that $\lambda_{\text{DSM}}(t)$, $\lambda_{\text{DLSM}'}(t)$, and $\lambda_{\text{CE}}(t)$ are adopted to balance the loss functions for different σ_t , and are different from the balancing coefficient λ in Eq. (11).

A.6.2 THE EXPERIMENTS ON THE INTER-TWINNING MOON DATASET

Dataset. We perform our motivational experiments on the inter-twinning moon dataset provided by the `sklearn` library in `python`. In this dataset, the data points are sampled from two interleaving crescents. The data points of the upper crescent are generated based on $(\cos(x), \sin(x)), \forall x \in [0, \pi]$ and are labeled as c_1 . On the other hand, the data points of the lower crescent are generated based on $(1 - \cos(x), 0.5 - \sin(x)), \forall x \in [0, \pi]$ and are labeled as c_2 . Before training, the data points are first normalized such that their mean become 0 and then their coordinates are scaled by a factor of 20.

Setups and Evaluation Method. In this experiment, the scores $\nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}})$ and $\nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}}|\tilde{\mathbf{y}})$ are obtained by calculating Eq. (2). The likelihood scores are computed according to $\nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}) = \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}}|\tilde{\mathbf{y}}) - \nabla_{\tilde{\mathbf{x}}} \log p_{\sigma}(\tilde{\mathbf{x}})$ in Eq. (5). As for the D_F and D_L metrics in Table 1, the errors between two score functions are measured using 1,225 data points uniformly sampled from the two-dimensional subspace $[-25, 25] \times [-40, 40]$. The scaling factor α adopted in method (b) is set to 10. The values of $\lambda_{\text{DSM}}(t)$ and $\lambda_{\text{DLSM}'}(t)$ are both set to $\frac{1}{\sigma_t^4}$, and the value of λ is set to 0.125.

Training and Implementation Details. The network architectures for the score model and the classifier are simple three-layer multilayer perceptrons (MLPs) with ReLu as the activation function. These networks are implemented using the `pytorch` library with `nn.Linear` layers that contain (128, 64, 32) neurons per layer for both of the score model and the classifier. The score model is trained using the Adam optimizer with a learning rate of 6.5×10^{-4} and a batch size of 4,000 for 40,000 iterations. The classifier is trained using the Adam optimizer with a learning rate of 2.0×10^{-5} and a batch size of 4,000 for 40,000 iterations.

Table A3: The experimental results on the inter-twinning moon dataset, where the classifiers are trained with L_{CE} , $L_{DLSM'}$, and L_{Total} , respectively. The errors of the score functions for different methods are measured using $\mathbb{E}_{U(\tilde{\mathbf{x}})}[D_P(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})]$ and $\mathbb{E}_{U(\tilde{\mathbf{x}})}[D_L(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})]$, where $U(\tilde{\mathbf{x}})$ is a uniform distribution over a two-dimensional subspace, $\tilde{\mathbf{y}}$ represents the classes specified in the second row of the table, and D_P and D_L are defined in Eq. (7). The arrow symbol \downarrow indicates that lower values correspond to better performances.

		L_{CE}		$L_{DLSM'}$		L_{Total}	
		c_1	c_2	c_1	c_2	c_1	c_2
$\tilde{\mathbf{y}}$							
$\mathbb{E}_{U(\tilde{\mathbf{x}})}[D_P(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})]$	\downarrow	0.198	0.197	0.066	0.065	0.066	0.064
$\mathbb{E}_{U(\tilde{\mathbf{x}})}[D_L(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})]$	\downarrow	0.190	0.181	0.053	0.053	0.052	0.052

A.6.3 THE EXPERIMENTS ON THE CIFAR-10 AND CIFAR-100 DATASETS

Dataset. We adopt the Cifar-10 and Cifar-100 datasets (Krizhevsky et al., 2009) for evaluating our design on image generation tasks. Cifar-10 consists of 60,000 32×32 RGB images with 10 classes, where each class contains 6,000 images. Cifar-100 is similar to Cifar-10, except it has 100 classes, where each class contains 600 images. The images are first normalized to $[-1, 1]^d$ before training.

Setups and Evaluation Method. The results on the CAS metric are evaluated using a classifier (ResNet-50 (He et al., 2016)) trained with the generated samples. This classifier is trained to minimize the cross-entropy loss with the Adam optimizer using a learning rate of 2.0×10^{-4} with a batch size of 150. The results on the P / R / D / C metrics are calculated using the officially released `code`. The hyperparameter k used in the P / R / D / C metrics is fixed to 3. The results on the FID and IS metrics are computed on 50,000 samples using the `tensorflow_gan` library in python. The value of $\lambda_{DSM}(t)$ is set to $\frac{1}{\sigma_t^2}$, while the value of $\lambda_{DLSM}(t)$ is set to 1. The value of λ is set to 1.

Training and Implementation Details. The network architecture for the score model is the same as (Song et al., 2021b). The network architecture for the classifier is a modified ResNet-18 (He et al., 2016) for Cifar-10 and a modified ResNet-34 for Cifar-100. The modification only involves an additional conditional branch for encoding the standard deviation σ_t . The score model is trained with the Adam optimizer using a learning rate of 2.0×10^{-4} and a batch size of 128. The classifier is trained with the Adam optimizer using a learning rate of 2.0×10^{-4} and a batch size of 100.

A.7 ADDITIONAL EXPERIMENTS

In this section, we present a number of additional experimental results and provide discussions on them. We first report the results of our method using different values of λ on Cifar-10 to demonstrate that L_{Total} is less sensitive to the choices of λ . Next, we apply the scaling technique on both the base method and our method with different choices of α to demonstrate the influence of the hyperparameter α on the evaluation results of several metrics. Finally, we provide uncurated examples on the Cifar-10 and Cifar-100 datasets.

A.7.1 ABLATION ANALYSIS ON THE INTER-TWINNING MOON DATASET

Table A3 compares the evaluation results on the $\mathbb{E}_{U(\tilde{\mathbf{x}})}[D_P(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})]$ and $\mathbb{E}_{U(\tilde{\mathbf{x}})}[D_L(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})]$ metrics when the classifiers are trained with L_{CE} , $L_{DLSM'}$, and L_{Total} , respectively. It is observed that the score errors $\mathbb{E}_{U(\tilde{\mathbf{x}})}[D_P(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})]$ and $\mathbb{E}_{U(\tilde{\mathbf{x}})}[D_L(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})]$ are the lowest when L_{Total} is adopted. The experimental results therefore validate the effectiveness of the proposed loss L_{Total} .

A.7.2 THE EXPERIMENTS ON THE BALANCING COEFFICIENT

Table A4 demonstrates the evaluation results on FID, IS, and P / R / D / C metrics using L_{Total} with different values of λ . It is observed that the evaluation results only slightly fluctuate on the six metrics with different choices of λ . The results thus experimentally validate that L_{Total} is less sensitive to the choice of λ .

Table A4: The experimental results of our method using L_{Total} with different value of λ . The arrow symbols \uparrow / \downarrow represent that a higher / lower evaluation result correspond to a better performance.

		$\lambda = 2.0$	$\lambda = 1.0$	$\lambda = 0.5$	$\lambda = 0.125$
FID	\downarrow	2.33	2.25	2.36	2.25
IS	\uparrow	9.90	9.90	9.98	9.79
Precision	\uparrow	0.64	0.65	0.64	0.64
Recall	\uparrow	0.62	0.62	0.62	0.62
Density	\uparrow	0.94	0.96	0.95	0.94
Coverage	\uparrow	0.79	0.81	0.79	0.78

Table A5: A comparison of the evaluation results of the base method, the scaling method with different α , ours method, and ours + scaling method with different α on the Cifar-10 dataset. The arrow symbols \uparrow / \downarrow represent that a higher / lower evaluation result correspond to a better performance.

		Base	Scaling					Ours	Ours + Scaling				
α	-	1.0	0.5	3.0	5.0	10.0	20.0	1.0	0.5	3.0	5.0	10.0	20.0
FID	\downarrow	4.10	3.22	6.16	8.06	12.48	16.59	2.25	2.57	3.96	6.40	11.06	14.92
IS	\uparrow	9.08	9.50	9.14	9.38	9.37	9.23	9.90	9.79	10.00	9.95	9.73	9.53
Precision	\uparrow	0.67	0.65	0.70	0.72	0.75	0.75	0.65	0.62	0.69	0.71	0.73	0.74
Recall	\uparrow	0.61	0.61	0.56	0.53	0.49	0.44	0.62	0.63	0.57	0.53	0.49	0.44
Density	\uparrow	1.05	0.98	1.17	1.25	1.36	1.36	0.96	0.87	1.15	1.23	1.28	1.30
Coverage	\uparrow	0.80	0.78	0.78	0.77	0.75	0.67	0.81	0.77	0.79	0.77	0.71	0.66
(CW) Precision	\uparrow	0.51	0.43	0.63	0.67	0.70	0.73	0.56	0.48	0.65	0.68	0.71	0.72
(CW) Recall	\uparrow	0.59	0.62	0.51	0.46	0.42	0.37	0.61	0.63	0.55	0.50	0.45	0.39
(CW) Density	\uparrow	0.63	0.46	0.97	1.10	1.23	1.30	0.76	0.58	1.05	1.16	1.23	1.26
(CW) Coverage	\uparrow	0.60	0.47	0.69	0.69	0.66	0.61	0.71	0.61	0.75	0.73	0.66	0.61

A.7.3 THE EXPERIMENTS ON THE SCALING TECHNIQUE

Table A5 demonstrates the evaluation results of the base method, the scaling method using different values of α , our method, and our method with the scaling technique using different values of α (i.e., ‘ours + scaling’) on the Cifar-10 dataset. It is observed that the performance of the ‘scaling method’ and ‘ours + scaling’ method on the fidelity metrics, i.e., precision and density, improves as the value of α grows. Conversely, the performance of the ‘scaling method’ and ‘ours + scaling’ method on the FID, recall, and coverage metrics degrades as the value of α grows. This results thus experimentally demonstrate the effects of the scaling technique on the performance of the base method and our method.

A.7.4 UNCURATED EXAMPLES

Figs. A3, A4, A5, and A6 depict a few uncurated examples that qualitatively compare the effectiveness of different methods.



Figure A3: The uncurated examples generated using the base method on the Cifar-10 dataset.



Figure A4: The uncurated examples generated using the scaling method on the Cifar-10 dataset.



Figure A5: The uncurated examples generated using our method on the Cifar-10 dataset.




Figure A6: The uncurated examples of classes ‘baby’, ‘apple’, ‘aquarium fish’, and ‘beargenerated’ using the base method, the scaling method, and our method on the Cifar-100 dataset.