
Distillation of Discrete Diffusion through Dimensional Correlations

Satoshi Hayakawa¹ Yuhta Takida² Masaaki Imaizumi³ Hiromi Wakaki¹ Yuki Mitsufuji^{1,2}

Abstract

Diffusion models have demonstrated exceptional performances in various fields of generative modeling, but suffer from slow sampling speed due to their iterative nature. While this issue is being addressed in continuous domains, discrete diffusion models face unique challenges, particularly in capturing dependencies between elements (e.g., pixel relationships in image, sequential dependencies in language) mainly due to the computational cost of processing high-dimensional joint distributions. In this paper, (i) we propose “mixture” models for discrete diffusion that are capable of treating dimensional correlations while remaining scalable, and (ii) we provide a set of loss functions for distilling the iterations of existing models. Two primary theoretical insights underpin our approach: First, conventional models with element-wise independence can well approximate the data distribution, but essentially require *many sampling steps*. Second, our loss functions enable the mixture models to distill such many-step conventional models into just a few steps by learning the dimensional correlations. Our experimental results show the effectiveness of the proposed method in distilling pretrained discrete diffusion models across image and language domains. The code used in the paper is available at <https://github.com/sony/di4c>.

1. Introduction

Diffusion models (Sohl-Dickstein et al., 2015; Ho et al., 2020; Song et al., 2021b) have demonstrated excellent performance in generative modeling, particularly for continuous data such as images (Nichol et al., 2021; Rombach et al., 2022; Saharia et al., 2022), audio (Kong et al., 2021; Chen et al., 2021; Evans et al., 2024), and video (Harvey et al.,

2022; Ho et al., 2022; Blattmann et al., 2023). Recent advancements in diffusion models often outperform traditional generative models, such as variational autoencoders (VAEs, Kingma & Welling, 2014; Higgins et al., 2017; Zhao et al., 2019) and generative adversarial networks (GANs, Goodfellow et al., 2014), in terms of sample quality and the controllability of the generated results. Furthermore, diffusion models are not limited to learning continuous data; they can also be applied to discrete or categorical data with modifications (Hoogeboom et al., 2021; Austin et al., 2021) and offer a promising approach for discrete generative modeling (Gu et al., 2022; Lou et al., 2024). Such discrete diffusion models are the main topic of this paper.

A notable drawback of diffusion models is their slow sampling speed due to requiring many sampling steps (Xiao et al., 2022; Zhang & Chen, 2023). In continuous domains, various approaches have been proposed to reduce the number of steps, including well-designed forward processes (Song et al., 2021a) and fast solvers of stochastic/ordinary differential equations (SDEs/ODEs, Lu et al., 2022a,b; Zheng et al., 2023b). Another notable approach is knowledge distillation, which significantly reduces the number of sampling steps compared with earlier attempts by compressing pretrained diffusion models into single- or few-step generative models (Luhman & Luhman, 2021; Salimans & Ho, 2022; Meng et al., 2023; Zheng et al., 2023a). An emerging sub-family of distillation is the consistency-type models (Song et al., 2023; Song & Dhariwal, 2023; Kim et al., 2024), which exploit the fact that samples generated via different paths from the same initial noise should coincide.

However, discrete diffusion models face a fundamental challenge when attempting to reduce the number of sampling steps. Conventional approaches use “product” models that treat each dimension independently as sampling distributions (Figure 1, center), since high-dimensional joint distributions are intractable. While this has been successful with hundreds of sampling steps, ignoring element-wise dependencies (which we refer to as *dimensional correlations*) causes non-negligible approximation errors (Figure 1). This ignoring is also pointed out in some concurrent works (Park et al., 2025; Liu et al., 2024; Xu et al., 2025).

In this paper, we propose **Di4C** (Distilling Discrete Diffu-

¹Sony Group Corporation, Tokyo, Japan ²Sony AI, Tokyo, Japan ³The University of Tokyo, Tokyo, Japan. Correspondence to: Satoshi Hayakawa <satoshi.a.hayakawa@sony.com>.

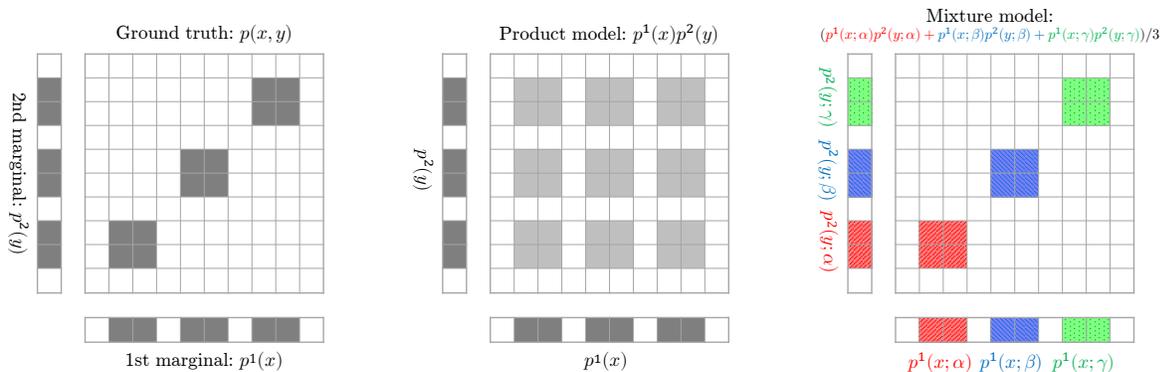


Figure 1. Illustration of dimensional correlations. (Left) Distribution $p(x, y)$ is two-dimensional categorical distribution. $p^1(x)$ and $p^2(y)$ are its marginals. (Center) Conventional denoiser in discrete diffusion uses *product model*, which is simply product of marginal distributions. It fails to approximate ground truth distribution. (Right) Our *mixture model* is given by expectation of product model $p(x, y; \lambda) = p^1(x; \lambda)p^2(y; \lambda)$ for random λ . In figure, λ takes α, β, γ in equal probabilities, and model reconstructs $p(x, y)$.

sion through Dimensional Correlations) to overcome this limitation. Our key insight is that while individual steps in conventional models are dimensionally independent, their composition over multiple steps can implicitly capture correlations (Section 3.1). On the basis of this observation, we develop (1) a “mixture” model that explicitly represents dimensional correlations while remaining computationally tractable (Figure 1, right), and (2) novel loss functions that effectively distill the many-step denoising of a product model into fewer steps. Our contribution is as follows:

Theoretical analysis: In Theorem 1, we show that N -step sampling with product models can approximate data distributions in $\mathcal{O}(1/N)$ total variation error. We also prove that this bound cannot be improved in a simple two-dimensional example. It underpins the empirical effectiveness of discrete diffusion models *with many steps* and, at the same time, shows the importance of modeling dimensional correlations to reduce the number of sampling steps.

Model and loss design: To capture the aforementioned dimensional correlations, we propose a “mixture” model that can represent dimensional correlations (Section 3.2). To distill a many-step discrete diffusion model (teacher) into a few-step model (student), we also propose Di4C loss functions for compressing the iterative process of the teacher (Section 3.3). In theory, we prove that the loss functions in Di4C can upper-bound the distance between the output distributions of the N -step teacher and the student with just one step (Theorem 2). In combination with Theorem 1, this provides an overall theoretical guarantee for Di4C.

Experiments: Finally, we demonstrate that our approach is general and applicable to multiple settings. (1) On CIFAR-10 with a pixel-based discretized Gaussian diffusion, we substantially improve the sample quality metrics of the teacher model (Campbell et al., 2022) in few-step sampling. (2) On

ImageNet class-conditional generation with masked diffusion, our method achieves a 2x speed-up while maintaining a comparable sample quality to the teacher model (Besnier & Chen, 2023). In addition, (3) on masked diffusion language modeling with OpenWebText, we show that Di4C can further distill a well-distilled model (Deschenaux & Gulcehre, 2025) by capturing dimensional correlations, without much harming of sampling diversity. These results consistently demonstrate that Di4C can effectively compress the sampling steps of discrete diffusion models while maintaining or improving generation quality.

Finally, the remainder of this paper is organized as follows: Section 2 gives preliminaries on discrete diffusion models and explains the dimensionality issue in discrete diffusion. We then explain the central idea of Di4C in Section 3 and show theoretical results in Section 4, which are partially described above as our contribution. In Section 5, we also provide experimental results with image and language tasks. After discussing related works in Section 6, we conclude the paper with discussions on its limitations and future work in Section 7.

2. Preliminaries

2.1. Discrete diffusion models

Suppose we have a data distribution $q_0 := q_{\text{data}}$ over the space \mathcal{X} . In diffusion models (Sohl-Dickstein et al., 2015; Ho et al., 2020), we consider a Markov process $(\mathbf{x}_t)_{0 \leq t \leq T}$ with $\mathbf{x}_0 \sim q_0$ and $\mathbf{x}_T \sim q_T$, where the time t can be either discrete or continuous. In this paper, we follow the notational convention that $q_{t|s}$ and $q_{s,t}$ represent the true conditional and joint distributions defined by this Markov process, respectively. This process is designed so that the terminal distribution q_T is a tractable distribution. Following the convention, we regard this forward process as adding

noise to the data distribution. Our aim is to generate samples approximately from the conditional distribution $q_{0|T}(\cdot|\mathbf{x}_T)$ with $\mathbf{x}_T \sim q_T$, which is a generative model for q_{data} . To this end, we introduce a *model* or *denoiser*, which is represented as $p_{s|t}$ (for $s < t$), to approximate $q_{s|t}$.

Our primary interest is in the discrete diffusion models (Austin et al., 2021; Campbell et al., 2022), where the space \mathcal{X} is a finite set. In this case, a probability distribution p on \mathcal{X} can be regarded as a function $p : \mathcal{X} \rightarrow \mathbb{R}$, and we will sometimes abuse the notation by treating p as just an ordinary function. We are given a finite set \mathcal{S} and consider a diffusion process over the product space $\mathcal{X} = \mathcal{S}^D$ for a large D . Each state $\mathbf{x} \in \mathcal{X}$ can thus be written as $\mathbf{x} = (x^d)_{d=1}^D$, where x^d indicates the entry of \mathbf{x} in the d -th dimension. Given a probability distribution $p = p(\mathbf{x})$ on \mathcal{X} , let $p^d = p^d(x^d)$ be its d -th marginal distribution, i.e., the distribution of x^d given $\mathbf{x} \sim p$. To enjoy scalability, the forward process is usually set to be factorized over dimensions, i.e., $q_{t|s}(\mathbf{x}_t|\mathbf{x}_s) = \prod_{d=1}^D q_{t|s}^d(x_t^d|x_s^d)$ holds for $s < t$ (Gu et al., 2022; Campbell et al., 2022).

2.2. Dimensional correlations in discrete diffusion

The common practices in modeling and training discrete diffusion models lead them to ignore the dimensional correlations within a data distribution. First, under the aforementioned problem setting, for the sake of scalability, the denoiser is usually defined as a *product model* that satisfies

$$p_{s|t}(\mathbf{x}_s|\mathbf{x}_t) = \prod_{d=1}^D p_{s|t}^d(x_s^d|x_t^d), \quad s < t. \quad (1)$$

Namely, the distribution $p_{s|t}(\cdot|\mathbf{x}_t)$ is dimensionally independent. This product modeling is common if not particularly highlighted (Campbell et al., 2022, Section G), due to the combinatorial explosion of the product discrete state. Indeed, adopting a product model significantly reduces the output length from $\mathcal{O}(D^{|\mathcal{S}|})$ to $\mathcal{O}(D|\mathcal{S}|)$ at the cost of representational capacity. This limited expressive power can be crucial for considering few-step discrete diffusion models. As an extreme example, consider doing one-step denoising in the case of masked (absorbing-state) diffusion (Austin et al., 2021); there is no chance we can approximate a complex joint distribution (as in Figure 1) in one step when \mathbf{x}_T is a completely masked sentence (i.e., following a delta distribution) and $p_{0|T}(\cdot|\mathbf{x}_T)$ is dimensionally independent. See Section F.1.1 for more examples.

Another potential factor making the learning of dimensional correlations infeasible in discrete diffusion models is that some of the existing loss functions are not well prepared for learning dimensional correlations. Most notably, in the continuous-time score-based discrete diffusion, we need only the marginal $p_{s|t}^d(\cdot|x_t^d)$ or its variant to compute the infinitesimal transition rate (see, e.g., Campbell et al. (2022,

Proposition 3) or Sun et al. (2023, Eq. 16)). Thus, learning backward transition rates does not lead to capturing dimensional correlations, even if we use a model capable of representing them.

3. Di4C for distilling discrete diffusion models

This section describes our proposed method, Di4C. We first show that the *composition* of well-trained discrete diffusion models can represent the dimensional correlation in Section 3.1, and in the later sections we discuss how to distill the multi-step denoising of a teacher model into a student model that requires fewer steps. In particular, Section 3.2 and Section 3.3 try to solve the existing limitation described in Section 2.2 in terms of modeling and loss design, respectively. See Section A for more technical details of Di4C.

3.1. Composition of diffusion denoisers for inducing dimensional correlations

We introduce the notion of composition, which plays a significant role in representing the dimensional correlations to be learned. Consider two general conditional distributions $p(\mathbf{x}|\mathbf{y})$ and $\tilde{p}(\mathbf{y}|\mathbf{z})$ over finite sets. We define their composition as

$$p \circ \tilde{p}(\mathbf{x}|\mathbf{z}) := \mathbb{E}_{\mathbf{y} \sim \tilde{p}(\cdot|\mathbf{z})} [p(\mathbf{x}|\mathbf{y})] = \sum_{\mathbf{y}} p(\mathbf{x}|\mathbf{y}) \tilde{p}(\mathbf{y}|\mathbf{z}),$$

where this definition can be extended to the continuous case in a straightforward way. Although this is just a convolution of two functions, it can be viewed as a composition of denoising operators in the context of diffusion models. Specifically, given a single-step denoiser $p_{s|t}$ and the finite timesteps $0 = t_0 < t_1 < \dots < t_N = T$, we typically use $p_{t_0|t_1} \circ \dots \circ p_{t_{N-1}|t_N}(\cdot|\mathbf{x}_T)$ with the terminal noise $\mathbf{x}_T \sim q_T$ as a generative sampler.

Notably, the composition can serve as a source of dimensional correlation in discrete diffusion models. Even if one-step denoisers, $p_{s|u}$ and $p_{u|t}$ ($s < u < t$), are dimensionally independent, their composition is generally not. Indeed, dimensionally independent denoisers are successful given hundreds of sampling steps (Austin et al., 2021; Gu et al., 2022). Our method aims at compressing the composition of well-trained denoisers into few-step sampling by learning dimensional correlations.

Let p^ψ be a pretrained teacher model with a product structure and $p_{0|t_1}^\psi \circ \dots \circ p_{t_{N-1}|t_N}^\psi$ be a sufficiently good approximation of $q_{0|T}$, where $0 < t_1 < \dots < t_N = T$ are timesteps. In our distillation, we would like the student model p^θ to approximate the teacher compositions as

$$p_{0|t_n}^\theta \approx p_{0|t_1}^\psi \circ \dots \circ p_{t_{n-1}|t_n}^\psi, \quad n = 1, \dots, N. \quad (2)$$

To achieve this, we provide a way of modeling p^θ that is

capable of representing dimensional correlations in Section 3.2, and we propose a set of loss functions to distill dimensional correlation represented by the compositions of a teacher model in Section 3.3.

3.2. Mixture models to capture dimensional correlations

As an effective instance to represent correlated multivariate categorical distributions, we propose a *mixture model*. We define it as a family of conditional probability distributions that have the following representation for $s < t$:

$$p_{s|t}^\theta(\mathbf{x}_s|\mathbf{x}_t) = \mathbb{E}_\lambda \left[p_{s|t}^\theta(\mathbf{x}_s|\mathbf{x}_t; \lambda) \right], \quad (3)$$

$$\text{where } p_{s|t}^\theta(\mathbf{x}_s|\mathbf{x}_t; \lambda) = \prod_{d=1}^D p_{s|t}^{\theta,d}(x_s^d|\mathbf{x}_t; \lambda).$$

Here, λ is a random variable with an arbitrary distribution. This distribution can be viewed as a convex mixture of product models indexed by λ . See Figure 1 (right) for an intuitive illustration. Despite the fact that $p_{0|t}^\theta(\mathbf{x}_0|\mathbf{x}_t; \lambda)$ is dimensionally independent for each given point λ , this mixture representation is universal in the following sense:

Proposition 1. *For any probability distribution p over \mathcal{S}^D , there exist a probability distribution π and a family of product distributions $p(\mathbf{x}; \lambda) = \prod_{d=1}^D p^d(x^d; \lambda)$ indexed by λ satisfying $p(\mathbf{x}) = \mathbb{E}_{\lambda \sim \pi}[p(\mathbf{x}; \lambda)]$ for all $\mathbf{x} \in \mathcal{S}^D$.*

Indeed, we have $p(\mathbf{x}) = \mathbb{E}_{\mathbf{z} \sim p}[\delta_{\mathbf{z}}(\mathbf{x})]$, where $\delta_{\mathbf{z}}$ is the delta distribution at \mathbf{z} , which is certainly a product distribution. Although the proof is not very informative, the assertion itself implies that the mixture model has sufficient expressive power to capture dimensional correlations. It should also be noted that sampling from this mixture model during the inference has almost no extra computational overhead compared with the conventional product model, since it just requires sampling and insertion of λ (see Section 5.4).

3.3. Consistency for distilling dimensional correlations

We present a set of (two) loss functions that take dimensional correlation into account. Consider we are given a product teacher model, which is denoted as p^ψ . Let p^θ be a general student model (with enough expressive power; an example is given in Section 3.2) that we want to train based on p^ψ .

Distillation loss. We first introduce a *distillation loss*, which forces the student model to approximate the teacher model at time δ ($\ll T$):

$$\mathcal{L}_{\text{distil}}(\theta; \psi, r_\delta, \delta) := \mathbb{E}_{\mathbf{x}_\delta \sim r_\delta} \left[D_{\text{KL}}(p_{0|\delta}^\psi(\cdot|\mathbf{x}_\delta) \| p_{0|\delta}^\theta(\cdot|\mathbf{x}_\delta)) \right],$$

where r_δ ($\approx q_\delta$) is a reference distribution over \mathcal{X} at time δ and D_{KL} is the Kullback–Leibler (KL) divergence. We expect that a single teacher denoising step is enough for a

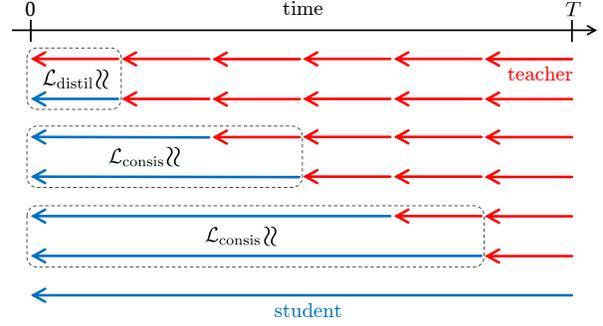


Figure 2. Illustration of how our loss functions work. Through $\mathcal{L}_{\text{distil}}$ and $\mathcal{L}_{\text{consist}}$, we distill multiple teacher denoising steps into fewer steps of student denoiser.

small δ ; the dimensional correlation is mainly incorporated in the following consistency loss (see also Section A.1).

Consistency loss. We then propose a *consistency loss*, which allows the student model to learn the dimensional correlation represented by the teacher denoiser compositions:

$$\begin{aligned} \mathcal{L}_{\text{consist}}(\theta; \psi, r_t, s, u, t) \\ := \mathbb{E}_{\mathbf{x}_t \sim r_t} \left[D_{\text{KL}}(p_{s|u}^\theta \circ p_{u|t}^\psi(\cdot|\mathbf{x}_t) \| p_{s|t}^\theta(\cdot|\mathbf{x}_t)) \right], \end{aligned}$$

where r_t is a reference distribution over \mathcal{X} at time t approximating q_t . While this loss is not straightforward to compute, we discuss how to approximate it in practice with Monte Carlo or control variates in Section A.2. Note that the idea of mixing the teacher denoiser and student denoiser in $\mathcal{L}_{\text{consist}}$ can also be found in the continuous-state setting regarding ODE trajectories (Kim et al., 2024, Fig. 3), but our loss is different in that we work on the compositions of conditional probabilities as in (2).

Figure 2 shows the intuition behind our loss functions. As reference distributions r_δ and r_t , we can either use q_t generated from data or the distribution obtained by applying multiple teacher denoising steps. See Section 4 for their roles and theoretical guarantees on $\mathcal{L}_{\text{distil}}$ and $\mathcal{L}_{\text{consist}}$.

4. Theoretical analysis

In this section, we present an overall theoretical analysis on our distillation method. In Section 4.1, we show that the conventional product model (1) can approximate a data distribution if the model’s marginal is perfectly trained and *given many steps* ($\mathcal{O}(1/N)$ upper bound). We also show that its N -step total variation error can be lower bounded by $\Omega(1/N)$ even for a simple two-dimensional example. Both bounds support the empirical evidences of existing models that work (only) under many steps. In Section 4.2, we prove that the proposed objective functions enable the many-step denoising with a teacher model to be distilled into a few-step

student model, provided that the student model has enough expressive power.

By combining the upper bounds in these results, we informally obtain the following estimate when distilling an N -step sampling process of a teacher model that learns the marginals perfectly:

$$\begin{aligned} d_{\text{TV}}(p(\text{1-step student}), q_0) & (\prec \text{Theorem 1, upper bound}) \\ & \leq d_{\text{TV}}(p(\text{1-step student}), p(\text{N-step teacher})) + \mathcal{O}(1/N) \\ & \leq (\text{Di4C losses in (6)}) + \mathcal{O}(1/N). \quad (\text{Theorem 2}) \end{aligned}$$

Here, $p(n\text{-step model})$ represents the resulting distribution of n -step sampling with the model starting from q_T , and d_{TV} denotes the total variation distance.

4.1. Product models with multi-step sampling can approximate data distribution

We first show that dimensionally independent denoisers with many steps are capable of approximately recovering a data distribution, which has already been empirically observed in existing studies. To consider varying the number of denoising steps, let us work on the continuous-time setting. Let $(\mathbf{x}_t)_{0 \leq t \leq T}$ follow a continuous-time Markov chain over $[0, T]$ and the space $\mathcal{X} = \mathcal{S}^D$ with a factorized forward process, i.e., $q_{t|s}(\mathbf{x}_t|\mathbf{x}_s) = \prod_{d=1}^D q_{t|s}^d(x_t^d|x_s^d)$ for $s < t$. See Section C for more details.

Theorem 1 shows the capability and limitation of a dimensionally independent sampling scheme called *analytical sampling* (Sun et al., 2023) or Tweedie τ -leaping (Lou et al., 2024; Ou et al., 2025), where we use a product-model denoiser $p_{s|t}(\mathbf{x}_s|\mathbf{x}_t) = \prod_{d=1}^D p_{s|t}^d(x_s^d|x_t^d)$ approximating the true marginal as $p_{s|t}^d(x_s^d|x_t^d) \approx q_{s|t}^d(x_s^d|x_t^d)$. Although commonly used, there has been only empirical evidence for the overall efficiency of this dimensionally independent method. Note that Campbell et al. (2022) provides a guarantee for another dimensionally independent method called τ -leaping.

Theorem 1 (N -step analytical sampling approximates data, informal). *Let $q_{t|s}$ be forward transition probabilities that factorize as above and $p_{s|t}$ be a product model with the correct marginals, i.e., $p_{s|t}(\mathbf{x}_s|\mathbf{x}_t) = \prod_{d=1}^D p_{s|t}^d(x_s^d|x_t^d)$ for $s < t$. Under regularity conditions, given timesteps $t_i = iT/N$ for $i = 0, \dots, N$, we have, as $N \rightarrow \infty$,*

$$d_{\text{TV}}(q_0, \mathbb{E}_{\mathbf{x}_T \sim q_T} [p_{t_0|t_1} \circ \dots \circ p_{t_{N-1}|t_N}(\cdot|\mathbf{x}_T)]) = \mathcal{O}(1/N). \quad (4)$$

Moreover, there is an example with $|\mathcal{S}| = D = 2$ such that the left-hand side of (4) is lower-bounded by c/N with some constant $c > 0$ for sufficiently large N .

Proof (sketch). We first prove the following estimate for $0 \leq t - \epsilon < t \leq T$ and $\mathbf{x} \in \mathcal{S}^D$ (Lemma 1):

$$d_{\text{TV}}(q_{t-\epsilon|t}(\cdot|\mathbf{x}), p_{t-\epsilon|t}(\cdot|\mathbf{x})) = \mathcal{O}(\epsilon^2), \quad (5)$$

as $\epsilon \rightarrow 0$. The proof exploits the factorization $q_{t|s}(\mathbf{x}_t|\mathbf{x}_s) = \prod_{d=1}^D q_{t|s}^d(x_t^d|x_s^d)$ and is based on additional continuity assumptions. We then decompose the left-hand side of (4) into N terms by using a triangle-like inequality on d_{TV} between compositional distributions (Proposition 4). The i -th term essentially measures the distance between $q_{t_{i-1}|t_i}$ and $p_{t_{i-1}|t_i}$ and thus is bounded by $\mathcal{O}(1/N^2)$ from (5) with $\epsilon = T/N$. By summing up the N terms, we obtain the desired estimate for the first part. For the second part, we actually construct a concrete example in Section C.3. \square

See Theorem 3 for a formal version. Theorem 1 is important as it underpins the use of dimensionally parallel denoising given sufficient steps, which has been claimed as an advantage of discrete diffusion over autoregressive models whose sampling is sequential (Lou et al., 2024). However, it still requires $\Omega(1/\epsilon)$ steps in order to have a uniform error bound ϵ , according to the latter half of the assertion. We show next that we can further reduce the number of steps with our loss functions, by distilling the distribution of an N -step teacher model into a few-step student model by learning dimensional correlations.

4.2. Our losses can distill multi-step denoising models

Let p^ψ and p^θ respectively be the teacher and student models given in Sections 3.1 & 3.3. The following statement gives a theoretical guarantee for using the proposed loss functions at the appropriate time and distribution settings.

Theorem 2 (Di4C student approximates N -step teacher). *Let $0 = t_0 < \dots < t_N = T$ be timesteps and r_T be a probability distribution on \mathcal{X} . For each n , let $r_{t_n} = \mathbb{E}_{\mathbf{x}_T \sim r_T} [p_{t_n|t_{n+1}}^\psi \circ \dots \circ p_{t_{N-1}|t_N}^\psi(\cdot|\mathbf{x}_T)]$. Then, we have*

$$\begin{aligned} d_{\text{TV}}(r_0, \mathbb{E}_{\mathbf{x}_T \sim r_T} [p_{0|T}^\theta(\cdot|\mathbf{x}_T)]) \\ \leq \frac{1}{\sqrt{2}} \left(\mathcal{L}_{\text{distil}}(\theta; \psi, r_{t_1}, t_1)^{1/2} \right. \\ \left. + \sum_{n=1}^{N-1} \mathcal{L}_{\text{consis}}(\theta; \psi, r_{t_{n+1}}, 0, t_n, t_{n+1})^{1/2} \right). \quad (6) \end{aligned}$$

We can prove this by formalizing the intuition behind Figure 2 (see Section D.1). Note that the right-hand side of inequality (6) becomes zero (so does the left-hand side) if the student model perfectly learns the composition of the teacher as in (2), so learning with these loss functions is feasible in theory if the student model has enough expressive power. Existing theoretical guarantees in consistency-based distillation of continuous-state diffusions typically discuss the case when consistency losses are exactly zero (Song et al., 2023; Daras et al., 2024; Lai et al., 2023), so our guarantee would be interesting in that it explicitly shows the

relationships between the magnitude of loss functions and the total variation bound between the teacher and student.

Regarding the choice of r_t , we should take $r_T = q_T$ if we would like to combine Theorem 2 with Theorem 1 to evaluate Di4C’s overall performance against the data distribution. For r_t with $t < T$, though we can generate samples $\mathbf{x}_t \sim r_t$ by using the teacher model, it might be expensive due to the multi-step inference required. Instead, we can use q_t if we have access to data, which is given by just one-step forward sampling from $q_{t|0}(\cdot|\mathbf{x}_0)$ with the data $\mathbf{x}_0 \sim q_0$. Since r_t is an approximation of q_t (Theorem 1), it would not harm the training quality as long as the teacher model is well-trained.

5. Experimental results

We evaluated our Di4C method in three different diffusion settings, each with distinct teacher models and state spaces. First, we examined continuous-time discrete-state diffusion with pixel-space representations using CIFAR-10, where we distilled from a well-trained U-Net teacher model (Section 5.1). Second, we explored masked diffusion on vector-quantized (VQ) space using ImageNet, working with a transformer-based teacher model designed for masked image generation (Section 5.2). Third, we tested our approach on masked diffusion language models trained on OpenWebText, demonstrating Di4C’s effectiveness in distilling transformer-based diffusion language models (Section 5.3). Finally, we demonstrate that our mixture modeling causes minimal latency overhead in all the above experiments (Section 5.4). These experimental results, spanning different domains, architectures, and diffusion processes, showcase the broad applicability of our method while highlighting its consistent ability to achieve faster sampling.

5.1. Discretized Gaussian diffusion on pixel space

In our first experiment, we adopted the same setting as Campbell et al. (2022): a continuous-time discrete-state Markov process (of discretized Gaussian transition) with the CIFAR-10 image dataset (Krizhevsky, 2009). We used the well-trained model checkpoint provided by Campbell et al. (2022) as our product teacher model (p^ψ). This model outperforms previous discrete-time discrete-state models such as Austin et al. (2021). As in the original paper, we worked directly with the discrete pixel channel values (0 to 255) on $32 \times 32 \times 3$ entries (i.e., $|\mathcal{S}| = 256$, $D = 3072$). The teacher model p^ψ has a U-net architecture (Ho et al., 2020) tailored for discrete diffusion, which is fed a time feature at each upsampling/downsampling stage.

To obtain an architecture for our student mixture model (3), we slightly extended the teacher’s architecture so that it accepts a conditioning with $\lambda \sim \text{Unif}([0, 1])$ (uniform distribution over $[0, 1]$) in this experiment, by following the

original implementation of time conditioning. In training, the student model was initialized by the teacher network parameters with additional zero-initialized subnetworks concerning λ . See Section F.1.3 for implementation details.

Table 1. Comparison of models on CIFAR-10 dataset. Fréchet inception distance (FID ↓) against training dataset and inception score (IS ↑) measured with 50,000 samples are shown in this order (FID / IS).

	10 steps	20 steps	40 steps
teacher	32.61 / 7.59	12.36 / 8.55	8.01 / 8.77
student	20.64 / 8.29	9.77 / 8.52	9.66 / 8.28
hybrid	25.54 / 8.00	9.47 / 8.56	8.02 / 8.43

Table 1 shows the results. The “hybrid” model used the student model for the first half of the denoising process and then switched to the teacher model for the remaining steps during inference. The student substantially improved the metrics compared with the teacher in 10-step sampling, while the benefits of our method diminished as the number of steps grew. In contrast, the hybrid model was the best at 20 steps and on par with the teacher in 40-step FID. We hypothesize that Di4C is particularly effective when using fewer sampling steps, as this is where capturing complex dimensional correlations is crucial (Theorem 1). We also give additional results for different sampling steps in Table 4.

Table 2. Comparison of teacher and student with PC steps. FID against CIFAR-10 training dataset measured with 10,000 generated samples are displayed (so there is discrepancy from Table 1). Number of function evaluations (NFEs) is adjusted to 20, with varying budget allocations between denoising and PC steps.

	NFE	6+7+7	10+5+5	14+3+3	20+0+0
teacher	57.57	32.74	21.32	14.42	
student	44.57	22.63	14.25	11.81	

In Table 2, we also compare the teacher and student models in combination with the predictor-corrector (PC, Campbell et al., 2022) steps. The notation $m+n+n$ indicates that we used n corrector steps before each of the final two out of m denoising steps, imitating the approach in Campbell et al. (2022). For each setting, the best PC step-size among $\{10^{-3}, 10^{-4}, 10^{-5}\}$ was chosen. Since a PC step requires an additional network evaluation (as expensive as one step of denoising), we adjusted the total NFE in the table. While adding PC steps shows improvement upon the vanilla m -step denoising (see also Table 4) in both teacher and student models, allocating all the NFE budget to denoising steps performs better in this small NFE regime; this observation aligns with Campbell et al. (2022, Fig. 4).

5.2. Masked image modeling on VQ space

Next, we evaluated our method by applying it to a larger-scale image generation model. For this purpose, we adopted the framework of MaskGIT (Chang et al., 2022) and worked on the ImageNet dataset (Deng et al., 2009) at 256×256 resolution. MaskGIT is one of the state-of-the-art image generation methods, based on masked diffusion modeling. Its generative (backward) process relies on heuristics including confidence-based sampling (Section F.2.2). The variant we used also uses discrete classifier-free guidance (CFG; Section F.2.3). We demonstrate that Di4C can enhance image generation even in combination with such heuristics.

In our setting, the model comprised two main components: a VQGAN (Esser et al., 2021) pre-trained on the ImageNet dataset and a masked diffusion model trained on the VQ space. The VQGAN encodes a 256×256 resolution image into $16 \times 16 = 256$ tokens, each drawn from a shared codebook \mathcal{S}^* of size 1024. The forward process of our diffusion modeling, denoted as $q_{t|0}(\cdot|\mathbf{x}_0)$, independently replaces x_0^d with [MASK] at a certain probability m_t for each d (i.e., $\mathcal{S} := \mathcal{S}^* \cup \{[\text{MASK}] \}$, $|\mathcal{S}^*| = 1025$, and $D = 256$). The masking probability m_t increases monotonically from $m_0 = 0$ to $m_1 = 1$, resulting in completely masked sequences at $t = 1$ regardless of \mathbf{x}_0 .

As the teacher model, we used the PyTorch-based implementation by Besnier & Chen (2023), which replicates the performance of the original MaskGIT. It uses a bidirectional transformer p^ψ that estimates the distribution of each token given an ImageNet label c and a partially masked sequence $\mathbf{x}_t \in \mathcal{S}^D$: $p_{0|t}^{\psi,d}(\cdot|\mathbf{x}_t, c) \approx q_{0|t}^d(\cdot|\mathbf{x}_t, c)$. The input sequence length is 257 including the embedding of c . Additionally, the model supports unconditional generation, enabling discrete CFG to match the original MaskGIT’s performance. In this configuration, the teacher model generates high-quality samples with only 8 steps, which is significantly faster than typical diffusion-based generative models (see, e.g., Chang et al., 2022, Table 1). In the implementation of our mixture (student) model, we simply added a single token embedding coming from $\lambda \sim \text{Unif}([0, 1])$, so the input sequence length is 258 instead of 257 in the student models.

Figure 3 shows FID-IS curves of 4-step sampling with various CFG coefficients: $w_{\text{cfg}} \in \{5, 6, 7, 8, 9, 10, 11, 12\}$ for 4 steps and $w_{\text{cfg}} \in \{2, 2.5, 3, 3.5, 4\}$ for 8 steps (teacher only). In the figure, **di4c** represents a model trained using the standard Di4C method, while **di4c-d** represents a model trained with an additional datapoint loss (see Section F.2.4 for details). As shown in the figure, the FID-IS performance of the 4-step Di4C models closely matched that of the 8-step sampling of the teacher model. This result indicates that Di4C can achieve an approximate 2x speed-up in the sampling process.

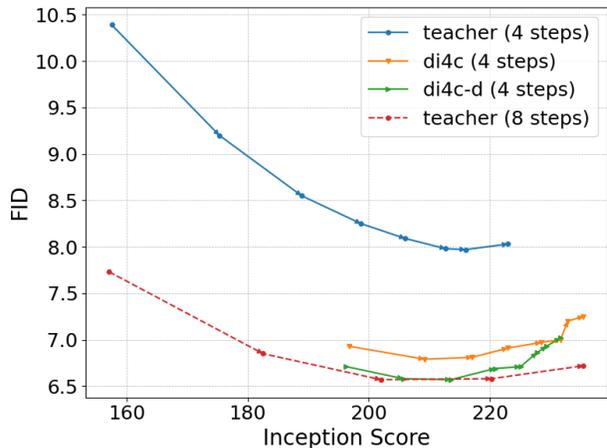


Figure 3. FID-IS curves of 4/8-step teacher and 4-step Di4C models on ImageNet 256×256 when varying CFG coefficients. Arrows connect experimental results (dots) at each CFG coefficient in ascending order.

See Section F.2 for further details of this experiment. In particular, Table 8 shows the detailed quantitative performance (including precision and recall) of each model at its best CFG coefficient.

5.3. Masked diffusion language models

Finally, we examined Di4C in language modeling. As teacher models, we adopted two versions of pretrained SDTT models (Deschenaux & Gulcehre, 2025, with a KL target), which we refer to as **sdtt-6** and **sdtt-7**. They were obtained after respectively 6 and 7 rounds of distillation of a masked diffusion language model (MDLM, Sahoo et al., 2024. See also its concurrent works: Shi et al. 2024; Ou et al. 2025) trained on the OpenWebText dataset (Gokaslan & Cohen, 2019). As a forward process, they use a masked diffusion with a GPT-2 tokenizer (Radford et al., 2019), which is essentially the same as the one explained in Section 5.2, while we had $D = 1024$ and $|\mathcal{S}^*| = 50257$ this time. In the mixture modeling for Di4C, we added one token from $\lambda \sim \text{Unif}([0, 1])$ to the transformer, similarly to the previous section. See Section F.3 for details.

Figure 4 shows the results of applying Di4C to the SDTT models. For $n = 6, 7$, **sdtt- n + di4c** represents the model obtained by fine-tuning **sdtt- n** using Di4C. Similarly, **sdtt- n + di4c²** represents the model obtained by applying another round of Di4C to **sdtt- n + di4c** (see Section A.4).

Unconditional generation. Figure 4(a) shows a performance comparison for unconditional generation. Each model generated 1024 samples of 1024 tokens, with the number of sampling steps in $\{8, 16, 32, 64, 128\}$. Following the previous work (Lou et al., 2024; Sahoo et al., 2024;

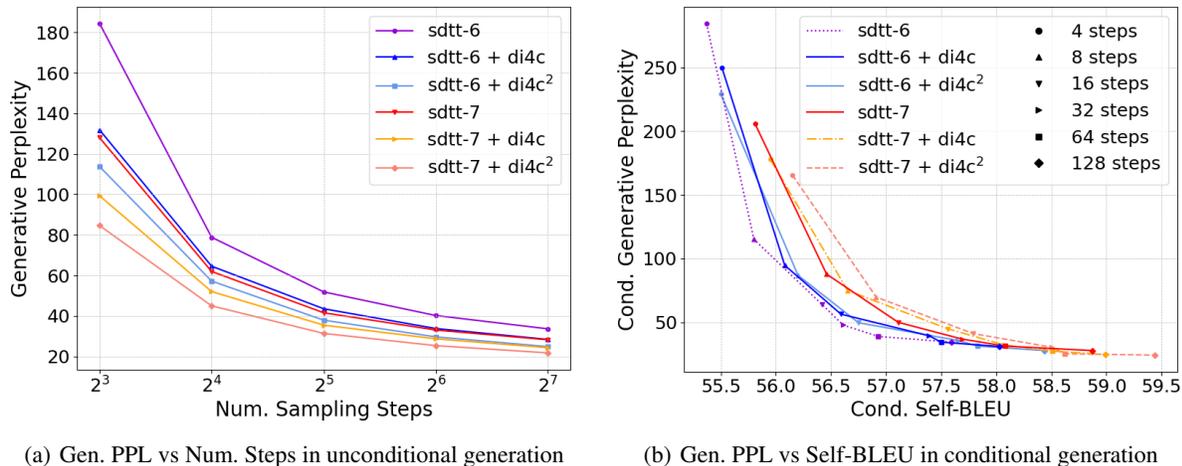


Figure 4. Comparison of SDTT checkpoints (Deschenaux & Gulcehre, 2025) and their Di4C distillations.

Deschenaux & Gulcehre, 2025), we used the GPT-2-large model (Radford et al., 2019) to evaluate the generative perplexity. The results show that applying one round of Di4C lead to a similar level of improvement as applying one round of SDTT (compare *sdtt-6 + di4c* with *sdtt-7*). Also, two rounds of Di4C (+ *di4c²*) showed speed-ups of more than 2x to match the teacher’s performance of 64 or 128 steps.

Conditional generation. We also tested the quality-diversity tradeoff in conditional generation. We used 256 samples from the WebText dataset (OpenAI, 2019), and each model, on the basis of the first 50 tokens of the sample, generated 5 continuations of 50 tokens, following Deschenaux & Gulcehre (2025). We computed the generative perplexity of the generated continuations and the Self-BLEU score (Zhu et al., 2018) of the 5 completed samples starting from the same prompt, and averaged them over the 256 prompts. Self-BLEU was computed by GPT-2 tokenizer with equal weights on n -gram for $n = 1, 2, 3, 4$ and scaled from 0 to 100 (lower Self-BLEU indicates higher diversity). Note that (Self-)BLEU has several parameters including tokenization and maximum length for n -grams (Post, 2018), so our numbers are not directly comparable to those from an existing work such as Agarwal et al. (2024).

Figure 4(b) shows the quality-diversity tradeoff curves of each model with various numbers of sampling steps. As shown with solid lines, both *sdtt-6 + di4c* and *sdtt-6 + di4c²* consistently achieved higher diversity than *sdtt-7* with the same number of sampling steps while maintaining comparable or better generative perplexity.

5.4. Latency overhead of mixture modeling

In this section, we demonstrate that the computational overhead in introducing mixture modeling is negligible com-

pared with its performance gain. Table 3 shows the quantitative results on the latency. The experiments listed in Table 3 correspond to those described in Sections 5.1–5.3, where the chosen student models are the best models (*di4c-d* in ImageNet and *sdtt-7 + di4c²* in OpenWebText). The table shows the average runtime and standard deviation, calculated over 10 batches with batch sizes of 50, 64, and 16, respectively. We tested unconditional generation for CIFAR-10 and OpenWebText, and a class-conditional generation with a uniform random class for ImageNet with a classifier-free guidance.

Table 3. Latency comparison between teacher (product) and student (mixture) models.

Experiment	# steps	teacher (sec.)	student (sec.)
CIFAR-10	10	0.5515±0.0024	0.5786±0.0017
ImageNet	4	2.0741±0.0035	2.0734±0.0043
OpenWebText	16	3.3409±0.0417	3.4817±0.0730

In the CIFAR-10 and OpenWebText experiments, the overhead from using mixture models was up to 5%, which is negligible compared with their performance gain. In the ImageNet experiment, the difference between the student and teacher was within the statistical noise. This is likely because the computational bottleneck in this case was not the logit inference itself but rather the sampling process that followed it.

6. Related work

Speeding up continuous diffusion models. In continuous diffusion models, knowledge distillation (Luhman & Luhman, 2021; Salimans & Ho, 2022; Meng et al., 2023; Zheng

et al., 2023a) and consistency-type techniques (Song et al., 2023; Song & Dhariwal, 2023; Kim et al., 2024) enable single- or few-step sampling, most of which are tailored for probability flow ODEs. Among the studies on continuous diffusion, the work by Li et al. (2024) is particularly relevant to our mixture modeling approach. They highlight the limited expressive power of unimodal Gaussian distributions in denoising continuous diffusions and demonstrate that using Gaussian mixtures can substantially reduce the number of sampling steps required in non-ODE diffusions.

Speeding up discrete diffusion models. For faster sample generation in discrete diffusion models, Park et al. (2025) propose a post-hoc optimization of sampling schedules and outperforms the uniform partitioning of $[0, T]$. To speed up MDLMs (Sahoo et al., 2024); SDTT (Deschenaux & Gulcehre, 2025) gathers the logits at unmasked tokens from a pretrained teacher MDLM throughout the denoising process, achieving a speed-up of approximately 32x compared with the teacher model. In addition, the confidence-based sampling by MaskGIT (Chang et al., 2022) is essentially a heuristic to enable faster sampling in masked diffusions. These approaches can be combined with Di4C; we indeed have demonstrated it for the latter two.

Two concurrent works are particularly relevant to our research, both of which attempt to incorporate dimensional correlations into discrete diffusion, specifically within the MDLM context. Liu et al. (2024) combine pretrained autoregressive and masked diffusion models, achieving superior performance compared with using either model alone. Meanwhile, Xu et al. (2025) propose an energy-based model to modify the dimensionally independent sampler. While both approaches improve upon the vanilla MDLM, their sampling processes can incur some time and memory overhead. This is due to the use of an additional non-diffusion model in both methods, and the reliance on importance sampling in the latter. In this regard, our mixture modeling allows for a simpler sampling process with minimal modification to the original diffusion. However, since the Di4C loss functions are model-agnostic, we should explore the combination of Di4C losses and other proposed models capturing dimensional correlations in future research.

7. Conclusion

In this paper, as current discrete diffusion models ignore the dimensional correlations that need to be incorporated to realize few-step models, we proposed Di4C, a method for distilling pretrained discrete diffusion models. Di4C provides a set of loss functions for models that can capture dimensional correlations, an example of which is the mixture model. As a theoretical contribution, we proved that the existing discrete diffusion models with many steps

can indeed recover a data distribution, even without modeling dimensional correlations. We also proved that such many-step models can be distilled into few-step ones, if we use the Di4C loss functions with a model that has enough expressive power, such as a mixture model. In numerical experiments, we confirmed the efficiency of our framework upon teacher models across multiple domains: improving sample quality in 10-step sampling on CIFAR-10 with a discretized Gaussian diffusion, achieving a 2x speed-up in ImageNet 256×256 generation with masked image modeling, and successfully distilling already-distilled masked diffusion language models on OpenWebText while maintaining generation diversity.

However, there are still problems to be solved. For example, although we can distill many-step models into one-step ones in theory (Theorem 2), our empirical results show only 2x fewer steps or so. To address this point, we need to further optimize the architecture (mainly concerning λ) and training hyperparameters, while the iterated Di4C training (Section A.4) can be a promising candidate.

Impact Statement

This paper presents a novel method for reducing the number of sampling steps required for high-quality generation in discrete data domains. While there are many potential social impacts as usual in the generative AI field, we do not foresee any specific problems particularly caused by this (mostly theoretical) work.

Acknowledgments

We would like to thank Masato Ishii, Chieh-Hsin Lai, Naoki Murata, Bac Nguyen, and Toshimitsu Uesaka for their constructive feedback. We also thank the anonymous reviewers for their valuable comments and suggestions.

References

- Agarwal, R., Vieillard, N., Zhou, Y., Stanczyk, P., Garea, S. R., Geist, M., and Bachem, O. On-policy distillation of language models: Learning from self-generated mistakes. In *Proceedings of the Twelfth International Conference on Learning Representations*, 2024.
- Anderson, W. J. *Continuous-time Markov chains: An applications-oriented approach*. Springer Science & Business Media, 2012.
- Austin, J., Johnson, D. D., Ho, J., Tarlow, D., and Van Den Berg, R. Structured denoising diffusion models in discrete state-spaces. In *Advances in Neural Information Processing Systems*, volume 34, pp. 17981–17993, 2021.
- Besnier, V. and Chen, M. A pytorch reproduction of

- masked generative image transformer. *arXiv preprint arXiv:2310.14400*, 2023.
- Bird, S., Klein, E., and Loper, E. *Natural language processing with Python: analyzing text with the natural language toolkit*. O’Reilly Media, Inc., 2009.
- Blanchard, P., Higham, D. J., and Higham, N. J. Accurately computing the log-sum-exp and softmax functions. *IMA Journal of Numerical Analysis*, 41(4):2311–2330, 2021.
- Blattmann, A., Dockhorn, T., Kulal, S., Mendelevitch, D., Kilian, M., Lorenz, D., Levi, Y., English, Z., Voleti, V., Letts, A., Jampani, V., and Rombach, R. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*, 2023.
- Campbell, A., Benton, J., De Bortoli, V., Rainforth, T., Deligiannidis, G., and Doucet, A. A continuous time framework for discrete denoising models. In *Advances in Neural Information Processing Systems*, volume 35, pp. 28266–28279, 2022.
- Canonne, C. L. A short note on an inequality between KL and TV. *arXiv preprint arXiv:2202.07198*, 2022.
- Chang, H., Zhang, H., Jiang, L., Liu, C., and Freeman, W. T. Maskgit: Masked generative image transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11315–11325, 2022.
- Chen, N., Zhang, Y., Zen, H., Weiss, R. J., Norouzi, M., and Chan, W. Wavegrad: Estimating gradients for waveform generation. In *Proceedings of the Ninth International Conference on Learning Representations*, 2021.
- Comunità, M., Zhong, Z., Takahashi, A., Yang, S., Zhao, M., Saito, K., Ikemiya, Y., Shibuya, T., Takahashi, S., and Mitsufuji, Y. Specmaskgit: Masked generative modeling of audio spectrograms for efficient audio synthesis and beyond. *arXiv preprint arXiv:2406.17672*, 2024.
- Cover, T. M. and Thomas, J. A. *Elements of information theory*. John Wiley & Sons, second edition, 2006.
- Daras, G., Dagan, Y., Dimakis, A., and Daskalakis, C. Consistent diffusion models: Mitigating sampling drift by learning to be consistent. In *Advances in Neural Information Processing Systems*, volume 36, 2024.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255. Ieee, 2009.
- Deschenaux, J. and Gulcehre, C. Beyond autoregression: Fast LLMs via self-distillation through time. In *Proceedings of the Thirteenth International Conference on Learning Representations*, 2025.
- Elfwing, S., Uchibe, E., and Doya, K. Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural Networks*, 107:3–11, 2018.
- Esser, P., Rombach, R., and Ommer, B. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12873–12883, 2021.
- Evans, Z., Carr, C., Taylor, J., Hawley, S. H., and Pons, J. Fast timing-conditioned latent audio diffusion. In *Proceedings of the Forty-First International Conference on Machine Learning*, pp. 12652–12665, 2024.
- Gillespie, D. T. Approximate accelerated stochastic simulation of chemically reacting systems. *The Journal of Chemical Physics*, 115(4):1716–1733, 2001.
- Glasserman, P. *Monte Carlo methods in financial engineering*, volume 53. Springer, 2004.
- Gokaslan, A. and Cohen, V. Openwebtext corpus. <http://Skylion007.github.io/OpenWebTextCorpus>, 2019.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pp. 2672–2680, 2014.
- Gu, S., Chen, D., Bao, J., Wen, F., Zhang, B., Chen, D., Yuan, L., and Guo, B. Vector quantized diffusion model for text-to-image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10696–10706, 2022.
- Harvey, W., Naderiparizi, S., Masrani, V., Weilbach, C., and Wood, F. Flexible diffusion modeling of long videos. In *Advances in Neural Information Processing Systems*, volume 35, pp. 27953–27965, 2022.
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A. beta-VAE: Learning basic visual concepts with a constrained variational framework. In *Proceedings of the Fifth International Conference on Learning Representations*, 2017.
- Ho, J., Jain, A., and Abbeel, P. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, pp. 6840–6851, 2020.
- Ho, J., Salimans, T., Gritsenko, A., Chan, W., Norouzi, M., and Fleet, D. J. Video diffusion models. In *Advances in Neural Information Processing Systems*, volume 35, pp. 8633–8646, 2022.

- Hoogeboom, E., Nielsen, D., Jaini, P., Forré, P., and Welling, M. Argmax flows and multinomial diffusion: Learning categorical distributions. In *Advances in Neural Information Processing Systems*, volume 34, pp. 12454–12465, 2021.
- Kim, D., Lai, C.-H., Liao, W.-H., Murata, N., Takida, Y., Uesaka, T., He, Y., Mitsufuji, Y., and Ermon, S. Consistency trajectory models: Learning probability flow ODE trajectory of diffusion. In *Proceedings of the Twelfth International Conference on Learning Representations*, 2024.
- Kingma, D. P. and Welling, M. Auto-encoding variational Bayes. In *Proceedings of the Second International Conference on Learning Representations*, 2014.
- Kong, Z., Ping, W., Huang, J., Zhao, K., and Catanzaro, B. Diffwave: A versatile diffusion model for audio synthesis. In *Proceedings of the Ninth International Conference on Learning Representations*, 2021.
- Krizhevsky, A. Learning multiple layers of features from tiny images. Master’s thesis, University of Toronto, 2009.
- Lai, C.-H., Takida, Y., Uesaka, T., Murata, N., Mitsufuji, Y., and Ermon, S. On the equivalence of consistency-type models: Consistency models, consistent diffusion models, and fokker-planck regularization. In *ICML 2023 Workshop on Structured Probabilistic Inference & Generative Modeling*, 2023.
- Li, Y., van Breugel, B., and van der Schaar, M. Soft mixture denoising: Beyond the expressive bottleneck of diffusion models. In *Proceedings of the Twelfth International Conference on Learning Representations*, 2024.
- Liu, A., Broadrick, O., Niepert, M., and Broeck, G. V. d. Discrete copula diffusion. *arXiv preprint arXiv:2410.01949*, 2024.
- Lou, A., Meng, C., and Ermon, S. Discrete diffusion modeling by estimating the ratios of the data distribution. In *Proceedings of the Forty-First International Conference on Machine Learning*, 2024.
- Lu, C., Zhou, Y., Bao, F., Chen, J., Li, C., and Zhu, J. DPM-solver: A fast ODE solver for diffusion probabilistic model sampling in around 10 steps. In *Advances in Neural Information Processing Systems*, volume 35, pp. 5775–5787, 2022a.
- Lu, C., Zhou, Y., Bao, F., Chen, J., Li, C., and Zhu, J. DPM-solver++: Fast solver for guided sampling of diffusion probabilistic models. *arXiv preprint arXiv:2211.01095*, 2022b.
- Luhman, E. and Luhman, T. Knowledge distillation in iterative generative models for improved sampling speed. *arXiv preprint arXiv:2101.02388*, 2021.
- Meng, C., Rombach, R., Gao, R., Kingma, D., Ermon, S., Ho, J., and Salimans, T. On distillation of guided diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14297–14306, 2023.
- Nichol, A., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., and Chen, M. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. In *Proceedings of the Thirty-Eighth International Conference on Machine Learning*, 2021.
- Oates, C. J., Girolami, M., and Chopin, N. Control functionals for monte carlo integration. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 79(3):695–718, 2017.
- OpenAI. GPT-2 output dataset, 2019. URL <https://github.com/openai/gpt-2-output-dataset>.
- Ou, J., Nie, S., Xue, K., Zhu, F., Sun, J., Li, Z., and Li, C. Your absorbing discrete diffusion secretly models the conditional distributions of clean data. In *Proceedings of the Thirteenth International Conference on Learning Representations*, 2025.
- Park, Y.-H., Lai, C.-H., Hayakawa, S., Takida, Y., and Mitsufuji, Y. Jump your steps: Optimizing sampling schedule of discrete diffusion models. In *Proceedings of the Thirteenth International Conference on Learning Representations*, 2025.
- Pillutla, K., Swayamdipta, S., Zellers, R., Thackstun, J., Welleck, S., Choi, Y., and Harchaoui, Z. Mauve: Measuring the gap between neural text and human text using divergence frontiers. In *Advances in Neural Information Processing Systems*, volume 34, pp. 4816–4828, 2021.
- Post, M. A call for clarity in reporting BLEU scores. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pp. 186–191, 2018.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., and Sutskever, I. Language models are unsupervised multitask learners. *OpenAI blog*, 2019.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Ommer, B. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10684–10695, 2022.

- Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E. L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., Ho, J., Fleet, D. J., and Norouzi, M. Photorealistic text-to-image diffusion models with deep language understanding. In *Advances in Neural Information Processing Systems*, volume 35, pp. 36479–36494, 2022.
- Sahoo, S. S., Arriola, M., Schiff, Y., Gokaslan, A., Marroquin, E., Chiu, J. T., Rush, A., and Kuleshov, V. Simple and effective masked diffusion language models. In *Advances in Neural Information Processing Systems*, volume 36, 2024.
- Salimans, T. and Ho, J. Progressive distillation for fast sampling of diffusion models. In *Proceedings of the Tenth International Conference on Learning Representations*, 2022.
- Shi, J., Han, K., Wang, Z., Doucet, A., and Titsias, M. Simplified and generalized masked diffusion for discrete data. In *Advances in Neural Information Processing Systems*, volume 37, pp. 103131–103167, 2024.
- Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., and Ganguli, S. Deep unsupervised learning using nonequilibrium thermodynamics. In *Proceedings of the Thirty-Sixth International Conference on Machine Learning*, pp. 2256–2265, 2015.
- Song, J., Meng, C., and Ermon, S. Denoising diffusion implicit models. In *Proceedings of the Ninth International Conference on Learning Representations*, 2021a.
- Song, Y. and Dhariwal, P. Improved techniques for training consistency models. In *Proceedings of the Eleventh International Conference on Learning Representations*, 2023.
- Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. Score-based generative modeling through stochastic differential equations. In *Proceedings of the Ninth International Conference on Learning Representations*, 2021b.
- Song, Y., Dhariwal, P., Chen, M., and Sutskever, I. Consistency models. In *Proceedings of the Fortieth International Conference on Machine Learning*, pp. 32211–32252. PMLR, 2023.
- Sun, H., Yu, L., Dai, B., Schuurmans, D., and Dai, H. Score-based continuous-time discrete diffusion models. In *Proceedings of the Eleventh International Conference on Learning Representations*, 2023.
- Tang, Z., Gu, S., Bao, J., Chen, D., and Wen, F. Improved vector quantized diffusion models. *arXiv preprint arXiv:2205.16007*, 2022.
- van den Oord, A., Vinyals, O., and Kavukcuoglu, K. Neural discrete representation learning. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- von Eitzen, H. Continuous function with continuous one-sided derivative. Mathematics Stack Exchange, 2014. URL <https://math.stackexchange.com/q/975094>. (version: 2014-10-16).
- Williams, R. J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8:229–256, 1992.
- Xiao, Z., Kreis, K., and Vahdat, A. Tackling the generative learning trilemma with denoising diffusion gans. In *Proceedings of the Tenth International Conference on Learning Representations*, 2022.
- Xu, M., Geffner, T., Kreis, K., Nie, W., Xu, Y., Leskovec, J., Ermon, S., and Vahdat, A. Energy-based diffusion language models for text generation. In *Proceedings of the Thirteenth International Conference on Learning Representations*, 2025.
- Zhang, Q. and Chen, Y. Fast sampling of diffusion models with exponential integrator. In *Proceedings of the Eleventh International Conference on Learning Representations*, 2023.
- Zhao, S., Song, J., and Ermon, S. InfoVAE: Balancing learning and inference in variational autoencoders. In *AAAI Conference on Artificial Intelligence*, pp. 5885–5892, 2019.
- Zheng, H., Nie, W., Vahdat, A., Azizzadenesheli, K., and Anandkumar, A. Fast sampling of diffusion models via operator learning. In *Proceedings of the Fortieth International Conference on Machine Learning*, pp. 42390–42402, 2023a.
- Zheng, K., Lu, C., Chen, J., and Zhu, J. DPM-solver-v3: Improved diffusion ODE solver with empirical model statistics. In *Advances in Neural Information Processing Systems*, volume 36, pp. 55502–55542, 2023b.
- Zhu, Y., Lu, S., Zheng, L., Guo, J., Zhang, W., Wang, J., and Yu, Y. Taxygen: A benchmarking platform for text generation models. In *The Forty-First International ACM SIGIR Conference on Research & Development in Information Retrieval*, pp. 1097–1100, 2018.

A. Training techniques for Di4C

In this section, we review the novel loss functions of Di4C and the mixture model given in Section 3.3 from an algorithmic perspective, and we provide a set of techniques to stably train it. Specifically, we introduce techniques to make the computation of the loss functions scalable through Monte Carlo integration and control variate methods.

Before going into the details of the training techniques, we introduce two auxiliary loss functions, which we can use in addition to $\mathcal{L}_{\text{distil}}$ and $\mathcal{L}_{\text{consis}}$ for practical improvements. One is the *datapoint loss* that directly computes the negative log-likelihood with respect to the data distribution (e.g., Austin et al., 2021, Eq. 5), which we can use when we have access to data q_0 :

$$\mathcal{L}_{\text{data}}(\theta; t) := \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_t) \sim q_{0,t}} \left[-\log p_{0|t}^\theta(\mathbf{x}_0 | \mathbf{x}_t) \right]. \quad (7)$$

The other is the following *marginal loss*, which is easier to compute, under the assumption that the teacher model sufficiently learns the true marginal, i.e., $p_{0|t}^{\psi,d} \approx q_{0|t}^d$:

$$\mathcal{L}_{\text{marginal}}(\theta; \psi, r_t, t) := \mathbb{E}_{\mathbf{x}_t \sim r_t} \left[\sum_{d=1}^D D_{\text{KL}}(p_{0|t}^{\psi,d}(\cdot | \mathbf{x}_t) \| p_{0|t}^{\theta,d}(\cdot | \mathbf{x}_t)) \right]. \quad (8)$$

A.1. Surrogate of distillation loss

Since the exact evaluation of $\mathcal{L}_{\text{distil}}$ with a mixture model seems intractable, we consider an upper bound of $\tilde{\mathcal{L}}_{\text{distil}}$ as a practical alternative:

$$\begin{aligned} \mathcal{L}_{\text{distil}}(\theta; \psi, r_\delta, \delta) &= \mathbb{E}_{\mathbf{x}_\delta \sim r_\delta} \left[D_{\text{KL}}(p_{0|\delta}^\psi(\cdot | \mathbf{x}_\delta) \| \mathbb{E}_\lambda [p_{0|\delta}^\theta(\cdot | \mathbf{x}_\delta; \lambda)]) \right] \\ &\leq \mathbb{E}_{\mathbf{x}_\delta \sim r_\delta} \mathbb{E}_\lambda \left[D_{\text{KL}}(p_{0|\delta}^\psi(\cdot | \mathbf{x}_\delta) \| p_{0|\delta}^\theta(\cdot | \mathbf{x}_\delta; \lambda)) \right] \\ &\leq \mathbb{E}_{\lambda, \mathbf{x}_\delta \sim r_\delta} \left[\sum_{d=1}^D D_{\text{KL}}(p_{0|\delta}^{\psi,d}(\cdot | \mathbf{x}_\delta) \| p_{0|\delta}^{\theta,d}(\cdot | \mathbf{x}_\delta; \lambda)) \right] =: \tilde{\mathcal{L}}_{\text{distil}}(\theta; \psi, r_\delta, \delta). \end{aligned}$$

Here, the inequality is given by the convexity of KL divergence (see Proposition 3). The upper bound $\tilde{\mathcal{L}}_{\text{distil}}$ (and then $\mathcal{L}_{\text{distil}}$) becomes zero if the student denoiser coincides with the teacher for the time interval $[0, \delta]$, regardless of λ . Therefore, the use of this upper bound is feasible if p^θ has enough expressive power.

A.2. Surrogate of consistency loss

We consider $\mathcal{L}_{\text{consis}}$ in this section. As $p_{s|u}^\theta$ is more “reliable” than $p_{s|t}^\theta$ (since $s < u < t$), we consider only the gradient of $\mathcal{L}_{\text{consis}}$ concerning $p_{s|t}^\theta$ and ignore the gradient coming from $p_{s|u}^\theta$. Therefore, we conduct stochastic gradient descent on θ with the loss

$$D_{\text{KL}}(p_{s|u}^{\text{sg}(\theta)} \circ p_{u|t}^\psi(\cdot | \mathbf{x}_t) \| p_{s|t}^\theta(\cdot | \mathbf{x}_t)) = H(p_{s|u}^{\text{sg}(\theta)} \circ p_{u|t}^\psi(\cdot | \mathbf{x}_t), p_{s|t}^\theta(\cdot | \mathbf{x}_t)) + \text{const.}, \quad (9)$$

where $\text{sg}(\cdot)$ is the stop-gradient operator (van den Oord et al., 2017) and $H(p, q) = \mathbb{E}_{\mathbf{x} \sim p} [-\log q(\mathbf{x})]$ is the cross entropy between p and q . We hereby ignore the constant term in (9) and consider how to efficiently compute the cross entropy term.

Most naively, by using finite samples $\mathbf{x}_s^{(1)}, \dots, \mathbf{x}_s^{(M)} \sim_{\text{iid}} p_{s|u}^{\text{sg}(\theta)} \circ p_{u|t}^\psi(\cdot | \mathbf{x}_t)$ and $\lambda_1, \dots, \lambda_N \sim_{\text{iid}} \lambda$, we can approximate this cross entropy by two-fold Monte Carlo:

$$\begin{aligned} &H(p_{s|u}^{\text{sg}(\theta)} \circ p_{u|t}^\psi(\cdot | \mathbf{x}_t), p_{s|t}^\theta(\cdot | \mathbf{x}_t)) \\ &\approx -\frac{1}{M} \sum_{j=1}^M \log p_{s|t}^\theta(\mathbf{x}_s^{(j)} | \mathbf{x}_t) \approx -\frac{1}{M} \sum_{j=1}^M \log \left(\frac{1}{N} \sum_{i=1}^N p_{s|t}^\theta(\mathbf{x}_s^{(j)} | \mathbf{x}_t; \lambda_i) \right). \end{aligned} \quad (10)$$

Although the value of each $p_{s|t}^\theta(\mathbf{x}_s^{(j)} | \mathbf{x}_t; \lambda_i) = \prod_{d=1}^D p_{s|t}^{\theta,d}(\mathbf{x}_s^{(j),d} | \mathbf{x}_t; \lambda_i)$ can be extremely small due to the D -fold product, we can exploit the log-sum-exp structure:

$$\log \left(\sum_{i=1}^N p_{s|t}^\theta(\mathbf{x}_s^{(j)} | \mathbf{x}_t; \lambda_i) \right) = \underbrace{\log \left(\sum_{i=1}^N \exp \left(\sum_{d=1}^D \log p_{s|t}^{\theta,d}(\mathbf{x}_s^{(j),d} | \mathbf{x}_t; \lambda_i) \right) \right)}_{\text{log-sum-exp}},$$

which is implemented as a function with additional stabilization to avoid under/overflows in some of the common numerical packages including PyTorch. See (Blanchard et al., 2021) for details of numerical properties associated with the log-sum-exp structure.

Dimensionally independent control variate. Although naive Monte Carlo sampling with a sufficiently large sample size can approximate the left-hand side of Eq. (10) well, a small batch can cause high variance in the evaluation of the expected values. An established way of stabilizing Monte Carlo integration is to use so-called *control variates* (Glasserman, 2004; Oates et al., 2017), also known as a *baseline* in reinforcement learning (Williams, 1992). To estimate an expectation $\mathbb{E}[f]$, we can subtract another function/random variable g , called a *control variate*, whose integral we know or can compute more precisely than Monte Carlo, and execute the Monte Carlo for $f - g$ by using the decomposition $\mathbb{E}[f] = \mathbb{E}[f - g] + \mathbb{E}[g]$. See Section E for a more detailed explanation. As a concrete application of this technique, we below propose the use of a dimensionally independent control variate.

We first exploit the compositional form of $p_{s|u}^{\text{sg}(\theta)} \circ p_{u|t}^\psi(\cdot|\mathbf{x}_t)$, which is more informative than $\mathbf{x}_s^{(j)}$, the pure samples in the Monte Carlo approach. We can write it in an expectation as follows:

$$p_{s|u}^{\text{sg}(\theta)} \circ p_{u|t}^\psi(\cdot|\mathbf{x}_t) = \mathbb{E}_{\lambda, \mathbf{x}_u \sim p_{u|t}^\psi(\cdot|\mathbf{x}_t)} \left[p_{s|u}^{\text{sg}(\theta)}(\cdot|\mathbf{x}_u; \lambda) \right]. \quad (11)$$

To simplify (11), let us denote $q^\eta := p_{s|u}^{\text{sg}(\theta)}(\cdot|\mathbf{x}_u; \lambda)$ and $q := \mathbb{E}_\eta[q^\eta]$ with $\eta = (\mathbf{x}_u, \lambda)$. To construct an efficient control variate given q , we need a function g such that (i) it reasonably approximates $p_{s|t}^\theta(\cdot|\mathbf{x}_t)$ and (ii) $\mathbb{E}_{\mathbf{x} \sim q}[g(\mathbf{x})]$ is easy to compute/approximate. One such example is the product model defined as

$$\bar{p}_{s|t}^\theta(\cdot|\mathbf{x}_t) := \prod_{d=1}^D \bar{p}_{s|t}^{\theta,d}(\cdot|\mathbf{x}_t), \quad \bar{p}_{s|t}^{\theta,d}(\cdot|\mathbf{x}_t) := p_{s|t}^{\theta,d}(\cdot|\mathbf{x}_t) = \mathbb{E}_\lambda \left[p_{s|t}^{\theta,d}(\cdot|\mathbf{x}_t; \lambda) \right]. \quad (12)$$

We defer the explanation of how (i) and (ii) are satisfied to Section E.1. Given a control variate $\bar{p}_{s|t}^\theta(\cdot|\mathbf{x}_t)$, we can decompose the loss computation:

$$H(q, p_{s|t}^\theta(\cdot|\mathbf{x}_t)) = \underbrace{\mathbb{E}_{\mathbf{x}_s \sim q} \left[-\log p_{s|t}^\theta(\mathbf{x}_s|\mathbf{x}_t) + \log \bar{p}_{s|t}^\theta(\mathbf{x}_s|\mathbf{x}_t) \right]}_{\text{Monte Carlo by sampling } \mathbf{x}_s} + \underbrace{\mathbb{E}_\eta \left[H(q^\eta, \bar{p}_{s|t}^\theta(\cdot|\mathbf{x}_t)) \right]}_{\text{Monte Carlo by sampling } \eta}. \quad (13)$$

Here, the first term can be treated similarly to (10), and we approximately compute the second term by sampling η and using the identity $H(q^\eta, \bar{p}_{s|t}^\theta(\cdot|\mathbf{x}_t)) = \sum_{d=1}^D H(q^{\eta,d}, \bar{p}_{s|t}^{\theta,d}(\cdot|\mathbf{x}_t))$ (see (64) in Section E.1). In this decomposition, we expect the mixture model to explicitly learn the dimensional correlation with the first term, while the second term stabilizes the overall approximation, as we use more detailed information on q than just its samples. See also Section E.2 for more background on how we derive \bar{p}^θ and another possible choice of control variate.

A.3. Auxiliary losses

While we can use a similar Monte Carlo estimate for $\mathcal{L}_{\text{data}}$ (with random samples of $\mathbf{x}_0, \mathbf{x}_t, \lambda$), we can regard $\mathcal{L}_{\text{marginal}}$ as a possible control variate for it. Indeed, if the teacher network is well-trained, we can expect that its marginal approximates the true marginal as $p^{\psi,d} \approx q^d$. Thus, for the marginal-matching product model \bar{p}^θ given in Eq. (12), we have

$$\mathbb{E}_{\mathbf{x}_t \sim q_t} \left[H(q_{0|t}(\cdot|\mathbf{x}_t), \bar{p}_{0|t}^\theta(\cdot|\mathbf{x}_t)) \right] \approx \mathcal{L}_{\text{marginal}}(\theta; \psi, q_t, t) + \text{const.}, \quad (14)$$

where the constant term is independent of θ . We give the derivation of (14) in Appendix E.3. We then obtain a decomposed formulation of $\mathcal{L}_{\text{data}}$ for given $\mathbf{x}_t \sim q_t$ as follows, by letting $q = q_{0|t}(\cdot|\mathbf{x}_t)$ and $s = 0$ in Eq. (13) and then using approximation (14):

$$\begin{aligned} \mathcal{L}_{\text{data}}(\theta; t) &\approx \mathcal{L}_{\text{corr}}(\theta; t) + \mathcal{L}_{\text{marginal}}(\theta; \psi, q_t, t) + \text{const.}, \\ \mathcal{L}_{\text{corr}}(\theta; t) &:= \mathbb{E}_{(\mathbf{x}_0, \mathbf{x}_t) \sim q_{0,t}} \left[-\log p_{0|t}^\theta(\mathbf{x}_0|\mathbf{x}_t) + \log \bar{p}_{0|t}^\theta(\mathbf{x}_0|\mathbf{x}_t) \right]. \end{aligned}$$

Here, $\mathcal{L}_{\text{corr}}$ measures the difference between p^θ and \bar{p}^θ and thus represents the dimensional correlation learned by the model p^θ . In the actual implementation for the first term $\mathcal{L}_{\text{corr}}$, we generate $\mathbf{x}_0 \sim q_0$ and then $\mathbf{x}_t \sim q_{t|0}(\cdot|\mathbf{x}_0)$, and regard them

as samples from $(x_0, x_t) \sim q_{0,t}$, which are required for conducting Monte Carlo. When combining $\mathcal{L}_{\text{data}}$ and $\mathcal{L}_{\text{marginal}}$ (both as a loss and control variate), we empirically find that mixing them as $\alpha_t \mathcal{L}_{\text{corr}}(\theta; t) + \mathcal{L}_{\text{marginal}}(\theta; \psi, q_t, t)$ with some $\alpha_t \in [0, 1]$ depending on t is more efficient than just using constant $\alpha_t = 0$ (pure marginal loss) or $\alpha_t = 1$ (pure data loss). See Section F.1 for details in this regard.

A.4. Iterated Di4C training

While we generally assume that the teacher model is given by a product model, the Di4C loss functions can also treat mixture teacher models. The only exception is the marginal loss (Section A.3), but we can just replace $q_{0|t}$ with $\bar{q}_{0|t}$ in (14) and conduct Monte Carlo estimates. Through this generalization, we can run multiple rounds of Di4C in a similar spirit as the multi-round SDTT (Deschenaux & Gulcehre, 2025).

B. Kullback–Leibler divergence and total variation distance

Let p and q be probability distributions on the same finite set \mathcal{X} . The KL divergence D_{KL} and the total variation distance d_{TV} are defined as follows:

$$D_{\text{KL}}(p \parallel q) := \sum_{x \in \mathcal{X}} p(x) \log \frac{p(x)}{q(x)}, \quad d_{\text{TV}}(p, q) := \sup_{A \subset \mathcal{X}} |p(A) - q(A)| = \frac{1}{2} \sum_{x \in \mathcal{X}} |p(x) - q(x)|.$$

Here, in the computation of D_{KL} , we ignore the term with $p(x) = 0$ and, if there is an x with $p(x) > 0$ and $q(x) = 0$, we then define $D_{\text{KL}}(p \parallel q) = 0$. These two error criteria between distributions are bridged by the following inequality (see, e.g., Canonne (2022)).

Proposition 2 (Pinsker’s inequality). *For probability distributions p and q on \mathcal{X} , we have*

$$d_{\text{TV}}(p, q) \leq \sqrt{\frac{1}{2} D_{\text{KL}}(p \parallel q)}.$$

The convexity of KL divergence in the following plays a role in the main body of the paper.

Proposition 3 (Cover & Thomas, 2006, Theorem 2.7.2). *$D_{\text{KL}}(p \parallel q)$ is convex with respect to the pair (p, q) . Namely, for $t \in [0, 1]$ and probability distributions p_1, p_2, q_1, q_2 on the same domain, we have*

$$D_{\text{KL}}(tp_1 + (1-t)p_2 \parallel tq_1 + (1-t)q_2) \leq tD_{\text{KL}}(p_1 \parallel q_1) + (1-t)D_{\text{KL}}(p_2 \parallel q_2).$$

We also use the following triangle-like inequality for the total variation distance of compositions.

Proposition 4. *For probability distributions $p_1(\cdot|y), p_2(\cdot|y)$ over \mathcal{X} conditioned on $y \in \mathcal{Y}$ and q_1, q_2 over \mathcal{Y} , we have*

$$d_{\text{TV}}(\mathbb{E}_{y \sim q_1} [p_1(\cdot|y)], \mathbb{E}_{y \sim q_2} [p_2(\cdot|y)]) \leq \mathbb{E}_{y \sim q_1} [d_{\text{TV}}(p_1(\cdot|y), p_2(\cdot|y))] + d_{\text{TV}}(q_1, q_2).$$

We give its proof in Section D.2.

C. Continuous-time Markov chains and Kolmogorov equations

Let us discuss the Kolmogorov forward/backward equations associated with continuous-time Markov chains. While the arguments below are mostly a reorganization of those given in previous studies (Campbell et al., 2022; Sun et al., 2023), we explicitly track the continuity/nonzero assumptions used in their derivations.

C.1. Kolmogorov equations in the general case

Let us consider a general Markov process over the continuous time interval $[0, T]$ and a discrete (finite) state space \mathcal{X} , which is called a continuous-time Markov chain (Anderson, 2012; Campbell et al., 2022). The starting block is the forward transition rate in a short-time interval. For $t < t + \epsilon$, assume the following equation for the infinitesimal forward transition:

$$q_{t+\epsilon|t}(y|x) = \delta_{y,x} + \epsilon Q_t(y, x) + o(\epsilon), \quad \epsilon > 0, \quad (15)$$

where $\delta_{y,x}$ is the Kronecker delta and Q_t is a function $\mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ called the transition rate. Here, for $s \leq t < t + \epsilon$, we have

$$\begin{aligned} q_{t+\epsilon|s}(y|x) &= \sum_z q_{t+\epsilon|t}(y|z)q_{t|s}(z|x) = \sum_z (\delta_{y,z} + Q_t(y, z)\epsilon)q_{t|s}(z|x) + o(\epsilon) \\ &= q_{t|s}(y|x) + \epsilon \sum_z Q_t(y, z)q_{t|s}(z|x) + o(\epsilon). \end{aligned}$$

This means that we have $\partial_t^+ q_{t|s}(y|x) = \sum_z Q_t(y, z)q_{t|s}(z|x)$, where ∂_t^+ is a right-derivative regarding t . Under the condition that Q_t is continuous over $[0, T]$ (assume it is continuously extended to $t = T$, though it is not necessary right now) and $q_{t|s}$ is continuous over $t \in [s, T]$, $q_{t|s}$ becomes differentiable over the open interval (from a general fact in analysis (von Eitzen, 2014)) and we have the Kolmogorov forward equation for $t \in (s, T)$:

$$\partial_t q_{t|s}(y|x) = \sum_z Q_t(y, z)q_{t|s}(z|x). \quad (16)$$

Now, let us derive the backward equation. For $s < s + \epsilon \leq t$, by using (15), we have

$$\begin{aligned} q_{t|s}(y|x) &= \sum_z q_{t|s+\epsilon}(y|z)q_{s+\epsilon|s}(z|x) = \sum_z q_{t|s+\epsilon}(y|z)(\delta_{z,x} + \epsilon Q_s(z, x)) + o(\epsilon) \\ &= q_{t|s+\epsilon}(y|x) + \epsilon \sum_z q_{t|s+\epsilon}(y|z)Q_s(z, x) + o(\epsilon). \end{aligned}$$

Thus, by additionally assuming the continuity of $q_{t|s}$ for $s \in [0, T]$, we obtain the one-sided derivative $\partial_s^+ q_{t|s}(y|x) = -\sum_z q_{t|s}(y|z)Q_s(z, x)$. When combined with the continuity of Q_s similarly to the above argument on the forward equation, it leads to the backward Kolmogorov equation for $s \in (0, t)$:

$$\partial_s q_{t|s}(y|x) = -\sum_z q_{t|s}(y|z)Q_s(z, x). \quad (17)$$

To summarize so far, under the assumption that $q_{t|s}$ is continuous for s, t with $0 \leq s \leq t \leq T$ and Q_t in (15) is continuous over $[0, T]$, we have the two Kolmogorov equations given by (16) and (17). Note that all the \sum_z are finite sums because of the finiteness of \mathcal{X} .

C.2. Kolmogorov equations for factorized forward processes

Let us now consider the case where $\mathcal{X} = \mathcal{S}^D$ and $\mathbf{x}_t = (x_t^d)_{d=1}^D$ follows a dimensionally independent forward process with transition rate Q_t^d . Namely, suppose

$$q_{t+\epsilon|t}^d(y^d|x^d) = \delta_{y^d, x^d} + \epsilon Q_t^d(y^d, x^d) + o(\epsilon) \quad (18)$$

for each $d = 1, \dots, D$ and $t < t + \epsilon$. In this case, we have

$$q_{t+\epsilon|t}(\mathbf{y}|\mathbf{x}) = \prod_{d=1}^D q_{t+\epsilon|t}^d(y^d|x^d) = \delta_{\mathbf{y}, \mathbf{x}} + \epsilon \sum_{d=1}^D Q_t^d(y^d, x^d) \delta_{\mathbf{y} \setminus d, \mathbf{x} \setminus d} + o(\epsilon) \quad (19)$$

by simply expanding the product, where $\mathbf{x} \setminus d \in \mathcal{S}^{D-1}$ is given by omitting the d -th entry of \mathbf{x} . From (19), the transition rate for \mathbf{x}_t is given by

$$Q_t(\mathbf{y}, \mathbf{x}) = \sum_{d=1}^D Q_t^d(y^d, x^d) \delta_{\mathbf{y} \setminus d, \mathbf{x} \setminus d} \quad (20)$$

as in Campbell et al. (2022, Proposition 3). Let us assume continuity regarding the forward process in each dimension:

Assumption A. For each $d = 1, \dots, D$, there exists a function $Q_t^d : \mathcal{S} \times \mathcal{S} \rightarrow \mathbb{R}$ indexed by $t \in [0, T]$ satisfying (18). Moreover, for any fixed $x, y \in \mathcal{S}$, $q_{t|s}^d(y|x)$ is continuous in $\{(s, t) \in [0, T]^2 \mid s \leq t\}$ and $Q_t^d(y, x)$ is continuous in $[0, T]$.

This can be satisfied by broad range of forward diffusion designs, including uniform, absorbing (masked), and discretized Gaussian diffusion (see, [Campbell et al., 2022](#), Section E).

Under Assumption A, $q_{t|s}$ and Q_t for the original process \mathbf{x}_t are also continuous since $q_{t|s}(\mathbf{y}|\mathbf{x}) = \prod_{d=1}^D q_{t|s}^d(y^d|x^d)$ and (20). Thus, we can apply the argument in Section C.1 to obtain Kolmogorov equations (16) & (17). Moreover, we can show a favorable property of the time-reversal process. This is just a re-formalization of a well-known fact (e.g., [Campbell et al. \(2022, Proposition 3\)](#) and [Sun et al. \(2023, Proposition 3.2\)](#)).

Proposition 5. *Let $\mathcal{S}_{t,+}^D := \{\mathbf{x} \in \mathcal{S}^D \mid q_t(\mathbf{x}) > 0\}$. Under Assumption A, there exists a function $R_t : \mathcal{S}^D \times \mathcal{S}_{t,+}^D \rightarrow \mathbb{R}$ indexed by $t \in (0, T]$ such that*

- (a) *we have $q_{t-\epsilon|t}(\mathbf{y}|\mathbf{x}) = \delta_{\mathbf{y},\mathbf{x}} + \epsilon R_t(\mathbf{y}, \mathbf{x}) + o(\epsilon)$ for $\mathbf{y} \in \mathcal{S}^D$, $\mathbf{x} \in \mathcal{S}_{t,+}^D$ and $\epsilon > 0$ with $t - \epsilon \geq 0$, and*
- (b) *$R_t(\mathbf{y}, \mathbf{x})$ can be nonzero only if \mathbf{x} and \mathbf{y} coincide in at least $D - 1$ entries.*

We give its proof in Section D.3. As one can see from the proof, the time-reversal transition rate R_t is given concretely by $R_t(\mathbf{y}, \mathbf{x}) = Q_t(\mathbf{x}, \mathbf{y})q_t(\mathbf{y})/q_t(\mathbf{x})$ when $\mathbf{x} \neq \mathbf{y}$ and $q_t(\mathbf{x}) > 0$. Note that the ratio $q_t(\mathbf{y})/q_t(\mathbf{x})$ is treated as a discrete counterpart of the score function ([Sun et al., 2023](#); [Lou et al., 2024](#)).

Let us add another regularity assumption:

Assumption B. *For each $d = 1, \dots, D$ and $x, y \in \mathcal{S}$, $Q_t^d(y, x)$ is differentiable for $t \in (0, T)$ and the derivative $\partial_t Q_t^d(y, x)$ can be continuously extended to $[0, T]$.*

Note that usual choices of Q_t^d regarding t including the time-homogeneous case $Q_t = Q$ and the noise scheduling $Q_t = \beta(t)Q$ with a smooth β ([Campbell et al., 2022](#); [Lou et al., 2024](#)) satisfy this assumption. Finally, under these two assumptions, we can formalize Theorem 1 as follows.

Theorem 3. *Suppose $(\mathbf{x}_t)_{0 \leq t \leq T}$ satisfies Assumptions A & B. Let $p_{s|t}$ be a product model with the correct marginals, i.e., $p_{s|t}(\mathbf{x}_s|\mathbf{x}_t) = \prod_{d=1}^D q_{s|t}^d(x_s^d|x_t^d)$ for $s < t$. Then, there exists a constant $C > 0$ such that, given timesteps $t_i = iT/N$ for $i = 0, \dots, N$, we have*

$$d_{\text{TV}}(q_0, \mathbb{E}_{\mathbf{x}_T \sim q_T} [p_{t_0|t_1} \circ p_{t_1|t_2} \circ \dots \circ p_{t_{N-1}|t_N}(\cdot|\mathbf{x}_T)]) \leq \frac{C}{N}. \quad (21)$$

Furthermore, there exists an example of $(\mathbf{x}_t)_{0 \leq t \leq T}$ satisfying $D = |\mathcal{S}| = 2$ and the same assumptions such that the left-hand side of (21) is lower-bounded by c/N with some constant $c > 0$ for sufficiently large N .

This theorem basically says the min-max convergence rate of the analytical sampling is $1/N$. We give the proof of the first half, i.e., Eq. (21), in Section D.4. For the latter half, we provide the concrete version in Proposition 6 in the following section.

C.3. Lower bound of Theorem 3

We shall provide an example that yields an $\Omega(1/N)$ error between the analytical and true denoisers. Our example even satisfies the following stronger assumption, which is often used in theoretical analysis (e.g., [Campbell et al., 2022](#), Assumption 1):

Assumption C. *For any $t \in [0, T]$ and $\mathbf{x} \in \mathcal{S}^D$, $q_t(\mathbf{x}) > 0$ holds.*

Consider $\mathcal{S} = \{a, b\}$ and $D = 2$, where the state-space is given by $\mathcal{X} = \{aa, ab, ba, bb\}$ by omitting parentheses. Consider the (forward) Markov process given by the initial distribution $q_0 = (\delta_{aa} + \delta_{bb})/2$ and the dimension-wise time-homogeneous transition rate $Q_t^d(y, x) = 1/2 - \delta_{yx}$ for $d = 1, 2$ and $x, y \in \mathcal{S}$. Under this setting, the forward transition probability is continuous and satisfies $\partial_t q_{t|s}^d(\cdot|a) = Q_t^d q_{t|s}^d(\cdot|a)$ as a vector-valued differential equation, so we have, for $t > s$,

$$\partial_t q_{t|s}^d(a|a) = -\frac{1}{2}q_{t|s}^d(a|a) + \frac{1}{2}q_{t|s}^d(b|a) = \frac{1}{2} - q_{t|s}^d(a|a).$$

By solving this, we obtain $q_{t|s}^d(a|a) = \frac{1}{2}(1 + e^{-(t-s)})$ for $t \geq s$. By symmetry, we generally have

$$q_{t|s}^d(a|a) = q_{t|s}^d(b|b) = \frac{1}{2}(1 + e^{-(t-s)}), \quad q_{t|s}^d(b|a) = q_{t|s}^d(a|b) = \frac{1}{2}(1 - e^{-(t-s)}) \quad (22)$$

This is a special case of uniform diffusion and clearly satisfies Assumptions A & B. Although the singularity of q_0 violates Assumption C at time zero, we can consider the time interval $[\delta, T]$ for some $\delta > 0$ instead of $[0, T]$ to ensure $q_t > 0$. We will, however, work with the singular q_0 for simplicity of computations. The following proposition gives the lower bound discussed in Theorem 3. If necessary, we can replace T with $T + \delta$ and consider $\mathbf{x}'_t = \mathbf{x}_{t+\delta}$ to match the time intervals.

Proposition 6. *Let $(\mathbf{x}_t)_{\delta \leq t \leq T}$ be the Markov process defined above and $p_{s|t}$ be the product model $p_{s|t}(\mathbf{x}_s|\mathbf{x}_t) = \prod_{d=1}^D q_{s|t}^d(x_s^d|\mathbf{x}_t)$ for $s < t$. If we let $N \geq 2(T - \delta)/\delta$ be an integer and $t_i = \delta + i(T - \delta)/N$ for $i = 0, \dots, N$ be timesteps, then there is a constant $c > 0$ such that*

$$d_{\text{TV}}(q_\delta, \mathbb{E}_{\mathbf{x}_T \sim q_T} [p_{t_0|t_1} \circ p_{t_1|t_2} \circ \dots \circ p_{t_{N-1}|t_N}(\cdot|\mathbf{x}_T)]) \geq \frac{c}{N}. \quad (23)$$

The proof is given in Section D.5.

D. Proofs

D.1. Proof of Theorem 2

Proof. For simplicity of notation, let $\tilde{p}_{t_n|T}^\psi$ be the denoiser given by the teacher with timesteps $t_n < t_{n+1} < \dots < t_N$, i.e.,

$$\tilde{p}_{t_n|T}^\psi := p_{t_n|t_{n+1}}^\psi \circ \dots \circ p_{t_{N-1}|t_N}^\psi,$$

so that we have $r_{t_n} = \mathbb{E}_{\mathbf{x}_T \sim r_T} [\tilde{p}_{t_n|T}^\psi(\cdot|\mathbf{x}_T)]$. Note that we can just set $\tilde{p}_{t_N|T}^\psi(\cdot|\mathbf{x}) = \tilde{p}_{T|T}^\psi(\cdot|\mathbf{x}) = \delta_{\mathbf{x}}$.

Also, let $p_{0,n} := \mathbb{E}_{\mathbf{x}_T \sim r_T} [p_{0|t_n}^\theta \circ \tilde{p}_{t_n|T}^\psi(\cdot|\mathbf{x}_T)]$ for $n = 1, \dots, N$, where $p_{0,N}$ is just given by $p_{0,N} = \mathbb{E}_{\mathbf{x}_T \sim r_T} [p_{0|T}^\theta(\cdot|\mathbf{x}_T)]$. We first compare $p_{0,n}$ and $p_{0,n+1}$ with the consistency loss.

For each $0 < u < t \leq T$, we have

$$\begin{aligned} \mathcal{L}_{\text{consis}}(\theta; \psi, r_t, 0, u, t) &= \mathbb{E}_{\mathbf{x}_t \sim r_t} [D_{\text{KL}}(p_{0|u}^\theta \circ p_{u|t}^\psi(\cdot|\mathbf{x}_t) \| p_{0|t}^\theta(\cdot|\mathbf{x}_t))] \\ &\geq D_{\text{KL}}\left(\mathbb{E}_{\mathbf{x}_t \sim r_t} [p_{0|u}^\theta \circ p_{u|t}^\psi(\cdot|\mathbf{x}_t)] \middle\| \mathbb{E}_{\mathbf{x}_t \sim r_t} [p_{0|t}^\theta(\cdot|\mathbf{x}_t)]\right) \end{aligned}$$

from the convexity (Proposition 3). If we let $u = t_n$ and $t = t_{n+1}$ for some $1 \leq n < N$, we can see

$$\mathbb{E}_{\mathbf{x}_t \sim r_t} [p_{0|u}^\theta \circ p_{u|t}^\psi(\cdot|\mathbf{x}_t)] = \mathbb{E}_{\mathbf{x}_T} [p_{0|t_n}^\theta \circ p_{t_n|t_{n+1}}^\psi \circ \tilde{p}_{t_{n+1}|T}^\psi(\cdot|\mathbf{x}_T)] = p_{0,n},$$

and $\mathbb{E}_{\mathbf{x}_t \sim r_t} [p_{0|t}^\theta(\cdot|\mathbf{x}_t)] = p_{0,n+1}$ hold. By using Pinsker's inequality (Proposition 2), we have

$$d_{\text{TV}}(p_{0,n}, p_{0,n+1}) \leq \frac{1}{\sqrt{2}} D_{\text{KL}}(p_{0,n} \| p_{0,n+1})^{1/2} \leq \frac{1}{\sqrt{2}} \mathcal{L}_{\text{consis}}(\theta; \psi, r_{t_{n+1}}, 0, t_n, t_{n+1})^{1/2}. \quad (24)$$

From a similar argument, we have

$$\mathcal{L}_{\text{distil}}(\theta; \psi, r_{t_1}, t_1) = \mathbb{E}_{\mathbf{x}_{t_1} \sim r_{t_1}} [D_{\text{KL}}(p_{0|t_1}^\psi(\cdot|\mathbf{x}_{t_1}) \| p_{0|t_1}^\theta(\cdot|\mathbf{x}_{t_1}))] \geq D_{\text{KL}}(r_0 \| p_{0,1}),$$

and thus

$$d_{\text{TV}}(r_0, p_{0,1}) \leq \frac{1}{\sqrt{2}} D_{\text{KL}}(r_0 \| p_{0,1})^{1/2} \leq \frac{1}{\sqrt{2}} \mathcal{L}_{\text{distil}}(\theta; \psi, r_{t_1}, t_1)^{1/2}. \quad (25)$$

By using the triangle inequality of the total variation distance, we obtain

$$d_{\text{TV}}(r_0, p_{0,N}) \leq d_{\text{TV}}(r_0, p_{0,1}) + \sum_{n=1}^{N-1} d_{\text{TV}}(p_{0,n}, p_{0,n+1}).$$

Finally, applying Eqs. (24) and (25) to its right-hand side yields the desired inequality. \square

D.2. Proof of Proposition 4

Proof. Let us first consider the case of $q_1 = q_2$. Then, we have

$$\begin{aligned}
 & d_{\text{TV}}(\mathbb{E}_{y \sim q_1}[p_1(\cdot|y)], \mathbb{E}_{y \sim q_1}[p_2(\cdot|y)]) \\
 &= \frac{1}{2} \sum_x \left| \sum_y p_1(x|y)q_1(y) - \sum_y p_2(x|y)q_1(y) \right| = \frac{1}{2} \sum_x \left| \sum_y (p_1(x|y) - p_2(x|y))q_1(y) \right| \\
 &\leq \frac{1}{2} \sum_x \sum_y |p_1(x|y) - p_2(x|y)| q_1(y) = \mathbb{E}_{y \sim q_1}[d_{\text{TV}}(p_1(\cdot|y), p_2(\cdot|y))], \tag{26}
 \end{aligned}$$

where we have used $q_1 \geq 0$ in the inequality. On the other hand, if $p_1 = p_2$, we have

$$\begin{aligned}
 & d_{\text{TV}}(\mathbb{E}_{y \sim q_1}[p_2(\cdot|y)], \mathbb{E}_{y \sim q_2}[p_2(\cdot|y)]) \\
 &= \frac{1}{2} \sum_x \left| \sum_y p_2(x|y)q_1(y) - \sum_y p_2(x|y)q_2(y) \right| = \frac{1}{2} \sum_x \left| \sum_y p_2(x|y)(q_1(y) - q_2(y)) \right| \\
 &\leq \frac{1}{2} \sum_x \sum_y p_2(x|y)|q_1(y) - q_2(y)| = \frac{1}{2} \sum_y |q_1(y) - q_2(y)| = d_{\text{TV}}(q_1, q_2), \tag{27}
 \end{aligned}$$

where we have used $p_2 \geq 0$ in the inequality and $\sum_x p_2(x|y) = 1$ in the last equality.

By utilizing the usual triangle inequality of d_{TV} and inequalities (26) & (27), we obtain

$$\begin{aligned}
 & d_{\text{TV}}(\mathbb{E}_{y \sim q_1}[p_1(\cdot|y)], \mathbb{E}_{y \sim q_2}[p_2(\cdot|y)]) \\
 &\leq d_{\text{TV}}(\mathbb{E}_{y \sim q_1}[p_1(\cdot|y)], \mathbb{E}_{y \sim q_1}[p_2(\cdot|y)]) + d_{\text{TV}}(\mathbb{E}_{y \sim q_1}[p_2(\cdot|y)], \mathbb{E}_{y \sim q_2}[p_2(\cdot|y)]) \\
 &\leq \mathbb{E}_{y \sim q_1}[d_{\text{TV}}(p_1(\cdot|y), p_2(\cdot|y))] + d_{\text{TV}}(q_1, q_2),
 \end{aligned}$$

which is the desired inequality. \square

D.3. Proof of Proposition 5

Proof. Note that, by Assumption A, $q_{t|s}$ is continuous over $\{(s, t) \in [0, T]^2 \mid s \leq t\}$, and Q_t given by (20) is continuous over $[0, T]$ and satisfies Eqs. (15)–(17), as mentioned in Section C.2 just after Assumption A.

Let us simply write $x \in \mathcal{X}$ instead of the bold style $\mathbf{x} \in \mathcal{S}^D$ in this paragraph. We consider only $x \in \mathcal{X}$ such that $q_t(x) > 0$. We follow the argument in Sun et al. (2023, Section B.2). Let us consider the conditional probability (namely, the true denoiser) $q_{s|t}(\cdot|x)$ for $s \leq t$, which is uniquely determined since $q_t(x) > 0$. Then, we have

$$\begin{aligned}
 \partial_s q_{s|t}(y|x) &= \partial_s \frac{q_s(y)q_{t|s}(x|y)}{q_t(x)} = \frac{(\partial_s q_s)(y)q_{t|s}(x|y) + q_s(y)(\partial_s q_{t|s})(x|y)}{q_t(x)} \\
 &= \frac{1}{q_t(x)} \left(q_{t|s}(x|y) \sum_z Q_s(y, z)q_s(z) - q_s(y) \sum_w q_{t|s}(x|w)Q_s(w, y) \right), \tag{28}
 \end{aligned}$$

where we have used the forward Kolmogorov equation of q_t given as

$$\partial_t q_t(x) = \sum_w \partial_t q_{t|0}(x|w)q_0(w) = \sum_w \sum_z Q_t(x, z)q_{t|0}(z|w)q_0(w) = \sum_z Q_t(x, z)q_t(z)$$

for computing $\partial_s q_s$ and the backward Kolmogorov equation for computing $\partial_s q_{t|s}$. By taking the limit $s \rightarrow t - 0$ in (28), we obtain $\lim_{s \rightarrow t-0} \partial_s q_{s|t}(y|x) = -\frac{q_t(y)}{q_t(x)}Q_t(x, y)$ if $y \neq x$, given the continuity of $q_{t|s}$ and Q_s . Then, from Taylor's theorem, we obtain a backward counterpart of (15) for $y \neq x$ as

$$q_{t-\epsilon|t}(y|x) = \epsilon \frac{q_t(y)}{q_t(x)}Q_t(x, y) + o(\epsilon), \quad \epsilon > 0. \tag{29}$$

Since $\sum_y q_{t-\epsilon|t}(y|x) = 1$ holds always, we also have that $q_{t-\epsilon|t}(x|x) = 1 + \epsilon R_{t,x} + o(\epsilon)$ for the coefficient $R_{t,x} = -\sum_{y \neq x} \frac{q_t(y)}{q_t(x)}Q_t(x, y)$. Therefore, we can prove (a) by letting $R_t(y, x) = \frac{q_t(y)}{q_t(x)}Q_t(x, y)$ for $y \neq x$ and $R_t(x, x) = R_{t,x}$.

We can see (b) from (20) and the concrete form of R_t . \square

D.4. Proof of first half of Theorem 3

We first prove the following auxiliary lemma replacing the $o(\epsilon)$ term in the backward transition by $O(\epsilon^2)$.

Lemma 1. *Under the same setting as in Theorem 3, there is a constant $C > 0$ such that, for any $t \in (0, T]$, $\epsilon \in (0, t]$, and $\mathbf{x} \in \mathcal{X}$ with $q_t(\mathbf{x}) > 0$, we have*

$$d_{\text{TV}}(q_{t-\epsilon|\cdot|\mathbf{x}}, p_{t-\epsilon|\cdot|\mathbf{x}}) \leq \frac{C\epsilon^2}{q_t(\mathbf{x})}. \quad (30)$$

Proof. From (28) and Assumption B, $q_{s|t}(\mathbf{y}|\mathbf{x})$ for $s < t$ is twice-differentiable with regard to s , and $q_t(\mathbf{x})\partial_s q_{s|t}(\mathbf{y}|\mathbf{x})$ can be represented as a polynomial of the function values of q_s , Q_s , $q_{t|s}$, and $\partial_s Q_s$. Thus, there is a constant C_1 depending on $|\mathcal{S}|$, D , $\sup_{s,z,w} Q_s(z, \mathbf{w})$ and $\sup_{s,z,w} \partial_s(z, \mathbf{w})$ such that $q_t(\mathbf{x})\partial_s^2 q_{s|t}(\mathbf{y}|\mathbf{x}) \leq C_1$ for any $s, t, \mathbf{y}, \mathbf{x}$ (note that $q_{t|s}$ and q_s are within $[0, 1]$).

Now that $\partial_s q_{s|t}$ can be continuously extended to $s \in [0, t]$ from (28), for each $t \in (0, T]$, $\epsilon \in (0, t]$ and $\mathbf{x}, \mathbf{y} \in \mathcal{S}^D$ with $q_t(\mathbf{x}) > 0$, Taylor's theorem yields that

$$|q_{t-\epsilon|t}(\mathbf{y}|\mathbf{x}) - \delta_{\mathbf{y},\mathbf{x}} - \epsilon R_t(\mathbf{y}, \mathbf{x})| = \left| \frac{(\partial_s^2 q_{s|t})(\mathbf{y}|\mathbf{x})|_{s=\theta}}{2} \epsilon^2 \right| \leq \frac{C_1}{2q_t(\mathbf{x})} \epsilon^2, \quad (31)$$

for a certain $\theta \in (t - \epsilon, t)$.

Let us next consider the marginal-matching product model $p_{t-\epsilon|t}$. For each d , if $y^d \neq x^d$, we have

$$\begin{aligned} \left| p_{t-\epsilon|t}^d(y^d|\mathbf{x}) - \epsilon R_t((y^d, \mathbf{x}^{\setminus d}), \mathbf{x}) \right| &= \left| \sum_{\mathbf{y}^{\setminus d} \in \mathcal{S}^{D-1}} q_{t-\epsilon|t}((y^d, \mathbf{y}^{\setminus d})|\mathbf{x}) - \epsilon R_t((y^d, \mathbf{x}^{\setminus d}), \mathbf{x}) \right| \\ &= \left| \sum_{\mathbf{y}^{\setminus d}} \left(q_{t-\epsilon|t}((y^d, \mathbf{y}^{\setminus d})|\mathbf{x}) - \epsilon R_t((y^d, \mathbf{y}^{\setminus d}), \mathbf{x}) \right) \right| \\ &\leq \frac{|\mathcal{S}|^{D-1} C_1}{2q_t(\mathbf{x})} \epsilon^2, \end{aligned} \quad (32)$$

where the second equality comes from Proposition 5(b) and the inequality is from (31). If $y^d = x^d$, since $p_{t-\epsilon|t}^d(x^d|\mathbf{x}) = 1 - \sum_{y^d \neq x^d} p_{t-\epsilon|t}^d(y^d|\mathbf{x})$ we can use (32) to obtain

$$\begin{aligned} \left| p_{t-\epsilon}^d(x^d|\mathbf{x}) - 1 + \epsilon \sum_{y^d \neq x^d} R_t((y^d, \mathbf{x}^{\setminus d}), \mathbf{x}) \right| &\leq \sum_{y^d \neq x^d} |p_{t-\epsilon|t}^d(y^d|\mathbf{x}) - \epsilon R_t((y^d, \mathbf{x}^{\setminus d}), \mathbf{x})| \\ &\leq \frac{|\mathcal{S}|^D C_1}{2q_t(\mathbf{x})} \epsilon^2. \end{aligned}$$

From (32) and this, by defining $R_t^d : \mathcal{S} \rightarrow \mathbb{R}$ as $R_t^d(y^d) = R_t((y^d, \mathbf{x}^{\setminus d}), \mathbf{x})$ for $y^d \neq x^d$ and $R_t^d(x^d) = -\sum_{y^d \neq x^d} R_t^d(y^d)$, there exists a constant $C_2 > 0$ and a function $A^d : \mathcal{S} \rightarrow \mathbb{R}$ (for fixed t and \mathbf{x}) such that

$$p_{t-\epsilon|t}^d(y^d|\mathbf{x}) = \delta_{y^d, x^d} - \epsilon R_t^d(y^d) + \frac{\epsilon^2}{q_t(\mathbf{x})} A^d(y^d, \epsilon), \quad \sup_{y^d \in \mathcal{S}, \epsilon} |A^d(y^d, \epsilon)| \leq C_2. \quad (33)$$

Therefore, we have

$$\begin{aligned} p_{t-\epsilon|t}(\mathbf{y}|\mathbf{x}) &= \prod_{d=1}^D \left(\delta_{y^d, x^d} - \epsilon R_t^d(y^d) + \frac{\epsilon^2}{q_t(\mathbf{x})} A^d(y^d, \epsilon) \right) \\ &= \delta_{\mathbf{y}, \mathbf{x}} + \epsilon \sum_{d=1}^D R_t^d(y^d) \delta_{\mathbf{y}^{\setminus d}, \mathbf{x}^{\setminus d}} + \underbrace{\sum_{k=1}^D P_k((\delta_{y^d, x^d}, R_t^d(y^d), A^d(y^d, \epsilon))_{d=1}^D)}_{\text{Remainder term}} \left(\frac{\epsilon^2}{q_t(\mathbf{x})} \right)^k, \end{aligned}$$

where P_k is a certain polynomial of $3D$ variables for each k . Note that, if $\mathbf{y} \neq \mathbf{x}$, $R_t^d(y^d)\delta_{\mathbf{y}^{\setminus d}, \mathbf{x}^{\setminus d}}$ can be nonzero only if $y^d \neq x^d$ and $\mathbf{y}^{\setminus d} = \mathbf{x}^{\setminus d}$. In that case, from the definition of $R_t^d(y^d)$, we have

$$p_{t-\epsilon|t}(\mathbf{y}|\mathbf{x}) = \epsilon R_t^d(y^d) + (\text{Remainder term}) = \epsilon R_t(\mathbf{y}, \mathbf{x}) + (\text{Remainder term}). \quad (34)$$

This equality also holds when \mathbf{y} and \mathbf{x} differ in more than one entry, since the coefficient of ϵ becomes zero in such a case, and $R_t(\mathbf{y}, \mathbf{x}) = 0$ from Proposition 5(b). Since the inputs for each P_k are all bounded, we have

$$(\text{Remainder term}) \leq C_3 \sum_{k=1}^D \left(\frac{\epsilon^2}{q_t(\mathbf{x})} \right)^k \leq C_3 D \left(\frac{\epsilon^2}{q_t(\mathbf{x})} + \frac{\epsilon^{2D}}{q_t(\mathbf{x})^D} \right) \quad (35)$$

for a constant $C_3 > 0$. By combining it with (31), for $\mathbf{y} \neq \mathbf{x}$, we have

$$\begin{aligned} |q_{t-\epsilon|t}(\mathbf{y}|\mathbf{x}) - p_{t-\epsilon|t}(\mathbf{y}|\mathbf{x})| &\leq |q_{t-\epsilon|t}(\mathbf{y}|\mathbf{x}) - \epsilon R_t(\mathbf{y}, \mathbf{x})| + |\epsilon R_t(\mathbf{y}, \mathbf{x}) - p_{t-\epsilon|t}(\mathbf{y}|\mathbf{x})| \\ &\leq \frac{C_1}{2q_t(\mathbf{x})} \epsilon^2 + C_3 D \left(\frac{\epsilon^2}{q_t(\mathbf{x})} + \frac{\epsilon^{2D}}{q_t(\mathbf{x})^D} \right) \leq C_4 \left(\frac{\epsilon^2}{q_t(\mathbf{x})} + \frac{\epsilon^{2D}}{q_t(\mathbf{x})^D} \right) \end{aligned}$$

for a constant $C_4 > 0$. In particular, we have

$$\begin{aligned} d_{\text{TV}}(q_{t-\epsilon|t}(\cdot|\mathbf{x}), p_{t-\epsilon|t}(\cdot|\mathbf{x})) &= \frac{1}{2} \sum_{\mathbf{y}} |q_{t-\epsilon|t}(\mathbf{y}|\mathbf{x}) - p_{t-\epsilon|t}(\mathbf{y}|\mathbf{x})| \\ &= \frac{1}{2} \left(\sum_{\mathbf{y} \neq \mathbf{x}} |q_{t-\epsilon|t}(\mathbf{y}|\mathbf{x}) - p_{t-\epsilon|t}(\mathbf{y}|\mathbf{x})| + \left| 1 - \sum_{\mathbf{y} \neq \mathbf{x}} q_{t-\epsilon|t}(\mathbf{y}|\mathbf{x}) - 1 + \sum_{\mathbf{y} \neq \mathbf{x}} p_{t-\epsilon|t}(\mathbf{y}|\mathbf{x}) \right| \right) \\ &\leq \sum_{\mathbf{y} \neq \mathbf{x}} |q_{t-\epsilon|t}(\mathbf{y}|\mathbf{x}) - p_{t-\epsilon|t}(\mathbf{y}|\mathbf{x})| \leq |\mathcal{S}|^D C_4 \left(\frac{\epsilon^2}{q_t(\mathbf{x})} + \frac{\epsilon^{2D}}{q_t(\mathbf{x})^D} \right) = C_5 \left(\frac{\epsilon^2}{q_t(\mathbf{x})} + \frac{\epsilon^{2D}}{q_t(\mathbf{x})^D} \right), \end{aligned}$$

for a constant $C_5 > 0$. Now, we can assume that $C_5 \geq 1/2$, by adding a positive number if necessary. Since d_{TV} is bounded above by 1 in general, we consider two cases:

- (a) If $\frac{\epsilon^2}{q_t(\mathbf{x})} \geq 1$, we have $d_{\text{TV}}(q_{t-\epsilon|t}(\cdot|\mathbf{x}), p_{t-\epsilon|t}(\cdot|\mathbf{x})) \leq 1 \leq 2C_5 \frac{\epsilon^2}{q_t(\mathbf{x})}$ since $2C_5 \geq 1$.
- (b) If $\frac{\epsilon^2}{q_t(\mathbf{x})} < 1$, we have $d_{\text{TV}}(q_{t-\epsilon|t}(\cdot|\mathbf{x}), p_{t-\epsilon|t}(\cdot|\mathbf{x})) \leq C_5 \left(\frac{\epsilon^2}{q_t(\mathbf{x})} + \frac{\epsilon^{2D}}{q_t(\mathbf{x})^D} \right) \leq 2C_5 \frac{\epsilon^2}{q_t(\mathbf{x})}$ since $\frac{\epsilon^2}{q_t(\mathbf{x})} \geq \frac{\epsilon^{2D}}{q_t(\mathbf{x})^D}$.

Therefore, we finally obtain (30). \square

By using the lemma and Proposition 4, we can prove the theorem.

Proof of Theorem 3. For each $i = 0, \dots, N$, let us define the compositions

$$\tilde{p}_{0|t_0}(\cdot|\mathbf{x}) = \delta_{\mathbf{x}}, \quad \tilde{p}_{0|t_i} := p_{t_0|t_1} \circ \dots \circ p_{t_{i-1}|t_i}, \quad i = 1, \dots, N.$$

Note also that, for \mathbf{x} with $q_T(\mathbf{x}) > 0$, we have $q_{t_i|T}(\cdot|\mathbf{x}) = q_{t_i|t_{i+1}} \circ \dots \circ q_{t_{N-1}|t_N}(\cdot|\mathbf{x})$ from the Markov property of the reverse process. Indeed, for $s < t < u$, we have $q_{u|t}(\mathbf{z}|\mathbf{y}) = q_{u|s,t}(\mathbf{z}|\mathbf{x}, \mathbf{y})$ from the Markov property of the forward process, so, for \mathbf{z} with $q_u(\mathbf{z}) > 0$,

$$\begin{aligned} \sum_{\mathbf{y}} q_{s|t}(\mathbf{x}|\mathbf{y}) q_{t|u}(\mathbf{y}|\mathbf{z}) &= \sum_{\mathbf{y}} \frac{q_{s,t}(\mathbf{x}, \mathbf{y})}{q_t(\mathbf{y})} \frac{q_{t,u}(\mathbf{y}, \mathbf{z})}{q_u(\mathbf{z})} \\ &= \sum_{\mathbf{y}} \frac{q_{s,t}(\mathbf{x}, \mathbf{y}) q_{u|t}(\mathbf{z}|\mathbf{y})}{q_u(\mathbf{z})} = \sum_{\mathbf{y}} \frac{q_{s,t}(\mathbf{x}, \mathbf{y}) q_{u|s,t}(\mathbf{z}|\mathbf{x}, \mathbf{y})}{q_u(\mathbf{z})} \\ &= \frac{\sum_{\mathbf{y}} q_{s,t,u}(\mathbf{x}, \mathbf{y}, \mathbf{z})}{q_u(\mathbf{z})} = \frac{q_{s,u}(\mathbf{x}, \mathbf{z})}{q_u(\mathbf{z})} = q_{s|u}(\mathbf{x}|\mathbf{z}), \end{aligned}$$

where we have implicitly used that $q_t(\mathbf{y}) > 0$ holds for \mathbf{y} satisfying $q_{t|u}(\mathbf{y}|\mathbf{z}) > 0$ (given $q_u(\mathbf{z}) > 0$). By using the inequality recursively, we can prove the aforementioned identity.

We prove the desired estimate by exploiting the compositions. Recall $q_0 = \mathbb{E}_{\mathbf{x}_T \sim q_T} [q_{0|t_N}(\cdot|\mathbf{x}_T)]$. What we want to estimate is $d_{\text{TV}}(\mathbb{E}_{\mathbf{x}_T \sim q_T} [q_{0|t_N}(\cdot|\mathbf{x}_T)], \mathbb{E}_{\mathbf{x}_T \sim q_T} [\tilde{p}_{0|t_N}(\cdot|\mathbf{x}_T)])$. We bound the distance with the following triangle inequality:

$$\begin{aligned} & d_{\text{TV}}(\mathbb{E}_{\mathbf{x}_T \sim q_T} [q_{0|t_N}(\cdot|\mathbf{x}_T)], \mathbb{E}_{\mathbf{x}_T \sim q_T} [\tilde{p}_{0|t_N}(\cdot|\mathbf{x}_T)]) \\ & \leq \sum_{i=0}^{N-1} d_{\text{TV}}(\mathbb{E}_{\mathbf{x}_T \sim q_T} [\tilde{p}_{0|t_i} \circ q_{t_i|t_N}(\cdot|\mathbf{x}_T)], \mathbb{E}_{\mathbf{x}_T \sim q_T} [\tilde{p}_{0|t_{i+1}} \circ q_{t_{i+1}|t_N}(\cdot|\mathbf{x}_T)]). \end{aligned} \quad (36)$$

Let us bound each term inside the summation by using Lemma 1 and Proposition 4. First, since $\tilde{p}_{0|t_{i+1}} = \tilde{p}_{0|t_i} \circ p_{t_i|t_{i+1}}$, by letting $p_1 = p_2 = \tilde{p}_{0|t_i}$ in Proposition 4, we have

$$\begin{aligned} & d_{\text{TV}}(\mathbb{E}_{\mathbf{x}_T \sim q_T} [\tilde{p}_{0|t_i} \circ q_{t_i|t_N}(\cdot|\mathbf{x}_T)], \mathbb{E}_{\mathbf{x}_T \sim q_T} [\tilde{p}_{0|t_{i+1}} \circ q_{t_{i+1}|t_N}(\cdot|\mathbf{x}_T)]) \\ & \leq d_{\text{TV}}(\mathbb{E}_{\mathbf{x}_T \sim q_T} [q_{t_i|t_N}(\cdot|\mathbf{x}_T)], \mathbb{E}_{\mathbf{x}_T \sim q_T} [p_{t_i|t_{i+1}} \circ q_{t_{i+1}|t_N}(\cdot|\mathbf{x}_T)]). \end{aligned} \quad (37)$$

Second, since $q_{t_{i+1}|t_N} = q_{t_i|t_{i+1}} \circ q_{t_{i+1}|t_N}$, by letting $q_1 = q_2 = q_{t_{i+1}} = \mathbb{E}_{\mathbf{x}_T \sim q_T} [q_{t_{i+1}|t_N}(\cdot|\mathbf{x}_T)]$ in Proposition 4 (note that the indices of q_1, q_2 here are different from time), we have

$$\begin{aligned} & d_{\text{TV}}(\mathbb{E}_{\mathbf{x}_T \sim q_T} [q_{t_i|t_N}(\cdot|\mathbf{x}_T)], \mathbb{E}_{\mathbf{x}_T \sim q_T} [p_{t_i|t_{i+1}} \circ q_{t_{i+1}|t_N}(\cdot|\mathbf{x}_T)]) \\ & \leq \mathbb{E}_{\mathbf{x} \sim q_{t_{i+1}}} [d_{\text{TV}}(q_{t_i|t_{i+1}}(\cdot|\mathbf{x}), p_{t_i|t_{i+1}}(\cdot|\mathbf{x}))] \\ & \leq \sum_{q_{t_{i+1}}(\mathbf{x}) > 0} q_{t_{i+1}}(\mathbf{x}) \cdot \frac{C(t_{i+1} - t_i)^2}{q_{t_{i+1}}(\mathbf{x})} \leq \frac{C|\mathcal{S}|^D T^2}{N^2}, \end{aligned} \quad (38)$$

where we have used (30) and $t_{i+1} - t_i = T/N$ in the last inequality. By combining estimates (36)–(38), we obtain

$$d_{\text{TV}}(\mathbb{E}_{\mathbf{x}_T \sim q_T} [q_{0|t_N}(\cdot|\mathbf{x}_T)], \mathbb{E}_{\mathbf{x}_T \sim q_T} [\tilde{p}_{0|t_N}(\cdot|\mathbf{x}_T)]) \leq \sum_{i=0}^{N-1} \frac{C|\mathcal{S}|^D T^2}{N^2} = \frac{C|\mathcal{S}|^D T^2}{N},$$

which completes the proof with a replacement of the constant factor. \square

D.5. Proof of Proposition 6

Proof. Consider the analytical sampler $p_{s|t}(zw|xy) = q_{s|t}^1(z|x)q_{s|t}^2(w|y)$ for $s < t$. Note that, because of the symmetry between a and b in q_0 and the forward transition, the distributions q_t or those given by the composition of $p_{s|t}$ are also symmetric. Thus, the probability of aa recovers all the information of the distributions we consider over \mathcal{X} .

Let us compute several probabilities regarding $q_{s|t}$ and the analytical sampler through (22). First, note that $q_{0|t}(ab|\cdot) = q_{0|t}(ba|\cdot) = 0$. Therefore, we have

$$\begin{aligned} q_{0|t}(aa|aa) &= \frac{q_{t|0}(aa|aa)q_0(aa)}{q_t(aa)} \\ &= \frac{q_{t|0}(aa|aa)q_0(aa)}{q_{t|0}(aa|aa)q_0(aa) + q_{t|0}(aa|bb)q_0(bb)} = \frac{\frac{1}{4}(1 + e^{-t})^2}{\frac{1}{4}(1 + e^{-t})^2 + \frac{1}{4}(1 - e^{-t})^2} = \frac{(1 + e^{-t})^2}{2(1 + e^{-2t})}, \end{aligned} \quad (39)$$

$$q_{0|t}(bb|aa) = 1 - q_{0|t}(aa|aa) = \frac{(1 - e^{-t})^2}{2(1 + e^{-2t})}, \quad (40)$$

$$q_{0|t}(aa|ab) = q_{0|t}(bb|ab) = \frac{1}{2}, \quad (41)$$

where (41) is derived from symmetry.

By using (39)–(41) and the general fact (for Markov processes)

$$q_{s|0,t}(\mathbf{x}_s|\mathbf{x}_0, \mathbf{x}_t) = \frac{q_{0,s,t}(\mathbf{x}_0, \mathbf{x}_s, \mathbf{x}_t)}{q_{0,t}(\mathbf{x}_0, \mathbf{x}_t)} = \frac{q_{s|0}(\mathbf{x}_s|\mathbf{x}_0)q_{t|0,s}(\mathbf{x}_t|\mathbf{x}_0, \mathbf{x}_s)}{q_{t|0}(\mathbf{x}_t|\mathbf{x}_0)} = \frac{q_{s|0}(\mathbf{x}_s|\mathbf{x}_0)q_{t|s}(\mathbf{x}_t|\mathbf{x}_s)}{q_{t|0}(\mathbf{x}_t|\mathbf{x}_0)} \quad (42)$$

for $0 \leq s \leq t$, we can compute $q_{s|t}(\cdot|aa)$ for any $s \in [0, t]$ as follows:

$$\begin{aligned} q_{s|t}(aa|aa) &= q_{0|t}(aa|aa)q_{s|0,t}(aa|aa, aa) + q_{0|t}(bb|aa)q_{s|0,t}(aa|bb, aa) \\ &= \frac{(1+e^{-t})^2}{2(1+e^{-2t})} \frac{\frac{1}{4}(1+e^{-s})^2 \frac{1}{4}(1+e^{-(t-s)})^2}{\frac{1}{4}(1+e^{-t})^2}} + \frac{(1-e^{-t})^2}{2(1+e^{-2t})} \frac{\frac{1}{4}(1-e^{-s})^2 \frac{1}{4}(1+e^{-(t-s)})^2}{\frac{1}{4}(1-e^{-t})^2}} \\ &= \frac{((1+e^{-s})^2 + (1-e^{-s})^2)(1+e^{-(t-s)})^2}{8(1+e^{-2t})} = \frac{(1+e^{-2s})(1+e^{-(t-s)})^2}{4(1+e^{-2t})}, \end{aligned} \quad (43)$$

$$\begin{aligned} q_{s|t}(bb|aa) &= q_{0|t}(aa|aa)q_{s|0,t}(bb|aa, aa) + q_{0|t}(bb|aa)q_{s|0,t}(bb|bb, aa) \\ &= \frac{(1+e^{-t})^2}{2(1+e^{-2t})} \frac{\frac{1}{4}(1-e^{-s})^2 \frac{1}{4}(1+e^{-(t-s)})^2}{\frac{1}{4}(1+e^{-t})^2}} + \frac{(1-e^{-t})^2}{2(1+e^{-2t})} \frac{\frac{1}{4}(1+e^{-s})^2 \frac{1}{4}(1+e^{-(t-s)})^2}{\frac{1}{4}(1-e^{-t})^2}} \\ &= \frac{((1-e^{-s})^2 + (1+e^{-s})^2)(1+e^{-(t-s)})^2}{8(1+e^{-2t})} = \frac{(1+e^{-2s})(1+e^{-(t-s)})^2}{4(1+e^{-2t})}, \end{aligned} \quad (44)$$

$$q_{s|t}(ab|aa) = q_{s|t}(ba|aa) = \frac{1}{2}(1 - q_{s|t}(aa|aa) - q_{s|t}(bb|aa)) \quad (45)$$

$$\begin{aligned} &= \frac{1}{2} - \frac{(1+e^{-2s})((1+e^{-(t-s)})^2 + (1+e^{-(t-s)})^2)}{8(1+e^{-2t})} \\ &= \frac{1}{2} - \frac{(1+e^{-2s})(1+e^{-2(t-s)})}{4(1+e^{-2t})} = \frac{1}{4} - \frac{e^{-2s} + e^{-2(t-s)}}{4(1+e^{-2t})}. \end{aligned} \quad (46)$$

We can also compute $q_{s|t}(aa|ab) = q_{s|t}(bb|ab)$ as

$$\begin{aligned} q_{s|t}(aa|ab) &= q_{0|t}(aa|ab)q_{s|0,t}(aa|aa, ab) + q_{0|t}(bb|ab)q_{s|0,t}(aa|bb, ab) \\ &= \frac{1}{2} \frac{\frac{1}{4}(1+e^{-s})^2 \frac{1}{4}(1+e^{-(t-s)})(1-e^{-(t-s)})}{\frac{1}{4}(1+e^{-t})(1-e^{-t})}} + \frac{1}{2} \frac{\frac{1}{4}(1-e^{-s})^2 \frac{1}{4}(1+e^{-(t-s)})(1-e^{-(t-s)})}{\frac{1}{4}(1-e^{-t})(1+e^{-t})}} \\ &= \frac{((1+e^{-s})^2 + (1-e^{-s})^2)(1-e^{-2(t-s)})}{8(1-e^{-2t})} = \frac{(1+e^{-2s})(1-e^{-2(t-s)})}{4(1-e^{-2t})} \\ &= \frac{1}{4} + \frac{e^{-2s} - e^{-2(t-s)}}{4(1-e^{-2t})}. \end{aligned} \quad (47)$$

Let us now compute the probabilities regarding the analytical sampler. To make it simple, let $q_{s|t}(x * \cdot) := q_{s|t}(xa|\cdot) + q_{s|t}(xb|\cdot)$ represent marginals; $q_{s|t}(*y|\cdot)$ is defined similarly. By using this notation and (43)–(47), we have

$$\begin{aligned} p_{s|t}(aa|aa) &= q_{s|t}(a * |aa)q_{s|t}(*a|aa) = q_{s|t}(a * |aa)^2 = (q_{s|t}(aa|aa) + q_{s|t}(ab|aa))^2 \\ &= \left(\frac{(1+e^{-2s})(1+e^{-(t-s)})^2}{4(1+e^{-2t})} + \frac{1}{2} - \frac{(1+e^{-2s})(1+e^{-2(t-s)})}{4(1+e^{-2t})} \right)^2 \\ &= \left(\frac{2(1+e^{-2t}) + (1+e^{-2s})((1+e^{-(t-s)})^2 - (1+e^{-2(t-s)}))}{4(1+e^{-2t})} \right)^2 \\ &= \left(\frac{(1+e^{-2t}) + (1+e^{-2s})e^{-(t-s)}}{2(1+e^{-2t})} \right)^2 = \left(\frac{(1+e^{-(t+s)})(1+e^{-(t-s)})}{2(1+e^{-2t})} \right)^2 \end{aligned} \quad (48)$$

$$\begin{aligned} p_{s|t}(bb|aa) &= q_{s|t}(b * |aa)q_{s|t}(*b|aa) = q_{s|t}(b * |aa)^2 = (q_{s|t}(bb|aa) + q_{s|t}(ba|aa))^2 \\ &= \left(\frac{(1+e^{-2s})(1+e^{-(t-s)})^2}{4(1+e^{-2t})} + \frac{1}{2} - \frac{(1+e^{-2s})(1+e^{-2(t-s)})}{4(1+e^{-2t})} \right)^2 \\ &= \left(\frac{2(1+e^{-2t}) + (1+e^{-2s})((1+e^{-(t-s)})^2 - (1+e^{-2(t-s)}))}{4(1+e^{-2t})} \right)^2 \\ &= \left(\frac{(1+e^{-2t}) - (1+e^{-2s})e^{-(t-s)}}{2(1+e^{-2t})} \right)^2 = \left(\frac{(1-e^{-(t+s)})(1+e^{-(t-s)})}{2(1+e^{-2t})} \right)^2 \end{aligned} \quad (49)$$

Let us compute the sum of (48) and (49) as we use it later:

$$\begin{aligned}
 & p_{s|t}(aa|aa) + p_{s|t}(bb|aa) \\
 &= \left(\frac{(1 + e^{-(t+s)})(1 + e^{-(t-s)})}{2(1 + e^{-2t})} \right)^2 + \left(\frac{(1 - e^{-(t+s)})(1 - e^{-(t-s)})}{2(1 + e^{-2t})} \right)^2 \\
 &= \frac{((1 + e^{-(t+s)})(1 + e^{-(t-s)}))^2 + ((1 - e^{-(t+s)})(1 - e^{-(t-s)}))^2}{4(1 + e^{-2t})^2} \\
 &= \frac{(1 + e^{-2t} + e^{-(t+s)} + e^{-(t-s)})^2 + (1 + e^{-2t} - e^{-(t+s)} - e^{-(t-s)})^2}{4(1 + e^{-2t})^2} \\
 &= \frac{(1 + e^{-2t})^2 + (e^{-(t+s)} + e^{-(t-s)})^2}{2(1 + e^{-2t})^2} = \frac{1}{2} + \frac{(e^{-(t+s)} + e^{-(t-s)})^2}{2(1 + e^{-2t})^2}. \tag{50}
 \end{aligned}$$

Next, $p_{s|t}(aa|ab)$ is the product of two marginals — $q_{s|t}(a * |ab)$ and $q_{s|t}(*a|ab)$, which can be computed as follows:

$$\begin{aligned}
 p_{s|t}(a * |ab) &= q_{0|t}(aa|ab)q_{s|0,t}^1(a|a, a) + q_{0|t}(bb|ab)q_{s|0,t}^1(a|b, a) \\
 &= \frac{1}{2} \frac{\frac{1}{2}(1 + e^{-s})\frac{1}{2}(1 + e^{-(t-s)})}{\frac{1}{2}(1 + e^{-t})} + \frac{1}{2} \frac{\frac{1}{2}(1 - e^{-s})\frac{1}{2}(1 + e^{-(t-s)})}{\frac{1}{2}(1 - e^{-t})} \\
 &= \frac{((1 + e^{-s})(1 - e^{-t}) + (1 - e^{-s})(1 + e^{-t}))(1 + e^{-(t-s)})}{4(1 - e^{-2t})} \\
 &= \frac{(1 - e^{-(t+s)})(1 + e^{-(t-s)})}{2(1 - e^{-2t})} = \frac{1}{2} + \frac{e^{-(t-s)} - e^{-(t+s)}}{2(1 - e^{-2t})}, \\
 p_{s|t}(*a|ab) &= p_{s|t}(a * |ba) = p_{s|t}(b * |ab) = 1 - p_{s|t}(a * |ab) = \frac{1}{2} - \frac{e^{-(t-s)} - e^{-(t+s)}}{2(1 - e^{-2t})},
 \end{aligned}$$

where the latter derivation is from the symmetries of the two dimensions and two characters. By using these, we have

$$p_{s|t}(aa|ab) = p_{s|t}(a * |ab)p_{s|t}(*a|ab) = \frac{1}{4} - \left(\frac{e^{-(t-s)} - e^{-(t+s)}}{2(1 - e^{-2t})} \right)^2. \tag{51}$$

Let us consider iteratively denoising from q_T by using $p_{s|t}$. For an $\epsilon > 0$ and nonnegative integers $n \leq T/\epsilon - 1$, define

$$p_T^\epsilon := p_T, \quad p_{T-(n+1)\epsilon}^\epsilon := \mathbb{E}_{\mathbf{x} \sim p_{T-n\epsilon}^\epsilon} [p_{T-(n+1)\epsilon|T-n\epsilon}(\cdot | \mathbf{x})], \quad n = 0, 1, \dots$$

Our goal is to estimate the difference between $p_{T-n\epsilon}^\epsilon$ and $q_{T-n\epsilon}$ for each n . Let us fix n and set $t = T - n\epsilon$ when computing $p_{t-\epsilon}^\epsilon$ in terms of p_t^ϵ . Because of the symmetry, $p_t^\epsilon(aa) = p_t^\epsilon(bb)$ and $p_t^\epsilon(ab) = p_t^\epsilon(ba) = \frac{1}{2} - p_t^\epsilon(aa)$ hold in general. Therefore, by using (50) and (51), we have

$$\begin{aligned}
 p_{t-\epsilon}^\epsilon(aa) &= p_{t-\epsilon|t}(aa|aa)p_t^\epsilon(aa) + p_{t-\epsilon|t}(aa|bb)p_t^\epsilon(bb) + p_{t-\epsilon|t}(aa|ab)p_t^\epsilon(ab) + p_{t-\epsilon|t}(aa|ba)p_t^\epsilon(ba) \\
 &= p_{t-\epsilon|t}(aa|aa)p_t^\epsilon(aa) + p_{t-\epsilon|t}(bb|aa)p_t^\epsilon(aa) + 2p_{t-\epsilon|t}(aa|ab) \left(\frac{1}{2} - p_t^\epsilon(aa) \right) \\
 &= p_{t-\epsilon|t}(aa|ab) + (p_{t-\epsilon|t}(aa|aa) + p_{t-\epsilon|t}(bb|aa) - 2p_{t-\epsilon|t}(aa|ab))p_t^\epsilon(aa) \\
 &= \frac{1}{4} - \frac{(e^{-\epsilon} - e^{-2(t-\epsilon)})^2}{4(1 - e^{-2t})^2} + \left(\frac{(e^{-\epsilon} + e^{-2(t-\epsilon)})^2}{2(1 + e^{-2t})^2} + \frac{(e^{-\epsilon} - e^{-2(t-\epsilon)})^2}{2(1 - e^{-2t})^2} \right) p_t^\epsilon(aa). \tag{52}
 \end{aligned}$$

To compare it with $q_{t-\epsilon}$, we also compute a similar recurrence equation by replacing p 's with q 's and using (45)–(47):

$$\begin{aligned}
 q_{t-\epsilon}(aa) &= q_{t-\epsilon|t}(aa|ab) + (q_{t-\epsilon|t}(aa|aa) + q_{t-\epsilon|t}(bb|aa) - 2q_{t-\epsilon|t}(aa|ab))q_t(aa) \\
 &= \frac{1}{4} - \frac{e^{-2\epsilon} - e^{-2(t-\epsilon)}}{4(1 - e^{-2t})} + \left(\frac{e^{-2\epsilon} + e^{-2(t-\epsilon)}}{2(1 + e^{-2t})} + \frac{e^{-2\epsilon} - e^{-2(t-\epsilon)}}{2(1 - e^{-2t})} \right) q_t(aa) \tag{53}
 \end{aligned}$$

Let us now compute quantities regarding the coefficients in (52) and (53).

$$\begin{aligned}
 & \frac{e^{-2\epsilon} - e^{-2(t-\epsilon)}}{1 - e^{-2t}} - \frac{(e^{-\epsilon} - e^{-(2t-\epsilon)})^2}{(1 - e^{-2t})^2} \\
 &= \frac{(e^{-2\epsilon} - e^{-2(t-\epsilon)})(1 - e^{-2t}) - (e^{-\epsilon} - e^{-(2t-\epsilon)})^2}{(1 - e^{-2t})^2} \\
 &= \frac{(e^{-2\epsilon} - e^{-2(t-\epsilon)} - e^{-2(t+\epsilon)} + e^{-2(2t-\epsilon)}) - (e^{-\epsilon} - e^{-(2t-\epsilon)})^2}{(1 - e^{-2t})^2} \\
 &= -\frac{(e^{-(t-\epsilon)} - e^{-(t+\epsilon)})^2}{(1 - e^{-2t})^2} = -\frac{e^{-2t}}{(1 - e^{-2t})^2} (e^\epsilon - e^{-\epsilon})^2, \tag{54}
 \end{aligned}$$

$$\begin{aligned}
 & \frac{e^{-2\epsilon} + e^{-2(t-\epsilon)}}{1 + e^{-2t}} - \frac{(e^{-\epsilon} + e^{-(2t-\epsilon)})^2}{(1 + e^{-2t})^2} \\
 &= \frac{(e^{-2\epsilon} + e^{-2(t-\epsilon)})(1 + e^{-2t}) - (e^{-\epsilon} + e^{-(2t-\epsilon)})^2}{(1 + e^{-2t})^2} \\
 &= \frac{(e^{-2\epsilon} + e^{-2(t-\epsilon)} + e^{-2(t+\epsilon)} + e^{-2(2t-\epsilon)}) - (e^{-\epsilon} + e^{-(2t-\epsilon)})^2}{(1 + e^{-2t})^2} \\
 &= \frac{(e^{-(t-\epsilon)} - e^{-(t+\epsilon)})^2}{(1 + e^{-2t})^2} = \frac{e^{-2t}}{(1 + e^{-2t})^2} (e^\epsilon - e^{-\epsilon})^2, \tag{55}
 \end{aligned}$$

$$\begin{aligned}
 & \frac{e^{-2\epsilon} + e^{-2(t-\epsilon)}}{1 + e^{-2t}} + \frac{e^{-2\epsilon} - e^{-2(t-\epsilon)}}{1 - e^{-2t}} \\
 &= \frac{(e^{-2\epsilon} + e^{-2(t-\epsilon)})(1 - e^{-2t}) + (e^{-2\epsilon} - e^{-2(t-\epsilon)})(1 + e^{-2t})}{1 - e^{-4t}} \\
 &= 2 + \frac{2(e^{-2\epsilon} - e^{-2(2t-\epsilon)}) - 2(1 - e^{-4t})}{1 - e^{-4t}} \\
 &= 2 + \frac{2(1 + e^{2(2t-\epsilon)})}{1 - e^{-4t}} (e^{-2\epsilon} - 1). \tag{56}
 \end{aligned}$$

We shall evaluate the difference $\Delta_t^\epsilon := q_t(aa) - p_t^\epsilon(aa)$ by using (52)–(56) as follows:

$$\begin{aligned}
 \Delta_{t-\epsilon}^\epsilon &= -\left(\frac{e^{-2\epsilon} - e^{-2(t-\epsilon)}}{4(1 - e^{-2t})} - \frac{(e^{-\epsilon} - e^{-(2t-\epsilon)})^2}{4(1 - e^{-2t})^2} \right) \\
 &+ \left(\frac{e^{-2\epsilon} + e^{-2(t-\epsilon)}}{2(1 + e^{-2t})} + \frac{e^{-2\epsilon} - e^{-2(t-\epsilon)}}{2(1 - e^{-2t})} \right) (p_t^\epsilon(aa) + \Delta_t^\epsilon) \\
 &- \left(\frac{(e^{-\epsilon} + e^{-(2t-\epsilon)})^2}{2(1 + e^{-2t})^2} + \frac{(e^{-\epsilon} - e^{-(2t-\epsilon)})^2}{2(1 - e^{-2t})^2} \right) p_t^\epsilon(aa) \\
 &= \frac{e^{-2t}}{4(1 - e^{-2t})^2} (e^\epsilon - e^{-\epsilon})^2 + \left(\frac{e^{-2\epsilon} + e^{-2(t-\epsilon)}}{2(1 + e^{-2t})} + \frac{e^{-2\epsilon} - e^{-2(t-\epsilon)}}{2(1 - e^{-2t})} \right) \Delta_t^\epsilon \\
 &+ \left(\frac{e^{-2\epsilon} + e^{-2(t-\epsilon)}}{2(1 + e^{-2t})} - \frac{(e^{-\epsilon} + e^{-(2t-\epsilon)})^2}{2(1 + e^{-2t})^2} + \frac{e^{-2\epsilon} - e^{-2(t-\epsilon)}}{2(1 - e^{-2t})} - \frac{(e^{-\epsilon} - e^{-(2t-\epsilon)})^2}{2(1 - e^{-2t})^2} \right) p_t^\epsilon(aa) \\
 &= \frac{e^{-2t}}{4(1 - e^{-2t})^2} (e^\epsilon - e^{-\epsilon})^2 + \left(1 + \frac{1 + e^{2(2t-\epsilon)}}{1 - e^{-4t}} (e^{-2\epsilon} - 1) \right) \Delta_t^\epsilon \\
 &+ \left(\frac{e^{-2t}}{2(1 + e^{-2t})^2} - \frac{e^{-2t}}{2(1 - e^{-2t})^2} \right) (e^\epsilon - e^{-\epsilon})^2 p_t^\epsilon(aa) \\
 &= \left(\frac{e^{-2t}}{2(1 + e^{-2t})^2} p_t^\epsilon(aa) + \frac{e^{-2t}}{2(1 - e^{-2t})^2} \left(\frac{1}{2} - p_t^\epsilon(aa) \right) \right) (e^\epsilon - e^{-\epsilon})^2 \\
 &+ \left(1 + \frac{1 + e^{2(2t-\epsilon)}}{1 - e^{-4t}} (e^{-2\epsilon} - 1) \right) \Delta_t^\epsilon. \tag{57}
 \end{aligned}$$

Since $p_t^\epsilon(aa) = p_t^\epsilon(bb) \leq 1/2$, we have

$$\begin{aligned} & \frac{e^{-2t}}{2(1+e^{-2t})^2} p_t^\epsilon(aa) + \frac{e^{-2t}}{2(1-e^{-2t})^2} \left(\frac{1}{2} - p_t^\epsilon(aa) \right) \\ & \geq \frac{1}{2} \min \left\{ \frac{e^{-2t}}{2(1+e^{-2t})^2}, \frac{e^{-2t}}{2(1-e^{-2t})^2} \right\} = \frac{e^{-2t}}{4(1-e^{-2t})^2}. \end{aligned}$$

Additionally, as the Taylor series of $(e^\epsilon - e^{-\epsilon})^2 = e^{2\epsilon} + e^{-2\epsilon} - 2$ is given by $\sum_{k=1}^{\infty} \frac{2}{(2k)!} (2\epsilon)^{2k}$, we especially have $(e^\epsilon - e^{-\epsilon})^2 \geq 4\epsilon^2$. Thus, we obtain

$$\begin{aligned} & \left(\frac{e^{-2t}}{2(1+e^{-2t})^2} p_t^\epsilon(aa) + \frac{e^{-2t}}{2(1-e^{-2t})^2} \left(\frac{1}{2} - p_t^\epsilon(aa) \right) \right) (e^\epsilon - e^{-\epsilon})^2 \\ & \geq \frac{e^{-2t}}{4(1-e^{-2t})^2} \cdot 4\epsilon^2 = \frac{e^{-2t}}{(1-e^{-2t})^2} \epsilon^2. \end{aligned} \quad (58)$$

Also, since $e^{-2\epsilon} \geq 1 - 2\epsilon$, we have

$$1 + \frac{1 + e^{2(2t-\epsilon)}}{1 - e^{-4t}} (e^{-2\epsilon} - 1) \geq 1 - \frac{2(1 + e^{2(2t-\epsilon)})}{1 - e^{-4t}} \epsilon \geq 1 - \frac{4}{1 - e^{-4t}} \epsilon. \quad (59)$$

Suppose we are working on the time interval $[\delta, T]$ for some $\delta, T > 0$. Let us take $\epsilon \leq \delta/2$; then we have

$$1 - \frac{4}{1 - e^{-4t}} \epsilon \geq 1 - \frac{4}{4t} \epsilon \geq 1 - \frac{\epsilon}{\delta} > 0. \quad (60)$$

For (58), we have

$$\frac{e^{-2t}}{(1 - e^{-2t})^2} \epsilon^2 \geq e^{-2t} \epsilon^2 \geq e^{-2T} \epsilon^2. \quad (61)$$

By combining (57)–(61), we first see that Δ_t^ϵ is nonnegative for all $t = T - n\epsilon$ by induction on $n = 0, 1, \dots$ (assuming $\epsilon \leq \delta/2$ and $t \in [\delta, T]$). Then, we obtain the following simple inequality:

$$\Delta_{t-\epsilon}^\epsilon \geq \left(1 - \frac{\epsilon}{\delta}\right) \Delta_t^\epsilon + e^{-2T} \epsilon^2$$

By recalling that $t = T - n\epsilon$, we can rewrite it as

$$\left(1 - \frac{\epsilon}{\delta}\right)^{-(n+1)} \Delta_{T-(n+1)\epsilon}^\epsilon \geq \left(1 - \frac{\epsilon}{\delta}\right)^{-n} \Delta_{T-n\epsilon}^\epsilon + \left(1 - \frac{\epsilon}{\delta}\right)^{-(n+1)} e^{-2T} \epsilon^2.$$

Since $\Delta_T^\epsilon = 0$, we have

$$\Delta_{T-n\epsilon}^\epsilon \geq \left(1 - \frac{\epsilon}{\delta}\right)^n \sum_{k=1}^n \left(1 - \frac{\epsilon}{\delta}\right)^{-k} e^{-2T} \epsilon^2 = \sum_{k=0}^{n-1} \left(1 - \frac{\epsilon}{\delta}\right)^k e^{-2T} \epsilon^2. \quad (62)$$

Since $n \leq T/\epsilon$ and $(1 - 1/x)^x$ is increasing over $x > 1$, for $k = 0, \dots, n-1$, we have

$$\left(1 - \frac{\epsilon}{\delta}\right)^k \geq \left(1 - \frac{\epsilon}{\delta}\right)^n \geq \left(1 - \frac{\epsilon}{\delta}\right)^{T/\epsilon} = \left(\left(1 - \frac{\epsilon}{\delta}\right)^{\delta/\epsilon} \right)^{T/\delta} \geq \left(\left(1 - \frac{1}{2}\right)^2 \right)^{T/\delta} = 2^{-2T/\delta},$$

where we have exploited the assumption $\epsilon \leq \delta/2$ (so that $\delta/\epsilon \geq 2$). By applying this to (62), we obtain

$$\Delta_{T-n\epsilon}^\epsilon \geq (2^{1/\delta} e)^{-2T} n \epsilon^2.$$

Now, let $\epsilon = (T - \delta)/N$ for the given N . Since $N \geq \frac{2(T-\delta)}{\delta}$ and thus $\epsilon \leq \delta/2$, we have

$$\Delta_\delta^\epsilon = \Delta_{T-N\epsilon}^\epsilon \geq (2^{1/\delta} e)^{-2T} N \epsilon^2 = (2^{1/\delta} e)^{-2T} \frac{(T - \delta)^2}{N}.$$

Finally, as $d_{\text{TV}}(q_\delta, p_\delta^{(T-\delta)/N}) \geq \Delta_\delta^\epsilon$, the constant $c = (2^{1/\delta} e)^{-2T} (T - \delta)^2$ satisfies (23). \square

E. Control variates

When we want to compute an expectation $\mathbb{E}[f(\mathbf{x})]$, instead of directly doing the Monte Carlo estimate $\frac{1}{N} \sum_{i=1}^N f(\mathbf{x}_i) \approx \mathbb{E}[f(\mathbf{x})]$, we can find a function $g \approx f$ such that $\mathbb{E}[g(\mathbf{x})]$ is tractable and then do the Monte Carlo estimate for the remainder term:

$$\frac{1}{N} \sum_{i=1}^N (f(\mathbf{x}_i) - g(\mathbf{x}_i)) + \mathbb{E}[g(\mathbf{x})] \approx \mathbb{E}[f(\mathbf{x})]. \quad (63)$$

This left-hand side is still an unbiased estimator of $\mathbb{E}[f(\mathbf{x})]$, and ideally has a lower variance than the vanilla Monte Carlo estimator $\frac{1}{N} \sum_{i=1}^N f(\mathbf{x}_i)$ if $g \approx f$ is a good function approximation. The role of g in (63) is called a *control variate* (Glasserman, 2004; Oates et al., 2017).

E.1. Marginal-matching product model as control variate

We briefly discuss how the product model \bar{p}^θ given in (12) satisfies the following favorable properties (already shown in Section A.2) for being a control variate:

- (i) it reasonably approximates $p_s^\theta(\cdot | \mathbf{x}_t)$, and
- (ii) $\mathbb{E}_{\mathbf{x} \sim q}[g(\mathbf{x})]$ is easy to compute/approximate.

For point (i), note that \bar{p}^θ is defined as a product model having the same marginal as p^θ . Since dimensionally independent modeling (when combined with multi-step sampling) works as in Theorem 1, \bar{p}^θ should approximate p^θ to a certain degree; see also Lemma 1 for a quantitative understanding. The remainder $p^\theta - \bar{p}^\theta$ can then be regarded as a dimensional correlation captured by p^θ , with which we conduct a usual Monte Carlo integration.

Regarding (ii), given a product distribution $\bar{p}(\mathbf{x}) = \prod_{d=1}^D \bar{p}^d(x^d)$ over $\mathcal{X} = \mathcal{S}^D$, we can indeed compute $H(q, \bar{p})$ with a Monte Carlo integral using samples of η as

$$\begin{aligned} H(q, \bar{p}) &= \mathbb{E}_{\mathbf{x}_s \sim q}[-\log \bar{p}(\mathbf{x}_s)] = \mathbb{E}_\eta \mathbb{E}_{\mathbf{x}_s \sim q^\eta}[-\log \bar{p}(\mathbf{x}_s)] \\ &= \mathbb{E}_\eta [H(q^\eta, \bar{p})] = \mathbb{E}_\eta \left[-\sum_{d=1}^D \sum_{x_s^d \in \mathcal{S}} q^\eta(x_s^d) \log \bar{p}^d(x_s^d) \right]. \end{aligned} \quad (64)$$

While it still requires Monte Carlo with η to estimate this, it utilizes the product structure of each q^η and \bar{p} to exactly compute $H(q^\eta, \bar{p})$. Thus, we heuristically expect it to be more accurate than the Monte Carlo estimate using samples from q .

E.2. Derivations of dimension-wise computable control variates for mixture model

Convex upper bound as control variate. To simplify the notation and situation, suppose we are given probability distributions $q = \mathbb{E}_\eta[q^\eta]$ and $p^\theta = \mathbb{E}_\lambda[p^{\theta, \lambda}]$, where q^η and $p^{\theta, \lambda}$ are product distributions, i.e., we have

$$q^\eta(\mathbf{x}) = \prod_{d=1}^D q^{\eta, d}(x^d), \quad p^{\theta, \lambda}(\mathbf{x}) = \prod_{d=1}^D p^{\theta, \lambda, d}(x^d).$$

By letting H be the (cross) entropy, we want to minimize

$$D_{\text{KL}}(q || p^\theta) = H(q, p^\theta) - H(q) = \mathbb{E}_{\mathbf{x} \sim q}[-\log p^\theta(\mathbf{x})] - \mathbb{E}_{\mathbf{x} \sim q}[-\log q(\mathbf{x})].$$

Since q is fixed, we simply want to minimize

$$H(q, p^\theta) = \mathbb{E}_{\mathbf{x} \sim q}[-\log p^\theta(\mathbf{x})] = \mathbb{E}_\eta \mathbb{E}_{\mathbf{x} \sim q^\eta}[-\log p^\theta(\mathbf{x})]$$

with regard to θ . However, it might have a high variance when we only sample $\mathbf{x} \sim q$ and execute Monte Carlo. One option is using the following upper bound like a negative ELBO given by Jensen's inequality (convex inequality) as a control variate:

$$-\log p^\theta(\mathbf{x}) = -\log \mathbb{E}_\lambda[p^{\theta, \lambda}(\mathbf{x})] \leq \mathbb{E}_\lambda[-\log p^{\theta, \lambda}(\mathbf{x})].$$

Indeed, its expectation regarding $\mathbf{x} \sim q$ is dimension-wise computable as

$$\begin{aligned} & \mathbb{E}_{\mathbf{x} \sim q} \mathbb{E}_\lambda [-\log p^{\theta, \lambda}(\mathbf{x})] \\ &= \mathbb{E}_\eta \mathbb{E}_{\mathbf{x} \sim q^\eta} \mathbb{E}_\lambda [-\log p^{\theta, \lambda}(\mathbf{x})] = \mathbb{E}_\eta \mathbb{E}_\lambda \mathbb{E}_{\mathbf{x} \sim q^\eta} [-\log p^{\theta, \lambda}(\mathbf{x})] \\ &= \mathbb{E}_\eta \mathbb{E}_\lambda \sum_{d=1}^D \mathbb{E}_{x^d \sim q^{\eta, d}} [-\log p^{\theta, \lambda, d}(x^d)] = \mathbb{E}_\eta \mathbb{E}_\lambda \left[-\sum_{d=1}^D \sum_{x^d} q^{\eta, d}(x^d) \log p^{\theta, \lambda, d}(x^d) \right], \end{aligned}$$

which does not require Monte Carlo sampling of \mathbf{x} . Overall, we can decompose the computation as

$$H(q, p^\theta) = \underbrace{\mathbb{E}_{\mathbf{x} \sim q} [-\log p^\theta(\mathbf{x}) + \mathbb{E}_\lambda [\log p^{\theta, \lambda}(\mathbf{x})]]}_{\text{Monte Carlo approximation}} + \underbrace{\mathbb{E}_{\mathbf{x} \sim q} \mathbb{E}_\lambda [-\log p^{\theta, \lambda}(\mathbf{x})]}_{\text{dim-wise computable}}.$$

Marginal control variate. The previous convex upper bound seems good, but since

$$\mathbb{E}_{\mathbf{x} \sim q} \mathbb{E}_\lambda [-\log p^{\theta, \lambda}(\mathbf{x})] = \mathbb{E}_\lambda [H(q, p^{\theta, \lambda})] \geq \inf_\lambda H(q, p^{\theta, \lambda}),$$

it might be a very loose bound (we want the mixture to outperform the best product distribution $p^{\theta, \lambda}$). To make it more practical, we can consider its dimension-wise tractable lower bound as follows:

$$\mathbb{E}_{\mathbf{x} \sim q} \mathbb{E}_\lambda [-\log p^{\theta, \lambda}(\mathbf{x})] = \mathbb{E}_\eta \sum_{d=1}^D \mathbb{E}_{x^d \sim q^{\eta, d}} \mathbb{E}_\lambda [-\log p^{\theta, \lambda, d}(x^d)] \geq -\mathbb{E}_\eta \sum_{d=1}^D \mathbb{E}_{x^d \sim q^{\eta, d}} \log \mathbb{E}_\lambda [p^{\theta, \lambda, d}(x^d)],$$

which is given by Jensen's inequality as well. Therefore, if we define the product distribution

$$\bar{p}^\theta(\mathbf{x}) = \prod_{d=1}^D \bar{p}^{\theta, d}(x_d), \quad \bar{p}^{\theta, d}(x_d) = \mathbb{E}_\lambda [\bar{p}^{\theta, d}(x_d)],$$

we have $\mathbb{E}_{\mathbf{x} \sim q} \mathbb{E}_\lambda [-\log p^{\theta, \lambda}(\mathbf{x})] \leq \mathbb{E}_{\mathbf{x} \sim q} [-\log \bar{p}^\theta(\mathbf{x})]$ and this alternative is also dimension-wise computable. Since p^θ and \bar{p}^θ coincide in each one-dimensional marginal, the difference between these two can be regarded as the result of dimensional correlation.

Therefore, we propose the following decomposition, which is also discussed in Section A.2:

$$H(q, p^\theta) = \underbrace{\mathbb{E}_{\mathbf{x} \sim q} [-\log p^\theta(\mathbf{x}) + \log \bar{p}^\theta(\mathbf{x})]}_{\text{Monte Carlo approximation}} + \underbrace{\mathbb{E}_{\mathbf{x} \sim q} [-\log \bar{p}^\theta(\mathbf{x})]}_{\text{dim-wise computable}}.$$

E.3. Product teacher model as control variate

For two models with the same marginals, we have the following proposition:

Proposition 7. *Let q, \tilde{q} be probability distributions on $\mathcal{X} = \mathcal{S}^D$ with the same marginals $q^d = \tilde{q}^d$. Then, for a product distribution $p(\mathbf{x}) = \prod_d p^d(x^d)$ over \mathcal{X} , we have $H(q, p) = H(\tilde{q}, p)$.*

Proof. It suffices to prove that $H(q, p)$ can be computed only by using the marginals q^d . Indeed, we have

$$\mathbb{E}_{\mathbf{x} \sim q} [\log p(\mathbf{x})] = \mathbb{E}_{\mathbf{x} \sim q} \left[\sum_{d=1}^D \log p^d(x^d) \right] = \sum_{d=1}^D \mathbb{E}_{\mathbf{x} \sim q} [\log p^d(x^d)] = \sum_{d=1}^D \sum_{x^d} q^d(x^d) \log p^d(x^d),$$

and it yields the desired conclusion. \square

From this proposition, under $p_{0|t}^{\psi, d} \approx q_{0|t}^d$ and the fact that \bar{p}^θ is a product model, we have

$$\mathbb{E}_{\mathbf{x}_t \sim q_t} \left[H(q_{0|t}(\cdot | \mathbf{x}_t), \bar{p}_{0|t}^\theta(\cdot | \mathbf{x}_t)) \right] \approx \mathbb{E}_{\mathbf{x}_t \sim q_t} \left[H(p_{0|t}^\psi(\cdot | \mathbf{x}_t), \bar{p}_{0|t}^\theta(\cdot | \mathbf{x}_t)) \right].$$

Since $H(p_1, p_2) = D_{\text{KL}}(p_1 \| p_2) + H(p_2, p_2)$ this right-hand side can be rewritten as

$$\mathbb{E}_{\mathbf{x}_t \sim q_t} \left[H(p_{0|t}^\psi(\cdot | \mathbf{x}_t), \bar{p}_{0|t}^\theta(\cdot | \mathbf{x}_t)) \right] = \mathbb{E}_{\mathbf{x}_t \sim q_t} \left[D_{\text{KL}}(p_{0|t}^\psi(\cdot | \mathbf{x}_t) \| \bar{p}_{0|t}^\theta(\cdot | \mathbf{x}_t)) \right] + \text{const.},$$

where the constant term is independent of θ . Since the KL divergence between two product distributions decomposes into the sum of the KL divergence between each marginal, we obtain approximation (14).

F. Experimental details

F.1. Discretized Gaussian diffusions

F.1.1. SAMPLING SCHEMES

In the experiments, we used the following two sampling schemes when evaluating the already trained product teacher model.

τ -leaping. In Campbell et al. (2022), the authors first approximate the infinitesimal transition rate by using each marginal $p_{0|t}^{\psi,d}$. Indeed, the transition rate can be represented only with $q_{0|t}^d$ and does not require a joint conditional distribution (Campbell et al., 2022, Proposition 3). After estimating the transition rate, they apply a dimensionally parallel sampling method called τ -leaping (Gillespie, 2001) coming from computational chemistry. Simply put, τ -leaping is a sort of generalization of the Euler method for solving the backward SDE, exploiting the ordinal structure of \mathcal{S} . We omit the corrector steps; the τ -leaping in Table 5 corresponds to τ LDR-0 in Campbell et al. (2022).

Analytical sampling. Although the τ -leaping (or Euler method) is efficient with a large number of sampling steps, we find that it deteriorates when we reduce the number of steps seemingly due to discretization error. Analytical sampling (Sun et al., 2023) (a.k.a. Tweedie τ -leaping; Lou et al., 2024), which is simply a parallel exact sampling of each dimension given as

$$q_{s|t}^d(x_s^d|\mathbf{x}_t) = \sum_{x_0^d} q_{s|0,t}^d(x_s^d|x_0^d, x_t^d) q_{0|t}^d(x_0^d|\mathbf{x}_t) \approx \sum_{x_0^d} q_{s|0,t}^d(x_s^d|x_0^d, x_t^d) p_{0|t}^{\psi,d}(x_0^d|\mathbf{x}_t), \quad (65)$$

does not suffer so much from the discretization. This is also mentioned in Gu et al. (2022) as a fast inference strategy, though they do not discuss dimensional correlations. See also (71) for the derivation of a dimensionally independent denoiser based on the product model $p_{0|t}$.

Note that these schemes are both dimensionally independent in the sense of (1) while not explicitly modeling $p_{s|t}$. Indeed, the dimensional independence is ubiquitous even when modeling $p_{s|t}$ implicitly. First, the reparametrization $p_{s|t}(\mathbf{x}_s|\mathbf{x}_t) = \sum_{\mathbf{x}_0} p_{0|t}(\mathbf{x}_0|\mathbf{x}_t) q_{s|0,t}(\mathbf{x}_s|\mathbf{x}_0, \mathbf{x}_t)$ (Austin et al., 2021; Gu et al., 2022), also used in analytical sampling, is dimensionally independent, provided that $p_{0|t}(\cdot|\mathbf{x}_t)$ is given by a product model and the forward diffusion is dimensionally independent. Second, we can apparently avoid the heuristic in the above modeling through the estimation of the transition rate in the continuous-time discrete diffusion (Campbell et al., 2022, Proposition 3), but the existing sampling schemes of \mathbf{x}_s given \mathbf{x}_t in continuous-time settings including τ -leaping (Campbell et al., 2022) and the Euler-based method (Sun et al., 2023; Lou et al., 2024) are still dimensionally independent.

N -step sampling in the actual experiment is given as follows. We first set the timesteps $0 = t_0 < t_1 < \dots < t_N = 1$, with $t_i = 0.01 + 0.99 \times \frac{i-1}{N-1}$ for $i \geq 1$. Given a terminal noise \mathbf{x}_{t_N} , we sample \mathbf{x}_{t_i} with our $p_{t_i|t_{i+1}}$ iteratively for $i = N-1, N-2, \dots, 1$. Finally, we sample $\mathbf{x}_0 \in \operatorname{argmax} p_{0|t_1}^{\psi}(\cdot|\mathbf{x}_{t_1})$ when using the teacher product model and $\mathbf{x}_0 \in \operatorname{argmax} p_{0|t_1}^{\theta}(\cdot|\mathbf{x}_{t_1}; \lambda)$ with a random λ when using the student mixture model.

F.1.2. ADDITIONAL EXPERIMENTAL RESULTS

To complement the main experimental results presented in Section 5.1, we provide additional details and analysis here.

FID/IS results of fewer sampling steps. For additional comparison, we computed FID/IS of the teacher and student models using 10K samples (fewer than the 50K samples used in the main body, so the numbers may be slightly worse) for 2-20 steps in Table 4. In terms of FID, our method achieves approximately 1.4 times acceleration in the 10–20 steps range. However, it does not perform well in very few steps (e.g., 2–4 steps).

More detailed results of Table 1. We evaluated two different sampling strategies with the teacher model p^{ψ} : (1) τ -leaping (Campbell et al., 2022) and (2) analytical sampling (Sun et al., 2023; Lou et al., 2024). The complete evaluation results are shown in Table 5, which extends the results presented in Section 5.1. A notable observation is that analytical sampling significantly outperforms τ -leaping in terms of sampling efficiency; 40-step analytical sampling achieves better FID scores than 1000-step τ -leaping.

Regarding distilled models, as we have highlighted in Section 5.1, p^{θ} works well in 10 steps, while it deteriorates as we grow the number of sampling steps. The hybrid model interestingly beats other models in 20-step FID and shows almost the

Table 4. Comparison of models on CIFAR-10 dataset in various sampling steps. Same setting as Table 1 except that 10,000 generated samples were used for computing FID/IS.

	#steps	2	4	6	8	10	12	14	16	18	20
FID	teacher	392.24	173.29	78.24	49.44	34.70	26.33	21.47	18.05	15.80	14.42
	student	411.70	147.67	59.62	33.85	22.57	17.46	14.37	12.86	12.28	11.81
IS	teacher	1.17	2.99	5.90	7.01	7.45	7.82	7.96	8.16	8.35	8.46
	student	1.25	3.48	6.71	7.68	8.17	8.33	8.37	8.50	8.38	8.39

Table 5. Comparison of models on CIFAR-10 dataset. Fréchet inception distance (FID \downarrow) against training dataset and inception score (IS \uparrow) are calculated using 50,000 generated samples. *: reported values from Campbell et al. (2022).

Method	10 steps		20 steps		40 steps		1000 steps	
	FID	IS	FID	IS	FID	IS	FID	IS
$p^{\psi} + \tau$ -leaping	-	-	-	-	315.75	1.66 \pm 0.01	8.10*	8.74*
$p^{\psi} +$ analytical	32.61	7.59 \pm 0.10	12.36	8.55 \pm 0.13	8.01	8.77\pm0.09	-	-
p^{θ} (student)	20.64	8.29\pm0.13	9.77	8.52 \pm 0.08	9.66	8.28 \pm 0.10	-	-
p^{θ} & p^{ψ} (hybrid)	25.54	8.00 \pm 0.11	9.47	8.56\pm0.14	8.02	8.43 \pm 0.11	-	-

same 40-step FID with the teacher, while using the student solely gets worse in 40 steps. We hypothesize (elaborating on the description in Section 5.1) that this is because the true denoiser $q_{s|t}$ ($s < t$) becomes more “dimensionally independent” as $t - s$ or t is small. The former condition (small $t - s$) explains the worse performance gain of the mixture model as the number of steps grows, and the latter partially explains the effectiveness of using the combined model. However, we should further consider different forward diffusion and/or noise schedule to investigate it.

F.1.3. IMPLEMENTATION AND TRAINING

Diffusion modeling. As explained in Section 5, the state-space has $D = 3 \times 32 \times 32$ dimensions, and each dimension has 256 possibilities of pixel values which corresponds to $\mathcal{S} = \{0, \dots, 255\}$. The forward diffusion process is defined through a discretized Gaussian transition rate with $T = 1$ (Campbell et al., 2022, Section E).

Network architecture. All the models are based on the implementation explained in Campbell et al. (2022, Section H.2), where $p_{0|t}^{\psi}$ is parameterized with a U-net (Ho et al., 2020) that has feature resolutions from 32×32 to 4×4 . Since the output of the original U-net architecture (Ho et al., 2020) is a D -dimensional sequence (in \mathcal{S}^D) rather than D marginal distributions, Campbell et al. (2022) adjusted the network so that it first outputs a Gaussian distribution over the real line for each marginal and then normalized it to obtain a distribution over \mathcal{S} . The time t in their implementation is passed to feature map used in Ho et al. (2020), and this embedding is fed to the upsampling/downsampling layers of the U-net after passing through SiLU-activated linear layers (Elfwing et al., 2018). See Campbell et al. (2022, Section H.2) and their GitHub repository for more details on the original implementation. All the models output the estimation of $q_{0|t}$, and we conduct denoising from time t to time s by using the dimension-wise analytical sampling (65), except for the τ -leaping benchmark in Table 5.

The only change we made on the architecture is the insertion of λ . We sample λ from the uniform distribution over $[0, 1]$, so we can basically use the same embedding architecture as the time t . For the downsampling layers, the embedding of λ is concatenated with the time embedding, and then fed to the linear layers. After the linear layers, similarly to the time embedding, it is added to the latent vector of the image. For the upsampling layers, we concatenate the embeddings of λ , t , and the pixel-wise average of the 4×4 resolution latent tensor, and the remaining process is the same as for the downsampling layers.

Training. Since our model is an expansion of the original model for p^{ψ} , we trained (finetuned) our student model p^{θ} from the checkpoint of p^{ψ} . The bias terms and the final layers concerning the embeddings of λ are zero-initialized, and the rest are randomly initialized following the default setting of the original model.

For the Di4C finetuning, we followed the original setting in terms of the use of the Adam optimizer and the learning rate 2×10^{-4} as well as other hyperparameters. The two primary differences in training are loss functions and the training steps/minibatch size (due to the Monte Carlo for λ). For the former point, we basically used

$$\mathcal{L}_{\text{distil}}(\theta; \psi, q_\delta, \delta) + \mathcal{L}_{\text{consis}}(\theta; \psi, q_t, 0, t - \Delta t, t) + \alpha_t \mathcal{L}_{\text{corr}}(\theta; t) + \mathcal{L}_{\text{marginal}}(\theta; \psi, q_t, t), \quad (66)$$

with techniques described in Section A. The following are additional details:

- Sampling from q_δ and q_t is based on the same sample of $\mathbf{x}_0 \sim q_0$.
- $\delta = 0.01$ with probability 1/2; otherwise, δ is taken uniformly from $[0.01, 0.02]$.
- Δt is sampled from a log-uniform distribution over $[0.001, 0.01]$; t is then sampled uniformly from $[0.01 + \Delta t, 1]$.
- We can use several α_t as in the ablation study in the following section. In the main model p^θ given in Table 5, we used the following sigmoid-based function as α_t :

$$g(t) = \frac{1}{1 + \exp(10 - 20t)}. \quad (67)$$

Regarding the training steps/minibatch details, the original teacher model checkpoint had been trained for 2M steps, where each step used 128 images from the CIFAR-10 dataset as a minibatch. In our finetuning, we stopped all the trainings in 320K steps (without warm-ups). Each step used a minibatch of $128/L$ images from the CIFAR-10 dataset, where L is a batch size for λ in the Monte Carlo estimates; we set $M = N = L$ in (10). $L = 16$ is adopted in our model in Table 5, while the ablation study in the following section compares various choices of L .

Evaluation. We measured FID and IS with the PyTorch-based implementation¹ following Campbell et al. (2022).

F.1.4. ABLATION STUDY

Table 6. Ablation study on α_t and use of control variates.

Method	10 steps		20 steps		40 steps	
	FID	IS	FID	IS	FID	IS
$p^\psi + \text{analytical}$	32.61	7.59 \pm 0.10	12.36	8.55 \pm 0.13	8.01	8.77\pm0.09
$\alpha_t = 0$	26.23	8.02 \pm 0.09	11.55	8.59\pm0.07	9.01	8.65 \pm 0.14
$\alpha_t = 0$, w/o CV	44.09	6.79 \pm 0.10	26.16	7.54 \pm 0.10	22.20	7.72 \pm 0.08
$\alpha_t = 1$	24.14	7.54 \pm 0.08	12.30	8.06 \pm 0.07	10.32	8.14 \pm 0.10
$\alpha_t = 1$, w/o CV	26.92	8.12 \pm 0.08	13.77	8.57 \pm 0.14	10.59	8.66 \pm 0.05
$\alpha_t = t$	24.21	8.10 \pm 0.11	10.85	8.55 \pm 0.08	9.27	8.51 \pm 0.10
$\alpha_t = g(t)$ (see (67))	22.77	8.19\pm0.08	10.07	8.54 \pm 0.12	9.01	8.42 \pm 0.11

As an ablation study, we compared several loss functions, mainly changing α_t , which controls the degree of dimensional correlations we aim to learn from datapoints. We also investigated whether the use of control variates is effective. The results are shown in Table 6, where “w/o CV” means that the control variates were not used in training. The efficiency of control variates was consistent, while $\alpha_t = 0$ and $\alpha_t = 1$ had pros and cons. Non-constant functions of α_t worked better, partially matching the hypothesis discussed at the end of Section 5.

Additionally, we compared different batch-sizes of λ in Table 7 (also see the end of the previous section). The non-constant $\alpha_t = g(t)$ was used in all the settings. L in the table represents the batch size of λ in Monte Carlo sampling. There is a certain tradeoff between FID and IS in 10- or 20-step sampling; we can expect a better FID with a larger L (smaller data batch), while a smaller L tends to result in a better IS.

¹<https://github.com/w86763777/pytorch-image-generation-metrics>, which got renamed from the original repository “pytorch-gan-metrics” to “pytorch-image-generation-metrics”.

Table 7. Ablation study on Monte Carlo sample size of λ .

Method	10 steps		20 steps		40 steps	
	FID	IS	FID	IS	FID	IS
p^ψ + analytical	32.61	7.59 \pm 0.10	12.36	8.55 \pm 0.13	8.01	8.77\pm0.09
$L = 2$	27.29	8.00 \pm 0.01	11.42	8.67\pm0.12	8.94	8.64 \pm 0.09
$L = 4$	24.94	8.05 \pm 0.14	10.66	8.60 \pm 0.11	8.90	8.59 \pm 0.07
$L = 8$	22.77	8.19 \pm 0.08	10.07	8.54 \pm 0.12	9.01	8.42 \pm 0.11
$L = 16$	20.64	8.29\pm0.13	9.77	8.52 \pm 0.08	9.66	8.28 \pm 0.10
$L = 32$	20.25	8.28 \pm 0.13	9.93	8.44 \pm 0.10	9.91	8.26 \pm 0.13
$L = 64$	19.26	8.13 \pm 0.10	10.13	8.26 \pm 0.11	10.59	8.02 \pm 0.15

F.1.5. GENERATED SAMPLES

Figure 5 shows image examples corresponding to Table 1, which were all generated with the analytical sampling.

F.2. Masked generative image modeling

F.2.1. MASKED DIFFUSION MODELING

As described in Section 5.2, we use the pretrained VQGAN codebook \mathcal{S}^* with $|\mathcal{S}^*| = 1024$ and add one [MASK] token to define \mathcal{S} . Also, $D = 256$ in this experiment. The only ingredient we need in masked diffusion is the *masking probability* m_t : a monotonically increasing function with $m_0 = 0$ and $m_1 = 1$. Following Chang et al. (2022) and Besnier & Chen (2023), we model the forward process of \mathbf{x}_t given $\mathbf{x}_0 \in (\mathcal{S}^*)^D$ as

$$x_t^d = \begin{cases} [\text{MASK}] & \text{with probability } m_t, \\ x_0^d & \text{with probability } 1 - m_t, \end{cases} \quad (68)$$

independently for each $d \in \{1, \dots, D\}$. Note that (68) does not necessarily determine a Markov process (and indeed an explicit Markov formulation is not needed for training). If one needs a Markov formulation, however, for $t > s > 0$ and $x \in \mathcal{S}^* = \mathcal{S} \setminus \{[\text{MASK}]\}$, we have

$$m_t = q_{t|0}^d([\text{MASK}]|x) = q_{s|0}^d([\text{MASK}]|x) + q_{t|s}^d([\text{MASK}]|x)q_{s|0}^d(x|x) = m_s + (1 - m_s)q_{t|s}^d([\text{MASK}]|x)$$

and thus we have

$$q_{t|s}^d([\text{MASK}]|x) = \frac{m_t - m_s}{1 - m_s}. \quad (69)$$

In the actual experiment, we used the arccos scheduler $m_t = 2 \arccos(1 - t)/\pi$.

F.2.2. CONFIDENCE-BASED SAMPLING

Given \mathbf{x}_t at time t , let $M_t := \{d \in \{1, \dots, D\} \mid x_t^d = [\text{MASK}]\}$. Suppose we have a *product* model $p_{0|t}(\cdot|\mathbf{x}_t) = \prod_{d=1}^D p_{0|t}^d(\cdot|\mathbf{x}_t)$ such that $p_{0|t}^d([\text{MASK}]|\mathbf{x}_t) = 0$ for all $d \in \{1, \dots, D\}$ and $p_{0|t}^d(x_t^d|x_t^d) = 1$ for all $d \notin M_t$. Let us explain how we sample \mathbf{x}_s with $s < t$ in one step of confidence-based sampling (Chang et al., 2022; Besnier & Chen, 2023). Following the original implementation of MaskGIT-PyTorch (Besnier & Chen, 2023) we conduct the sampling as follows:

1. Sample $\tilde{\mathbf{x}}_0 = (\tilde{x}_0^d)_{d=1}^D \sim p_{0|t}(\cdot|\mathbf{x}_t)$. Note that we have $\tilde{x}_0^d \neq [\text{MASK}]$ for each d and $\tilde{x}_0^d = x_t^d$ for $d \notin M_t$.
2. Calculate the confidence for each chosen \tilde{x}_0^d for $d \notin M_t$ as

$$\text{conf}(d) = \log p_{0|t}^d(\tilde{x}_0^d|\mathbf{x}_t) + c_{\text{gb}}(t) \cdot \epsilon_{\text{gb}}^d,$$

where the second term is given by constant multiplication of Gumbel noise to add stochasticity in confidence-based sampling (see, e.g., Comunità et al. (2024, Section 3.4.4) for a concise explanation). To be concrete, ϵ_{gb}^d for each $d \notin M_t$ is an independent standard Gumbel noise, and c_{gb} is a scale factor given by $c_{\text{gb}}(t) = \frac{9}{2} \frac{t-1/N}{1-1/N}$ in our experiments, where N is the number of steps in the whole sampling process.

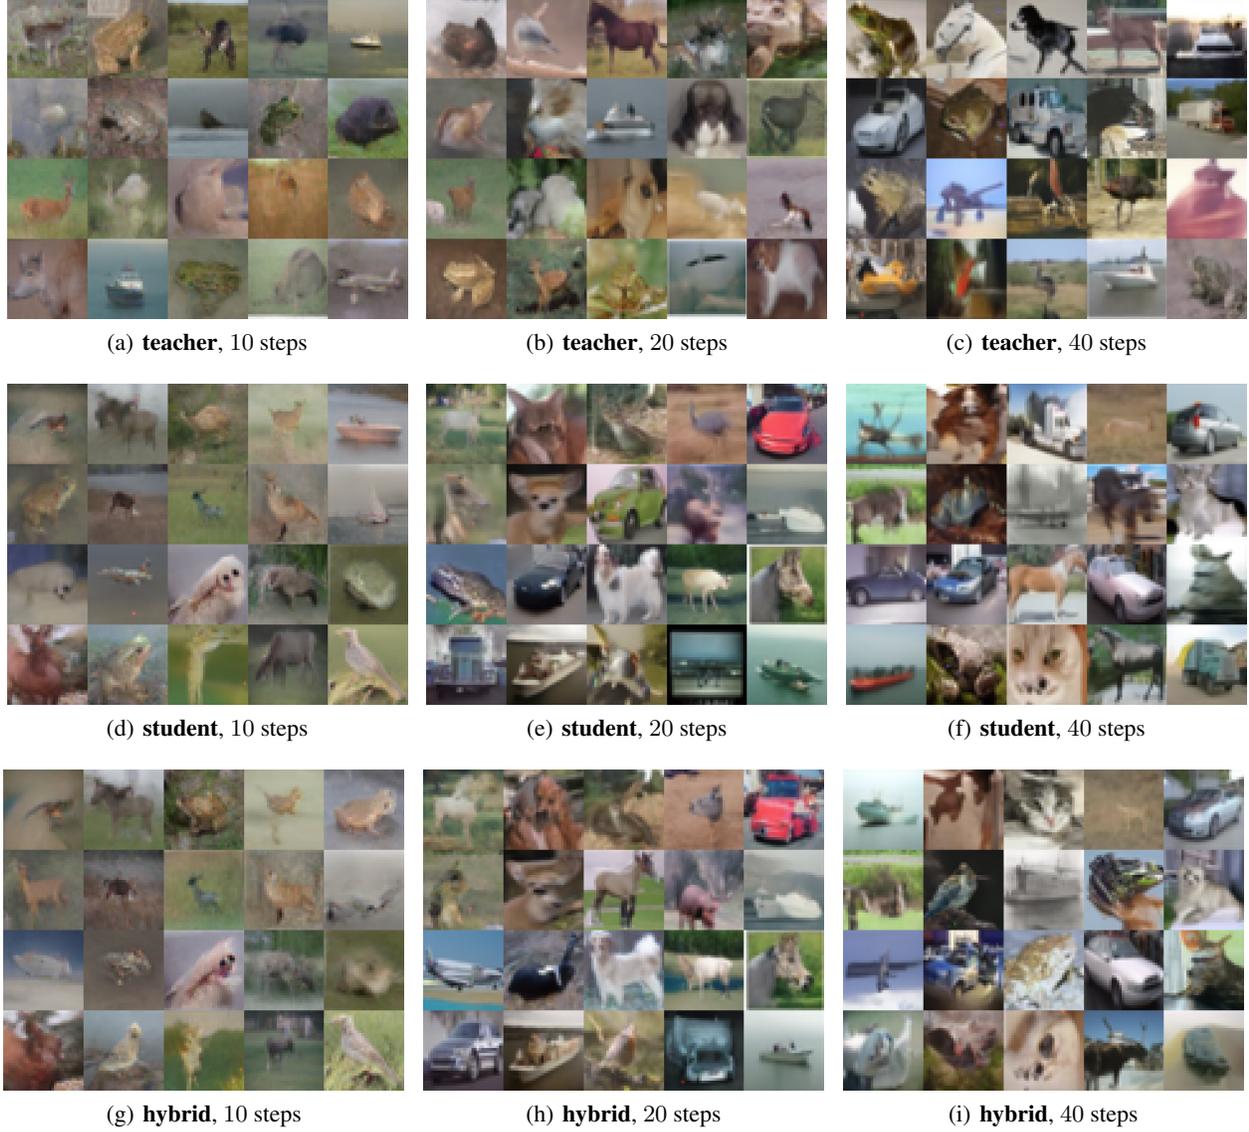


Figure 5. Comparison of generated samples in CIFAR-10 experiment.

3. Let $n(t, s)$ be the number of tokens we unmask in this single sampling step from t to s . Let $d_{n(t,s)}^* \notin M_t$ be the index with the $n(t, s)$ -th largest $\text{conf}(d)$. We define $x_s = (x_s^d)_{d=1}^D$ as follows:

$$x_s^d = \begin{cases} \tilde{x}_0^d & \text{for } d \notin M_t \text{ with } \text{conf}(d) \geq \text{conf}(d_{n(t,s)}^*), \\ x_t^d & \text{otherwise.} \end{cases}$$

Note that $\text{conf}(d)$ coincides with probability zero thanks to the Gumbel noise.

This is for one step of confidence-based denoising. When we use a mixture model, we first sample λ and conduct the confidence-based sampling for the product model conditioned by λ . In the actual sampling process, we specify n (the number of unmasked tokens) as we explain below².

Let N be the number of sampling steps and $t_i = i/N$ for $i = 0, \dots, N$. We jump from t_i to t_{i-1} at each sampling step

²The definition of $n(t_i, t_{i-1})$ described below is different from the implementation of MaskGIT-PyTorch, since the original implementation did not correspond to forward processes.

for $i = N, N - 1, \dots, 1$, so we set $t = t_i$ and $s = t_{i-1}$. Let m_t be the masking probability at time t , so that the expected number of [MASK] tokens equals Dm_t given the forward process design, where $m_0 = 0$ and $m_1 = 1$. By rounding them into integers, we define $n(t_i, t_{i-1}) = \text{round}(Dm_{t_i}) - \text{round}(Dm_{t_{i-1}})$ for each i . In the code, we further add 1 when $n(t_i, t_{i-1}) = 0$ and deduct the added ones from $n(t_1, t_0)$ so that at least one token gets unmasked, following the original implementation.

F.2.3. DISCRETE CLASSIFIER-FREE GUIDANCE

Following the original implementation, we use the discrete classifier-free guidance (discrete CFG; Tang et al., 2022). Given an unconditional denoiser $p_{0|t}(\cdot|\mathbf{x}_t)$ and a conditional denoiser $p_{0|t}(\cdot|\mathbf{x}_t, c)$ with a class label c , we sample from the distribution

$$p_{0|t}[w](\cdot|\mathbf{x}_t, c) \propto p_{0|t}(\cdot|\mathbf{x}_t, c)^{1+w} p_{0|t}(\cdot|\mathbf{x}_t)^{-w}$$

for a CFG guidance scale w ; $w = 0$ corresponds to sampling from the conditional denoiser. In our implementation, given the number of steps N , we let $w = w_{\text{cfg}} \cdot \frac{(1-t)N}{N-1}$ to linearly increase the guidance scale in the sampling process, where w_{cfg} is a user-selected CFG coefficient, mentioned in Section 5.2. When we use a mixture model, we first sample a single λ and sample from the CFG-guided distribution given by the λ -conditioned unconditional/conditional denoisers $p_{0|t}(\cdot|\mathbf{x}_t; \lambda)$ and $p_{0|t}(\cdot|\mathbf{x}_t, c; \lambda)$ as a heuristic.

F.2.4. IMPLEMENTATION AND TRAINING

Network architecture. For the teacher model p^ψ , we just used the implementation of Besnier & Chen (2023, Table 1), which uses a 24-layer transformer with 16 attention heads to compute the logits for each of $D = 256$ visual tokens. As inputs, p^ψ gets the D visual tokens $\mathbf{x}_t = (x_t^d)_{d=1}^D$ and a class token c (when unconditional, it is replaced by an ‘‘unconditional token’’), and these $D + 1$ tokens are added positional encodings. Note that time conditioning is not fed to the model. Each learned embedding (of 1024 VQGAN codebook elements and 1000 ImageNet classes) is of 768 dimensions. Since the output of the transformer has $D + 1$ vectors in \mathbb{R}^{768} , we discard the vector corresponding to c and compute the logits by using the similarity with the embeddings of codebook elements.

To realize a mixture model $p^\theta(\cdot|\mathbf{x}_t; \lambda)$ upon this implementation, we first sample $\lambda \sim \text{Unif}([0, 1])$, pass it to the timestep embedding used in the CIFAR-10 experiment (Appendix F.1.3), and transform its output into a 768-dimensional vector using a two-layer MLP with a single GELU activation after the first layer. Thus, we obtain $D + 2$ vectors in \mathbb{R}^{768} consisting of D visual embeddings, a class embedding, and an embedding of λ . We simply feed them to a transformer having the same architecture as p^ψ and discard the final output vectors concerning c and λ to calculate the logits. The logit calculation is done in the same way as the teacher model.

Training. We basically used the same formulation as (66) with a slight modification to reduce the computational burden as follows:

$$\frac{\mathbf{1}_{\{t \leq \Delta t\}}}{\Delta t} \mathcal{L}_{\text{distil}}(\theta; \psi, q_t, t) + \mathbf{1}_{\{t > \Delta t\}} \mathcal{L}_{\text{consis}}(\theta; \psi, q_t, 0, t - \Delta t, t) + \alpha_t \mathcal{L}_{\text{corr}}(\theta; t) + \mathcal{L}_{\text{marginal}}(\theta; \psi, q_t, t), \quad (70)$$

with $\Delta t = 0.05$ and the following details:

- We sampled $t \sim \text{Unif}([0, 1])$ and \mathbf{x}_t according to (68) using \mathbf{x}_0 from data.
- We used $\alpha_t = 0$ for **di4c** and $\alpha_t = 0.1 \cdot g(t)$ for **di4c-d** in the experiment (Figure 3), where $g(t)$ is given in (67).
- Control variates were used in both variants of Di4C.

For optimization, we followed the original implementation, i.e., the AdamW optimizer with a learning rate 10^{-5} , $(\beta_1, \beta_2) = (0.9, 0.96)$ and a weight decay 10^{-5} . The teacher model was trained for 300 epochs with a minibatch size of 512 using eight A100 GPUs (adding up to 768 GPU hours; Besnier & Chen, 2023). Our finetuning used two A6000 GPUs with a minibatch size 4 (2 for each of two GPUs) and a λ -batch size of 32. It was trained for 30K iterations (so only 120K out of 1.28M ImageNet training images were used), which amounts to approximately 50 GPU hours.

Evaluation. In each experiment, we generated 50,000 samples (50 images for each ImageNet class) and then computed FID (against test data) and IS with the original implementation of Besnier & Chen (2023).

F.2.5. GENERATED SAMPLES

Samples generated by each model in Section 5.2 are shown in Figure 6. The best CFG coefficient in Figure 3 in terms of FID was chosen for each model, as shown in Table 8. The table also shows the FID/IS performance, together with the Precision/Recall values computed by the code of Besnier & Chen (2023). Regarding the choice of ImageNet labels (Ostrich, Burger, Volcano), we followed Besnier & Chen (2023, Figure 3).

Table 8. Chosen CFG coefficients and corresponding FID/IS values.

Model	# steps	w_{cfg}	FID (\downarrow)	IS (\uparrow)	Precision (\uparrow)	Recall (\uparrow)
teacher	8	3.0	6.57	202.0	0.7939	0.5499
teacher	4	11.0	7.97	216.0	0.7737	0.5057
di4c	4	6.0	6.79	209.2	0.7910	0.5363
di4c-d	4	7.0	6.57	213.6	0.7866	0.5391

F.3. Masked diffusion language models

F.3.1. DIFFUSION MODELING AND SAMPLING

Diffusion modeling. We basically worked under the same setting as in Section F.2.4 except for the following:

- $D = 1024$.
- The non-mask codebook \mathcal{S}^* is given by the set of GPT-2 tokenizers, which satisfies $|\mathcal{S}^*| = 50257$.
- The masking probability is given by $m_t = t$.

While the above apparently linear noise scheduling is called *log-linear* (Sahoo et al., 2024, Section E.1), the naming comes from their parametrization: $m_t = 1 - e^{-\sigma(t)}$ with $\sigma(t) = -\log(1 - t)$.

Analytical sampling in masked diffusion. The sampling algorithm used for the MDLMs is essentially the same as the analytical sampling in Section F.1.1, tailored for masked diffusions. What we have is a dimensionally independent denoiser $p_{0|t}(\mathbf{x}_0|\mathbf{x}_t) = \prod_{d=1}^D p_{0|t}^d(x_0^d|\mathbf{x}_t)$. By using (42), we can deduce the resulting *product* denoiser $p_{s|t}$ based on $p_{0|t}$ as follows:

$$\begin{aligned}
 p_{s|t}(\mathbf{x}_s|\mathbf{x}_t) &= \sum_{\mathbf{x}_0} q_{s|0,t}(\mathbf{x}_s|\mathbf{x}_0, \mathbf{x}_t) p_{0|t}(\mathbf{x}_0|\mathbf{x}_t) = \sum_{\mathbf{x}_0} \frac{q_{s|0}(\mathbf{x}_s|\mathbf{x}_0) q_{t|s}(\mathbf{x}_t|\mathbf{x}_s)}{q_{t|0}(\mathbf{x}_t|\mathbf{x}_0)} p_{0|t}(\mathbf{x}_0|\mathbf{x}_t) \\
 &= \sum_{\mathbf{x}_0} \prod_{d=1}^D \frac{q_{s|0}^d(x_s^d|x_0^d) q_{t|s}^d(x_t^d|x_s^d)}{q_{t|0}^d(x_t^d|x_0^d)} p_{0|t}^d(x_0^d|\mathbf{x}_t) \\
 &= \prod_{d=1}^D \sum_{x_0^d} \frac{q_{s|0}^d(x_s^d|x_0^d) q_{t|s}^d(x_t^d|x_s^d)}{q_{t|0}^d(x_t^d|x_0^d)} p_{0|t}^d(x_0^d|\mathbf{x}_t) =: \prod_{d=1}^D p_{s|t}^d(x_s^d|\mathbf{x}_t). \tag{71}
 \end{aligned}$$

Note that, once a token (dimension) is unmasked, we do not need to further change that token in the backward process: this property is incorporated in $p_{0|t}$ as $p_{0|t}^d(x_t^d|\mathbf{x}_t) = 1$ for $x_t^d \neq [\text{MASK}]$ (Sahoo et al., 2024, Section 3.2.3). Thus, we just need to consider the case $x_t^d = [\text{MASK}]$. By using the fact that the masking probability is given by $m_t = t$ and $q_{t|s}^d([\text{MASK}]|x) = \frac{m_t - m_s}{1 - m_s}$ for $x \neq [\text{MASK}]$ in (69), if $x_t^d = [\text{MASK}]$ and $x \neq [\text{MASK}]$, we simply have

$$p_{s|t}^d(x|\mathbf{x}_t) = \frac{q_{s|0}^d(x|x) q_{t|s}^d([\text{MASK}]|x)}{q_{t|0}^d([\text{MASK}]|x)} p_{0|t}^d(x|\mathbf{x}_t) = \frac{(1 - m_s) \frac{m_t - m_s}{1 - m_s}}{m_t} p_{0|t}^d(x|\mathbf{x}_t) = \frac{m_t - m_s}{m_t} p_{0|t}^d(x|\mathbf{x}_t),$$

which corresponds to Sahoo et al. (2024, Eq. 7).



Figure 6. Comparison of generated samples. (a)–(d): Conditioned with ImageNet label Ostrich (009). (e)–(h): Conditioned with Burger (933). (i)–(l): Conditioned with Volcano (980).

F.3.2. IMPLEMENTATION AND TRAINING

Network architecture. For the teacher model, we used the two “small” checkpoints (round 6 and 7) of Deschenaux & Gulcehre (2025)³, which uses a transformer architecture to compute the logits of $p_{0|t}(\cdot|\mathbf{x}_t)$ for each token. The transformer architecture is originally from (Sahoo et al., 2024), and has 169M parameters with 12 layers, the embedding dimension of 768, and 12 attention heads. The model receives $D = 1024$ tokens and does not depend on the timestep. For adaptation to mixture modeling, we just applied the same modification as given in Section F.2.4.

Training. While the model accepts a continuous time training, we followed Deschenaux & Gulcehre (2025) to digitize the timesteps to $\mathbb{T} = \{\Delta t \cdot n \mid n = 0, 1, 2, 3, \dots, 1024\}$, with $\Delta t = 1/1024$. We used the following loss function:

$$\mathbf{1}_{\{t \leq \delta\}} \mathcal{L}_{\text{distil}}(\theta; \psi, q_t, t) + \mathbf{1}_{\{t > \delta\}} \mathcal{L}_{\text{consis}}(\theta; \psi, q_t, 0, t - \Delta t, t) + \alpha_t \mathcal{L}_{\text{corr}}(\theta; t) + \mathcal{L}_{\text{marginal}}(\theta; \psi, q_t, t), \quad (72)$$

where the details are as follows:

- We set $\delta = 0.02$.
- We sample $r \sim \text{Unif}([0, 1])$ and let

$$t = 2^{-10}(1 + \lfloor 2^{10} t^*(r) \rfloor), \quad \text{where } t^*(r) = \begin{cases} \delta r & \text{with probability 0.3,} \\ \delta + (1 - \delta)r & \text{with probability 0.7.} \end{cases}$$

Therefore, $t \in \mathbb{T} \setminus \{0\}$ almost surely.

- We sampled \mathbf{x}_t according to (68) using \mathbf{x}_0 from data (OpenWebText).
- We used $\alpha_t = 0$ in the experimental results shown in Section 5.3. In Section F.3.3, we also report the result of setting $\alpha_t = 0.1 \cdot g(t)$ with $g(t)$ from (67).
- Control variates were used in all the experiments.

For training, we mostly followed the original setting of Deschenaux & Gulcehre (2025): We used the Adam optimizer (but with a learning rate of 3×10^{-5}) with EMA (weight decay 0.9999) and did a constant warm-up (increasing the learning rate linearly for the first 500 iterations and setting it constant after that). For each experiment, the Di4C training (one round) was run for 100K iterations over 2x A6000 GPUs, where the minibatch size was 2 (1 for each device) and λ -batch size was 16.

Self-BLEU computation. As described in Section 5.3, in the conditional generation experiment, we generated $M = 5$ continuations $C^{(i)} = \{X_1^{(i)}, \dots, X_5^{(i)}\}$ conditioned on the first 50 tokens (prompt) of each WebText datapoint $X^{(i)}$. Each continuation was of 100 tokens including the prompt, and we used $M = 256$ prompts from the WebText dataset in total. To quantify the diversity of continuations, we followed Deschenaux & Gulcehre (2025) and computed the Self-BLEU score as

$$\frac{1}{N} \sum_{i=1}^N \frac{1}{M} \sum_{j=1}^M \text{BLEU}(X_j^{(i)}; C^{(i)} \setminus X_j^{(i)}),$$

where $\text{BLEU}(X; C)$ is the BLEU score of a sentence X against the set of reference sentences C . To actually compute this, we tokenized the inputs with the GPT-2 tokenizer, and then utilized the implementation of Zhu et al. (2018)⁴ with `ngram = 4`, which internally calls the `sentence_bleu` function of the NLTK library (Bird et al., 2009)⁵ with equal weighting and the `method1` smoothing function.

Evaluation. Except for the Self-BLEU computation, all the evaluations were done using the code of SDTT (Deschenaux & Gulcehre, 2025)⁶.

³It is loaded by `load_small_student(loss='kld', round=n)` with $n = 6, 7$, from the library `sdt` in <https://github.com/jdeschena/sdt>.

⁴<https://github.com/geek-ai/Texygen/blob/master/utils/metrics/SelfBleu.py>.

⁵https://github.com/nltk/nltk/blob/3.7/nltk/translate/bleu_score.py.

⁶<https://github.com/jdeschena/sdt/tree/main>.

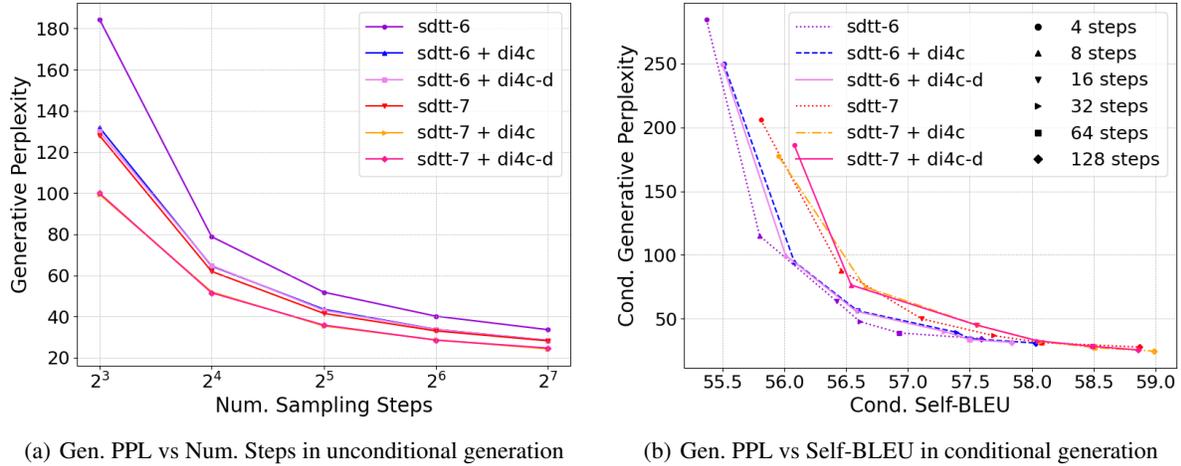


Figure 7. Comparison of Di4C distillations of SDTT checkpoints with and without data loss.

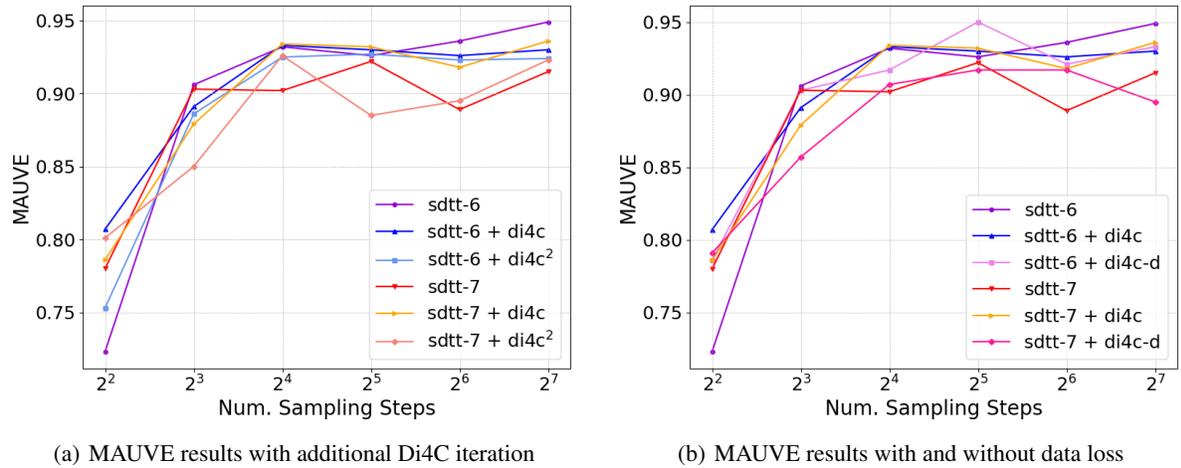


Figure 8. Comparison of Di4C distillations of SDTT checkpoints with and without data loss.

F.3.3. ADDITIONAL EXPERIMENTAL RESULTS

Data loss ablation. In Figure 7, we show the result for **sdtt-6/7 + di4c-d**, which used loss (72) with $\alpha_t = 0.1g(t)$ as mentioned in Section F.3.2. The results are almost the same as those without data loss (especially in Figure 7, where **sdtt-7 + di4c** is hidden behind the curve of **sdtt-7 + di4c-d**). Since the ones without data loss showed slightly better generative perplexities, we presented them as the main model in Section 5.3.

MAUVE results. We also tested our models with the MAUVE score (Pillutla et al., 2021). The setting is the same as the unconditional generation in Section 5.3, and the MAUVE computation is done by using the code of SDTT (Deschenaux & Gulcehre, 2025). As shown in Figure 8, no significant performance decay from the teacher model was observed.

F.3.4. GENERATED SAMPLES

Let us qualitatively compare our best model (**sdtt-7 + di4c²**) with the SDTT checkpoint which our model is based on (**sdtt-7**), conditioned on the first 50 tokens from this paper’s abstract (old version; highlighted in blue):

- (**sdtt-7**, 4 steps) Diffusion models have demonstrated exceptional performances in various fields of generative modeling. While they often outperform competitors including VAEs and GANs in sample quality and diversity, they suffer from slow sampling speed due to their iterative nature. Recently, districting the image spl Image\n\ndistriction the size spl

Image size The stable image spl consists of thin-arr 2D-frame data that supports band- construction data and generates the physical representation of image presentation. Although much is known

- **(sdtt-7 + di4c², 4 steps)** Diffusion models have demonstrated exceptional performances in various fields of generative modeling. While they often outperform competitors including VAEs and GANs in sample quality and diversity, they suffer from slow sampling speed due to their iterative nature. Recently, distilling the high spl Image Image \n\ndistribution the whole spl Image Image generating an image Image is a thin-narr 2D-layer architecture that exhibits banding-resolving effects in the lowizing of a low resolution. Little is known
- **(sdtt-7, 16 steps)** Diffusion models have demonstrated exceptional performances in various fields of generative modeling. While they often outperform competitors including VAEs and GANs in sample quality and diversity, they suffer from slow sampling speed due to their iterative nature. Recently, distillilusion models, which use a dataset from the sample and using inferred model data, have become standard. self-arrained three-block sampling that utilizes a subset of the data and combining Bayesianesian and inferred model data that combines Bayesian
- **(sdtt-7 + di4c², 16 steps)** Diffusion models have demonstrated exceptional performances in various fields of generative modeling. While they often outperform competitors including VAEs and GANs in sample quality and diversity, they suffer from slow sampling speed due to their iterative nature. Recently, distorting Diffusion models, nested within a conventional model architecture, and using different model architectures, has become a flexible self-arrative three-model architecture that supports intensive problem-solving and mature Bayesianesian model and model development. This approach has
- **(sdtt-7, 64 steps)** Diffusion models have demonstrated exceptional performances in various fields of generative modeling. While they often outperform competitors including VAEs and GANs in sample quality and diversity, they suffer from slow sampling speed due to their iterative nature. Recently, distancing from the Diffusion model has allowed developers to construct new models using rapid-processing, supervised learning supervised (GCI supervised) software-drawing that improves the ability to identify discriminant parameters, functional model depth, and is the process of rapidly
- **(sdtt-7 + di4c², 64 steps)** Diffusion models have demonstrated exceptional performances in various fields of generative modeling. While they often outperform competitors including VAEs and GANs in sample quality and diversity, they suffer from slow sampling speed due to their iterative nature. Recently, distancing from the Diffusion model has allowed scientists to analyze the model using discriml, a software that utilizes to discriml. This software can generate images that offset the time to be discrimlative, and reduce the time to be discrimlative.
- **(sdtt-7, 256 steps)** Diffusion models have demonstrated exceptional performances in various fields of generative modeling. While they often outperform competitors including VAEs and GANs in sample quality and diversity, they suffer from slow sampling speed due to their iterative nature. Recently, distancing from the traditional resource management process has led to a change in the type of user experience, producing significant advances in the types of software-developing operations, while facilitating the adoption of parallel programming and functional programming approaches. However,, it is increasingly
- **(sdtt-7 + di4c², 256 steps)** Diffusion models have demonstrated exceptional performances in various fields of generative modeling. While they often outperform competitors including VAEs and GANs in sample quality and diversity, they suffer from slow sampling speed due to their iterative nature. Recently, distorting cloud images are in the process of processing to high standards in both sample quality and storage, producing important advances in the development of software-developing operations, while facilitating the development of parallel programming and functional efficiency. GPUs are in the process of processing