
Estimating the Optimal Covariance with Imperfect Mean in Diffusion Probabilistic Models

Fan Bao¹ Chongxuan Li² Jiacheng Sun³ Jun Zhu¹ Bo Zhang¹

Abstract

Diffusion probabilistic models (DPMs) are a class of powerful deep generative models (DGMs). Despite their success, the iterative generation process over the full timesteps is much less efficient than other DGMs such as GANs. Thus, the generation performance on a subset of timesteps is crucial, which is greatly influenced by the covariance design in DPMs. In this work, we consider diagonal and full covariances to improve the expressive power of DPMs. We derive the optimal result for such covariances, and then correct it when the mean of DPMs is imperfect. Both the optimal and the corrected ones can be decomposed into terms of conditional expectations over functions of noise. Building upon it, we propose to estimate the optimal covariance and its correction given imperfect mean by learning these conditional expectations. Our method can be applied to DPMs with both discrete and continuous timesteps. We consider the diagonal covariance in our implementation for computational efficiency. For an efficient practical implementation, we adopt a parameter sharing scheme and a two-stage training process. Empirically, our method outperforms a wide variety of covariance design on likelihood results, and improves the sample quality especially on a small number of timesteps.

1. Introduction

Recently, diffusion probabilistic models (DPMs) (Sohl-Dickstein et al., 2015; Ho et al., 2020; Song et al., 2020b)

¹Dept. of Comp. Sci. & Tech., Institute for AI, Tsinghua-Huawei Joint Center for AI, BNRist Center, State Key Lab for Intell. Tech. & Sys., Tsinghua University ²Gaoling School of AI, Renmin University of China; Beijing Key Lab of Big Data Management and Analysis Methods, Beijing, China ³Huawei Noah’s Ark Lab. Correspondence to: Chongxuan Li <chongxuanli@ruc.edu.cn>, Jun Zhu <dczsj@tsinghua.edu.cn>.

have shown great promise for generative modeling. Such models smoothly inject noise to the data distribution, which forms a diffusion process. By learning to reverse the process using a Markov model, DPMs are able to generate high quality images (Ho et al., 2020; Song et al., 2020b; Dhariwal & Nichol, 2021) and audios (Chen et al., 2020; Kong et al., 2020), which are comparable or even superior than the current state-of-the-art generative models (Brock et al., 2018; Wu et al., 2019; Karras et al., 2020; Bińkowski et al., 2019; Kalchbrenner et al., 2018).

However, the iterative generation over the full timesteps of DPMs makes them much less efficient than generative models such as GANs (Goodfellow et al., 2014). Thus, the generation performance on a subset of timesteps is crucial. In this case, the transition of the reversed diffusion process becomes more complex (Xiao et al., 2021), and the covariance design in DPMs matters (Nichol & Dhariwal, 2021; Bao et al., 2022). Most prior works (Ho et al., 2020; Song et al., 2020a; Bao et al., 2022) use an isotropic covariance that only depends on the timestep without considering the state. A notable recent progress is Analytic-DPM (Bao et al., 2022), which estimates the optimal isotropic covariance (in the sense of maximum likelihood) instead of using handcrafted values (Ho et al., 2020; Song et al., 2020a) and shows significant improvement on likelihood estimation and sampling efficiency. However, the isotropic covariance sacrifices the expressive power of DPMs for simplicity. Moreover, the optimality of the covariance estimate in Bao et al. (2022) holds by assuming that the optimal mean is known, which is not the case in practice.

To overcome the aforementioned limitations and further improve DPMs for likelihood estimation and sampling efficiency, we consider diagonal and full covariances to improve the expressive power of DPMs. We derive the optimal mean and covariance from the perspective of the maximum likelihood. We also correct the optimal covariance given an imperfect mean (i.e., considering the approximation and optimization errors) in terms of maximum likelihood. Both the optimal and the corrected ones can be decomposed into terms of conditional expectations over functions of noise, which can be estimated by minimizing mean squared error (MSE) losses. Although our theory applies to the full covariance case, in our implementation, we consider the diagonal

covariance for computational efficiency. Besides, we adopt a parameter sharing scheme for inference efficiency and a two-stage training process motivated by our theory.

Our method is applicable to DPMs with both discrete (Ho et al., 2020) and continuous (Song et al., 2020b) timesteps. In our experiments, we directly compare our method with a variety of baselines (Ho et al., 2020; Song et al., 2020a; Bao et al., 2022; Song et al., 2020b) in terms of sample quality and likelihood estimation in DPMs with both discrete and continuous timesteps. At nearly the same computation cost, our method consistently outperforms these baselines on likelihood estimation. Besides, our method also outperforms these baselines on the sample quality in most cases, especially when the number of timesteps is small.

2. Background

Diffusion probabilistic models (DPMs) are special Markov models with Gaussian transitions:

$$p(\mathbf{x}_{0:N}) = p(\mathbf{x}_N) \prod_{n=1}^N p(\mathbf{x}_{n-1}|\mathbf{x}_n), \quad (1)$$

$$p(\mathbf{x}_{n-1}|\mathbf{x}_n) = \mathcal{N}(\mathbf{x}_{n-1}|\boldsymbol{\mu}_n(\mathbf{x}_n), \boldsymbol{\Sigma}_n(\mathbf{x}_n)),$$

which aim to reverse a forward process $q(\mathbf{x}_{1:N}|\mathbf{x}_0)$ that gradually injects noise to a data distribution $q(\mathbf{x}_0)$. Song et al. (2020a) consider a family of forward processes indexed by a non-negative vector $\lambda = (\lambda_1, \dots, \lambda_N) \in \mathbb{R}_{\geq 0}^N$:

$$q(\mathbf{x}_{1:N}|\mathbf{x}_0) = q(\mathbf{x}_N|\mathbf{x}_0) \prod_{n=2}^N q(\mathbf{x}_{n-1}|\mathbf{x}_n, \mathbf{x}_0), \quad (2)$$

$$q(\mathbf{x}_N|\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_N|\sqrt{\bar{\alpha}_N}\mathbf{x}_0, \bar{\beta}_N\mathbf{I}),$$

$$q(\mathbf{x}_{n-1}|\mathbf{x}_n, \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_{n-1}|\tilde{\boldsymbol{\mu}}_n(\mathbf{x}_n, \mathbf{x}_0), \lambda_n^2\mathbf{I}),$$

$$\tilde{\boldsymbol{\mu}}_n(\mathbf{x}_n, \mathbf{x}_0) = \sqrt{\bar{\alpha}_{n-1}}\mathbf{x}_0 + \sqrt{\bar{\beta}_{n-1} - \lambda_n^2} \cdot \frac{\mathbf{x}_n - \sqrt{\bar{\alpha}_n}\mathbf{x}_0}{\sqrt{\bar{\beta}_n}},$$

where $\bar{\alpha}_1, \bar{\alpha}_2, \dots, \bar{\alpha}_N \in (0, 1)$ is a strictly decreasing sequence, $\bar{\beta}_n := 1 - \bar{\alpha}_n$ and \mathbf{I} is the identity matrix. Observing that $q(\mathbf{x}_n|\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_n|\sqrt{\bar{\alpha}_n}\mathbf{x}_0, \bar{\beta}_n\mathbf{I})$, we can quickly sample \mathbf{x}_n given \mathbf{x}_0 by

$$\mathbf{x}_n = \sqrt{\bar{\alpha}_n}\mathbf{x}_0 + \sqrt{\bar{\beta}_n}\boldsymbol{\epsilon}_n, \quad \boldsymbol{\epsilon}_n \sim \mathcal{N}(\mathbf{0}, \mathbf{I}).$$

Let $\alpha_n := \bar{\alpha}_n/\bar{\alpha}_{n-1}$, $\beta_n := 1 - \alpha_n$ and $\tilde{\beta}_n := \frac{\bar{\beta}_{n-1}}{\bar{\beta}_n}\beta_n$. Two commonly used forward processes are the denoising diffusion probabilistic model (DDPM) forward process (corresponding to $\lambda_n^2 = \tilde{\beta}_n$) and the denoising diffusion implicit model (DDIM) forward process (corresponding to $\lambda_n^2 = 0$). In particular, the DDPM forward process is Markovian with linear Gaussian transition:

$$q(\mathbf{x}_n|\mathbf{x}_{n-1}) = \mathcal{N}(\mathbf{x}_n|\sqrt{\bar{\alpha}_n}\mathbf{x}_{n-1}, \beta_n\mathbf{I}).$$

The reverse process is learned by maximizing the evidence lower bound (ELBO) $L_{\text{elbo}} = \mathbb{E}_q \log \frac{p(\mathbf{x}_{0:N})}{q(\mathbf{x}_{1:N}|\mathbf{x}_0)}$ on log-likelihood, or equivalently, by minimizing the KL divergence between the forward and the reverse process

$$\max_{\{\boldsymbol{\mu}_n, \boldsymbol{\Sigma}_n\}_{n=1}^N} L_{\text{elbo}} \Leftrightarrow \min_{\{\boldsymbol{\mu}_n, \boldsymbol{\Sigma}_n\}_{n=1}^N} D_{\text{KL}}(q(\mathbf{x}_{0:N})\|p(\mathbf{x}_{0:N})). \quad (3)$$

Note that problem (3) optimizes the mean and the covariance jointly. Thus, we term it *joint optimization*.

For simplicity and at the cost of the flexibility, prior works (Ho et al., 2020; Song et al., 2020a; Bao et al., 2022) set $\boldsymbol{\Sigma}_n(\mathbf{x}_n) = \sigma_n^2\mathbf{I}$ to an isotropic covariance that only depends on the timestep n , where σ_n^2 is the variance for each component. Under this restriction, Analytic-DPM (Bao et al., 2022) shows that both the optimal mean $\boldsymbol{\mu}_n^*(\mathbf{x}_n)$ and variance σ_n^{*2} have analytic forms w.r.t. the conditional expectation of the noise¹ $\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]$:

$$\boldsymbol{\mu}_n^*(\mathbf{x}_n) = \tilde{\boldsymbol{\mu}}_n\left(\mathbf{x}_n, \frac{1}{\sqrt{\bar{\alpha}_n}}(\mathbf{x}_n - \sqrt{\bar{\beta}_n}\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n])\right), \quad (4)$$

$$\sigma_n^{*2} = \lambda_n^2 + \gamma_n^2 \frac{\bar{\beta}_n}{\bar{\alpha}_n} \left(1 - \mathbb{E}_{q(\mathbf{x}_n)} \frac{\|\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]\|_2^2}{d}\right), \quad (5)$$

where $\boldsymbol{\epsilon}_n = \frac{\mathbf{x}_n - \sqrt{\bar{\alpha}_n}\mathbf{x}_0}{\sqrt{\bar{\beta}_n}}$ is the noise used to generate \mathbf{x}_n from \mathbf{x}_0 , d is the dimension of the data \mathbf{x}_0 and $\gamma_n = \sqrt{\bar{\alpha}_{n-1} - \bar{\beta}_{n-1} - \lambda_n^2} \sqrt{\frac{\bar{\alpha}_n}{\bar{\beta}_n}}$. Ho et al. (2020) estimate $\boldsymbol{\mu}_n^*(\mathbf{x}_n)$ by a noise prediction network $\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)$

$$\hat{\boldsymbol{\mu}}_n(\mathbf{x}_n) = \tilde{\boldsymbol{\mu}}_n\left(\mathbf{x}_n, \frac{1}{\sqrt{\bar{\alpha}_n}}(\mathbf{x}_n - \sqrt{\bar{\beta}_n}\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n))\right). \quad (6)$$

Here $\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)$ aims to learn $\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]$ by minimizing the following MSE loss:

$$\min_{\{\hat{\boldsymbol{\epsilon}}_n\}_{n=1}^N} \mathbb{E}_n \mathbb{E}_{q(\mathbf{x}_0, \mathbf{x}_n)} \|\boldsymbol{\epsilon}_n - \hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)\|_2^2, \quad (7)$$

where n is uniformly sampled from $\{1, 2, \dots, N\}$. Eq. (7) admits an optimal solution when $\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n) = \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]$ for all $n \in \{1, 2, \dots, N\}$.

The optimal variance in Eq. (5) can also be estimated using $\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)$ as following (Bao et al., 2022)

$$\hat{\sigma}_n^2 = \lambda_n^2 + \gamma_n^2 \frac{\bar{\beta}_n}{\bar{\alpha}_n} \left(1 - \mathbb{E}_{q(\mathbf{x}_n)} \frac{\|\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)\|_2^2}{d}\right). \quad (8)$$

¹The original work (Bao et al., 2022) uses the score function $\nabla_{\mathbf{x}_n} \log q(\mathbf{x}_n)$ to represent $\boldsymbol{\mu}_n^*(\mathbf{x}_n)$ and σ_n^{*2} , which is equivalent to the conditional expectation of the noise up to multiplying a constant: $\nabla_{\mathbf{x}_n} \log q(\mathbf{x}_n) = -\frac{1}{\sqrt{\bar{\beta}_n}} \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]$.

Forward process	$\bar{\alpha}_n, \bar{\beta}_n$	represent the cumulative amount of Gaussian noise at timestep n , i.e., $q(\mathbf{x}_n \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_n \sqrt{\bar{\alpha}_n}\mathbf{x}_0, \bar{\beta}_n\mathbf{I})$.
	λ_n^2	represents the variance of the transition $q(\mathbf{x}_{n-1} \mathbf{x}_n, \mathbf{x}_0)$.
	$\tilde{\boldsymbol{\mu}}_n(\mathbf{x}_n, \mathbf{x}_0)$	represents the mean of the transition $q(\mathbf{x}_{n-1} \mathbf{x}_n, \mathbf{x}_0)$.
	γ_n	represents the coefficient of \mathbf{x}_0 in $\tilde{\boldsymbol{\mu}}_n(\mathbf{x}_n, \mathbf{x}_0)$.
Reverse mean	$\boldsymbol{\mu}_n(\mathbf{x}_n), \boldsymbol{\mu}_n^*(\mathbf{x}_n)$	represent the mean of $p(\mathbf{x}_{n-1} \mathbf{x}_n)$ and the optimal one.
	$\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n), \hat{\boldsymbol{\mu}}_n(\mathbf{x}_n)$	represent the noise prediction network and the estimate of $\boldsymbol{\mu}_n^*(\mathbf{x}_n)$.
Reverse covariance	$\boldsymbol{\Sigma}_n(\mathbf{x}_n), \boldsymbol{\sigma}_n^*(\mathbf{x}_n)^2$	represent the covariance of $p(\mathbf{x}_{n-1} \mathbf{x}_n)$ and the optimal diagonal one.
	$\mathbf{h}_n(\mathbf{x}_n), \hat{\boldsymbol{\sigma}}_n(\mathbf{x}_n)^2$	represent the SN prediction network and the estimate of $\boldsymbol{\sigma}_n^*(\mathbf{x}_n)^2$.
	$\tilde{\boldsymbol{\sigma}}_n^*(\mathbf{x}_n)^2$	represents the corrected optimal diagonal covariance with an imperfect mean.
	$\mathbf{g}_n(\mathbf{x}_n), \hat{\tilde{\boldsymbol{\sigma}}}_n^*(\mathbf{x}_n)^2$	represent the NPR prediction network and the estimate of $\tilde{\boldsymbol{\sigma}}_n^*(\mathbf{x}_n)^2$.

Table 1. Notations used in this paper.

3. Method

We improve the covariance estimate of DPMs building upon the recent progress (Bao et al., 2022). First, we consider diagonal and full covariances instead of using an isotropic one to improve the expressive power of DPMs and obtain the expression of the optimal solution, as detailed in Section 3.1. Second, we correct the optimal covariance given an imperfect mean (i.e., with approximation and optimization errors), as presented in Section 3.2. For clarity, we refer the readers to Appendix A for all proofs.

Although our theoretical results are applicable to the full covariance case (see Appendix B), it is often time consuming to obtain samples from Gaussian transitions with full covariances (e.g., via Cholesky decomposition). To balance the flexibility and the time cost, we focus on diagonal covariances throughout the main text including the experiments.

3.1. The Optimal Solution for Covariances beyond Isotropic Ones

Instead of using an isotropic covariance that only depends on the timestep without considering the state, we consider a diagonal covariance to improve the expressive power of DPMs. To be clear and rigorous, we first construct an example where the data distribution is a mixture of Gaussian. Even in such a simple case, we can prove that the optimal ELBO (in Eq. (3)) with a diagonal covariance is *strictly greater* than that with an isotropic covariance. See a formal description and a detailed proof in Proposition A.5.

Formally, we consider the covariance in the form of $\boldsymbol{\Sigma}_n(\mathbf{x}_n) = \text{diag}(\boldsymbol{\sigma}_n(\mathbf{x}_n)^2)$, where $\boldsymbol{\sigma}_n(\mathbf{x}_n)^2 : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is the diagonal of the covariance matrix and $\text{diag}(\cdot)$ denotes

the vector-to-matrix diag operator. Naturally, we optimize problem (3) w.r.t. the mean and covariance jointly and derive the the optimal solution, as summarized in Theorem 3.1. See the result for full covariances in Appendix B.

Theorem 3.1. (Optimal solution to joint optimization) *Suppose $\boldsymbol{\Sigma}_n(\mathbf{x}_n) = \text{diag}(\boldsymbol{\sigma}_n(\mathbf{x}_n)^2)$. Then the optimal mean to problem (3) is $\boldsymbol{\mu}_n^*(\mathbf{x}_n)$ as in Eq. (4), and the optimal covariance to problem (3) is*

$$\boldsymbol{\sigma}_n^*(\mathbf{x}_n)^2 = \lambda_n^2 \mathbf{1} + \gamma_n^2 \frac{\bar{\beta}_n}{\bar{\alpha}_n} (\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n^2] - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]^2),$$

where $\boldsymbol{\epsilon}_n = \frac{\mathbf{x}_n - \sqrt{\bar{\alpha}_n}\mathbf{x}_0}{\sqrt{\bar{\beta}_n}}$ is the noise used to generate \mathbf{x}_n from \mathbf{x}_0 , $(\cdot)^2$ is the element-wise square, $\mathbf{1}$ is the vector of ones and $\gamma_n = \sqrt{\bar{\alpha}_{n-1}} - \sqrt{\bar{\beta}_{n-1}} - \lambda_n^2 \sqrt{\frac{\bar{\alpha}_n}{\bar{\beta}_n}}$.

The proof idea is similar to Theorem 1 in Bao et al. (2022), which only considers the isotropic covariance case.

To estimate the optimal covariance $\boldsymbol{\sigma}_n^*(\mathbf{x}_n)^2$ in Theorem 3.1, we need to estimate both $\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n^2]$ and $\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]^2$.

As mentioned in Section 2, the noise prediction network $\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)$ aims to estimate $\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]$ by minimizing the MSE loss in Eq. (7). Thus, we use $\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)^2$ to estimate $\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]^2$. Note that the error of $\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)^2$ is likely to be amplified compared to that of $\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)$ (see Appendix E.1 for details). However, we still find it performs well on sample quality (see Section 6.1).

As for $\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n^2]$, we note that it is the expectation of the *squared noise (SN)* $\boldsymbol{\epsilon}_n^2$ conditioned on \mathbf{x}_n . Such a conditional expectation can be learned using a neural network

$\mathbf{h}_n(\mathbf{x}_n) \in \mathbb{R}^d$ trained on the following MSE loss

$$\min_{\{\mathbf{h}_n\}_{n=1}^N} \mathbb{E}_n \mathbb{E}_{q(\mathbf{x}_0, \mathbf{x}_n)} \|\epsilon_n^2 - \mathbf{h}_n(\mathbf{x}_n)\|_2^2, \quad (9)$$

where n is uniformly sampled from $\{1, 2, \dots, N\}$ and \mathbf{h}_n attempts to predict ϵ_n^2 given \mathbf{x}_n . Here we call $\mathbf{h}_n(\mathbf{x}_n)$ the *SN prediction network*. Eq. (9) admits an optimal solution when $\mathbf{h}_n(\mathbf{x}_n) = \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\epsilon_n^2]$ for all $n \in \{1, 2, \dots, N\}$.

By estimating $\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\epsilon_n^2]$ with $\hat{\epsilon}_n(\mathbf{x}_n)^2$ and estimating $\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\epsilon_n^2]$ with $\mathbf{h}_n(\mathbf{x}_n)$, we can estimate $\sigma_n^*(\mathbf{x}_n)^2$ by

$$\hat{\sigma}_n(\mathbf{x}_n)^2 = \lambda_n^2 \mathbf{1} + \gamma_n^2 \frac{\bar{\beta}_n}{\bar{\alpha}_n} (\mathbf{h}_n(\mathbf{x}_n) - \hat{\epsilon}_n(\mathbf{x}_n)^2). \quad (10)$$

In the paper, we refer to the above estimate as *SN-DPM*.

3.2. The Optimal Covariance with an Imperfect Mean

It is plausible to obtain the estimates $\hat{\boldsymbol{\mu}}_n(\mathbf{x}_n)$ and $\hat{\sigma}_n(\mathbf{x}_n)^2$, according to Theorem 3.1. However, we argue that the exact optimal solution $\boldsymbol{\mu}_n^*(\mathbf{x}_n)$, $\sigma_n^*(\mathbf{x}_n)^2$ in Theorem 3.1 cannot be achieved due to the approximation and optimization errors, making it possible to further improve the estimates considering the errors. In fact, below we present how to correct the optimal covariance given a potentially imperfect mean $\hat{\boldsymbol{\mu}}_n$.

We note that, in Theorem 3.1, the optimal mean is irrelevant to the covariance, while the optimal covariance is expressed by the optimal mean. Naturally, we can first optimize the mean solely to obtain the mean estimate $\hat{\boldsymbol{\mu}}_n$ and then optimize the covariance solely given $\hat{\boldsymbol{\mu}}_n$. Such a two-stage approach is at least not worse than the estimate in Section 3.1 under the maximum likelihood, because we have

$$L_{\text{elbo}}(\hat{\boldsymbol{\mu}}_n, \hat{\sigma}_n(\cdot)^2) \leq \max_{\sigma_n(\cdot)^2} L_{\text{elbo}}(\hat{\boldsymbol{\mu}}_n, \sigma_n(\cdot)^2), \quad (11)$$

where the equality does not hold in general.

To formalize the idea, we obtain the optimal solutions to the following two problems

$$\text{Given arbitrary covariance } \Sigma_n, \max_{\{\boldsymbol{\mu}_n\}_{n=1}^N} L_{\text{elbo}}, \quad (12)$$

$$\text{Given arbitrary mean } \boldsymbol{\mu}_n, \max_{\{\Sigma_n\}_{n=1}^N} L_{\text{elbo}}, \quad (13)$$

as summarized in Theorem 3.2 and Theorem 3.3.

Theorem 3.2. (Optimal solution to optimization solely w.r.t. mean) For any covariance Σ_n , the optimal mean to problem (12) is always Eq. (4), i.e.,

$$\boldsymbol{\mu}_n^*(\mathbf{x}_n) = \tilde{\boldsymbol{\mu}}_n \left(\mathbf{x}_n, \frac{1}{\sqrt{\bar{\alpha}_n}} (\mathbf{x}_n - \sqrt{\bar{\beta}_n} \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\epsilon_n]) \right),$$

which is irrelevant to Σ_n .

Theorem 3.2 reveals that the optimal mean doesn't depend on the covariance and therefore we can learn the mean on Eq. (7) without knowing information about the covariance.

In contrast to the mean, the optimal covariance to problem (13) depends on the mean, as shown in Theorem 3.3. See the result for full covariances in Appendix B.

Theorem 3.3. (Optimal solution to optimization w.r.t. covariance solely) Suppose $\Sigma_n(\mathbf{x}_n) = \text{diag}(\sigma_n(\mathbf{x}_n)^2)$. For any mean $\boldsymbol{\mu}_n(\mathbf{x}_n)$ that is parameterized by a noise prediction network $\hat{\epsilon}_n(\mathbf{x}_n)$ as in Eq. (6), the optimal covariance $\tilde{\sigma}_n^*(\mathbf{x}_n)^2$ to problem (13) is

$$\begin{aligned} \tilde{\sigma}_n^*(\mathbf{x}_n)^2 &= \sigma_n^*(\mathbf{x}_n)^2 + \gamma_n^2 \frac{\bar{\beta}_n}{\bar{\alpha}_n} \underbrace{(\hat{\epsilon}_n(\mathbf{x}_n) - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\epsilon_n])^2}_{\text{error}} \\ &= \lambda_n^2 \mathbf{1} + \gamma_n^2 \frac{\bar{\beta}_n}{\bar{\alpha}_n} \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)} [(\epsilon_n - \hat{\epsilon}_n(\mathbf{x}_n))^2], \end{aligned} \quad (14)$$

where $\sigma_n^*(\mathbf{x}_n)^2$ is the optimal covariance to the joint optimization problem in Theorem 3.1, $\epsilon_n = \frac{\mathbf{x}_n - \sqrt{\bar{\alpha}_n} \mathbf{x}_0}{\sqrt{\bar{\beta}_n}}$ is the noise used to generate \mathbf{x}_n from \mathbf{x}_0 , $(\cdot)^2$ is the element-wise square, $\mathbf{1}$ is the vector of ones and $\gamma_n = \sqrt{\bar{\alpha}_{n-1} - \sqrt{\bar{\beta}_{n-1} - \lambda_n^2 \sqrt{\frac{\bar{\alpha}_n}{\bar{\beta}_n}}}}$.

The proof of Theorem 3.3 mainly builds on this fact: minimizing the KL divergence between a target density and a Gaussian density conditioned on a fixed mean is equivalent to firstly obtaining the second moment of the target density, and then correcting it by the difference between the first moments of the two densities. We term it *conditioned moment matching* (see Lemma A.1).

Theorem 3.3 shows that the optimal covariance to the joint optimization problem in Theorem 3.1 should be corrected by the error of $\hat{\epsilon}_n(\mathbf{x}_n)$ when $\hat{\epsilon}_n(\mathbf{x}_n) \neq \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\epsilon]$, which is inevitable in practice due to a non-realizable model family and the optimization error of gradient-based methods.

According to Eq. (14), the corrected covariance $\tilde{\sigma}_n^*(\mathbf{x}_n)^2$ is determined by $\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)} [(\epsilon_n - \hat{\epsilon}_n(\mathbf{x}_n))^2]$. We learn it by training a neural network $\mathbf{g}_n(\mathbf{x}_n) \in \mathbb{R}^d$ that predicts $(\epsilon_n - \hat{\epsilon}_n(\mathbf{x}_n))^2$ given \mathbf{x}_n on the following MSE loss

$$\min_{\{\mathbf{g}_n\}_{n=1}^N} \mathbb{E}_n \mathbb{E}_{q(\mathbf{x}_0, \mathbf{x}_n)} \|(\epsilon_n - \hat{\epsilon}_n(\mathbf{x}_n))^2 - \mathbf{g}_n(\mathbf{x}_n)\|_2^2, \quad (15)$$

where n is uniformly sampled from $\{1, 2, \dots, N\}$. Here we term $(\epsilon_n - \hat{\epsilon}_n(\mathbf{x}_n))^2$ the *noise prediction residual (NPR)*, which is the residual between the ground truth noise and the noise predicted by the noise prediction network, and we call $\mathbf{g}_n(\mathbf{x}_n)$ the *NPR prediction network*. Eq. (15) admits an optimal solution when $\mathbf{g}_n(\mathbf{x}_n) = \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)} [(\epsilon_n - \hat{\epsilon}_n(\mathbf{x}_n))^2]$ for all $n \in \{1, 2, \dots, N\}$. Note that in contrast to $\hat{\epsilon}_n(\mathbf{x}_n)^2$ used in SN-DPM, $\mathbf{g}_n(\mathbf{x}_n)$ doesn't have the error amplification problem (see Appendix E.1 for details). With a NPR

prediction network $g_n(\mathbf{x}_n)$ trained on Eq. (15), we estimate $\tilde{\sigma}_n^*(\mathbf{x}_n)^2$ by

$$\hat{\sigma}_n(\mathbf{x}_n)^2 = \lambda_n^2 \mathbf{1} + \gamma_n^2 \frac{\bar{\beta}_n}{\bar{\alpha}_n} g_n(\mathbf{x}_n). \quad (16)$$

In the paper, we refer to the above estimate as *NPR-DPM*.

4. Implementation

We provide more details of implementing SN-DPM and NPR-DPM. Specifically, we design a parameter sharing scheme for inference efficiency in Section 4.1. Motivated by Theorem 3.2 & 3.3, we employ a two-stage training process for both SN-DPM and NPR-DPM in Section 4.2.

4.1. Parameter Sharing and Inference Efficiency

During inference, both SN-DPM and NPR-DPM need to evaluate two networks. To reduce the computation cost, we let the two neural networks share parameters as follows

$$\begin{aligned} \hat{\epsilon}_n(\mathbf{x}_n) &= \text{NN}_1(\text{UNet}(\mathbf{x}_n, n; \boldsymbol{\theta}); \phi_1), \\ \mathbf{h}_n(\mathbf{x}_n) \text{ or } \mathbf{g}_n(\mathbf{x}_n) &= \text{NN}_2(\text{UNet}(\mathbf{x}_n, n; \boldsymbol{\theta}); \phi_2), \end{aligned} \quad (17)$$

where UNet is the most commonly used architecture in DPMs (Ho et al., 2020; Song et al., 2020b) and is parameterized by the shared parameter $\boldsymbol{\theta}$, NN_1 and NN_2 are two small neural networks such as convolution layers and are parameterized by ϕ_1 and ϕ_2 respectively.

In comparison to original DPMs with the same UNet, both SN-DPM and NPR-DPM require negligible additional memory cost and at most 10% more computation time (see details in Appendix F.3).

4.2. Two-Stage Learning and Pretraining

Recall that the optimal mean to problem (12) is irrelevant to the covariance. Thus, the learning of the mean and the covariance can naturally be divided into two stages. In the first stage, we learn the mean by training a noise prediction network on Eq. (7), or just use a pretrained one from prior works. In the second stage, we fix the parameter of the noise prediction network, take the UNet as a pretrained model, and only learn the covariance by tuning the parameter ϕ_2 (in Eq. (17)) of the NPR (or SN) prediction network using the gradient of Eq. (15) (or Eq. (9)):

$$\begin{aligned} \nabla_{\phi_2} \mathbb{E}_n \mathbb{E}_{q(\mathbf{x}_0, \mathbf{x}_n)} \|(\epsilon_n - \hat{\epsilon}_n(\mathbf{x}_n))^2 - g_n(\mathbf{x}_n)\|_2^2, \text{ or} \\ \nabla_{\phi_2} \mathbb{E}_n \mathbb{E}_{q(\mathbf{x}_0, \mathbf{x}_n)} \|\epsilon_n^2 - \mathbf{h}_n(\mathbf{x}_n)\|_2^2. \end{aligned}$$

Note that this pretraining paradigm for the second stage only needs to tune the small neural network NN_2 , which does not need to keep the computation graph of the UNet. Algorithm 1&2 present the second stage learning procedure.

5. Extension to DPMs with Continuous Timesteps

Song et al. (2020b) generalize DPMs to continuous timesteps by introducing a SDE $d\mathbf{x} = f(t)\mathbf{x}dt + g(t)d\mathbf{w}$, where $f(t)$ and $g(t)$ are two pre-specified scalar functions and \mathbf{w} is the standard Wiener process. The SDE can be viewed as a continuous version of the forward process in Eq. (2). It constructs a diffusion process $\{\mathbf{x}_t\}_{t=0}^T$ indexed by a continuous timestep $t \in [0, T]$, where \mathbf{x}_0 obeys the data distribution $q(\mathbf{x}_0)$. For $0 \leq s < t \leq T$, the conditional distribution of \mathbf{x}_t given \mathbf{x}_s is $q(\mathbf{x}_t|\mathbf{x}_s) = \mathcal{N}(\mathbf{x}_t|\sqrt{\alpha_{t|s}}\mathbf{x}_s, \beta_{t|s}\mathbf{I})$, where $\alpha_{t|s} = e^{2\int_s^t f(\tau)d\tau}$ and $\beta_{t|s} = \int_s^t g(\tau)^2 \alpha_{t|\tau} d\tau$. Following Bao et al. (2022); Kingma et al. (2021), we reverse the diffusion process from timesteps t to s by $p(\mathbf{x}_s|\mathbf{x}_t) = \mathcal{N}(\mathbf{x}_s|\boldsymbol{\mu}_{s|t}(\mathbf{x}_t), \boldsymbol{\Sigma}_{s|t}(\mathbf{x}_t))$ ($0 \leq s < t \leq T$), which is learned by

$$\min_{\boldsymbol{\mu}_{s|t}, \boldsymbol{\Sigma}_{s|t}} \ell_{s,t} := \mathbb{E}_{q(\mathbf{x}_t)} D_{\text{KL}}(q(\mathbf{x}_s|\mathbf{x}_t) \| p(\mathbf{x}_s|\mathbf{x}_t)), \quad (18)$$

where $q(\mathbf{x}_t)$ is the marginal distribution of \mathbf{x}_t . Similarly to the optimal solution in Eq. (4) & (5), Bao et al. (2022) derive the optimal mean

$$\boldsymbol{\mu}_{s|t}^*(\mathbf{x}_t) = \frac{1}{\sqrt{\alpha_{t|s}}} (\mathbf{x}_t - \frac{\beta_{t|s}}{\sqrt{\beta_{t|0}}} \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\epsilon_t]), \quad (19)$$

where $\epsilon_t = \frac{\mathbf{x}_t - \sqrt{\alpha_{t|0}}\mathbf{x}_0}{\sqrt{\beta_{t|0}}}$, as well as the optimal isotropic covariance $\sigma_{s|t}^{*2} \mathbf{I}$. Similarly to Eq. (7), a noise prediction network $\hat{\epsilon}_t(\mathbf{x}_t)$ is used to learn $\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\epsilon_t]$, inducing an estimate of the optimal mean

$$\hat{\boldsymbol{\mu}}_{s|t}(\mathbf{x}_t) = \frac{1}{\sqrt{\alpha_{t|s}}} (\mathbf{x}_t - \frac{\beta_{t|s}}{\sqrt{\beta_{t|0}}} \hat{\epsilon}_t(\mathbf{x}_t)). \quad (20)$$

Firstly, we extend Theorem 3.1 to DPMs with continuous timesteps. We derive the optimal solution to problem (18) when $\boldsymbol{\Sigma}_{s|t}(\mathbf{x}_t) = \text{diag}(\sigma_{s|t}(\mathbf{x}_t)^2)$, as shown in Proposition 5.1.

Proposition 5.1. *Suppose $\boldsymbol{\Sigma}_{s|t}(\mathbf{x}_t) = \text{diag}(\sigma_{s|t}(\mathbf{x}_t)^2)$. Then the optimal mean to problem (18) is $\boldsymbol{\mu}_{s|t}^*(\mathbf{x}_t)$ as in Eq. (19), and the optimal covariance to problem (18) is*

$$\sigma_{s|t}^*(\mathbf{x}_t)^2 = \tilde{\beta}_{s|t} \mathbf{1} + \frac{\beta_{t|s}^2}{\beta_{t|0} \alpha_{t|s}} (\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\epsilon_t^2] - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\epsilon_t]^2),$$

where $\epsilon_t = \frac{\mathbf{x}_t - \sqrt{\alpha_{t|0}}\mathbf{x}_0}{\sqrt{\beta_{t|0}}}$ is the noise used to generate \mathbf{x}_t from \mathbf{x}_0 and $\tilde{\beta}_{s|t} = \frac{\beta_{s|0}}{\beta_{t|0}} \beta_{t|s}$.

Then, we extend Theorem 3.3 to DPMs with continuous timesteps. Similarly to problem (12) & (13), we consider

Algorithm 1 Learning of the NPR prediction network

- 1: **Input:** The NPR prediction network $g_n(\mathbf{x}_n)$ and its tunable parameter ϕ_2 ; the pretrained noise prediction network $\hat{\epsilon}_n(\mathbf{x}_n)$
- 2: **repeat**
- 3: $\mathbf{x}_0 \sim q(\mathbf{x}_0)$
- 4: $n \sim \text{Uniform}(\{1, \dots, N\})$
- 5: $\epsilon_n \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
- 6: $\mathbf{x}_n = \sqrt{\alpha_n} \mathbf{x}_0 + \sqrt{\beta_n} \epsilon_n$
- 7: Take gradient descent step on $\nabla_{\phi_2} \|(\epsilon_n - \hat{\epsilon}_n(\mathbf{x}_n))^2 - g_n(\mathbf{x}_n)\|_2^2$
- 8: **until** converged

Algorithm 2 Learning of the SN prediction network

- 1: **Input:** The SN prediction network $h_n(\mathbf{x}_n)$ and its tunable parameter ϕ_2
- 2: **repeat**
- 3: $\mathbf{x}_0 \sim q(\mathbf{x}_0)$
- 4: $n \sim \text{Uniform}(\{1, \dots, N\})$
- 5: $\epsilon_n \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
- 6: $\mathbf{x}_n = \sqrt{\alpha_n} \mathbf{x}_0 + \sqrt{\beta_n} \epsilon_n$
- 7: Take gradient descent step on $\nabla_{\phi_2} \|\epsilon_n^2 - h_n(\mathbf{x}_n)\|_2^2$
- 8: **until** converged

the following two optimizations w.r.t. the mean or the covariance solely:

$$\text{Given arbitrary covariance } \Sigma_{s|t}, \min_{\mu_{s|t}} \ell_{s,t}, \quad (21)$$

$$\text{Given arbitrary mean } \mu_{s|t}, \min_{\Sigma_{s|t}} \ell_{s,t}. \quad (22)$$

The optimal mean to problem (21) is also Eq. (19), which is irrelevant to $\Sigma_{s|t}$. The optimal covariance to problem (22) is given in Proposition 5.2.

Proposition 5.2. *Suppose $\Sigma_{s|t}(\mathbf{x}_t) = \text{diag}(\sigma_{s|t}(\mathbf{x}_t)^2)$. For any mean $\mu_{s|t}(\mathbf{x}_t)$ that is parameterized by a noise prediction network $\hat{\epsilon}_t(\mathbf{x}_t)$ as in Eq. (20), the optimal covariance $\tilde{\sigma}_{s|t}^*(\mathbf{x}_t)^2$ to problem (22) is*

$$\begin{aligned} \tilde{\sigma}_{s|t}^*(\mathbf{x}_t)^2 &= \sigma_{s|t}^*(\mathbf{x}_t)^2 + \frac{\beta_{t|s}^2}{\beta_{t|0} \alpha_{t|s}} \underbrace{(\hat{\epsilon}_t(\mathbf{x}_t) - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\epsilon_t])^2}_{\text{error}} \\ &= \tilde{\beta}_{s|t} \mathbf{1} + \frac{\beta_{t|s}^2}{\beta_{t|0} \alpha_{t|s}} \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)} [(\epsilon_t - \hat{\epsilon}_t(\mathbf{x}_t))^2], \end{aligned}$$

where $\epsilon_t = \frac{\mathbf{x}_t - \sqrt{\alpha_{t|0}} \mathbf{x}_0}{\sqrt{\beta_{t|0}}}$ is the noise used to generate \mathbf{x}_t from \mathbf{x}_0 and $\tilde{\beta}_{s|t} = \frac{\beta_{s|0}}{\beta_{t|0}} \beta_{t|s}$.

Similarly to Eq. (10) & (16), to obtain estimates of $\sigma_{s|t}^*(\mathbf{x}_t)^2$ and $\tilde{\sigma}_{s|t}^*(\mathbf{x}_t)^2$, we can learn conditional expectations appeared in Proposition 5.1 & 5.2 by minimizing MSE losses. Please see details in Appendix C.

6. Experiments

We evaluate SN-DPM and NPR-DPM in DPMs with both discrete and continuous timesteps².

As for DPMs with discrete timesteps (see Section 2), we consider the DDPM forward process (corresponding to

²We provide our implementation at <https://github.com/baofff/Extended-Analytic-DPM>.

$\lambda_n^2 = \tilde{\beta}_n$ in Eq. (2)), and the DDIM forward process (corresponding to $\lambda_n^2 = 0$ in Eq. (2)). We explicitly call our method *NPR-DDPM* and *NPR-DDIM* (or *SN-DDPM* and *SN-DDIM*) in these two cases respectively. We compare our method with the following baselines: (i) the *original DDPM* (Ho et al., 2020), which uses handcrafted values $\Sigma_n(\mathbf{x}_n) = \tilde{\beta}_n \mathbf{I}$ or $\Sigma_n(\mathbf{x}_n) = \beta_n \mathbf{I}$, (ii) the *original DDIM* (Song et al., 2020a), which uses handcrafted values $\Sigma_n(\mathbf{x}_n) = \lambda_n^2 \mathbf{I} = \mathbf{0}$, and (iii) the *Analytic-DPM* (Bao et al., 2022), as introduced in Section 2.

As for DPMs with continuous timesteps (see Section 5), we consider the VP SDE (Song et al., 2020b) as the forward process. We compare our method with the following baselines: (i) the *Euler-Maruyama solver* (Song et al., 2020b), which firstly reverses the VP SDE and then discretizes the reverse one using the Euler-Maruyama solver, (ii) the *ancestral sampling* (Song et al., 2020b), which designs a Markov chain similar to Eq. (1) and samples from the Markov chain, (iii) the *probability flow* (Song et al., 2020b), which firstly derives an equivalent ODE to the SDE and then discretizes the ODE, (iv) the *Analytic-DPM* (Bao et al., 2022), as introduced in Section 2, and (v) the “*Gotta Go Fast*” SDE solver (Jolicoeur-Martineau et al., 2021), which uses adaptive step sizes.

Since the performance on a subset of timesteps is crucial, we compare our method and baselines constrained on trajectories $1 \leq \tau_1 < \dots < \tau_K = N$ with different number of timesteps K (Song et al., 2020a; Bao et al., 2022) (see Appendix D for details of how to apply our method to trajectories). Following Bao et al. (2022), we consider two kinds of trajectories. The first one is the even trajectory (ET) (Nichol & Dhariwal, 2021), where the timesteps are evenly spaced. The second one is the optimal trajectory (OT) (Watson et al., 2021), where the timesteps are determined by dynamic programming that maximizes ELBO.

We evaluate our method on six pretrained noise prediction networks provided by prior works (Ho et al., 2020; Song

et al., 2020b;a; Nichol & Dhariwal, 2021; Bao et al., 2022). Three of them are trained on CIFAR10 (Krizhevsky et al., 2009) with the linear schedule (LS) of β_n (Ho et al., 2020), the cosine schedule (CS) of β_n (Nichol & Dhariwal, 2021) and the VP SDE (Song et al., 2020b) respectively. We refer to the first two settings as CIFAR10 (LS) and CIFAR10 (CS) respectively, which have discrete timesteps, and refer to the last setting as CIFAR10 (VP SDE), which has continuous timesteps. The others are trained with discrete timesteps on CelebA 64x64 (Liu et al., 2015), ImageNet 64x64 (Deng et al., 2009) and LSUN Bedroom (Yu et al., 2015) respectively. We train NPR (or SN) prediction networks for all of them, following the implementation in Section 4. See more experimental details in Appendix F.

6.1. Sample Quality

In this part, we compare the sample quality quantitatively, as measured by the commonly used FID score (Heusel et al., 2017). We evaluate DPMs with both the discrete and continuous timesteps. As for DPMs with discrete timesteps, we report results under both the DDPM and DDIM forward processes. Since Watson et al. (2021); Bao et al. (2022) report that a DPM with the optimal trajectory is likely to have a worse FID score, we only include results under the even trajectory.

As shown in Table 2, both NPR-DPM and SN-DPM outperform baselines in most cases. In particular, SN-DPM improves the sample quality remarkably on CIFAR10 and CelebA 64x64 when the number of timesteps is small. Moreover, our methods perform well in DPMs with both discrete and continuous timesteps.

Following Bao et al. (2022), we also compare the least number of timesteps required to achieve a FID around 6, as shown in Table 3. Remarkably, on CelebA 64x64, SN-DPM requires less than half the number of timesteps compared to baselines. In Appendix G.3, we show FID comparison after normalizing for the extra cost, as well as the comparison with the ‘‘Gotta Go Fast’’ SDE solver. In Appendix G.4, we show generated samples of our methods.

6.2. Likelihood Results

In this part, we evaluate the upper bound $-L_{\text{elbo}}$ (see Eq. (3)) on the negative log-likelihood of DPMs with discrete timesteps. Since Bao et al. (2022) claim that $-L_{\text{elbo}} = \infty$ in the DDIM forward process, we only include results under the DDPM forward process.

As shown in Table 4, our NPR-DPM consistently outperforms all baselines on all datasets, all number of timesteps and both kinds of trajectories.

In initial experiments, we find that SN-DPM doesn’t perform well on the likelihood results, potentially due to not

considering the imperfect mean and the amplified error of $\hat{\epsilon}_n(x_n)^2$ as mentioned in Section 3.1. The inconsistency with its FID results in Table 2 roots in the different natures of the two metrics, which has been observed and explored extensively in prior works (Ho et al., 2020; Nichol & Dhariwal, 2021; Song et al., 2021; Vahdat et al., 2021; Watson et al., 2021; Kingma et al., 2021; Bao et al., 2022). We also note that while NPR-DPM performs better than SN-DPM w.r.t. likelihood, it does not perform better w.r.t. FID. Please see Appendix E.2 for a discussion.

We also compare to Improved DDPM (Nichol & Dhariwal, 2021) in terms of likelihood estimation. With a small number of timesteps, our method can get a much better likelihood result than Nichol & Dhariwal (2021) on ImageNet 64x64. With the full timesteps, two methods show a similar likelihood performance (the upper bound on NLL of our method is 3.60 and that of Improved DDPM is 3.57), which implies that using an alternative objective such as the MSE loss can achieve a similar likelihood performance to directly optimizing L_{elbo} . See details in Appendix G.1.

7. Related Work

DPMs and their variants. The idea that a data generation process can be constructed by reversing a diffusion process is initially introduced by Sohl-Dickstein et al. (2015). Specifically, Sohl-Dickstein et al. (2015) reverse the diffusion process using a Markov chain with discrete timesteps, which is learned on the ELBO objective. Ho et al. (2020) propose to parameterize the mean of the Markov chain by the noise prediction network, which share parameters between different timesteps and is learned on a MSE loss. Song et al. (2020b) consider DPMs with infinitesimal timesteps, which can be represented by stochastic differential equations (SDEs). Recently, some variants of DPMs are proposed. Kingma et al. (2021) additionally learn the forward diffusion process, Vahdat et al. (2021) learn DPMs in latent space and Dockhorn et al. (2021) introduce auxiliary velocity variables to the diffusion process.

Based on the powerful generation performance, DPMs have shown great promise in a wide variety of applications, such as controllable generation (Choi et al., 2021; Meng et al., 2021; Sinha et al., 2021; Nichol et al., 2021), voice conversion (Popov et al., 2021), image super-resolution (Saharia et al., 2021; Li et al., 2021), image-to-image translation (Sasaki et al., 2021), shape generation (Zhou et al., 2021), 3d point cloud generation (Luo & Hu, 2021) and time series forecasting (Rasul et al., 2021).

Covariance design in DPMs. As mentioned in Section 1, the covariance design in DPMs is crucial for the performance on a subset of timesteps, and some prior works (Ho et al., 2020; Song et al., 2020a; Bao et al., 2022) use

Table 2. FID score \downarrow . All results are evaluated under the ET. Here A-DPM denotes the baseline Analytic-DPM. Note that the extra time cost of the NPR (or SN) prediction network is negligible on CIFAR10, CelebA 64x64, and at most 4.5% on ImageNet 64x64 (see Appendix F.3 for details). Thus, we can use the number of timesteps to compare the efficiency of these methods.

# TIMESTEPS K	CIFAR10 (LS)						CIFAR10 (CS)					
	10	25	50	100	200	1000	10	25	50	100	200	1000
DDPM, $\tilde{\beta}_n$	44.45	21.83	15.21	10.94	8.23	5.11	34.76	16.18	11.11	8.38	6.66	4.92
DDPM, β_n	233.41	125.05	66.28	31.36	12.96	3.04	205.31	84.71	37.35	14.81	5.74	3.34
A-DDPM	34.26	11.60	7.25	5.40	4.01	4.03	22.94	8.50	5.50	4.45	4.04	4.31
NPR-DDPM	32.35	10.55	6.18	4.52	3.57	4.10	19.94	7.99	5.31	4.52	4.10	4.27
SN-DDPM	24.06	6.91	4.63	3.67	3.31	3.65	16.33	6.05	4.17	3.83	3.72	4.07
DDIM	21.31	10.70	7.74	6.08	5.07	4.13	34.34	16.68	10.48	7.94	6.69	4.89
A-DDIM	14.00	5.81	4.04	3.55	3.39	3.74	26.43	9.96	6.02	4.88	4.92	4.66
NPR-DDIM	13.34	5.38	3.95	3.53	3.42	3.72	22.81	9.47	6.04	5.02	5.06	4.62
SN-DDIM	12.19	4.28	3.39	3.23	3.22	3.65	17.90	7.36	5.16	4.63	4.63	4.51

# TIMESTEPS K	CELEBA 64X64						IMAGENET 64X64					
	10	25	50	100	200	1000	25	50	100	200	400	4000
DDPM, $\tilde{\beta}_n$	36.69	24.46	18.96	14.31	10.48	5.95	29.21	21.71	19.12	17.81	17.48	16.55
DDPM, β_n	294.79	115.69	53.39	25.65	9.72	3.16	170.28	83.86	45.04	28.39	21.38	16.38
A-DDPM	28.99	16.01	11.23	8.08	6.51	5.21	32.56	22.45	18.80	17.16	16.40	16.34
NPR-DDPM	28.37	15.74	10.89	8.23	7.03	5.33	28.27	20.89	18.06	16.96	16.32	16.38
SN-DDPM	20.60	12.00	7.88	5.89	5.02	4.42	27.58	20.74	18.04	16.61	16.37	16.22
DDIM	20.54	13.45	9.33	6.60	4.96	3.40	26.06	20.10	18.09	17.84	17.74	19.00
A-DDIM	15.62	9.22	6.13	4.29	3.46	3.13	25.98	19.23	17.73	17.49	17.44	18.98
NPR-DDIM	14.98	8.93	6.04	4.27	3.59	3.15	28.84	19.62	17.63	17.42	17.30	18.91
SN-DDIM	10.20	5.48	3.83	3.04	2.85	2.90	28.07	19.38	17.53	17.23	17.23	18.89

# TIMESTEPS K	CIFAR10 (VP SDE)					
	10	25	50	100	200	1000
EULER-MARUYAMA	292.20	170.17	90.79	47.46	21.92	2.55
ANCESTRAL SAMPLING	235.28	129.29	68.52	31.99	12.81	2.72
PROBABILITY FLOW	107.74	21.34	7.78	4.33	3.27	2.82
A-DPM	35.10	11.57	6.54	4.71	3.61	2.98
NPR-DPM	33.70	10.44	5.83	3.97	3.05	3.04
SN-DPM	25.30	7.34	4.46	3.27	2.83	2.71

isotropic covariances that only depends on the timestep without considering the state, which is less flexible compared to ours. Besides, we note that [Nichol & Dhariwal \(2021\)](#) also use state-dependent diagonal covariances. The difference is that [Nichol & Dhariwal \(2021\)](#) directly train a covariance network on L_{elbo} , and our one is learned by optimizing a MSE loss based on the form of the optimal covariance. As a result, our method can get a reasonable likelihood result on a small number of timesteps, which is contrast to [Nichol & Dhariwal \(2021\)](#) (see Appendix G.1).

Faster DPMs. In addition to exploring covariance design in DPMs, there are also other attempts toward faster DPMs. Several works explore to find short trajectories by applying searching algorithms, e.g., grid search ([Chen et al., 2020](#)), dynamic programming ([Watson et al., 2021](#)) and differentiable search ([Anonymous, 2022a](#)). These searching algo-

gorithms can be further combined with our method to get a better performance. For example, the combination of the dynamic programming and our method leads to a better likelihood performance, as shown by the OT results in Table 4. Several works change the model family in the reverse process, replacing the Gaussian model. [Luhman & Luhman \(2021\)](#); [Anonymous \(2022b\)](#) distill the original reverse process to a model of a single or a few timesteps, and [Xiao et al. \(2021\)](#) model the reverse process as a conditional generator and propose to train with an adversarial loss. However, their likelihood estimation is non-trivial due to the change of the model family. Several works design faster solvers ([Jolicœur-Martineau et al., 2021](#); [Popov et al., 2021](#)) for DPMs with continuous timesteps. However, these solvers only consider isotropic covariances as in [Ho et al. \(2020\)](#); [Bao et al. \(2022\)](#), which are less flexible compared to ours.

Table 3. The least number of timesteps \downarrow required to achieve a FID around 6 (with the corresponding FID). The time cost in a single timestep is the same for all baselines. Since our methods have extra prediction networks which require extra time cost, we multiply their results by the ratio of the time cost in a single timestep in them to that in baselines. See experimental details in Appendix F.7.

METHOD	CIFAR10	CELEBA 64X64	LSUN BEDROOM
DDPM (HO ET AL., 2020)	90 (6.12)	> 200	130 (6.06)
DDIM (SONG ET AL., 2020A)	30 (5.85)	> 100	BEST FID > 6
IMPROVED DDPM (NICHOL & DHARIWAL, 2021)	45 (5.96)	MISSING MODEL	90 (6.02)
ANALYTIC-DPM (BAO ET AL., 2022)	25 (5.81)	55 (5.98)	100 (6.05)
NPR-DPM (OURS)	1.002×23 (5.76)	1.011×50 (6.04)	1.017× 90 (6.01)
SN-DPM (OURS)	1.005× 17 (5.81)	1.012× 22 (5.96)	1.100×92 (6.02)

Table 4. Upper bound ($-L_{\text{elbo}}$) on the negative log-likelihood (bits/dim) \downarrow . Here A-DPM denotes the baseline Analytic-DPM. Note that the extra time cost of the NPR prediction network is negligible on CIFAR10, CelebA 64x64, and at most 4.5% on ImageNet 64x64 (see Appendix F.3 for details). Thus, we can use the number of timesteps to compare the efficiency of these methods. To reduce the effect of randomness, we also repeat the experiments under the full timesteps and report the variance in Appendix G.2.

# TIMESTEPS K		CIFAR10 (LS)						CIFAR10 (CS)					
		10	25	50	100	200	1000	10	25	50	100	200	1000
ET	DDPM, $\tilde{\beta}_n$	74.95	24.98	12.01	7.08	5.03	3.73	75.96	24.94	11.96	7.04	4.95	3.60
	DDPM, β_n	6.99	6.11	5.44	4.86	4.39	3.75	6.51	5.55	4.92	4.41	4.03	3.54
	A-DDPM	5.47	4.79	4.38	4.07	3.84	3.59	5.08	4.45	4.09	3.83	3.64	3.42
	NPR-DDPM	5.40	4.64	4.25	3.98	3.79	3.57	5.03	4.33	3.99	3.76	3.59	3.41
OT	DDPM, β_n	5.38	4.34	3.97	3.82	3.77	3.75	5.51	4.30	3.86	3.65	3.57	3.54
	A-DDPM	4.11	3.68	3.61	3.59	3.59	3.59	3.99	3.56	3.47	3.44	3.43	3.42
	NPR-DDPM	3.91	3.64	3.59	3.58	3.57	3.57	3.88	3.52	3.45	3.42	3.41	3.41

# TIMESTEPS K		CELEBA 64X64						IMAGENET 64X64					
		10	25	50	100	200	1000	25	50	100	200	400	4000
ET	DDPM, $\tilde{\beta}_n$	33.42	13.09	7.14	4.60	3.45	2.71	105.87	46.25	22.02	12.10	7.59	3.89
	DDPM, β_n	6.67	5.72	4.98	4.31	3.74	2.93	5.81	5.20	4.70	4.31	4.04	3.65
	A-DDPM	4.54	3.89	3.48	3.16	2.92	2.66	4.78	4.42	4.15	3.95	3.81	3.61
	NPR-DDPM	4.46	3.78	3.40	3.11	2.89	2.65	4.66	4.22	3.96	3.80	3.71	3.60
OT	DDPM, β_n	4.76	3.58	3.16	2.99	2.94	2.93	4.56	4.09	3.84	3.73	3.68	3.65
	A-DDPM	2.97	2.71	2.67	2.66	2.66	2.66	3.83	3.70	3.64	3.62	3.62	3.61
	NPR-DDPM	2.88	2.69	2.66	2.66	2.65	2.65	3.73	3.65	3.62	3.60	3.60	3.60

8. Conclusion

We consider diagonal and full covariances to improve the expressive power of DPMs. We derive the optimal mean and covariance in terms of the maximum likelihood, and also correct it given an imperfect mean. Both the optimal and the corrected ones can be decomposed into terms of conditional expectations over functions of noise, which can be estimated by minimizing MSE losses. We consider diagonal covariances in our implementation for computational efficiency. Besides, we adopt a parameter sharing scheme for inference efficiency and a two-stage training process motivated by our theory. Our method is applicable to DPMs with both discrete and continuous timesteps. Empirically, our method consistently outperforms a wide variety of baselines on likelihood estimation, and improves the sample quality especially on a small number of timesteps.

Acknowledgements

This work was supported by NSF of China Projects (Nos. 62061136001, 61620106010, 62076145, U19B2034, U1811461, U19A2081, 6197222); Beijing NSF Project (No. JQ19016); Beijing Outstanding Young Scientist Program NO. BJJWZYJH012019100020098; a grant from Tsinghua Institute for Guo Qiang; the NVIDIA NVAIL Program with GPU/DGX Acceleration; the High Performance Computing Center, Tsinghua University; and Major Innovation & Planning Interdisciplinary Platform for the ‘‘Double-First Class’’ Initiative, Renmin University of China.

References

Anonymous. Optimizing few-step diffusion samplers by gradient descent. In *Submitted to The Tenth In-*

- ternational Conference on Learning Representations, 2022a. URL <https://openreview.net/forum?id=VFBjuF8HEp>. under review.
- Anonymous. Progressive distillation for fast sampling of diffusion models. In *Submitted to The Tenth International Conference on Learning Representations, 2022b*. URL <https://openreview.net/forum?id=TIIdIXIpzhoI>. under review.
- Bao, F., Li, C., Zhu, J., and Zhang, B. Analytic-dpm: an analytic estimate of the optimal reverse variance in diffusion probabilistic models. *arXiv preprint arXiv:2201.06503*, 2022.
- Bińkowski, M., Donahue, J., Dieleman, S., Clark, A., Elsen, E., Casagrande, N., Cobo, L. C., and Simonyan, K. High fidelity speech synthesis with adversarial networks. *arXiv preprint arXiv:1909.11646*, 2019.
- Brock, A., Donahue, J., and Simonyan, K. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018.
- Chen, N., Zhang, Y., Zen, H., Weiss, R. J., Norouzi, M., and Chan, W. Wavegrad: Estimating gradients for waveform generation. *arXiv preprint arXiv:2009.00713*, 2020.
- Choi, J., Kim, S., Jeong, Y., Gwon, Y., and Yoon, S. Ilvr: Conditioning method for denoising diffusion probabilistic models. *arXiv preprint arXiv:2108.02938*, 2021.
- 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.
- Dhariwal, P. and Nichol, A. Diffusion models beat gans on image synthesis. *arXiv preprint arXiv:2105.05233*, 2021.
- Dockhorn, T., Vahdat, A., and Kreis, K. Score-based generative modeling with critically-damped langevin diffusion. *arXiv preprint arXiv:2112.07068*, 2021.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.
- Ho, J., Jain, A., and Abbeel, P. Denoising diffusion probabilistic models. *arXiv preprint arXiv:2006.11239*, 2020.
- Jolicoeur-Martineau, A., Li, K., Piché-Taillefer, R., Kachman, T., and Mitliagkas, I. Gotta go fast when generating data with score-based models. *arXiv preprint arXiv:2105.14080*, 2021.
- Kalchbrenner, N., Elsen, E., Simonyan, K., Noury, S., Casagrande, N., Lockhart, E., Stimberg, F., Oord, A., Dieleman, S., and Kavukcuoglu, K. Efficient neural audio synthesis. In *International Conference on Machine Learning*, pp. 2410–2419. PMLR, 2018.
- Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., and Aila, T. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8110–8119, 2020.
- Kingma, D. P., Salimans, T., Poole, B., and Ho, J. Variational diffusion models. *arXiv preprint arXiv:2107.00630*, 2021.
- Kong, Z., Ping, W., Huang, J., Zhao, K., and Catanzaro, B. Diffwave: A versatile diffusion model for audio synthesis. *arXiv preprint arXiv:2009.09761*, 2020.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. 2009.
- Li, H., Yang, Y., Chang, M., Feng, H., Xu, Z., Li, Q., and Chen, Y. Srdiff: Single image super-resolution with diffusion probabilistic models. *arXiv preprint arXiv:2104.14951*, 2021.
- Liu, Z., Luo, P., Wang, X., and Tang, X. Deep learning face attributes in the wild. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pp. 3730–3738. IEEE Computer Society, 2015.
- Loshchilov, I. and Hutter, F. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- Luhman, E. and Luhman, T. Knowledge distillation in iterative generative models for improved sampling speed. *arXiv preprint arXiv:2101.02388*, 2021.
- Luo, S. and Hu, W. Diffusion probabilistic models for 3d point cloud generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2837–2845, 2021.
- Meng, C., Song, Y., Song, J., Wu, J., Zhu, J.-Y., and Ermon, S. Sdedit: Image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*, 2021.
- Nichol, A. and Dhariwal, P. Improved denoising diffusion probabilistic models. *arXiv preprint arXiv:2102.09672*, 2021.

- 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. *arXiv preprint arXiv:2112.10741*, 2021.
- Popov, V., Vovk, I., Gogoryan, V., Sadekova, T., Kudinov, M., and Wei, J. Diffusion-based voice conversion with fast maximum likelihood sampling scheme. *arXiv preprint arXiv:2109.13821*, 2021.
- Rasul, K., Seward, C., Schuster, I., and Vollgraf, R. Autoregressive denoising diffusion models for multivariate probabilistic time series forecasting. *arXiv preprint arXiv:2101.12072*, 2021.
- Saharia, C., Ho, J., Chan, W., Salimans, T., Fleet, D. J., and Norouzi, M. Image super-resolution via iterative refinement. *arXiv preprint arXiv:2104.07636*, 2021.
- Sasaki, H., Willcocks, C. G., and Breckon, T. P. Unit-ddpm: Unpaired image translation with denoising diffusion probabilistic models. *arXiv preprint arXiv:2104.05358*, 2021.
- Sinha, A., Song, J., Meng, C., and Ermon, S. D2c: Diffusion-denoising models for few-shot conditional generation. *arXiv preprint arXiv:2106.06819*, 2021.
- Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., and Ganguli, S. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pp. 2256–2265. PMLR, 2015.
- Song, J., Meng, C., and Ermon, S. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020a.
- Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020b.
- Song, Y., Durkan, C., Murray, I., and Ermon, S. Maximum likelihood training of score-based diffusion models. *arXiv e-prints*, pp. arXiv–2101, 2021.
- Theis, L., Oord, A. v. d., and Bethge, M. A note on the evaluation of generative models. *arXiv preprint arXiv:1511.01844*, 2015.
- Vahdat, A., Kreis, K., and Kautz, J. Score-based generative modeling in latent space. *arXiv preprint arXiv:2106.05931*, 2021.
- Watson, D., Ho, J., Norouzi, M., and Chan, W. Learning to efficiently sample from diffusion probabilistic models. *arXiv preprint arXiv:2106.03802*, 2021.
- Wu, Y., Donahue, J., Balduzzi, D., Simonyan, K., and Lillcrap, T. Logan: Latent optimisation for generative adversarial networks. *arXiv preprint arXiv:1912.00953*, 2019.
- Xiao, Z., Kreis, K., and Vahdat, A. Tackling the generative learning trilemma with denoising diffusion gans. *arXiv preprint arXiv:2112.07804*, 2021.
- Yu, F., Zhang, Y., Song, S., Seff, A., and Xiao, J. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015.
- Zhou, L., Du, Y., and Wu, J. 3d shape generation and completion through point-voxel diffusion. *arXiv preprint arXiv:2104.03670*, 2021.

A. Proof

A.1. Proof of Theorem 3.1, 3.2 & 3.3

Lemma A.1. (Conditioned moment matching) Suppose $q(\mathbf{x})$ is a probability density function with mean $\boldsymbol{\mu}_q$ and covariance matrix $\boldsymbol{\Sigma}_q$ and $p(\mathbf{x}) = \mathcal{N}(\mathbf{x}|\boldsymbol{\mu}, \text{diag}(\boldsymbol{\sigma}^2))$ is a Gaussian distribution. Then

1. for any $\boldsymbol{\sigma}^2$, the optimization problem $\min_{\boldsymbol{\mu}} D_{\text{KL}}(q||p)$ admits an optimal solution $\boldsymbol{\mu}^* = \boldsymbol{\mu}_q$;
2. for any $\boldsymbol{\mu}$, the optimization problem $\min_{\boldsymbol{\sigma}^2} D_{\text{KL}}(q||p)$ admits an optimal solution $\tilde{\boldsymbol{\sigma}}^{*2} = \text{diag}(\boldsymbol{\Sigma}_q) + (\boldsymbol{\mu} - \boldsymbol{\mu}_q)^2$, where $(\cdot)^2$ is the element-wise square;
3. the optimization problem $\min_{\boldsymbol{\mu}, \boldsymbol{\sigma}^2} D_{\text{KL}}(q||p)$ admits an optimal solution $\boldsymbol{\mu}^* = \boldsymbol{\mu}_q$ and $\boldsymbol{\sigma}^{*2} = \text{diag}(\boldsymbol{\Sigma}_q)$.

Proof. Bao et al. (2022) show that the KL divergence between a probability density function q and a Gaussian distribution p can be written as $D_{\text{KL}}(q||p) = D_{\text{KL}}(\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_q, \boldsymbol{\Sigma}_q)||p) + c$, where c is a constant only related to q (see Lemma 2 of Bao et al. (2022)). Thus, the optimizations of $D_{\text{KL}}(q||p)$ and $D_{\text{KL}}(\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_q, \boldsymbol{\Sigma}_q)||p)$ w.r.t. $\boldsymbol{\mu}$ or $\boldsymbol{\sigma}^2$ are equivalent.

We expand $D_{\text{KL}}(\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_q, \boldsymbol{\Sigma}_q)||p)$ according to the definition of KL divergence:

$$2D_{\text{KL}}(\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_q, \boldsymbol{\Sigma}_q)||p) = \text{tr}(\text{diag}(\boldsymbol{\sigma}^{-2})\boldsymbol{\Sigma}_q) - d + \log \frac{|\text{diag}(\boldsymbol{\sigma}^2)|}{|\boldsymbol{\Sigma}_q|} + (\boldsymbol{\mu} - \boldsymbol{\mu}_q)^\top \text{diag}(\boldsymbol{\sigma}^{-2})(\boldsymbol{\mu} - \boldsymbol{\mu}_q) \quad (23)$$

$$= \sum_{i=1}^d \{ \sigma_i^{-2} [(\boldsymbol{\Sigma}_q)_{i,i} + (\boldsymbol{\mu}_i - (\boldsymbol{\mu}_q)_i)^2] + \log \sigma_i^2 \} - \log |\boldsymbol{\Sigma}_q| - d. \quad (24)$$

According to Eq. (23), we know that

$$\min_{\boldsymbol{\mu}} D_{\text{KL}}(q||p) \Leftrightarrow \min_{\boldsymbol{\mu}} (\boldsymbol{\mu} - \boldsymbol{\mu}_q)^\top \text{diag}(\boldsymbol{\sigma}^{-2})(\boldsymbol{\mu} - \boldsymbol{\mu}_q),$$

which admits an optimal solution $\boldsymbol{\mu}^* = \boldsymbol{\mu}_q$.

According to Eq. (24), we know that

$$\min_{\boldsymbol{\sigma}^2} D_{\text{KL}}(q||p) \Leftrightarrow \min_{\boldsymbol{\sigma}^2} \sum_{i=1}^d \{ \sigma_i^{-2} [(\boldsymbol{\Sigma}_q)_{i,i} + (\boldsymbol{\mu}_i - (\boldsymbol{\mu}_q)_i)^2] + \log \sigma_i^2 \}$$

By taking gradients, we know $\sigma_i^{-2} [(\boldsymbol{\Sigma}_q)_{i,i} + (\boldsymbol{\mu}_i - (\boldsymbol{\mu}_q)_i)^2] + \log \sigma_i^2$ as a function of σ_i^2 gets its minimum at

$$\tilde{\sigma}_i^{*2} = (\boldsymbol{\Sigma}_q)_{i,i} + (\boldsymbol{\mu}_i - (\boldsymbol{\mu}_q)_i)^2.$$

Thus, $\min_{\boldsymbol{\sigma}^2} D_{\text{KL}}(q||p)$ admits an optimal solution $\tilde{\boldsymbol{\sigma}}^{*2} = \text{diag}(\boldsymbol{\Sigma}_q) + (\boldsymbol{\mu} - \boldsymbol{\mu}_q)^2$.

Combining $\boldsymbol{\mu}^*$ and $\tilde{\boldsymbol{\sigma}}^{*2}$, we know that $\min_{\boldsymbol{\mu}, \boldsymbol{\sigma}^2} D_{\text{KL}}(q||p)$ admits an optimal solution $\boldsymbol{\mu}^* = \boldsymbol{\mu}_q$ and $\boldsymbol{\sigma}^{*2} = \text{diag}(\boldsymbol{\Sigma}_q) + (\boldsymbol{\mu}^* - \boldsymbol{\mu}_q)^2 = \text{diag}(\boldsymbol{\Sigma}_q)$. \square

Lemma A.2. Suppose $q(\mathbf{x}_{0:N})$ is defined as Eq. (2). Then we have

$$\begin{aligned} \mathbb{E}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}] &= \tilde{\boldsymbol{\mu}}_n(\mathbf{x}_n, \frac{1}{\sqrt{\tilde{\alpha}_n}}(\mathbf{x}_n - \sqrt{\tilde{\beta}_n} \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n])), \\ \text{Cov}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}] &= \lambda_n^2 \mathbf{I} + \gamma_n^2 \frac{\tilde{\beta}_n}{\tilde{\alpha}_n} (\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n \boldsymbol{\epsilon}_n^\top] - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n] \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]^\top), \\ \text{diag}(\text{Cov}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}]) &= \lambda_n^2 \mathbf{1} + \gamma_n^2 \frac{\tilde{\beta}_n}{\tilde{\alpha}_n} (\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n^2] - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]^2). \end{aligned}$$

where $\boldsymbol{\epsilon}_n = \frac{\mathbf{x}_n - \sqrt{\tilde{\alpha}_n} \mathbf{x}_0}{\sqrt{\tilde{\beta}_n}}$ is the noise used to generate \mathbf{x}_n from \mathbf{x}_0 and $\gamma_n = \sqrt{\tilde{\alpha}_{n-1}} - \sqrt{\tilde{\beta}_{n-1} - \lambda_n^2} \sqrt{\frac{\tilde{\alpha}_n}{\tilde{\beta}_n}}$.

Proof. Since $\tilde{\boldsymbol{\mu}}_n(\mathbf{x}_n, \mathbf{x}_0)$ is linear w.r.t. \mathbf{x}_0 , we have

$$\begin{aligned} \mathbb{E}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}] &= \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}\mathbb{E}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n, \mathbf{x}_0)}[\mathbf{x}_{n-1}] \\ &= \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}\tilde{\boldsymbol{\mu}}_n(\mathbf{x}_n, \mathbf{x}_0) = \tilde{\boldsymbol{\mu}}_n(\mathbf{x}_n, \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}\mathbf{x}_0) = \tilde{\boldsymbol{\mu}}_n(\mathbf{x}_n, \frac{1}{\sqrt{\bar{\alpha}_n}}(\mathbf{x}_n - \sqrt{\bar{\beta}_n}\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n])). \end{aligned}$$

As for $\text{Cov}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}]$, Bao et al. (2022) show that it can be expanded as $\text{Cov}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}] = \lambda_n^2 \mathbf{I} + \gamma_n^2 \text{Cov}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\mathbf{x}_0]$, where $\gamma_n = \sqrt{\bar{\alpha}_{n-1}} - \sqrt{\bar{\beta}_{n-1}} - \lambda_n^2 \sqrt{\frac{\bar{\alpha}_n}{\bar{\beta}_n}}$ (see Lemma 13 in Bao et al. (2022)).

Besides, we have $\text{Cov}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n] = \text{Cov}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\frac{\mathbf{x}_n - \sqrt{\bar{\alpha}_n}\mathbf{x}_0}{\sqrt{\bar{\beta}_n}}] = \frac{\bar{\alpha}_n}{\bar{\beta}_n} \text{Cov}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\mathbf{x}_0]$. Thus,

$$\begin{aligned} \text{Cov}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}] &= \lambda_n^2 \mathbf{I} + \gamma_n^2 \frac{\bar{\beta}_n}{\bar{\alpha}_n} \text{Cov}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n] \\ &= \lambda_n^2 \mathbf{I} + \gamma_n^2 \frac{\bar{\beta}_n}{\bar{\alpha}_n} (\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n \boldsymbol{\epsilon}_n^\top] - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n] \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]^\top), \\ \text{diag}(\text{Cov}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}]) &= \lambda_n^2 \mathbf{1} + \gamma_n^2 \frac{\bar{\beta}_n}{\bar{\alpha}_n} (\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n^2] - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]^2). \end{aligned}$$

□

Theorem 3.1. (Optimal solution to joint optimization) Suppose $\boldsymbol{\Sigma}_n(\mathbf{x}_n) = \text{diag}(\boldsymbol{\sigma}_n(\mathbf{x}_n)^2)$. Then the optimal mean to problem (3) is $\boldsymbol{\mu}_n^*(\mathbf{x}_n)$ as in Eq. (4), and the optimal covariance to problem (3) is

$$\boldsymbol{\sigma}_n^*(\mathbf{x}_n)^2 = \lambda_n^2 \mathbf{1} + \gamma_n^2 \frac{\bar{\beta}_n}{\bar{\alpha}_n} (\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n^2] - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]^2),$$

where $\boldsymbol{\epsilon}_n = \frac{\mathbf{x}_n - \sqrt{\bar{\alpha}_n}\mathbf{x}_0}{\sqrt{\bar{\beta}_n}}$ is the noise used to generate \mathbf{x}_n from \mathbf{x}_0 , $(\cdot)^2$ is the element-wise square, $\mathbf{1}$ is the vector of ones and $\gamma_n = \sqrt{\bar{\alpha}_{n-1}} - \sqrt{\bar{\beta}_{n-1}} - \lambda_n^2 \sqrt{\frac{\bar{\alpha}_n}{\bar{\beta}_n}}$.

Proof. Bao et al. (2022) show that the KL divergence between a joint probability density function $q(\mathbf{x}_{0:N})$ and a Markov chain $p(\mathbf{x}_{0:N})$ can be written as

$$D_{\text{KL}}(q(\mathbf{x}_{0:N})\|p(\mathbf{x}_{0:N})) = \sum_{n=1}^N \mathbb{E}_q D_{\text{KL}}(q(\mathbf{x}_{n-1}|\mathbf{x}_n)\|p(\mathbf{x}_{n-1}|\mathbf{x}_n)) + c,$$

where c is a constant only related to q (see Lemma 8 in Bao et al. (2022)). As a result, the optimization of the KL divergence is decomposed into n independent optimization sub-problems:

$$\min_{\boldsymbol{\mu}_n(\cdot), \boldsymbol{\sigma}_n(\cdot)^2} \mathbb{E}_q D_{\text{KL}}(q(\mathbf{x}_{n-1}|\mathbf{x}_n)\|p(\mathbf{x}_{n-1}|\mathbf{x}_n)), \quad 1 \leq n \leq N.$$

Since $q(\mathbf{x}_{n-1}|\mathbf{x}_n)$ is a probability density function with mean $\mathbb{E}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}]$ and covariance matrix $\text{Cov}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}]$ and $p(\mathbf{x}_{n-1}|\mathbf{x}_n) = \mathcal{N}(\mathbf{x}_{n-1}|\boldsymbol{\mu}_n(\mathbf{x}_n), \text{diag}(\boldsymbol{\sigma}_n(\mathbf{x}_n)^2))$ is a Gaussian distribution, according to Lemma A.1 and Lemma A.2, we know the optimal solution is

$$\begin{aligned} \boldsymbol{\mu}_n^*(\mathbf{x}_n) &= \mathbb{E}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}] = \tilde{\boldsymbol{\mu}}_n(\mathbf{x}_n, \frac{1}{\sqrt{\bar{\alpha}_n}}(\mathbf{x}_n - \sqrt{\bar{\beta}_n}\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n])), \\ \boldsymbol{\sigma}_n^*(\mathbf{x}_n)^2 &= \text{diag}(\text{Cov}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}]) = \lambda_n^2 \mathbf{1} + \gamma_n^2 \frac{\bar{\beta}_n}{\bar{\alpha}_n} (\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n^2] - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]^2). \end{aligned}$$

□

Theorem 3.2. (Optimal solution to optimization solely w.r.t. mean) For any covariance Σ_n , the optimal mean to problem (12) is always Eq. (4), i.e.,

$$\boldsymbol{\mu}_n^*(\mathbf{x}_n) = \tilde{\boldsymbol{\mu}}_n \left(\mathbf{x}_n, \frac{1}{\sqrt{\bar{\alpha}_n}} (\mathbf{x}_n - \sqrt{\bar{\beta}_n} \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]) \right),$$

which is irrelevant to Σ_n .

Proof. Similarly to the proof in Theorem 3.1, these optimizations are equivalent

$$\max_{\{\boldsymbol{\mu}_n\}_{n=1}^N} L_{\text{elbo}} \Leftrightarrow \min_{\{\boldsymbol{\mu}_n\}_{n=1}^N} D_{\text{KL}}(q(\mathbf{x}_{0:N}) \| p(\mathbf{x}_{0:N})) \Leftrightarrow \min_{\boldsymbol{\mu}_n} \mathbb{E}_q D_{\text{KL}}(q(\mathbf{x}_{n-1}|\mathbf{x}_n) \| p(\mathbf{x}_{n-1}|\mathbf{x}_n)), \quad 1 \leq n \leq N.$$

According to Lemma A.1 and Lemma A.2, we know the optimal $\boldsymbol{\mu}_n^*(\mathbf{x}_n)$ is

$$\boldsymbol{\mu}_n^*(\mathbf{x}_n) = \mathbb{E}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}] = \tilde{\boldsymbol{\mu}}_n \left(\mathbf{x}_n, \frac{1}{\sqrt{\bar{\alpha}_n}} (\mathbf{x}_n - \sqrt{\bar{\beta}_n} \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]) \right).$$

□

Theorem 3.3. (Optimal solution to optimization w.r.t. covariance solely) Suppose $\Sigma_n(\mathbf{x}_n) = \text{diag}(\boldsymbol{\sigma}_n(\mathbf{x}_n)^2)$. For any mean $\boldsymbol{\mu}_n(\mathbf{x}_n)$ that is parameterized by a noise prediction network $\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)$ as in Eq. (6), the optimal covariance $\tilde{\boldsymbol{\sigma}}_n^*(\mathbf{x}_n)^2$ to problem (13) is

$$\begin{aligned} \tilde{\boldsymbol{\sigma}}_n^*(\mathbf{x}_n)^2 &= \boldsymbol{\sigma}_n^*(\mathbf{x}_n)^2 + \gamma_n^2 \frac{\bar{\beta}_n}{\bar{\alpha}_n} \underbrace{(\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n) - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n])^2}_{\text{error}} \\ &= \lambda_n^2 \mathbf{1} + \gamma_n^2 \frac{\bar{\beta}_n}{\bar{\alpha}_n} \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[(\boldsymbol{\epsilon}_n - \hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n))^2], \end{aligned} \quad (14)$$

where $\boldsymbol{\sigma}_n^*(\mathbf{x}_n)^2$ is the optimal covariance to the joint optimization problem in Theorem 3.1, $\boldsymbol{\epsilon}_n = \frac{\mathbf{x}_n - \sqrt{\bar{\alpha}_n} \mathbf{x}_0}{\sqrt{\bar{\beta}_n}}$ is the noise used to generate \mathbf{x}_n from \mathbf{x}_0 , $(\cdot)^2$ is the element-wise square, $\mathbf{1}$ is the vector of ones and $\gamma_n = \sqrt{\bar{\alpha}_{n-1}} - \sqrt{\bar{\beta}_{n-1}} - \lambda_n^2 \sqrt{\frac{\bar{\alpha}_n}{\bar{\beta}_n}}$.

Proof. Similarly to the proof in Theorem 3.1, these optimizations are equivalent

$$\max_{\{\Sigma_n\}_{n=1}^N} L_{\text{elbo}} \Leftrightarrow \min_{\{\Sigma_n\}_{n=1}^N} D_{\text{KL}}(q(\mathbf{x}_{0:N}) \| p(\mathbf{x}_{0:N})) \Leftrightarrow \min_{\Sigma_n} \mathbb{E}_q D_{\text{KL}}(q(\mathbf{x}_{n-1}|\mathbf{x}_n) \| p(\mathbf{x}_{n-1}|\mathbf{x}_n)), \quad 1 \leq n \leq N.$$

According to Lemma A.1, we know the optimal $\tilde{\boldsymbol{\sigma}}_n^*(\mathbf{x}_n)^2$ is

$$\tilde{\boldsymbol{\sigma}}_n^*(\mathbf{x}_n)^2 = \text{diag}(\text{Cov}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}]) + (\boldsymbol{\mu}_n(\mathbf{x}_n) - \mathbb{E}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}])^2.$$

According to Theorem 3.1, $\text{diag}(\text{Cov}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}]) = \boldsymbol{\sigma}_n^*(\mathbf{x}_n)^2$.

According to Lemma A.2 and the parameterization of $\boldsymbol{\mu}_n(\mathbf{x}_n)$ in Eq. (6), we have

$$\boldsymbol{\mu}_n(\mathbf{x}_n) - \mathbb{E}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}] = \gamma_n \sqrt{\frac{\bar{\beta}_n}{\bar{\alpha}_n}} (\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n] - \hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)).$$

Finally,

$$\begin{aligned}
 \tilde{\sigma}_n^*(\mathbf{x}_n)^2 &= \text{diag}(\text{Cov}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}]) + (\boldsymbol{\mu}_n(\mathbf{x}_n) - \mathbb{E}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}])^2 \\
 &= \sigma_n^*(\mathbf{x}_n)^2 + \gamma_n^2 \frac{\bar{\beta}_n}{\bar{\alpha}_n} \underbrace{(\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n) - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n])^2}_{\text{error}} \\
 &= \lambda_n^2 \mathbf{1} + \gamma_n^2 \frac{\bar{\beta}_n}{\bar{\alpha}_n} \{ \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n^2] - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]^2 + \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]^2 - 2\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n) + \hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)^2 \} \\
 &= \lambda_n^2 \mathbf{1} + \gamma_n^2 \frac{\bar{\beta}_n}{\bar{\alpha}_n} \{ \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n^2] - 2\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n) + \hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)^2 \} \\
 &= \lambda_n^2 \mathbf{1} + \gamma_n^2 \frac{\bar{\beta}_n}{\bar{\alpha}_n} \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[(\boldsymbol{\epsilon}_n - \hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n))^2].
 \end{aligned}$$

□

A.2. Proof of Proposition 5.1 and Proposition 5.2

Lemma A.3. *Let $0 \leq s < t \leq T$, and $q(\mathbf{x}_s|\mathbf{x}_t)$ be the conditional distribution of \mathbf{x}_s given \mathbf{x}_t determined by the SDE $d\mathbf{x} = f(t)\mathbf{x}dt + g(t)d\mathbf{w}$. Then we have*

$$\begin{aligned}
 \mathbb{E}_{q(\mathbf{x}_s|\mathbf{x}_t)}[\mathbf{x}_s] &= \frac{1}{\sqrt{\alpha_{t|s}}} \left(\mathbf{x}_t - \frac{\beta_{t|s}}{\sqrt{\beta_{t|0}}} \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\boldsymbol{\epsilon}_t] \right), \\
 \text{Cov}_{q(\mathbf{x}_s|\mathbf{x}_t)}[\mathbf{x}_s] &= \tilde{\beta}_{s|t} \mathbf{I} + \frac{\beta_{t|s}^2}{\beta_{t|0}\alpha_{t|s}} \left(\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\boldsymbol{\epsilon}_t \boldsymbol{\epsilon}_t^\top] - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\boldsymbol{\epsilon}_t] \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\boldsymbol{\epsilon}_t]^\top \right), \\
 \text{diag}(\text{Cov}_{q(\mathbf{x}_s|\mathbf{x}_t)}[\mathbf{x}_s]) &= \tilde{\beta}_{s|t} \mathbf{1} + \frac{\beta_{t|s}^2}{\beta_{t|0}\alpha_{t|s}} \left(\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\boldsymbol{\epsilon}_t^2] - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\boldsymbol{\epsilon}_t]^2 \right),
 \end{aligned}$$

where $\boldsymbol{\epsilon}_t = \frac{\mathbf{x}_t - \sqrt{\alpha_{t|0}}\mathbf{x}_0}{\sqrt{\beta_{t|0}}}$ is the noise used to generate \mathbf{x}_t from \mathbf{x}_0 and $\tilde{\beta}_{s|t} = \frac{\beta_{s|0}}{\beta_{t|0}}\beta_{t|s}$.

Proof. The conditional distribution of \mathbf{x}_s given \mathbf{x}_t and \mathbf{x}_0 determined by the SDE is

$$q(\mathbf{x}_s|\mathbf{x}_t, \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_s | \frac{\sqrt{\alpha_{t|s}}\beta_{s|0}}{\beta_{t|0}}\mathbf{x}_t + \frac{\sqrt{\alpha_{s|0}}\beta_{t|s}}{\beta_{t|0}}\mathbf{x}_0, \tilde{\beta}_{s|t}\mathbf{I}).$$

As for $\mathbb{E}_{q(\mathbf{x}_s|\mathbf{x}_t)}[\mathbf{x}_s]$, we have

$$\mathbb{E}_{q(\mathbf{x}_s|\mathbf{x}_t)}[\mathbf{x}_s] = \frac{1}{\sqrt{\alpha_{t|s}}} (\mathbf{x}_t + \beta_{t|s} \nabla \log q_t(\mathbf{x}_t)) = \frac{1}{\sqrt{\alpha_{t|s}}} (\mathbf{x}_t - \frac{\beta_{t|s}}{\sqrt{\beta_{t|0}}} \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\boldsymbol{\epsilon}_t]).$$

As for $\text{Cov}_{q(\mathbf{x}_s|\mathbf{x}_t)}[\mathbf{x}_s]$, we can expand it using the law of total variance

$$\begin{aligned}
 \text{Cov}_{q(\mathbf{x}_s|\mathbf{x}_t)}[\mathbf{x}_s] &= \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)} \text{Cov}_{q(\mathbf{x}_s|\mathbf{x}_t, \mathbf{x}_0)}[\mathbf{x}_s] + \text{Cov}_{q(\mathbf{x}_0|\mathbf{x}_t)} \mathbb{E}_{q(\mathbf{x}_s|\mathbf{x}_t, \mathbf{x}_0)}[\mathbf{x}_s] \\
 &= \tilde{\beta}_{s|t} \mathbf{I} + \frac{\alpha_{s|0}\beta_{t|s}^2}{\beta_{t|0}^2} \text{Cov}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\mathbf{x}_0] \\
 &= \tilde{\beta}_{s|t} \mathbf{I} + \frac{\alpha_{s|0}\beta_{t|s}^2}{\beta_{t|0}^2} \frac{\beta_{t|0}}{\alpha_{t|0}} \text{Cov}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\boldsymbol{\epsilon}_t] \\
 &= \tilde{\beta}_{s|t} \mathbf{I} + \frac{\beta_{t|s}^2}{\beta_{t|0}\alpha_{t|s}} \left(\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\boldsymbol{\epsilon}_t \boldsymbol{\epsilon}_t^\top] - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\boldsymbol{\epsilon}_t] \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\boldsymbol{\epsilon}_t]^\top \right).
 \end{aligned}$$

Its diagonal is

$$\text{diag}(\text{Cov}_{q(\mathbf{x}_s|\mathbf{x}_t)}[\mathbf{x}_s]) = \tilde{\beta}_{s|t}\mathbf{1} + \frac{\beta_{t|s}^2}{\beta_{t|0}\alpha_{t|s}} (\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\boldsymbol{\epsilon}_t^2] - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\boldsymbol{\epsilon}_t]^2).$$

□

Proposition 5.1. Suppose $\Sigma_{s|t}(\mathbf{x}_t) = \text{diag}(\boldsymbol{\sigma}_{s|t}(\mathbf{x}_t)^2)$. Then the optimal mean to problem (18) is $\boldsymbol{\mu}_{s|t}^*(\mathbf{x}_t)$ as in Eq. (19), and the optimal covariance to problem (18) is

$$\boldsymbol{\sigma}_{s|t}^*(\mathbf{x}_t)^2 = \tilde{\beta}_{s|t}\mathbf{1} + \frac{\beta_{t|s}^2}{\beta_{t|0}\alpha_{t|s}} (\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\boldsymbol{\epsilon}_t^2] - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\boldsymbol{\epsilon}_t]^2),$$

where $\boldsymbol{\epsilon}_t = \frac{\mathbf{x}_t - \sqrt{\alpha_{t|0}}\mathbf{x}_0}{\sqrt{\beta_{t|0}}}$ is the noise used to generate \mathbf{x}_t from \mathbf{x}_0 and $\tilde{\beta}_{s|t} = \frac{\beta_{s|0}}{\beta_{t|0}}\beta_{t|s}$.

Proof. According to Lemma A.1 and Lemma A.3, the optimal solution is

$$\begin{aligned} \boldsymbol{\mu}_{s|t}^*(\mathbf{x}_t) &= \mathbb{E}_{q(\mathbf{x}_s|\mathbf{x}_t)}[\mathbf{x}_s] = \frac{1}{\sqrt{\alpha_{t|s}}} \left(\mathbf{x}_t - \frac{\beta_{t|s}}{\sqrt{\beta_{t|0}}} \mathbb{E}_{q_0|t}(\mathbf{x}_0|\mathbf{x}_t)[\boldsymbol{\epsilon}_t] \right), \\ \boldsymbol{\sigma}_{s|t}^*(\mathbf{x}_t)^2 &= \text{diag}(\text{Cov}_{q(\mathbf{x}_s|\mathbf{x}_t)}[\mathbf{x}_s]) = \tilde{\beta}_{s|t}\mathbf{1} + \frac{\beta_{t|s}^2}{\beta_{t|0}\alpha_{t|s}} (\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\boldsymbol{\epsilon}_t^2] - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\boldsymbol{\epsilon}_t]^2). \end{aligned}$$

□

Proposition 5.2. Suppose $\Sigma_{s|t}(\mathbf{x}_t) = \text{diag}(\boldsymbol{\sigma}_{s|t}(\mathbf{x}_t)^2)$. For any mean $\boldsymbol{\mu}_{s|t}(\mathbf{x}_t)$ that is parameterized by a noise prediction network $\hat{\boldsymbol{\epsilon}}_t(\mathbf{x}_t)$ as in Eq. (20), the optimal covariance $\tilde{\boldsymbol{\sigma}}_{s|t}^*(\mathbf{x}_t)^2$ to problem (22) is

$$\begin{aligned} \tilde{\boldsymbol{\sigma}}_{s|t}^*(\mathbf{x}_t)^2 &= \boldsymbol{\sigma}_{s|t}^*(\mathbf{x}_t)^2 + \frac{\beta_{t|s}^2}{\beta_{t|0}\alpha_{t|s}} \underbrace{(\hat{\boldsymbol{\epsilon}}_t(\mathbf{x}_t) - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\boldsymbol{\epsilon}_t])^2}_{\text{error}} \\ &= \tilde{\beta}_{s|t}\mathbf{1} + \frac{\beta_{t|s}^2}{\beta_{t|0}\alpha_{t|s}} \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[(\boldsymbol{\epsilon}_t - \hat{\boldsymbol{\epsilon}}_t(\mathbf{x}_t))^2], \end{aligned}$$

where $\boldsymbol{\epsilon}_t = \frac{\mathbf{x}_t - \sqrt{\alpha_{t|0}}\mathbf{x}_0}{\sqrt{\beta_{t|0}}}$ is the noise used to generate \mathbf{x}_t from \mathbf{x}_0 and $\tilde{\beta}_{s|t} = \frac{\beta_{s|0}}{\beta_{t|0}}\beta_{t|s}$.

Proof. According to Lemma A.1 and Proposition 5.1, the optimal covariance $\tilde{\boldsymbol{\sigma}}_{s|t}^*(\mathbf{x}_t)^2$ is

$$\begin{aligned} \tilde{\boldsymbol{\sigma}}_{s|t}^*(\mathbf{x}_t)^2 &= \text{diag}(\text{Cov}_{q(\mathbf{x}_s|\mathbf{x}_t)}[\mathbf{x}_s]) + (\boldsymbol{\mu}_{s|t}(\mathbf{x}_t) - \mathbb{E}_{q(\mathbf{x}_s|\mathbf{x}_t)}[\mathbf{x}_s])^2 \\ &= \boldsymbol{\sigma}_{s|t}^*(\mathbf{x}_t)^2 + \frac{\beta_{t|s}^2}{\beta_{t|0}\alpha_{t|s}} \underbrace{(\hat{\boldsymbol{\epsilon}}_t(\mathbf{x}_t) - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\boldsymbol{\epsilon}_t])^2}_{\text{error}} \\ &= \tilde{\beta}_{s|t}\mathbf{1} + \frac{\beta_{t|s}^2}{\beta_{t|0}\alpha_{t|s}} \{ \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\boldsymbol{\epsilon}_t^2] - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\boldsymbol{\epsilon}_t]^2 + \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\boldsymbol{\epsilon}_t]^2 - 2\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\boldsymbol{\epsilon}_t]\hat{\boldsymbol{\epsilon}}_t(\mathbf{x}_t) + \hat{\boldsymbol{\epsilon}}_t(\mathbf{x}_t)^2 \} \\ &= \tilde{\beta}_{s|t}\mathbf{1} + \frac{\beta_{t|s}^2}{\beta_{t|0}\alpha_{t|s}} \{ \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\boldsymbol{\epsilon}_t^2] - 2\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\boldsymbol{\epsilon}_t]\hat{\boldsymbol{\epsilon}}_t(\mathbf{x}_t) + \hat{\boldsymbol{\epsilon}}_t(\mathbf{x}_t)^2 \} \\ &= \tilde{\beta}_{s|t}\mathbf{1} + \frac{\beta_{t|s}^2}{\beta_{t|0}\alpha_{t|s}} \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[(\boldsymbol{\epsilon}_t - \hat{\boldsymbol{\epsilon}}_t(\mathbf{x}_t))^2]. \end{aligned}$$

□

A.3. Strictly Improved Expressive Power when the Data Distribution is a Mixture of Gaussian

Lemma A.4. Suppose $q(\mathbf{x}, \mathbf{y})$ is a probability density function, s.t., $q(\mathbf{x}) = \sum_{j=1}^J \gamma_j \mathcal{N}(\boldsymbol{\mu}_j, c\mathbf{I})$ with $J \geq 2$ and pairwise distinct $\{\boldsymbol{\mu}_j\}_{j=1}^J$, and $q(\mathbf{y}|\mathbf{x}) = \mathcal{N}(\mathbf{y}|\sqrt{\alpha}\mathbf{x}, \beta\mathbf{I})$ for some $\alpha, \beta > 0$. Suppose $p(\mathbf{x}|\mathbf{y}) = \mathcal{N}(\mathbf{x}|\boldsymbol{\mu}(\mathbf{y}), \boldsymbol{\Sigma}(\mathbf{y}))$ is a Gaussian distribution conditioned on \mathbf{y} . Then we have,

$$\min_{\boldsymbol{\mu}, \boldsymbol{\Sigma}: \boldsymbol{\Sigma}(\mathbf{y}) = \text{diag}(\boldsymbol{\sigma}(\mathbf{y})^2)} \mathbb{E}_{q(\mathbf{y})} D_{\text{KL}}(q(\mathbf{x}|\mathbf{y}) \| p(\mathbf{x}|\mathbf{y})) < \min_{\boldsymbol{\mu}, \boldsymbol{\Sigma}: \boldsymbol{\Sigma}(\mathbf{y}) = \sigma^2 \mathbf{I}} \mathbb{E}_{q(\mathbf{y})} D_{\text{KL}}(q(\mathbf{x}|\mathbf{y}) \| p(\mathbf{x}|\mathbf{y})).$$

Proof. As mentioned in the proof of Lemma A.1, the KL divergence between a probability density function q and a Gaussian distribution p can be written as $D_{\text{KL}}(q \| p) = D_{\text{KL}}(\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_q, \boldsymbol{\Sigma}_q) \| p) + \text{const}$, where $\boldsymbol{\mu}_q$ is the mean of q , $\boldsymbol{\Sigma}_q$ is the covariance matrix of q and const is only related to q . Thus, we only need to prove

$$\begin{aligned} & \min_{\boldsymbol{\mu}, \boldsymbol{\Sigma}: \boldsymbol{\Sigma}(\mathbf{y}) = \text{diag}(\boldsymbol{\sigma}(\mathbf{y})^2)} \mathbb{E}_{q(\mathbf{y})} D_{\text{KL}}(\mathcal{N}(\mathbf{x}|\mathbb{E}_{q(\mathbf{x}|\mathbf{y})}[\mathbf{x}], \text{Cov}_{q(\mathbf{x}|\mathbf{y})}[\mathbf{x}]) \| p(\mathbf{x}|\mathbf{y})) \\ & < \min_{\boldsymbol{\mu}, \boldsymbol{\Sigma}: \boldsymbol{\Sigma}(\mathbf{y}) = \sigma^2 \mathbf{I}} \mathbb{E}_{q(\mathbf{y})} D_{\text{KL}}(\mathcal{N}(\mathbf{x}|\mathbb{E}_{q(\mathbf{x}|\mathbf{y})}[\mathbf{x}], \text{Cov}_{q(\mathbf{x}|\mathbf{y})}[\mathbf{x}]) \| p(\mathbf{x}|\mathbf{y})). \end{aligned}$$

We know the optimal mean is $\boldsymbol{\mu}^*(\mathbf{y}) = \mathbb{E}_{q(\mathbf{x}|\mathbf{y})}[\mathbf{x}]$ for both problems. Let $\mathbf{M}(\mathbf{y}) = \text{Cov}_{q(\mathbf{x}|\mathbf{y})}[\mathbf{x}]$. We expand the KL divergence, and substitute $\boldsymbol{\mu}^*$ into $\boldsymbol{\mu}$

$$2D_{\text{KL}}(\mathcal{N}(\mathbf{x}|\mathbb{E}_{q(\mathbf{x}|\mathbf{y})}[\mathbf{x}], \text{Cov}_{q(\mathbf{x}|\mathbf{y})}[\mathbf{x}]) \| p(\mathbf{x}|\mathbf{y})) = \text{tr}(\boldsymbol{\Sigma}(\mathbf{y})^{-1} \mathbf{M}(\mathbf{y})) - d + \log \frac{|\boldsymbol{\Sigma}(\mathbf{y})|}{|\text{Cov}_{q(\mathbf{x}|\mathbf{y})}[\mathbf{x}]|}.$$

Thus, we only need to prove

$$\min_{\boldsymbol{\Sigma}: \boldsymbol{\Sigma}(\mathbf{y}) = \text{diag}(\boldsymbol{\sigma}(\mathbf{y})^2)} \mathbb{E}_{q(\mathbf{y})} [\text{tr}(\boldsymbol{\Sigma}(\mathbf{y})^{-1} \mathbf{M}(\mathbf{y})) + \log |\boldsymbol{\Sigma}(\mathbf{y})|] < \min_{\boldsymbol{\Sigma}: \boldsymbol{\Sigma}(\mathbf{y}) = \sigma^2 \mathbf{I}} \mathbb{E}_{q(\mathbf{y})} [\text{tr}(\boldsymbol{\Sigma}(\mathbf{y})^{-1} \mathbf{M}(\mathbf{y})) + \log |\boldsymbol{\Sigma}(\mathbf{y})|],$$

which is equivalent to

$$\min_{\boldsymbol{\sigma}(\cdot)^2} \mathbb{E}_{q(\mathbf{y})} \left[\sum_{i=1}^d (\boldsymbol{\sigma}(\mathbf{y})_i^{-2} \mathbf{M}(\mathbf{y})_{ii} + \log \boldsymbol{\sigma}(\mathbf{y})_i^2) \right] < \min_{\sigma^2} \mathbb{E}_{q(\mathbf{y})} \left[\sum_{i=1}^d (\sigma^{-2} \mathbf{M}(\mathbf{y})_{ii} + \log \sigma^2) \right].$$

The left problem admits the optimal solution $\boldsymbol{\sigma}^*(\mathbf{y})_i^2 = \mathbf{M}(\mathbf{y})_{ii}$, and the right one admits the optimal solution $\sigma^{*2} = \mathbb{E}_{q(\mathbf{y})} \frac{\text{tr}(\mathbf{M}(\mathbf{y}))}{d}$. Thus, we only need to prove

$$\mathbb{E}_{q(\mathbf{y})} \left[\sum_{i=1}^d (\boldsymbol{\sigma}^*(\mathbf{y})_i^{-2} \mathbf{M}(\mathbf{y})_{ii} + \log \boldsymbol{\sigma}^*(\mathbf{y})_i^2) \right] < \mathbb{E}_{q(\mathbf{y})} \left[\sum_{i=1}^d (\sigma^{*-2} \mathbf{M}(\mathbf{y})_{ii} + \log \sigma^{*2}) \right]. \quad (25)$$

We prove it by contradiction. Suppose

$$\mathbb{E}_{q(\mathbf{y})} \left[\sum_{i=1}^d (\boldsymbol{\sigma}^*(\mathbf{y})_i^{-2} \mathbf{M}(\mathbf{y})_{ii} + \log \boldsymbol{\sigma}^*(\mathbf{y})_i^2) \right] = \mathbb{E}_{q(\mathbf{y})} \left[\sum_{i=1}^d (\sigma^{*-2} \mathbf{M}(\mathbf{y})_{ii} + \log \sigma^{*2}) \right].$$

Since $q(\mathbf{y})$ is fully supported, it implies that

$$\forall \mathbf{y}, \sum_{i=1}^d (\boldsymbol{\sigma}^*(\mathbf{y})_i^{-2} \mathbf{M}(\mathbf{y})_{ii} + \log \boldsymbol{\sigma}^*(\mathbf{y})_i^2) = \sum_{i=1}^d (\sigma^{*-2} \mathbf{M}(\mathbf{y})_{ii} + \log \sigma^{*2}),$$

which holds if and only if $\forall \mathbf{y}, \forall i, \boldsymbol{\sigma}^*(\mathbf{y})_i^2 = \sigma^{*2}$. This means $\forall \mathbf{y}, \forall i, \mathbf{M}(\mathbf{y})_{ii} = \mathbb{E}_{q(\mathbf{y})} \frac{\text{tr}(\mathbf{M}(\mathbf{y}))}{d}$, which implies that there exists $m \in \mathbb{R}$ irrelevant to \mathbf{y} , s.t., $\text{diag}(\text{Cov}_{q(\mathbf{x}|\mathbf{y})}[\mathbf{x}]) = m\mathbf{I}$.

Now we derive $q(\mathbf{x}|\mathbf{y})$,

$$\begin{aligned}
 q(\mathbf{x}|\mathbf{y}) &\propto q(\mathbf{x})q(\mathbf{y}|\mathbf{x}) \propto \sum_{j=1}^J \gamma_j \exp\left(-\frac{\|\mathbf{x} - \boldsymbol{\mu}_j\|_2^2}{2c}\right) \exp\left(-\frac{\|\mathbf{y} - \sqrt{\alpha}\mathbf{x}\|_2^2}{2\beta}\right) \\
 &\propto \sum_{j=1}^J \gamma_j \exp\left(-\frac{\|\boldsymbol{\mu}_j\|_2^2}{2c}\right) \exp\left(-\frac{(c\alpha + \beta)\|\mathbf{x}\|_2^2 - 2\langle \mathbf{x}, \beta\boldsymbol{\mu}_j + c\sqrt{\alpha}\mathbf{y} \rangle}{2c\beta}\right) \\
 &= \sum_{j=1}^J \gamma_j \exp\left(\frac{\|\beta\boldsymbol{\mu}_j + c\sqrt{\alpha}\mathbf{y}\|_2^2}{2c\beta(c\alpha + \beta)} - \frac{\|\boldsymbol{\mu}_j\|_2^2}{2c}\right) \exp\left(-\frac{\|\mathbf{x} - (\beta\boldsymbol{\mu}_j + c\sqrt{\alpha}\mathbf{y})/(c\alpha + \beta)\|_2^2}{2c\beta/(c\alpha + \beta)}\right) \\
 &\propto \sum_{j=1}^J \exp\left(\log \gamma_j - \frac{\alpha}{c\alpha + \beta} \frac{\|\boldsymbol{\mu}_j\|_2^2}{2} + \frac{\sqrt{\alpha}}{c\alpha + \beta} \boldsymbol{\mu}_j^\top \mathbf{y}\right) \mathcal{N}\left(\frac{(\beta\boldsymbol{\mu}_j + c\sqrt{\alpha}\mathbf{y})}{(c\alpha + \beta)}, c\beta/(c\alpha + \beta) \mathbf{I}\right).
 \end{aligned}$$

Let $\xi_j = \log \gamma_j - \frac{\alpha}{c\alpha + \beta} \frac{\|\boldsymbol{\mu}_j\|_2^2}{2}$. Let $\boldsymbol{\phi}(\mathbf{y}) = (\phi_1(\mathbf{y}), \dots, \phi_J(\mathbf{y}))$, where $\phi_j(\mathbf{y}) = \xi_j + \frac{\sqrt{\alpha}}{c\alpha + \beta} \boldsymbol{\mu}_j^\top \mathbf{y}$, and $\boldsymbol{\eta}(\mathbf{y}) = \text{softmax}(\boldsymbol{\phi}(\mathbf{y}))$. Let $\boldsymbol{\nu}_j(\mathbf{y}) = (\beta\boldsymbol{\mu}_j + c\sqrt{\alpha}\mathbf{y})/(c\alpha + \beta)$, $\bar{\boldsymbol{\mu}}(\mathbf{y}) = \sum_{j=1}^J \eta_j(\mathbf{y})\boldsymbol{\mu}_j$ and $\bar{\boldsymbol{\nu}}(\mathbf{y}) = \sum_{j=1}^J \eta_j(\mathbf{y})\boldsymbol{\nu}_j(\mathbf{y}) = (\beta\bar{\boldsymbol{\mu}}(\mathbf{y}) + c\sqrt{\alpha}\mathbf{y})/(c\alpha + \beta)$. Then

$$q(\mathbf{x}|\mathbf{y}) = \sum_{j=1}^J \eta_j(\mathbf{y}) \mathcal{N}(\boldsymbol{\nu}_j(\mathbf{y}), c\beta/(c\alpha + \beta) \mathbf{I}),$$

which is also a mixture of Gaussian. According to the property of mixture of Gaussian, the diagonal of its covariance is

$$\text{diag}(\text{Cov}_{q(\mathbf{x}|\mathbf{y})}[\mathbf{x}]) = c\beta/(c\alpha + \beta) \mathbf{1} + \sum_{j=1}^J \eta_j(\mathbf{y}) \boldsymbol{\nu}_j(\mathbf{y})^2 - \bar{\boldsymbol{\nu}}(\mathbf{y})^2.$$

Meanwhile, $\text{diag}(\text{Cov}_{q(\mathbf{x}|\mathbf{y})}[\mathbf{x}]) = m\mathbf{I}$. Thus,

$$c\beta/(c\alpha + \beta)d + \sum_{j=1}^J \eta_j(\mathbf{y}) \|\boldsymbol{\nu}_j(\mathbf{y})\|_2^2 - \|\bar{\boldsymbol{\nu}}(\mathbf{y})\|_2^2 = dm,$$

which is irrelevant to \mathbf{y} . Since

$$\sum_{j=1}^J \eta_j(\mathbf{y}) \|\boldsymbol{\nu}_j(\mathbf{y})\|_2^2 - \|\bar{\boldsymbol{\nu}}(\mathbf{y})\|_2^2 = \beta^2/(c\alpha + \beta)^2 \left[\sum_{j=1}^J \eta_j(\mathbf{y}) \|\boldsymbol{\mu}_j\|_2^2 - \|\bar{\boldsymbol{\mu}}(\mathbf{y})\|_2^2 \right],$$

$\sum_{j=1}^J \eta_j(\mathbf{y}) \|\boldsymbol{\mu}_j\|_2^2 - \|\bar{\boldsymbol{\mu}}(\mathbf{y})\|_2^2$ is also irrelevant to \mathbf{y} . Let $\mathbf{y} = t\mathbf{y}_0$, $j_0 = \arg \max_j \boldsymbol{\mu}_j^\top \mathbf{y}_0$. Then $\lim_{t \rightarrow \infty} \eta_{j_0}(t\mathbf{y}_0) = 1$, and

$$\lim_{t \rightarrow \infty} \sum_{j=1}^J \eta_j(t\mathbf{y}_0) \|\boldsymbol{\mu}_j\|_2^2 - \|\bar{\boldsymbol{\mu}}(t\mathbf{y}_0)\|_2^2 = 0.$$

Thus, $\sum_{j=1}^J \eta_j(\mathbf{y}) \|\boldsymbol{\mu}_j\|_2^2 - \|\bar{\boldsymbol{\mu}}(\mathbf{y})\|_2^2 = 0$ for all \mathbf{y} . This means

$$\sum_{j=1}^J \eta_j(\mathbf{y}) \|\boldsymbol{\mu}_j\|_2^2 - \|\bar{\boldsymbol{\mu}}(\mathbf{y})\|_2^2 = \sum_{j=1}^J \eta_j(\mathbf{y}) \|\boldsymbol{\mu}_j - \bar{\boldsymbol{\mu}}(\mathbf{y})\|_2^2 = 0.$$

Thus, $\boldsymbol{\mu}_1 = \boldsymbol{\mu}_2 = \dots = \boldsymbol{\mu}_J = \bar{\boldsymbol{\mu}}(\mathbf{y})$. This contradicts the assumption that $\{\boldsymbol{\mu}_j\}_{j=1}^J$ are pairwise distinct. Therefore, Eq. (25) must hold. \square

Proposition A.5. *If the data distribution $q(\mathbf{x}_0) = \sum_{j=1}^J \gamma_j \mathcal{N}(\boldsymbol{\mu}_j, c\mathbf{I})$ is a mixture of $J \geq 2$ Gaussian, and $\{\boldsymbol{\mu}_j\}_{j=1}^J$ are pairwise distinct, then we have*

$$\max_{\boldsymbol{\mu}_n, \boldsymbol{\Sigma}_n: \boldsymbol{\Sigma}_n(\mathbf{x}_n) = \text{diag}(\sigma_n(\mathbf{x}_n)^2)} L_{\text{elbo}} > \max_{\boldsymbol{\mu}_n, \boldsymbol{\Sigma}_n: \boldsymbol{\Sigma}_n(\mathbf{x}_n) = \sigma_n^2 \mathbf{I}} L_{\text{elbo}}.$$

Proof. We only need to prove

$$\min_{\boldsymbol{\mu}_n, \boldsymbol{\Sigma}_n: \boldsymbol{\Sigma}_n(\mathbf{x}_n) = \text{diag}(\sigma_n(\mathbf{x}_n)^2)} D_{\text{KL}}(q(\mathbf{x}_{0:N}) \| p(\mathbf{x}_{0:N})) < \min_{\boldsymbol{\mu}_n, \boldsymbol{\Sigma}_n: \boldsymbol{\Sigma}_n(\mathbf{x}_n) = \sigma_n^2 \mathbf{I}} D_{\text{KL}}(q(\mathbf{x}_{0:N}) \| p(\mathbf{x}_{0:N})).$$

As mentioned in the proof of Theorem 3.1, $D_{\text{KL}}(q(\mathbf{x}_{0:N}) \| p(\mathbf{x}_{0:N}))$ can be decomposed as

$$D_{\text{KL}}(q(\mathbf{x}_{0:N}) \| p(\mathbf{x}_{0:N})) = \sum_{n=1}^N \mathbb{E}_q D_{\text{KL}}(q(\mathbf{x}_{n-1} | \mathbf{x}_n) \| p(\mathbf{x}_{n-1} | \mathbf{x}_n)) + c,$$

where c is a constant only related to q . Thus, we only need to prove

$$\begin{aligned} & \sum_{n=1}^N \min_{\boldsymbol{\mu}_n, \boldsymbol{\Sigma}_n: \boldsymbol{\Sigma}_n(\mathbf{x}_n) = \text{diag}(\sigma_n(\mathbf{x}_n)^2)} \mathbb{E}_q D_{\text{KL}}(q(\mathbf{x}_{n-1} | \mathbf{x}_n) \| p(\mathbf{x}_{n-1} | \mathbf{x}_n)) \\ & < \sum_{n=1}^N \min_{\boldsymbol{\mu}_n, \boldsymbol{\Sigma}_n: \boldsymbol{\Sigma}_n(\mathbf{x}_n) = \sigma_n^2 \mathbf{I}} \mathbb{E}_q D_{\text{KL}}(q(\mathbf{x}_{n-1} | \mathbf{x}_n) \| p(\mathbf{x}_{n-1} | \mathbf{x}_n)). \end{aligned}$$

It is sufficient to prove when $n = 1$,

$$\begin{aligned} & \min_{\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1: \boldsymbol{\Sigma}_1(\mathbf{x}_1) = \text{diag}(\sigma_1(\mathbf{x}_1)^2)} \mathbb{E}_q D_{\text{KL}}(q(\mathbf{x}_0 | \mathbf{x}_1) \| p(\mathbf{x}_0 | \mathbf{x}_1)) \\ & < \min_{\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1: \boldsymbol{\Sigma}_1(\mathbf{x}_1) = \sigma_1^2 \mathbf{I}} \mathbb{E}_q D_{\text{KL}}(q(\mathbf{x}_0 | \mathbf{x}_1) \| p(\mathbf{x}_0 | \mathbf{x}_1)), \end{aligned}$$

which holds according to Lemma A.4. □

B. Results for Full Covariances

In this part, we derive results on full covariances. In Proposition B.2, we derive the optimal solution to the joint optimization problem (3). In Proposition B.3, we derive the optimal covariance to the optimization problem (13) w.r.t. the covariance solely. Their proof requires Lemma B.1, which is very similar to Lemma A.1.

Lemma B.1. *Suppose $q(\mathbf{x})$ is a probability density function with mean $\boldsymbol{\mu}_q$ and covariance matrix $\boldsymbol{\Sigma}_q$ and $p(\mathbf{x}) = \mathcal{N}(\mathbf{x} | \boldsymbol{\mu}, \boldsymbol{\Sigma})$ is a Gaussian distribution. Then*

1. *for any $\boldsymbol{\Sigma}$, the optimization problem $\min_{\boldsymbol{\mu}} D_{\text{KL}}(q \| p)$ admits an optimal solution $\boldsymbol{\mu}^* = \boldsymbol{\mu}_q$;*
2. *for any $\boldsymbol{\mu}$, the optimization problem $\min_{\boldsymbol{\Sigma}} D_{\text{KL}}(q \| p)$ admits an optimal solution $\tilde{\boldsymbol{\Sigma}}^* = \boldsymbol{\Sigma}_q + (\boldsymbol{\mu} - \boldsymbol{\mu}_q)(\boldsymbol{\mu} - \boldsymbol{\mu}_q)^\top$*
3. *the optimization problem $\min_{\boldsymbol{\mu}, \boldsymbol{\Sigma}} D_{\text{KL}}(q \| p)$ admits an optimal solution $\boldsymbol{\mu}^* = \boldsymbol{\mu}_q$ and $\boldsymbol{\Sigma}^* = \boldsymbol{\Sigma}_q$.*

Proof. Similarly to the proof of Lemma A.1, the optimizations of $D_{\text{KL}}(q \| p)$ and $D_{\text{KL}}(\mathcal{N}(\mathbf{x} | \boldsymbol{\mu}_q, \boldsymbol{\Sigma}_q) \| p)$ w.r.t. $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are equivalent. We expand $D_{\text{KL}}(\mathcal{N}(\mathbf{x} | \boldsymbol{\mu}_q, \boldsymbol{\Sigma}_q) \| p)$ according to the definition of KL divergence:

$$\begin{aligned} 2D_{\text{KL}}(\mathcal{N}(\mathbf{x} | \boldsymbol{\mu}_q, \boldsymbol{\Sigma}_q) \| p) &= \text{tr}(\boldsymbol{\Sigma}^{-1} \boldsymbol{\Sigma}_q) - d + \log \frac{|\boldsymbol{\Sigma}|}{|\boldsymbol{\Sigma}_q|} + (\boldsymbol{\mu} - \boldsymbol{\mu}_q)^\top \boldsymbol{\Sigma}^{-1} (\boldsymbol{\mu} - \boldsymbol{\mu}_q) \\ &= \text{tr}(\boldsymbol{\Sigma}^{-1} (\boldsymbol{\Sigma}_q + (\boldsymbol{\mu} - \boldsymbol{\mu}_q)(\boldsymbol{\mu} - \boldsymbol{\mu}_q)^\top)) + \log \frac{|\boldsymbol{\Sigma}|}{|\boldsymbol{\Sigma}_q|} - d. \end{aligned}$$

Thus, $\min_{\boldsymbol{\mu}} D_{\text{KL}}(q||p) \Leftrightarrow \min_{\boldsymbol{\mu}} (\boldsymbol{\mu} - \boldsymbol{\mu}_q)^\top \boldsymbol{\Sigma}^{-1} (\boldsymbol{\mu} - \boldsymbol{\mu}_q)$, which admits an optimal solution $\boldsymbol{\mu}^* = \boldsymbol{\mu}_q$, and $\min_{\boldsymbol{\Sigma}} D_{\text{KL}}(q||p) \Leftrightarrow \min_{\boldsymbol{\Sigma}} \text{tr}(\boldsymbol{\Sigma}^{-1} (\boldsymbol{\Sigma}_q + (\boldsymbol{\mu} - \boldsymbol{\mu}_q)(\boldsymbol{\mu} - \boldsymbol{\mu}_q)^\top)) + \log |\boldsymbol{\Sigma}|$, which admits an optimal solution $\tilde{\boldsymbol{\Sigma}}^* = \boldsymbol{\Sigma}_q + (\boldsymbol{\mu} - \boldsymbol{\mu}_q)(\boldsymbol{\mu} - \boldsymbol{\mu}_q)^\top$. Combining $\boldsymbol{\mu}^*$ and $\tilde{\boldsymbol{\Sigma}}^*$, we know that $\min_{\boldsymbol{\mu}, \boldsymbol{\Sigma}} D_{\text{KL}}(q||p)$ admits an optimal solution $\boldsymbol{\mu}^* = \boldsymbol{\mu}_q$ and $\boldsymbol{\Sigma}^* = \boldsymbol{\Sigma}_q$. \square

Proposition B.2. *Suppose $\boldsymbol{\Sigma}_n(\mathbf{x}_n)$ is a full covariance. Then the optimal mean to problem (3) is $\boldsymbol{\mu}_n^*(\mathbf{x}_n)$ as in Eq. (4), and the optimal covariance to problem (3) is*

$$\boldsymbol{\Sigma}_n^*(\mathbf{x}_n) = \lambda_n^2 \mathbf{I} + \gamma_n^2 \frac{\bar{\beta}_n}{\bar{\alpha}_n} \left(\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n \boldsymbol{\epsilon}_n^\top] - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n] \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]^\top \right).$$

Proof. Bao et al. (2022) show that the KL divergence between a joint probability density function $q(\mathbf{x}_{0:N})$ and a Markov chain $p(\mathbf{x}_{0:N})$ can be written as

$$D_{\text{KL}}(q(\mathbf{x}_{0:N})||p(\mathbf{x}_{0:N})) = \sum_{n=1}^N \mathbb{E}_q D_{\text{KL}}(q(\mathbf{x}_{n-1}|\mathbf{x}_n)||p(\mathbf{x}_{n-1}|\mathbf{x}_n)) + c,$$

where c is a constant only related to q (see Lemma 8 in Bao et al. (2022)). As a result, the optimization of the KL divergence is decomposed into n independent optimization sub-problems:

$$\min_{\boldsymbol{\mu}_n(\cdot), \boldsymbol{\Sigma}_n(\cdot)} \mathbb{E}_q D_{\text{KL}}(q(\mathbf{x}_{n-1}|\mathbf{x}_n)||p(\mathbf{x}_{n-1}|\mathbf{x}_n)), \quad 1 \leq n \leq N.$$

Since $q(\mathbf{x}_{n-1}|\mathbf{x}_n)$ is a probability density function with mean $\mathbb{E}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}]$ and covariance matrix $\text{Cov}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}]$ and $p(\mathbf{x}_{n-1}|\mathbf{x}_n) = \mathcal{N}(\mathbf{x}_{n-1}|\boldsymbol{\mu}_n(\mathbf{x}_n), \boldsymbol{\Sigma}_n(\mathbf{x}_n))$ is a Gaussian distribution, according to Lemma B.1 and Lemma A.2, we know the optimal solution is

$$\begin{aligned} \boldsymbol{\mu}_n^*(\mathbf{x}_n) &= \mathbb{E}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}] = \tilde{\boldsymbol{\mu}}_n(\mathbf{x}_n, \frac{1}{\sqrt{\bar{\alpha}_n}}(\mathbf{x}_n - \sqrt{\bar{\beta}_n} \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n])), \\ \boldsymbol{\Sigma}_n^*(\mathbf{x}_n) &= \text{Cov}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}] = \lambda_n^2 \mathbf{I} + \gamma_n^2 \frac{\bar{\beta}_n}{\bar{\alpha}_n} \left(\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n \boldsymbol{\epsilon}_n^\top] - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n] \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]^\top \right). \end{aligned}$$

\square

Proposition B.3. *Suppose $\boldsymbol{\Sigma}_n(\mathbf{x}_n)$ is a full covariance. For any mean $\boldsymbol{\mu}_n(\mathbf{x}_n)$ that is parameterized by a noise prediction network $\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)$ as in Eq. (6), the optimal covariance $\tilde{\boldsymbol{\Sigma}}_n^*(\mathbf{x}_n)$ to problem (13) is*

$$\begin{aligned} \tilde{\boldsymbol{\Sigma}}_n^*(\mathbf{x}_n) &= \boldsymbol{\Sigma}_n^*(\mathbf{x}_n) + \gamma_n^2 \frac{\bar{\beta}_n}{\bar{\alpha}_n} \underbrace{\left(\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n] - \hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n) \right) \left(\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n] - \hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n) \right)^\top}_{\text{error}} \\ &= \lambda_n^2 \mathbf{I} + \gamma_n^2 \frac{\bar{\beta}_n}{\bar{\alpha}_n} \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)} [(\boldsymbol{\epsilon}_n - \hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n))(\boldsymbol{\epsilon}_n - \hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n))^\top]. \end{aligned}$$

Proof. According to Lemma B.1, we know the optimal $\tilde{\boldsymbol{\Sigma}}_n^*(\mathbf{x}_n)$ is

$$\tilde{\boldsymbol{\Sigma}}_n^*(\mathbf{x}_n) = \text{Cov}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}] + (\boldsymbol{\mu}_n(\mathbf{x}_n) - \mathbb{E}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}])(\boldsymbol{\mu}_n(\mathbf{x}_n) - \mathbb{E}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}])^\top.$$

According to Proposition B.2, $\text{Cov}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}] = \boldsymbol{\Sigma}_n^*(\mathbf{x}_n)$.

According to Lemma A.2 and the parameterization of $\boldsymbol{\mu}_n(\mathbf{x}_n)$ in Eq. (6), we have

$$\boldsymbol{\mu}_n(\mathbf{x}_n) - \mathbb{E}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}] = \gamma_n \sqrt{\frac{\bar{\beta}_n}{\bar{\alpha}_n}} (\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n] - \hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)). \quad (26)$$

Finally,

$$\begin{aligned}
 \tilde{\Sigma}_n^*(\mathbf{x}_n) &= \text{Cov}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}] + (\boldsymbol{\mu}_n(\mathbf{x}_n) - \mathbb{E}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}])(\boldsymbol{\mu}_n(\mathbf{x}_n) - \mathbb{E}_{q(\mathbf{x}_{n-1}|\mathbf{x}_n)}[\mathbf{x}_{n-1}])^\top \\
 &= \Sigma_n^*(\mathbf{x}_n) + \gamma_n^2 \frac{\bar{\beta}_n}{\bar{\alpha}_n} \underbrace{(\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n] - \hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n))(\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n] - \hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n))^\top}_{\text{error}} \\
 &= \lambda_n^2 \mathbf{I} + \gamma_n^2 \frac{\bar{\beta}_n}{\bar{\alpha}_n} \{ \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n \boldsymbol{\epsilon}_n^\top] - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n] \hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)^\top - \hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n) \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]^\top + \hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n) \hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)^\top \} \\
 &= \lambda_n^2 \mathbf{I} + \gamma_n^2 \frac{\bar{\beta}_n}{\bar{\alpha}_n} \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n \boldsymbol{\epsilon}_n^\top - \boldsymbol{\epsilon}_n \hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)^\top - \hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n) \boldsymbol{\epsilon}_n^\top + \hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n) \hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)^\top] \\
 &= \lambda_n^2 \mathbf{I} + \gamma_n^2 \frac{\bar{\beta}_n}{\bar{\alpha}_n} \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[(\boldsymbol{\epsilon}_n - \hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n))(\boldsymbol{\epsilon}_n - \hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n))^\top].
 \end{aligned}$$

□

Note that $\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n \boldsymbol{\epsilon}_n^\top]$ in Proposition B.2 and $\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[(\boldsymbol{\epsilon}_n - \hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n))(\boldsymbol{\epsilon}_n - \hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n))^\top]$ in Proposition B.3 are matrices of shape $d \times d$. We can also learn these conditional expectations by training neural networks that output matrices of shape $d \times d$ on MSE losses, and then estimate $\Sigma_n^*(\mathbf{x}_n)$ and $\tilde{\Sigma}_n^*(\mathbf{x}_n)$ using these neural networks.

However, obtaining samples from Gaussian transitions with full covariances requires a decomposition of the covariance matrix, e.g., Cholesky decomposition, which has a $\mathcal{O}(d^3)$ time complexity. In practice, d can be large (e.g., high resolution images), and such a decomposition can be time consuming. Therefore, in practical settings, it is more suitable to consider diagonal covariances as in the main text, which do not involve a matrix decomposition.

C. Details on Extension to DPMs with Continuous Timesteps

To obtain the estimate of $\sigma_{s|t}^*(\mathbf{x}_t)^2$ appeared in Proposition 5.1, we use a network $\mathbf{h}_t(\mathbf{x}_t) \in \mathbb{R}^d$ to learn $\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[\boldsymbol{\epsilon}_t^2]$ by minimizing the following MSE loss

$$\min_{\mathbf{h}_t} \mathbb{E}_t \mathbb{E}_{q(\mathbf{x}_0, \mathbf{x}_t)} \|\boldsymbol{\epsilon}_t^2 - \mathbf{h}_t(\mathbf{x}_t)\|_2^2,$$

where t is uniformly sampled from $[0, T]$. Then we obtain the estimate of $\sigma_{s|t}^*(\mathbf{x}_t)^2$:

$$\hat{\sigma}_{s|t}(\mathbf{x}_t)^2 = \tilde{\beta}_{s|t} \mathbf{1} + \frac{\beta_{t|s}^2}{\beta_{t|0} \alpha_{t|s}} (\mathbf{h}_t(\mathbf{x}_t) - \hat{\boldsymbol{\epsilon}}_t(\mathbf{x}_t))^2.$$

To obtain the estimate of $\tilde{\sigma}_{s|t}^*(\mathbf{x}_t)^2$ appeared in Proposition 5.2, we use a network $\mathbf{g}_t(\mathbf{x}_t) \in \mathbb{R}^d$ to learn $\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_t)}[(\boldsymbol{\epsilon}_t - \hat{\boldsymbol{\epsilon}}_t(\mathbf{x}_t))^2]$ by minimizing the following MSE loss

$$\min_{\mathbf{g}_t} \mathbb{E}_t \mathbb{E}_{q(\mathbf{x}_0, \mathbf{x}_t)} \|(\boldsymbol{\epsilon}_t - \hat{\boldsymbol{\epsilon}}_t(\mathbf{x}_t))^2 - \mathbf{g}_t(\mathbf{x}_t)\|_2^2,$$

where t is uniformly sampled from $[0, T]$. Then we obtain the estimate of $\tilde{\sigma}_{s|t}^*(\mathbf{x}_t)^2$:

$$\hat{\tilde{\sigma}}_{s|t}(\mathbf{x}_t)^2 = \tilde{\beta}_{s|t} \mathbf{1} + \frac{\beta_{t|s}^2}{\beta_{t|0} \alpha_{t|s}} \mathbf{g}_t(\mathbf{x}_t).$$

D. Inference on Trajectories

To speed up the inference, we can reverse a shorter forward process $q(\mathbf{x}_{\tau_1}, \dots, \mathbf{x}_{\tau_k} | \mathbf{x}_0)$ constrained on a trajectory $1 \leq \tau_1 < \dots < \tau_k = N$ of K timesteps (Song et al., 2020a; Bao et al., 2022). Following the notation of Bao et al. (2022),

the shorter forward process is defined as

$$\begin{aligned} q(\mathbf{x}_{\tau_1}, \dots, \mathbf{x}_{\tau_K} | \mathbf{x}_0) &= q(\mathbf{x}_{\tau_K} | \mathbf{x}_0) \prod_{k=2}^K q(\mathbf{x}_{\tau_{k-1}} | \mathbf{x}_{\tau_k}, \mathbf{x}_0), \\ q(\mathbf{x}_{\tau_{k-1}} | \mathbf{x}_{\tau_k}, \mathbf{x}_0) &= \mathcal{N}(\mathbf{x}_{\tau_{k-1}} | \tilde{\boldsymbol{\mu}}_{\tau_{k-1}|\tau_k}(\mathbf{x}_{\tau_k}, \mathbf{x}_0), \lambda_{\tau_{k-1}|\tau_k}^2 \mathbf{I}), \\ \tilde{\boldsymbol{\mu}}_{\tau_{k-1}|\tau_k}(\mathbf{x}_{\tau_k}, \mathbf{x}_0) &= \sqrt{\bar{\alpha}_{\tau_{k-1}}} \mathbf{x}_0 + \sqrt{\bar{\beta}_{\tau_{k-1}} - \lambda_{\tau_{k-1}|\tau_k}^2} \cdot \frac{\mathbf{x}_{\tau_k} - \sqrt{\bar{\alpha}_{\tau_k}} \mathbf{x}_0}{\sqrt{\bar{\beta}_{\tau_k}}}. \end{aligned}$$

Similarly to Eq. (1), the shorter Markov model to reverse the forward process is defined as

$$\begin{aligned} p(\mathbf{x}_0, \mathbf{x}_{\tau_1}, \dots, \mathbf{x}_{\tau_K}) &= p(\mathbf{x}_{\tau_K}) \prod_{k=1}^K p(\mathbf{x}_{\tau_{k-1}} | \mathbf{x}_{\tau_k}), \\ p(\mathbf{x}_{\tau_{k-1}} | \mathbf{x}_{\tau_k}) &= \mathcal{N}(\mathbf{x}_{\tau_{k-1}} | \boldsymbol{\mu}_{\tau_{k-1}|\tau_k}(\mathbf{x}_{\tau_k}), \boldsymbol{\Sigma}_{\tau_{k-1}|\tau_k}(\mathbf{x}_{\tau_k}) \mathbf{I}). \end{aligned}$$

According to Theorem 3.1, when $\boldsymbol{\Sigma}_{\tau_{k-1}|\tau_k}(\mathbf{x}_{\tau_k}) = \text{diag}(\boldsymbol{\sigma}_{\tau_{k-1}|\tau_k}(\mathbf{x}_{\tau_k})^2)$, the optimal covariance $\boldsymbol{\sigma}_{\tau_{k-1}|\tau_k}^*(\mathbf{x}_{\tau_k})^2$ is

$$\boldsymbol{\sigma}_{\tau_{k-1}|\tau_k}^*(\mathbf{x}_{\tau_k})^2 = \lambda_{\tau_{k-1}|\tau_k}^2 \mathbf{1} + \gamma_{\tau_{k-1}|\tau_k}^2 \frac{\bar{\beta}_{\tau_k}}{\bar{\alpha}_{\tau_k}} \left(\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_{\tau_k})}[\boldsymbol{\epsilon}_{\tau_k}^2] - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_{\tau_k})}[\boldsymbol{\epsilon}_{\tau_k}]^2 \right),$$

where $\boldsymbol{\epsilon}_{\tau_k} = \frac{\mathbf{x}_{\tau_k} - \sqrt{\bar{\alpha}_{\tau_k}} \mathbf{x}_0}{\sqrt{\bar{\beta}_{\tau_k}}}$ is the noise used to generate \mathbf{x}_{τ_k} from \mathbf{x}_0 and $\gamma_{\tau_{k-1}|\tau_k} = \sqrt{\bar{\alpha}_{\tau_{k-1}} - \lambda_{\tau_{k-1}|\tau_k}^2} \sqrt{\frac{\bar{\alpha}_{\tau_k}}{\bar{\beta}_{\tau_k}}}$.

Thus, we can reuse the SN prediction network $\mathbf{h}_n(\mathbf{x}_n)$ trained on the loss in Eq. (9), and estimate $\boldsymbol{\sigma}_{\tau_{k-1}|\tau_k}^*(\mathbf{x}_{\tau_k})^2$ as follows:

$$\hat{\boldsymbol{\sigma}}_{\tau_{k-1}|\tau_k}(\mathbf{x}_{\tau_k})^2 = \lambda_{\tau_{k-1}|\tau_k}^2 \mathbf{1} + \gamma_{\tau_{k-1}|\tau_k}^2 \frac{\bar{\beta}_{\tau_k}}{\bar{\alpha}_{\tau_k}} (\mathbf{h}_{\tau_k}(\mathbf{x}_{\tau_k}) - \hat{\boldsymbol{\epsilon}}_{\tau_k}(\mathbf{x}_{\tau_k}))^2.$$

When considering an imperfect mean (i.e., considering the error of $\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)$), according to Theorem 3.3, the corrected optimal covariance $\tilde{\boldsymbol{\sigma}}_{\tau_{k-1}|\tau_k}^*(\mathbf{x}_{\tau_k})^2$ is

$$\tilde{\boldsymbol{\sigma}}_{\tau_{k-1}|\tau_k}^*(\mathbf{x}_{\tau_k})^2 = \lambda_{\tau_{k-1}|\tau_k}^2 \mathbf{1} + \gamma_{\tau_{k-1}|\tau_k}^2 \frac{\bar{\beta}_{\tau_k}}{\bar{\alpha}_{\tau_k}} \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_{\tau_k})}[(\boldsymbol{\epsilon}_{\tau_k} - \hat{\boldsymbol{\epsilon}}_{\tau_k}(\mathbf{x}_{\tau_k}))^2].$$

Thus, we can reuse the NPR prediction network $\mathbf{g}_n(\mathbf{x}_n)$ trained on the loss in Eq. (15), and estimate $\tilde{\boldsymbol{\sigma}}_{\tau_{k-1}|\tau_k}^*(\mathbf{x}_{\tau_k})^2$ as follows:

$$\hat{\tilde{\boldsymbol{\sigma}}}_{\tau_{k-1}|\tau_k}(\mathbf{x}_{\tau_k})^2 = \lambda_{\tau_{k-1}|\tau_k}^2 \mathbf{1} + \gamma_{\tau_{k-1}|\tau_k}^2 \frac{\bar{\beta}_{\tau_k}}{\bar{\alpha}_{\tau_k}} \mathbf{g}_{\tau_k}(\mathbf{x}_{\tau_k}).$$

We emphasize that the inference on a shorter trajectory doesn't need to train a new NPR (or SN) prediction network on the corresponding shorter process, since the shorter process and the original one share the same marginal distribution. We reuse the one trained on the full timesteps.

E. Additional Discussion

E.1. Error Amplification

Note that in SN-DPM, the error of $\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)^2$ is likely to be amplified compared to that of $\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)$. SN-DPM estimates $\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]^2$ using $\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)^2$, which is the square of the model. However, the training objective for $\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)$ in Eq. (7) is the MSE loss between $\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)$ and $\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]$ instead of their squares. Therefore, the error between $\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)^2$ and $\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]^2$ is likely to be amplified. Specifically, the error of $\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)^2$ at i -th component is $|\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)_i^2 - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]_i^2| = A_i |\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)_i - \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]_i|$, where $A_i = |\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)_i + \mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[\boldsymbol{\epsilon}_n]_i|$. Thus, the error of $\hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n)^2$ at i -th component is amplified by a factor of A_i . In practice, the value of A_i is not guaranteed to be bounded, and the error is likely to be amplified by a large factor.

In contrast, NPR-DPM doesn't have the error amplification problem, since it uses the model $\mathbf{g}_n(\mathbf{x}_n)$ itself to estimate the target $\mathbb{E}_{q(\mathbf{x}_0|\mathbf{x}_n)}[(\boldsymbol{\epsilon}_n - \hat{\boldsymbol{\epsilon}}_n(\mathbf{x}_n))^2]$, and directly minimizes the MSE loss between them.

E.2. NPR-DPM is Not Always Better Than SN-DPM

We note that SN-DPM is better w.r.t. FID and NPR-DPM is better w.r.t. likelihood. The performance of NPR-DPM on likelihood is guaranteed, since it aims to optimize the log-likelihood with the actual imperfect mean. Note the sample quality and likelihood are not necessarily consistent (see Section 3.2 in Theis et al. (2015)), and NPR-DPM does not directly optimize the sample quality. Thus, the performance of NPR-DPM on FID is not guaranteed.

In practice, the NPR-DPM is useful in applications such as lossless compression (see Section 7.3 in Kingma et al. (2021)), where the likelihood performance and the computation efficiency play a central role.

F. Experimental Details

F.1. Details of Pretrained Noise Prediction Networks

In Table 5, we list details of pretrained noise prediction networks used in our experiments. The ImageNet 64x64 pretrained model includes a noise prediction network and additionally a covariance network, and we only use the former one.

Table 5. Details of noise prediction networks used in our experiments. LS means the linear schedule of β_n (Ho et al., 2020) in the forward process of discrete timesteps (see Eq. (2)). CS means the cosine schedule of β_n (Nichol & Dhariwal, 2021) in the forward process of discrete timesteps (see Eq. (2)). The VP SDE is defined as $d\mathbf{x} = -\frac{1}{2}\beta(t)\mathbf{x}dt + \sqrt{\beta(t)}d\mathbf{w}$, where $\beta(t)$ is a linear function of t with $\beta(0) = 0.1$ and $\beta(1) = 20$ (Song et al., 2020b).

	TIMESTEPS TYPE	# TIMESTEPS N	FORWARD PROCESS	PROVIDED BY
CIFAR10 (LS)	DISCRETE	1000	LS	BAO ET AL. (2022)
CIFAR10 (CS)	DISCRETE	1000	CS	BAO ET AL. (2022)
CIFAR10 (VP SDE)	CONTINUOUS	–	VP SDE	SONG ET AL. (2020B)
CELEBA 64X64	DISCRETE	1000	LS	SONG ET AL. (2020A)
IMAGENET 64X64	DISCRETE	4000	CS	NICHOL & DHARIWAL (2021)
LSUN BEDROOM	DISCRETE	1000	LS	HO ET AL. (2020)

F.2. Structure Details of Prediction Networks

In Table 6, we list structure details of NN_1 and NN_2 of prediction networks used in our experiments.

Table 6. NN_1 of pretrained noise prediction networks and NN_2 of SN & NPR prediction networks used in our experiments. Conv denotes the convolution layer. Res denotes the residual block.

	NN_1	NN_2 (SN)	NN_2 (NPR)
CIFAR10 (LS)	CONV	CONV	CONV
CIFAR10 (CS)	CONV	CONV	CONV
CIFAR10 (VP SDE)	CONV	CONV	CONV
CELEBA 64X64	CONV	CONV	CONV
IMAGENET 64X64	CONV	RES+CONV	CONV
LSUN BEDROOM	CONV	RES+CONV	CONV

F.3. Details of Memory and Time Cost

In Table 7, we list the memory and time cost of models (with the corresponding methods) used in our experiments. The extra memory cost of the NPR (or SN) prediction network is negligible. The extra time cost of the NPR (or SN) prediction network is negligible on CIFAR10, CelebA 64x64, at most 4.5% on ImageNet 64x64, and at most 10% on LSUN Bedroom.

Table 7. Model size (MB) and the averaged time (ms) to run a model function evaluation (i.e., the time to compute noise prediction network for baselines, and the time to compute noise & NPR (or SN) prediction networks that share parameters for our method). We also show the extra time cost of the NPR (or SN) prediction network relative to that of the noise prediction network in parentheses. All are evaluated with a batch size of 10 on one GeForce RTX 2080 Ti.

	NOISE PREDICTION NETWORK (ALL BASELINES)	NOISE & SN PREDICTION NETWORKS (SN-DPM)	NOISE & NPR PREDICTION NETWORKS (NPR-DPM)
CIFAR10 (LS)	200.44 MB / 25.65 MS	200.45 MB / 25.77 MS (+0.5%)	200.45 MB / 25.71 MS (+0.2%)
CIFAR10 (CS)	200.44 MB / 25.26 MS	200.45 MB / 25.62 MS (+1.4%)	200.45 MB / 25.64 MS (+1.5%)
CIFAR10 (VP SDE)	235.77 MB / 26.98 MS	235.78 MB / 27.17 MS (+0.7%)	235.78 MB / 27.17 MS (+0.7%)
CELEBA 64X64	300.22 MB / 48.64 MS	300.23 MB / 49.20 MS (+1.2%)	300.23 MB / 49.18 MS (+1.1%)
IMAGENET 64X64	461.82 MB / 72.32 MS	463.47 MB / 75.56 MS (+4.5%)	461.84 MB / 72.80 MS (+0.7%)
LSUN BEDROOM	433.63 MB / 441.41 MS	435.02 MB / 485.54 MS (+10.0%)	433.64 MB / 448.74 MS (+1.7%)

F.4. Training Details

We use a similar training setting to that of the noise prediction network in Bao et al. (2022). On all datasets, we use the AdamW optimizer (Loshchilov & Hutter, 2017) with a learning rate of 0.0001; we train 500K iterations; we use an exponential moving average (EMA) with a rate of 0.9999. We use a batch size of 64 on LSUN Bedroom, and 128 on other datasets. We save a checkpoint every 10K iterations and select the one with the best FID on 1000 generated samples. By default, these samples are generated with full timesteps, except for LSUN Bedroom. On LSUN Bedroom, these samples are generated with 100 timesteps for an acceptable time cost (the time cost on LSUN Bedroom is much larger than other datasets as shown in Table 7).

Training a NPR (or SN) network on CIFAR10 takes about 32 hours on one GeForce RTX 2080 Ti. Training on CelebA 64x64 takes about 72 hours on two GeForce RTX 2080 Ti. Training on ImageNet 64x64 takes about 83 hours on two GeForce RTX 2080 Ti. Training on LSUN Bedroom takes about 171 hours on four GeForce RTX 2080 Ti.

F.5. Log-Likelihood and Sampling

The image data is a vector of integers in $\{0, 1, \dots, 255\}$ and we linearly scale it to $[-1, 1]$ following Ho et al. (2020); Bao et al. (2022).

Log-likelihood. Following Ho et al. (2020); Bao et al. (2022), we discretize the last Markov transition $p(\mathbf{x}_0|\mathbf{x}_1)$ to get discrete NLL and its upper bound. The likelihood results are evaluated on the whole test dataset by default.

Sampling. Following Ho et al. (2020); Bao et al. (2022), we only display the mean of $p(\mathbf{x}_0|\mathbf{x}_1)$ noiselessly at the end of sampling. Following Bao et al. (2022), we clip the covariance $\sigma_2(\mathbf{x}_2)^2$ of $p(\mathbf{x}_1|\mathbf{x}_2)$ by infinity norm, such that $\|\sigma_2(\mathbf{x}_2)\|_\infty \mathbb{E}|\epsilon| \leq \frac{2}{255}y$, where $\|\sigma_2(\mathbf{x}_2)\|_\infty$ denotes the infinity norm, ϵ is the standard Gaussian noise and y is the maximum tolerated perturbation of a channel. Following Bao et al. (2022), we use $y = 2$ on CIFAR10 (LS) and CelebA 64x64 under the DDPM forward process, and use $y = 1$ for other cases. Following Bao et al. (2022), we calculate the FID score on 50K generated samples, using the official implementation of FID to pytorch (<https://github.com/mseitzer/pytorch-fid>). Following Nichol & Dhariwal (2021); Bao et al. (2022), the reference distribution statistics of FID are computed on the full training set on CIFAR10 and ImageNet 64x64, and 50K training samples on CelebA 64x64 and LSUN Bedroom.

F.6. Number of Monte Carlo Samples to Calculate the Optimal Trajectory

Calculating the optimal trajectory needs to estimate every decomposed term appeared in L_{elbo} , which involves a Monte Carlo estimate calculated on a set of training samples (Watson et al., 2021; Bao et al., 2022). Following Bao et al. (2022), we use 50K samples on CIFAR10, 10K samples on CelebA 64x64 and ImageNet 64x64.

F.7. Experimental Details of Table 2, 3 & 4

In Table 2 & 4, the results of baselines and our methods are based on the same noise prediction networks (i.e., those listed in Table 5). The baseline results in DPMs with discrete timesteps are provided by Bao et al. (2022). As for DPMs with

continuous timesteps, we get results of the Euler-Maruyama solver and the probability flow by running the pytorch official code (https://github.com/yang-song/score_sde_pytorch) of Song et al. (2020b), and get results of the ancestral sampling and Analytic-DPM using our implementation.

In Table 3, the results of DDPM, DDIM, Analytic-DPM and our methods are based on the same noise prediction networks (i.e., those listed in Table 5). The time cost in a single timestep is the same for all baselines (DDPM, DDIM, Improved DDPM and Analytic-DPM), based on the statistics reported in Table 4 in Bao et al. (2022). The ratio of the time cost in a single timestep in NPR-DPM and SN-DPM to that in baselines is based on Table 7. The result of baselines are provided by Bao et al. (2022). The results of NPR-DPM and SN-DPM are based on the DDPM forward process on LSUN Bedroom, and based on the DDIM forward process on other datasets. These choices require smaller number of timesteps than their alternatives.

G. Additional Experiments

G.1. Additional Likelihood Comparison

We compare our NPR-DDPM and Improved DDPM (Nichol & Dhariwal, 2021) based on the ImageNet 64x64 model mentioned in Appendix F.1. Improved DDPM parameterizes the diagonal covariance using an interpolation between β_n and $\tilde{\beta}_n$ and learn the covariance by directly optimizing L_{elbo} . As shown in Table 8, our NPR-DDPM can get a reasonable likelihood result with a small number of timesteps, which is contrast to Improved DDPM. With the full timesteps, two methods show a similar likelihood performance. This implies that using an alternative objective such as the MSE loss can achieve a similar likelihood performance to directly optimizing L_{elbo} .

Table 8. Upper bound ($-L_{\text{elbo}}$) on the negative log-likelihood (bits/dim) \downarrow , under the setting of the DDPM forward process on ImageNet 64x64 and the even trajectory.

MODEL \ # Timesteps K	25	50	100	200	400	1000	4000
IMPROVED DDPM	18.91	8.46	5.27	4.24	3.86	3.68	3.57
NPR-DDPM	4.66	4.22	3.96	3.80	3.71	3.64	3.60

G.2. Variance of Likelihood Results

We note that under the full timesteps, the likelihood results of Analytic-DDPM and NPR-DDPM in Table 4 are close. To reduce the effect of randomness, we report the standard deviation of the likelihood results under the full timesteps in Table 9.

Table 9. Upper bound ($-L_{\text{elbo}}$) on the negative log-likelihood (bits/dim) \downarrow . We repeat 5 times with different seeds and report the mean and the standard deviation. Each time we randomly select 10K test samples.

	CIFAR10 (LS)	CIFAR10 (CS)	CELEBA 64X64	IMAGENET 64X64
ANALYTIC-DDPM	3.58593±0.00028	3.42229±0.00019	2.65657±0.00126	3.61749±0.00745
NPR-DDPM	3.57204±0.00039	3.41013±0.00031	2.65257±0.00113	3.59660±0.00793

G.3. Comparison After Normalizing for the Extra Cost

We present results after normalizing for the extra inference cost. As shown in Table 10, our method still outperforms the strong baseline Analytic-DPM. We also compare with the ‘‘Gotta Go Fast’’ SDE solver (Jolicœur-Martineau et al., 2021), and our method performs better when the NFE is small (e.g., ≤ 49), and thereby is more efficient. When the NFE is large (e.g., ≥ 147), different methods have similar performance.

We also present results after normalizing for the extra training cost. On CIFAR10 (LS), we train a noise prediction network with 300K iterations and train a NPR (or SN) prediction network with 200K iterations. This ensures that the total training cost is strictly less than that of the original CIFAR10 (LS) model provided by Bao et al. (2022), which is trained with 500K iterations. With 24 and 49 timesteps respectively, NPR-DDPM gets a FID of 10.56 and 6.24; SN-DDPM gets a FID of 7.51

and 4.64. The results are not affected much by the normalization for the extra training cost, and our method still outperforms the strong baseline Analytic-DPM.

Table 10. FID comparison after normalizing for the extra inference cost. The normalized number of timesteps and model function evaluations (NFE) are displayed in parentheses. The normalization ensures the time cost of our method is *strictly less* than that of the baselines.

# TIMESTEPS	CIFAR10 (LS)		CIFAR10 (CS)		CELEBA 64X64		IMAGENET 64X64		CIFAR10 (VP SDE)	
	25 (24)	50 (49)	25 (24)	50 (49)	25 (24)	50 (49)	25 (23)	50 (47)	25 (24)	50 (49)
ANALYTIC-DDPM	11.60	7.25	8.50	5.50	16.01	11.23	32.56	22.45	11.57	6.54
NPR-DDPM	10.82	6.28	7.98	5.33	16.00	11.06	29.72	21.52	10.88	5.80
SN-DDPM	7.42	4.60	6.15	4.20	12.11	8.08	29.05	20.80	7.94	4.47
NFE ON CIFAR10 (VP SDE)		29 (28)	39 (38)	49 (48)	147 (145)	179 (177)	274 (272)	329 (326)		
GOTTA GO FAST (OFFICIAL CODE)		247.79	116.91	72.29	2.95	2.59	2.74	2.70		
NPR-DDPM		9.22	6.97	5.88	3.44	3.19	2.91	2.94		
SN-DDPM		6.62	5.21	4.55	2.96	2.85	2.67	2.64		

G.4. Samples

Recall that Table 3 reports the least number of timesteps required to achieve a FID around 6 on CIFAR10, CelebA 64x64 and LSUN Bedroom. In Figure 1-3, we show generated samples of NPR-DPM and SN-DPM under these number of timesteps on these datasets. Here we use K to denote the number of timesteps. In Figure 4, we also show generated samples on ImageNet 64x64 under $K = 25$ timesteps and the DDPM forward process.

In Figure 5-9, we show more generated samples from both NPR-DPM and SN-DPM under both the DDPM and DDIM forward processes, as well as the VP SDE.



Figure 1. Generated samples on CIFAR10.

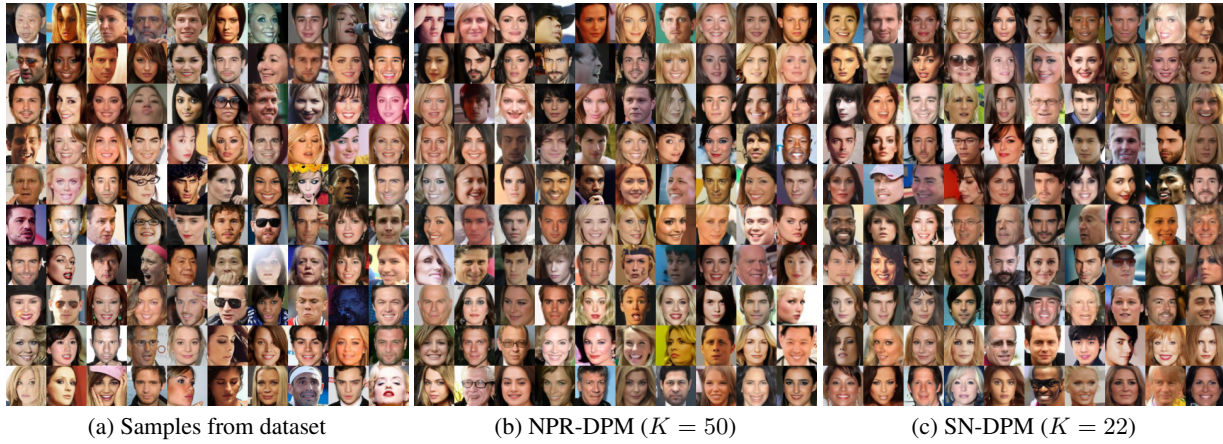


Figure 2. Generated samples on CelebA 64x64.



Figure 3. Generated samples on LSUN Bedroom.

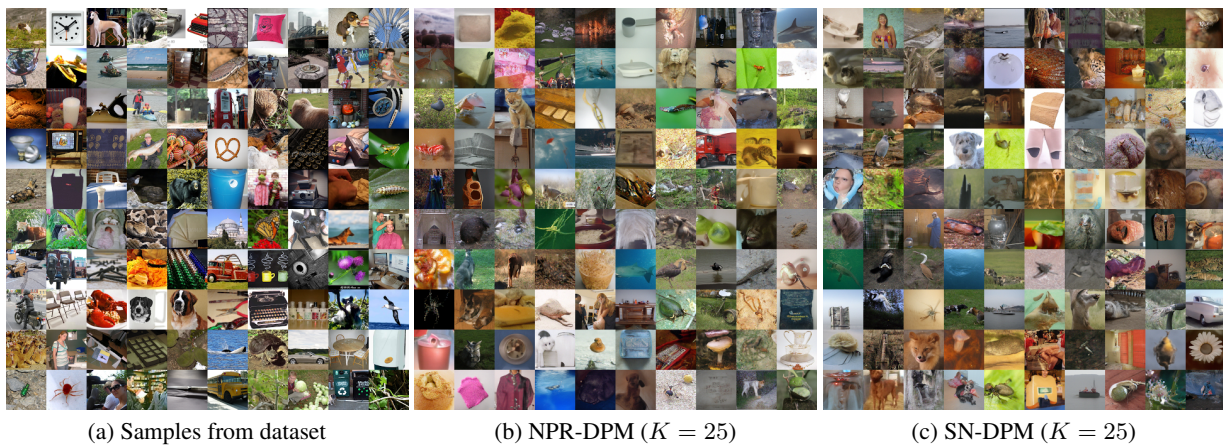


Figure 4. Generated samples on ImageNet 64x64.

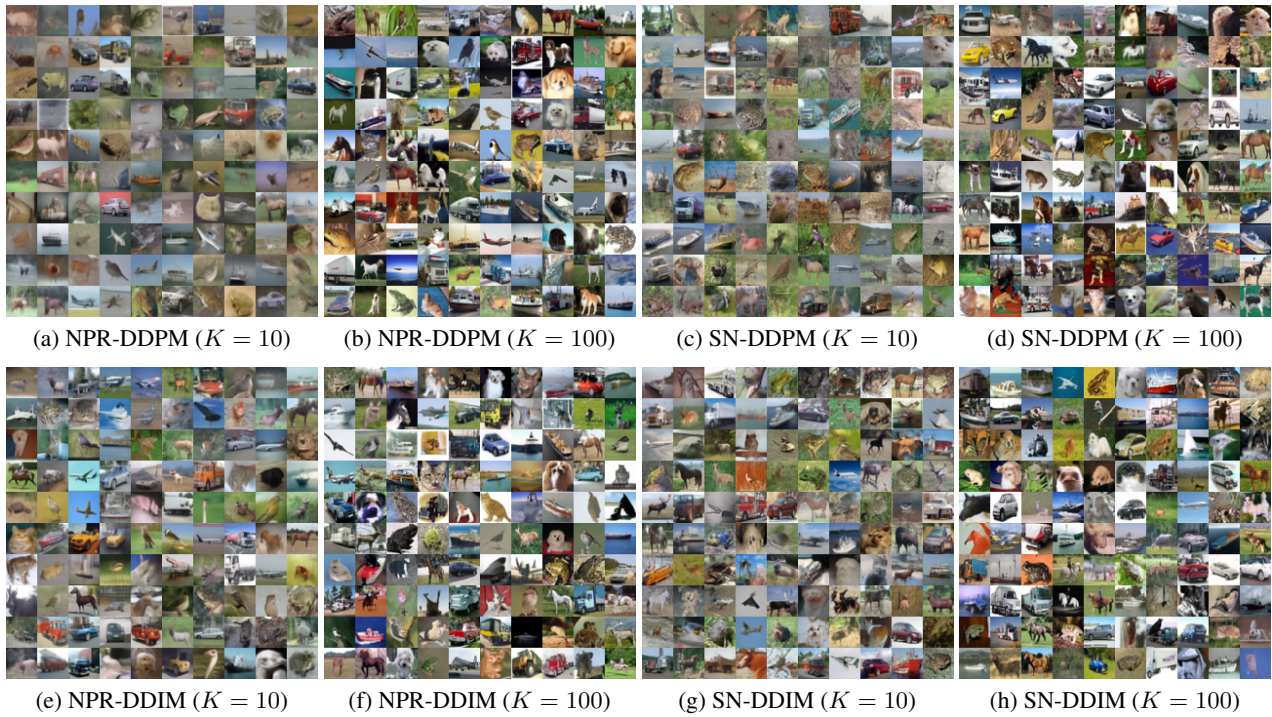


Figure 5. Generated samples on CIFAR10 (LS).



Figure 6. Generated samples on CIFAR10 (CS).

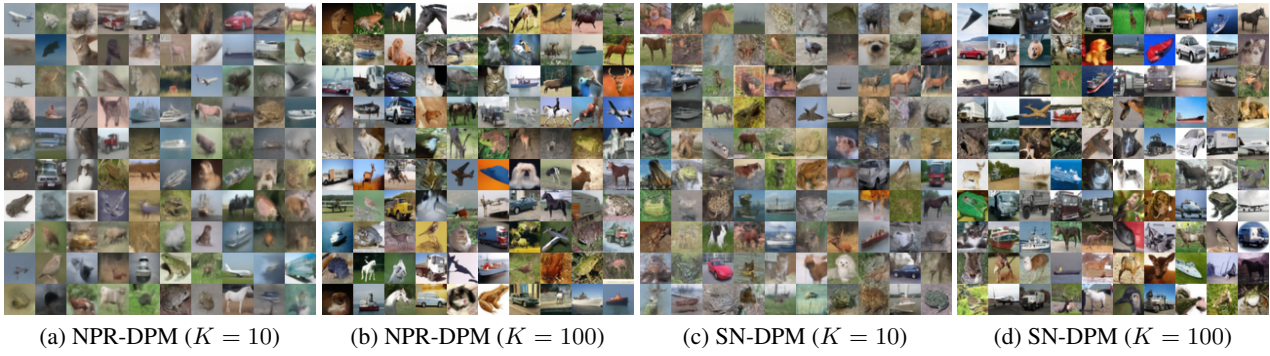


Figure 7. Generated samples on CIFAR10 (VP SDE).



Figure 8. Generated samples on CelebA 64x64.



Figure 9. Generated samples on ImageNet 64x64.