ALTERNATING DIFFUSION FOR PROXIMAL SAMPLING WITH ZEROTH ORDER QUERIES

Anonymous authors
Paper under double-blind review

ABSTRACT

This work introduces a new approximate proximal sampler that operates solely with zeroth-order information of the potential function. Prior theoretical analyses have revealed that proximal sampling corresponds to alternating forward and backward steps of the heat flow. The backward step was originally implemented by rejection sampling, whereas we directly simulate the dynamics. Unlike diffusion-based sampling methods that estimate scores via learned models or by invoking auxiliary samplers, our method treats the intermediate particle distribution as a Gaussian mixture, thereby yielding a Monte Carlo score estimator from directly samplable distributions. Theoretically, when the score estimation error is sufficiently controlled, our method inherits the exponential convergence of proximal sampling under isoperimetric conditions on the target distribution. In practice, the algorithm avoids rejection sampling, permits flexible step sizes, and runs with a deterministic runtime budget. Numerical experiments demonstrate that our approach converges rapidly to the target distribution, driven by interactions among multiple particles and by exploiting parallel computation.

1 INTRODUCTION

Sampling from probability distributions $\pi(x) \propto e^{-f(x)}$ is a fundamental task in statistics and machine learning, with applications in Bayesian posterior inference and score-based generative modeling. Methods like the Unadjusted Langevin Algorithm (ULA) (Roberts & Tweedie, 1996; Durmus & Moulines, 2017) and the Metropolis-Adjusted Langevin Algorithm (MALA) (Roberts & Rosenthal, 1998; Roberts & Stramer, 2002) are widely used. Their convergence under strong convexity assumptions has been established in sharp nonasymptotic terms (Dalalyan, 2017; Dwivedi et al., 2018), and it has further been shown that ULA enjoys exponential convergence in KL divergence under functional inequalities (Cheng & Bartlett, 2018; Vempala & Wibisono, 2019).

Beyond Langevin-type approaches, there has been growing interest in alternative sampling schemes with nonasymptotic guarantees. Among them, *proximal sampling* (Liang & Chen, 2023b) introduces an auxiliary distribution close to the target—typically the Gaussian convolution of π (Lee et al., 2021)—and alternates conditional updates between the target and auxiliary variables. From a theoretical perspective, each update can be interpreted as alternating a forward heat-flow step and a reverse denoising step, which yields exponential convergence under suitable functional-inequality assumptions on the target distributions (Chen et al., 2022; Kook et al., 2024).

Despite this line of analysis, scalable implementations of proximal samplers remain challenging. Existing implementations rely on local optimization of f combined with rejection sampling (Liang & Chen, 2023a;b; Fan et al., 2023), which necessitates small step sizes (i.e., weak convolution) to maintain acceptance and thus incurs many iterations and high overall cost. These bottlenecks have spurred diffusion-based Monte Carlo, which simulates denoising stochastic differential equations (SDEs) (Huang et al., 2024a;b; He et al., 2024). This raises a natural question: *can proximal sampling be implemented in its theoretical form and in a scalable way, directly through Gaussian convolutions and diffusion processes, without relying on rejection sampling?*

Another practical consideration is the efficient sampling of many particles in parallel. In such settings, particle-based algorithms that introduce additional interaction terms or gradient-flow structures can promote faster mixing while maintaining diversity (Liu & Wang, 2016; Futami et al., 2020;

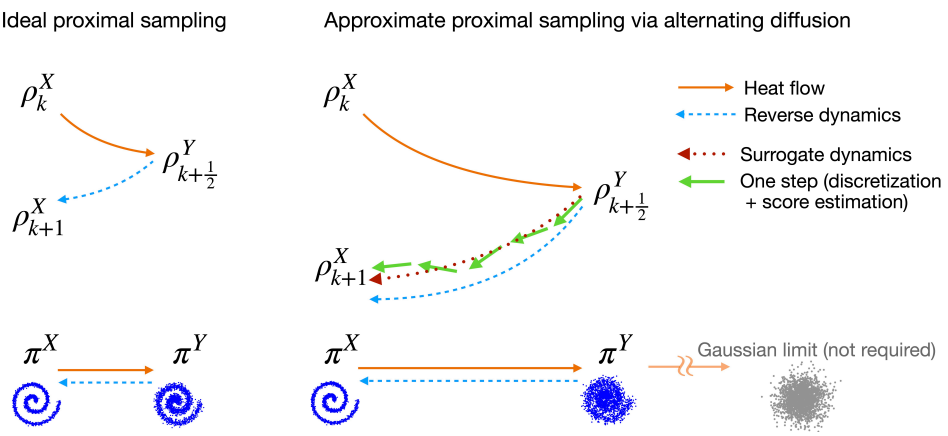

Figure 1: Illustration of the ideal proximal sampling (left) and our approximation (right). Heat flow and reverse dynamics are defined between π^X and π^Y, but applied to intermediate ρ. Although these do not reach their targets in one step, the ideal version attains exponential convergence. Compared to the rejection sampling-based implementation of proximal samplers, our approach allows for larger step sizes (i.e., stronger convolution), which reduce the iterations to reach the target distribution.

Boffi & Vanden-Eijnden, 2023; Lu et al., 2024; Ilin et al., 2025). However, proximal sampling was originally formulated as a single-particle iterative framework, leaving a gap between its theoretical appeal and the practical demands of multi-particle sampling.

1.1 Contributions

In this work, we propose and analyze a diffusion-based approximation of proximal sampling. The key idea is to interpret the auxiliary samples as forming a Gaussian-mixture approximation and to exploit this structure for approximate time-dependent score estimation in the denoising dynamics. Our main contributions are as follows:

- We introduce a new algorithm that serves as an approximate proximal sampler. It is learning-free, gradient-free with respect to f, and rejection-free with a fixed runtime.
- We extend the theory of proximal sampling to show that our diffusion-based approximation inherits comparable convergence rates under suitable assumptions. This implies that our algorithms which alternate perturbation and denoising converge to the target distribution without requiring initialization from the Gaussian limit, unlike standard diffusion models.
- We provide empirical evidence that our method converges rapidly to representative targets compared to existing implementations of proximal sampler, and that its multi-particle extension improves both wall-clock efficiency and sample diversity.

2 Preliminaries

In this section, we provide a brief overview of proximal sampling, which underlies our proposed method, along with its convergence properties.

2.1 Definitions

Let μ and ν be two probability measures on $\mathbb{R}^d$ with $\mu \ll \nu$. We define the Kullback–Leibler (KL) divergence, the Rényi divergence of order $q \geq 1$, and the relative Fisher information as

$$H_\nu(\mu) = \int \log \frac{d\mu}{d\nu}\, \mathrm{d}\mu, \quad \mathcal{R}_{q,\nu}(\mu) = \frac{1}{q-1} \log \int \left(\frac{\mathrm{d}\mu}{\mathrm{d}\nu}\right)^q \mathrm{d}\nu, \quad J_\nu(\mu) = \int \left\| \nabla \log \frac{\mathrm{d}\mu}{\mathrm{d}\nu} \right\|^2 \mathrm{d}\mu.$$

Note that setting $q = 1$ yields $\mathcal{R}_{1,\nu}(\mu) = H_\nu(\mu)$. We say that ν satisfies a log-Sobolev inequality (LSI) with constant $C_{\mathrm{LSI}} > 0$ if $H_\nu(\mu) \leq \frac{1}{2} C_{\mathrm{LSI}} J_\nu(\mu)$.

2.2 PROXIMAL SAMPLING

Let $f : \mathbb{R}^d \to \mathbb{R}$ be a potential function, and consider the target distribution $\pi^X(x) \propto \exp(-f(x))$. For a given step size $h > 0$, following Lee et al. (2021), we define a joint distribution on $\mathbb{R}^d \times \mathbb{R}^d$ as

$$\pi^{X,Y}(x, y) \propto \exp\left(-f(x) - \frac{\|x - y\|^2}{2h}\right). \tag{1}$$

Marginalizing over y yields $\pi^X(x) \propto \int \pi^{X,Y}(x, y)\,\mathrm{d}y$. On the other hand, marginalizing over x defines a new distribution π^Y as $\pi^Y(y) \propto \int \exp\left(-f(x) - \frac{1}{2h}\|x - y\|^2\right)\mathrm{d}x$. This corresponds to the Gaussian convolution $\pi^Y = \pi^X * \mathcal{N}(0, hI_d)$.

The proximal sampler with step size h iterates updating a current particle $x_k \in \mathbb{R}^d$ as follows:

$$y_{k+\frac{1}{2}} \sim \mathcal{N}(x_k, hI_d) \tag{2}$$

$$x_{k+1} \sim \pi^{X|Y=y_{k+\frac{1}{2}}} \tag{3}$$

where $\pi^{X|Y=y_{k+1/2}}(x) \propto \exp(-f(x) - \frac{\|x-y_{k+1/2}\|^2}{2h})$. The restricted Gaussian oracle (RGO) (Lee et al., 2021) assumes access to an oracle which enables sampling (3) from $\pi^{X|Y=y_{k+1/2}}(x)$.

Several works have implemented the RGO (Liang & Chen, 2023a;b; Fan et al., 2023) through rejection sampling, achieving an expected complexity of $\tilde{O}(1)$ under regularity conditions on f. However, this requires choosing h sufficiently small depending on f and d, which in practice leads to many iterations, and the computational cost further fluctuates due to randomness.

2.3 CONNECTIONS TO DIFFUSION PROCESSES

Chen et al. (2022) established an improved convergence analysis of proximal sampling by interpreting each update through dynamics interpolated by an SDE, as illustrated in the left panel of Figure 1. The forward step (2) amounts to Gaussian convolution, which corresponds to evolving the heat equation or its associated SDE, for $t \in [0, h]$,

$$\mathrm{d}Z_t = \mathrm{d}B_t, \quad \partial_t \mu_t = \tfrac{1}{2}\Delta\mu_t, \tag{4}$$

where B_t is a standard Brownian motion. Writing P_t for the Gaussian convolution kernel $\mathcal{N}(0, tI_d)$, we have $\mu_t = \mu_0 P_t$, and in particular $\pi^Y = \pi^X P_h$. The backward step (3) can be viewed as the reverse operation of the forward step conditioned on Z_h. It is governed by the following SDE starting from $Z_0^{\leftarrow} = Z_h$ and the corresponding Fokker–Planck equation; for $t \in [0, h]$,

$$\mathrm{d}Z_t^{\leftarrow} = \nabla \log(\pi^X P_{h-t})(Z_t^{\leftarrow})\,\mathrm{d}t + \mathrm{d}B_t^{\leftarrow}, \quad \partial_t \mu_t^{\leftarrow} = -\mathrm{div}\big(\mu_t^{\leftarrow} \nabla \log(\pi^X P_{h-t})\big) + \tfrac{1}{2}\Delta\mu_t^{\leftarrow}, \tag{5}$$

where $B_t^{\leftarrow}$ denotes the backward Brownian motion. Letting Q_t denote its transition kernel, we obtain $\mu_h^{\leftarrow} = \mu_0^{\leftarrow} Q_h = \int \pi^{X|Y=y}(x)\,\mu_0^{\leftarrow}(y)\,\mathrm{d}y$. In particular, we have $\pi^X = \pi^Y Q_h$.

Hence the two steps in proximal sampling can be viewed as SDEs between π^X and π^Y, with P_t and Q_t as their transition kernels. The proximal sampler applies the dynamics (4) and (5) with the start distributions $\rho_k^X := \mathrm{law}(x_k)$ and $\rho_{k+1/2}^Y := \mathrm{law}(y_{k+1/2})$ at k-th iteration, respectively, which then evolve through Gaussian convolution or conditional integration. Although a single step does not reach π^Y or π^X, iterating the forward–backward scheme leads to contraction towards the target distribution, as made precise in the following theorem.

Theorem 1 (Chen et al. (2022), Theorem 3)**.** *Assume that π^X satisfies LSI with constant C_{LSI}. For any $h > 0$ and any initial distribution ρ_0^X, the k-th iterate ρ_k^X of the proximal sampler with step size h satisfies, for $q \geq 1$,*

$$\mathcal{R}_{q,\pi^X}(\rho_k^X) \leq \frac{\mathcal{R}_{q,\pi^X}(\rho_0^X)}{(1 + h/C_{LSI})^{2k/q}}. \tag{6}$$

Theorem 1 provides the exponential convergence guarantee of proximal sampling. Its proof follows by combining the forward and backward contraction properties at each step stated in the next lemma.

Lemma 1 (Chen et al. (2022), Appendix A.4)**.** *Assume that π^X satisfies LSI with constant C_{LSI} and $\rho \ll \pi^X$, $\rho' \ll \pi^Y$. For $q \geq 1$, we have*

$$\mathcal{R}_{q,\pi^Y}(\rho P_h) = \mathcal{R}_{q,\pi^X P_h}(\rho P_h) \leq \frac{\mathcal{R}_{q,\pi^X}(\rho)}{(1+h/C_{LSI})^{1/q}}, \tag{7}$$

$$\mathcal{R}_{q,\pi^X}(\rho' Q_h) = \mathcal{R}_{q,\pi^Y Q_h}(\rho' Q_h) \leq \frac{\mathcal{R}_{q,\pi^Y}(\rho')}{(1+h/C_{LSI})^{1/q}}. \tag{8}$$

In this paper, we aim to realize the proximal sampler by directly simulating the associated SDEs. In the forward direction, convolution with π^X can be simulated exactly by injecting Gaussian noise, so the bound (7) is directly applicable. In contrast, the backward dynamics (5) cannot be simulated exactly, and our surrogate version of (5) will provide an approximation in place of (8).

3 Approximate Multi-particle Proximal Sampling

Building on the SDE interpretation in Section 2, we propose to approximate the backward dynamics by replacing π^X with a surrogate distribution constructed from the current particles. This enables Monte Carlo estimation of the score function using only evaluations of f, without requiring gradient or rejection sampling. At a high level, each iteration of our diffusion-based proximal sampler consists of (i) evolving particles by the forward heat flow (4) and (ii) applying the surrogate dynamics, which time-discretizes and approximates the reverse flow (5).

3.1 Score Estimation for Surrogate Dynamics

We derive the update rule in the backward step, given the particles at k-th iteration: $X_k = \{x_i\}_{i=1}^N$ and $Y_{k+1/2} = \{y_j\}_{j=1}^N$. We approximate the backward dynamics (5) by replacing π^X with a surrogate distribution $\hat{q}_{k+1}(\cdot \mid Y_{k+1/2}, X_k)$ constructed from these particles, as follows:

$$\hat{q}_{k+1}(x \mid Y_{k+1/2}, X_k) \propto \frac{1}{N}\sum_{j=1}^N \pi^{X|Y=y_j}(x) \frac{\pi^Y(y_j)}{\hat{q}_{k+1/2}(y_j \mid X_k)}, \tag{9}$$

where $\hat{q}_{k+1/2}(y \mid X_k) = \frac{1}{N}\sum_i \mathcal{N}(y; x_i, hI_d)$. The conditional distribution with the reweighting term ensures that, as $N \to \infty$, sampling $y \sim \hat{q}_{k+1/2}$ recovers the target distribution π^X. By substituting the explicit form of $\pi^{X|Y=y_j}$ and π^Y we obtain

$$\begin{aligned}\hat{q}_{k+1}(x \mid Y_{k+1/2}, X_k) &\propto \frac{1}{N}\sum_{j=1}^N \frac{\exp\left(-f(x) - \frac{1}{2h}\|x-y_j\|^2\right)}{\hat{q}_{k+1/2}(y_j \mid X_k)} \\ &\propto \frac{1}{N}\sum_{j=1}^N \frac{\mathcal{N}(x; y_j, hI_d)}{\hat{q}_{k+1/2}(y_j \mid X_k)} \exp(-f(x)) =: g_N^{k+1/2}(x)\exp(-f(x)), \end{aligned} \tag{10}$$

where $g_N^{k+1/2}$ is the unnormalized density of an N-component weighted Gaussian mixture (see Appendix A.1 for the full derivation). $\hat{q}_{k+1}$ involves an inverse reweighting with respect to a Gaussian mixture $\hat{q}_{k+1/2}$. This reduces weights when y_j is located in areas where X_k is overly concentrated, while amplifying their importance in sparse regions. The resulting term, derived from the empirical particle system, may help prevent particle collapse and promote exploration. We later verify this effect in our experiments.

The surrogate score is defined as $\hat{s}_t(x_t) \coloneqq \nabla \log(\hat{q}_{k+1}P_{h-t})(x_t)$, which serves as the drift term in the surrogate reverse dynamics, replacing $\nabla \log(\pi^X P_{h-t})$ in the backward dynamics (5). This expression can be rewritten in expectation form using Bayes' rule:

$$\hat{s}_t(x_t) = \mathbb{E}_{g_N^{k+1/2}(x_0|x_t)}\left[\frac{x_0 - x_t}{\sigma_t^2} e^{-f(x_0)}/C_t^k\right], \tag{11}$$

$$\text{where} \quad \sigma_t^2 = h - t, \quad C_t = \int g_N^{k+\frac{1}{2}}(x_0|x_t) e^{-f(x_0)}\,\mathrm{d}x_0. \tag{12}$$

Algorithm 1 Zeroth-Order Diffusive Proximal Sampler

Input: potential function $f:\mathbb{R}^d \to \mathbb{R}$, initial samples $\{x_0^{(i)}\}_{i=1}^N$ step size h, iterations K, diffusion steps T, number of interim samples M, noise schedule $\{\sigma_t^2\}_{t=0}^T$ with $\sigma_T^2 = h$ and $\sigma_0^2 = \sigma_{\min}^2$.

▷ All operations for $i = 1, \dots, N$, $j = 1, \dots, N$, and $l = 1, \dots, M$ are evaluated in parallel.
for $k = 0, 1, \dots, K-1$ **do**
  `# Step 1: Forward heat flow (4)`
  $y_{k+\frac{1}{2}}^{(j)} \leftarrow x_k^{(j)} + \sqrt{h}\,\xi_k^{(j)}$ with $\xi_k^{(j)} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_d)$.
  Initialize $z_T^{(i)} \leftarrow x_k^{(i)} + \sqrt{h}\,\xi_k'^{(i)}$ with $\xi_k'^{(i)} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_d)$.
  `# Step 2: Surrogate version of reverse dynamics (5)`
  for $t = T, T-1, \dots, 1$ **do**
    Set $\Delta t \leftarrow \sigma_t^2 - \sigma_{t-1}^2$.
    Compute $\bar{\sigma}, m_{i,j} \coloneqq m_j(z_t^{(i)}), w_{i,j} \coloneqq w_j(z_t^{(i)})$ by (14) and set $w_{i,j} \leftarrow w_{i,j} / \sum_j w_{i,j}$
    Draw M samples $z_0^{(i,l)} \overset{\text{i.i.d.}}{\sim} \sum_j w_{i,j}\mathcal{N}(m_{i,j}, \bar{\sigma}^2)$.
    Compute $c_{i,l} \coloneqq \exp\big(-f(z_0^{(i,l)})\big)$ and set $c_{i,l} \leftarrow c_{i,l} / \sum_l c_{i,l}$.
    `# Update each particle using the Euler-Maruyama step`
    $z_{t-1}^{(i)} \leftarrow z_t^{(i)} + \Delta t \sum_l c_{i,l}\big(z_0^{(i,l)} - z_t^{(i)}\big)/\sigma_t^2 + \sqrt{\Delta t}\,\xi_k''^{(i)}$ with $\xi_k''^{(i)} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_d)$.
  end for
  Set $x_{k+1}^{(i)} \leftarrow z_0^{(i)}$.
end for
return $\{x_K^{(i)}\}_{i=1}^N$

See Section A.1 for the derivation. Since $g_N^{k+1/2}(x_0)$ is a Gaussian mixture and the conditional law $x_t \mid x_0$ under the perturbation kernel is $\mathcal{N}(x_0, \sigma_t^2 I_d)$, Bayes' rule implies that the posterior distribution $g_N^{k+1/2}(x_0 \mid x_t)$ is itself a Gaussian mixture:

$$g_N^{k+\frac{1}{2}}(x_0|x_t) \propto \sum_{j=1}^N w_j(x_t)\,\mathcal{N}(x_0; m_j(x_t), \bar{\sigma}^2 I), \tag{13}$$

$$\text{where } \bar{\sigma}^2 = (h^{-1} + \sigma_t^{-2})^{-1},\ m_j(x_t) = \bar{\sigma}^2(h^{-1}y_j + \sigma_t^{-2}x_t),\ w_j(x_t) = \frac{\mathcal{N}(x_t; y_j, (h+\sigma_t^2)I_d)}{\hat{q}^{k+\frac{1}{2}}(y_j \mid X_k)}. \tag{14}$$

The detailed derivation can be found in Section A.2. This formulation enables practical score estimation via Monte Carlo sampling from the Gaussian mixture without requiring model training.

3.2 Algorithm and Complexity

The surrogate score (11) involves an expectation with respect to the Gaussian-mixture posterior (13). In practice, we approximate this expectation using Monte Carlo estimation by drawing M samples for each particle. Combining the forward step (Gaussian perturbation) and the discretized surrogate reverse dynamics with T denoising steps, we obtain the iterative sampling procedure that approximates the proximal sampling. It is worth emphasizing that the proposed method relies only on a zeroth-order oracle of f and does not require gradient information, unlike gradient-based sampling methods such as Langevin Monte Carlo. A complete description of the algorithm is provided in Algorithm 1.

Regarding computational cost, a straightforward implementation would require $KTMN$ evaluations of the potential f for K outer iterations. This complexity can be further reduced to KT under the assumption of a parallel oracle that can evaluate f simultaneously on multiple samples, Such a design naturally exploits parallel computation, making the method efficient in modern computing environments.

4 THEORETICAL ANALYSIS

In this section, we provide a theoretical analysis that essentially quantifies the deviation from the ideal proximal sampling caused by time and space discretization. Our method splits the macro step size h into T substeps of length $\eta = h/T$ for the backward dynamics. For each $\ell \in \{1, \dots, T\}$, let $t_\ell := (\ell - 1)\eta$ be the substep start time, and let μ_{t_ℓ} denote the current law at time t_ℓ. The following lemma establishes the contraction of the KL divergence in a single outer iteration of Algorithm 1, up to an error due to time discretization. We assume that π^X satisfies an LSI with constant C_{LSI}. Focusing on iteration k, we write $\rho_k^X := \mathrm{law}(x_k^{(i)})$ and $\rho_{k+1/2}^Y := \mathrm{law}(y_{k+1/2}^{(j)})$. A complete description of the assumptions and proofs is deferred to Appendix A.

Lemma 2 (Discretization-only one-step bound)**.** *Assume $N, M \to \infty$, so the backward update uses the exact score. Suppose further that within each substep ℓ, the temporal variation of the score between t_ℓ and $t_\ell + t$ ($t \in [0, \eta]$) integrating over $\mu_{t_\ell + t}$ is bounded by $L_{\ell,t}^2 t^2$, the score of the reference measure ν_{t_ℓ} is Lipschitz in space with constant $L_{\nu,\ell}$, and there exists a uniform entropy bound $\bar{H}^{(k)} > 0$ along the substep path. Then for $u \geq 1$,*

$$H_{\pi^X}(\rho_{k+1}^X) \leq \frac{H_{\pi^Y}(\rho_{k+1/2}^Y)}{(1 + h/C_{LSI})^{1-1/(2u^2)}} + 2u^4\,\Lambda_1^{(k)},$$

$$\textit{where} \quad \Lambda_1^{(k)} := C_{LSI}\Big((1 + h/C_{LSI})^{1/(2u^2)} - 1\Big)\,\{4\eta^2 L_{\nu,*}^2 (C_{LSI} + h)\bar{H}^{(k)} + \eta d C\},$$

$$C := 2\eta L_{\nu,*}^3 + L_{\nu,*}^2 + \eta L_s^2, \ \textit{with } L_{\nu,*} := \sup_\ell L_{\nu,\ell} \textit{ and } L_s := \sup_{\ell,t} L_{\ell,t}.$$

The additional discretization error term in Lemma 2 scales as $O(h/T)$ with respect to the number of substeps T. Therefore, if T is chosen appropriately, the backward step essentially inherits the convergence rate of the ideal proximal sampler in (8).

When $N', M < \infty$, we assume that at each substep start t_ℓ the score estimator admits the bound $\Lambda_2^{(k)} := 2\sup_{\ell=1,\dots,T} \mathbb{E}_{\mu_{t_\ell}}[\|\hat{s}_{N',M}^{(k)}(\cdot, t_\ell) - s_{t_\ell}^{(k)}(\cdot)\|^2]$, where s_{t_ℓ} is the true score at t_ℓ and $\hat{s}_{N',M}^{(k)}(\cdot, t_\ell)$ is the Monte Carlo estimator. The following result incorporates this condition.

Proposition 1 (Main one-step bound with split errors)**.** *Let $u \geq 1$ and there exists C_{LSI} and $\Lambda_2^{(k)}$. Then the k-th iterate of Algorithm 1 satisfies*

$$H_{\pi^X}(\rho_{k+1}^X) \leq \frac{H_{\pi^Y}(\rho_{k+1/2}^Y)}{(1 + h/C_{LSI})^{1-1/(2u^2)}} + 2u^4 C_{LSI}\Big((1 + h/C_{LSI})^{1/(2u^2)} - 1\Big)\,(\Lambda_1^{(k)} + 2\Lambda_2^{(k)}).$$

Proposition 1 directly implies that, as shown in Corollary 3, with $r := (1 + h/C_{LSI})^{\,2-1/(2u^2)}$, the iteration complexity for achieving $H_{\pi^X}(\rho_K^X) \leq \varepsilon$ is $K = O\big(\log(H_{\pi^X}(\rho_0^X)/\varepsilon)/\log r\big)$, provided that, each error term satisfies $\Lambda_1^{(k)} + \Lambda_2^{(k)} \asymp O(\varepsilon/C_{LSI})$.

$\Lambda_2^{(k)}$ reflects Monte Carlo fluctuations due to the finite numbers of samples N and M. Under bounded conditional variances of π^X we obtain $O(1/N)$ error, and under bounded expectations of e^{4f} and finite fourth moments under the Gaussian mixture proposals we obtain $O(1/M)$ error, yielding $\Lambda_2^{(k)} = O(1/N + 1/M)$.

From Theorem 1, the ideal proximal sampler requires about $(2\log(1 + h/C_{LSI}))^{-1}$ outer iterations times a logarithmic factor in the initial divergence. When $C_{LSI} \gg h$, this factor is approximated by $\tilde{O}(C_{LSI}/h)$, suggesting that larger h is favorable. While rejection-sampling implementations of RGO suffer from an upper bound on h, our method can take h large as long as $T = O(h)$ to control the discretization error. This does not change the overall computational cost (outer iterations $\times$ T steps), but allows h to reflect global structure of π^X such as inter-mode distances rather than only local smoothness of f, which we confirm to be practically advantageous in experiments.

5 EXPERIMENTS

We design two experiments to showcase the strengths of our algorithm as an approximation to proximal sampling. First, we revisit the Gaussian Lasso mixture (Liang & Chen, 2023b) to test

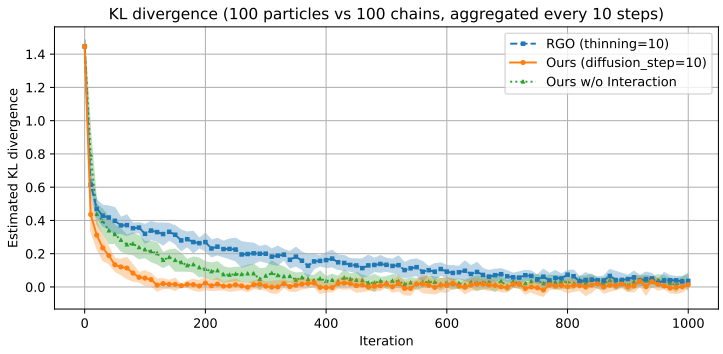

Figure 2: Convergence of estimated KL divergence, averaged over 10 random seeds with shaded areas indicating variances. Our method (orange) outperforms both the proximal sampler with RGO (blue) and an ablated variant of our algorithm without particle interactions (green). It achieves the same accuracy as RGO in about $10\times$ fewer iterations ($100\times$ faster when accounting for thinning).

whether our method accelerates convergence beyond RGO-based proximal sampling, especially via parallel particle updates. Second, we study uniform distributions over bounded, nonconvex, and disjoint domains, comparing our method with In-and-Out (Kook et al., 2024), the proximal sampler for uniform distributions originally analyzed for convex bodies. These experiments respectively demonstrate the benefits of larger step sizes and the applicability of our approach when gradients are unavailable.

5.1 Gaussian Lasso Mixture

Setup. Following the experimental setting in Liang & Chen (2023b), we set the target distribution as a Gaussian Lasso mixture:

$$\pi^X(x) = \frac{\sqrt{\det Q}}{2\sqrt{(2\pi)^d}} \exp\left(-\frac{1}{2}(x-1)^\top Q(x-1)\right) + 2^{d-1}\exp\left(-\|4x\|_1\right)$$

where $Q = USU^\top$, $d = 5$, $S = \mathrm{diag}(14, 15, 16, 17, 18)$, and U is an arbitrary orthogonal matrix.

We compare against the RGO baseline with 100 chains and step size $h{=}1/135$, exactly matching the experimental setting of Liang & Chen (2023b), where it outperformed ULA and MALA. Our method uses step size $h{=}1/10$ under two settings: $N{=}100$ interacting particles, or $N{=}1$ run with 100 chains. For comparability, we thin the RGO baseline by grouping every 10 single-step updates into one iteration, so that the parallel f-evaluation cost is comparable to ours.

Convergence is measured by KL divergence to the target distribution, estimated using Büth et al. (2025) with a fixed reference of 1000 particles from a long RGO run. At each evaluation, 1000 particles are aggregated from 10 successive iterations. The histogram in the right panel of Figure 3 shows this reference distribution. All experiments are repeated with 10 random seeds, reporting the mean and variance of the KL estimate. Detailed settings are given in Table 1 and Section B.

Results. Figure 2 shows the convergence curves. Our method converges substantially faster than the RGO baseline. With step size 13.5 times larger, our method reaches comparable KL divergence in $\sim$100 iterations, while the proximal sampler needs $\sim$950 iterations ($\approx$ 9500 RGO updates). Our method also surpasses the ablated variant with $N{=}1$ independent parallel chains, indicating that particle interactions are essential to accelerate mixing. The marginals in Figure 3 indicate that our method is already approaching the target by $\sim$100 iterations (vs. $\sim$950 for RGO), and it closely matches the long-run reference by $\sim$250 iterations. A more detailed observation is deferred to Section B (Figure 5), which shows that increasing the step size h leads to faster convergence.

Discussion. These results demonstrate that relaxing the stringent step size restriction of proximal sampling yields a practical gain of nearly an order of magnitude in convergence speed. The benefit is amplified when leveraging multiple interacting particles, which facilitate more efficient exploration

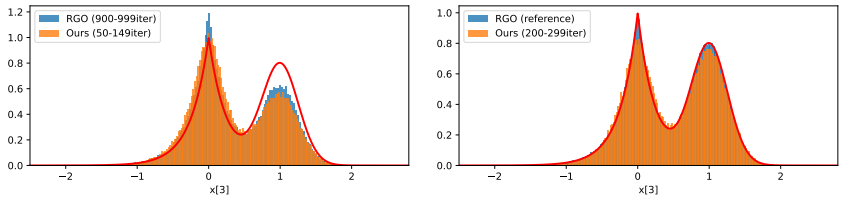

Figure 3: One-dimensional marginals of π^X along the third coordinate (seed $= 0$). The red curve is the ground-truth marginal. Our method around 100 iterations (left, orange) already matches the proximal sampler with RGO in $\sim$1000 iterations (left, blue), and after 200–300 iterations it closely aligns with the reference obtained from a long run of the proximal sampler with RGO.

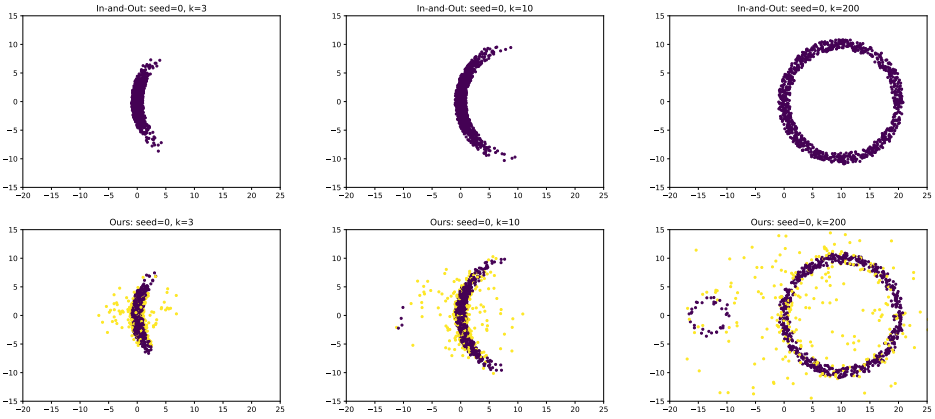

Figure 4: Empirical distributions at $k = 3, 10, 200$ iterations with seed $= 0$ on the two-tori domain. In-and-Out (top row) finds T_1, which overlaps with the initial standard Gaussian distribution, but fails to reach T_2. Our method (bottom row) generates some particles outside the domain but gradually drives particles toward T_2, demonstrating its ability to explore both components.

of the mixture components. This effect can be explained by the surrogate target distribution in (9), which re-weights via $(\hat{q}_{k+1/2}(y \mid X_k))^{-1}$ to down-weight overpopulated regions and encourage exploration of sparser ones, thereby accelerating convergence beyond the step-size effect alone.

5.2 Uniform Distributions on Bounded Domains

Setup. We evaluate our method on uniform sampling over a bounded, nonconvex, and disjoint domain $K \subset \mathbb{R}^3$, and compare it with the In-and-Out (Kook et al., 2024). In-and-Out proposes $y_{i+1} \sim \mathcal{N}(x_i, hI_d)$ and repeatedly resamples $x_{i+1} \sim \mathcal{N}(y_{i+1}, hI_d)$ until $x_{i+1} \in K$ or a retry threshold R is reached. It converges under warm starts, but practical efficiency requires convex K.

We take K as the union of two disjoint solid tori in $\mathbb{R}^3$. T_1 is a torus centered at $(10, 0, 0)$ with major radius 10 and minor radius 1, and T_2 is a torus centered at $(-13, 0, 0)$ with major radius 3 and minor radius 1. We denote $K = T_1 \cup T_2$. Particles start from $\mathcal{N}(0, I_3)$. Our sampler uses the potential $f(x) = 0$ if $x \in K$ and 100 otherwise, inducing an approximate uniform law. We run 1,000 particles and monitor their occupancies, visualized from the third coordinate direction. Details of the experimental setting are given in Table 2.

Results. Figure 4 shows that In-and-Out converges to uniformity inside the near torus T_1, consistent with its guarantee under warm starts, but fails to reach the distant torus T_2. Our method first fills the near torus, then transitions to the remote one, where many particles eventually accumulate.

Discussion. This experiment confirms that our algorithm works without gradients and can explore disconnected modes where exact proximal steps stagnate. The noisy score approximation facilitates such transitions, akin to effects observed in diffusion-based black-box optimization (Lyu et al.,

2025). Other constrained-domain samplers include Projected Langevin Monte Carlo (Bubeck et al., 2018), MYULA (Brosse et al., 2017), and penalized Langevin dynamics (Gurbuzbalaban et al., 2024) based on distance to the constraint set, all of which rely on projection-type operations and are therefore limited to simple convex bodies where such mappings are computable. By contrast, our sampler only requires a membership oracle or a simple outside penalty, making the zeroth-order oracle framework applicable to a broader range of domains.

6 Connections to Other Sampling Methods

Diffusion-based Monte Carlo methods transport samples from a Gaussian initialization to the target distribution. Some rely on training models for score estimation (Vargas et al., 2023; Richter & Berner, 2024), while others drive auxiliary samplers (Huang et al., 2024a; He et al., 2024). Several works also develop acceleration and correction techniques within this paradigm (Lu et al., 2022; Kim & Ye, 2023; Li et al., 2024). Beyond diffusion-based pushforwards, related approaches construct explicit density paths (Fan et al., 2024; Guo et al., 2025).

Our approach also specifies and employs SDE dynamics between two fixed distributions, but unlike the one-way pushforward paradigm, it repeatedly applies these dynamics. This removes the restriction of restarting from a Gaussian base by using a variance-expanding (VE) diffusion. In the context of variance-preserving (VP) diffusions, sampling error partly arises from the discrepancy between the Gaussian equilibrium and the distribution obtained after finite-time mixing (Li et al., 2023; Pierret & Galerne, 2025). Our analysis suggests that, even though the VE process does not reach the Gaussian equilibrium in any bounded horizon, convergence can still be ensured by alternating finite-time noise addition and denoising dynamics through repeated cycles.

A particularly relevant comparison is with ZOD-MC (He et al., 2024), a diffusion-based method that performs denoising from a Gaussian initialization. In ZOD-MC, the outer loop transports samples while an inner proximal sampler is used for score estimation. Our method inverts this structure: instead of nesting a sampler inside a pushforward loop, we simulate proximal-style SDE dynamics directly. Moreover, unlike ZOD-MC, which requires access to the minimizer of the potential—or, in practice, gradient information to locate local solutions for proximal updates—our algorithm operates solely with zeroth-order oracle information. Another related line is SLIPS (Grenioux et al., 2024), which alternates between denoising noisy observations and updating auxiliary variables. This alternating scheme resembles ours, but the overall iteration follows a one-way dynamics that converges to the target distribution as the time horizon grows.

Finally, our method is also close in spirit to Diffusive Gibbs Sampling (DiGS) (Chen et al., 2024), which performs Gibbs updates by alternating perturbation and denoising, updating both the state and an auxiliary variable at each step. However, the designs differ: DiGS employs VP diffusion with auxiliary samplers such as MALA and provides no convergence guarantees, whereas our method uses VE diffusion consistent with proximal sampling, yields provable contraction guarantees, and exploits parallel multi-particle computation for lightweight score estimation.

7 Conclusion

We have introduced a diffusion-based approximation of proximal sampling, which simulates the backward SDE using only zeroth-order information of f. The key idea is to approximate the intermediate particle distribution by a Gaussian mixture, enabling score estimation without auxiliary samplers or additional model training.

Our method remains within the proximal sampling framework, while also being closely connected to diffusion-based Monte Carlo methods. Our theoretical analysis shows that it achieves comparable convergence rates when discretization and score estimation errors are properly controlled. Unlike rejection-sampling implementations of RGO, which are constrained to small step sizes by local properties of f, our algorithm may accommodate larger step sizes that reflect global features such as inter-mode distances.

Finally, we incorporate interaction terms among multiple particles, which empirically accelerate convergence. A complete theoretical characterization of these interactions, together with efficient parameter tuning across iterations, remains an important direction for future work.

Reproducibility Statement

We provide complete assumptions and full proofs of all theoretical results in Appendix A. The detailed settings of our experimental evaluations are described in Section 5 and Section B.

References

Nicholas M Boffi and Eric Vanden-Eijnden. Probability flow solution of the fokker–planck equation. *Machine Learning: Science and Technology*, 4(3):035012, jul 2023. doi: 10.1088/2632-2153/ace2aa. URL https://dx.doi.org/10.1088/2632-2153/ace2aa.

Nicolas Brosse, Alain Durmus, Éric Moulines, and Marcelo Pereyra. Sampling from a log-concave distribution with compact support with proximal langevin monte carlo. In Satyen Kale and Ohad Shamir (eds.), *Proceedings of the 2017 Conference on Learning Theory*, volume 65 of *Proceedings of Machine Learning Research*, pp. 319–342. PMLR, 07–10 Jul 2017. URL https://proceedings.mlr.press/v65/brosse17a.html.

Sébastien Bubeck, Ronen Eldan, and Joseph Lehec. Sampling from a log-concave distribution with projected langevin monte carlo. *Discrete & Computational Geometry*, 59(4):757–783, Jun 2018. ISSN 1432-0444. doi: 10.1007/s00454-018-9992-1. URL https://doi.org/10.1007/s00454-018-9992-1.

Carlson Moses Büth, Kishor Acharya, and Massimiliano Zanin. infomeasure: a comprehensive python package for information theory measures and estimators. *Scientific Reports*, 15(1):29323, Aug 2025. ISSN 2045-2322. doi: 10.1038/s41598-025-14053-5. URL https://doi.org/10.1038/s41598-025-14053-5.

Djalil Chafaï. Entropies, convexity, and functional inequalities, On Φ-entropies and Φ-Sobolev inequalities. *Journal of Mathematics of Kyoto University*, 44(2):325 – 363, 2004. doi: 10.1215/kjm/1250283556. URL https://doi.org/10.1215/kjm/1250283556.

Wenlin Chen, Mingtian Zhang, Brooks Paige, José Miguel Hernández-Lobato, and David Barber. Diffusive Gibbs sampling. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 7731–7747. PMLR, 21–27 Jul 2024. URL https://proceedings.mlr.press/v235/chen24be.html.

Yongxin Chen, Sinho Chewi, Adil Salim, and Andre Wibisono. Improved analysis for a proximal algorithm for sampling. In Po-Ling Loh and Maxim Raginsky (eds.), *Proceedings of Thirty Fifth Conference on Learning Theory*, volume 178 of *Proceedings of Machine Learning Research*, pp. 2984–3014. PMLR, 02–05 Jul 2022. URL https://proceedings.mlr.press/v178/chen22c.html.

Xiang Cheng and Peter Bartlett. Convergence of langevin mcmc in kl-divergence. In Firdaus Janoos, Mehryar Mohri, and Karthik Sridharan (eds.), *Proceedings of Algorithmic Learning Theory*, volume 83 of *Proceedings of Machine Learning Research*, pp. 186–211. PMLR, 07–09 Apr 2018. URL https://proceedings.mlr.press/v83/cheng18a.html.

Arnak Dalalyan. Further and stronger analogy between sampling and optimization: Langevin monte carlo and gradient descent. In Satyen Kale and Ohad Shamir (eds.), *Proceedings of the 2017 Conference on Learning Theory*, volume 65 of *Proceedings of Machine Learning Research*, pp. 678–689. PMLR, 07–10 Jul 2017. URL https://proceedings.mlr.press/v65/dalalyan17a.html.

Alain Durmus and Éric Moulines. Nonasymptotic convergence analysis for the unadjusted Langevin algorithm. *The Annals of Applied Probability*, 27(3):1551 – 1587, 2017. doi: 10.1214/16-AAP1238. URL https://doi.org/10.1214/16-AAP1238.

Raaz Dwivedi, Yuansi Chen, Martin J Wainwright, and Bin Yu. Log-concave sampling: Metropolis-hastings algorithms are fast! In Sébastien Bubeck, Vianney Perchet, and Philippe Rigollet (eds.),

Proceedings of the 31st Conference On Learning Theory, volume 75 of *Proceedings of Machine Learning Research*, pp. 793–797. PMLR, 06–09 Jul 2018. URL https://proceedings.mlr.press/v75/dwivedi18a.html.

Jiaojiao Fan, Bo Yuan, and Yongxin Chen. Improved dimension dependence of a proximal algorithm for sampling. In Gergely Neu and Lorenzo Rosasco (eds.), *Proceedings of Thirty Sixth Conference on Learning Theory*, volume 195 of *Proceedings of Machine Learning Research*, pp. 1473–1521. PMLR, 12–15 Jul 2023. URL https://proceedings.mlr.press/v195/fan23a.html.

Mingzhou Fan, Ruida Zhou, Chao Tian, and Xiaoning Qian. Path-guided particle-based sampling. In *Proceedings of the 41st International Conference on Machine Learning*, ICML'24. JMLR.org, 2024.

Futoshi Futami, Issei Sato, and Masashi Sugiyama. Accelerating the diffusion-based ensemble sampling by non-reversible dynamics. In Hal Daumé III and Aarti Singh (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 3337–3347. PMLR, 13–18 Jul 2020. URL https://proceedings.mlr.press/v119/futami20a.html.

Louis Grenioux, Maxence Noble, Marylou Gabrié, and Alain Oliviero Durmus. Stochastic localization via iterative posterior sampling. In *Forty-first International Conference on Machine Learning*, 2024. URL https://openreview.net/forum?id=2Gr5wZR6uc.

Wei Guo, Molei Tao, and Yongxin Chen. Provable benefit of annealed langevin monte carlo for non-log-concave sampling. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=P6IVIoGRRg.

Mert Gurbuzbalaban, Yuanhan Hu, and Lingjiong Zhu. Penalized overdamped and underdamped langevin monte carlo algorithms for constrained sampling. *Journal of Machine Learning Research*, 25(263):1–67, 2024. URL http://jmlr.org/papers/v25/22-1443.html.

Ye He, Kevin Rojas, and Molei Tao. Zeroth-order sampling methods for non-log-concave distributions: Alleviating metastability by denoising diffusion. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 71122–71161. Curran Associates, Inc., 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/8337b176767d948275d0b6d17cf4221e-Paper-Conference.pdf.

Xunpeng Huang, Hanze Dong, Yifan HAO, Yian Ma, and Tong Zhang. Reverse diffusion monte carlo. In *The Twelfth International Conference on Learning Representations*, 2024a. URL https://openreview.net/forum?id=kIPEyMSdFV.

Xunpeng Huang, Difan Zou, Hanze Dong, Yi-An Ma, and Tong Zhang. Faster sampling without isoperimetry via diffusion-based monte carlo. In Shipra Agrawal and Aaron Roth (eds.), *Proceedings of Thirty Seventh Conference on Learning Theory*, volume 247 of *Proceedings of Machine Learning Research*, pp. 2438–2493. PMLR, 30 Jun–03 Jul 2024b. URL https://proceedings.mlr.press/v247/huang24a.html.

Vasily Ilin, Peter Sushko, and Jingwei Hu. Score-based deterministic density sampling. *arXiv preprint arXiv:2504.18130*, 2025.

Beomsu Kim and Jong Chul Ye. Denoising MCMC for accelerating diffusion-based generative models. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 16955–16977. PMLR, 23–29 Jul 2023. URL https://proceedings.mlr.press/v202/kim23z.html.

Yunbum Kook, Santosh S. Vempala, and Matthew S. Zhang. In-and-out: Algorithmic diffusion for sampling convex bodies. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 108354–108388. Curran Associates, Inc., 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/c4006ff54a7bbda74c09bad6f7586f5b-Paper-Conference.pdf.

LF Kozachenko and Nikolai N Leonenko. Sample estimate of the entropy of a random vector. *Problemy Peredachi Informatsii*, 23(2):9–16, 1987.

Alexander Kraskov, Harald Stögbauer, and Peter Grassberger. Estimating mutual information. *Phys. Rev. E*, 69:066138, Jun 2004. doi: 10.1103/PhysRevE.69.066138. URL https://link.aps.org/doi/10.1103/PhysRevE.69.066138.

Yin Tat Lee, Ruoqi Shen, and Kevin Tian. Structured logconcave sampling with a restricted gaussian oracle. In Mikhail Belkin and Samory Kpotufe (eds.), *Proceedings of Thirty Fourth Conference on Learning Theory*, volume 134 of *Proceedings of Machine Learning Research*, pp. 2993–3050. PMLR, 15–19 Aug 2021. URL https://proceedings.mlr.press/v134/lee21a.html.

Gen Li, Yu Huang, Timofey Efimov, Yuting Wei, Yuejie Chi, and Yuxin Chen. Accelerating convergence of score-based diffusion models, provably. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 27942–27954. PMLR, 21–27 Jul 2024. URL https://proceedings.mlr.press/v235/li24ad.html.

Puheng Li, Zhong Li, Huishuai Zhang, and Jiang Bian. On the generalization properties of diffusion models. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 2097–2127. Curran Associates, Inc., 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/06abed94583030dd50abe6767bd643b1-Paper-Conference.pdf.

Jiaming Liang and Yongxin Chen. A proximal algorithm for sampling from non-smooth potentials. In *Proceedings of the Winter Simulation Conference*, WSC '22, pp. 3229–3240. IEEE Press, 2023a.

Jiaming Liang and Yongxin Chen. A proximal algorithm for sampling. *Transactions on Machine Learning Research*, 2023b. ISSN 2835-8856. URL https://openreview.net/forum?id=CkXOwlhf27.

Qiang Liu and Dilin Wang. Stein variational gradient descent: A general purpose bayesian inference algorithm. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016. URL https://proceedings.neurips.cc/paper_files/paper/2016/file/b3ba8f1bee1238a2f37603d90b58898d-Paper.pdf.

Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan LI, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 5775–5787. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/260a14acce2a89dad36adc8eefe7c59e-Paper-Conference.pdf.

Jianfeng Lu, Yue Wu, and Yang Xiang. Score-based transport modeling for mean-field fokker-planck equations. *Journal of Computational Physics*, 503:112859, 2024. ISSN 0021-9991. doi: https://doi.org/10.1016/j.jcp.2024.112859. URL https://www.sciencedirect.com/science/article/pii/S0021999124001086.

Yueming Lyu, Atsushi Nitanda, and Ivor Tsang. Nonparametric distributional black-box optimization via diffusion process. In *ICLR 2025 Workshop on Deep Generative Model in Machine Learning: Theory, Principle and Efficacy*, 2025. URL https://openreview.net/forum?id=DOXsyHprRN.

Emile Pierret and Bruno Galerne. Diffusion models for gaussian distributions: Exact solutions and wasserstein errors. In *Forty-second International Conference on Machine Learning*, 2025. URL https://openreview.net/forum?id=bxYbxzCI2R.

Lorenz Richter and Julius Berner. Improved sampling via learned diffusions. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=h4pNROsO06.

G. O. Roberts and O. Stramer. Langevin diffusions and metropolis-hastings algorithms. *Methodology And Computing In Applied Probability*, 4(4):337–357, Dec 2002. ISSN 1573-7713. doi: 10.1023/A:1023562417138. URL https://doi.org/10.1023/A:1023562417138.

Gareth O. Roberts and Jeffrey S. Rosenthal. Optimal scaling of discrete approximations to langevin diffusions. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 60(1): 255–268, 1998. ISSN 13697412, 14679868. URL http://www.jstor.org/stable/2985986.

Gareth O. Roberts and Richard L. Tweedie. Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli*, 2(4):341 – 363, 1996.

Francisco Vargas, Will Sussman Grathwohl, and Arnaud Doucet. Denoising diffusion samplers. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=8pvnfTAbu1f.

Santosh Vempala and Andre Wibisono. Rapid convergence of the unadjusted langevin algorithm: Isoperimetry suffices. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/65a99bb7a3115fdede20da98b08a370f-Paper.pdf.

A Proofs and Technical Details

A.1 Derivation of the Surrogate Score Estimator

We begin by substituting π^X with

$$\hat{q}_{k+1}(x \mid Y_{k+\frac{1}{2}}, X_k) \propto \frac{1}{N}\sum_{j=1}^{N} \pi^{X|Y=y_j}(x)\, \frac{\pi^Y(y_j)}{\hat{q}_{k+\frac{1}{2}}(y_j \mid X_k)},$$

where $\hat{q}_{k+1/2}(y \mid X_k) = \frac{1}{N}\sum_i \mathcal{N}(y; x_i, hI_d)$. This follows from the fact that the forward heat flow (4) transports the empirical Dirac mixture $\frac{1}{N}\sum_i \delta_{x_i}$ to the Gaussian mixture $\hat{q}_{k+1/2}$.

Using

$$\pi^{X|Y=y_j}(x) = \frac{\exp\!\big(-f(x) - \frac{1}{2h}\|x-y_j\|^2\big)}{\int \exp\!\big(-f(z) - \frac{1}{2h}\|z-y_j\|^2\big)\,\mathrm{d}z},$$
$$\pi^Y(y_j) \propto \int \exp\!\big(-f(z) - \frac{1}{2h}\|z-y_j\|^2\big)\,\mathrm{d}z,$$

we obtain

$$\begin{aligned}\hat{q}_{k+1}(x \mid Y_{k+\frac{1}{2}}, X_k) &\propto \frac{1}{N}\sum_{j=1}^{N} \frac{\exp\!\big(-f(x) - \frac{1}{2h}\|x-y_j\|^2\big)}{\hat{q}_{k+\frac{1}{2}}(y_j \mid X_k)} \\ &\propto \left(\frac{1}{N}\sum_{j=1}^{N} \frac{\mathcal{N}(x; y_j, hI_d)}{\hat{q}_{k+1/2}(y_j \mid X_k)}\right)\exp(-f(x)) =: g_N^{k+1/2}(x)\exp(-f(x)),\end{aligned}$$

where $g_N^{k+1/2}$ is the unnormalized density of an N-component weighted Gaussian mixture. As $N \to \infty$,

$$\lim_{N\to\infty} g_N^{k+1/2}(x) = \int \frac{\mathcal{N}(x; y, hI_d)}{\hat{q}_{k+1/2}(y \mid X_k)}\hat{q}_{k+1/2}(y \mid X_k)\mathrm{d}y = \text{const},$$

and therefore $\hat{q}_{k+1}(\cdot \mid Y_{k+1/2}, X_k) \to \pi^X(\cdot)$.

Since $\hat{q}_{k+1}$ can be regarded as the product of (i) a weighted Gaussian mixture term and (ii) the exponential factor $\exp(-f)$, the surrogate score function $\hat{s}_t(x_t) := \nabla \log(\hat{q}_{k+1}P_{h-t})(x_t)$ can be computed using only the particles and evaluations of f. Following the derivation in Lyu et al. (2025), we obtain

$$\begin{aligned}\hat{s}_t(x_t) &= \nabla \log\left(\int g_N^{k+1/2}(x_0)\, p(x_t \mid x_0)\, e^{-f(x_0)}\mathrm{d}x_0\right) \\ &= \frac{\int g_N^{k+1/2}(x_0)\, \nabla p(x_t \mid x_0)\, e^{-f(x_0)}\mathrm{d}x_0}{\int g_N^{k+1/2}(x_0)\, p(x_t \mid x_0)\, e^{-f(x_0)}\mathrm{d}x_0}.\end{aligned}$$

Here

$$p(x_t \mid x_0) = \frac{1}{(2\pi\sigma_t^2)^{d/2}}\exp\!\left(-\frac{\|x_t - x_0\|^2}{2\sigma_t^2}\right), \quad \sigma_t^2 = h - t,$$

so that

$$\nabla p(x_t \mid x_0) = p(x_t \mid x_0)\,\frac{x_0 - x_t}{\sigma_t^2}.$$

Substituting into the above expression yields

$$\begin{aligned}\hat{s}_t(x_t) &= \frac{\int g_N^{k+1/2}(x_0)\, p(x_t \mid x_0)\, \frac{x_0-x_t}{\sigma_t^2} e^{-f(x_0)}\mathrm{d}x_0}{\int g_N^{k+1/2}(x_0)\, p(x_t \mid x_0)\, e^{-f(x_0)}\mathrm{d}x_0} \\ &= \frac{\int g_N^{k+1/2}(x_0 \mid x_t)\, \frac{x_0-x_t}{\sigma_t^2} e^{-f(x_0)}\mathrm{d}x_0}{\int g_N^{k+1/2}(x_0 \mid x_t)\, e^{-f(x_0)}\mathrm{d}x_0},\end{aligned}$$

which is exactly the score function of surrogate dynamics (11).

Remark 1. Lyu et al. (2025) consider distributional black–box optimization for maximizing f. When $p^k(x_0)$ is a Gaussian mixture, they show that setting $p^{k+1}(x_0) \propto p^k(x_0)\exp(-f(x_0)/\lambda)$ as the denoised distribution at time 0 in VP diffusion, the score at each time t can be computed directly from samples of the posterior Gaussian mixture, without any auxiliary sampler.

A structurally similar mechanism appears in our particle-based proximal sampling. To sample from $\pi(x) \propto \exp(-f(x))$ (in practice, by moving samples from $\hat{q}_{k+1/2}(\, y \mid X_k\,)$ toward $\hat{q}_{k+1}(\, x \mid Y_{k+1/2}, X_k\,)$), we place the denoised distribution $g_N^{k+1/2}(x)\exp(-f(x))$, where $g_N^{k+1/2}$ is the Gaussian mixture obtained by smoothing the empirical particles $Y_{k+1/2}$. The score identity yields a backward update that uses only zeroth-order evaluations of f.

A.2 Derivation of the (Unnormalized) Gaussian Mixture Posterior

When sampling from a N-component Gaussian mixture, we first draw a component index according to the relative weights and then sample from the corresponding Gaussian component. For this reason, we often ignore constant multiplicative factors of the mixture distribution for simplicity.

We write $g_N^{k+1/2}(x_0) \propto \sum_j \alpha_j \mathcal{N}(x_0; y_j, hI_d)$ with $\alpha_j \coloneqq 1/\hat{q}^{k+1/2}(y_j)$. By Bayes' rule,

$$\begin{aligned} g_N^{k+\frac{1}{2}}(x_0 \mid x_t) &= \frac{g_N^{k+\frac{1}{2}}(x_0)\, p(x_t \mid x_0)}{\int g_N^{k+\frac{1}{2}}(u)\, p(x_t \mid u)\,\mathrm{d}u} \\ &= \frac{\sum_{j=1}^N \alpha_j \mathcal{N}(x_0; y_j, hI_d)\,\mathcal{N}(x_t; x_0, \sigma_t^2 I_d)}{\sum_{j=1}^N \alpha_j \int \mathcal{N}(u; y_j, hI_d)\,\mathcal{N}(x_t; u, \sigma_t^2 I_d)\,\mathrm{d}u} \\ &\propto \sum_{j=1}^N \alpha_j \mathcal{N}(x_0; y_j, hI_d)\,\mathcal{N}(x_t; x_0, \sigma_t^2 I_d). \end{aligned} \tag{15}$$

We now complete the square in the exponent of the product $\mathcal{N}(x_0; y_j, hI_d)\,\mathcal{N}(x_t; x_0, \sigma_t^2 I_d)$ viewed as a function of x_0. Using the identity

$$\frac{1}{2h}\|x_0 - y_j\|^2 + \frac{1}{2\sigma_t^2}\|x_t - x_0\|^2 = \frac{1}{2\bar{\sigma}^2}\|x_0 - m_j(x_t)\|^2 + \frac{1}{2(h+\sigma_t^2)}\|x_t - y_j\|^2,$$

with

$$\bar{\sigma}^2 \coloneqq \Big(h^{-1} + \sigma_t^{-2}\Big)^{-1} = \frac{h\,\sigma_t^2}{h+\sigma_t^2}, \quad m_j(x_t) \coloneqq \bar{\sigma}^2\Big(h^{-1} y_j + \sigma_t^{-2} x_t\Big) = \frac{\sigma_t^2}{h+\sigma_t^2} y_j + \frac{h}{h+\sigma_t^2} x_t,$$

we obtain the product-of-Gaussians factorization

$$\mathcal{N}(x_0; y_j, hI_d)\,\mathcal{N}(x_t; x_0, \sigma_t^2 I_d) = \mathcal{N}(x_t; y_j, (h+\sigma_t^2) I_d)\,\mathcal{N}(x_0; m_j(x_t), \bar{\sigma}^2 I_d).$$

Substituting this into (15) yields

$$g_N^{k+\frac{1}{2}}(x_0 \mid x_t) \propto \sum_{j=1}^N \alpha_j\, \mathcal{N}(x_t; y_j, (h+\sigma_t^2) I_d)\,\mathcal{N}(x_0; m_j(x_t), \bar{\sigma}^2 I_d).$$

Therefore $g_N^{k+1/2}(x_0 \mid x_t)$ is again a Gaussian mixture with a common covariance $\bar{\sigma}^2 I_d$ and updated means $m_j(x_t)$. Writing the relative weights as

$$w_j(x_t) \coloneqq \alpha_j\, \mathcal{N}(x_t; y_j, (h+\sigma_t^2) I_d) = \frac{1}{\hat{q}^{k+\frac{1}{2}}(y_j | X_k)}\, \mathcal{N}(x_t; y_j, (h+\sigma_t^2) I_d),$$

we obtain the desired posterior decomposition

$$g_N^{k+\frac{1}{2}}(x_0 \mid x_t) \propto \sum_{j=1}^N w_j(x_t)\, \mathcal{N}(x_0; m_j(x_t), \bar{\sigma}^2 I_d).$$

Finally, note that using $W(x_t) = \sum_j w_j(x_t)$, the normalized posterior distribution can be written as

$$g_N^{k+\frac{1}{2}}(x_0 \mid x_t) = \frac{1}{W(x_t)} \sum_{j=1}^N \tilde{w}_j(x_t)\, \mathcal{N}(x_0; m_j(x_t), \bar{\sigma}^2 I_d).$$

A.3 Analysis for Time-discretization Error

Our algorithm simulates the surrogate version of diffusion process (5) by dividing the dynamics over horizon h into T steps. We first analyze the discretization error under the assumption that the score function is perfectly estimated (i.e., $N, M \to \infty$). Each step simulates the time evolution over an interval of length $\eta := h/T$, corresponding to the segment $(\ell-1)\eta + t \in [(\ell-1)\eta, \ell\eta]$ with $t \in [0, \eta]$ for $\ell = 1, \dots, T$. In the ideal dynamics (5), the drift term is time-dependent, whereas in the discretized scheme we approximate it by fixing the drift at the beginning of each step. Following an argument similar to that in Vempala & Wibisono (2019), we obtain the following result.

Lemma 3. *Fix a step index $\ell \in \{1, \dots, T\}$ and let $\eta := h/T$. Let $(\nu_t)_{t\in[0,\eta]}$ denote the ideal backward marginals in (5) at elapsed time $h-(\ell-1)\eta-t$ for $t \in [0, \eta]$, and $(\mu_t)_{t\in[0,\eta]}$ denote the frozen-drift approximation within this step. Then for $u \geq 1$ and all $t \in [0, \eta]$,*

$$\frac{\mathrm{d}}{\mathrm{d}t} H_{\nu_t}(\mu_t) \;\leq\; -\Big(\frac{1}{2} - \frac{1}{4u^2}\Big)\, J_{\nu_t}(\mu_t) \;+\; u^2\, \mathbb{E}_{\mu_{0,t}}\big[\|\tilde{s}_0(z_0) - \tilde{s}_t(z_t)\|^2\big]. \tag{16}$$

Equivalently, integrating over $t \in [0, \eta]$ yields

$$H_{\nu_\eta}(\mu_\eta) - H_{\nu_0}(\mu_0) \;\leq\; -\Big(\frac{1}{2} - \frac{1}{4u^2}\Big) \int_0^\eta J_{\nu_t}(\mu_t)\,\mathrm{d}t \;+\; u^2 \int_0^\eta \mathbb{E}_{\mu_{0,t}}\big[\|\tilde{s}_0(z_0) - \tilde{s}_t(z_t)\|^2\big]\,\mathrm{d}t.$$

Proof. For convenience to analyze the behavior at $(\ell-1)\eta + t$ with $t \in [0, \eta]$ which corresponds to the time interval $[(\ell-1)\eta, \ell\eta]$ in original backward process defined in (5), we restate the SDE and the associated Fokker–Plank equation as: for $t \in [0, h]$,

$$\mathrm{d}Z_t^{\leftarrow} = \tilde{s}_t(Z_t^{\leftarrow})\,\mathrm{d}t + \mathrm{d}B_t, \quad \mathrm{law}(Z_t^{\leftarrow}) = \nu_t,$$
$$\partial_t \nu_t = -\mathrm{div}\big(\nu_t \tilde{s}_t\big) + \frac{1}{2}\Delta\nu_t = \nabla \cdot \Big(\nu_t\big(-\tilde{s}_t + \frac{1}{2}\nabla \log \nu_t\big)\Big),$$

where we define the score function $\tilde{s}_t := \nabla \log\big(\pi^X P_{h-(\ell-1)\eta-t}\big)$.

In contrast, the time-discretized approximation corresponds to the process

$$\mathrm{d}z_t = \tilde{s}_0(z_0)\,\mathrm{d}t + \mathrm{d}B_t, \quad \mathrm{law}(z_t) = \mu_t, \tag{17}$$

where the drift is frozen at the beginning of the step.

The associated Fokker–Planck equation is

$$\begin{aligned}\partial_t \mu_t(z_t \mid z_0) &= -\mathrm{div}\big(\mu_t(z_t \mid z_0)\, \tilde{s}_0(z_0)\big) + \frac{1}{2}\Delta\mu_t(z_t \mid z_0) \\ &= \nabla \cdot (\mu_t(z_t \mid z_0)(-\tilde{s}_0(z_0) + \frac{1}{2}\nabla \log \mu_t(z_t \mid z_0))).\end{aligned}$$

Averaging over $z_0 \sim \mu_0$, this becomes

$$\partial_t \mu_t = -\nabla \cdot \Big(\mu_t\big(-\mathbb{E}_{\mu_{0|t}}[\tilde{s}_0(z_0) \mid z_t] + \frac{1}{2}\nabla \log \mu_t\big)\Big).$$

The time derivative of the KL divergence is then

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}t} H_{\nu_t}(\mu_t) &= \frac{\mathrm{d}}{\mathrm{d}t} \int \mu_t \log \frac{\mu_t}{\nu_t}\,\mathrm{d}z \\ &= \int \Big[(\partial_t \mu_t) \log \frac{\mu_t}{\nu_t} + \Big(\partial_t \mu_t - \mu_t \frac{\partial_t \nu_t}{\nu_t}\Big)\Big]\,\mathrm{d}z \\ &= \int \Big[\nabla \cdot \Big(\mu_t\big(-\mathbb{E}_{\mu_{0|t}}[\tilde{s}_0(z_0) \mid z_t] + \frac{1}{2}\nabla \log \mu_t\big)\Big)\, \log \frac{\mu_t}{\nu_t} \\ &\qquad - \frac{\mu_t}{\nu_t}\, \nabla \cdot \Big(\nu_t\big(-\tilde{s}_t + \frac{1}{2}\nabla \log \nu_t\big)\Big)\Big]\,\mathrm{d}z \\ &= \int \Big[-\big\langle \mu_t\big(-\mathbb{E}_{\mu_{0|t}}[\tilde{s}_0(z_0) \mid z_t] + \tfrac{1}{2}\nabla \log \mu_t\big), \nabla \log \tfrac{\mu_t}{\nu_t}\big\rangle \\ &\qquad + \big\langle \nu_t\big(-\tilde{s}_t + \frac{1}{2}\nabla \log \nu_t\big), \frac{\mu_t}{\nu_t}\nabla \log \frac{\mu_t}{\nu_t}\big\rangle\Big]\,\mathrm{d}z \\ &= \int \Big\langle -\frac{1}{2}\nabla \log \frac{\mu_t}{\nu_t} + \mathbb{E}_{\mu_{0|t}}[\tilde{s}_0(z_0) \mid z_t] - \tilde{s}_t,\, \nabla \log \frac{\mu_t}{\nu_t}\Big\rangle\, \mu_t \mathrm{d}z.\end{aligned}$$

Here, the third equality substitutes the Fokker–Planck equations for $\partial_t \mu_t$ and $\partial_t \nu_t$ and also uses that the integral of $\partial_t \mu_t$ vanishes due to mass conservation. The forth equality uses integration by parts (assuming sufficiently fast decay at infinity) and the identity.

Simplifying, we obtain

$$\frac{\mathrm{d}}{\mathrm{d}t} H_{\nu_t}(\mu_t) = -\frac{1}{2} J_{\nu_t}(\mu_t) + \mathbb{E}_{\mu_t}\Big[\big\langle \mathbb{E}_{\mu_{0|t}}[\tilde{s}_0(z_0) \mid z_t] - \tilde{s}_t, \nabla \log \frac{\mu_t}{\nu_t} \big\rangle\Big].$$

Applying $\langle a, b\rangle \le u^2\|a\|^2 + \frac{1}{4u^2}\|b\|^2$, we conclude

$$\frac{\mathrm{d}}{\mathrm{d}t} H_{\nu_t}(\mu_t) \le -\left(\frac{1}{2} - \frac{1}{4u^2}\right) J_{\nu_t}(\mu_t) + u^2 \mathbb{E}_{\mu_{0,t}(z_0,z_t)}\big[\|\tilde{s}_0(z_0) - \tilde{s}_t(z_t)\|^2\big].$$

□

This establishes the discretization error bound of one diffution step in terms of the deviation between the frozen score $\tilde{s}_0$ and the ideal time-dependent score $\tilde{s}_t$.

Assumption 1 (Smoothness of the interim distribution)**.** *The interim distribution $\pi^X P_{h-(\ell-1)\eta}$ is L_ℓ-smooth, i.e., the potential gradient $\nabla \log(\pi^X P_{h-(\ell-1)\eta})$ is L_ℓ-Lipschitz.*

Assumption 2 (Lipschitz condition along the time direction)**.** *The expected score satisfies a Lipschitz-type condition along the time direction:*

$$\mathbb{E}_{\mu_t}\big[\|\tilde{s}_0(z_t) - \tilde{s}_t(z_t)\|^2\big] \le C_{t,\ell}^2\, t^2.$$

Corollary 1. *Assume $\nu_0 \coloneqq \pi^X P_{(\ell-1)\eta}$ satisfies LSI with constant $C_{LSI}(\nu_0)$ and under Assumption 1 and Assumption 2,*

$$\frac{\mathrm{d}}{\mathrm{d}t} H_{\nu_t}(\mu_t) \le -\left(\frac{1}{2} - \frac{1}{4u^2}\right) J_{\nu_t}(\mu_t) + u^2\left(4\eta^2 L_\ell^4\, C_{LSI}(\nu_0)\, H_{\nu_0}(\mu_0) + \eta d C\right), \tag{18}$$

$$\textit{where} \quad C = \sup_l \sup_t 2tL_\ell^3 + L_\ell^2 + tC_{t,\ell}^2/d. \tag{19}$$

Proof. We bound the second term of the right hand side in the inequality in Lemma 3 using Assumption 1 and Assumption 2,

$$\begin{aligned}\mathbb{E}_{\mu_{0,t}(z_0,z_t)}\big[\|\tilde{s}_0(z_0) - \tilde{s}_t(z_t)\|^2\big] &\le \mathbb{E}_{\mu_{0,t}(z_0,z_t)}\big[\|\tilde{s}_0(z_0) - \tilde{s}_0(z_t)\|^2\big] + \mathbb{E}_{\mu_t}\big[\|\tilde{s}_0(z_t) - \tilde{s}_t(z_t)\|^2\big] \\ &\le L_\ell^2 \mathbb{E}_{\mu_{0,t}(z_0,z_t)}\big[\|z_0 - z_t\|^2\big] + C_{t,\ell}^2\, t^2.\end{aligned}$$

Under the discretization update

$$z_t = z_0 + \tilde{s}_0(x_0)\, t + \sqrt{t}\,\xi, \quad \xi \sim \mathcal{N}(0, I_d),$$

we have

$$\begin{aligned}\mathbb{E}_{\mu_{0,t}(z_0,z_t)}\big[\|z_0 - z_t\|^2\big] &= \mathbb{E}_{\mu_0}[\|\tilde{s}_0(x_0)\, t + \sqrt{t}\,\xi\|^2] \\ &= t^2\, \mathbb{E}_{\mu_0}[\|\tilde{s}_0(x_0)\|^2] + td \\ &\le t^2(4L_\ell^2\, C_{\mathrm{LSI}}(\nu_0)\, H_{\nu_0}(\mu_0) + 2dL_\ell) + td.\end{aligned}$$

The last inequality comes from Lemma 12 in Vempala & Wibisono (2019) with Assumption 1 and ν_0 satisfying Talagrand's inequality with constant $C_{\mathrm{LSI}}(\nu_0)^{-1}$. Putting them altogether, we obtain (18) where C is independent on ℓ by taking supremum as (19). □

Proposition 2 (One-step bound for the diffusion-approximated proximal sampler without score estimation error)**.** *Assume Assumption 1 and Assumption 2 hold, and that π^X satisfies an LSI with constant $C_{LSI}(\pi^X)$. Let C be as in (19). Let overall step size $h > 0$ be split into T steps with $\eta \coloneqq h/T$. In the regime $N, M \to \infty$ where the score estimation error vanishes, distribution at k-th iteration ρ_k^X satisfies*

$$\begin{aligned}H_{\pi^X}(\rho_{k+1}^X) \le\ & \frac{H_{\pi^X}(\rho_k^X)}{(1 + h/C_{LSI}(\pi^X))^{2-1/(2u^2)}} \\ &+ 2u^4 C_{LSI}(\pi^X)\big((1 + h/C_{LSI}(\pi^X))^{1/(2u^2)} - 1\big)\Big\{4\eta^2 \bar{L}^2 (C_{LSI}(\pi^X) + h)\bar{H} + \eta d C\Big\}\end{aligned}$$

for $u \ge 1$, where $\bar{H}$ is the supremum of KL divergence between the interim distribution of the backward denoised distribution at timestep $t = \ell\eta$ and $\pi^Y Q_{\ell\eta}$ and $\bar{L} = \sup_\ell L_\ell$

Proof. As shown in Corollary 13 of Chafaï (2004), the logarithmic Sobolev constant satisfies

$$C_{LSI}(\pi^X P_t) \le C_{LSI}(\pi^X) + t.$$

For $t \in [0, \eta]$, the law ν_t corresponds to the ideal denoising path at time $(\ell-1)\eta + t$, namely $\nu_t = \pi^X P_{h-(\ell-1)\eta-t}$. Hence its logarithmic Sobolev constant can be bounded as

$$C_{LSI}(\nu_t) \le C_{LSI}(\pi^X) + h - (\ell-1)\eta - t =: \alpha_\ell - t.$$

From Corollary 1 combined with LSI, we obtain

$$\frac{\mathrm{d}}{\mathrm{d}t} H_{\nu_t}(\mu_t) \le -\frac{1 - 1/(2u^2)}{\alpha_\ell - t} H_{\nu_t}(\mu_t) + u^2[4\eta^2 L_\ell^4 C_{LSI}(\nu_0) H_{\nu_0}(\mu_0) + \eta d C].$$

Applying Gronwall's inequality, where we multiply by the integrating factor: $(\alpha_\ell - t)^{-1+1/(2u^2)}$,

$$\frac{\mathrm{d}}{\mathrm{d}t}\Big\{(\alpha_{\nu_t} - t)^{-1+1/(2u^2)} H_{\nu_t}(\mu_t)\Big\} \le (\alpha_{\nu_t} - t)^{-1+1/(2u^2)} u^2 \Big[4\eta^2 L_{v_0}^2 C_{LSI}(\nu_0) H_{\nu_0}(\mu_0) + \eta d C\Big]$$

and integrating over $t \in [0, \eta]$, we obtain

$$\begin{aligned}
&(\alpha_\ell - \eta)^{-1+1/(2u^2)} H_{\nu_\eta}(\mu_\eta) - \alpha_\ell^{-1+1/(2u^2)} H_{\nu_0}(\mu_0) \\
&\quad \le 2u^4 (\alpha_\ell^{1/(2u^2)} - (\alpha_\ell - \eta)^{1/(2u^2)}) \Big\{4\eta^2 L_{v_0}^2 C_{LSI}(\nu_0) H_{\nu_0}(\mu_0) + \eta d C\Big\}.
\end{aligned}$$

Recall that we μ_t, ν_t is in ℓ-th step in time discretized diffusion. Using the fact that μ_η, ν_η is equivalent to μ_0, ν_0 in $\ell+1$-th step and $\alpha_\ell - \eta = \alpha_{\ell+1}$, we iterating this inequality over $\ell = 1, \dots, T$ to obtain

$$\begin{aligned}
&C_{\mathrm{LSI}}(\pi^X)^{-1+1/(2u^2)} H_{\pi^X}(p_{k+1}^X) - (C_{\mathrm{LSI}}(\pi^X) + h)^{-1+1/(2u^2)} H_{\pi^X}(p_{k+1/2}^Y) \\
&\le 2u^4 \big((C_{\mathrm{LSI}}(\pi^X) + h)^{1/(2u^2)} - C_{\mathrm{LSI}}(\pi^X)^{1/(2u^2)}\big) \Big\{4\eta^2 \bar{L}^2 (C_{\mathrm{LSI}}(\pi^X) + h)\bar{H} + \eta d C\Big\},
\end{aligned}$$

where $\bar{H} \coloneqq \sup_l H_{\nu_0}(\mu_0)$, the supremum of the KL divergence between the updated distribution in each diffusion step and corresponding ideal distribution without distretization error and $\bar{L} \coloneqq L_\ell$.

Equivalently,

$$\begin{aligned}
H_{\pi^X}(p_{k+1}^X) \le\ & \frac{H_{\pi^Y}(p_{k+1/2}^Y)}{(1 + h/C_{\mathrm{LSI}}(\pi^X))^{1-1/(2u^2)}} \\
&+ 2u^4 C_{\mathrm{LSI}}(\pi^X) \big((1 + h/C_{LSI}(\pi^X))^{1/(2u^2)} - 1\big) \Big\{4\eta^2 L_\ell^2 (C_{LSI}(\pi^X) + h)\bar{H} + \eta d C\Big\}.
\end{aligned}$$

This corresponds to the time discretized version of the KL contraction in (8) with $q = 1$. We finalize this proof by applying inequality (7) in Lemma 1 with $q = 1$.

□

https://anonymous.4open.science/r/zod-ps-662B

A.4 Analysis for Score Estimation Error

In this section, we evaluate the error of score estimation under a set of assumptions. We do not necessarily generate $Y_{k+1/2}$ by evolving particles from X_k, but may instead sample $Y_{k+1/2}$ directly from the Gaussian mixture. In this case, we denote the number of $Y_{k+1/2}$ particles by N'. We prove that the estimation error scales as $O(1/N' + 1/M)$, which converges to zero in the limit of $N', M \to 0$.

Recall that $\hat{q}_{k+1/2}(y|X_k) = \frac{1}{N}\sum_i \mathcal{N}(x_i, hI_d)$ and we further define (normalized) Gaussian mixture density

$$\hat{g}_{N'}^{k+\frac{1}{2}}(x \mid X_k) \propto \frac{1}{N'} \sum_{j=1}^{N'} \frac{\mathcal{N}(x; y_j, hI_d)}{\hat{q}^{k+\frac{1}{2}}(y_j \mid X_k)}.$$

At time t, the true score in the surrogate dynamics is expressed as

$$\hat{s}_{N'}(z_t, t) = \nabla \log\big((\hat{g}_{N'}^{k+\frac{1}{2}} \mid_{X_k} e^{-f}) P_{h-t}\big)(z_t).$$

We approximate this by

$$\hat{s}_{N',M}(z_t, t) = \frac{\frac{1}{M}\sum_{m=1}^{M} \exp\big(-f(\hat{z}^{(m)})\big)(\hat{z}^{(m)} - z_t)/(h-t)}{\frac{1}{M}\sum_{m=1}^{M} \exp\big(-f(\hat{z}^{(m)})\big)}, \quad \hat{z}^{(m)} \sim \hat{g}_{N'}^{k+\frac{1}{2}}|_{X_k}(\cdot \mid z_t). \tag{20}$$

Compared to the setting with only time discretization in Section A.3, the drift term $\tilde{s}_0$ in one denoising step of the discretized dynamics (17) is replaced by the estimator $\hat{s}_{N',M}(\cdot, (\ell-1)\eta)$ given in (20). Using this substitution, the statement of Lemma 3, originally expressed as (16), is modified accordingly as follows:

$$\frac{\mathrm{d}}{\mathrm{d}t} H_{\nu_t}(\mu_t) \;\le\; -\Big(\frac{1}{2} - \frac{1}{4u^2}\Big) J_{\nu_t}(\mu_t) \;+\; u^2\, \mathbb{E}_{\mu_{0,t}}\big[\|\hat{s}_{N',M}(z_0, (\ell-1)\eta) - \tilde{s}_t(z_t)\|^2 \mid U, \hat{Z}|_U\big]$$

where $U \sim_{N'} \hat{q}_{k+1/2}(\cdot \mid X_k)$ and $\hat{Z}|_U \sim_M \hat{g}_{N'}^{k+\frac{1}{2}}(\cdot \mid z_0)$.

The expectation in the second term of right hand side is evaluated as:

$$\begin{aligned}
&\mathbb{E}_{\mu_{0,t}}\big[\|\hat{s}_{N',M}(z_0, (\ell-1)\eta) - \tilde{s}_t(z_t)\|^2 \mid U, \hat{Z}|_U\big] \\
&\quad\le 2\mathbb{E}_{\mu_{0,t}}\big[\|\tilde{s}_0(z_0) - \tilde{s}_t(z_t)\|^2\big] + 2\mathbb{E}_{\mu_0}\big[\|\hat{s}_{N',M}(z_0, (\ell-1)\eta) - \tilde{s}_0(z_0)\|^2 \mid U, \hat{Z}|_U\big] \\
&\quad\le 2\mathbb{E}_{\mu_{0,t}}\big[\|\tilde{s}_0(z_0) - \tilde{s}_t(z_t)\|^2\big] + 4\mathbb{E}_{\mu_0}\big[\|\hat{s}_{N'}(z_0, (\ell-1)\eta) - \tilde{s}_0(z_0)\|^2 \mid U\big] \\
&\qquad + 4\mathbb{E}_{\mu_0}\big[\|\hat{s}_{N',M}(z_0, (\ell-1)\eta) - \hat{s}_{N'}(z_0, (\ell-1)\eta)\|^2 \mid U, \hat{Z}|_U\big],
\end{aligned}$$

which means we can independently evaluate the errors from time discretization (the first term), from finite N' (the second term) and from finite M (the third term).

Score error from finite N'. Let $\tau \coloneqq h - (\ell-1)\eta$ and z_τ equivalent to z_0, initial sample of ℓ-th denoising step. For $z_0 \in \mathbb{R}^d$, define the conditional law, where denoised $\hat{z}$ is conditioned by z_τ,

$$\pi^X(\hat{z} \mid z_\tau) \;\propto\; \exp\Big(-f(\hat{z}) - \frac{\|\hat{z} - z_\tau\|^2}{2\tau}\Big),$$

with normalizer

$$Z(z_\tau) \coloneqq \int \exp\Big(-f(\hat{z}) - \frac{\|\hat{z} - z_\tau\|^2}{2\tau}\Big)\,\mathrm{d}\hat{z}.$$

Let

$$G(z_\tau) \;\coloneqq\; \frac{1}{Z(z_\tau)} \int \hat{g}_{N'}^{k+\frac{1}{2}}(\hat{z} \mid X_k)\, \exp\Big(-f(\hat{z}) - \frac{\|\hat{z} - z_\tau\|^2}{2\tau}\Big)\,\mathrm{d}\hat{z}.$$

Here we assume:

Assumption 3 (Bounded Covariance of the conditioned target distribution)**.** *For all $\ell = 1, \cdots L$ and z_τ, there exists $C_v > 0$ s.t. $\sup \mathrm{Var}_{\pi^X(\hat{z}|z_\tau)}[\hat{z}] \;\le\; C_V$.*

Assumption 4 (Lower bound of density ratio between the perturbed distributions)**.**

$$\frac{\pi^X(\hat{z} \mid z_\tau) * \mathcal{N}(0, hI_d)(u)}{\hat{\rho}_{k+1/2}^Y(u)} \;\ge\; K \;>\; 0$$

for a.e. u, hence $G(z_\tau) \ge K$.

Lemma 4 (finite-N' error is $O(1/N')$)**.** *Under Assumption 3 and Assumption 4, conditioned on $U \sim_{N'} \hat{q}_{k+1/2}(\cdot \mid X_k)$, the error term induced by the finite base sample size N' satisfies*

$$\begin{aligned}
\mathbb{E}_{\mu_0}\Big[\|\hat{s}_{N'}(z_\tau, h-\tau) - s_{h-\tau}(z_\tau)\|^2 \,\Big|\, U\Big] \le \frac{2K^{-2}C_V}{N'\tau^2} \mathbb{E}_{\mu_0, \pi^X(\hat{z}|z_\tau)} \Big[\chi^2\big(\mathcal{N}(\hat{z}, hI_d) \,\|\, \hat{\rho}_{k+1/2}^Y\big) \\
+ \chi^2\big(\pi^X(\hat{z} \mid z_\tau) * \mathcal{N}(0, hI_d) \,\|\, \hat{\rho}_{k+1/2}^Y\big)\Big].
\end{aligned}$$

In particular, under bounded χ^2-divergences, the second term on the right-hand side of the inequality in the context is $O(1/N)$.

Proof. By Tweedie's formula at time $\tau = h - (\ell - 1)\eta$, we have

$$\mathbb{E}_{\mu_0}\Big[\|\hat{s}_{N'}(z_\tau, h-\tau) - s_{h-\tau}(z_\tau)\|^2 \,\Big|\, U\Big] = \frac{1}{\tau^2}\,\mathbb{E}_{\mu_0}\left[\Big\|\mathbb{E}_{\hat{z}\sim\hat{g}_{N'}^{k+\frac{1}{2}}(\cdot|z_\tau)}[\hat{z}] - \mathbb{E}_{\hat{z}\sim\pi^X(\cdot|z_\tau)}[\hat{z}]\Big\|^2\right]$$
$$= \frac{1}{\tau^2}\,\mathbb{E}_{\mu_0}\left[\Big\|\int \hat{z}\Big(\tfrac{\hat{g}_{N'}^{k+\frac{1}{2}}(\hat{z}|X_k)}{G(z_\tau)} - 1\Big)\,\pi^X(\hat{z}\mid z_\tau)\,d\hat{z}\Big\|^2\right].$$

Let $C(z_\tau) \coloneqq \int \hat{z}\,\pi^X(\hat{z}\mid z_\tau)\,\mathrm{d}\hat{z}$. Since $\int(\frac{\hat{g}_{N'}^{k+\frac{1}{2}}(\hat{z}|X_k)}{G(z_\tau)} - 1)\,\pi^X(\hat{z}\mid z_\tau)\,d\hat{z} = 0$, we can center the integrand and apply Cauchy–Schwarz to obtain

$$\mathbb{E}_{\mu_0}\Big[\|\hat{s}_{N'}(z_\tau, h-\tau) - s_{h-\tau}(z_\tau)\|^2 \,\Big|\, U\Big] \le \frac{1}{\tau^2}\,\mathbb{E}_{\mu_0}\big[\mathrm{Var}_{\pi^X(\hat{z}|z_\tau)}(\hat{z})\cdot \mathrm{Var}_{\pi^X(\hat{z}|z_\tau)}\big(\tfrac{\hat{g}_{N'}^{k+\frac{1}{2}}(\hat{z}|X_k)}{G(z_\tau)}\big)\big]$$
$$\le \frac{C_V}{\tau^2}\,\mathbb{E}_{\mu_0}\left[\int \frac{1}{G(z_\tau)^2}\big(\hat{g}_{N'}^{k+\frac{1}{2}}(\hat{z}\mid X_k) - G(z_\tau)\big)^2\,\pi^X(\hat{z}\mid z_\tau)\,\mathrm{d}\hat{z}\right]$$
$$\le \frac{K^{-2}C_V}{\tau^2}\,\mathbb{E}_{\mu_0}\left[\int \big(\hat{g}_{N'}^{k+\frac{1}{2}}(\hat{z}\mid X_k) - G(z_\tau)\big)^2\,\pi^X(\hat{z}\mid z_\tau)\,\mathrm{d}\hat{z}\right].$$

The second inequality is from Assumption 3 and Assumption 4 yields the last inequality. Using the unbiasedness identities with $u_j \sim \hat{\rho}^Y_{k+1/2}$,

$$\mathbb{E}_{u_j\sim\hat{\rho}^Y_{k+1/2}}\left[\frac{1}{N'}\sum_{j=1}^{N'}\frac{\mathcal{N}(\hat{z};u_j,hI_d)}{\hat{\rho}^Y_{k+1/2}(u_j)}\right] = 1, \quad \mathbb{E}_{u_j\sim\hat{\rho}^Y_{k+1/2}}\left[\frac{\pi^X(\hat{z}\mid z_\tau) * \mathcal{N}(0,hI_d)(u_j)}{\hat{\rho}^Y_{k+1/2}(u_j)}\right] = 1,$$

and the decomposition $\big((\hat{g}-1)+(1-G)\big)^2 \le 2(\hat{g}-1)^2 + 2(G-1)^2$, we get

$$\le \frac{2K^{-2}C_V}{\tau^2}\,\mathbb{E}_{\mu_0}\left[\int\Big((\hat{g}_{N'}^{k+\frac{1}{2}}(\hat{z}\mid X_k)-1)^2 + (G(z_\tau)-1)^2\Big)\,\pi^X(\hat{z}\mid z_\tau)\,\mathrm{d}\hat{z}\right].$$

Standard variance calculations for importance-weighted kernel mixtures yield

$$\int(\hat{g}_{N'}^{k+\frac{1}{2}}(\hat{z}\mid X_k)-1)^2\,\pi^X(\hat{z}\mid z_\tau)\,\mathrm{d}\hat{z} = \frac{1}{N'}\,\mathrm{Var}_{u_j\sim\hat{\rho}^Y_k}\Big(\frac{\mathcal{N}(\hat{z};u_j,hI_d)}{\hat{\rho}^Y_{k+1/2}(u_j)}\Big),$$
$$(G(z_\tau)-1)^2 = \frac{1}{N'}\,\mathrm{Var}_{u_j\sim\hat{\rho}^Y_{k+1/2}}\Big(\frac{\pi^X(\hat{z}\mid z_\tau) * \mathcal{N}(0,hI_d)(u_j)}{\hat{\rho}^Y_{k+1/2}(u_j)}\Big).$$

Each variance is upper bounded by the corresponding χ^2-divergence, i.e., $\mathrm{Var}_q(\frac{p}{q}) \le \chi^2(p\|q)$, yielding the claim. □

Score estimation error from finite M. Here we assume

Assumption 5. *The following fourth moments under $\hat{g}_{N'}^{k+\frac{1}{2}}(\cdot\mid z_\tau)$ are finite:*

$$\mathbb{E}\left[\exp\big(4f(\hat{z})\big)\right] < \infty, \quad \mathbb{E}\left[\|\hat{z}-z_\tau\|^4\exp\big(4f(\hat{z})\big)\right] < \infty.$$

This can be satisfied, for instance, the variance of each component in Gaussian mixture $\hat{g}_{N'}^{k+1/2}$ is sufficiently small compared to the divergence of f.

Lemma 5 (finite-M error is $O(1/M)$)**.** *Fix z_τ and $U \sim_{N'} \hat{q}_{k+\frac{1}{2}}(\cdot\mid X_k)$. Under Assumption 5, there exists a constant C_M such that*

$$\mathbb{E}_{\mu_0}\big[\|\hat{s}_{N',M}(z_\tau,h-\tau) - \hat{s}_{N'}(z_\tau,h-\tau)\|^2\,\big|\,U\big] \;\le\; \frac{C_M}{M\,\tau^2}.$$

Proof. Write the estimator (20) at time τ as a ratio of empirical means as following:

$$\hat{s}_{N',M}(z_\tau,\tau) = \frac{A_M}{\tau\,B_M}, \quad A_M \coloneqq \frac{1}{M}\sum_{m=1}^{M}(\hat{z}^{(m)}-z_\tau)\,e^{-f(\hat{z}^{(m)})}, \quad B_M \coloneqq \frac{1}{M}\sum_{m=1}^{M}e^{-f(\hat{z}^{(m)})}.$$

Let $A \coloneqq \mathbb{E}[A_M]$ and $B \coloneqq \mathbb{E}[B_M]$ under $\hat{z}^{(m)} \overset{\text{i.i.d.}}{\sim} \hat{g}_{N'}^{k+\frac{1}{2}}(\cdot \mid z_\tau)$. Then $\hat{s}_{N'}(z_\tau, h-\tau) = A/(\tau B)$ and

$$\|\hat{s}_{N',M} - \hat{s}_{N'}\|^2 = \frac{1}{\tau^2}\left\|\frac{A_M}{B_M} - \frac{A}{B}\right\|^2 = \frac{1}{\tau^2}\left\|\frac{A_M B - A B_M}{B_M B}\right\|^2.$$

Use the decomposition $A_M B - A B_M = (A_M - A)B + A(B - A B_M/B)$ to obtain the crude bound

$$\left\|\frac{A_M}{B_M} - \frac{A}{B}\right\|^2 \le \frac{2\|A_M - A\|^2}{B_M^2} + \frac{2\|A\|^2\,\|B_M - B\|^2}{B_M^2 B^2}.$$

Since $e^{-f(\hat{z})} > 0$, Jensen implies $B_M^{-2} \le \frac{1}{M}\sum_{m=1}^{M} e^{2f(\hat{z}^{(m)})}$. Hence

$$\begin{aligned}&\|\hat{s}_{N',M} - \hat{s}_{N'}\|^2 \\ &\quad\le \frac{2}{\tau^2}\left\{\Big(\frac{1}{M}\sum_{m=1}^{M} e^{2f(\hat{z}^{(m)})}\Big)\|A_M - A\|^2 + \frac{\|A\|^2}{B^2}\Big(\frac{1}{M}\sum_{m=1}^{M} e^{2f(\hat{z}^{(m)})}\Big)\|B_M - B\|^2\right\}.\end{aligned}$$

Taking conditional expectation with respect to $\hat{Z}|_U$ and applying Hölder's inequality with exponents $(2, 2)$ to each term, we have

$$\begin{aligned}&\mathbb{E}\big[\|\hat{s}_{N',M} - \hat{s}_{N'}\|^2 \,\big|\, U\big] \\ &\quad\le \frac{2}{\tau^2}\Big\{\big(\mathbb{E}[X_M^2]\big)^{\frac{1}{2}}\big(\mathbb{E}[\|A_M - A\|^4]\big)^{\frac{1}{2}} + \frac{\|A\|^2}{B^2}\big(\mathbb{E}[X_M^2]\big)^{\frac{1}{2}}\big(\mathbb{E}[(B_M - B)^4]\big)^{\frac{1}{2}}\Big\},\end{aligned}$$

where $X_M \coloneqq \frac{1}{M}\sum_{m=1}^{M} e^{2f(\hat{z}^{(m)})}$. We compute the second moment of X_M explicitly:

$$\mathbb{E}\big[X_M^2\big] = \frac{1}{M}\,\mathbb{E}\Big[e^{4f(\hat{z})}\Big] \;+\; \frac{M-1}{M}\,\big(\mathbb{E}[e^{2f(\hat{z})}]\big)^2.$$

For the fourth moments of the centered empirical means, using the standard 4th-moment expansion for i.i.d. averages, we obtain

$$\begin{aligned}&\mathbb{E}\left[\|A_M - A\|^4\right] \\ &\quad\le \frac{1}{M^3}\,\mathbb{E}\left[\Big\|\big((\hat{z} - z_\tau)e^{-f(\hat{z})} - \mathbb{E}[(\hat{z} - z_\tau)e^{-f(\hat{z})}]\big)\Big\|^4\right] + \frac{3}{M^2}\,\mathrm{Var}\left((\hat{z} - z_\tau)e^{-f(\hat{z})}\right)^2, \\ &\mathbb{E}\left[(B_M - B)^4\right] \le \frac{1}{M^3}\,\mathbb{E}\left[\big(e^{-f(\hat{z})} - \mathbb{E}[e^{-f(\hat{z})}]\big)^4\right] + \frac{3}{M^2}\,\mathrm{Var}\left(e^{-f(\hat{z})}\right)^2.\end{aligned}$$

Substituting these bounds back gives

$$\begin{aligned}\mathbb{E}\big[\|\hat{s}_{N',M} - \hat{s}_{N'}\|^2 \,\big|\, U\big] \le \frac{2}{\tau^2}\Big(\mathbb{E}[X_M^2]\Big)^{\frac{1}{2}}\Bigg\{&\left(\frac{1}{M^3}\mathbb{E}\big[\|\cdot\|^4\big] + \frac{3}{M^2}\mathrm{Var}\Big(\big(\hat{z} - z_\tau\big)e^{-f(\hat{z})}\Big)^2\right)^{\frac{1}{2}} \\ &+ \frac{\|A\|^2}{B^2}\left(\frac{1}{M^3}\mathbb{E}\Big[\big(e^{2f(\hat{z})} - \mathbb{E}[e^{2f(\hat{z})}]\big)^4\Big] + \frac{3}{M^2}\mathrm{Var}\Big(e^{2f(\hat{z})}\Big)^2\right)^{\frac{1}{2}}\Bigg\}.\end{aligned}$$

If the 4-th moments are finite under Assumption 5, $\mathbb{E}[X_M^2] = O(1)$ and each rooted bracket is $O(M^{-1})$. Therefore,

$$\mathbb{E}\big[\|\hat{s}_{N',M}(z_\tau, \tau) - \hat{s}_{N'}(z_\tau, \tau)\|^2 \,\big|\, U\big] \le \frac{C_M}{M\,\tau^2},$$

for a constant C_M depending only on the above moments. □

Combining Lemma 4 and Lemma 5, we conclude the overall score estimation error is $O(1/N' + 1/M)$ under several appropriate assumptions.

A.5 Overall Iteration Complexity

We begin by combining the KL contraction in the forward step (Lemma 1, (7)) with the approximate KL contraction in the backward step subject to estimation errors (Proposition 1).

Corollary 2 (Main one-iteration bound with split errors)**.** *Let $u \geq 1$ and suppose there exist constants C_{LSI} and $\Lambda_2^{(k)}$. Then the k-th iterate of Algorithm 1 satisfies*

$$H_{\pi^X}(\rho_{k+1}^X) \;\leq\; \frac{H_{\pi^X}(\rho_{k+1/2}^X)}{r} \;+\; c\,\big(\Lambda_1^{(k)} + 2\,\Lambda_2^{(k)}\big),$$

where $r \coloneqq (1 + h/C_{LSI})^{\,2-\frac{1}{2u^2}}$ and $c \coloneqq 2u^4 C_{LSI}\big((1 + h/C_{LSI})^{\frac{1}{2u^2}} - 1\big)$.

Applying the standard iteration argument to this yields

$$H_{\pi^X}(\rho_k^X) \;\leq\; \frac{1}{r^k} H_{\pi^X}(\rho_0^X) \;+\; c\,\sum_{j=0}^{k-1} \frac{\Lambda_1^{(j)} + 2\,\Lambda_2^{(j)}}{r^{\,k-1-j}}.$$

Corollary 3 (Iteration complexity under uniform error bound)**.** *Suppose that the per-step error satisfies*

$$\Lambda_1^{(j)} + 2\Lambda_2^{(j)} \;\leq\; \Lambda \quad \textit{for all } j \geq 0.$$

Then the iterates of Algorithm 1 obey

$$H_{\pi^X}(\rho_k^X) \;\leq\; \frac{1}{r^k} H_{\pi^X}(\rho_0^X) \;+\; \frac{cr}{r-1}\,\Lambda.$$

Consequently, to guarantee $H_{\pi^X}(\rho_k^X) \leq \varepsilon$, it is sufficient that

$$k \;\geq\; \frac{\log\big(2H_{\pi^X}(\rho_0^X)/\varepsilon\big)}{\log r} \quad \textit{and} \quad \Lambda \;\leq\; \frac{\varepsilon(r-1)}{2cr}.$$

Furthermore, the order evaluation with r and c substituted is

$$k \;=\; O\!\left(\frac{C_{LSI}}{h}\,\log\frac{H_{\pi^X}(\rho_0^X)}{\varepsilon}\right), \qquad \Lambda \;=\; O(\varepsilon/C_{LSI}),$$

where we used the asymptotic equivalences

$$\log r \;\asymp\; \tfrac{h}{C_{LSI}}, \qquad c \;\asymp\; u^2 h, \quad \textit{as } h/C_{LSI} \to 0.$$

B EXPERIMENT DETAILS

We provide parameters of each method for our experiments in Table 1 and Table 2. All experiments were conducted on an Intel Xeon CPU Max 9480 without GPU acceleration.

KL divergence estimation for the Gaussian Lasso experiment. Following the setting of Liang & Chen (2023b), we first ran the proximal sampler for 100,000 burn-in iterations, and then continued for 400,000 iterations. From this long trajectory we randomly collected 1000 samples to serve as reference particles. At each evaluation of the experiment, the KL divergence was estimated using 1000 current particles and these 1000 reference particles. For estimation we used the k-nearest-neighbor estimator (Kozachenko & Leonenko, 1987) implemented in Büth et al. (2025), with $k = 4$ as recommended by Kraskov et al. (2004). Although this estimation involves sampling error, we repeated the entire experiment with 10 different random seeds and, for each seed, computed the KL divergence based on the corresponding set of particles. We then reported the mean and variance of these estimates across the 10 runs to provide a more robust evaluation.

Table 1: Parameter setting for the experiment in Section 5.1.

Method	**Parameters**
Proximal Sampler with RGO	Initial distribution: $y_{1/2}^{(j)} \sim \mathcal{N}(0, I_d)$ Step size: $\eta = 1/135$ Number of independent chains: 100 Thinning: 10
Ours	Initial distribution: $x_0^{(i)} \sim \mathcal{N}(0, I_d)$ Step size: $h = 1/10$ Diffusion steps: $T = 10$ Noise schedule: linear interpolation between 0 and h Number of particles: $N = 100$ Number of interim samples: $M = 4000$
Ours without interaction	Same as *Ours*, except: Number of particles: $N = 1$ Number of independent chains: 100

Table 2: Parameter setting for the experiment in Section 5.2.

Method	**Parameters**
In-and-Out	Initial distribution: $x_0^{(i)} \sim \mathcal{N}(0, I_d)$ Step size: $h = 1$ Number of proposals for rejection sampling: 10000 (Particles are discarded if not accepted) Number of independent chains: 1000
Ours	Initial distribution: $x_0^{(i)} \sim \mathcal{N}(0, I_d)$ Step size: $h = 1$ Diffusion steps: $T = 10$ Noise schedule: linear interpolation from 0.01 to 1 Number of particles: $N = 1000$ Number of interim samples: $M = 300$

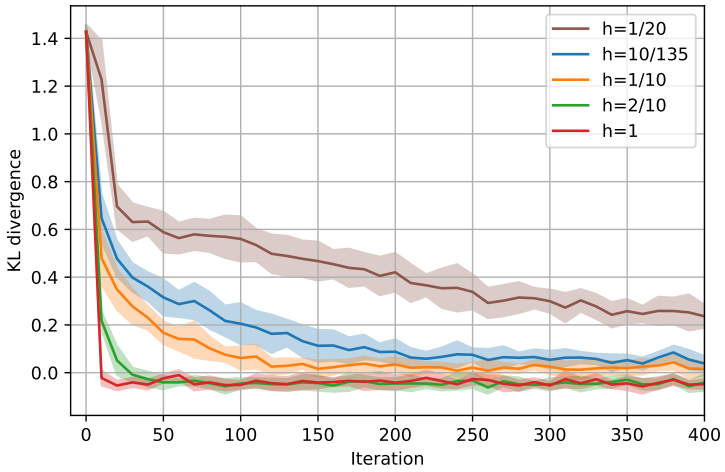

Figure 5: Convergence of KL divergence for different step sizes h, with all other parameters set as *Ours* in Table 1. Each curve shows the mean over 10 random seeds, with shaded areas indicating variances. Larger step sizes lead to faster convergence.

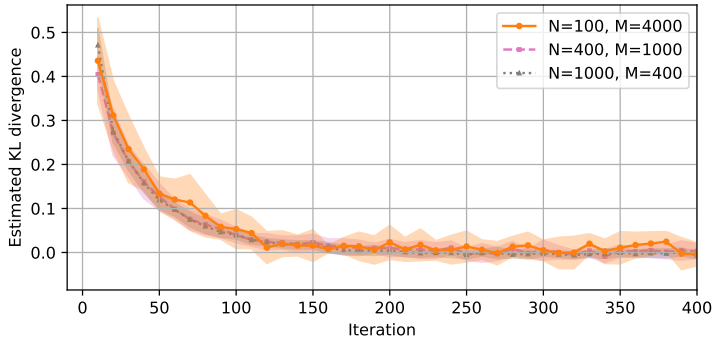

Figure 6: Comparison of hyperparameter tuning while keeping the number of particles used during the algorithm's execution (i.e., $M \times N$). For readability, the first iteration has been omitted While increasing the number of particles N used for approximating the target distribution stabilizes the KL divergence, changing both M and N results in very similar convergence patterns.

Effect of step size. In addition to the default choice $h = 1/10$, we conducted experiments with other values of h (Figure 5). We observe that larger step sizes accelerate convergence, which can be attributed to the heat flow more rapidly bridging the two modes. This phenomenon resembles diffusion-based Monte Carlo methods, where pushforward dynamics from a Gaussian initialization cover the target distribution.

Hyperparameter sensitivity under fixed computational cost. Furthermore, Figure 6 shows that when the total number of particles used during the algorithm's execution (i.e., $M \times N$) is fixed, the algorithm exhibits comparable behavior across different choices of N and M. Increasing N improves the approximation of the target distribution, thereby stabilizing the estimated KL divergence. Conversely, increasing M enhances the Monte Carlo estimation of the diffusion scores. Thus, under sufficient computational budgets, the method remains robust to these hyperparameters.