

PROVABLE BENEFIT OF ANNEALED LANGEVIN MONTE CARLO FOR NON-LOG-CONCAVE SAMPLING

Anonymous authors

Paper under double-blind review

ABSTRACT

We consider the outstanding problem of sampling from an unnormalized density that may be non-log-concave and multimodal. To enhance the performance of simple Markov chain Monte Carlo (MCMC) methods, techniques of annealing type have been widely used. However, quantitative theoretical guarantees of these techniques are under-explored. This study takes a first step toward providing a non-asymptotic analysis of annealed MCMC. Specifically, we establish, for the first time, an oracle complexity of $\tilde{O}\left(\frac{d\beta^2\mathcal{A}^2}{\varepsilon^6}\right)$ for the simple annealed Langevin Monte Carlo algorithm to achieve ε^2 accuracy in Kullback-Leibler divergence to the target distribution $\pi \propto e^{-V}$ on \mathbb{R}^d with β -smooth potential V . Here, \mathcal{A} represents the action of a curve of probability measures interpolating the target distribution π and a readily sampleable distribution.

1 INTRODUCTION

We study the task of efficient sampling from a probability distribution $\pi \propto e^{-V}$ on \mathbb{R}^d . This fundamental problem is pivotal across various fields including computational statistics (Liu, 2004; Brooks et al., 2011), Bayesian inference (Gelman et al., 2013), statistical physics (Newman & Barkema, 1999), and finance (Dagpunar, 2007), and has been extensively studied in the literature (Chewi, 2024). The most common approach to this problem is Markov Chain Monte Carlo (MCMC), among which Langevin Monte Carlo (LMC) (Durmus et al., 2019; Vempala & Wibisono, 2019; Chewi et al., 2022; Mousavi-Hosseini et al., 2023) is a particularly popular choice. LMC can be understood as a time-discretization of a diffusion process, known as Langevin diffusion (LD), whose stationary distribution is the target distribution π , and has been attractive partly due to its robust performance despite conceptual simplicity.

Although LMC and its variants converge rapidly when the target distribution π is strongly log-concave or satisfies isoperimetric inequalities such as the log-Sobolev inequality (LSI) (Durmus et al., 2019; Vempala & Wibisono, 2019; Chewi et al., 2022), its effectiveness diminishes when dealing with target distributions that are not strongly log-concave or are multimodal, such as mixtures of Gaussians. In such scenarios, the sampler often becomes confined to a single mode, severely limiting its ability to explore the entire distribution effectively. This results in significant challenges in transitioning between modes, which can dramatically increase the mixing time, making it exponential in problem parameters such as dimension, distance between modes, etc. (Dong & Tong, 2022; Ma et al., 2019). Such limitations highlight the need for enhanced MCMC methodologies that can efficiently navigate the complex landscapes of multimodal distributions, thereby improving convergence rates and overall sampling efficiency.

To address the challenges posed by multimodality, techniques around the notion of annealing have been widely employed (Gelfand et al., 1990; Neal, 2001). The general philosophy involves constructing a sequence of intermediate distributions $\pi_0, \pi_1, \dots, \pi_M$ that bridge the gap between an easily samplable distribution π_0 (e.g., Gaussian or Dirac-like), and the target distribution $\pi_M = \pi$. The process starts with sampling from π_0 and progressively samples from each subsequent distribution until π_M is reached. When π_i and π_{i+1} are close enough, approximate samples from π_i can serve as a warm start for sampling from π_{i+1} , thereby facilitating this transition. Employing LMC within this framework gives rise to what is known as the annealed LMC algorithm, which is the focus of our study. Despite its empirical success (Song & Ermon, 2019; 2020; Zilberstein et al., 2023; 2024),

Table 1: Comparison of oracle complexities in terms of d, ε , and the LSI constant for sampling from $\pi \propto e^{-V}$. “poly(\cdot)” indicates a polynomial dependence on the specified parameters.

Algorithm	Isoperimetric Assumptions	Other Assumptions	Criterion	Complexity
LMC	C -LSI	Potential smooth	ε^2 , KL ($\cdot \ \pi$)	$\tilde{O}(C^2 d \varepsilon^{-2})$
PS	C -LSI	Potential smooth	ε , TV	$\tilde{O}(C d^{1/2} \log \varepsilon^{-1})$
STLMC	/	Translated mixture of a well-conditioned distribution	ε , TV	$O(\text{poly}(d, \varepsilon^{-1}))$
RDMC	/	Potential smooth, nearly convex at ∞	ε , TV	$O(\text{poly}(d) e^{\text{poly}(\varepsilon^{-1})})$
RS-DMC	/	Potential smooth	ε^2 , KL ($\pi \ \cdot$)	$\exp(O(\log^3 d \varepsilon^{-2}))$
ZOD-MC	/	Potential growing at most quadratically	ε , TV + W_2	$\exp(\tilde{O}(d) O(\log \varepsilon^{-1}))$
ALMC (ours)	/	Potential smooth	ε^2 , KL ($\pi \ \cdot$)	$\tilde{O}(d \mathcal{A}(d)^2 \varepsilon^{-6})$

a thorough theoretical understanding of annealed LMC, particularly its non-asymptotic complexity bounds, remains elusive.

In this work, we take a first step toward developing a non-asymptotic analysis of annealed MCMC. Utilizing the Girsanov theorem to quantify the differences between the sampling dynamics and a reference process, we derive an upper bound on the error of annealed MCMC, which consists of two key terms. The first term is the ratio of an action functional in Wasserstein geometry, induced by optimal transport, to the duration of the process. This term decreases when the annealing schedule is sufficiently slow. The second term captures the discretization error inherent in practical implementations. Our approach challenges the traditional view that annealed MCMC is simply a series of warm starts. Analyses based on this perspective typically require assumptions of log-concavity or isoperimetric inequalities. In contrast, our theoretical framework for annealed MCMC dispenses with these assumptions, marking a significant shift in the understanding and analysis of annealed MCMC.

Our key technical contributions are summarized as follows.

1. We propose a novel strategy to analyze the non-asymptotic complexity bounds of annealed MCMC algorithms, bypassing the need for assumptions such as log-concavity or isoperimetry.
2. In Section 4, we investigate the annealed LD, which involves running LD with a dynamically changing target distribution. We derive a notable bound on the time required to simulate the SDE for achieving ε^2 -accuracy in KL divergence.
3. Building on the insights from the analysis of the continuous dynamic and incorporating discretization errors, we establish a non-asymptotic oracle complexity bound for annealed LMC in Section 5, which is applicable to a wide range of annealing schemes.

The quantitative results are summarized and compared to other sampling algorithms in Table 1. For algorithms requiring isoperimetric assumptions, we include Langevin Monte Carlo (LMC, Vempala & Wibisono (2019)) and Proximal Sampler (PS, Fan et al. (2023)), which converges rapidly but do not have theoretical guarantees without isoperimetric assumptions. For isoperimetry-free samplers, we include Simulated Tempering Langevin Monte Carlo (STLMC, Ge et al. (2018b)), a tempering-based sampler converging rapidly for a specific family of non-log-concave distributions, and three algorithms inspired by score-based generative models: Reverse Diffusion Monte Carlo (RDMC, Huang et al. (2024a)), Recursive Score Diffusion-based Monte Carlo (RS-DMC, Huang et al. (2024b)), and Zeroth Order Diffusion-Monte Carlo (ZOD-MC, He et al. (2024)), which involve simulating the time-reversal of Ornstein-Uhlenbeck (OU) process and require estimating the scores of the intermediate distributions. Notably, our approach operates under the least stringent assumptions and exhibits the most favorable ε -dependence among all isoperimetry-free sampling methods.

Related works and comparison. We provide a brief overview of the literature, mainly focusing on the algorithms for non-log-concave sampling and their theoretical analysis.

1. Samplers based on tempering. The fundamental concept of tempering involves sampling the system at various temperatures simultaneously: at higher temperatures, the distribution flattens, allowing particles to easily transition between modes, while at lower temperatures, particles can more effectively explore local structures. In simulated tempering (Marinari & Parisi, 1992; Woodard et al., 2009), the system’s temperature is randomly switched, while in parallel tempering (also known as replica exchange) (Swendsen & Wang, 1986; Lee & Shen, 2023), the temperatures of two particles are swapped according to a specific rule. However, quantitative theoretical results for tempering are limited, and the existing results (e.g., Ge et al. (2018a;b); Dong & Tong (2022)) apply primarily to certain special classes of non-log-concave distributions.

2. Samplers based on general diffusions. Inspired by score-based diffusion models (Ho et al., 2020; Song et al., 2021a;b), recent advances have introduced sampling methods that reverse the OU process, as detailed in Huang et al. (2024a;b); He et al. (2024). These samplers exhibit reduced sensitivity to isoperimetric conditions, but rely on estimating score functions (gradients of log-density) via importance sampling, which poses significant challenges in high-dimensional settings. Concurrently, studies such as Zhang & Chen (2022); Vargas et al. (2023; 2024); Richter & Berner (2024) have employed neural networks to approximate unknown drift terms, enabling an SDE to transport a readily sampleable distribution to the target distribution. This approach has shown excellent performance in handling complex distributions, albeit at the expense of significant computational resources required for neural network training. In contrast, annealed LMC runs on a known interpolation of probability distributions, thus simplifying sampling by obviating the need for intensive score estimation or neural network training.

3. Non-asymptotic analysis for non-log-concave sampling. Drawing upon the stationary-point analysis in non-convex optimization, the seminal work Balasubramanian et al. (2022) characterizes the convergence of non-log-concave sampling via Fisher divergence. Subsequently, Cheng et al. (2023) applies this methodology to examine the local mixing of LMC. However, Fisher divergence is a relatively weak criterion compared to more commonly employed metrics such as total-variational distance or Wasserstein distances. In contrast, our study provides a convergence guarantee in terms of KL divergence, which implies convergence in total-variation distance and offers a stronger result.

Notations and definitions. For $a, b \in \mathbb{R}$, we define $[a, b] := [a, b] \cap \mathbb{Z}$, $a \wedge b := \min(a, b)$, and $a \vee b := \max(a, b)$. For $a, b > 0$, the notations $a \lesssim b$, $a = O(b)$, and $b = \Omega(a)$ indicate that $a \leq Cb$ for some universal constant $C > 0$, and the notations $a \asymp b$ and $a = \Theta(b)$ stand for both $a = O(b)$ and $b = O(a)$. $\tilde{O}(\cdot)$ hides logarithmic dependence in $O(\cdot)$. A function $U \in C^2(\mathbb{R}^d)$ is $\alpha(> 0)$ -strongly-convex if $\nabla^2 U \succeq \alpha I$, and is $\beta(> 0)$ -smooth¹ if $-\beta I \preceq \nabla^2 U \preceq \beta I$. The total-variation (TV) distance is defined as $\text{TV}(\mu, \nu) = \sup_{A \subseteq \mathbb{R}^d} |\mu(A) - \nu(A)|$, and the Kullback-Leibler (KL) divergence is defined as $\text{KL}(\mu \parallel \nu) = \mathbb{E}_\mu \log \frac{d\mu}{d\nu}$. $\|\cdot\|$ represents the ℓ^2 norm on \mathbb{R}^d .

For $f : \mathbb{R}^d \rightarrow \mathbb{R}^{d'}$ and a probability measure μ on \mathbb{R}^d , $\|f\|_{L^2(\mu)} := \left(\int \|f\|^2 d\mu \right)^{1/2}$, and the second-order moment of μ is defined as $\mathbb{E}_\mu \|\cdot\|^2$.

2 PRELIMINARIES

2.1 STOCHASTIC DIFFERENTIAL EQUATIONS AND GIRSANOV THEOREM

A stochastic differential equation (SDE) $X = (X_t)_{t \in [0, T]}$ is a stochastic process on $\Omega = C([0, T]; \mathbb{R}^d)$, the space of continuous functions from $[0, T]$ to \mathbb{R}^d . The dynamics of X are typically represented by the equation $dX_t = b_t(X)dt + \sigma_t(X)dB_t$, $t \in [0, T]$, where $(B_t)_{t \in [0, T]}$ is a standard Brownian motion in \mathbb{R}^d , and $b_t(X) \in \mathbb{R}^d$, $\sigma_t(X) \in \mathbb{R}^{d \times d}$ depends on $(X_s)_{s \in [0, t]}$. The **path measure** of X , denoted \mathbb{P}^X , characterizes the distribution of X over Ω and is defined by $\mathbb{P}^X(A) = \mathbb{P}(X \in A)$ for all measurable subset A of Ω . The following lemma, as a corollary

¹Smoothness has several different definitions such as being C^1 or C^∞ , and the one used here is also a common one appearing in many standard textbooks in optimization and sampling, such as Chewi (2024).

of the Girsanov theorem (Üstünel & Zakai, 2013), provides a methodology for computing the KL divergence between two path measures and serves as a crucial technical tool in our proof.

Lemma 1. Assume we have the following two SDEs on Ω :

$$dX_t = a_t(X)dt + \sqrt{2}dB_t, \quad X_0 \sim \mu; \quad dY_t = b_t(Y)dt + \sqrt{2}dB_t, \quad Y_0 \sim \nu.$$

Let \mathbb{P}^X and \mathbb{P}^Y denote the path measures of X and Y , respectively. Then

$$\text{KL}(\mathbb{P}^X \parallel \mathbb{P}^Y) = \text{KL}(\mu \parallel \nu) + \frac{1}{4} \mathbb{E}_{X \sim \mathbb{P}^X} \int_0^T \|a_t(X) - b_t(X)\|^2 dt.$$

2.2 LANGEVIN DIFFUSION AND LANGEVIN MONTE CARLO

The **Langevin diffusion (LD)** with target distribution $\pi \propto e^{-V}$ is the solution to the SDE

$$dX_t = -\nabla V(X_t)dt + \sqrt{2}dB_t, \quad t \in [0, \infty); \quad X_0 \sim \mu_0. \quad (1)$$

It is well-known that under mild conditions, π is the unique stationary distribution this SDE, and when π has good regularity properties, the marginal distribution of X_t converges to π as $t \rightarrow +\infty$, so we can sample from π by simulating Equation (1) for a long time. However, in most of the cases, LD is intractable to simulate exactly, and the Euler-Maruyama discretization of Equation (1) leads to the **Langevin Monte Carlo (LMC)** algorithm. LMC with step size $h > 0$ and target distribution $\pi \propto e^{-V}$ is a Markov chain $\{X_{kh}\}_{k=0,1,\dots}$ constructed by iterating the following update rule:

$$X_{(k+1)h} = X_{kh} - h\nabla V(X_{kh}) + \sqrt{2}(B_{(k+1)h} - B_{kh}), \quad k = 0, 1, \dots; \quad X_0 \sim \mu_0, \quad (2)$$

where $\{B_{(k+1)h} - B_{kh}\}_{k=0,1,\dots} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, hI)$.

2.3 ISOPERIMETRIC INEQUALITIES

A probability measure π on \mathbb{R}^d satisfies a **log-Sobolev inequality (LSI)** with constant C , or C -LSI, if for all $f \in C^1(\mathbb{R}^d)$ with $\mathbb{E}_\pi f^2 > 0$,

$$\mathbb{E}_\pi f^2 \log \frac{f^2}{\mathbb{E}_\pi f^2} \leq 2C \mathbb{E}_\pi \|\nabla f\|^2.$$

A probability measure π on \mathbb{R}^d satisfies a **Poincaré inequality (PI)** with constant C , or C -PI, if for all $f \in C^1(\mathbb{R}^d)$,

$$\text{Var}_\pi f \leq C \mathbb{E}_\pi \|\nabla f\|^2.$$

It is worth noting that α -strongly-log-concave distributions satisfy $\frac{1}{\alpha}$ -LSI, and C -LSI implies C -PI (Bakry et al., 2014). It is established in Vempala & Wibisono (2019) that when $\pi \propto e^{-V}$ satisfies C -LSI, the LD converges exponentially fast in KL divergence; furthermore, when the potential V is β -smooth, the LMC also converges exponentially with a bias that vanishes when the step size approaches 0.

2.4 WASSERSTEIN DISTANCE AND CURVES OF PROBABILITY MEASURES

We briefly introduce several fundamental concepts in optimal transport, and direct readers to authoritative textbooks (Villani, 2008; 2021; Ambrosio et al., 2008; 2021) for an in-depth exploration.

For two probability measures μ, ν on \mathbb{R}^d with finite second-order moments, the **Wasserstein-2 (W_2) distance** between μ and ν is defined as

$$W_2(\mu, \nu) = \inf_{\gamma \in \Pi(\mu, \nu)} \left(\int \|x - y\|^2 \gamma(dx, dy) \right)^{\frac{1}{2}},$$

where $\Pi(\mu, \nu)$ is the set of all couplings of (μ, ν) , i.e., probability measure γ on $\mathbb{R}^d \times \mathbb{R}^d$ with $\gamma(A \times \mathbb{R}^d) = \mu(A)$ and $\gamma(\mathbb{R}^d \times A) = \nu(A)$, for all measurable set $A \subset \mathbb{R}^d$.

Given a vector field $v = (v_t : \mathbb{R}^d \rightarrow \mathbb{R}^d)_{t \in [a,b]}$ and a curve of probability measures $\rho = (\rho_t)_{t \in [a,b]}$ on \mathbb{R}^d with finite second-order moments, we say that v **generates** ρ if the continuity equation $\partial_t \rho_t + \nabla \cdot (\rho_t v_t) = 0$, $t \in [a, b]$ holds. The **metric derivative** of ρ at $t \in [a, b]$ is defined as

$$|\dot{\rho}|_t := \lim_{\delta \rightarrow 0} \frac{W_2(\rho_{t+\delta}, \rho_t)}{|\delta|},$$

which can be interpreted as the “speed” of this curve. If $|\dot{\rho}|_t$ exists and is finite for all $t \in [a, b]$, we say that ρ is **absolutely continuous (AC)**. **AC is a fairly weak regularity condition on the curve of probability measures. In Lemma 4, we present a sufficient condition for AC as a formal statement.**

The metric derivative and the continuity equation are closely related by the following important fact (Ambrosio et al., 2008, Theorems 8.3.1 and 8.4.5):

Lemma 2. *For an AC curve of probability measures $(\rho_t)_{t \in [a,b]}$, any vector field $(v_t)_{t \in [a,b]}$ that generates $(\rho_t)_{t \in [a,b]}$ satisfies $|\dot{\rho}|_t \leq \|v_t\|_{L^2(\rho_t)}$ for a.e. $t \in [a, b]$. Moreover, **there exists a unique vector field $(v_t^*)_{t \in [a,b]}$ generating $(\rho_t)_{t \in [a,b]}$ that satisfies $|\dot{\rho}|_t = \|v_t^*\|_{L^2(\rho_t)}$ for a.e. $t \in [a, b]$.***

Finally, we define the **action** of an AC curve of probability measures $(\rho_t)_{t \in [a,b]}$ as $\int_a^b |\dot{\rho}|_t^2 dt$. **As will be shown in the next section, the action is a key property characterizing the effectiveness of a curve in annealed sampling. The following lemma provides additional intuition about the action and may be helpful for interested readers.**

Lemma 3. *Given an AC curve of probability measures $(\rho_t)_{t \in [0,1]}$, and let \mathcal{A} be its action. Then*

1. $\mathcal{A} \geq W_2^2(\rho_0, \rho_1)$, and the equality is attained when $(\rho_t)_{t \in [0,1]}$ is a constant-speed Wasserstein geodesic, i.e., let two random variables (X_0, X_1) follow the optimal coupling of (ρ_0, ρ_1) such that $\mathbb{E}\|X_0 - X_1\|^2 = W_2^2(\rho_0, \rho_1)$, and define ρ_t as the law of $(1-t)X_0 + tX_1$.
2. If ρ_t satisfies $C_{\text{LSI}}(\rho_t)$ -LSI for all t , then $\mathcal{A} \leq \int_0^1 C_{\text{LSI}}(\rho_t)^2 \|\partial_t \nabla \log \rho_t\|_{L^2(\rho_t)}^2 dt$.
3. If ρ_t satisfies $C_{\text{PI}}(\rho_t)$ -PI for all t , then $\mathcal{A} \leq \int_0^1 2C_{\text{PI}}(\rho_t) \|\partial_t \log \rho_t\|_{L^2(\rho_t)}^2 dt$.

The proof of Lemma 3 can be found in Appendix A. The lower bound above can be derived by definition or via variational representations of action and W_2 distance, while the two upper bounds relate the action to the isoperimetric inequalities, indicating that finite action is a weaker assumption than requiring the LSI or PI constants along the trajectory. In fact, these bounds may not be tight, as demonstrated by the following surprising example, in which the action is polynomial with respect to certain problem parameters and even independent of some of them, whereas the LSI and PI constants exhibit exponential dependence. Its proof is detailed in Appendix A.

Example 1. *Consider the curve that bridges a target distribution ρ_0 to a readily sampleable distribution ρ_1 , in the form $\rho_t = \rho_0 * \mathcal{N}(0, 2StI)$, $t \in [0, 1]$, for some large positive number S . Let \mathcal{A} be the action of this curve. Then*

1. $\mathcal{A} = S \int_0^S \|\nabla \log p_s\|_{L^2(p_s)}^2 ds$, where $p_s = \rho_{s/S} = \rho_0 * \mathcal{N}(0, 2sI)$, $s \in [0, S]$.
2. In particular, when ρ_0 is a mixture of Gaussian distribution $\sum_{i=1}^N w_i \mathcal{N}(\mu_i, \sigma^2 I)$, then $\mathcal{A} \leq \frac{Sd}{2} \log(1 + \frac{2S}{\sigma^2})$ **is independent of the number of components N , the weights w_i , and the means μ_i . However, in general, the LSI and PI constants may depend exponentially on $\max_{i,j} \|\mu_i - \mu_j\|$, the maximum distance between the means.**

3 PROBLEM SETUP

Recall that the rationale behind annealing involves a gradual transition from π_0 , a simple distribution that is easy to sample from, to $\pi_1 = \pi$, the more complex target distribution. Throughout this paper, we define a curve of probability measures $(\pi_\theta)_{\theta \in [0,1]}$, along which we will apply annealing-based MCMC for sampling from π . For now, we do not specify the exact form of this curve, but instead introduce the following mild regularity assumption:

Assumption 1. Each π_θ has a finite second-order moment, and the curve $(\pi_\theta)_{\theta \in [0,1]}$ is AC with finite action $\mathcal{A} = \int_0^1 |\dot{\pi}|_\theta^2 d\theta$.

For the purpose of non-asymptotic analysis, we further introduce the following mild assumption on the target distribution $\pi \propto e^{-V}$ on \mathbb{R}^d , which is widely used in the field of sampling (see Chewi (2024) for an overview):

Assumption 2. The potential V is β -smooth, and there exists a global minimizer x_* of V such that $\|x_*\| \leq R$. Moreover, π has finite second-order moment.

With this foundational setup, we now proceed to introduce the annealed LD and annealed LMC algorithms. Our goal is to characterize the non-asymptotic complexity of annealed MCMC algorithms for sampling from possibly non-log-concave distributions.

4 ANALYSIS OF ANNEALED LANGEVIN DIFFUSION

To elucidate the concept of annealing more clearly, we first consider the **annealed Langevin diffusion (ALD)** algorithm, which samples from the $\pi \propto e^{-V}$ by running LD with a dynamically changing target distribution. For the sake of this discussion, we assume the following: (i) we can exactly sample from π_0 , (ii) the scores $(\nabla \log \pi_\theta)_{\theta \in [0,1]}$ are known in closed form, and (iii) we can exactly simulate any SDE with known drift and diffusion terms.

Fix a sufficiently long time T . We define a reparametrized curve of probability measures $(\tilde{\pi}_t := \pi_{t/T})_{t \in [0,T]}$. Starting with an initial sample $X_0 \sim \pi_0 = \tilde{\pi}_0$, we run the SDE

$$dX_t = \nabla \log \tilde{\pi}_t(X_t)dt + \sqrt{2}dB_t, \quad t \in [0, T], \quad (3)$$

and ultimately output $X_T \sim \nu^{\text{ALD}}$ as an approximate sample from the target distribution π . Intuitively, when $\tilde{\pi}_t$ is changing slowly, the distribution of X_t should closely resemble $\tilde{\pi}_t$, leading to an output distribution ν^{ALD} that approximates the target distribution. This turns out to be true, as is confirmed by the following theorem, which provides a convergence guarantee for the ALD process.

Theorem 1. When choosing $T = \frac{\mathcal{A}}{4\varepsilon^2}$, it follows that $\text{KL}(\pi \| \nu^{\text{ALD}}) \leq \varepsilon^2$.

Proof. Let \mathbb{Q} be the path measure of ALD (Equation (3)) initialized at $X_0 \sim \tilde{\pi}_0$, and define \mathbb{P} as the path measure corresponding to the following reference SDE:

$$dX_t = (\nabla \log \tilde{\pi}_t + v_t)(X_t)dt + \sqrt{2}dB_t, \quad X_0 \sim \tilde{\pi}_0, \quad t \in [0, T]. \quad (4)$$

The vector field $v = (v_t)_{t \in [0,T]}$ is designed such that $X_t \sim \tilde{\pi}_t$ for all $t \in [0, T]$. According to the Fokker-Planck equation², this is equivalent to the following PDE:

$$\partial_t \tilde{\pi}_t = -\nabla \cdot (\tilde{\pi}_t (\nabla \log \tilde{\pi}_t + v_t)) + \Delta \tilde{\pi}_t = -\nabla \cdot (\tilde{\pi}_t v_t), \quad t \in [0, T],$$

which means that v generates $(\tilde{\pi}_t := \pi_{t/T})_{t \in [0,T]}$. We can compute $\text{KL}(\mathbb{P} \| \mathbb{Q})$ using Lemma 1:

$$\text{KL}(\mathbb{P} \| \mathbb{Q}) = \frac{1}{4} \mathbb{E}_{\mathbb{P}} \int_0^T \|v_t(X_t)\|^2 dt = \frac{1}{4} \int_0^T \|v_t\|_{L^2(\tilde{\pi}_t)}^2 dt.$$

Leveraging Lemma 2, among all vector fields v that generate $(\tilde{\pi}_t := \pi_{t/T})_{t \in [0,T]}$, we can choose the one that minimizes $\|v_t\|_{L^2(\tilde{\pi}_t)}$, thereby making $\|v_t\|_{L^2(\tilde{\pi}_t)} = \left| \dot{\tilde{\pi}} \right|_t$, the metric derivative. With the reparameterization $\tilde{\pi}_t = \pi_{t/T}$, we have the following relation by chain rule:

$$\left| \dot{\tilde{\pi}} \right|_t = \lim_{\delta \rightarrow 0} \frac{W_2(\tilde{\pi}_{t+\delta}, \tilde{\pi}_t)}{|\delta|} = \lim_{\delta \rightarrow 0} \frac{W_2(\pi_{(t+\delta)/T}, \pi_{t/T})}{T|\delta/T|} = \frac{1}{T} |\dot{\pi}|_{t/T}.$$

Employing the change-of-variable formula leads to

$$\text{KL}(\mathbb{P} \| \mathbb{Q}) = \frac{1}{4} \int_0^T \left| \dot{\tilde{\pi}} \right|_t^2 dt = \frac{1}{4T} \int_0^1 |\dot{\pi}|_\theta^2 d\theta = \frac{\mathcal{A}}{4T}.$$

²Here, we assume that the solution of the Fokker-Planck equation exists and is unique, which holds under weak regularity conditions.

Finally, using data-processing inequality (see, e.g., (Chewi, 2024, Theorem 1.5.3)), with $T = \frac{A}{4\varepsilon^2}$,

$$\text{KL}(\pi \parallel \nu^{\text{ALD}}) = \text{KL}(\mathbb{P}_T \parallel \mathbb{Q}_T) \leq \text{KL}(\mathbb{P} \parallel \mathbb{Q}) = \varepsilon^2,$$

where \mathbb{P}_T and \mathbb{Q}_T stand for the marginal distributions of \mathbb{P} and \mathbb{Q} at time T , respectively. \square

Let us delve deeper into the mechanics of the ALD. Although at time t the SDE (Equation (3)) targets the distribution $\tilde{\pi}_t$, the distribution of X_t does not precisely align with $\tilde{\pi}_t$. Nevertheless, by choosing a sufficiently long time T , we actually move on the curve $(\pi_\theta)_{\theta \in [0,1]}$ sufficiently slowly, thus minimizing the discrepancy between the path measure of $(X_t)_{t \in [0,T]}$ and the reference curve $(\tilde{\pi}_t)_{t \in [0,T]}$. Using the data-processing inequality, we can upper bound the error between the marginal distributions at time T by the joint distributions of the two paths. In essence, *moving more slowly, sampling more precisely*.

Our analysis primarily addresses the *global* error across the entire curve of probability measures, rather than focusing solely on the *local* error at time T . This approach is inspired by Dalalyan & Tsybakov (2012); Chen et al. (2023), and stands in contrast to the isoperimetry-based analyses of LD (e.g., Vempala & Wibisono (2019); Chewi et al. (2022); Balasubramanian et al. (2022)), which focus on the decay of the KL divergence from the distribution of X_t to the target distribution, and require LSI to bound the time derivative of this quantity. Notably, the total time T needed to run the SDE depends solely on the action of the curve $(\pi_\theta)_{\theta \in [0,1]}$, obviating the need for assumptions related to log-concavity or isoperimetry.

It is also worth noting that ALD plays a critical role in the field of non-equilibrium stochastic thermodynamics (Seifert, 2012). Recently, a refinement of the fluctuation theorem was discovered in Chen et al. (2020); Fu et al. (2021), showing that the irreversible entropy production in a stochastic thermodynamic system is equal to the ratio of a similar action integral and the duration of the process T , closely resembling Theorem 1.

5 ANALYSIS OF ANNEALED LANGEVIN MONTE CARLO

It is crucial to recognize that, in practice, running the algorithm requires knowledge of the score functions $(\nabla \log \pi_\theta)_{\theta \in [0,1]}$ in closed form. Additionally, simulating ALD (Equation (3)) necessitates discretization schemes, which introduces further errors. This section presents a detailed non-asymptotic convergence analysis for the **annealed Langevin Monte Carlo (ALMC)** algorithm, which is a practical approach for real-world implementations.

The outline of this section is as follows. We will first introduce a family of the interpolation curve $(\pi_\theta)_{\theta \in [0,1]}$ that serves as the basis for our analysis. Next, we propose a discretization scheme tailored to the special structure of $(\pi_\theta)_{\theta \in [0,1]}$. Finally, we state the complexity analysis in Theorem 2, and provide an example demonstrating the improvement of our bounds over existing results.

5.1 CHOICE OF THE INTERPOLATION CURVE

We consider the following curve of probability measures on \mathbb{R}^d :

$$\pi_\theta \propto \exp \left(-\eta(\theta)V - \frac{\lambda(\theta)}{2} \|\cdot\|^2 \right), \quad \theta \in [0,1], \quad (5)$$

where the functions $\eta(\cdot)$ and $\lambda(\cdot)$ are called the **annealing schedule**. These functions must be differentiable and monotonic, satisfying the boundary conditions $\eta_0 = \eta(0) \nearrow \eta(1) = 1$ and $\lambda_0 = \lambda(0) \searrow \lambda(1) = 0$. The values of $\eta_0 \in [0,1]$ and $\lambda_0 \in (1, +\infty)$ are chosen so that π_0 corresponds a Gaussian distribution $\mathcal{N}\left(0, \frac{1}{\lambda_0}I\right)$ or can be efficiently sampled within $\tilde{O}(1)$ steps of rejection sampling (see Lemma 5 for details). We also remark that, under Assumption 2, π_θ has a finite second-order moment (see Lemma 8 the proof), ensuring the W_2 distance between π_θ 's is well-defined.

This flexible interpolation scheme includes many general cases. For example, Brosse et al. (2018) and Ge et al. (2020) used the schedule $\eta(\cdot) \equiv 1$, while Neal (2001) used the schedule $\lambda(\theta) = c(1 - \eta(\theta))$. The key motivation for this interpolation is that when $\theta = 0$, the quadratic term

predominates, making the potential of π_0 strongly-convex with a moderate condition number, thus π_0 is easy to sample from; on the other hand, when $\theta = 1$, π_1 is just the target distribution π . As θ gradually increases from 0 to 1, the readily sampleable distribution π_0 slowly transforms into the target distribution π_1 .

5.2 THE DISCRETIZATION ALGORITHM

With the interpolation curve $(\pi_\theta)_{\theta \in [0,1]}$ chosen as above, a straightforward yet non-optimal method to discretize Equation (3) involves employing the Euler-Maruyama scheme, i.e.,

$$X_{t+\Delta t} \approx X_t + \nabla \log \tilde{\pi}_t(X_t) \Delta t + \sqrt{2}(B_{t+\Delta t} - B_t), \quad 0 \leq t < t + \Delta t \leq T.$$

However, considering that $\nabla \log \tilde{\pi}_t(x) = -\eta\left(\frac{t}{T}\right) \nabla V(x) - \lambda\left(\frac{t}{T}\right) x$, the integral of the linear term can be computed in closed form, so we can use the exponential-integrator scheme (Zhang & Chen, 2023; Zhang et al., 2023) to further reduce the discretization error. Given the total time T , we define a sequence of points $0 = \theta_0 < \theta_1 < \dots < \theta_M = 1$, and set $T_\ell = T\theta_\ell$, $h_\ell = T(\theta_\ell - \theta_{\ell-1})$. The exponential-integrator scheme is then expressed as

$$dX_t = \left(-\eta\left(\frac{t}{T}\right) \nabla V(X_{t-}) - \lambda\left(\frac{t}{T}\right) X_t \right) dt + \sqrt{2}dB_t, \quad t \in [0, T], \quad X_0 \sim \pi_0, \quad (6)$$

where $t_- := T_{\ell-1}$ when $t \in [T_{\ell-1}, T_\ell)$, $\ell \in \llbracket 1, M \rrbracket$. The explicit update rule is detailed in Algorithm 1, with x_ℓ denoting X_{T_ℓ} , and the derivation of Equation (7) is presented in Appendix C.1. Notably, replacing $\nabla V(X_{t-})$ with $\nabla V(X_t)$ recovers the ALD (Equation (3)), and setting $\eta \equiv 1$ and $\lambda \equiv 0$ reduces to the classical LMC iterations.

Algorithm 1: Annealed LMC Algorithm

Input: Target distribution $\pi \propto e^{-V}$, total time T , annealing schedule $\eta(\cdot)$ and $\lambda(\cdot)$, discrete points $\theta_0, \dots, \theta_M$.

- 1 For $0 \leq \theta < \theta' \leq 1$, define $\Lambda_0(\theta', \theta) = \exp\left(-T \int_\theta^{\theta'} \lambda(u) du\right)$,
 $H(\theta', \theta) = T \int_\theta^{\theta'} \eta(u) \Lambda_0(u, \theta') du$, and $\Lambda_1(\theta', \theta) = \sqrt{2T \int_\theta^{\theta'} \Lambda_0^2(u, \theta') du}$;
 - 2 Obtain a sample $x_0 \sim \pi_0$ (e.g., using rejection sampling (Algorithm 2));
 - 3 **for** $\ell = 1, 2, \dots, M$ **do**
 - 4 Sample an independent Gaussian noise $\xi_\ell \sim \mathcal{N}(0, I)$;
 - 5 Update

$$x_\ell = \Lambda_0(\theta_\ell, \theta_{\ell-1})x_{\ell-1} - H(\theta_\ell, \theta_{\ell-1})\nabla V(x_{\ell-1}) + \Lambda_1(\theta_\ell, \theta_{\ell-1})\xi_\ell. \quad (7)$$
 - 6 **end**
- Output:** $x_M \sim \nu^{\text{ALMC}}$, an approximate sample from π .
-

We illustrate the ALMC algorithm in Figure 1. The underlying intuition behind this algorithm is that by setting a sufficiently long total time T , the trajectory of the continuous dynamic (i.e., annealed LD) approaches the reference trajectory closely, as established in Theorem 1. Additionally, by adopting sufficiently small step sizes h_1, \dots, h_M , the discretization error can be substantially reduced. Unlike traditional annealed LMC methods, which require multiple LMC update steps for each intermediate distribution π_1, \dots, π_M , our approach views the annealed LMC as a discretization of a continuous-time process. Consequently, it is adequate to perform *a single step* of LMC towards each intermediate distribution $\tilde{\pi}_{T_1}, \dots, \tilde{\pi}_{T_M}$, provided that T is sufficiently large and the step sizes are appropriately small.

5.3 CONVERGENCE ANALYSIS OF THE ALGORITHM

The subsequent theorem provides a convergence guarantee for the annealed LMC algorithm, with a detailed proof available in Appendix C.2.

Theorem 2. *Under Assumptions 1 and 2, Algorithm 1 can generate a distribution ν^{ALMC} satisfying $\text{KL}(\pi \| \nu^{\text{ALMC}}) \leq \varepsilon^2$ within*

$$\tilde{O}\left(\frac{d\beta^2\mathcal{A}^2}{\varepsilon^6}\right)$$




Figure 1: Illustration of the ALMC algorithm. The ℓ -th green arrow, proceeding from left to right, represents one step of LMC towards $\tilde{\pi}_{T_\ell}$ with step size h_ℓ , while each red arrow corresponds to the application of the same transition kernel, initialized at $\tilde{\pi}_{T_{\ell-1}}$ on the reference trajectory \mathbb{P} , which is depicted in purple. To evaluate $\text{KL}(\mathbb{P} \parallel \mathbb{Q})$, the Girsanov theorem implies that we only need to bound the aggregate KL divergence across each small interval (i.e., the sum of the blue “distances”).

calls to the oracle of V and ∇V in expectation.

Sketch of Proof. Let \mathbb{Q} be the path measure of ALMC (Equation (3)), the time-discretized sampling process, whose marginal distribution at time T is the output distribution ν^{ALMC} . Again, let \mathbb{P} denote the reference path measure of Equation (4) used in the proof of Theorem 1, in which the same vector field $(v_t)_{t \in [0, T]}$ ensures that $X_t \sim \tilde{\pi}_t$ under \mathbb{P} for all $t \in [0, T]$. Applying Girsanov theorem (Lemma 1) and carefully dealing with the discretization error, we can upper bound $\text{KL}(\mathbb{P} \parallel \mathbb{Q})$ by

$$\text{KL}(\mathbb{P} \parallel \mathbb{Q}) \lesssim \sum_{\ell=1}^M \left(\frac{1 + \eta(\theta_\ell)^2 \beta^2 h_\ell^2}{T} \int_{\theta_{\ell-1}}^{\theta_\ell} |\dot{\pi}|_\theta^2 d\theta + \eta(\theta_\ell)^2 \beta^2 d h_\ell^2 (1 + h_\ell (\beta \eta(\theta_\ell) + \lambda(\theta_{\ell-1}))) \right).$$

The first summation is governed by the total time T , which pertains to the convergence of the continuous dynamic (i.e., ALD). Setting $T \asymp \frac{A}{\varepsilon^2}$ ensures that the first summation remains $O(\varepsilon^2)$, provided that the step size h_ℓ is sufficiently small. The second summation addresses the discretization error, and it suffices to determine the appropriate value of θ_ℓ to minimize M , the total number of calls to the oracle of ∇V for discretizing the SDE. Combining M with the complexity of sampling from π_0 determines the overall complexity of the algorithm. \square

Once again, our analysis relies on bounding the global error between two path measures by Girsanov theorem. The annealing schedule $\eta(\cdot)$ and $\lambda(\cdot)$ influences the complexity exclusively through the action \mathcal{A} . This crucial identity, based on the choice of interpolation curve $(\pi_\theta)_{\theta \in [0, 1]}$, significantly affects the effectiveness of ALMC. Our specific choice (Equation (5)), among all interpolation curves with closed-form scores, aims at simplifying the discretization error analysis. However, alternative annealing schemes could be explored to find the optimal one in practice. For instance, a recent paper (Fan et al., 2024) proposed an annealing scheme $\pi_\theta(x) \propto \pi_0((1 - a\theta)x)^{1-\theta} \pi_1\left(\frac{x}{b+(1-b)\theta}\right)^\theta$ for some $a \in [0, 1]$ and $b \in (0, 1]$. The extra scaling factor $\frac{1}{b+(1-b)\theta}$ might improve mixing speed. We leave the choice of the interpolation curve as a topic for future research.

We also note that our Assumption 2 encompasses strongly-log-concave target distributions. For sampling from these well-conditioned distributions via LMC, the complexity required to achieve ε -accuracy in TV distance scales as $\tilde{O}(\varepsilon^{-2})$ (Chewi, 2024). However, using Pinsker inequality $\text{KL} \geq 2\text{TV}^2$, our complexity to meet the same error criterion is $O(\varepsilon^{-6})$, indicating a significantly higher computational demand. While our discretization error is as sharp as existing works, the main reason for our worse ε -dependence is due to the time required for the continuous dynamics to converge, which is based on Girsanov theorem rather than LSI. While sharpening the ε -dependence based on Girsanov theorem remains a challenge, our $O(\varepsilon^{-6})$ dependence still outperforms all the other bounds in Table 1 without isoperimetry assumptions.

Finally, we conclude by demonstrating an example of a class of mixture of Gaussian distributions, which illustrates how our analysis can improve the complexity bound from exponential to polynomial. The detailed proof is provided in Appendix D.

Example 2. Consider a mixture of Gaussian distributions in \mathbb{R}^d defined by $\pi = \sum_{i=1}^N p_i \mathcal{N}(y_i, \beta^{-1}I)$, where the weights $p_i > 0$, $\sum_{i=1}^N p_i = 1$, and $\|y_i\| = r$ for all $i \in [1, N]$.

Consequently, the potential $V = -\log \pi$ is B -smooth, where $B = \beta(4r^2\beta + 1)$. With an annealing schedule defined by $\eta(\cdot) \equiv 1$ and $\lambda(\theta) = dB(1 - \theta)^\gamma$ for some $1 \leq \gamma = O(1)$, it follows that $\mathcal{A} = O\left(d(r^2\beta + 1)\left(r^2 + \frac{d}{\beta}\right)\right)$.

To show the superiority of our theoretical results, consider a simplified scenario with $N = 2$, $y_1 = -y_2$, and $r^2 \gg \frac{1}{\beta}$. By applying Example 2 and the Pinsker inequality, the total complexity required to obtain an ε -accurate sample in TV distance is $\tilde{O}\left(d^3\beta^2r^4(r^4\beta^2 \vee d^2)\varepsilon^{-6}\right)$. In contrast, studies such as Schlichting (2019); Chen et al. (2021); Dong & Tong (2022) indicate that the LSI constant of π is $\Omega\left(e^{\Theta(\beta r^2)}\right)$, so existing bounds in Table 1 suggest that LMC, to achieve the same accuracy, can only provide an exponential complexity guarantee of $\tilde{O}\left(e^{\Theta(\beta r^2)}d\varepsilon^{-2}\right)$.

6 EXPERIMENTS

We conduct simple numerical experiments to verify the findings in Example 2, which demonstrate that for a certain class of mixture of Gaussian distributions, ALMC achieves polynomial convergence with respect to r . Specifically, we consider a target distribution comprising a mixture of 6 Gaussians with uniform weights $\frac{1}{6}$, means $\left\{ \left(r \cos \frac{k\pi}{3}, r \sin \frac{k\pi}{3} \right) : k \in \llbracket 0, 5 \rrbracket \right\}$, and the same covariance $0.1I$ (corresponding to $\beta = 10$ in Example 2). We experimented $r \in \{2, 5, 10, 15, 20, 25, 30\}$, and compute the number of iterations required for the empirical KL divergence from the target distribution to the sampled distribution to fall below 0.2 (blue curve) and 0.1 (orange curve). The results, displayed in Figure 2, use a log-log scale for both axes. The near-linear behavior of the curves in this plot confirms that the complexity depends polynomially on r , validating our theory. Further details of the experimental setup and implementation can be found in Appendix F.




Figure 2: Relationship between norm of means r and the number of iterations to reach 0.2 and 0.1 accuracy in KL divergence, both in log scale.

7 CONCLUSIONS AND FUTURE WORK

In this paper, we have explored the complexity of ALMC for sampling from a non-log-concave probability measure, circumventing the reliance on log-concavity or isoperimetric inequalities. Central to our analysis are the Girsanov theorem and optimal transport techniques, providing a novel approach. While our proof primarily focuses on the annealing scheme as described in Equation (5), it can potentially be adapted to more general interpolations. Further exploration of these applications will be a key direction for future research. Technically, our proof methodology could be expanded to a broader range of target distributions beyond those with smooth potentials, such as those with Hölder-continuous gradients (Chatterji et al., 2020; Liang & Chen, 2022; Fan et al., 2023). Eliminating the assumption of global minimizers of the potential function would further enhance the practical applicability of our algorithm. Finally, while this work emphasizes the upper bounds for non-log-concave sampling, exploring the tightness of these bounds and investigating potential lower bounds for this problem (Chewi et al., 2023a;b) are intriguing avenues for future research.

REFERENCES

- Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient Flows: In Metric Spaces and in the Space of Probability Measures*. Lectures in Mathematics. ETH Zürich. Birkhäuser Basel, 2 edition, 2008. doi: 10.1007/978-3-7643-8722-8.
- Luigi Ambrosio, Elia Brué, and Daniele Semola. *Lectures on optimal transport*, volume 130 of *UNITEXT*. Springer Cham, 2021. doi: 10.1007/978-3-030-72162-6. URL <https://link.springer.com/book/10.1007/978-3-030-72162-6>.
- Dominique Bakry, Ivan Gentil, and Michel Ledoux. *Analysis and geometry of Markov diffusion operators*, volume 103 of *Grundlehren der mathematischen Wissenschaften*. Springer Cham, 1 edition, 2014. doi: 10.1007/978-3-319-00227-9.
- Krishna Balasubramanian, Sinho Chewi, Murat A Erdogdu, Adil Salim, and Shunshi Zhang. Towards a theory of non-log-concave sampling: First-order stationarity guarantees for Langevin Monte Carlo. 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. 2896–2923. PMLR, 02–05 Jul 2022. URL <https://proceedings.mlr.press/v178/balasubramanian22a.html>.
- Erhan Bayraktar and Gaoyue Guo. Strong equivalence between metrics of Wasserstein type. *Electronic Communications in Probability*, 26(none):1 – 13, 2021. doi: 10.1214/21-ECP383. URL <https://doi.org/10.1214/21-ECP383>.
- Steve Brooks, Andrew Gelman, Galin Jones, and Xiao-Li Meng. *Handbook of Markov chain Monte Carlo*. CRC press, 2011.
- Nicolas Brosse, Alain Durmus, and Éric Moulines. Normalizing constants of log-concave densities. *Electronic Journal of Statistics*, 12(1):851 – 889, 2018. doi: 10.1214/18-EJS1411. URL <https://doi.org/10.1214/18-EJS1411>.
- Niladri Chatterji, Jelena Diakonikolas, Michael I. Jordan, and Peter Bartlett. Langevin Monte Carlo without smoothness. In Silvia Chiappa and Roberto Calandra (eds.), *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 1716–1726. PMLR, 26–28 Aug 2020. URL <https://proceedings.mlr.press/v108/chatterji20a.html>.
- Hong-Bin Chen, Sinho Chewi, and Jonathan Niles-Weed. Dimension-free log-Sobolev inequalities for mixture distributions. *Journal of Functional Analysis*, 281(11):109236, 2021. ISSN 0022-1236. doi: <https://doi.org/10.1016/j.jfa.2021.109236>. URL <https://www.sciencedirect.com/science/article/pii/S0022123621003189>.
- Sitan Chen, Sinho Chewi, Jerry Li, Yuanzhi Li, Adil Salim, and Anru Zhang. Sampling is as easy as learning the score: theory for diffusion models with minimal data assumptions. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=zyLVMgsZ0U_.
- Yongxin Chen, Tryphon T. Georgiou, and Allen Tannenbaum. Stochastic control and nonequilibrium thermodynamics: Fundamental limits. *IEEE Transactions on Automatic Control*, 65(7):2979–2991, 2020. doi: 10.1109/TAC.2019.2939625.
- Xiang Cheng, Bohan Wang, Jingzhao Zhang, and Yusong Zhu. Fast conditional mixing of MCMC algorithms for non-log-concave distributions. In A. Oh, T. Neumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 13374–13394. Curran Associates, Inc., 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/2b00b3331bd0f5fbfdd966ac06338f6d-Paper-Conference.pdf.
- Sinho Chewi. *Log-Concave Sampling*. Book draft, in preparation, 2024. URL <https://chewisinho.github.io>.

- Sinho Chewi, Murat A Erdogdu, Mufan Li, Ruoqi Shen, and Shunshi Zhang. Analysis of Langevin Monte Carlo from Poincaré to log-Sobolev. 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. 1–2. PMLR, 02–05 Jul 2022. URL <https://proceedings.mlr.press/v178/chewi22a.html>.
- Sinho Chewi, Jaume De Dios Pont, Jerry Li, Chen Lu, and Shyam Narayanan. Query lower bounds for log-concave sampling. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 2139–2148, 2023a. doi: 10.1109/FOCS57990.2023.00131.
- Sinho Chewi, Patrik Gerber, Holden Lee, and Chen Lu. Fisher information lower bounds for sampling. In Shipra Agrawal and Francesco Orabona (eds.), *Proceedings of The 34th International Conference on Algorithmic Learning Theory*, volume 201 of *Proceedings of Machine Learning Research*, pp. 375–410. PMLR, 20 Feb–23 Feb 2023b. URL <https://proceedings.mlr.press/v201/chewi23b.html>.
- John S. Dagpunar. *Simulation and Monte Carlo: With Applications in Finance and MCMC*. John Wiley & Sons, Ltd, 2007. doi: 10.1002/9780470061336.
- Arnak Dalalyan and Alexandre B. Tsybakov. Sparse regression learning by aggregation and Langevin Monte-Carlo. *Journal of Computer and System Sciences*, 78(5):1423–1443, 2012. ISSN 0022-0000. doi: <https://doi.org/10.1016/j.jcss.2011.12.023>. URL <https://www.sciencedirect.com/science/article/pii/S0022000012000220>. JCSS Special Issue: Cloud Computing 2011.
- Jing Dong and Xin T. Tong. Spectral gap of replica exchange langevin diffusion on mixture distributions. *Stochastic Processes and their Applications*, 151:451–489, 2022. ISSN 0304-4149. doi: <https://doi.org/10.1016/j.spa.2022.06.006>. URL <https://www.sciencedirect.com/science/article/pii/S0304414922001314>.
- Alain Durmus, Szymon Majewski, and Błażej Miasojedow. Analysis of langevin monte carlo via convex optimization. *Journal of Machine Learning Research*, 20(73):1–46, 2019. URL <http://jmlr.org/papers/v20/18-173.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 *Forty-first International Conference on Machine Learning*, 2024. URL <https://openreview.net/forum?id=Kt4fwuKqf>.
- Rui Fu, Amirhossein Taghvaei, Yongxin Chen, and Tryphon T Georgiou. Maximal power output of a stochastic thermodynamic engine. *Automatica*, 123:109366, 2021.
- Guillaume Garrigos and Robert M Gower. Handbook of convergence theorems for (stochastic) gradient methods. *arXiv preprint arXiv:2301.11235*, 2023.
- Rong Ge, Holden Lee, and Andrej Risteski. Beyond log-concavity: Provable guarantees for sampling multi-modal distributions using simulated tempering Langevin Monte Carlo. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018a. URL https://proceedings.neurips.cc/paper_files/paper/2018/file/c6ede20e6f597abf4b3f6bb30cee16c7-Paper.pdf.
- Rong Ge, Holden Lee, and Andrej Risteski. Simulated tempering Langevin Monte Carlo II: An improved proof using soft Markov chain decomposition. *arXiv preprint arXiv:1812.00793*, 2018b.
- Rong Ge, Holden Lee, and Jianfeng Lu. Estimating normalizing constants for log-concave distributions: Algorithms and lower bounds. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2020, pp. 579–586, New York, NY, USA, 2020.

- Association for Computing Machinery. ISBN 9781450369794. doi: 10.1145/3357713.3384289. URL <https://doi.org/10.1145/3357713.3384289>.
- Saul Brian Gelfand, Sanjoy K Mitter, et al. On sampling methods and annealing algorithms. *technical report*, 1990.
- Andrew Gelman, John B. Carlin, Hal S. Stern, and Donald B. Rubin. *Bayesian data analysis*. Chapman and Hall/CRC, 3 edition, 2013.
- Ye He, Kevin Rojas, and Molei Tao. Zeroth-order sampling methods for non-log-concave distributions: Alleviating metastability by denoising diffusion. *arXiv preprint arXiv:2402.17886*, 2024.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.
- 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 <http://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>.
- Holden Lee and Zeyu Shen. Improved bound for mixing time of parallel tempering. *arXiv preprint arXiv:2304.01303*, 2023.
- Jiaming Liang and Yongxin Chen. A proximal algorithm for sampling from non-smooth potentials. In *2022 Winter Simulation Conference (WSC)*, pp. 3229–3240, 2022. doi: 10.1109/WSC57314.2022.10015293.
- Jiaming Liang and Yongxin Chen. A proximal algorithm for sampling. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL <https://openreview.net/forum?id=CkXOwlhf27>.
- Jun S. Liu. *Monte Carlo strategies in scientific computing*. Springer Series in Statistics. Springer New York, NY, 2004. doi: 10.1007/978-0-387-76371-2.
- Yuan Liu. The Poincaré inequality and quadratic transportation-variance inequalities. *Electronic Journal of Probability*, 25(none):1 – 16, 2020. doi: 10.1214/19-EJP403. URL <https://doi.org/10.1214/19-EJP403>.
- Yi-An Ma, Yuansi Chen, Chi Jin, Nicolas Flammarion, and Michael I. Jordan. Sampling can be faster than optimization. *Proceedings of the National Academy of Sciences*, 116(42):20881–20885, 2019. doi: 10.1073/pnas.1820003116. URL <https://www.pnas.org/doi/abs/10.1073/pnas.1820003116>.
- Enzo Marinari and Giorgio Parisi. Simulated tempering: A new Monte Carlo scheme. *Europhysics Letters*, 19(6):451, jul 1992. doi: 10.1209/0295-5075/19/6/002. URL <https://dx.doi.org/10.1209/0295-5075/19/6/002>.
- Alireza Mousavi-Hosseini, Tyler K. Farghly, Ye He, Krishna Balasubramanian, and Murat A. Erdogdu. Towards a complete analysis of Langevin Monte Carlo: Beyond Poincaré inequality. In Gergely Neu and Lorenzo Rosasco (eds.), *Proceedings of Thirty Sixth Conference on Learning Theory*, volume 195 of *Proceedings of Machine Learning Research*, pp. 1–35. PMLR, 12–15 Jul 2023. URL <https://proceedings.mlr.press/v195/mousavi-hosseini23a.html>.
- Radford M. Neal. Annealed importance sampling. *Statistics and Computing*, 11(2):125–139, April 2001. ISSN 1573-1375. doi: 10.1023/A:1008923215028. URL <https://doi.org/10.1023/A:1008923215028>.

- Mark E. J. Newman and Gerard T. Barkema. *Monte Carlo Methods in Statistical Physics*. Oxford University Press, 02 1999. ISBN 9780198517962. doi: 10.1093/oso/9780198517962.001.0001. URL <https://doi.org/10.1093/oso/9780198517962.001.0001>.
- 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>.
- André Schlichting. Poincaré and log–Sobolev inequalities for mixtures. *Entropy*, 21(1), 2019. ISSN 1099-4300. doi: 10.3390/e21010089. URL <https://www.mdpi.com/1099-4300/21/1/89>.
- Udo Seifert. Stochastic thermodynamics, fluctuation theorems and molecular machines. *Reports on progress in physics*, 75(12):126001, 2012.
- Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021a. URL <https://openreview.net/forum?id=StlgiaRCHLP>.
- Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. 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/3001ef257407d5a371a96dcd947c7d93-Paper.pdf.
- Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 12438–12448. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/92c3b916311a5517d9290576e3ea37ad-Paper.pdf.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021b. URL <https://openreview.net/forum?id=PxtIG12RRHS>.
- Robert H. Swendsen and Jian-Sheng Wang. Replica Monte Carlo simulation of spin-glasses. *Phys. Rev. Lett.*, 57:2607–2609, Nov 1986. doi: 10.1103/PhysRevLett.57.2607. URL <https://link.aps.org/doi/10.1103/PhysRevLett.57.2607>.
- Zoltán Szabó. Information theoretical estimators toolbox. *Journal of Machine Learning Research*, 15(9):283–287, 2014. URL <http://jmlr.org/papers/v15/szabo14a.html>.
- Ali Süleyman Üstünel and Moshe Zakai. *Transformation of Measure on Wiener Space*. Springer Monographs in Mathematics. Springer Berlin, Heidelberg, 1 edition, 2013. doi: 10.1007/978-3-662-13225-8.
- 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=8pvnfTAbulf>.
- Francisco Vargas, Shreyas Padhy, Denis Blessing, and Nikolas Nüsken. Transport meets variational inference: Controlled Monte Carlo diffusions. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=PPlrudnxiW>.
- Santosh S. Vempala and Andre Wibisono. *Rapid Convergence of the Unadjusted Langevin Algorithm: Isoperimetry Suffices*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper_files/paper/2019/file/65a99bb7a3115fdede20da98b08a370f-Paper.pdf.
- Cédric Villani. *Optimal Transport: Old and New*. Grundlehren der mathematischen Wissenschaften. Springer Berlin, Heidelberg, 1 edition, 2008. doi: 10.1007/978-3-540-71050-9. URL <https://link.springer.com/book/10.1007/978-3-540-71050-9>.

- Cédric Villani. *Topics in optimal transportation*, volume 58. American Mathematical Society, 2021.
- Dawn B. Woodard, Scott C. Schmidler, and Mark Huber. Conditions for rapid mixing of parallel and simulated tempering on multimodal distributions. *The Annals of Applied Probability*, 19(2):617–640, 2009. doi: 10.1214/08-AAP555. URL <https://doi.org/10.1214/08-AAP555>.
- Qinsheng Zhang and Yongxin Chen. Path integral sampler: A stochastic control approach for sampling. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=_uCb2ynRu7Y.
- Qinsheng Zhang and Yongxin Chen. Fast sampling of diffusion models with exponential integrator. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=Loek7hfb46P>.
- Qinsheng Zhang, Molei Tao, and Yongxin Chen. gDDIM: Generalized denoising diffusion implicit models. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=1hKE9qjvz->.
- Nicolas Zilberstein, Chris Dick, Rahman Doost-Mohammady, Ashutosh Sabharwal, and Santiago Segarra. Annealed Langevin dynamics for massive MIMO detection. *IEEE Transactions on Wireless Communications*, 22(6):3762–3776, 2023. doi: 10.1109/TWC.2022.3221057.
- Nicolas Zilberstein, Ashutosh Sabharwal, and Santiago Segarra. Solving linear inverse problems using higher-order Annealed langevin diffusion. *IEEE Transactions on Signal Processing*, 72:492–505, 2024. doi: 10.1109/TSP.2023.3348202.

A PROOF OF RESULTS ON ACTION

Lemma 4 (Sufficient Condition for Absolute Continuity (Informal)). *Assume that a curve of probability distributions $(\pi_\theta)_{\theta \in [0,1]}$ on \mathbb{R}^d has a density $\pi : [0, 1] \times \mathbb{R}^d \ni (\theta, x) \mapsto \pi_\theta(x) > 0$ that is jointly C^1 . Then, this curve is AC.*

Proof. The following proof is informal. We leave the task of formalizing this statement for future work.

We first prove the result in one-dimension. Let $F_\theta(x) = \int_{-\infty}^x \pi_\theta(u) du$ be the c.d.f. of π_θ , which is strictly increasing as $\pi_\theta(x) > 0$ everywhere. As a result, its inverse F_θ^{-1} is well-defined, and $F_\theta^{-1}(\cdot)$ is also jointly C^1 . We know from (Ambrosio et al., 2008, Theorem 6.0.2) that

$$W_2^2(\pi_\theta, \pi_{\theta+\delta}) = \int_0^1 |F_{\theta+\delta}^{-1}(q) - F_\theta^{-1}(q)|^2 dq.$$

The above quantity should be $O(\delta^2)$ as $\delta \rightarrow 0$ due to the C^1 property, and hence the curve is AC.

In d -dimensions, consider the sliced Wasserstein-2 distance

$$\overline{W}_2(\mu, \nu) = \sup_{v \in \mathbb{S}^{d-1}} W_2(v_\# \mu, v_\# \nu),$$

where $\mathbb{S}^{d-1} = \{v \in \mathbb{R}^d : \|v\| = 1\}$ is the set of unit vectors in \mathbb{R}^d , and $v_\# \mu$ (respectively, $v_\# \nu$) is the law of $v^\top X$ when $X \sim \mu$ (respectively, $X \sim \nu$), which is a probability measure in \mathbb{R} . By (Bayraktar & Guo, 2021, Section 2.3), $W_2(\mu, \nu) \leq \sqrt{d} \overline{W}_2(\mu, \nu)$. Similar argument as above shows that $W_2^2(v_\# \pi_\theta, v_\# \pi_{\theta+\delta}) = O(\delta^2)$, and hence $W_2^2(\pi_\theta, \pi_{\theta+\delta}) = O(\delta^2)$. \square

Proof of Lemma 3.

Proof. 1. By Cauchy-Schwarz inequality,

$$\mathcal{A} = \int_0^1 |\dot{\rho}|_t^2 dt \int_0^1 1 dt \geq \left(\int_0^1 |\dot{\rho}|_t dt \right)^2.$$

It remains to prove that $W_2(\rho_0, \rho_1) \leq \int_0^1 |\dot{\rho}|_t dt$. In fact, take any sequence $0 = t_0 < t_1 < \dots < t_N = 1$, and let $\Delta t := \max_{1 \leq i \leq N} (t_i - t_{i-1})$. By triangle inequality of Wasserstein distance,

$$\begin{aligned} W_2(\rho_0, \rho_1) &\leq \sum_{i=1}^N W_2(\rho_{t_i}, \rho_{t_{i-1}}) \\ &= \sum_{i=1}^N \frac{W_2(\rho_{t_i}, \rho_{t_{i-1}})}{t_i - t_{i-1}} (t_i - t_{i-1}) \\ &= \sum_{i=1}^N (|\dot{\rho}|_{t_i} + o(\Delta t)) (t_i - t_{i-1}) \rightarrow \int_0^1 |\dot{\rho}|_t dt. \end{aligned}$$

Another derivation is by investigating the variational representations of these two quantities. First, by Lemma 2, one can write

$$\mathcal{A} = \int_0^1 |\dot{\rho}|_t^2 dt = \inf_{v_t: \partial_t \rho_t + \nabla \cdot (\rho_t v_t) = 0} \int_0^1 \|v_t\|_{L^2(\rho_t)}^2 dt.$$

On the other hand, the Benamou-Brenier formula (Ambrosio et al., 2008, Equation 8.0.3) implies

$$W_2^2(\rho_0, \rho_1) = \inf_{(\rho_t, v_t): \partial_t \rho_t + \nabla \cdot (\rho_t v_t) = 0; \rho_{t=0} = \rho_0, \rho_{t=1} = \rho_1} \int_0^1 \|v_t\|_{L^2(\rho_t)}^2 dt.$$

Thus, one can easily conclude the desired inequality.

To show that the inequality is attained when $(\rho_t)_{t \in [0,1]}$ is the constant-speed Wasserstein geodesic, note that the derivation implies that the inequality is attained when $|\dot{\rho}|_t$ is a constant for all $t \in [0, 1]$, and for any $0 \leq t_1 < t_2 < t_3 \leq 1$, $W_2(\rho_{t_1}, \rho_{t_2}) + W_2(\rho_{t_2}, \rho_{t_3}) = W_2(\rho_{t_1}, \rho_{t_3})$. We refer readers to (Ambrosio et al., 2008, Chapter 7.2) for the construction of the constant-speed Wasserstein geodesic.

2. It is known from Bakry et al. (2014) that π satisfies C -LSI implies

$$\text{KL}(\mu \|\pi) \leq \frac{C}{2} \mathbb{E}_\mu \left\| \nabla \log \frac{d\mu}{d\pi} \right\|^2, \quad \text{KL}(\mu \|\pi) \geq \frac{1}{2C} W_2^2(\mu, \pi), \quad \forall \mu.$$

Therefore,

$$\frac{1}{\delta^2} W_2^2(\rho_{t+\delta}, \rho_t) \leq C_{\text{LSI}}(\rho_t)^2 \int \frac{\|\nabla \log \rho_{t+\delta} - \nabla \log \rho_t\|^2}{\delta^2} d\rho_{t+\delta}.$$

By letting $\delta \rightarrow 0$ and assuming regularity conditions, we have

$$|\dot{\rho}|_t^2 \leq C_{\text{LSI}}(\rho_t)^2 \int \|\partial_t \nabla \log \rho_t\|^2 d\rho_t = C_{\text{LSI}}(\rho_t)^2 \|\partial_t \nabla \log \rho_t\|_{L^2(\rho_t)}^2$$

3. By Liu (2020), we know that π satisfies C -PI implies

$$W_2^2(\mu, \pi) \leq 2C \left(\int \frac{d\mu}{d\pi} d\mu - 1 \right), \quad \forall \mu.$$

Therefore,

$$\frac{1}{\delta^2} W_2^2(\rho_{t+\delta}, \rho_t) \leq 2C_{\text{PI}}(\rho_t) \int \frac{(\rho_{t+\delta} - \rho_t)^2}{\delta^2 \rho_t} dx.$$

By letting $\delta \rightarrow 0$, we have

$$|\dot{\rho}|_t^2 \leq 2C_{\text{PI}}(\rho_t) \int \frac{(\partial_t \rho_t)^2}{\rho_t} dx = 2C_{\text{PI}}(\rho_t) \|\partial_t \log \rho_t\|_{L^2(\rho_t)}^2.$$

□

Proof of Example 1.

Proof. 1. The reparameterized curve $(p_s)_{s \in [0, S]}$ satisfies the standard heat equation:

$$\partial_s p_s = \Delta p_s \implies \partial_s p_s + \nabla \cdot (p_s (-\nabla \log p_s)) = 0,$$

which means $(-\nabla \log p_s)_{s \in [0, S]}$ generates $(p_s)_{s \in [0, S]}$. According to the uniqueness result in Lemma 2 and also (Ambrosio et al., 2008, Theorem 8.3.1), the Lebesgue-a.e. unique vector field $(v_t^*)_{t \in [a, b]}$ that generates $(\rho_t)_{t \in [a, b]}$ and satisfies $|\dot{\rho}|_t = \|v_t^*\|_{L^2(\rho_t)}$ for Lebesgue-a.e. $t \in [a, b]$ can be written in a gradient field $v_t^* = \nabla \phi_t$ for some $\phi_t : \mathbb{R}^d \rightarrow \mathbb{R}$. This implies that

$$|\dot{p}|_s = \|\nabla \log p_s\|_{L^2(p_s)}^2.$$

By the time change $[0, S] \ni s \mapsto \theta = s/S \in [0, 1]$ and taking integral, we obtain the desired result.

2. Since $p_s = \sum_{i=1}^N w_i \mathcal{N}(\mu_i, (\sigma^2 + 2s)I)$, by Lemma 7, we know that $-\nabla^2 \log p_s \preceq \frac{1}{\sigma^2 + 2s} I$. Thus Lemma 6 implies $\|\nabla \log p_s\|_{L^2(p_s)}^2 \leq \frac{d}{\sigma^2 + 2s}$, so

$$\mathcal{A} \leq S \int_0^S \frac{d}{\sigma^2 + 2s} ds = \frac{Sd}{2} \log \left(1 + \frac{2S}{\sigma^2} \right).$$

Finally, the LSI and PI constant of mixture of Gaussian distributions are, in general, exponential with respect to the maximum distance between the means. To illustrate this, consider a simpler case where there are only 2 Gaussian components: $\rho = \frac{1}{2} \mathcal{N}(0, \sigma^2 I) + \frac{1}{2} \mathcal{N}(y, \sigma^2 I)$. (Schlichting, 2019, Section 4.1) has shown that $C_{\text{PI}}(\rho) \leq C_{\text{LSI}}(\rho) \leq \frac{\sigma^2}{2} (e^{\|y\|^2/\sigma^2} + 3)$, while (Dong & Tong, 2022, Proposition 1) has shown that when $\|y\| \gtrsim \sigma \sqrt{d}$, $C_{\text{LSI}}(\rho) \geq C_{\text{PI}}(\rho) \gtrsim \frac{\sigma^4}{\|y\|^2} e^{\Omega(\|y\|^2/\sigma^2)}$.

□

B SAMPLING FROM π_0

Lemma 5. When $\eta_0 = 0$, $\pi_0 = \mathcal{N}(0, \frac{1}{\lambda_0} I)$ can be sampled directly; when $\eta_0 \in (0, 1]$, we choose $\lambda_0 = \eta_0 d \beta$, so that under Assumption 2, it takes $O\left(1 \vee \log \frac{\eta_0 \beta R^2}{d^2}\right) = \tilde{O}(1)$ queries to the oracle of V and ∇V in expectation to precisely sample from π_0 via rejection sampling.

Proof. Our proof is inspired by Liang & Chen (2023). We only consider the nontrivial case $\eta_0 \in (0, 1]$. The potential of π_0 is $V_0 := \eta_0 V + \frac{\lambda_0}{2} \|\cdot\|^2$, which is $(\lambda_0 - \eta_0 \beta)$ -strongly-convex and $(\lambda_0 + \eta_0 \beta)$ -smooth. Note that for any fixed point $x' \in \mathbb{R}^d$,

$$\pi_0(x) \propto \exp(-V_0(x)) \leq \exp\left(-V_0(x') - \langle \nabla V_0(x'), x - x' \rangle - \frac{\lambda_0 - \eta_0 \beta}{2} \|x - x'\|^2\right).$$

The right-hand side is the unnormalized density of $\pi'_0 := \mathcal{N}\left(x' - \frac{\nabla V_0(x')}{\lambda_0 - \eta_0 \beta}, \frac{1}{\lambda_0 - \eta_0 \beta} I\right)$. The rejection sampling algorithm is as follows: sample $X \sim \pi'_0$, and accept X as a sample from π_0 with probability

$$\exp\left(-V_0(X) + V_0(x') + \langle \nabla V_0(x'), X - x' \rangle + \frac{\lambda_0 - \eta_0 \beta}{2} \|X - x'\|^2\right) \in (0, 1].$$

By (Chewi, 2024, Theorem 7.1.1), the number of queries to the oracle of V until acceptance follows a geometric distribution with mean

$$\begin{aligned}
& \frac{\int \exp \left(-V_0(x') - \langle \nabla V_0(x'), x - x' \rangle - \frac{\lambda_0 - \eta_0 \beta}{2} \|x - x'\|^2 \right) dx}{\int \exp(-V_0(x)) dx} \\
& \leq \frac{\int \exp \left(-V_0(x') - \langle \nabla V_0(x'), x - x' \rangle - \frac{\lambda_0 - \eta_0 \beta}{2} \|x - x'\|^2 \right) dx}{\int \exp \left(-V_0(x') - \langle \nabla V_0(x'), x - x' \rangle - \frac{\lambda_0 + \eta_0 \beta}{2} \|x - x'\|^2 \right) dx} \\
& = \frac{\int \exp \left(-\langle \nabla V_0(x'), x \rangle - \frac{\lambda_0 - \eta_0 \beta}{2} \|x\|^2 \right) dx}{\int \exp \left(-\langle \nabla V_0(x'), x \rangle - \frac{\lambda_0 + \eta_0 \beta}{2} \|x\|^2 \right) dx} \\
& = \frac{\int \exp \left(-\frac{\lambda_0 - \eta_0 \beta}{2} \left\| x + \frac{\nabla V_0(x')}{\lambda_0 - \eta_0 \beta} \right\|^2 + \frac{\|\nabla V_0(x')\|^2}{2(\lambda_0 - \eta_0 \beta)} \right) dx}{\int \exp \left(-\frac{\lambda_0 + \eta_0 \beta}{2} \left\| x + \frac{\nabla V_0(x')}{\lambda_0 + \eta_0 \beta} \right\|^2 + \frac{\|\nabla V_0(x')\|^2}{2(\lambda_0 + \eta_0 \beta)} \right) dx} \\
& = \left(\frac{\lambda_0 + \eta_0 \beta}{\lambda_0 - \eta_0 \beta} \right)^{\frac{d}{2}} \exp \left(\frac{\eta_0 \beta}{\lambda_0^2 - \eta_0^2 \beta^2} \|\nabla V_0(x')\|^2 \right) \\
& \leq \exp \left(\frac{\eta_0 \beta d}{\lambda_0 - \eta_0 \beta} \right) \exp \left(\frac{\eta_0 \beta}{\lambda_0^2 - \eta_0^2 \beta^2} \|\nabla V_0(x')\|^2 \right)
\end{aligned}$$

We choose $\lambda_0 = \eta_0 \beta d$ such that $\exp \left(\frac{\eta_0 \beta d}{\lambda_0 - \eta_0 \beta} \right) \lesssim 1$. With this λ_0 , $\exp \left(\frac{\eta_0 \beta}{\lambda_0^2 - \eta_0^2 \beta^2} \|\nabla V_0(x')\|^2 \right)$ is also $O(1)$ as long as $\|\nabla V_0(x')\|^2 \lesssim \eta_0 \beta d^2$.

Let x'' be the global minimizer of the strongly convex potential function V_0 , which satisfies

$$0 = \nabla V_0(x'') = \eta_0 \nabla V(x'') + \lambda_0 x'' = \eta_0 \nabla V(x'') + \eta_0 \beta d x''.$$

Given the smoothness of V_0 ,

$$\|\nabla V_0(x')\|^2 = \|\nabla V_0(x') - \nabla V_0(x'')\|^2 \leq (\eta_0 \beta d + \eta_0 \beta)^2 \|x' - x''\|^2.$$

Therefore, to guarantee $\|\nabla V_0(x')\|^2 \lesssim \eta_0 \beta d^2$, it suffices to find an x' such that $\|x' - x''\| \lesssim \frac{1}{\sqrt{\eta_0 \beta}}$.

We first derive an upper bound of $\|x''\|$ under Assumption 2. Since x_* is a global minimizer of V ,

$$\begin{aligned}
\beta d \|x''\| &= \|\nabla V(x'')\| = \|\nabla V(x'') - \nabla V(x_*)\| \\
&\leq \beta \|x'' - x_*\| \leq \beta (\|x''\| + R), \\
\Rightarrow \|x''\| &\leq \frac{R}{d-1} \lesssim \frac{R}{d}.
\end{aligned}$$

When $R = 0$, i.e., 0 is a known global minimizer of both V and V_0 , we can directly set $x' = 0$; otherwise, we need to run optimization algorithms to find such an x' . According to standard results in convex optimization (see, e.g., (Garrigos & Gower, 2023, Theorem 3.6)), running gradient descent on function V_0 with step size $\frac{1}{\lambda_0 + \eta_0 \beta}$ yields the following convergence rate: starting from $x_0 = 0$, the t -th iterate x_t satisfies

$$\|x_t - x''\|^2 \leq \left(1 - \frac{\lambda_0 - \eta_0 \beta}{\lambda_0 + \eta_0 \beta} \right)^t \|0 - x''\|^2 \lesssim \left(\frac{2}{d} \right)^t \frac{R^2}{d^2} \leq \frac{R^2}{e^t d^2},$$

where the last inequality holds when $d \geq 6$. Thus, $\log \frac{\eta_0 \beta R^2}{d^2} + O(1)$ iterations are sufficient to find a desired x' . We summarize the rejection sampling in Algorithm 2. In conclusion, precisely sampling from π_0 requires

$$O \left(1 \vee \log \frac{\eta_0 \beta R^2}{d^2} \right) = \tilde{O}(1)$$

calls to the oracle of V and ∇V in expectation. \square

Algorithm 2: Rejection Sampling for π_0 .**Input:** $\pi \propto \exp(-V)$, $\eta_0 \in (0, 1]$, $\lambda_0 = \eta_0 \beta d$.

- 1 Let $V_0 := \eta_0 V + \frac{\lambda_0}{2} \|\cdot\|^2$;
 - 2 Use gradient descent to find an x' that is $O\left(\frac{1}{\sqrt{\eta_0 \beta}}\right)$ -close to the global minimizer of V_0 ;
 - 3 Let $\pi'_0 := \mathcal{N}\left(x' - \frac{\nabla V_0(x')}{\lambda_0 - \eta_0 \beta}, \frac{1}{\lambda_0 - \eta_0 \beta} I\right)$;
 - 4 **while** *True* **do**
 - 5 Sample $X \sim \pi'_0$ and $U \in [0, 1]$ independently;
 - 6 Accept X as a sample of π_0 when
$$U \leq \exp\left(-V_0(X) + V_0(x') + \langle \nabla V_0(x'), X - x' \rangle + \frac{\lambda_0 - \eta_0 \beta}{2} \|X - x'\|^2\right).$$
 - 7 **end**
- Output:** $X \sim \pi_0$.

Remark. The parameter R reflects prior knowledge about global minimizer(s) of the potential function V . Unless it is exceptionally large, indicating scarce prior information about the global minimizer(s) of V , this $\tilde{O}(1)$ complexity is negligible compared to the overall complexity of sampling. In particular, when the exact location of a global minimizer of V is known, we can adjust the potential V so that 0 becomes a global minimizer, thereby reducing the complexity to $O(1)$.

C PROOFS FOR ANNEALED LMC**C.1 PROOF OF EQUATION (7)**

Define $\Lambda(t) := \int_0^t \lambda\left(\frac{\tau}{T}\right) d\tau$, whose derivative is $\Lambda'(t) = \lambda\left(\frac{t}{T}\right)$. By Itô's formula, we have

$$d\left(e^{\Lambda(t)} X_t\right) = e^{\Lambda(t)} \left(\lambda\left(\frac{t}{T}\right) X_t dt + dB_t \right) = e^{\Lambda(t)} \left(-\eta \left(\frac{t}{T}\right) \nabla V(X_{t-}) dt + \sqrt{2} dB_t \right).$$

Integrating over $t \in [T_{\ell-1}, T_\ell]$ (note that in this case $t_- = T_{\ell-1}$), we obtain

$$e^{\Lambda(T_\ell)} X_{T_\ell} - e^{\Lambda(T_{\ell-1})} X_{T_{\ell-1}} = - \left(\int_{T_{\ell-1}}^{T_\ell} \eta \left(\frac{t}{T}\right) e^{\Lambda(t)} dt \right) \nabla V(X_{T_{\ell-1}}) + \sqrt{2} \int_{T_{\ell-1}}^{T_\ell} e^{\Lambda(t)} dB_t,$$

i.e.,

$$\begin{aligned} X_{T_\ell} &= e^{-(\Lambda(T_\ell) - \Lambda(T_{\ell-1}))} X_{T_{\ell-1}} - \left(\int_{T_{\ell-1}}^{T_\ell} \eta \left(\frac{t}{T}\right) e^{-(\Lambda(T_\ell) - \Lambda(t))} dt \right) \nabla V(X_{T_{\ell-1}}) \\ &\quad + \sqrt{2} \int_{T_{\ell-1}}^{T_\ell} e^{-(\Lambda(T_\ell) - \Lambda(t))} dB_t. \end{aligned}$$

Since

$$\Lambda(T_\ell) - \Lambda(t) = \int_t^{T_\ell} \lambda\left(\frac{\tau}{T}\right) d\tau = T \int_{t/T}^{T_\ell/T} \lambda(u) du,$$

by defining $\Lambda_0(\theta', \theta) = \exp\left(-T \int_\theta^{\theta'} \lambda(u) du\right)$, we have $e^{-(\Lambda(T_\ell) - \Lambda(T_{\ell-1}))} = \Lambda_0(\theta_\ell, \theta_{\ell-1})$. Similarly, we can show that

$$\int_{T_{\ell-1}}^{T_\ell} \eta \left(\frac{t}{T}\right) e^{-(\Lambda(T_\ell) - \Lambda(t))} dt = \int_{T_{\ell-1}}^{T_\ell} \eta \left(\frac{t}{T}\right) \Lambda_0 \left(\frac{T_\ell}{T}, \frac{t}{T}\right) dt = T \int_{\theta_{\ell-1}}^{\theta_\ell} \eta(u) \Lambda_0(\theta_\ell, u) du,$$

and $\sqrt{2} \int_{T_{\ell-1}}^{T_\ell} e^{-(\Lambda(T_\ell) - \Lambda(t))} dB_t$ is a zero-mean Gaussian random vector with covariance

$$2 \int_{T_{\ell-1}}^{T_\ell} e^{-2(\Lambda(T_\ell) - \Lambda(t))} dt \cdot I = 2 \int_{T_{\ell-1}}^{T_\ell} \Lambda_0^2 \left(\frac{T_\ell}{T}, \frac{t}{T}\right) dt \cdot I = 2T \int_{\theta_{\ell-1}}^{\theta_\ell} \Lambda_0^2(\theta_\ell, u) du \cdot I.$$

C.2 PROOF OF THEOREM 2

We denote the path measure of ALMC (Equation (6)) by \mathbb{Q} . Then, \mathbb{Q}_T , the marginal distribution of X_T , serves as the output distribution ν^{ALMC} . Similar to the methodology in the proof of Theorem 1, we use \mathbb{P} to denote the reference path measure of Equation (4), in which the vector field $(v_t)_{t \in [0, T]}$ generates the cure of probability distributions $(\tilde{\pi}_t)_{t \in [0, T]}$.

Using the data-processing inequality, it suffices to demonstrate that $\text{KL}(\mathbb{P} \parallel \mathbb{Q}) \leq \varepsilon^2$. By Girsanov theorem (Lemma 1) and triangle inequality, we have

$$\begin{aligned} \text{KL}(\mathbb{P} \parallel \mathbb{Q}) &= \frac{1}{4} \int_0^T \mathbb{E}_{\mathbb{P}} \left\| \eta \left(\frac{t}{T} \right) (\nabla V(X_t) - \nabla V(X_{t-})) - v_t(X_t) \right\|^2 dt \\ &\lesssim \int_0^T \mathbb{E}_{\mathbb{P}} \left[\eta \left(\frac{t}{T} \right)^2 \left\| \nabla V(X_t) - \nabla V(X_{t-}) \right\|^2 + \|v_t(X_t)\|^2 \right] dt \\ &\leq \sum_{\ell=1}^M \eta \left(\frac{T_\ell}{T} \right)^2 \beta^2 \int_{T_{\ell-1}}^{T_\ell} \mathbb{E}_{\mathbb{P}} \|X_t - X_{t-}\|^2 dt + \int_0^T \|v_t\|_{L^2(\tilde{\pi}_t)}^2 dt. \end{aligned}$$

The last inequality arises from the smoothness of V and the increasing property of $\eta(\cdot)$. To bound $\mathbb{E}_{\mathbb{P}} \|X_t - X_{t-}\|^2$, note that under \mathbb{P} , for $t \in [T_{\ell-1}, T_\ell)$, we have

$$X_t - X_{t-} = \int_{T_{\ell-1}}^t (\nabla \log \tilde{\pi}_\tau + v_\tau)(X_\tau) d\tau + \sqrt{2}(B_t - B_{T_{\ell-1}}).$$

Thanks to the fact that $X_t \sim \tilde{\pi}_t$ under \mathbb{P} ,

$$\begin{aligned} \mathbb{E}_{\mathbb{P}} \|X_t - X_{t-}\|^2 &\lesssim \mathbb{E}_{\mathbb{P}} \left\| \int_{T_{\ell-1}}^t (\nabla \log \tilde{\pi}_\tau + v_\tau)(X_\tau) d\tau \right\|^2 + \mathbb{E} \left\| \sqrt{2}(B_t - B_{T_{\ell-1}}) \right\|^2 \\ &\lesssim (t - T_{\ell-1}) \int_{T_{\ell-1}}^t \mathbb{E}_{\mathbb{P}} \|(\nabla \log \tilde{\pi}_\tau + v_\tau)(X_\tau)\|^2 d\tau + d(t - T_{\ell-1}) \\ &\lesssim (t - T_{\ell-1}) \int_{T_{\ell-1}}^t \left(\|\nabla \log \tilde{\pi}_\tau\|_{L^2(\tilde{\pi}_\tau)}^2 + \|v_\tau\|_{L^2(\tilde{\pi}_\tau)}^2 \right) d\tau + d(t - T_{\ell-1}) \\ &\lesssim h_\ell \int_{T_{\ell-1}}^{T_\ell} \left(\|\nabla \log \tilde{\pi}_\tau\|_{L^2(\tilde{\pi}_\tau)}^2 + \|v_\tau\|_{L^2(\tilde{\pi}_\tau)}^2 \right) d\tau + dh_\ell. \end{aligned}$$

The second inequality arises from the application of the Cauchy-Schwarz inequality, and the last inequality is due to the definition $h_\ell = T_\ell - T_{\ell-1}$. Taking integral over $t \in [T_{\ell-1}, T_\ell]$,

$$\int_{T_{\ell-1}}^{T_\ell} \mathbb{E} \|X_t - X_{t-}\|^2 dt \lesssim h_\ell^2 \int_{T_{\ell-1}}^{T_\ell} \left(\|\nabla \log \tilde{\pi}_t\|_{L^2(\tilde{\pi}_t)}^2 + \|v_t\|_{L^2(\tilde{\pi}_t)}^2 \right) dt + dh_\ell^2.$$

Recall that the potential of $\tilde{\pi}_t$ is $(\eta(\frac{t}{T})\beta + \lambda(\frac{t}{T}))$ -smooth. By Lemma 6 and the monotonicity of $\eta(\cdot)$ and $\lambda(\cdot)$, we have

$$\begin{aligned} \int_{T_{\ell-1}}^{T_\ell} \|\nabla \log \tilde{\pi}_t\|_{L^2(\tilde{\pi}_t)}^2 dt &\leq \int_{T_{\ell-1}}^{T_\ell} d \left(\eta \left(\frac{t}{T} \right) \beta + \lambda \left(\frac{t}{T} \right) \right) dt \\ &\leq dh_\ell \left(\beta \eta \left(\frac{T_\ell}{T} \right) + \lambda \left(\frac{T_{\ell-1}}{T} \right) \right) \\ &= dh_\ell (\beta \eta(\theta_\ell) + \lambda(\theta_{\ell-1})). \end{aligned}$$

Therefore, $\text{KL}(\mathbb{P} \parallel \mathbb{Q})$ is upper bounded by

$$\begin{aligned} &\lesssim \sum_{\ell=1}^M \eta(\theta_\ell)^2 \beta^2 \int_{T_{\ell-1}}^{T_\ell} \mathbb{E}_{\mathbb{P}} \|X_t - X_{t-}\|^2 dt + \int_0^T \|v_t\|_{L^2(\tilde{\pi}_t)}^2 dt \\ &\lesssim \sum_{\ell=1}^M \eta(\theta_\ell)^2 \beta^2 \left(dh_\ell^3 (\beta \eta(\theta_\ell) + \lambda(\theta_{\ell-1})) + h_\ell^2 \int_{T_{\ell-1}}^{T_\ell} \|v_t\|_{L^2(\tilde{\pi}_t)}^2 dt + dh_\ell^2 \right) + \int_0^T \|v_t\|_{L^2(\tilde{\pi}_t)}^2 dt \\ &= \sum_{\ell=1}^M \left((1 + \eta(\theta_\ell)^2 \beta^2 h_\ell^2) \int_{T_{\ell-1}}^{T_\ell} \|v_t\|_{L^2(\tilde{\pi}_t)}^2 dt + \eta(\theta_\ell)^2 \beta^2 dh_\ell^2 (1 + h_\ell (\beta \eta(\theta_\ell) + \lambda(\theta_{\ell-1}))) \right). \end{aligned}$$

For the remaining integral, given that $(v_t)_{t \in [0, T]}$ generates $(\tilde{\pi}_t)_{t \in [0, T]}$, according to Lemma 2, we may choose v_t such that $\|v_t\|_{L^2(\tilde{\pi}_t)} = \left| \dot{\tilde{\pi}} \right|_t$. Thus,

$$\int_{T_{\ell-1}}^{T_\ell} \|v_t\|_{L^2(\tilde{\pi}_t)}^2 dt = \int_{T_{\ell-1}}^{T_\ell} \left| \dot{\tilde{\pi}} \right|_t^2 dt = \int_{T_{\ell-1}}^{T_\ell} \frac{1}{T^2} |\dot{\tilde{\pi}}|_{t/T}^2 dt = \frac{1}{T} \int_{\theta_{\ell-1}}^{\theta_\ell} |\dot{\tilde{\pi}}|_\theta^2 d\theta,$$

through a change-of-variable analogous to that used in the proof of Theorem 1. Therefore,

$$\text{KL}(\mathbb{P} \parallel \mathbb{Q}) \lesssim \sum_{\ell=1}^M \left(\frac{1 + \eta(\theta_\ell)^2 \beta^2 h_\ell^2}{T} \int_{\theta_{\ell-1}}^{\theta_\ell} |\dot{\tilde{\pi}}|_\theta^2 d\theta + \eta(\theta_\ell)^2 \beta^2 dh_\ell^2 (1 + h_\ell (\beta \eta(\theta_\ell) + \lambda(\theta_{\ell-1}))) \right).$$

Assume $h_\ell \lesssim \frac{1}{\beta d}$ (which will be verified later), so we can further simplify the above expression to

$$\begin{aligned} \text{KL}(\mathbb{P} \parallel \mathbb{Q}) &\lesssim \sum_{\ell=1}^M \left(\frac{1}{T} \int_{\theta_{\ell-1}}^{\theta_\ell} |\dot{\tilde{\pi}}|_\theta^2 d\theta + \eta(\theta_\ell)^2 \beta^2 dh_\ell^2 \right) \\ &= \frac{\mathcal{A}}{T} + \beta^2 d \sum_{\ell=1}^M \eta(\theta_\ell)^2 h_\ell^2 \\ &= \frac{\mathcal{A}}{T} + \beta^2 d \sum_{\ell=1}^M \eta(\theta_\ell)^2 T^2 (\theta_\ell - \theta_{\ell-1})^2. \end{aligned}$$

To bound the above expression by ε^2 , we first select $T \asymp \frac{\mathcal{A}}{\varepsilon^2}$, mirroring the total time T required for running annealed LD as specified in Theorem 1. This guarantees that the continuous dynamics closely approximate the reference path measure. Given that $\eta(\cdot) \leq 1$, it remains only to ensure

$$\beta^2 d \sum_{\ell=1}^M \eta(\theta_\ell)^2 T^2 (\theta_\ell - \theta_{\ell-1})^2 \leq \beta^2 d \frac{\mathcal{A}^2}{\varepsilon^4} \sum_{\ell=1}^M (\theta_\ell - \theta_{\ell-1})^2 \lesssim \varepsilon^2,$$

which is equivalent to

$$\sum_{\ell=1}^M (\theta_\ell - \theta_{\ell-1})^2 \lesssim \frac{\varepsilon^6}{d \beta^2 \mathcal{A}^2}. \quad (8)$$

To minimize M , we apply Cauchy-Schwarz inequality:

$$\left(\sum_{\ell=1}^M 1 \right) \left(\sum_{\ell=1}^M (\theta_\ell - \theta_{\ell-1})^2 \right) \geq \left(\sum_{\ell=1}^M (\theta_\ell - \theta_{\ell-1}) \right)^2 = 1,$$

which establishes a lower bound for M . The equality is achieved when $\theta_\ell - \theta_{\ell-1} = \frac{1}{M}$ for all $\ell \in [1, M]$. Thus, selecting

$$M \asymp \frac{d \beta^2 \mathcal{A}^2}{\varepsilon^6}$$

satisfies the constraint given in Equation (8). In this case, the step size $h_\ell = \frac{T}{M} \asymp \frac{\varepsilon^4}{d \beta^2} \lesssim \frac{1}{\beta d}$. Combining this with the $\tilde{O}(1)$ complexity for sampling from π_0 , we have completed the proof.

D PROOF OF EXAMPLE 2

The smoothness of V comes from Lemma 7.

Note that $\pi(x) \propto \sum_{i=1}^N p_i \exp\left(-\frac{\beta}{2} \|x - y_i\|^2\right)$, and define

$$\begin{aligned}\hat{\pi}_\lambda(x) &\propto \pi(x) \exp\left(-\frac{\lambda}{2} \|x\|^2\right) \\ &= \sum_{i=1}^N p_i \exp\left(-\frac{\lambda\beta}{2(\lambda+\beta)} \|y_i\|^2 - \frac{\lambda+\beta}{2} \left\|x - \frac{\beta}{\lambda+\beta} y_i\right\|^2\right) \\ &\propto \sum_{i=1}^N p_i \exp\left(-\frac{\lambda+\beta}{2} \left\|x - \frac{\beta}{\lambda+\beta} y_i\right\|^2\right) \\ &= \sum_{i=1}^N p_i \mathcal{N}\left(\frac{\beta}{\lambda+\beta} y_i, \frac{1}{\lambda+\beta} I\right).\end{aligned}$$

We use coupling method to upper bound $W_2^2(\hat{\pi}_\lambda, \hat{\pi}_{\lambda+\delta})$. We first sample I with distribution $\mathbb{P}(I = i) = p_i, i \in [1, N]$, and then independently sample $\eta \sim \mathcal{N}(0, I)$. Then,

$$\begin{aligned}X &:= \frac{\beta}{\lambda+\beta} y_I + \frac{1}{\sqrt{\lambda+\beta}} \eta \sim \hat{\pi}_\lambda, \\ Y &:= \frac{\beta}{\lambda+\delta+\beta} y_I + \frac{1}{\sqrt{\lambda+\delta+\beta}} \eta \sim \hat{\pi}_{\lambda+\delta}.\end{aligned}$$

By definition of W_2 distance, we have

$$\begin{aligned}W_2^2(\hat{\pi}_\lambda, \hat{\pi}_{\lambda+\delta}) &\leq \mathbb{E} \|X - Y\|^2 \\ &= \mathbb{E}_I \mathbb{E}_\eta \left\| \left(\frac{\beta}{\lambda+\beta} - \frac{\beta}{\lambda+\delta+\beta} \right) y_I + \left(\frac{1}{\sqrt{\lambda+\beta}} - \frac{1}{\sqrt{\lambda+\delta+\beta}} \right) \eta \right\|^2 \\ &= \mathbb{E}_I \left(\frac{\beta}{\lambda+\beta} - \frac{\beta}{\lambda+\delta+\beta} \right)^2 \|y_I\|^2 + \left(\frac{1}{\sqrt{\lambda+\beta}} - \frac{1}{\sqrt{\lambda+\delta+\beta}} \right)^2 d \\ &= \left(\frac{\beta}{\lambda+\beta} - \frac{\beta}{\lambda+\delta+\beta} \right)^2 r^2 + \left(\frac{1}{\sqrt{\lambda+\beta}} - \frac{1}{\sqrt{\lambda+\delta+\beta}} \right)^2 d.\end{aligned}$$

This implies

$$\left| \dot{\hat{\pi}} \right|_\lambda^2 = \lim_{\delta \rightarrow 0} \frac{W_2^2(\hat{\pi}_\lambda, \hat{\pi}_{\lambda+\delta})}{\delta^2} \leq \frac{\beta^2 r^2}{(\lambda+\beta)^4} + \frac{d}{4(\lambda+\beta)^3}.$$

By time reparameterization $\pi_\theta = \hat{\pi}_{\lambda(\theta)}$, $|\dot{\pi}|_\theta = \left| \dot{\hat{\pi}} \right|_{\lambda(\theta)} |\dot{\lambda}(\theta)|$. With $\lambda(\theta) = \lambda_0(1-\theta)^\gamma$, $|\dot{\lambda}(\theta)| = |\gamma \lambda_0(1-\theta)^{\gamma-1}| \leq \gamma \lambda_0 \lesssim \lambda_0$. Therefore,

$$\begin{aligned}\mathcal{A} &= \int_0^1 |\dot{\pi}|_\theta^2 d\theta = \int_0^1 \left| \dot{\hat{\pi}} \right|_{\lambda(\theta)}^2 |\dot{\lambda}(\theta)|^2 d\theta \\ &\lesssim \lambda_0 \int_0^1 \left| \dot{\hat{\pi}} \right|_{\lambda(\theta)}^2 |\dot{\lambda}(\theta)| d\theta = \lambda_0 \int_0^{\lambda_0} \left| \dot{\hat{\pi}} \right|_\lambda^2 d\lambda \\ &= \lambda_0 \int_0^{\lambda_0} \left(\frac{\beta^2 r^2}{(\lambda+\beta)^4} + \frac{d}{4(\lambda+\beta)^3} \right) d\lambda \\ &\lesssim \lambda_0 \left(\beta^2 r^2 \frac{1}{\beta^3} + d \frac{1}{\beta^2} \right) \\ &\lesssim d\beta(r^2\beta + 1) \left(\frac{r^2}{\beta} + \frac{d}{\beta^2} \right) \\ &= d(r^2\beta + 1) \left(r^2 + \frac{d}{\beta} \right).\end{aligned}$$

It follows from the proof that as long as $\max_{\theta \in [0,1]} |\dot{\lambda}(\theta)|$ is bounded by a polynomial function of β , d and r , so is the action \mathcal{A} .

E SUPPLEMENTARY LEMMAS

Lemma 6 ((Chewi, 2024, Lemma 4.E.1)). *Consider a probability measure $\mu \propto e^{-U}$ on \mathbb{R}^d . If $\nabla^2 U \preceq \beta I$ for some $\beta > 0$, then $\mathbb{E}_\mu \|\nabla U\|^2 \leq \beta d$.*

Lemma 7 ((Cheng et al., 2023, Lemma 4)). *For a mixture of Gaussian distribution $\pi = \sum_i w_i \mathcal{N}(\mu_i, \sigma^2 I)$, we have*

$$\left(\frac{1}{\sigma^2} - \frac{\max_{i,j} \|\mu_i - \mu_j\|^2}{2\sigma^4} \right) I \preceq -\nabla^2 \log \pi \preceq \frac{1}{\sigma^2} I.$$

Remark. *This is a slightly improved version, as in (Cheng et al., 2023, Lemma 4) the authors omitted the 2 in the denominator on the left-hand-side.*

Lemma 8. *Under Assumption 2, π_θ defined in Equation (5) has finite second-order moment when $\eta(\theta) \in [0, 1]$.*

Proof. When $\eta(\theta) = 1$, $\pi_\theta \propto \exp\left(-V - \frac{\lambda(\theta)}{2} \|\cdot\|^2\right) \leq \exp(-V)$, and the claim is straightforward. Otherwise, by convexity of $u \mapsto e^{-u}$, we have

$$\begin{aligned} \exp\left(-\eta(\theta)V - \frac{\lambda(\theta)}{2} \|\cdot\|^2\right) &= \exp\left(-\eta(\theta)V - (1 - \eta(\theta)) \frac{\lambda(\theta)}{2(1 - \eta(\theta))} \|\cdot\|^2\right) \\ &\leq \eta(\theta) \exp(-V) + (1 - \eta(\theta)) \exp\left(-\frac{\lambda(\theta)}{2(1 - \eta(\theta))} \|\cdot\|^2\right). \end{aligned}$$

Multiplying both sides by $\|\cdot\|^2$ and taking integral over \mathbb{R}^d , we see that $\mathbb{E}_{\pi_\theta} \|\cdot\|^2 < +\infty$. \square

F FURTHER DETAILS OF EXPERIMENTS IN SECTION 6

Hyperparameter settings: For all experiments, we use the annealing schedule $\lambda(\theta) = 5(1 - \theta)^{10}$ and $\eta(\theta) \equiv 1$. The step size for ALMC is designed to follow a quadratic schedule: the step size at the ℓ -th iteration (out of M total iterations, $\ell \in [1, M]$) is given by $-\frac{s_{\max} - s_{\min}}{M^2/4} \left(\ell - \frac{M}{2}\right)^2 + s_{\max}$, which increases on $[0, \frac{M}{2}]$ and decreases on $[\frac{M}{2}, M]$, with a maximum of s_{\max} and a minimum of s_{\min} . For $r = 2, 5, 10, 15, 20, 25, 30$, the M we choose are 200, 500, 2500, 10000, 20000, 40000, 60000, respectively, and we set $s_{\max} = 0.05$ and $s_{\min} = 0.01$, which achieve the best performance across all settings.

KL divergence estimation: We use the Information Theoretical Estimators (ITE) toolbox (Szabó, 2014) to empirically estimate the KL divergence. In all cases, we generate 1000 samples using ALMC, and an additional 1000 samples from the target distribution. The KL divergence is estimated using the `ite.cost.BDKL_KnnK()` function, which leverages k -nearest-neighbor techniques to compute the divergence from the target distribution to the sampled distribution.

Regression coefficients: We also compute the linear regression approximation of the curves in Figure 2 via `sklearn.linear_model.LinearRegression`. The blue curve has a slope of 2.841 and an intercept of 1.257, while the orange one has a slope of 2.890 and an intercept of 0.904. The R^2 scores, calculated using `sklearn.metrics.r2_score`, are 0.995 and 0.997, respectively.