

COMPLEXITY ANALYSIS OF NORMALIZING CONSTANT ESTIMATION: FROM JARZYNSKI EQUALITY TO ANNEALED IMPORTANCE SAMPLING AND BEYOND

Wei Guo, Molei Tao, Yongxin Chen

Georgia Institute of Technology

{wei.guo, mtao, yongchen}@gatech.edu

ABSTRACT

Given an unnormalized probability density $\pi \propto e^{-V}$, estimating its normalizing constant $Z = \int_{\mathbb{R}^d} e^{-V(x)} dx$ or free energy $F = -\log Z$ is a crucial problem in Bayesian statistics, statistical mechanics, and machine learning. It is challenging especially in high dimensions or when π is multimodal. To mitigate the high variance of conventional importance sampling estimators, annealing-based methods such as Jarzynski equality and annealed importance sampling are commonly adopted, yet their quantitative complexity guarantees remain largely unexplored. We take a first step toward a non-asymptotic analysis of annealed importance sampling. In particular, we derive an oracle complexity of $\tilde{O}\left(\frac{d\beta^2\mathcal{A}^2}{\varepsilon^4}\right)$ for estimating Z within ε relative error with high probability, where β is the smoothness of V and \mathcal{A} denotes the action of a curve of probability measures interpolating π and a tractable reference distribution. Our analysis, leveraging Girsanov theorem and optimal transport, does not explicitly require isoperimetric assumptions on the target distribution. Finally, to tackle the large action of the widely used geometric interpolation of probability distributions, we propose a new normalizing constant estimation algorithm based on reverse diffusion samplers and establish a framework for analyzing its complexity.

1 INTRODUCTION

We study the problem of estimating the normalizing constant $Z = \int_{\mathbb{R}^d} e^{-V(x)} dx$ of an unnormalized probability density function (p.d.f.) $\pi \propto e^{-V}$ on \mathbb{R}^d , so that $\pi(x) = \frac{1}{Z} e^{-V(x)}$. The normalizing constant appears in various fields: in Bayesian statistics, when e^{-V} is the product of likelihood and prior, Z is also referred to as the marginal likelihood or evidence (Gelman et al., 2013); in statistical mechanics, when V is the Hamiltonian¹, Z is known as the partition function, and $F := -\log Z$ is called the free energy (Chipot & Pohorille, 2007; Lelièvre et al., 2010; Pohorille et al., 2010). The task of normalizing constant estimation has numerous applications, including computing log-likelihoods in probabilistic models (Sohl-Dickstein & Culpepper, 2012), estimating free energy differences (Lelièvre et al., 2010), and training energy-based models in generative modeling (Song & Kingma, 2021; Carbone et al., 2023; Sander et al., 2025). It is challenging in high dimensions or when π is multimodal (i.e., V has a complex landscape).

Conventional approaches based on importance sampling (Meng & Wong, 1996) are widely adopted to tackle this problem, but they suffer from high variance due to the mismatch between target and proposal distributions when the target distribution is complicated (Chatterjee & Diaconis, 2018). To alleviate this issue, the technique of annealing tries constructing a sequence of intermediate distributions that bridge these two distributions, which motivates several popular methods including path sampling (Chen & Shao, 1997; Gelman & Meng, 1998), annealed importance sampling (AIS, Neal (2001)), and sequential Monte Carlo (SMC, Doucet et al. (2000); Del Moral et al. (2006); Syed

¹Up to a multiplicative constant $\beta = \frac{1}{k_B T}$ known as the thermodynamic beta, where k_B is the Boltzmann constant and T is the temperature. When borrowing terminologies from physics, we ignore this quantity for simplicity.

et al. (2024)) in statistics literature, as well as thermodynamic integration (TI, Kirkwood (1935)) and Jarzynski equality (JE, Jarzynski (1997); Ge & Jiang (2008); Hartmann et al. (2019)) in statistical mechanics literature. In particular, JE points out the connection between the free energy difference between two states and the work done over a series of trajectories linking these two states, while AIS constructs a sequence of intermediate distributions and estimates the normalizing constant by importance sampling over these distributions. These are our primary focus in this paper.

Despite the empirical success of annealing-based methods (Ma et al., 2013; Krause et al., 2020; Mazzanti & Romero, 2020; Yasuda & Takahashi, 2022; Chen & Ying, 2024; Schönle et al., 2025), the theoretical understanding of their performance is still limited. Existing works for importance sampling mainly focus on the asymptotic bias and variance of the estimator (Meng & Wong, 1996; Gelman & Meng, 1998), while works on JE usually simplify the problem by assuming the work follows simple distributions (e.g., Gaussian or gamma) (Echeverria & Amzel, 2012; Arrar et al., 2019). Moreover, only analyses asymptotic in the number of particles derived from central limit theorem exist (Lelièvre et al., 2010, Sec. 4.1). In this paper, we aim to establish a rigorous non-asymptotic analysis of estimators based on JE and AIS, while introducing minimal assumptions on the target distribution. Moreover, we also propose a new algorithm based on reverse diffusion samplers to tackle a potential shortcoming of AIS.

Contributions. Our key technical contributions are summarized as follows.

- We discover a novel strategy for analyzing the complexity of normalizing constant estimation, applicable to a wide range of target distributions (see Assumps. 1 and 2) that may not satisfy isoperimetric conditions such as log-concavity.
- In Sec. 4, we study JE and prove an upper bound on the time required for running the annealed Langevin dynamics to estimate the normalizing constant within ε relative error with high probability. The final bound depends on the action of the curve, specifically the integral of the squared metric derivative in Wasserstein-2 distance.
- Building on the insights from the analysis of the continuous dynamics, in Sec. 5 we establish the first non-asymptotic oracle complexity bound for AIS, representing the first analysis of normalizing constant estimation algorithms without assuming a log-concave target distribution.
- Finally, in Sec. 6, we point out a potential limitation of the geometric interpolation commonly used in annealing. To address this issue, we propose a novel algorithm based on reverse diffusion samplers and build up a framework for analyzing its oracle complexity.

2 PRELIMINARIES

2.1 STOCHASTIC DIFFERENTIAL EQUATIONS AND GIRSANOV THEOREM

Throughout this paper, (B_t) and (W_t) represent standard Brownian motions (BM) on \mathbb{R}^d . For a stochastic differential equation (SDE) $X = (X_t)_{t \in [0, T]}$ defined on $\Omega = C([0, T]; \mathbb{R}^d)$, the distribution of X over Ω is called the **path measure** of X , defined by \mathbb{P}^X : measurable $A \subset \Omega \mapsto \Pr(X \in A)$. The following lemma, as a corollary of the Girsanov theorem (Üstünel & Zakai, 2013, Prop. 2.3.1 & Cor. 2.3.1), provides a method for computing the Radon-Nikodým (RN) derivative and KL divergence between two path measures, which serves as a key technical tool in our proof.

Lemma 1. *Assume we have the following two SDEs with $t \in [0, T]$:*

$$dX_t = a_t(X_t)dt + \sigma dB_t, X_0 \sim \mu; \quad dY_t = b_t(Y_t)dt + \sigma dB_t, Y_0 \sim \nu.$$

Denote the path measures of X and Y as \mathbb{P}^X and \mathbb{P}^Y , respectively. Then for any trajectory $\xi \in \Omega$,

$$\log \frac{d\mathbb{P}^X}{d\mathbb{P}^Y}(\xi) = \log \frac{d\mu}{d\nu}(\xi_0) + \frac{1}{\sigma^2} \int_0^T \langle a_t(\xi_t) - b_t(\xi_t), d\xi_t \rangle - \frac{1}{2\sigma^2} \int_0^T (\|a_t(\xi_t)\|^2 - \|b_t(\xi_t)\|^2) dt.$$

In particular, plugging in $\xi \leftarrow X \sim \mathbb{P}^X$, we can compute the KL divergence:

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

We now define the **backward SDE**, which can be perceived as the time-reversal of a forward SDE. Given a BM $(B_t)_{t \in [0, T]}$, let its time-reversal be $(B_t^{\leftarrow} := B_{T-t})_{t \in [0, T]}$. We say that a process

$(X_t^\leftarrow)_{t \in [0, T]}$ satisfies the backward SDE $dX_t^\leftarrow = a_t(X_t^\leftarrow)dt + \sigma dB_t^\leftarrow$, $t \in [0, T]$; $X_T^\leftarrow \sim \nu$ if its time-reversal $(X_t = X_{T-t}^\leftarrow)_{t \in [0, T]}$ satisfies the following forward SDE: $dX_t = -a_{T-t}(X_t)dt + \sigma dB_t$, $t \in [0, T]$; $X_0 \sim \nu$.

The forward and backward SDEs are related through the Nelson’s relation (Lem. 3), which also allows us to calculate the RN derivative between path measures of forward and backward SDEs (Lem. 4). We postpone the detailed derivations to App. A.

2.2 WASSERSTEIN DISTANCE, METRIC DERIVATIVE, AND ACTION

We provide a concise overview of essential concepts in optimal transport (OT) that will be used in the paper. See standard textbooks (Villani, 2003; 2008; Ambrosio et al., 2008; 2021) for details.

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) . The Brenier’s theorem states that when μ has a Lebesgue density, then there exists a unique coupling $(\text{id} \times T_{\mu \rightarrow \nu})_{\#} \mu$ that reaches the infimum. Here, $\#$ stands for the push-forward of a measure ($T_{\#} \mu(\cdot) = \mu(\{\omega : T(\omega) \in \cdot\})$), and $T_{\mu \rightarrow \nu}$ is known as the **OT map** from μ to ν and can be written as the gradient of a convex function.

Given a vector field $v = (v_t)_{t \in [a, b]}$ and a curve of probability measures $\rho = (\rho_t)_{t \in [a, b]}$ with finite second-order moment on \mathbb{R}^d , 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 in the weak sense. 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. We say ρ is **absolutely continuous (AC)** if $|\dot{\rho}|_t$ exists and is finite for Lebesgue-a.e. $t \in [a, b]$. The metric derivative and the continuity equation are related through the following fact (Ambrosio et al., 2008, Thm. 8.3.1 & Prop. 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 Lebesgue-a.e. $t \in [a, b]$. Moreover, there exists an a.s. unique vector field $(v_t^* \in L^2(\rho_t))_{t \in [a, b]}$ that generates $(\rho_t)_{t \in [a, b]}$ and satisfies $|\dot{\rho}|_t = \|v_t^*\|_{L^2(\rho_t)}$ for Lebesgue-a.e. $t \in [a, b]$, which is $v_t^* = \lim_{\delta \rightarrow 0} \frac{T_{\rho_t \rightarrow \rho_{t+\delta}} - \text{id}}{\delta}$.*

Finally, we define the **action** of an AC curve of probability measures $(\rho_t)_{t \in [a, b]}$ as $\int_a^b |\dot{\rho}|_t^2 dt$, which plays a key role in characterizing the efficiency of a curve for normalizing constant estimation. For basic properties of the action and its relation to isoperimetric inequalities such as log-Sobolev and Poincaré inequalities, we refer the reader to Guo et al. (2025, Lem. 3 & Ex. 1).

2.3 LANGEVIN DIFFUSION AND LANGEVIN MONTE CARLO

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

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

Under mild regularity conditions, π is the unique stationary distribution of this SDE, and when π has good properties such as strong log-concavity, X_t converges to π in probability rapidly. In practice, when the closed-form solution of this SDE is unavailable, one usually leverages the Euler-Maruyama scheme to discretize Eq. (1), leading to the (overdamped) **Langevin Monte Carlo (LMC)** algorithm: with step size $h > 0$, iterate the following update rule for $k = 0, 1, \dots$:

$$X_{(k+1)h} = X_{kh} - h\nabla V(X_{kh}) + \sqrt{2}(B_{(k+1)h} - B_{kh}), \quad \text{where } B_{(k+1)h} - B_{kh} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, hI). \quad (2)$$

2.4 REVERSE DIFFUSION SAMPLERS

Inspired by score-based generative models (Song et al., 2021), recent advancements have led to the development of multimodal samplers based on reversing the Ornstein-Uhlenbeck (OU) process (Huang et al., 2024a;b; He et al., 2024; Vacher et al., 2025). In this paper, we collectively refer to these methods as the **reverse diffusion samplers (RDS)**.

The following OU process transforms any target distribution π into $\phi := \mathcal{N}(0, I)$ as $T \rightarrow \infty$:

$$dY_t = -Y_t dt + \sqrt{2} dB_t, \quad t \in [0, T]; \quad Y_0 \sim \pi, \quad (3)$$

We denote the law of Y_t by $\bar{\pi}_t$. The time-reversal $(Y_t^{\leftarrow} := Y_{T-t} \sim \bar{\pi}_{T-t})_{t \in [0, T]}$ satisfies the SDE

$$dY_t^{\leftarrow} = (Y_t^{\leftarrow} + 2\nabla \log \bar{\pi}_{T-t}(Y_t^{\leftarrow})) dt + \sqrt{2} dW_t, \quad t \in [0, T]; \quad Y_0^{\leftarrow} \sim \bar{\pi}_T (\approx \phi). \quad (4)$$

Hence, to draw samples from π , it suffices to approximate the scores $\nabla \log \bar{\pi}_t$ and discretize Eq. (4), which can be implemented in various ways. For example, by Tweedie’s formula (Robbins, 1992), $\nabla \log \bar{\pi}_t$ is an affine function of $\mathbb{E}(Y_0 | Y_t = \cdot)$ (Eq. (34)), while the law of $Y_0 | Y_t = \cdot$ is analytically tractable (Eq. (35)) and provably easier to sample from than the target π (Huang et al., 2024a).

3 PROBLEM SETTING

Motivated by Brosse et al. (2018); Ge et al. (2020), given an accuracy threshold $\varepsilon \ll 1$, our goal is to study the complexity (measured by the number of calls to the oracles V and ∇V) required to obtain an estimator \widehat{Z} of Z such that with $\Omega(1)$ probability, the relative error is within ε :

$$\Pr \left(\left| \frac{\widehat{Z}}{Z} - 1 \right| \leq \varepsilon \right) \geq \frac{3}{4}. \quad (5)$$

Remark. We make two remarks regarding this criterion. First, similar to how taking the mean of i.i.d. estimates reduces variance, we show in Lem. 11 that the probability above can be boosted to any $1 - \zeta$ using the median trick: obtaining $O\left(\log \frac{1}{\zeta}\right)$ i.i.d. estimates satisfying Eq. (5) and taking their median. Therefore, we focus on the task of obtaining a single estimate satisfying Eq. (5) hereafter. Second, Eq. (5) also allows us to quantify the complexity of estimating the free energy $F = -\log Z$, which is often of greater interest in statistical mechanics than the partition function Z . We show in App. G that estimating Z with $O(\varepsilon)$ relative error and estimating F with $O(\varepsilon)$ absolute error share the same complexity up to constants. Further discussion of this guarantee, including a literature review and the comparison with bias and variance, is deferred to App. G.

Recall that the rationale behind annealing involves a gradual transition from π_0 , a simple distribution that is easy to sample from and estimate the normalizing constant, to $\pi_1 = \pi$, the more complicated target distribution. Throughout this paper, we define a curve of probability measures $(\pi_\theta = \frac{1}{Z_\theta} e^{-V_\theta})_{\theta \in [0, 1]}$, where $V_1 = V$ is the potential of π , and the normalizing constant $Z_1 = Z$ is what we need to estimate. We do not specify the exact form of this curve now, but only introduce the following mild regularity assumption on the curve, as assumed in classical textbooks such as Ambrosio et al. (2008; 2021); Santambrogio (2015):

Assumption 1. The potential $[0, 1] \times \mathbb{R}^d \ni (\theta, x) \mapsto V_\theta(x) \in \mathbb{R}$ is jointly C^1 , 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:

Assumption 2. V is β -smooth, and there exists x_* , with $\|x_*\| =: R \lesssim \frac{1}{\sqrt{\beta}}$ such that $\nabla V(x_*) = 0$. Moreover, let $m := \sqrt{\mathbb{E}_\pi \|\cdot\|^2} < +\infty$.

Remark. Finding a global minimum of (possibly non-convex) V is challenging, but it is always feasible to find some x_+ close to a stationary point x_* using optimization algorithms (e.g., Allen-Zhu & Li (2018)) within negligible cost compared with the complexity for normalizing constant estimation. By considering the translated distribution $\pi(\cdot - x_+)$, we can assume the existence of a stationary point near 0. The assumption $R \lesssim \frac{1}{\sqrt{\beta}}$ is for technical purposes in our proof.

Equipped with this foundational setup, we now proceed to introduce the annealed LD and annealed LMC algorithms, and establish an analysis for JE and AIS.

4 ANALYSIS OF THE JARZYNSKI EQUALITY

To elucidate how annealing works in the task of normalizing constant estimation, we first consider **annealed Langevin diffusion (ALD)**, which runs LD with a dynamically changing target distribution. We introduce a reparameterized curve $(\tilde{\pi}_t = \pi_{\frac{t}{T}})_{t \in [0, T]}$ for some large T to be determined later, and define the ALD as the following SDE:

$$dX_t = \nabla \log \tilde{\pi}_t(X_t) dt + \sqrt{2} dB_t, \quad t \in [0, T]; \quad X_0 \sim \tilde{\pi}_0. \quad (6)$$

The following Jarzynski equality provides a connection between the work functional and the free energy difference, which naturally yields a method for normalizing constant estimation.

Theorem 1 (Jarzynski equality (Jarzynski, 1997)). *Let \mathbb{P}^\rightarrow be the path measure of Eq. (6), and define the work functional W and the free energy difference ΔF as*

$$W(X) := \frac{1}{T} \int_0^T \partial_\theta V_\theta|_{\theta=\frac{t}{T}}(X_t) dt, \quad \Delta F := -\log \frac{Z_1}{Z_0}.$$

Then we have the following relation: $\mathbb{E}_{\mathbb{P}^\rightarrow} e^{-W} = e^{-\Delta F}$.

Below, we sketch the proof from Vargas et al. (2024, Prop. 3.3), which offers a crucial aspect for our analysis. The complete proof is detailed in App. C.1.

Sketch of Proof. Let \mathbb{P}^\leftarrow be the path measure of the following backward SDE:

$$dX_t = -\nabla \log \tilde{\pi}_t(X_t) dt + \sqrt{2} dB_t^\leftarrow, \quad t \in [0, T]; \quad X_T \sim \tilde{\pi}_T. \quad (7)$$

Leveraging Girsanov theorem (Lem. 1) and Itô's formula, one can establish the following identity of the RN derivative, known as the *Crooks fluctuation theorem* (Crooks, 1998; 1999):

$$\log \frac{d\mathbb{P}^\rightarrow}{d\mathbb{P}^\leftarrow}(X) = -\int_0^T (\partial_t \log \tilde{\pi}_t)(X_t) dt = W(X) - \Delta F, \quad \text{a.s. } X \sim \mathbb{P}^\rightarrow, \quad (8)$$

which directly implies JE by the identity $\mathbb{E}_{\mathbb{P}^\rightarrow} \frac{d\mathbb{P}^\leftarrow}{d\mathbb{P}^\rightarrow} = 1$.

Assume for the moment that (i) Z_0 is known, (ii) we can exactly simulate Eq. (6), and (iii) we can calculate the work functional $W(X)$ given any continuous trajectory X . According to Thm. 1, $\hat{Z} := Z_0 e^{-W(X)}$ with $X \sim \mathbb{P}^\rightarrow$ is an unbiased estimator of $Z = Z_0 e^{-\Delta F}$. We establish an upper bound on the time T required to run the ALD in order to satisfy the accuracy criterion Eq. (5) in the following theorem, whose proof is detailed in App. C.2.

Theorem 2. *Under Assump. 1, it suffices to choose $T = \frac{32A}{\varepsilon^2}$ to obtain $\Pr\left(\left|\frac{\hat{Z}}{Z} - 1\right| \leq \varepsilon\right) \geq \frac{3}{4}$.*

To illustrate the proof idea of Thm. 2, note that while the ALD (Eq. (6)) targets the distribution $\tilde{\pi}_t$ at time t , there is always a lag between $\tilde{\pi}_t$ and the actual law of X_t . Similarly, the backward SDE (Eq. (7)) can also be seen as a time-reversed ALD which targets $\tilde{\pi}_t$ at time t , and the same lag exists. This lag turns out to be the source of the error in the estimator \hat{Z} .

To alleviate the issue of high variance in estimating free energy differences, Vaikuntanathan & Jarzynski (2008) proposed adding a compensatory drift term $v_t(X_t)$ to the ALD (Eq. (6)). Ideally, the optimal choice would eliminate the lag entirely, ensuring $X_t \sim \tilde{\pi}_t$ for all $t \in [0, T]$. Inspired by this, we compare the path measure of ALD \mathbb{P}^\rightarrow to the SDE having the perfect compensatory drift term, whose path measure \mathbb{P} has marginal distribution $\tilde{\pi}_t$ at time t . To make possible the perfect match, v_t must satisfy the Fokker-Planck equation. The Girsanov theorem (Lem. 1) enables the computation of $\text{KL}(\mathbb{P} \parallel \mathbb{P}^\rightarrow)$ and $\text{KL}(\mathbb{P} \parallel \mathbb{P}^\leftarrow)$, which are related to $\|v_t\|_{L^2(\tilde{\pi}_t)}^2$. Finally, among all admissible drift terms v_t , Lem. 2 suggests the optimal choice of $v_t^* = \lim_{\delta \rightarrow 0} \frac{T_{\tilde{\pi}_t \rightarrow \tilde{\pi}_{t+\delta}} - \text{id}}{\delta}$ to minimize this norm, thereby leading to the metric derivative $|\dot{\tilde{\pi}}|_t$ and the action \mathcal{A} . Through this approach, we derive a bound not explicitly relying on isoperimetric assumptions.

A similar connection between free energy and action integral was discovered in stochastic thermodynamics (Sekimoto, 2010; Seifert, 2012), one paradigm for non-equilibrium thermodynamics. By the second law of thermodynamics, the averaged dissipated work, defined as the averaged work minus the free energy difference, i.e., $\mathcal{W}_{\text{diss}} := \mathcal{W} - \Delta F := \mathbb{E}_{\mathbb{P}^\rightarrow} W - \Delta F$, is non-negative. When the

underlying process is modeled by an overdamped LD, $\mathcal{W}_{\text{diss}}$ can be quantified by an action integral divided by the length of the process (Aurell et al., 2011; Chen et al., 2020). This follows from the observation that $\mathcal{W}_{\text{diss}} = \text{KL}(\mathbb{P}^{\rightarrow} \parallel \mathbb{P}^{\leftarrow})$ and then a similar argument to that above. This connection provides a finer description of the second law of thermodynamics (Aurell et al., 2012) over a finite time horizon. Finally, we also observe that our bound aligns with the $O\left(\frac{1}{T}\right)$ decay rate of the variance of the work in Mazonka & Jarzynski (1999) (see also Lelièvre et al. (2010, Chap. 4.1.4)), computed when the curve consists of Gaussian distributions with linearly varying means.

5 ANALYSIS OF THE ANNEALED IMPORTANCE SAMPLING

In practice, it is not feasible to simulate the ALD precisely, nor is it possible to evaluate the exact value of the work $W(X)$. Therefore, discretization and approximation are required. To address this, we first outline the following annealed importance sampling (AIS) equality akin to JE.

Theorem 3 (Annealed importance sampling equality (Neal, 2001)). *Suppose we have probability distributions $\pi_\ell = \frac{1}{Z_\ell} f_\ell$, $\ell \in \llbracket 0, M \rrbracket$ and transition kernels $F_\ell(x, \cdot)$, $\ell \in \llbracket 1, M \rrbracket$, and assume that each π_ℓ is an invariant distribution of F_ℓ , $\ell \in \llbracket 1, M \rrbracket$. Define the path measure*

$$\mathbb{P}^{\rightarrow}(x_{0:M}) = \pi_0(x_0) \prod_{\ell=1}^M F_\ell(x_{\ell-1}, x_\ell). \quad (9)$$

Then the same relation between the work function W and free energy difference ΔF holds:

$$\mathbb{E}_{\mathbb{P}^{\rightarrow}} e^{-W} = e^{-\Delta F}, \quad \text{where } W(x_{0:M}) := \log \prod_{\ell=0}^{M-1} \frac{f_\ell(x_\ell)}{f_{\ell+1}(x_\ell)} \text{ and } \Delta F := -\log \frac{Z_M}{Z_0}.$$

Proof. Since π_ℓ is invariant for F_ℓ , the following backward transition kernels are well-defined:

$$B_\ell(x, x') = \frac{\pi_\ell(x')}{\pi_\ell(x)} F_\ell(x', x), \quad \ell \in \llbracket 1, M \rrbracket.$$

By applying these backward transition kernels sequentially, we define the backward path measure

$$\mathbb{P}^{\leftarrow}(x_{0:M}) = \pi_M(x_M) \prod_{\ell=1}^M B_\ell(x_\ell, x_{\ell-1}). \quad (10)$$

It can be easily demonstrated, as in Eq. (8), that $\log \frac{d\mathbb{P}^{\rightarrow}}{d\mathbb{P}^{\leftarrow}}(x_{0:M}) = W(x_{0:M}) - \Delta F$. Consequently, the identity $\mathbb{E}_{\mathbb{P}^{\rightarrow}} \frac{d\mathbb{P}^{\leftarrow}}{d\mathbb{P}^{\rightarrow}} = 1$ implies the desired equality. \square

To study non-asymptotic complexity guarantees, we focus on a widely used curve in theoretical analysis (Brosse et al., 2018; Ge et al., 2020), which we refer to as the *geometric interpolation*²:

$$\pi_\theta = \frac{1}{Z_\theta} f_\theta = \frac{1}{Z_\theta} \exp\left(-V - \frac{\lambda(\theta)}{2} \|\cdot\|^2\right), \quad \theta \in [0, 1], \quad (11)$$

where $\lambda(\cdot)$ is a decreasing function with $\lambda(0) = 2\beta$ and $\lambda(1) = 0$, referred to as the *annealing schedule*. With this choice of $\lambda(0)$, by Assump. 2, the potential of π_0 is β -strongly-convex and 3β -smooth, making sampling and normalizing constant estimation relatively easy. To estimate Z_0 ,

we use the TI algorithm from Ge et al. (2020), which requires $\tilde{O}\left(\frac{d^{\frac{3}{2}}}{\varepsilon^2}\right)$ gradient oracle calls. In a nutshell, TI is an equilibrium method that constructs a series of intermediate distributions and estimates adjacent normalizing constant ratios via expectation under these intermediate distributions, realized through MCMC sampling from each intermediate distribution. As TI is peripheral to our primary focus, we defer its full description and complexity analysis to App. H.1 and Lem. 6.

²Eq. (11) differs slightly from a widely used curve in applications (Gelman & Meng, 1998; Neal, 2001): $\pi_\theta \propto \pi^{1-\lambda(\theta)} \phi^{\lambda(\theta)}$, where ϕ is a prior distribution (typically Gaussian). We refer to both as *geometric interpolation*.

Given the curve Eq. (11), we introduce discrete time points $0 = \theta_0 < \theta_1 < \dots < \theta_M = 1$ to be specified later, and adopt the framework outlined in Thm. 3 by setting $\pi_\ell = \frac{1}{Z_\ell} f_\ell$ to correspond to $\pi_{\theta_\ell} = \frac{1}{Z_{\theta_\ell}} f_{\theta_\ell}$, albeit with a slight abuse of notation. To estimate the normalizing constant, we need to sample from the forward path measure \mathbb{P}^\rightarrow , and calculate the work function along the trajectory. Since π_{θ_ℓ} must be an invariant distribution of the transition kernel F_ℓ in \mathbb{P}^\rightarrow , we define F_ℓ via running LD targeting π_{θ_ℓ} for a short time T_ℓ , i.e., $F_\ell(x, \cdot)$ is given by the law of X_{T_ℓ} in the following SDE initialized at $X_0 = x$:

$$dX_t = \nabla \log \pi_{\theta_\ell}(X_t) dt + \sqrt{2} dB_t, \quad t \in [0, T_\ell]. \quad (12)$$

In this setting, AIS can be interpreted as a discretized version of JE (Lelièvre et al., 2010, Remark 4.5). However, in practice, exact samples from π_0 are often unavailable, and the simulation of LD cannot be performed perfectly. To capture these practical considerations, we define the following sampling path measure:

$$\widehat{\mathbb{P}}^\rightarrow(x_{0:M}) = \widehat{\pi}_0(x_0) \prod_{\ell=1}^M \widehat{F}_\ell(x_{\ell-1}, x_\ell), \quad (13)$$

where $\widehat{\pi}_0$ is the law of an approximate sample from π_0 , and the transition kernel \widehat{F}_ℓ is a discretization of the LD in F_ℓ , defined as running *one step* of **annealed Langevin Monte Carlo (ALMC)** using the exponential integrator discretization scheme (Zhang & Chen, 2023; Zhang et al., 2023b;a) with step size T_ℓ . Formally, $\widehat{F}_\ell(x, \cdot)$ is the law of X_{T_ℓ} in the following SDE initialized at $X_0 = x$:

$$dX_t = - \left(\nabla V(X_0) + \lambda \left(\theta_{\ell-1} + \frac{t}{T_\ell} (\theta_\ell - \theta_{\ell-1}) \right) X_t \right) dt + \sqrt{2} dB_t, \quad t \in [0, T_\ell]. \quad (14)$$

Here, instead of simply setting \widehat{F}_ℓ as one step of LMC targeting π_{θ_ℓ} , the dynamically changing $\lambda(\cdot)$ helps reduce the discretization error, as will be shown in our proof. Furthermore, with a sufficiently small step size, the overall discretization error can also be minimized, motivating us to apply just one update step in each transition kernel.

We refer readers to Line 18 for a summary of the detailed implementation of our proposed AIS algorithm, including the TI procedure and the update rules in Eq. (14). The following theorem delineates the oracle complexity of the algorithm required to obtain an estimate \widehat{Z} meeting the desired accuracy criterion (Eq. (5)), whose detailed proof can be located in App. D.

Theorem 4. *Let \widehat{Z} be the AIS estimator described as in Line 18, i.e., $\widehat{Z} := \widehat{Z}_0 e^{-W(x_{0:M})}$ where \widehat{Z}_0 is estimated by TI and $x_{0:M} \sim \widehat{\mathbb{P}}^\rightarrow$. Under Assumps. 1 and 2, consider the annealing schedule $\lambda(\theta) = 2\beta(1-\theta)^r$ for some $1 \leq r \lesssim 1$. Use \mathcal{A}_r to denote the action of $(\pi_\theta)_{\theta \in [0,1]}$ to emphasize the dependence on r . Then, the oracle complexity for obtaining an estimate \widehat{Z} that satisfies the criterion $\Pr \left(\left| \frac{\widehat{Z}}{Z} - 1 \right| \leq \varepsilon \right) \geq \frac{3}{4}$ is*

$$\widetilde{O} \left(\frac{d^{\frac{3}{2}}}{\varepsilon^2} \vee \frac{m\beta\mathcal{A}_r^{\frac{1}{2}}}{\varepsilon^2} \vee \frac{d\beta^2\mathcal{A}_r^2}{\varepsilon^4} \right). \quad (15)$$

Our proposed algorithm consists of two phases: first, estimating Z_0 by TI, which is provably efficient for well-conditioned distributions, and second, estimating Z by AIS, which is better suited for handling non-log-concave distributions. The three terms in Eq. (15) arise from (i) ensuring the accuracy of \widehat{Z}_0 , (ii) controlling $\text{KL}(\mathbb{P} \parallel \mathbb{P}^\leftarrow)$, and (iii) controlling $\text{KL}(\mathbb{P} \parallel \widehat{\mathbb{P}}^\rightarrow)$, respectively, as discussed in 2. above. Due to the non-log-concavity of π , the action \mathcal{A} is typically large, making (iii), the cost for controlling the discretization error, the dominant complexity. Finally, the ε -dependence can be interpreted as the total duration $T = \Theta \left(\frac{1}{\varepsilon^2} \right)$ required for the continuous dynamics to converge (as in Thm. 2) divided by the step size $\widetilde{\Theta}(\varepsilon^2)$ to control the discretization error.

6 NORMALIZING CONSTANT ESTIMATION VIA REVERSE DIFFUSION SAMPLER

From the analysis of JE and AIS (Thms. 2 and 4), the choice of the interpolation curve $(\pi_\theta)_{\theta \in [0,1]}$ is crucial for the complexity of AIS. The geometric interpolation (Eq. (11)) is widely adopted in

practice due to the availability of closed-form scores of the intermediate distributions π_θ . For certain structured non-log-concave distributions, the associated action is polynomial in the problem parameters, enabling efficient AIS. For instance, Guo et al. (2025, Ex. 2) analyzed a Gaussian mixture target distribution with identical covariance, means having the *same* norm, and arbitrary weights. However, for general target distributions, the action of the related curve can grow prohibitively large. To illustrate this, we establish an exponential lower bound on the action of a curve starting from a Gaussian mixture, highlighting the potential inefficiency of AIS under geometric interpolation. Our key technical tool is a closed-form expression of the W_2 distance in \mathbb{R} expressed by the inverse c.d.f.s of the involved distributions (A similar approach was used independently in Chemseddine et al. (2025)). We then lower bound the metric derivative near the target distribution, where the curve changes the most drastically. The proof of this result is detailed in App. E.1.

Proposition 1. *Consider the Gaussian mixture target distribution $\pi = \frac{1}{2}\mathcal{N}(0, 1) + \frac{1}{2}\mathcal{N}(m, 1)$ on \mathbb{R} for some sufficiently large $m \gtrsim 1$, whose potential is $\frac{m^2}{2}$ -smooth. Under the setting in AIS (Thm. 4), define $\pi_\theta(x) \propto \pi(x)e^{-\frac{\lambda(\theta)}{2}x^2}$, $\theta \in [0, 1]$, where $\lambda(\theta) = m^2(1 - \theta)^r$ for some $1 \leq r \lesssim 1$. Then, the action of the curve $(\pi_\theta)_{\theta \in [0, 1]}$ is lower bounded by $\mathcal{A}_r \gtrsim m^4 e^{\frac{m^2}{40}}$.*

Motivated by RDS, we propose leveraging the curve along the OU process in AIS. To support this idea, we first present the following proposition, whose proof is available in App. E.2.

Proposition 2. *Let $\bar{\pi}_t$ be the law of Y_t in the OU process (Eq. (3)) initialized from $Y_0 \sim \pi \propto e^{-V}$, where V is β -smooth and let $m^2 := \mathbb{E}_\pi \|\cdot\|^2 < \infty$. Then, $\int_0^\infty |\bar{\pi}_t|^2 dt \leq d\beta + m^2$.*

This proposition shows that under fairly weak conditions on the target distribution, the action of the curve along the OU process, $(\bar{\pi}_{T-t})_{t \in [0, T]}$, behaves much better than Eq. (11). Hence, our analysis of JE (Thm. 2) suggests that this curve is likely to yield more efficient normalizing constant estimation. Furthermore, recall that in our earlier proof, we introduced a compensatory drift term v_t to eliminate the lag in ALD. The same principle applies here: ensuring X_t precisely following the reference trajectory is advantageous, which results in the time-reversal of OU process (Eq. (4)). Building on this insight, we propose an RDS-based algorithm for normalizing constant estimation, and establish a framework for analyzing its oracle complexity. The proof is in App. E.3.

Theorem 5. *Assume a total time duration T , an early stopping time $\delta \geq 0$, and discrete time points $0 = t_0 < t_1 < \dots < t_N = T - \delta \leq T$. For $t \in [0, T - \delta)$, let t_- denote t_k if $t \in [t_k, t_{k+1})$. Let $s \approx \nabla \log \bar{\pi}$ be a score estimator, and $\phi = \mathcal{N}(0, I)$. Consider the following two SDEs on $[0, T - \delta]$ representing the sampling trajectory and the time-reversed OU process, respectively:*

$$\begin{aligned} \mathbb{Q}^\dagger : \quad dX_t &= (X_t + 2s_{T-t_-}(X_{t_-}))dt + \sqrt{2}dB_t, & X_0 &\sim \phi; \\ \mathbb{Q} : \quad dX_t &= (X_t + 2\nabla \log \bar{\pi}_{T-t}(X_t))dt + \sqrt{2}dB_t, & X_0 &\sim \bar{\pi}_T. \end{aligned} \quad (16)$$

Let $\hat{Z} := e^{-W(X)}$, $X \sim \mathbb{Q}^\dagger$ be the estimator of Z , where the functional $X \mapsto W(X)$ is defined as

$$\log \phi(X_0) + V(X_{T-\delta}) + (T - \delta)d + \int_0^{T-\delta} \left(\|s_{T-t_-}(X_{t_-})\|^2 dt + \sqrt{2} \langle s_{T-t_-}(X_{t_-}), dB_t \rangle \right).$$

Then, to ensure \hat{Z} satisfies Eq. (5), it suffices that $\text{KL}(\mathbb{Q} \parallel \mathbb{Q}^\dagger) \lesssim \varepsilon^2$ and $\text{TV}(\pi, \bar{\pi}_\delta) \lesssim \varepsilon$.

For detailed implementation of this algorithm including the update rule in Eq. (16) and the computation of $W(X)$, see Line 10. To determine the overall complexity, we leverage existing results for RDS (Huang et al., 2024a;b; He et al., 2024; Vacher et al., 2025) to derive the oracle complexity to achieve $\text{KL}(\mathbb{Q} \parallel \mathbb{Q}^\dagger) \lesssim \varepsilon^2$. When early stopping is needed (i.e., $\delta > 0$), we establish in Lem. 8 that choosing $\delta \asymp \frac{\varepsilon^2}{\beta^2 d^2}$ suffices to ensure ε -closeness in TV distance between $\bar{\pi}_\delta$ and π , under weak assumptions similar to Assump. 2. The detailed complexity analysis is deferred to App. E.5.

As discussed, RDS can be viewed as an *optimally compensated* ALD using the OU process as the trajectory. We conclude this section by contrasting these two approaches. On the one hand, analytically-tractable curves such as the geometric interpolation offer closed-form drift terms at all time points, but may exhibit poor action properties (Prop. 1) or bad isoperimetric constants (Chehab et al., 2025), making annealed sampling challenging. On the other hand, alternative curves like the OU process may have better properties in action and isoperimetric constants, but their drift terms, often related to the scores of the intermediate distributions, lack closed-form expressions, and estimating these terms is also non-trivial. This highlights a fundamental trade-off between the complexity of the drift term estimation and the property of the interpolation curve.

REFERENCES

- Michael S Albergo and Eric Vanden-Eijnden. NETS: A non-equilibrium transport sampler. *arXiv preprint arXiv:2410.02711*, 2024.
- Zeyuan Allen-Zhu and Yuanzhi Li. NEON2: Finding local minima via first-order oracles. 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., 2018. URL https://proceedings.neurips.cc/paper_files/paper/2018/file/d4b2aeb2453bdadaa45cbe9882ffefcf-Paper.pdf.
- 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>.
- Brian D.O. Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982. ISSN 0304-4149. doi: 10.1016/0304-4149(82)90051-5. URL <https://www.sciencedirect.com/science/article/pii/0304414982900515>.
- Christophe Andrieu, James Ridgway, and Nick Whiteley. Sampling normalizing constants in high dimensions using inhomogeneous diffusions. *arXiv preprint arXiv:1612.07583*, 2016.
- Mehrnoosh Arrar, Fernando Martín Boubeta, Maria Eugenia Szretter, Mariela Sued, Leonardo Boechi, and Daniela Rodriguez. On the accurate estimation of free energies using the Jarzynski equality. *Journal of Computational Chemistry*, 40(4):688–696, 2019. doi: 10.1002/jcc.25754. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/jcc.25754>.
- Erik Aurell, Carlos Mejía-Monasterio, and Paolo Muratore-Ginanneschi. Optimal protocols and optimal transport in stochastic thermodynamics. *Physical review letters*, 106(25):250601, 2011.
- Erik Aurell, Krzysztof Gawędzki, Carlos Mejía-Monasterio, Roya Mohayaee, and Paolo Muratore-Ginanneschi. Refined second law of thermodynamics for fast random processes. *Journal of statistical physics*, 147:487–505, 2012.
- 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.
- Denis Blessing, Julius Berner, Lorenz Richter, and Gerhard Neumann. Underdamped diffusion bridges with applications to sampling. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=Q1QTxFm0Is>.
- 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>.
- Davide Carbone, Mengjian Hua, Simon Coste, and Eric Vanden-Eijnden. Efficient training of energy-based models using Jarzynski equality. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 52583–52614. Curran Associates, Inc., 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/a4ddb865e0a8ca3cca43fd7387b4b0da-Paper-Conference.pdf.
- Sourav Chatterjee and Persi Diaconis. The sample size required in importance sampling. *The Annals of Applied Probability*, 28(2):1099 – 1135, 2018. doi: 10.1214/17-AAP1326. URL <https://doi.org/10.1214/17-AAP1326>.

- Omar Chehab, Aapo Hyvärinen, and Andrej Risteski. Provable benefits of annealing for estimating normalizing constants: Importance sampling, noise-contrastive estimation, and beyond. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=iWGC0Ns9i>.
- Omar Chehab, Anna Korba, Austin J Stromme, and Adrien Vacher. Provable convergence and limitations of geometric tempering for Langevin dynamics. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=DZcmz9wU0i>.
- Jannis Chemseddine, Christian Wald, Richard Duong, and Gabriele Steidl. Neural sampling from Boltzmann densities: Fisher-Rao curves in the Wasserstein geometry. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=TUvg5uwdeG>.
- Haoxuan Chen and Lexing Ying. Ensemble-based annealed importance sampling. *arXiv preprint arXiv:2401.15645*, 2024.
- Junhua Chen, Lorenz Richter, Julius Berner, Denis Blessing, Gerhard Neumann, and Anima Anandkumar. Sequential controlled Langevin Diffusions. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=dImD2sgy86>.
- Ming-Hui Chen and Qi-Man Shao. On Monte Carlo methods for estimating ratios of normalizing constants. *The Annals of Statistics*, 25(4):1563 – 1594, 1997. doi: 10.1214/aos/1031594732. URL <https://doi.org/10.1214/aos/1031594732>.
- 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 Michele Pavon. On the relation between optimal transport and Schrödinger bridges: A stochastic control viewpoint. *Journal of Optimization Theory and Applications*, 169:671–691, 2016. doi: 10.1007/s10957-015-0803-z.
- 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.
- Yongxin Chen, Tryphon T. Georgiou, and Michele Pavon. Stochastic control liaisons: Richard Sinkhorn meets Gaspard Monge on a Schrödinger bridge. *SIAM Review*, 63(2):249–313, 2021. doi: 10.1137/20M1339982. URL <https://doi.org/10.1137/20M1339982>.
- Xiang Cheng, Niladri S. Chatterji, Peter L. Bartlett, and Michael I. Jordan. Underdamped Langevin MCMC: A non-asymptotic analysis. 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. 300–323. PMLR, 06–09 Jul 2018. URL <https://proceedings.mlr.press/v75/cheng18a.html>.
- Sinho Chewi. *Log-Concave Sampling*. Book draft, in preparation, 2022. 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>.
- Christophe Chipot and Andrew Pohorille (eds.). *Free Energy Calculations: Theory and Applications in Chemistry and Biology*. Springer Series in Chemical Physics. Springer Berlin, Heidelberg, 2007. doi: 10.1007/978-3-540-38448-9.

- Giovanni Conforti and Luca Tamanini. A formula for the time derivative of the entropic cost and applications. *Journal of Functional Analysis*, 280(11):108964, 2021. ISSN 0022-1236. doi: 10.1016/j.jfa.2021.108964. URL <https://www.sciencedirect.com/science/article/pii/S002212362100046X>.
- Ben Cousins and Santosh Vempala. Gaussian cooling and $O^*(n^3)$ algorithms for volume and Gaussian volume. *SIAM Journal on Computing*, 47(3):1237–1273, 2018. doi: 10.1137/15M1054250. URL <https://doi.org/10.1137/15M1054250>.
- Gavin E. Crooks. Nonequilibrium measurements of free energy differences for microscopically reversible Markovian systems. *Journal of Statistical Physics*, 90:1481–1487, 1998. doi: 10.1023/A:1023208217925.
- Gavin E. Crooks. Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences. *Phys. Rev. E*, 60:2721–2726, Sep 1999. doi: 10.1103/PhysRevE.60.2721. URL <https://link.aps.org/doi/10.1103/PhysRevE.60.2721>.
- Pierre Del Moral, Arnaud Doucet, and Ajay Jasra. Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 68(3):411–436, 5 2006.
- Arnaud Doucet, Simon Godsill, and Christophe Andrieu. On sequential Monte Carlo sampling methods for bayesian filtering. *Statistics and computing*, 10:197–208, 2000. doi: 10.1023/A:1008935410038.
- Arnaud Doucet, Will Grathwohl, Alexander G Matthews, and Heiko Strathmann. Score-based diffusion meets annealed importance sampling. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp. 21482–21494. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/86b7128efa3950df7c0f6c0342e6dcc1-Paper-Conference.pdf.
- Ignacia Echeverria and L. Mario Amzel. Estimation of free-energy differences from computed work distributions: An application of Jarzynski’s equality. *The Journal of Physical Chemistry B*, 116(36):10977–11396, 2012. doi: 10.1021/jp300527q.
- 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>.
- Hao Ge and Da-Quan Jiang. Generalized Jarzynski’s equality of inhomogeneous multidimensional diffusion processes. *Journal of Statistical Physics*, 131:675–689, 3 2008. ISSN 1572-9613. doi: 10.1007/s10955-008-9520-4.
- 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>.
- Andrew Gelman and Xiao-Li Meng. Simulating normalizing constants: from importance sampling to bridge sampling to path sampling. *Statistical Science*, 13(2):163 – 185, 1998. doi: 10.1214/ss/1028905934. URL <https://doi.org/10.1214/ss/1028905934>.
- Andrew Gelman, John B. Carlin, Hal S. Stern, and Donald B. Rubin. *Bayesian data analysis*. Chapman and Hall/CRC, 3 edition, 2013.
- 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>.

- Carsten Hartmann and Lorenz Richter. Nonasymptotic bounds for suboptimal importance sampling. *SIAM/ASA Journal on Uncertainty Quantification*, 12(2):309–346, 2024. doi: 10.1137/21M1427760. URL <https://doi.org/10.1137/21M1427760>.
- Carsten Hartmann, Lorenz Richter, Christof Schütte, and Wei Zhang. Variational characterization of free energy: Theory and algorithms. *Entropy*, 19(11), 2017. ISSN 1099-4300. doi: 10.3390/e19110626. URL <https://www.mdpi.com/1099-4300/19/11/626>.
- Carsten Hartmann, Christof Schütte, and Wei Zhang. Jarzynski’s equality, fluctuation theorems, and variance reduction: Mathematical analysis and numerical algorithms. *Journal of Statistical Physics*, 175:1214–1261, 2019. doi: 10.1007/s10955-019-02286-4. URL <https://doi.org/10.1007/s10955-019-02286-4>.
- Ye He, Kevin Rojas, and Molei Tao. Zeroth-order sampling methods for non-log-concave distributions: Alleviating metastability by denoising diffusion. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=X3Aljulsw5>.
- 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>.
- Mark Huber. Approximation algorithms for the normalizing constant of Gibbs distributions. *The Annals of Applied Probability*, 25(2):974 – 985, 2015. doi: 10.1214/14-AAP1015. URL <https://doi.org/10.1214/14-AAP1015>.
- Christopher Jarzynski. Nonequilibrium equality for free energy differences. *Phys. Rev. Lett.*, 78: 2690–2693, Apr 1997. doi: 10.1103/PhysRevLett.78.2690. URL <https://link.aps.org/doi/10.1103/PhysRevLett.78.2690>.
- Ajay Jasra, Kengo Kamatani, Prince Peprah Osei, and Yan Zhou. Multilevel particle filters: normalizing constant estimation. *Statistics and Computing*, 28:47–60, 2018. doi: 10.1007/s11222-016-9715-5.
- Mark R. Jerrum, Leslie G. Valiant, and Vijay V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoretical Computer Science*, 43:169–188, 1986. ISSN 0304-3975. doi: 10.1016/0304-3975(86)90174-X. URL <https://www.sciencedirect.com/science/article/pii/030439758690174X>.
- Ioannis Karatzas and Steven E. Shreve. *Brownian Motion and Stochastic Calculus*. Graduate Texts in Mathematics. Springer New York, NY, 2 edition, 1991. doi: 10.1007/978-1-4612-0949-2.
- John G. Kirkwood. Statistical mechanics of fluid mixtures. *The Journal of Chemical Physics*, 3(5): 300–313, 05 1935. ISSN 0021-9606. doi: 10.1063/1.1749657. URL <https://doi.org/10.1063/1.1749657>.
- Yunbum Kook and Santosh S Vempala. Sampling and integration of logconcave functions by algorithmic diffusion. *arXiv preprint arXiv:2411.13462*, 2024.
- Svetoslav Kostov and Nick Whiteley. An algorithm for approximating the second moment of the normalizing constant estimate from a particle filter. *Methodology and Computing in Applied Probability*, 19:799–818, 2017. doi: 10.1007/s11009-016-9513-8.
- Oswin Krause, Asja Fischer, and Christian Igel. Algorithms for estimating the partition function of restricted Boltzmann machines. *Artificial Intelligence*, 278:103195, 2020. ISSN 0004-3702. doi: 10.1016/j.artint.2019.103195. URL <https://www.sciencedirect.com/science/article/pii/S0004370219301948>.

- C. Le Bris and P.-L. Lions. Existence and uniqueness of solutions to Fokker–Planck type equations with irregular coefficients. *Communications in Partial Differential Equations*, 33(7): 1272–1317, 2008. doi: 10.1080/03605300801970952. URL <https://doi.org/10.1080/03605300801970952>.
- Holden Lee, Jianfeng Lu, and Yixin Tan. Convergence of score-based generative modeling for general data distributions. 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. 946–985. PMLR, 20 Feb–23 Feb 2023. URL <https://proceedings.mlr.press/v201/lee23a.html>.
- Tony Lelièvre, Mathias Rousset, and Gabriel Stoltz. *Free Energy Computations: A Mathematical Perspective*. Imperial College Press, 2010. doi: 10.1142/p579.
- Christian Léonard. A survey of the Schrödinger problem and some of its connections with optimal transport. *Discrete and Continuous Dynamical Systems - Series A*, 34(4):1533–1574, 2014. URL <https://hal.science/hal-00849930>.
- Jianzhu Ma, Jin Peng, Sheng Wang, and Jinbo Xu. Estimating the partition function of graphical models using langevin importance sampling. In Carlos M. Carvalho and Pradeep Ravikumar (eds.), *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics*, volume 31 of *Proceedings of Machine Learning Research*, pp. 433–441, Scottsdale, Arizona, USA, 29 Apr–01 May 2013. PMLR. URL <https://proceedings.mlr.press/v31/ma13a.html>.
- O Mazonka and C Jarzynski. Exactly solvable model illustrating far-from-equilibrium predictions. *arXiv preprint cond-mat/9912121*, 1999.
- Ferran Mazzanti and Enrique Romero. Efficient evaluation of the partition function of RBMs with annealed importance sampling. *arXiv preprint arXiv:2007.11926*, 2020.
- Xiao-Li Meng and Wing Hung Wong. Simulating ratios of normalizing constants via a simple identity: a theoretical exploration. *Statistica Sinica*, 6(4):831–860, 1996. ISSN 10170405, 19968507. URL <http://www.jstor.org/stable/24306045>.
- 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>.
- Edward Nelson. *Dynamical Theories of Brownian Motion*. Princeton University Press, 1967. ISBN 9780691079509. URL <http://www.jstor.org/stable/j.ctv15r57jg>.
- Nikolas Nüsken and Lorenz Richter. Solving high-dimensional Hamilton–Jacobi–Bellman PDEs using neural networks: perspectives from the theory of controlled diffusions and measures on path space. *Partial differential equations and applications*, 2(4):48, 2021. doi: 10.1007/s42985-021-00102-x.
- Andrew Pohorille, Christopher Jarzynski, and Christophe Chipot. Good practices in free-energy calculations. *The Journal of Physical Chemistry B*, 114(32):10235–10253, 2010. ISSN 1520-6106. doi: 10.1021/jp102971x. URL <https://doi.org/10.1021/jp102971x>.
- Yinuo Ren, Haoxuan Chen, Grant M. Rotskoff, and Lexing Ying. How discrete and continuous diffusion meet: Comprehensive analysis of discrete diffusion models via a stochastic integral framework. In *The Thirteenth International Conference on Learning Representations*, 2025a. URL <https://openreview.net/forum?id=6awxwQE182>.
- Yinuo Ren, Haoxuan Chen, Yuchen Zhu, Wei Guo, Yongxin Chen, Grant M. Rotskoff, Molei Tao, and Lexing Ying. Fast solvers for discrete diffusion models: Theory and applications of high-order algorithms. *arXiv preprint arXiv:2502.00234*, 2025b.
- 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=h4pNROs006>.

- Herbert E. Robbins. *An Empirical Bayes Approach to Statistics*, pp. 388–394. Springer New York, New York, NY, 1992. ISBN 978-1-4612-0919-5. doi: 10.1007/978-1-4612-0919-5_26. URL https://doi.org/10.1007/978-1-4612-0919-5_26".
- Hamza Ruzayqat, Neil K Chada, and Ajay Jasra. Multilevel estimation of normalization constants using ensemble Kalman–Bucy filters. *Statistics and Computing*, 32(3):38, 2022.
- Michael E. Sander, Vincent Roulet, Tianlin Liu, and Mathieu Blondel. Joint learning of energy-based models and their partition function. *arXiv preprint arXiv:2501.18528*, 2025.
- Filippo Santambrogio. *Optimal Transport for Applied Mathematicians*. Birkhäuser Cham, 2015. doi: 10.1007/978-3-319-20828-2.
- Christoph Schönle, Marylou Gabrié, Tony Lelièvre, and Gabriel Stoltz. Sampling metastable systems using collective variables and Jarzynski-Crooks paths. *Journal of Computational Physics*, 527:113806, 2025. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2025.113806>. URL <https://www.sciencedirect.com/science/article/pii/S0021999125000890>.
- Udo Seifert. Stochastic thermodynamics, fluctuation theorems and molecular machines. *Reports on progress in physics*, 75(12):126001, 2012.
- Ken Sekimoto. Stochastic energetics, 2010.
- Jascha Sohl-Dickstein and Benjamin J Culpepper. Hamiltonian annealed importance sampling for partition function estimation. Technical report, Redwood Technical Report, 2012. URL <http://arxiv.org/abs/1205.1925>.
- Yang Song and Diederik P Kingma. How to train your energy-based models. *arXiv preprint arXiv:2101.03288*, 2021.
- 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*, 2021. URL <https://openreview.net/forum?id=PxtIG12RRHS>.
- Jingtong Sun, Julius Berner, Lorenz Richter, Marius Zeinhofer, Johannes Müller, Kamyar Aziz-zadenesheli, and Anima Anandkumar. Dynamical measure transport and neural PDE solvers for sampling. *arXiv preprint arXiv:2407.07873*, 2024.
- Saifuddin Syed, Alexandre Bouchard-Côté, Kevin Chern, and Arnaud Doucet. Optimised annealed sequential Monte Carlo samplers. *arXiv preprint arXiv:2408.12057*, 2024.
- 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.
- Adrien Vacher, Omar Chehab, and Anna Korba. Polynomial time sampling from log-smooth distributions in fixed dimension under semi-log-concavity of the forward diffusion with application to strongly dissipative distributions. *arXiv preprint arXiv:2501.00565*, 2025.
- Suriyanarayanan Vaikuntanathan and Christopher Jarzynski. Escorted free energy simulations: Improving convergence by reducing dissipation. *Phys. Rev. Lett.*, 100:190601, May 2008. doi: 10.1103/PhysRevLett.100.190601. URL <https://link.aps.org/doi/10.1103/PhysRevLett.100.190601>.
- 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. *Topics in optimal transportation*, volume 58. American Mathematical Society, 2003.
- 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>.
- Muneki Yasuda and Chako Takahashi. Free energy evaluation using marginalized annealed importance sampling. *Phys. Rev. E*, 106:024127, Aug 2022. doi: 10.1103/PhysRevE.106.024127. URL <https://link.aps.org/doi/10.1103/PhysRevE.106.024127>.
- 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, Jiaming Song, and Yongxin Chen. Improved order analysis and design of exponential integrator for diffusion models sampling. *arXiv preprint arXiv:2308.02157*, 2023a.
- Qinsheng Zhang, Molei Tao, and Yongxin Chen. gDDIM: Generalized denoising diffusion implicit models. In *The Eleventh International Conference on Learning Representations*, 2023b. URL <https://openreview.net/forum?id=1hKE9qjvz->.
- Daniel M. Zuckerman and Thomas B. Woolf. Theory of a systematic computational error in free energy differences. *Phys. Rev. Lett.*, 89:180602, Oct 2002. doi: 10.1103/PhysRevLett.89.180602. URL <https://link.aps.org/doi/10.1103/PhysRevLett.89.180602>.
- Daniel M. Zuckerman and Thomas B. Woolf. Systematic finite-sampling inaccuracy in free energy differences and other nonlinear quantities. *Journal of Statistical Physics*, 114:1303–1323, 3 2004. doi: 10.1023/B:JOSS.0000013961.84860.5b.

CONTENTS

1	Introduction	1
2	Preliminaries	2
2.1	Stochastic Differential Equations and Girsanov Theorem	2
2.2	Wasserstein Distance, Metric Derivative, and Action	3
2.3	Langevin Diffusion and Langevin Monte Carlo	3
2.4	Reverse Diffusion Samplers	3
3	Problem Setting	4
4	Analysis of the Jarzynski Equality	5
5	Analysis of the Annealed Importance Sampling	6
6	Normalizing Constant Estimation via Reverse Diffusion Sampler	7
A	Preliminaries (Continued)	18
B	Pseudo-codes of the Algorithms	20
C	Proofs for Sec. 4	20
C.1	A Complete Proof of Thm. 1	20
C.2	Proof of Thm. 2	21
D	Proof of Thm. 4	23
E	Proofs for Sec. 6	30
E.1	Proof of Prop. 1	30
E.2	Proof of Prop. 2	33
E.3	Proof of Thm. 5	33
E.4	An Upper Bound of the TV Distance along the OU Process	34
E.5	Discussion on the Overall Complexity of RDS	35
F	Supplementary Lemmas	36
G	Review and Discussion on the Error Guarantee (Eq. (5))	38
G.1	Literature Review of Existing Bounds	38
G.2	Equivalence in Complexities for Estimating Z and F	39
G.3	Eq. (5) is Weaker than Bias and Variance	39
G.4	An Upper Bound on the Normalized Variance of \hat{Z} in Jarzynski Equality	39
H	Further Discussion on Related Works	40

H.1	Thermodynamic Integration	41
H.2	Proof of the Second Part of Lem. 6	42
H.3	Path Integral Sampler and Controlled Monte Carlo Diffusion	45
I	Conclusion and Future Work	47

A PRELIMINARIES (CONTINUED)

Notations and definitions. For $a, b \in \mathbb{R}$, let $\llbracket a, b \rrbracket := [a, b] \cap \mathbb{Z}$, $a \wedge b := \min(a, b)$, and $a \vee b := \max(a, b)$. For $a, b > 0$, the notations $a \lesssim b$, $b \gtrsim a$, $a = O(b)$, $b = \Omega(a)$ indicate that $a \leq Cb$ for some constant $C > 0$, and the notations $a \asymp b$, $a = \Theta(b)$ stand for $a \lesssim b \lesssim a$. $\tilde{O}(\cdot)$, $\tilde{\Theta}(\cdot)$ hide logarithmic dependence in $O(\cdot)$, $\Theta(\cdot)$. A function $U \in C^2(\mathbb{R}^d)$ is $\alpha (> 0)$ -strongly-convex if $\nabla^2 U \succeq \alpha I$, and is $\beta (> 0)$ -smooth if $-\beta I \preceq \nabla^2 U \preceq \beta I$. We do not distinguish probability measures on \mathbb{R}^d from their Lebesgue densities. For two probability measures μ, ν , the total-variation (TV) distance is $\text{TV}(\mu, \nu) = \sup_{\text{measurable } A} |\mu(A) - \nu(A)|$, and the Kullback-Leibler (KL) divergence is $\text{KL}(\mu \parallel \nu) = \int \log \frac{d\mu}{d\nu} d\mu$. We call $\mathbb{E}_\mu \|\cdot\|^2$ the second-order moment of μ . Finally, a function $T : \mathbb{R}^d \times \mathbb{R}^d \rightarrow [0, +\infty)$ is a transition kernel if for any x , $T(x, \cdot)$ is a p.d.f.

The theories of backward stochastic integral and the Girsanov theorem are adapted from Vargas et al. (2024). Here, we include relevant results and proofs to ensure a self-contained presentation.

Lemma 3 (Nelson’s relation (Nelson, 1967; Anderson, 1982)). *Given a BM $(B_t)_{t \in [0, T]}$ and its time-reversal $(B_t^{\leftarrow} = B_{T-t})_{t \in [0, T]}$, the following two SDEs*

$$dX_t = a_t(X_t)dt + \sigma dB_t, \quad X_0 \sim p_0; \quad dY_t = b_t(Y_t)dt + \sigma dB_t^{\leftarrow}, \quad Y_T \sim q$$

have the same path measure if and only if

$$q = p_T, \quad \text{and} \quad b_t = a_t - \sigma^2 \nabla \log p_t, \quad \forall t \in [0, T],$$

where p_t is the p.d.f. of X_t .

Proof. The proof is by verifying the Fokker-Planck equation. For X , we have

$$\partial_t p_t = -\nabla \cdot (a_t p_t) + \frac{\sigma^2}{2} \Delta p_t.$$

Let $\star_t^{\leftarrow} := \star_{T-t}$. Then p_t^{\leftarrow} satisfies

$$\partial_t p_t^{\leftarrow} = \nabla \cdot (a_t^{\leftarrow} p_t^{\leftarrow}) - \frac{\sigma^2}{2} \Delta p_t^{\leftarrow} = -\nabla \cdot ((-a_t^{\leftarrow} + \sigma^2 \nabla \log p_t^{\leftarrow}) p_t^{\leftarrow}) + \frac{\sigma^2}{2} \Delta p_t^{\leftarrow},$$

which means $(X_t^{\leftarrow})_{t \in [0, T]}$ has the same path measure as the following SDE:

$$dZ_t = -(a_t^{\leftarrow} - \sigma^2 \nabla \log p_t^{\leftarrow})(Z_t)dt + \sigma dB_t, \quad Z_t \sim p_t^{\leftarrow}.$$

On the other hand, by definition, $(Y_t^{\leftarrow})_{t \in [0, T]}$ satisfies the forward SDE

$$dY_t^{\leftarrow} = -b_t^{\leftarrow}(Y_t^{\leftarrow})dt + \sigma dB_t, \quad Y_0 \sim q,$$

and thus the claim is evident. \square

Definition 1 (Backward stochastic integral). *For two continuous stochastic processes X and Y on $C([0, T]; \mathbb{R}^d)$, the **backward stochastic integral** of Y with respect to X is defined as*

$$\int_0^T \langle Y_t, *dX_t \rangle := \Pr - \lim_{\|\Pi\| \rightarrow 0} \sum_{i=0}^{n-1} \langle Y_{t_{i+1}}, X_{t_{i+1}} - X_{t_i} \rangle,$$

where $\Pi = \{0 = t_0 < t_1 < \dots < t_n = T\}$ is a partition of $[0, T]$, $\|\Pi\| := \max_{i \in \llbracket 1, n \rrbracket} (t_{i+1} - t_i)$, and the convergence is in the probability sense. When both X and Y are continuous semi-martingales, one can equivalently define

$$\int_0^T \langle Y_t, *dX_t \rangle := \int_0^T \langle Y_t, dX_t \rangle + [X, Y]_T, \quad (17)$$

where $[X, Y]$ is the cross quadratic variation process³ of the local martingale parts of X and Y .

³The notation used in Karatzas & Shreve (1991) is $\langle \cdot, \cdot \rangle$. We use square brackets here to avoid conflict with the notation for inner product.

Remark. Although rarely used in practice, the backward stochastic integral is sometimes referred to as the Hänggi-Klimontovich integral in the literature. Recall that the Itô integral is defined as the limit of Riemann sums when the leftmost point of each interval is used, while the Stratonovich integral is based on the midpoint and the backward integral uses the rightmost point. The equivalence in Eq. (17) can be justified in Karatzas & Shreve (1991, Chap. 3.3).

Lemma 4 (Continuation of Lem. 1). *1. If we replace the SDEs in Lem. 1 with*

$$dX_t = a_t(X_t)dt + \sigma dB_t^{\leftarrow}, X_T \sim \mu; \quad dY_t = b_t(Y_t)dt + \sigma dB_t^{\leftarrow}, Y_T \sim \nu,$$

while keeping other assumptions and notations unchanged, then for any trajectory $\xi \in \Omega$,

$$\log \frac{d\mathbb{P}^X}{d\mathbb{P}^Y}(\xi) = \log \frac{d\mu}{d\nu}(\xi_T) + \frac{1}{\sigma^2} \int_0^T \langle a_t(\xi_t) - b_t(\xi_t), *d\xi_t \rangle - \frac{1}{2\sigma^2} \int_0^T (\|a_t(\xi_t)\|^2 - \|b_t(\xi_t)\|^2)dt,$$

and consequently,

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

2. Define the following two SDEs from 0 to T :

$$dX_t = a_t(X_t)dt + \sigma dB_t, X_0 \sim \mu; \quad dY_t = b_t(Y_t)dt + \sigma dB_t^{\leftarrow}, Y_T \sim \nu.$$

Denote the path measures of X and Y as \mathbb{P}^X and \mathbb{P}^Y , respectively. Then for any trajectory $\xi \in \Omega$,

$$\log \frac{d\mathbb{P}^X}{d\mathbb{P}^Y}(\xi) = \log \frac{\mu(\xi_0)}{\nu(\xi_T)} + \frac{1}{\sigma^2} \int_0^T (\langle a_t(\xi_t), d\xi_t \rangle - \langle b_t(\xi_t), *d\xi_t \rangle) - \frac{1}{2\sigma^2} \int_0^T (\|a_t(\xi_t)\|^2 - \|b_t(\xi_t)\|^2)dt.$$

Proof. 1. Let $\star_t^{\leftarrow} := \star_{T-t}$. We know that

$$dX_t^{\leftarrow} = -a_t^{\leftarrow}(X_t^{\leftarrow})dt + \sigma dB_t, X_0^{\leftarrow} \sim \mu; \quad dY_t^{\leftarrow} = -b_t^{\leftarrow}(Y_t^{\leftarrow})dt + \sigma dB_t, Y_0^{\leftarrow} \sim \nu.$$

Let $\mathbb{P}^{X^{\leftarrow}}$ and $\mathbb{P}^{Y^{\leftarrow}}$ be the path measures of X^{\leftarrow} and Y^{\leftarrow} , respectively. From Lem. 1, we know that

$$\log \frac{d\mathbb{P}^{X^{\leftarrow}}}{d\mathbb{P}^{Y^{\leftarrow}}}(\xi) = \log \frac{d\mu}{d\nu}(\xi_0) - \frac{1}{\sigma^2} \int_0^T \langle a_t^{\leftarrow}(\xi_t) - b_t^{\leftarrow}(\xi_t), d\xi_t \rangle - \frac{1}{2\sigma^2} \int_0^T (\|a_t^{\leftarrow}(\xi_t)\|^2 - \|b_t^{\leftarrow}(\xi_t)\|^2)dt.$$

Since $\mathbb{P}^{X^{\leftarrow}}(d\xi) = \Pr(X^{\leftarrow} \in d\xi) = \Pr(X \in d\xi^{\leftarrow}) = \mathbb{P}^X(d\xi^{\leftarrow})$, we obtain

$$\begin{aligned} \log \frac{d\mathbb{P}^X}{d\mathbb{P}^Y}(\xi) &= \log \frac{d\mathbb{P}^{X^{\leftarrow}}}{d\mathbb{P}^{Y^{\leftarrow}}}(\xi^{\leftarrow}) \\ &= \log \frac{d\mu}{d\nu}(\xi_0^{\leftarrow}) - \frac{1}{\sigma^2} \int_0^T \langle a_t^{\leftarrow}(\xi_t^{\leftarrow}) - b_t^{\leftarrow}(\xi_t^{\leftarrow}), d\xi_t^{\leftarrow} \rangle - \frac{1}{2\sigma^2} \int_0^T (\|a_t^{\leftarrow}(\xi_t^{\leftarrow})\|^2 - \|b_t^{\leftarrow}(\xi_t^{\leftarrow})\|^2)dt \\ &= \log \frac{d\mu}{d\nu}(\xi_T) + \frac{1}{\sigma^2} \int_0^T \langle a_t(\xi_t) - b_t(\xi_t), *d\xi_t \rangle - \frac{1}{2\sigma^2} \int_0^T (\|a_t(\xi_t)\|^2 - \|b_t(\xi_t)\|^2)dt. \end{aligned}$$

To justify the last equality, if ξ, η are two continuous stochastic processes, then by definition,

$$\begin{aligned} \int_0^T \langle \xi_t^{\leftarrow}, d\eta_t^{\leftarrow} \rangle &= \Pr - \lim_{\|\Pi\| \rightarrow 0} \sum_{i=0}^{n-1} \langle \xi_{t_{i-1}}^{\leftarrow}, \eta_{t_i}^{\leftarrow} - \eta_{t_{i-1}}^{\leftarrow} \rangle \\ &= \Pr - \lim_{\|\Pi\| \rightarrow 0} \sum_{i=0}^{n-1} \langle \xi_{T-t_{i-1}}, \eta_{T-t_i} - \eta_{T-t_{i-1}} \rangle \\ &= \Pr - \lim_{\|\Pi\| \rightarrow 0} - \sum_{i=0}^{n-1} \langle \xi_{T-t_{i-1}}, \eta_{T-t_{i-1}} - \eta_{T-t_i} \rangle \\ &= - \int_0^T \langle \xi_t, *d\eta_t \rangle. \end{aligned} \tag{18}$$

On the other hand,

$$\int_0^T \xi_t^{\leftarrow} dt = \int_0^T \xi_{T-t} dt = \int_0^T \xi_t dt.$$

Therefore, the equality of RN derivative holds. Plugging in $\xi \leftarrow X$, we have

$$\log \frac{d\mathbb{P}^X}{d\mathbb{P}^Y}(X) = \log \frac{d\mu}{d\nu}(X_T) + \frac{1}{\sigma} \int_0^T \langle a_t(X_t) - b_t(X_t), *dB_t^{\leftarrow} \rangle + \frac{1}{2\sigma^2} \int_0^T \|a_t(X_t) - b_t(X_t)\|^2 dt.$$

To obtain the KL divergence, it suffices to show the expectation of the second term is zero. Let

$$M_t := \int_t^T \langle a_r(X_r) - b_r(X_r), *dB_r^{\leftarrow} \rangle, \quad t \in [0, T].$$

By Eq. (18), we have

$$M_t^{\leftarrow} = - \int_0^t \langle a_r^{\leftarrow}(X_r^{\leftarrow}) - b_r^{\leftarrow}(X_r^{\leftarrow}), dB_r \rangle.$$

Since $dX_t^{\leftarrow} = -a_t^{\leftarrow}(X_t^{\leftarrow})dt + \sigma dB_t$, we conclude that M_t^{\leftarrow} is a (forward) martingale, and thus M is a *backward* martingale with $\mathbb{E} M_t = \mathbb{E} M_{T-t}^{\leftarrow} = 0$.

2. We present a formal proof by considering the process $dZ_t = \sigma dB_t$ and $Z_0 \sim \lambda$, the Lebesgue measure. As a result, formally $Z_t \sim \lambda$ for all t , so it can also be written as $dZ_t = \sigma dB_t^{\leftarrow}$, $Z_T \sim \lambda$. The result follows by applying Lem. 1 to X and Z and **1.** to Y and Z . □

Remark. *The Girsanov theorem requires a technical condition ensuring that a local martingale is a true martingale, typically verified via the Novikov condition (Karatzas & Shreve, 1991, Chap. 3, Cor. 5.13), which can be challenging to establish. However, when only an upper bound of the KL divergence is needed, the approximation argument from Chen et al. (2023, App. B.2) circumvents the verification of the Novikov condition. For additional context, see Chewi (2022, Sec. 3.2). In this paper, we omit these technical details and always assume that the Novikov condition holds.*

Definition 2 (Isoperimetric inequalities). *A probability measure π on \mathbb{R}^d satisfies a **Poincaré inequality (PI)** with constant C , or C -PI, if for all $f \in C_c^\infty(\mathbb{R}^d)$,*

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

*It satisfies a **log-Sobolev inequality (LSI)** with constant C , or C -LSI, if for all $0 \neq f \in C_c^\infty(\mathbb{R}^d)$,*

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

Furthermore, α -strong-log-concavity implies $\frac{1}{\alpha}$ -LSI, and C -LSI implies C -PI (Bakry et al., 2014).

B PSEUDO-CODES OF THE ALGORITHMS

See Lines 10 and 18 for the detailed implementation of the AIS and RDS algorithms, respectively.

C PROOFS FOR SEC. 4

C.1 A COMPLETE PROOF OF THM. 1

Proof. By Girsanov theorem (Lem. 4), we have

$$\log \frac{d\mathbb{P}^{\rightarrow}}{d\mathbb{P}^{\leftarrow}}(\xi) = \log \frac{\tilde{\pi}_0(\xi_0)}{\tilde{\pi}_T(\xi_T)} + \frac{1}{2} \int_0^T (\langle \nabla \log \tilde{\pi}_t(\xi_t), d\xi_t \rangle + \langle \nabla \log \tilde{\pi}_t(\xi_t), *d\xi_t \rangle).$$

We first prove the following result (Vargas et al., 2024, Eq. (15)): if $dx_t = a_t(x_t)dt + \sqrt{2}dB_t$, then

$$\int_0^T \langle a_t(x_t), *dx_t \rangle = \int_0^T \langle a_t(x_t), dx_t \rangle + 2 \int_0^T \text{tr} \nabla a_t(X_t) dt.$$

Proof. Due to Eq. (17), it suffices to calculate $[a(X), X]_T$. By Itô's formula, we have

$$da_t(x_t) = (\partial_t a_t(x_t) + \langle \nabla a_t(x_t), a_t(x_t) \rangle + \Delta a_t(x_t))dt + \sqrt{2} \nabla a_t dB_t,$$

Algorithm 1: Normalizing constant estimation via AIS.

Input: The target distribution $\pi \propto e^{-V}$, smoothness parameter β , total time T ; TI annealing schedule $\lambda_0 > \dots > \lambda_K = 0$; AIS annealing schedule $\lambda(\cdot)$ with $\lambda(0) = 2\beta$, AIS time points $0 = \theta_0 < \dots < \theta_M = 1$.

Output: \widehat{Z} , an estimation of $Z = \int_{\mathbb{R}^d} e^{-V(x)} dx$.

- 1 // Phase 1: estimate Z_0 via TI.
- 2 Define $V_0 := V + \beta \|\cdot\|^2$, $\rho_k := \propto \exp(-V_0 - \frac{\lambda_k}{2} \|\cdot\|^2)$, and $g_k := \exp\left(\frac{\lambda_k - \lambda_{k+1}}{2} \|\cdot\|^2\right)$, for $k \in \llbracket 0, K-1 \rrbracket$;
- 3 Initialize $\widehat{Z}_0 \leftarrow \exp\left(-V_0(0) + \frac{\|\nabla V_0(0)\|^2}{2(3\beta + \lambda_0)}\right) \left(\frac{2\pi}{3\beta + \lambda_0}\right)^{\frac{d}{2}}$;
- 4 **for** $k = 0$ **to** $K-1$ **do**
- 5 Obtain N i.i.d. approximate samples $x_1^{(k)}, \dots, x_N^{(k)}$ from ρ_k (e.g., using LMC or proximal sampler);
- 6 Update $\widehat{Z}_0 \leftarrow \left(\frac{1}{N} \sum_{n=1}^N g_k(X_n^{(k)})\right) \widehat{Z}_0$;
- 7 **end**
- 8 // Phase 2: estimate Z via AIS.
- 9 Approximately sample x_0 from π_0 (e.g., using LMC or proximal sampler);
- 10 Initialize $W \leftarrow -\frac{1}{2}(\lambda(\theta_0) - \lambda(\theta_1))\|x_0\|^2$;
- 11 **for** $\ell = 1$ **to** $M-1$ **do**
- 12 Sample an independent $\xi \sim \mathcal{N}(0, I_d)$;
- 13 Define $\Lambda(t) := \int_0^t \lambda\left(\theta_{\ell-1} + \frac{\tau}{T_\ell}(\theta_\ell - \theta_{\ell-1})\right) d\tau$, where $T_\ell := T(\theta_\ell - \theta_{\ell-1})$;
- 14 Update

$$x_\ell \leftarrow e^{-\Lambda(T_\ell)} x_{\ell-1} - \left(\int_0^{T_\ell} e^{-(\Lambda(T_\ell) - \Lambda(t))} dt\right) \nabla V(x_{\ell-1}) + \left(2 \int_0^{T_\ell} e^{-2(\Lambda(T_\ell) - \Lambda(t))} dt\right)^{\frac{1}{2}} \xi$$

// see Lem. 12 for the derivation
- 15 ;
- 16 Update $W \leftarrow W - \frac{1}{2}(\lambda(\theta_\ell) - \lambda(\theta_{\ell+1}))\|x_\ell\|^2$;
- 17 **end**
- 18 **return** $\widehat{Z} = \widehat{Z}_0 e^{-W}$

and hence

$$[a(X), X]_T = \left[\int_0^\cdot \sqrt{2} \nabla a_t(x_t) dB_t, \int_0^\cdot \sqrt{2} dB_t \right]_T = \text{tr} \int_0^T 2 \nabla a_t(x_t) dt.$$

□

Therefore, for $X \sim \mathbb{P}^\rightarrow$, we have

$$\log \frac{d\mathbb{P}^\rightarrow}{d\mathbb{P}^\leftarrow}(X) = \log \frac{\tilde{\pi}_0(X_0)}{\tilde{\pi}_T(X_T)} + \int_0^T (\langle \nabla \log \tilde{\pi}_t(X_t), dX_t \rangle + \Delta \log \tilde{\pi}_t(X_t) dt).$$

On the other hand, by Itô's formula, we have

$$d \log \tilde{\pi}_t(X_t) = \partial_t \log \tilde{\pi}_t(X_t) + \langle \nabla \log \tilde{\pi}_t(X_t), dX_t \rangle + \Delta \log \tilde{\pi}_t(X_t) dt.$$

Taking the integral, we immediately obtain Eq. (8), and the proof is complete. □

C.2 PROOF OF THM. 2

Proof. The proof builds on the techniques developed in Guo et al. (2025, Thm. 1). We define \mathbb{P} as the path measure of the following SDE:

$$dX_t = (\nabla \log \tilde{\pi}_t + v_t)(X_t) dt + \sqrt{2} dB_t, \quad t \in [0, T]; \quad X_0 \sim \tilde{\pi}_0, \quad (19)$$

Algorithm 2: Normalizing constant estimation via RDS.

Input: The target distribution $\pi \propto e^{-V}$, total time duration T , early stopping time $\delta \geq 0$, time points $0 = t_0 < t_1 < \dots < t_N = T - \delta$; score estimator $s. \approx \nabla \log \bar{\pi}.$

Output: \hat{Z} , an estimation of $Z = \int_{\mathbb{R}^d} e^{-V(x)} dx.$

- 1 Sample $X_0 \sim \mathcal{N}(0, I)$, and initialize $W := -\frac{\|X_0\|^2}{2} - \frac{d}{2} \log 2\pi;$
- 2 **for** $k = 0$ **to** $N - 1$ **do**
- 3 Sample an independent pair of $\begin{pmatrix} \xi_1 \\ \xi_2 \end{pmatrix} \sim \mathcal{N}\left(0, \begin{pmatrix} 1 & \rho_k \\ \rho_k & 1 \end{pmatrix} \otimes I\right)$, where the correlation is $\rho_k = \frac{\sqrt{2}(e^{t_{k+1}-t_k} - 1)}}{\sqrt{(e^{2(t_{k+1}-t_k)} - 1)(t_{k+1}-t_k)}}$, and \otimes stands for the Kronecker product;
- 4 Update $X_{t_{k+1}} \leftarrow e^{t_{k+1}-t_k} X_{t_k} + 2(e^{t_{k+1}-t_k} - 1)s_{T-t_k}(X_{t_k}) + \sqrt{e^{2(t_{k+1}-t_k)} - 1}\xi_1$; // see Lem. 13 for the derivation
- 5 ;
- 6 Update $W \leftarrow W + (t_{k+1} - t_k)\|s_{T-t_k}(X_{t_k})\|^2 + \sqrt{2(t_{k+1} - t_k)} \langle s_{T-t_k}(X_{t_k}), \xi_2 \rangle$; // see Lem. 13 for the derivation
- 7 ;
- 8 **end**
- 9 Update $W \leftarrow W + V(X_{t_N}) + (T - \delta)d;$
- 10 **return** $\hat{Z} = e^{-W}.$

where the vector field $(v_t)_{t \in [0, T]}$ is chosen such that $X_t \sim \tilde{\pi}_t$ under \mathbb{P} for all $t \in [0, T]$. According to the Fokker-Planck equation⁴, $(v_t)_{t \in [0, T]}$ must satisfy the 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 $(v_t)_{t \in [0, T]}$ generates $(\tilde{\pi}_t)_{t \in [0, T]}$. The Nelson's relation (Lem. 3) implies an equivalent definition of \mathbb{P} as the path measure of

$$dX_t = (-\nabla \log \tilde{\pi}_t + v_t)(X_t)dt + \sqrt{2}dB_t^{\leftarrow}, \quad t \in [0, T]; \quad X_T \sim \tilde{\pi}_T.$$

Now we bound the probability of ε relative error:

$$\begin{aligned} \Pr\left(\left|\frac{\hat{Z}}{Z} - 1\right| \geq \varepsilon\right) &= \mathbb{P}^{\rightarrow}\left(\left|\frac{e^{-W}}{e^{-\Delta F}} - 1\right| \geq \varepsilon\right) = \mathbb{P}^{\rightarrow}\left(\left|\frac{d\mathbb{P}^{\leftarrow}}{d\mathbb{P}^{\rightarrow}} - 1\right| \geq \varepsilon\right) \\ &\leq \frac{1}{\varepsilon} \mathbb{E}_{\mathbb{P}^{\rightarrow}}\left|\frac{d\mathbb{P}^{\leftarrow}}{d\mathbb{P}^{\rightarrow}} - 1\right| = \frac{2}{\varepsilon} \text{TV}(\mathbb{P}^{\leftarrow}, \mathbb{P}^{\rightarrow}) \\ &\leq \frac{2}{\varepsilon} (\text{TV}(\mathbb{P}, \mathbb{P}^{\rightarrow}) + \text{TV}(\mathbb{P}, \mathbb{P}^{\leftarrow})) \\ &\leq \frac{\sqrt{2}}{\varepsilon} \left(\sqrt{\text{KL}(\mathbb{P} \parallel \mathbb{P}^{\rightarrow})} + \sqrt{\text{KL}(\mathbb{P} \parallel \mathbb{P}^{\leftarrow})}\right). \end{aligned} \quad (20)$$

In the second line above, we apply Markov inequality along with an equivalent definition of the TV distance: $\text{TV}(\mu, \nu) = \frac{1}{2} \int \left| \frac{d\mu}{d\lambda} - \frac{d\nu}{d\lambda} \right| d\lambda$, where λ is a measure that dominates both μ and ν . The third line follows from the triangle inequality for TV distance, while the final line is a consequence of Pinsker's inequality $\text{KL} \geq 2 \text{TV}^2$.

By Girsanov theorem (Lems. 1 and 4), it is straightforward to see that

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

Leveraging the relation between metric derivative and continuity equation (Lem. 2), among all vector fields $(v_t)_{t \in [0, T]}$ that generate $(\tilde{\pi}_t)_{t \in [0, T]}$, we can choose the one that minimizes $\|v_t\|_{L^2(\tilde{\pi}_t)}$, thereby

⁴We assume the existence of a unique curve of probability measures solving the Fokker-Planck equation given the drift and diffusion terms, guaranteed under mild regularity conditions (Le Bris & Lions, 2008).

making $\|v_t\|_{L^2(\tilde{\pi}_t)} = |\dot{\tilde{\pi}}|_t$, the metric derivative. With the reparameterization $\tilde{\pi}_t = \pi_{t/T}$, we have the following relation by chain rule:

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

Employing the change-of-variable formula leads to

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

Therefore, it suffices to choose $T = \frac{32\mathcal{A}}{\varepsilon^2}$ to make the r.h.s. of Eq. (20) less than $\frac{1}{4}$. \square

D PROOF OF THM. 4

A sketch of proof. We present a high-level proof sketch using Fig. 1. The continuous dynamics, comprising the forward path \mathbb{P}^{\rightarrow} , the backward path \mathbb{P}^{\leftarrow} , and the reference path \mathbb{P} , are depicted as three black curves. To address discretization error, the ℓ -th red (purple) arrow proceeding from left to right represents the transition kernel $\hat{F}_\ell(B_\ell)$, whose composition forms $\hat{\mathbb{P}}^{\rightarrow}(\mathbb{P}^{\leftarrow})$.

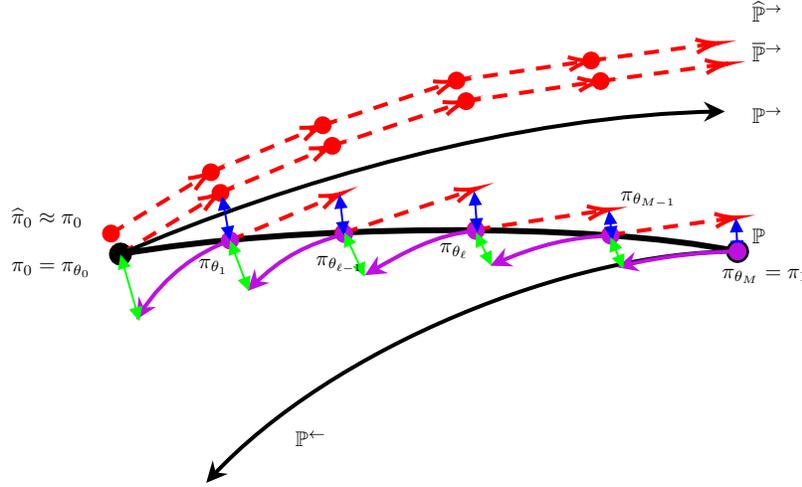


Figure 1: Illustration of the proof idea for Thm. 4.

1. Analogously to the analysis of JE (Thm. 2), define the reference path measure \mathbb{P} with transition kernels F_ℓ^* such that $x_\ell \sim \pi_{\theta_\ell}$. Given the sampling path measure $\hat{\mathbb{P}}^{\rightarrow}$, define $\bar{\mathbb{P}}^{\rightarrow}$ as the version of $\hat{\mathbb{P}}^{\rightarrow}$ without the initialization error, i.e., by replacing $\hat{\pi}_0$ with π_0 in Eq. (13).
2. Show that it suffices to obtain an accurate estimate \hat{Z}_0 and initialization distribution $\hat{\pi}_0$, together with sufficiently small KL divergences $\text{KL}(\mathbb{P} \parallel \mathbb{P}^{\leftarrow})$ and $\text{KL}(\mathbb{P} \parallel \bar{\mathbb{P}}^{\rightarrow})$, which quantify the closeness between the continuous dynamics and the discretization error in implementation, respectively.
3. Using the chain rule, decompose $\text{KL}(\mathbb{P} \parallel \mathbb{P}^{\leftarrow})$ into the sum of KL divergences between each pair of transition kernels \hat{F}_ℓ and F_ℓ^* (i.e., the sum of green “distances”). As in the proof of the convergence of JE (Thm. 2), F_ℓ^* , a transition kernel from $\pi_{\theta_{\ell-1}}$ to π_{θ_ℓ} , is realized by ALD with a compensatory vector field, ensuring the SDE exactly follows the trajectory $(\pi_\theta)_{\theta \in [\theta_{\ell-1}, \theta_\ell]}$. Similarly, by applying the chain rule and Girsanov theorem, we can express $\text{KL}(\mathbb{P} \parallel \bar{\mathbb{P}}^{\rightarrow})$ as the sum of the blue “distances”, allowing for a similar analysis.
4. Finally, derive three necessary conditions on the time steps θ_ℓ to control both $\text{KL}(\mathbb{P} \parallel \mathbb{P}^{\leftarrow})$ and $\text{KL}(\mathbb{P} \parallel \bar{\mathbb{P}}^{\rightarrow})$. Choosing a proper schedule yields the desired complexity bound.

The full proof. With the forward and backward path measures \mathbb{P}^{\rightarrow} and \mathbb{P}^{\leftarrow} defined in Eqs. (9) and (10), we further define the reference path measure

$$\mathbb{P}(x_{0:M}) = \pi_0(x_0) \prod_{\ell=1}^M F_\ell^*(x_{\ell-1}, x_\ell), \quad (21)$$

where F_ℓ^* can be an arbitrary transition kernel transporting $\pi_{\theta_{\ell-1}}$ to π_{θ_ℓ} , i.e., it satisfies

$$\pi_{\theta_\ell}(y) = \int F_\ell^*(x, y) \pi_{\theta_{\ell-1}}(x) dx, \forall y \in \mathbb{R}^d \implies x_\ell \sim \pi_{\theta_\ell}, \forall \ell \in \llbracket 0, M \rrbracket.$$

Define the backward transition kernel of F_ℓ^* as

$$B_\ell^*(x, x') = \frac{\pi_{\theta_{\ell-1}}(x')}{\pi_{\theta_\ell}(x)} F_\ell^*(x', x), \ell \in \llbracket 1, M \rrbracket,$$

which transports π_{θ_ℓ} to $\pi_{\theta_{\ell-1}}$. Equivalently, we can write

$$\mathbb{P}(x_{0:M}) = \pi_1(x_M) \prod_{\ell=1}^M B_\ell^*(x_\ell, x_{\ell-1}).$$

Straightforward calculations yield

$$\begin{aligned} \text{KL}(\mathbb{P} \parallel \mathbb{P}^{\rightarrow}) &= \sum_{\ell=1}^M \mathbb{E}_{\pi_{\theta_{\ell-1}}(x_{\ell-1})} \text{KL}(F_\ell^*(x_{\ell-1}, \cdot) \parallel F_\ell(x_{\ell-1}, \cdot)), \\ \text{KL}(\mathbb{P} \parallel \mathbb{P}^{\leftarrow}) &= \sum_{\ell=1}^M \mathbb{E}_{\pi_{\theta_\ell}(x_\ell)} \text{KL}(B_\ell^*(x_\ell, \cdot) \parallel B_\ell(x_\ell, \cdot)) \\ &= \sum_{\ell=1}^M \text{KL}(\pi_{\theta_{\ell-1}}(x_{\ell-1}) F_\ell^*(x_{\ell-1}, x_\ell) \parallel \pi_{\theta_\ell}(x_{\ell-1}) F_\ell(x_{\ell-1}, x_\ell)) \end{aligned} \quad (22)$$

$$= \text{KL}(\mathbb{P} \parallel \mathbb{P}^{\rightarrow}) + \sum_{\ell=1}^M \text{KL}(\pi_{\theta_{\ell-1}} \parallel \pi_{\theta_\ell}). \quad (23)$$

Also, recall that the sampling path measure $\widehat{\mathbb{P}}^{\rightarrow}$ is defined in Eq. (13) starts at $\widehat{\pi}_0$, the distribution of an approximate sample of π_0 . For convenience, we define the following path measure, which differs from $\widehat{\mathbb{P}}^{\rightarrow}$ only from the initial distribution:

$$\overline{\mathbb{P}}^{\rightarrow}(x_{0:M}) = \pi_0(x_0) \prod_{\ell=1}^M \widehat{F}_\ell(x_{\ell-1}, x_\ell). \quad (24)$$

Equipped with these definitions, we first prove a lemma about a necessary condition for the estimator \widehat{Z} to satisfy the desired accuracy Eq. (5).

Lemma 5. *Define the estimator $\widehat{Z} := \widehat{Z}_0 e^{-W(x_{0:M})}$, where $x_{0:M} \sim \widehat{\mathbb{P}}^{\rightarrow}$, and \widehat{Z}_0 is independent of $x_{0:M}$. To make \widehat{Z} satisfy the criterion Eq. (5), it suffices to meet the following four conditions:*

$$\Pr \left(\left| \frac{\widehat{Z}_0}{Z_0} - 1 \right| \geq \frac{\varepsilon}{8} \right) \leq \frac{1}{8}, \quad (25)$$

$$\text{TV}(\widehat{\pi}_0, \pi_0) \lesssim 1, \quad (26)$$

$$\text{KL}(\mathbb{P} \parallel \mathbb{P}^{\leftarrow}) \lesssim \varepsilon^2, \quad (27)$$

$$\text{KL}(\mathbb{P} \parallel \overline{\mathbb{P}}^{\rightarrow}) \lesssim 1. \quad (28)$$

Proof. Recall that $Z = Z_0 e^{-\Delta F}$. Using Lem. 9, we have

$$\begin{aligned} \Pr \left(\left| \frac{\widehat{Z}}{Z} - 1 \right| \geq \varepsilon \right) &\leq \Pr \left(\left| \log \frac{\widehat{Z}}{Z} \right| \geq \frac{\varepsilon}{2} \right) = \Pr_{x_{0:M} \sim \widehat{\mathbb{P}}^{\rightarrow}} \left(\left| \log \frac{\widehat{Z}_0}{Z_0} + \log \frac{e^{-W(x_{0:M})}}{e^{-\Delta F}} \right| \geq \frac{\varepsilon}{2} \right) \\ &\leq \Pr \left(\left| \log \frac{\widehat{Z}_0}{Z_0} \right| \geq \frac{\varepsilon}{4} \right) + \widehat{\mathbb{P}}^{\rightarrow} \left(\left| \log \frac{e^{-W}}{e^{-\Delta F}} \right| \geq \frac{\varepsilon}{4} \right) \\ &\leq \Pr \left(\left| \frac{\widehat{Z}_0}{Z_0} - 1 \right| \geq \frac{\varepsilon}{8} \right) + \widehat{\mathbb{P}}^{\rightarrow} \left(\left| \frac{e^{-W}}{e^{-\Delta F}} - 1 \right| \geq \frac{\varepsilon}{8} \right). \end{aligned}$$

The first term is $\leq \frac{1}{8}$ if Eq. (25) holds. To bound the second term, using the definition of TV distance and the triangle inequality, we have

$$\begin{aligned} & \widehat{\mathbb{P}}^{\rightarrow} \left(\left| \frac{e^{-W}}{e^{-\Delta F}} - 1 \right| \geq \frac{\varepsilon}{8} \right) \\ & \leq \text{TV}(\widehat{\mathbb{P}}^{\rightarrow}, \mathbb{P}^{\rightarrow}) + \mathbb{P}^{\rightarrow} \left(\left| \frac{e^{-W}}{e^{-\Delta F}} - 1 \right| \geq \frac{\varepsilon}{8} \right) \\ & \leq \text{TV}(\widehat{\mathbb{P}}^{\rightarrow}, \overline{\mathbb{P}}^{\rightarrow}) + \text{TV}(\overline{\mathbb{P}}^{\rightarrow}, \mathbb{P}) + \text{TV}(\mathbb{P}, \mathbb{P}^{\rightarrow}) + \mathbb{P}^{\rightarrow} \left(\left| \frac{d\mathbb{P}^{\leftarrow}}{d\mathbb{P}^{\rightarrow}} - 1 \right| \geq \frac{\varepsilon}{8} \right). \end{aligned}$$

Recall that by definition (Eqs. (13) and (24)), the distributions of $x_{1:M}$ conditional on x_0 are the same under $\widehat{\mathbb{P}}^{\rightarrow}$ and $\overline{\mathbb{P}}^{\rightarrow}$. Hence, $\text{TV}(\widehat{\mathbb{P}}^{\rightarrow}, \overline{\mathbb{P}}^{\rightarrow}) = \text{TV}(\widehat{\pi}_0, \pi_0)$. Applying Pinsker's inequality and leveraging Eq. (20), we have

$$\begin{aligned} & \widehat{\mathbb{P}}^{\rightarrow} \left(\left| \frac{e^{-W}}{e^{-\Delta F}} - 1 \right| \geq \frac{\varepsilon}{8} \right) \\ & \lesssim \text{TV}(\widehat{\pi}_0, \pi_0) + \sqrt{\text{KL}(\mathbb{P} \parallel \overline{\mathbb{P}}^{\rightarrow})} + \sqrt{\text{KL}(\mathbb{P} \parallel \mathbb{P}^{\rightarrow})} + \frac{\sqrt{\text{KL}(\mathbb{P} \parallel \mathbb{P}^{\rightarrow})} + \sqrt{\text{KL}(\mathbb{P} \parallel \mathbb{P}^{\leftarrow})}}{\varepsilon}. \end{aligned}$$

Note that from Eq. (23) we know that $\text{KL}(\mathbb{P} \parallel \mathbb{P}^{\rightarrow}) \leq \text{KL}(\mathbb{P} \parallel \mathbb{P}^{\leftarrow})$, so if Eqs. (26) to (28) hold up to some small enough absolute constants, we can achieve $\widehat{\mathbb{P}}^{\rightarrow} \left(\left| \frac{e^{-W}}{e^{-\Delta F}} - 1 \right| \geq \frac{\varepsilon}{8} \right) \leq \frac{1}{8}$, and therefore $\Pr \left(\left| \frac{\widehat{Z}}{Z} - 1 \right| \geq \varepsilon \right) \leq \frac{1}{4}$. \square

In the next lemma, we show how to sample from π_0 and estimate \widehat{Z}_0 within the desired accuracy.

Lemma 6. 1. *Using LMC initialized at $\mu_0 = \mathcal{N}(0, \beta^{-1}I)$, the oracle complexity for obtaining a sample following a distribution $\widehat{\pi}_0$ that is $O(1)$ -close in TV distance to π_0 is $\widetilde{O}(d)$.*

2. *The oracle complexity for obtaining an estimator \widehat{Z}_0 of Z_0 such that Eq. (25) holds is $\widetilde{O}\left(\frac{d^{\frac{3}{2}}}{\varepsilon^2}\right)$.*

Remark. *Since $R \lesssim \frac{1}{\sqrt{\beta}}$, for both cases the dependence on R is negligible.*

Proof. 1. The bound comes from Vempala & Wibisono (2019, Theorem 2) (see also Chewi (2022, Theorem 4.2.5)). In particular, the bound there depends on $\log \text{KL}(\mu_0 \parallel \pi_0)$. We show that $\text{KL}(\mu_0 \parallel \pi_0)$ has a uniform upper bound over all $R \lesssim 1$. The proof is as follows.

Note that π_0 's potential $V_0 = V + \frac{2\beta}{2} \|\cdot\|^2$ is β -strongly-convex and 3β -smooth. Let x' be its global minimizer, which satisfies $\nabla V(x') + 2\beta x' = 0$. Recall from Assump. 2 that $\nabla V(x_*) = 0$, $\|x_*\| \leq R$. So we have

$$2\beta \|x'\| = \|\nabla V(x') - \nabla V(x_*)\| \leq \beta \|x' - x_*\| \leq \beta (\|x'\| + R) \implies \|x'\| \leq R.$$

Therefore,

$$\begin{aligned} \text{KL}(\mu_0 \parallel \pi_0) &= \mathbb{E}_{\mu_0} [\log \mu_0 - \log \pi_0] \\ &= \mathbb{E}_{\mu_0} \left[-\frac{\beta}{2} \|\cdot\|^2 + \frac{d}{2} \log \frac{\beta}{2\pi} + V_0 + \log Z_0 \right] \\ &= -\frac{d}{2} + \frac{d}{2} \log \frac{d}{2\pi} + \mathbb{E}_{\mu_0} V_0 + \log Z_0. \end{aligned}$$

By strong-convexity and smoothness,

$$\begin{aligned} \mathbb{E}_{\mu_0} V_0 &\leq \mathbb{E}_{\mu_0} \left[V_0(x') + \frac{3\beta}{2} \|\cdot - x'\|^2 \right] = V_0(x') + \frac{3\beta}{2} \left(\frac{d}{\beta} + R^2 \right); \\ \log Z_0 &= \log \int e^{-V_0(x)} dx \leq \log \int \exp \left(-V_0(x') - \frac{\beta}{2} \|x - x'\|^2 \right) dx \\ &= -V_0(x') + \frac{d}{2} \log \frac{\beta}{2\pi}, \end{aligned}$$

so we conclude that

$$\text{KL}(\mu_0 \parallel \pi_0) \leq d + d \log \frac{\beta}{2\pi} + \frac{3\beta R^2}{2}.$$

2. The result is adapted from Ge et al. (2020, Section 3), with two key modifications. First, we relax their assumption that the global minimizer is at zero, requiring instead that the global minimizer x' satisfies $\|x'\| \leq R \lesssim \frac{1}{\sqrt{\beta}}$. Second, we use replace their Metropolis-Hasting adjusted Langevin algorithm (MALA) with the proximal sampler (Fan et al., 2023), which achieves improved dimensional dependence. For completeness, we include a proof sketch in App. H.2 and refer the readers to the original work for full technical details. Our analysis confirms that these relaxations have negligible impact on the final bounds. \square

Next, we study how to satisfy the conditions in Eqs. (27) and (28) while minimizing oracle complexity. Given that we already have an approximate sample from π_0 and an accurate estimate of Z_0 , we proceed to the next step of the AIS algorithm. Since each transition kernel requires one call to the oracle of ∇V , and by plugging in $f_\theta \leftarrow V + \frac{\lambda(\theta)}{2} \|\cdot\|^2$ in AIS (Thm. 3), the work function $W(x_{0:M})$ is independent of V , it follows that the remaining oracle complexity is M . The result is formalized in the following lemma.

Lemma 7. *To minimize the oracle complexity, it suffices to find the minimal M such that there exists a sequence $0 = \theta_0 < \theta_1 < \dots < \theta_M = 1$ satisfying the following three constraints:*

$$\sum_{\ell=1}^M \int_{\theta_{\ell-1}}^{\theta_\ell} (\lambda(\theta) - \lambda(\theta_\ell))^2 d\theta \lesssim \frac{\varepsilon^4}{m^2 \mathcal{A}}, \quad (29)$$

$$\sum_{\ell=1}^M (\theta_\ell - \theta_{\ell-1})^2 \lesssim \frac{\varepsilon^4}{d\beta^2 \mathcal{A}^2}, \quad (30)$$

$$\max_{\ell \in [1, M]} (\theta_\ell - \theta_{\ell-1}) \lesssim \frac{\varepsilon^2}{\beta \mathcal{A}}. \quad (31)$$

Proof. We break down the argument into two steps.

Step 1. We first consider Eq. (27).

Note that when defining the reference path measure \mathbb{P} , the only requirement for the transition kernel F_ℓ^* is that it should transport $\pi_{\theta_{\ell-1}}$ to π_{θ_ℓ} . Our aim is to find the “optimal” F_ℓ^* ’s in order to minimize the sum of KL divergences, which can be viewed as a *static Schrödinger bridge problem* (Léonard, 2014; Chen et al., 2016; 2021). By data-processing inequality,

$$C_\ell := \inf_{F_\ell^*} \text{KL}(\pi_{\theta_{\ell-1}}(x_{\ell-1}) F_\ell^*(x_{\ell-1}, x_\ell) \| \pi_{\theta_\ell}(x_{\ell-1}) F_\ell(x_{\ell-1}, x_\ell)) \leq \inf_{\mathbf{P}^\ell} \text{KL}(\mathbf{P}^\ell \| \mathbf{Q}^\ell),$$

where the infimum is taken among all path measures from 0 to T_ℓ with the marginal constraints $\mathbf{P}_0^\ell = \pi_{\theta_{\ell-1}}$ and $\mathbf{P}_{T_\ell}^\ell = \pi_{\theta_\ell}$; \mathbf{Q}^ℓ is the path measure of Eq. (12) (i.e., LD with target distribution π_{θ_ℓ}) initialized at $X_0 \sim \pi_{\theta_{\ell-1}}$.

For each $\ell \in [1, M]$, define the following interpolation between $\pi_{\theta_{\ell-1}}$ and π_{θ_ℓ} :

$$\mu_t^\ell := \pi_{\theta_{\ell-1} + \frac{t}{T_\ell}(\theta_\ell - \theta_{\ell-1})}, \quad t \in [0, T_\ell].$$

Let \mathbf{P}^ℓ be the path measure of

$$dX_t = (\nabla \log \mu_t^\ell + u_t^\ell)(X_t) dt + \sqrt{2} dB_t, \quad t \in [0, T_\ell]; \quad X_0 \sim \pi_{\theta_{\ell-1}},$$

where the vector field $(u_t^\ell)_{t \in [0, T_\ell]}$ is chosen such that $X_t \sim \mu_t^\ell$ under \mathbf{P}^ℓ , and in particular, the marginal distributions at 0 and T_ℓ are $\pi_{\theta_{\ell-1}}$ and π_{θ_ℓ} , respectively. By verifying the Fokker-Planck equation, the following PDE needs to be satisfied:

$$\partial_t \mu_t^\ell = -\nabla \cdot (\mu_t^\ell (\nabla \log \mu_t^\ell + u_t^\ell)) + \Delta \mu_t^\ell = -\nabla \cdot (\mu_t^\ell u_t^\ell), \quad t \in [0, T_\ell],$$

meaning that $(u_t^\ell)_{t \in [0, T_\ell]}$ generates $(\mu_t^\ell)_{t \in [0, T_\ell]}$. Similar to the proof of JE (Thm. 2), using the relation between metric derivative and continuity equation (Lem. 2), among all vector fields generating

$(\mu_t^\ell)_{t \in [0, T_\ell]}$, we choose $(u_t^\ell)_{t \in [0, T_\ell]}$ to be the a.s.-unique vector field that satisfies $\|u_t^\ell\|_{L^2(\mu_t^\ell)} = |\dot{\mu}^\ell|_t$ for Lebesgue-a.e. $t \in [0, T_\ell]$, which implies (using the chain rule)

$$\begin{aligned} \int_0^{T_\ell} \|u_t^\ell\|_{L^2(\mu_t^\ell)}^2 dt &= \int_0^{T_\ell} |\dot{\mu}^\ell|_t^2 dt \\ &= \int_0^{T_\ell} \left(\frac{\theta_\ell - \theta_{\ell-1}}{T_\ell} |\dot{\pi}|_{\theta_{\ell-1} + \frac{t}{T_\ell}(\theta_\ell - \theta_{\ell-1})} \right)^2 dt = \frac{\theta_\ell - \theta_{\ell-1}}{T_\ell} \int_{\theta_{\ell-1}}^{\theta_\ell} |\dot{\pi}|_\theta^2 d\theta. \end{aligned}$$

By Lem. 3, we can equivalently write \mathbf{P}^ℓ as the path measure of the following backward SDE:

$$dX_t = (-\nabla \log \mu_t^\ell + u_t^\ell)(X_t)dt + \sqrt{2}dB_t^{\leftarrow}, \quad t \in [0, T_\ell]; \quad X_T \sim \pi_{\theta_\ell}.$$

Recall that \mathbf{Q}^ℓ is the path measure of Eq. (12) initialized at $X_0 \sim \pi_{\theta_\ell}$, so $X_t \sim \pi_{\theta_\ell}$ for all $t \in [0, T_\ell]$. By Nelson's relation (Lem. 3), we can equivalently write \mathbf{Q}^ℓ as the path measure of

$$dX_t = -\nabla \log \pi_{\theta_\ell}(X_t)dt + \sqrt{2}dB_t^{\leftarrow}, \quad t \in [0, T_\ell]; \quad X_{T_\ell} \sim \pi_{\theta_\ell}.$$

The purpose of writing these two path measures in the way of backward SDEs is to avoid the extra term of the KL divergence between the initialization distributions $\pi_{\theta_{\ell-1}}$ and π_{θ_ℓ} at time 0 when calculating $\text{KL}(\mathbf{P}^\ell \|\mathbf{Q}^\ell)$. To see this, by Girsanov theorem (Lem. 4), the triangle inequality, and the change-of-variable formula, we have

$$\begin{aligned} C_\ell \leq \text{KL}(\mathbf{P}^\ell \|\mathbf{Q}^\ell) &= \frac{1}{4} \int_0^{T_\ell} \left\| u_t^\ell - \nabla \log \frac{\mu_t^\ell}{\pi_{\theta_\ell}} \right\|_{L^2(\mu_t^\ell)}^2 dt \\ &\lesssim \int_0^{T_\ell} \|u_t^\ell\|_{L^2(\mu_t^\ell)}^2 dt + \int_0^{T_\ell} \left\| \nabla \log \frac{\mu_t^\ell}{\pi_{\theta_\ell}} \right\|_{L^2(\mu_t^\ell)}^2 dt \\ &= \frac{\theta_\ell - \theta_{\ell-1}}{T_\ell} \int_{\theta_{\ell-1}}^{\theta_\ell} |\dot{\pi}|_\theta^2 d\theta + \frac{T_\ell}{\theta_\ell - \theta_{\ell-1}} \int_{\theta_{\ell-1}}^{\theta_\ell} \left\| \nabla \log \frac{\pi_\theta}{\pi_{\theta_\ell}} \right\|_{L^2(\pi_\theta)}^2 d\theta. \end{aligned}$$

Remark. Our bound above is based on a specific interpolation between $\pi_{\theta_{\ell-1}}$ and π_{θ_ℓ} along the curve $(\pi_\theta)_{\theta \in [\theta_{\ell-1}, \theta_\ell]}$. This approach is inspired by, yet slightly differs from, Conforti & Tamanini (2021, Theorem 1.6), where the interpolation is based on the Wasserstein geodesic. As we will demonstrate shortly, our formulation simplifies the analysis of the second term (the Fisher divergence), making it more straightforward to bound.

Now, summing over all $\ell \in \llbracket 1, M \rrbracket$, we can see that in order to ensure $\text{KL}(\mathbb{P} \|\mathbb{P}^{\leftarrow}) \leq \sum_{\ell=1}^M C_\ell \leq \varepsilon^2$, we only need the following two conditions to hold:

$$\sum_{\ell=1}^M \frac{\theta_\ell - \theta_{\ell-1}}{T_\ell} \int_{\theta_{\ell-1}}^{\theta_\ell} |\dot{\pi}|_\theta^2 d\theta \lesssim \varepsilon^2, \quad (32)$$

$$\sum_{\ell=1}^M \frac{T_\ell}{\theta_\ell - \theta_{\ell-1}} \int_{\theta_{\ell-1}}^{\theta_\ell} \left\| \nabla \log \frac{\pi_\theta}{\pi_{\theta_\ell}} \right\|_{L^2(\pi_\theta)}^2 d\theta \lesssim \varepsilon^2. \quad (33)$$

Since $\sum_{\ell=1}^M \int_{\theta_{\ell-1}}^{\theta_\ell} |\dot{\pi}|_\theta^2 d\theta = \mathcal{A}$, it suffices to choose

$$\frac{T_\ell}{\theta_\ell - \theta_{\ell-1}} =: T \asymp \frac{\mathcal{A}}{\varepsilon^2}, \quad \forall \ell \in \llbracket 1, M \rrbracket$$

to make the l.h.s. of Eq. (32) $O(\varepsilon^2)$. Notably, T is the summation over all T_ℓ 's, which has the same order as the total time T for running JE (Eq. (6)) in the continuous scenario, in Thm. 1. Plugging this T_ℓ into the second summation, and noticing that by Eq. (11) and Lem. 15,

$$\left\| \nabla \log \frac{\pi_\theta}{\pi_{\theta'}} \right\|_{L^2(\pi_\theta)}^2 = \mathbb{E}_{x \sim \pi_\theta} \|(\lambda(\theta) - \lambda(\theta'))x\|^2 \leq (\lambda(\theta) - \lambda(\theta'))^2 m^2,$$

we conclude that Eq. (29) implies Eq. (33).

Step 2. Now consider the other constraint Eq. (28). By chain rule and data-processing inequality,

$$\text{KL}(\mathbb{P} \|\widehat{\mathbb{P}}^\rightarrow) = \sum_{\ell=1}^M \text{KL}(\pi_{\theta_{\ell-1}}(x_{\ell-1}) F_\ell^*(x_{\ell-1}, x_\ell) \|\pi_{\theta_{\ell-1}}(x_{\ell-1}) \widehat{F}_\ell(x_{\ell-1}, x_\ell)) \leq \sum_{\ell=1}^M \text{KL}(\mathbf{P}^\ell \|\widehat{\mathbf{Q}}^\ell),$$

where \mathbf{P}^ℓ is the previously defined path measure of the SDE

$$\begin{aligned} dX_t &= (\nabla \log \mu_t^\ell + u_t^\ell)(X_t) dt + \sqrt{2} dB_t \\ &= \left(-\nabla V(X_t) - \lambda \left(\theta_{\ell-1} + \frac{t}{T_\ell} (\theta_\ell - \theta_{\ell-1}) \right) X_t + u_t^\ell(X_t) \right) dt + \sqrt{2} dB_t, \quad t \in [0, T_\ell]; \quad X_0 \sim \pi_{\theta_{\ell-1}}, \end{aligned}$$

and $\widehat{\mathbf{Q}}^\ell$ is the path measure of Eq. (14) initialized at $X_0 \sim \pi_{\theta_{\ell-1}}$, i.e.,

$$dX_t = \left(-\nabla V(X_0) - \lambda \left(\theta_{\ell-1} + \frac{t}{T_\ell} (\theta_\ell - \theta_{\ell-1}) \right) X_t \right) dt + \sqrt{2} dB_t, \quad t \in [0, T_\ell]; \quad X_0 \sim \pi_{\theta_{\ell-1}}.$$

By Lem. 1, triangle inequality, and the smoothness of V , we have

$$\begin{aligned} \text{KL}(\mathbf{P}^\ell \|\widehat{\mathbf{Q}}^\ell) &= \frac{1}{4} \int_0^{T_\ell} \mathbb{E}_{\mathbf{P}^\ell} \|\nabla V(X_t) - \nabla V(X_0) - u_t^\ell(X_t)\|^2 dt \\ &\lesssim \int_0^{T_\ell} \mathbb{E}_{\mathbf{P}^\ell} [\|\nabla V(X_t) - \nabla V(X_0)\|^2 + \|u_t^\ell(X_t)\|^2] dt \\ &\leq \beta^2 \int_0^{T_\ell} \mathbb{E}_{\mathbf{P}^\ell} \|X_t - X_0\|^2 dt + \int_0^{T_\ell} \|u_t^\ell\|_{L^2(\mu_t^\ell)}^2 dt \end{aligned}$$

To bound the first part, note that under \mathbf{P}^ℓ , we have

$$X_t - X_0 = \int_0^t (\nabla \log \mu_\tau^\ell + u_\tau^\ell)(X_\tau) d\tau + \sqrt{2} B_t.$$

Thanks to the fact that $X_t \sim \mu_t^\ell$ under \mathbf{P}^ℓ ,

$$\begin{aligned} \mathbb{E}_{\mathbf{P}^\ell} \|X_t - X_0\|^2 &\lesssim \mathbb{E}_{\mathbf{P}^\ell} \left\| \int_0^t (\nabla \log \mu_\tau^\ell + u_\tau^\ell)(X_\tau) d\tau \right\|^2 + \mathbb{E} \|\sqrt{2} B_t\|^2 \\ &\lesssim t \int_0^t \mathbb{E}_{\mathbf{P}^\ell} \|(\nabla \log \mu_\tau^\ell + u_\tau^\ell)(X_\tau)\|^2 d\tau + dt \\ &\lesssim t \int_0^t \left(\|\nabla \log \mu_\tau^\ell\|_{L^2(\mu_\tau^\ell)}^2 + \|u_\tau^\ell\|_{L^2(\mu_\tau^\ell)}^2 \right) d\tau + dt \\ &\lesssim T_\ell \int_0^{T_\ell} \left(\|\nabla \log \mu_\tau^\ell\|_{L^2(\mu_\tau^\ell)}^2 + \|u_\tau^\ell\|_{L^2(\mu_\tau^\ell)}^2 \right) d\tau + dT_\ell, \quad \forall t \in [0, T_\ell], \end{aligned}$$

where the second inequality follows from Jensen's inequality (Cheng et al., 2018, Sec. 4):

$$\left\| \int_0^t f_\tau d\tau \right\|^2 = t^2 \|\mathbb{E}_{\tau \sim \text{Unif}(0,t)} f_\tau\|^2 \leq t^2 \mathbb{E}_{\tau \sim \text{Unif}(0,t)} \|f_\tau\|^2 = t \int_0^t \|f_\tau\|^2 d\tau.$$

Therefore,

$$\begin{aligned} &\text{KL}(\mathbf{P}^\ell \|\widehat{\mathbf{Q}}^\ell) \\ &\leq \beta^2 \int_0^{T_\ell} \mathbb{E}_{\mathbf{P}^\ell} \|X_t - X_0\|^2 dt + \int_0^{T_\ell} \|u_t^\ell\|_{L^2(\mu_t^\ell)}^2 dt \\ &\leq \beta^2 T_\ell^2 \int_0^{T_\ell} \|\nabla \log \mu_\tau^\ell\|_{L^2(\mu_\tau^\ell)}^2 d\tau + (\beta^2 T_\ell^2 + 1) \int_0^{T_\ell} \|u_\tau^\ell\|_{L^2(\mu_\tau^\ell)}^2 d\tau + d\beta^2 T_\ell^2 \\ &= \beta^2 T_\ell^2 \frac{T_\ell}{\theta_\ell - \theta_{\ell-1}} \int_{\theta_{\ell-1}}^{\theta_\ell} \|\nabla \log \pi_\theta\|_{L^2(\pi_\theta)}^2 d\theta + (\beta^2 T_\ell^2 + 1) \frac{\theta_\ell - \theta_{\ell-1}}{T_\ell} \int_{\theta_{\ell-1}}^{\theta_\ell} |\dot{\pi}|_\theta^2 d\theta + d\beta^2 T_\ell^2. \end{aligned}$$

Recall that the potential of π_θ is $(\beta + \lambda(\theta))$ -smooth. By Lem. 14 and the monotonicity of $\lambda(\cdot)$,

$$\int_{\theta_{\ell-1}}^{\theta_\ell} \|\nabla \log \pi_\theta\|_{L^2(\pi_\theta)}^2 d\theta \leq \int_{\theta_{\ell-1}}^{\theta_\ell} d(\beta + \lambda(\theta)) d\theta \leq d(\theta_\ell - \theta_{\ell-1})(\beta + \lambda(\theta_{\ell-1})).$$

Thus,

$$\begin{aligned} \text{KL}(\mathbb{P} \|\overline{\mathbb{P}}^{\rightarrow}) &\leq \sum_{\ell=1}^M \left(\beta^2 T_\ell^3 d(\beta + \lambda(\theta_{\ell-1})) + (\beta^2 T_\ell^2 + 1) \frac{\theta_\ell - \theta_{\ell-1}}{T_\ell} \int_{\theta_{\ell-1}}^{\theta_\ell} |\dot{\pi}|_\theta^2 d\theta + d\beta^2 T_\ell^2 \right) \\ &= \sum_{\ell=1}^M \left(\beta^2 dT_\ell^2 (T_\ell(\beta + \lambda(\theta_{\ell-1})) + 1) + (\beta^2 T_\ell^2 + 1) \frac{\theta_\ell - \theta_{\ell-1}}{T_\ell} \int_{\theta_{\ell-1}}^{\theta_\ell} |\dot{\pi}|_\theta^2 d\theta \right) \end{aligned}$$

Assume $\max_{\ell \in \llbracket 1, M \rrbracket} T_\ell \lesssim \frac{1}{\beta}$, i.e., Eq. (31). so $\max_{\ell \in \llbracket 1, M \rrbracket} T_\ell(\beta + \lambda(\theta_{\ell-1})) \lesssim 1$, due to $\lambda(\cdot) \leq 2\beta$. We can further simplify the above expression to

$$\begin{aligned} \text{KL}(\mathbb{P} \|\overline{\mathbb{P}}^{\rightarrow}) &\leq \sum_{\ell=1}^M \left(\beta^2 dT_\ell^2 + \frac{\theta_\ell - \theta_{\ell-1}}{T_\ell} \int_{\theta_{\ell-1}}^{\theta_\ell} |\dot{\pi}|_\theta^2 d\theta \right) \lesssim \beta^2 d \left(\sum_{\ell=1}^M T_\ell^2 \right) + \varepsilon^2 \\ &= \beta^2 dT^2 \sum_{\ell=1}^M (\theta_\ell - \theta_{\ell-1})^2 + \varepsilon^2 \lesssim \beta^2 d \frac{\mathcal{A}^2}{\varepsilon^4} \sum_{\ell=1}^M (\theta_\ell - \theta_{\ell-1})^2 + \varepsilon^2. \end{aligned}$$

So Eq. (31) implies that the r.h.s. of the above equation is $O(1)$. □

Finally, we have arrived at the last step of proving Thm. 4, that is to decide the schedule of θ_ℓ 's.

Define $\vartheta_\ell := 1 - \theta_\ell$, $\ell \in \llbracket 0, M \rrbracket$. We consider the annealing schedule $\lambda(\theta) = 2\beta(1 - \theta)^r$ for some $1 \leq r \lesssim 1$, and to emphasize the dependence on r , we use \mathcal{A}_r to represent the action of $(\pi_\theta)_{\theta \in [0, 1]}$. The l.h.s. of Eq. (29) is

$$\begin{aligned} \sum_{\ell=1}^M \int_{\theta_{\ell-1}}^{\theta_\ell} (\lambda(\theta) - \lambda(\theta_\ell))^2 d\theta &\leq \sum_{\ell=1}^M (\theta_\ell - \theta_{\ell-1}) (2\beta(1 - \theta_{\ell-1})^r - 2\beta(1 - \theta_\ell)^r)^2 \\ &= \sum_{\ell=1}^M (\vartheta_{\ell-1} - \vartheta_\ell) (2\beta\vartheta_{\ell-1}^r - 2\beta\vartheta_\ell^r)^2 \\ &\lesssim \beta^2 \sum_{\ell=1}^M (\vartheta_{\ell-1} - \vartheta_\ell) (\vartheta_{\ell-1}^r - \vartheta_\ell^r)^2 \\ &\lesssim \beta^2 \sum_{\ell=1}^M (\vartheta_{\ell-1} - \vartheta_\ell) (\vartheta_{\ell-1} - \vartheta_\ell)^2 = \beta^2 \sum_{\ell=1}^M (\vartheta_{\ell-1} - \vartheta_\ell)^3, \end{aligned}$$

where the last inequality comes from Lem. 10. So to satisfy Eq. (29), it suffices to ensure

$$\sum_{\ell=1}^M (\vartheta_{\ell-1} - \vartheta_\ell)^3 \lesssim \frac{\varepsilon^4}{m^2 \beta^2 \mathcal{A}_r},$$

while Eq. (30) and Eq. (31) are equivalent to

$$\sum_{\ell=1}^M (\vartheta_{\ell-1} - \vartheta_\ell)^2 \lesssim \frac{\varepsilon^4}{d\beta^2 \mathcal{A}_r^2}, \quad \max_{\ell \in \llbracket 1, M \rrbracket} (\vartheta_{\ell-1} - \vartheta_\ell) \lesssim \frac{\varepsilon^2}{\beta \mathcal{A}_r}.$$

Since we are minimizing the total number of oracle calls M , the Hölder's inequality implies that the optimal schedule of ϑ_ℓ 's is an arithmetic sequence, i.e., $\vartheta_\ell = 1 - \frac{\ell}{M}$. We need to ensure

$$\frac{1}{M^2} \lesssim \frac{\varepsilon^4}{m^2 \beta^2 \mathcal{A}_r}, \quad \frac{1}{M} \lesssim \frac{\varepsilon^4}{d\beta^2 \mathcal{A}_r^2}, \quad \frac{1}{M} \lesssim \frac{\varepsilon^2}{\beta \mathcal{A}_r}.$$

So it suffices to choose $\frac{1}{M} \asymp \frac{\varepsilon^2}{m\beta\mathcal{A}_r^{\frac{1}{2}}} \wedge \frac{\varepsilon^4}{d\beta^2\mathcal{A}_r^2}$, which implies the oracle complexity

$$M \asymp \frac{m\beta\mathcal{A}_r^{\frac{1}{2}}}{\varepsilon^2} \vee \frac{d\beta^2\mathcal{A}_r^2}{\varepsilon^4}.$$

□

E PROOFS FOR SEC. 6

E.1 PROOF OF PROP. 1

Proof. The claim of smoothness follows from Guo et al. (2025, Lem. 7). Throughout this proof, let ϕ and Φ denote the p.d.f. and c.d.f. of the standard normal distribution $\mathcal{N}(0, 1)$, respectively. Unless otherwise specified, the integration ranges are assumed to be $(-\infty, \infty)$.

Note that

$$\begin{aligned} \pi(x)e^{-\frac{\lambda}{2}x^2} &\propto \left(e^{-\frac{x^2}{2}} + e^{-\frac{(x-m)^2}{2}} \right) e^{-\frac{\lambda}{2}x^2} \\ &= e^{-\frac{\lambda+1}{2}x^2} + e^{-\frac{\lambda m^2}{2(\lambda+1)}} e^{-\frac{\lambda+1}{2}(x-\frac{m}{\lambda+1})^2} \\ &= \frac{1}{1 + e^{-\frac{\lambda m^2}{2(\lambda+1)}}} \mathcal{N}\left(x \mid 0, \frac{1}{\lambda+1}\right) + \frac{e^{-\frac{\lambda m^2}{2(\lambda+1)}}}{1 + e^{-\frac{\lambda m^2}{2(\lambda+1)}}} \mathcal{N}\left(x \mid \frac{m}{\lambda+1}, \frac{1}{\lambda+1}\right). \end{aligned}$$

Define $S(\theta) := \frac{1}{1+m^2(1-\theta)^r}$, and let

$$\underline{\pi}_s(x) := \pi(x)e^{-\frac{1/s-1}{2}x^2} = w(s)\mathcal{N}(x|0, s) + (1-w(s))\mathcal{N}(x|sm, s),$$

where

$$w(s) = \frac{1}{1 + e^{-(1-s)m^2/2}}, \quad w'(s) = -\frac{e^{-(1-s)m^2/2}m^2/2}{(1 + e^{-(1-s)m^2/2})^2}.$$

By definition, $\pi_\theta = \underline{\pi}_{S(\theta)}$. The p.d.f. of $\underline{\pi}_s$ is

$$f_s(x) = \frac{w(s)}{\sqrt{s}}\phi\left(\frac{x}{\sqrt{s}}\right) + \frac{1-w(s)}{\sqrt{s}}\phi\left(\frac{x-sm}{\sqrt{s}}\right),$$

and the c.d.f. of $\underline{\pi}_s$ is

$$F_s(x) = w(s)\Phi\left(\frac{x}{\sqrt{s}}\right) + (1-w(s))\Phi\left(\frac{x-sm}{\sqrt{s}}\right).$$

We now derive a formula for calculating the metric derivative. From Villani (2003, Thm. 2.18), $W_2^2(\mu, \nu) = \int_0^1 (F_\mu^{-1}(y) - F_\nu^{-1}(y))^2 dy$, where F_μ, F_ν stand for the c.d.f.s of μ, ν . Assuming regularity conditions hold, we have

$$\lim_{\delta \rightarrow 0} \frac{W_2^2(\underline{\pi}_s, \underline{\pi}_{s+\delta})}{\delta^2} = \lim_{\delta \rightarrow 0} \int_0^1 \left(\frac{F_{s+\delta}^{-1}(y) - F_s^{-1}(y)}{\delta} \right)^2 dy = \int_0^1 (\partial_s F_s^{-1}(y))^2 dy.$$

Consider change of variable $y = F_s(x)$, then $\frac{dy}{dx} = f_s(x)$. As $x = F_s^{-1}(y)$, $(F_s^{-1})'(y) = \frac{dx}{dy} = \frac{1}{f_s(x)}$. Taking derivation of s on both sides of the equation $x = F_s^{-1}(F_s(x))$ yields

$$\begin{aligned} 0 &= \partial_s F_s^{-1}(F_s(x)) + (F_s^{-1})'(F_s(x))\partial_s F_s(x) \\ &= \partial_s F_s^{-1}(y) + \frac{1}{f_s(x)}\partial_s F_s(x). \end{aligned}$$

Therefore,

$$\int_0^1 (\partial_s F_s^{-1}(y))^2 dy = \int \left(\frac{\partial_s F_s(x)}{f_s(x)} \right)^2 f_s(x) dx = \int \frac{(\partial_s F_s(x))^2}{f_s(x)} dx.$$

Consider the interval $x \in [\frac{m}{2} - 0.1, \frac{m}{2} + 0.1]$, and fix the range of s to be $[0.9, 0.99]$. We have

$$\begin{cases} 1 - w(s) = \frac{1}{1 + e^{(1-s)m^2/2}} \asymp \frac{1}{e^{(1-s)m^2/2}}, & \forall m \gtrsim 1 \\ -w'(s) = \frac{e^{(1-s)m^2/2} m^2/2}{(1 + e^{(1-s)m^2/2})^2} \asymp \frac{m^2}{e^{(1-s)m^2/2}}, & \forall m \gtrsim 1 \end{cases}$$

First consider upper bounding $f_s(x)$. We have the following two bounds:

$$\begin{aligned} \frac{w(s)}{\sqrt{s}} \phi\left(\frac{x}{\sqrt{s}}\right) &\lesssim e^{-\frac{x^2}{2s}} \leq e^{-\frac{(m/2-0.1)^2}{2 \times 0.99}} \leq e^{-\frac{m^2}{8}}, \quad \forall m \gtrsim 1, \\ \frac{1-w(s)}{\sqrt{s}} \phi\left(\frac{x-sm}{\sqrt{s}}\right) &\lesssim \frac{1}{e^{(1-s)m^2/2}} e^{-\frac{(sm-x)^2}{2s}} = \exp\left(-\frac{1}{2} \left[\frac{(sm-x)^2}{s} + (1-s)m^2 \right]\right). \end{aligned}$$

The term in the square brackets above is

$$\begin{aligned} \frac{(sm-x)^2}{s} + (1-s)m^2 &\geq \frac{1}{s} \left(sm - \frac{m}{2} - 0.1 \right)^2 + (1-s)m^2 \\ &= \frac{m^2}{4s} - 0.2 \left(1 - \frac{1}{2s} \right) m + \frac{0.01}{s} \\ &\geq \frac{m^2}{4 \times 0.99} - 0.1m + 0.1 \geq \frac{m^2}{4}, \quad \forall m \gtrsim 1. \end{aligned}$$

Hence, we conclude that $f_s(x) \lesssim e^{-\frac{m^2}{8}}$.

Next, we consider lower bounding the term $(\partial_s F_s(x))^2$. Note that

$$\begin{aligned} -\partial_s F_s(x) &= -w'(s) \left(\Phi\left(\frac{x}{\sqrt{s}}\right) - \Phi\left(\frac{x-sm}{\sqrt{s}}\right) \right) \\ &\quad + w(s) \phi\left(\frac{x}{\sqrt{s}}\right) \frac{x}{2s^{\frac{3}{2}}} + (1-w(s)) \phi\left(\frac{x-sm}{\sqrt{s}}\right) \left(\frac{x}{2s^{\frac{3}{2}}} + \frac{m}{2s^{\frac{1}{2}}} \right). \end{aligned}$$

As $x \in [\frac{m}{2} - 0.1, \frac{m}{2} + 0.1]$ and $s \in [0.9, 0.99]$, all these three terms are positive. We only focus on the first term. Note the following two bounds:

$$\begin{cases} \Phi\left(\frac{x}{\sqrt{s}}\right) \geq \Phi\left(\frac{m}{2} - 0.1\right) \geq \frac{3}{4}, & \forall m \gtrsim 1, \\ \Phi\left(\frac{x-sm}{\sqrt{s}}\right) \leq \Phi\left(\frac{m/2+0.1-sm}{\sqrt{s}}\right) \leq \Phi(-0.4m+0.1) \leq \frac{1}{4}, & \forall m \gtrsim 1. \end{cases}$$

Therefore, we have

$$-\partial_s F_s(x) \gtrsim \frac{m^2}{e^{(1-s)m^2/2}}.$$

To summarize, we derive the following lower bound on the metric derivative:

$$\begin{aligned} |\dot{\pi}|_s^2 &= \int \frac{(\partial_s F_s(x))^2}{f_s(x)} dx \geq \int_{\frac{m}{2}-0.1}^{\frac{m}{2}+0.1} \frac{(\partial_s F_s(x))^2}{f_s(x)} dx \\ &\gtrsim \int_{\frac{m}{2}-0.1}^{\frac{m}{2}+0.1} \frac{m^4 e^{-(1-s)m^2}}{e^{-m^2/8}} dx \\ &\gtrsim m^4 e^{(s-\frac{7}{8})m^2} \geq m^4 e^{\frac{m^2}{40}}, \quad \forall s \in [0.9, 0.99]. \end{aligned}$$

Finally, recall that $S(\theta) := \frac{1}{1+m^2(1-\theta)^r}$, and $\pi_\theta = \pi_{S(\theta)}$. Hence, by chain rule of derivative, $|\dot{\pi}|_\theta = |\dot{\pi}|_{S(\theta)} |S'(\theta)|$. Let

$$\Theta := \{\theta \in [0, 1] : S(\theta) \in [0.9, 0.99]\} = \left[1 - \left(\frac{1/0.9 - 1}{m^2} \right)^{\frac{1}{r}}, 1 - \left(\frac{1/0.99 - 1}{m^2} \right)^{\frac{1}{r}} \right].$$

We have

$$\begin{aligned} \mathcal{A}_r &= \int_0^1 |\dot{\pi}|_\theta^2 d\theta = \int_0^1 |\dot{\pi}|_{S'(\theta)}^2 |S'(\theta)|^2 d\theta \geq \int_\Theta |\dot{\pi}|_{S'(\theta)}^2 |S'(\theta)|^2 d\theta \\ &\geq \min_{\theta \in \Theta} |S'(\theta)| \cdot \int_\Theta |\dot{\pi}|_{S'(\theta)}^2 |S'(\theta)| d\theta = \min_{\theta \in \Theta} |S'(\theta)| \cdot \int_{0.9}^{0.99} |\dot{\pi}|_s^2 ds. \end{aligned}$$

Since

$$|S'(\theta)| = \frac{m^2 r (1-\theta)^{r-1}}{(1+m^2(1-\theta)^r)^2} \geq \frac{m^2 r \left(\frac{1/0.99-1}{m^2}\right)^{1-1/r}}{\left(1+m^2 \left(\frac{1/0.9-1}{m^2}\right)\right)^2} \gtrsim m^{2/r} \gtrsim 1, \forall \theta \in \Theta,$$

the proof is complete. \square

Remark. In the above theorem, we established an exponential lower bound on the metric derivative of the W_2 distance, given by $\lim_{\delta \rightarrow 0} \frac{W_2(\pi_s, \pi_{s+\delta})}{|\delta|}$. In OT, another useful distance, the **Wasserstein-1 (W_1) distance**, defined as $W_1(\mu, \nu) = \inf_{\gamma \in \Pi(\mu, \nu)} \int \|x - y\| \gamma(dx, dy)$, is a lower bound of the W_2 distance. Below, we present a surprising result regarding the metric derivative of W_1 distance on the same curve of probability distributions. This result reveals an exponentially large gap between the W_1 and W_2 metric derivatives on the same curve, which is of independent interest.

Theorem 6. Define the probability distributions π_s as in the proof of Prop. 1, for some large enough $m \gtrsim 1$. Then, for all $s \in [0.9, 0.99]$, we have

$$\lim_{\delta \rightarrow 0} \frac{W_1(\pi_s, \pi_{s+\delta})}{|\delta|} \lesssim 1.$$

Proof. Since $W_1(\mu, \nu) = \int |F_\mu(x) - F_\nu(x)| dx$ (Villani, 2003, Thm. 2.18), by assuming regularity conditions, we have

$$\begin{aligned} \lim_{\delta \rightarrow 0} \frac{W_1(\pi_s, \pi_{s+\delta})}{|\delta|} &= \int |\partial_s F_s(x)| dx \\ &\leq \int \left| w'(s) \left(\Phi\left(\frac{x}{\sqrt{s}}\right) - \Phi\left(\frac{x-sm}{\sqrt{s}}\right) \right) \right| dx \\ &\quad + \int \left| w(s) \phi\left(\frac{x}{\sqrt{s}}\right) \frac{x}{2s^{\frac{3}{2}}} \right| dx \\ &\quad + \int \left| (1-w(s)) \phi\left(\frac{x-sm}{\sqrt{s}}\right) \left(\frac{x}{2s^{\frac{3}{2}}} + \frac{m}{2s^{\frac{1}{2}}} \right) \right| dx. \end{aligned}$$

To bound the first term, notice that for any $\lambda > 0$,

$$\Phi\left(\frac{x}{\sqrt{s}}\right) - \Phi\left(\frac{x-sm}{\sqrt{s}}\right) \lesssim \begin{cases} \sqrt{s} m e^{-\frac{(x-sm)^2}{2s}}, & \frac{x-sm}{\sqrt{s}} \geq \lambda; \\ \sqrt{s} m e^{-\frac{x^2}{2s}}, & \frac{x}{\sqrt{s}} \leq -\lambda; \\ 1, & \text{otherwise.} \end{cases}$$

Therefore, using Gaussian tail bound $1 - \Phi(\lambda) \leq \frac{1}{2} e^{-\frac{\lambda^2}{2}}$, the first term is bounded by

$$\begin{aligned} &\lesssim \frac{m^2}{e^{(1-s)m^2/2}} [2\sqrt{s}\lambda + sm + sm(1 - \Phi(\lambda)) + sm\Phi(-\lambda)] \\ &\lesssim \frac{m^2}{e^{(1-s)m^2/2}} [\lambda + m + e^{-\frac{\lambda^2}{2}}] \stackrel{\lambda \leftarrow \Theta(m)}{\lesssim} \frac{m^3}{e^{(1-s)m^2/2}} = o(1). \end{aligned}$$

The second term is bounded by

$$\lesssim \int \phi\left(\frac{x}{\sqrt{s}}\right) |x| dx = s \int \phi(u) |u| du \lesssim 1.$$

Finally, the third term is bounded by

$$\begin{aligned} &\lesssim \frac{1}{e^{(1-s)m^2/2}} \int \phi\left(\frac{x-sm}{\sqrt{s}}\right) (|x|+m) dx \\ &\lesssim \frac{1}{e^{(1-s)m^2/2}} \int \phi(u) (|u|+m) du \lesssim \frac{m}{e^{(1-s)m^2/2}} = o(1). \end{aligned}$$

□

E.2 PROOF OF PROP. 2

Proof. Note that $(\pi_t)_{t \in [0, \infty)}$ satisfies the Fokker-Planck equation $\partial_t \pi_t = \nabla \cdot (\pi_t \nabla \log \frac{\pi_t}{\gamma})$. Hence, the vector field $(v_t := -\nabla \log \frac{\pi_t}{\gamma})_{t \in [0, \infty)}$ generates $(\pi_t)_{t \in [0, \infty)}$, and each v_t can be written as a gradient field of a potential function. Thus, by the uniqueness result in Lem. 2, we conclude that

$$|\dot{\pi}_t|_t^2 = \left\| \nabla \log \frac{\pi_t}{\gamma} \right\|_{L^2(\pi_t)}^2 = \text{FI}(\pi_t \|\gamma) \leq e^{-2t} \text{FI}(\pi \|\gamma),$$

where FI is the Fisher divergence, and the last inequality is due to Villani (2003, Eq. 9.34). Finally, using the smoothness of V and Lem. 14, we have

$$\text{FI}(\pi \|\gamma) = \mathbb{E}_{\pi(x)} \|\nabla V(x) + x\|^2 \leq 2(\mathbb{E}_{\pi} \|\nabla V\|^2 + \mathbb{E}_{\pi} \|\cdot\|^2) \leq 2(d\beta + m^2),$$

□

E.3 PROOF OF THM. 5

Proof. By Nelson's relation (Lem. 3), \mathbb{Q} is equivalent to the path measure of the following SDE:

$$dX_t = X_t dt + \sqrt{2} dB_t^+, \quad t \in [0, T - \delta]; \quad X_{T-\delta} \sim \underline{\pi}_\delta.$$

Leveraging Girsanov theorem (Lem. 4), we know that for a.s. $X \sim \mathbb{Q}^\dagger$:

$$\begin{aligned} &\log \frac{d\mathbb{Q}^\dagger}{d\mathbb{Q}}(X) \\ &= \log \frac{\phi(X_0)}{\underline{\pi}_\delta(X_{T-\delta})} + (T - \delta)d + \int_0^{T-\delta} \left(\|s_{T-t_-}(X_{t_-})\|^2 dt + \sqrt{2} \langle s_{T-t_-}(X_{t_-}), dB_t \rangle \right) \\ &= \log Z + W(X) + \log \frac{d\pi}{d\underline{\pi}_\delta}(X_{T-\delta}). \end{aligned}$$

Thus, the equation $\mathbb{E}_{\mathbb{Q}^\dagger} \frac{d\mathbb{Q}}{d\mathbb{Q}^\dagger} = 1$ implies

$$Z = \mathbb{E}_{\mathbb{Q}^\dagger(X)} e^{-W(X)} \frac{d\underline{\pi}_\delta}{d\pi}(X_{T-\delta}).$$

Since $\frac{\hat{Z}}{Z} = \frac{d\mathbb{Q}}{d\mathbb{Q}^\dagger}(X) \frac{d\pi}{d\underline{\pi}_\delta}(X_{T-\delta})$, we have

$$\begin{aligned} \Pr \left(\left| \frac{\hat{Z}}{Z} - 1 \right| \geq \varepsilon \right) &= \Pr_{X \sim \mathbb{Q}^\dagger} \left(\left| \frac{d\mathbb{Q}}{d\mathbb{Q}^\dagger}(X) \frac{d\pi}{d\underline{\pi}_\delta}(X_{T-\delta}) - 1 \right| \geq \varepsilon \right) \\ &\leq \Pr_{X \sim \mathbb{Q}^\dagger} \left(\left| \frac{d\mathbb{Q}}{d\mathbb{Q}^\dagger}(X) - 1 \right| \gtrsim \varepsilon \right) + \Pr_{X \sim \mathbb{Q}^\dagger} \left(\left| \frac{d\pi}{d\underline{\pi}_\delta}(X_{T-\delta}) - 1 \right| \gtrsim \varepsilon \right). \end{aligned}$$

The inequality is due to the fact that $|ab - 1| \geq \varepsilon$ implies $|a - 1| \geq \frac{\varepsilon}{3}$ or $|b - 1| \geq \frac{\varepsilon}{3}$ for $\varepsilon \in [0, 1]$. It suffices to make both terms above $O(1)$. To bound the first term, we use the similar approach as in the proof of Eq. (20) in Thm. 2:

$$\Pr_{X \sim \mathbb{Q}^\dagger} \left(\left| \frac{d\mathbb{Q}}{d\mathbb{Q}^\dagger}(X) - 1 \right| \gtrsim \varepsilon \right) = \mathbb{Q}^\dagger \left(\left| \frac{d\mathbb{Q}}{d\mathbb{Q}^\dagger} - 1 \right| \gtrsim \varepsilon \right) \lesssim \frac{\text{TV}(\mathbb{Q}, \mathbb{Q}^\dagger)}{\varepsilon} \lesssim \frac{\sqrt{\text{KL}(\mathbb{Q} \|\mathbb{Q}^\dagger)}}{\varepsilon}.$$

Hence, it suffices to let $\text{TV}(\mathbb{Q}, \mathbb{Q}^\dagger)^2 \lesssim \text{KL}(\mathbb{Q} \parallel \mathbb{Q}^\dagger) \lesssim \varepsilon^2$. To bound the second term, we have

$$\begin{aligned} \Pr_{X \sim \mathbb{Q}^\dagger} \left(\left| \frac{d\pi}{d\pi_\delta}(X_{T-\delta}) - 1 \right| \gtrsim \varepsilon \right) &\leq \Pr_{X \sim \mathbb{Q}} \left(\left| \frac{d\pi}{d\pi_\delta}(X_{T-\delta}) - 1 \right| \gtrsim \varepsilon \right) + \text{TV}(\mathbb{Q}, \mathbb{Q}^\dagger) \\ &\leq \pi_\delta \left(\left| \frac{d\pi}{d\pi_\delta} - 1 \right| \gtrsim \varepsilon \right) + \text{TV}(\mathbb{Q}, \mathbb{Q}^\dagger) \\ &\lesssim \frac{\text{TV}(\pi_\delta, \pi)}{\varepsilon} + \varepsilon. \end{aligned}$$

Therefore, it suffices to make $\text{TV}(\pi_\delta, \pi) \lesssim \varepsilon$. \square

E.4 AN UPPER BOUND OF THE TV DISTANCE ALONG THE OU PROCESS

Lemma 8. *Assume that the target distribution $\pi \propto e^{-V}$ satisfies Assump. 2, with the exception that $R \lesssim \frac{1}{\sqrt{\beta}}$. Let $\bar{\pi}_\delta$ be the distribution of Y_δ in the OU process (Eq. (3)) initialized at $Y_0 \sim \pi$, for some $\delta \lesssim 1$. Then,*

$$\text{TV}(\pi, \bar{\pi}_\delta) \lesssim \delta(\beta m^2 + \beta Rm + d + \beta) + \delta^{\frac{1}{2}} d^{\frac{1}{2}} \beta(m + R).$$

Remark. *Consider a simplified case where $R \ll 1$, $\beta \gtrsim 1$, and $m^2 \gtrsim d$. Then it suffices to choose $\delta \lesssim \frac{\varepsilon^2}{\beta^2 d^2}$ in order to guarantee $\text{TV}(\pi, \bar{\pi}_\delta) \lesssim \varepsilon$.*

Proof. Our proof is inspired by Lee et al. (2023, Lem. 6.4), which addresses the case where V is Lipschitz.

Without loss of generality, suppose $\pi = e^{-V}$. Let ϕ be the p.d.f. of $\mathcal{N}(0, I)$, and define $\sigma^2 := 1 - e^{-2\delta} \asymp \delta$. We will use the following inequality: $|e^a - e^b| \leq (e^a + e^b)|a - b|$, which is due to the convexity of the exponential function. By the smoothness of V ,

$$\|\nabla V(x)\| = \|\nabla V(x) - \nabla V(x_*)\| \leq \beta \|x - x_*\| \leq \beta(\|x\| + R).$$

Define $\pi'(x) = e^{d\delta} \pi(e^\delta x)$, and thus $\bar{\pi}_\delta(x) = \int \pi'(x + \sigma u) \phi(u) du$. Using triangle inequality, we bound $\text{TV}(\pi, \pi')$ and $\text{TV}(\pi', \bar{\pi}_\delta)$ separately. First,

$$\begin{aligned} \text{TV}(\pi, \pi') &= \frac{1}{2} \int |e^{-V(x)} - e^{-V(e^\delta x) + d\delta}| dx \\ &\lesssim \int (\pi(x) + \pi'(x)) (|V(e^\delta x) - V(x)| + d\delta) dx. \end{aligned}$$

By the smoothness,

$$\begin{aligned} |V(e^\delta x) - V(x)| &\leq \|\nabla V(x)\| (e^\delta - 1) \|x\| + \frac{\beta}{2} (e^\delta - 1)^2 \|x\|^2 \\ &\lesssim \beta(\|x\| + R) \delta \|x\| + \beta \delta^2 \|x\|^2 \\ &\lesssim \beta \delta \|x\|^2 + \beta \delta R \|x\|. \\ \implies \text{TV}(\pi, \pi') &\lesssim \delta \int (\pi(x) + \pi'(x)) (\beta \|x\|^2 + \beta R \|x\| + d) dx. \end{aligned}$$

Note that

$$\int \pi(x) (\beta \|x\|^2 + \beta R \|x\| + d) dx \leq \beta m^2 + \beta Rm + d.$$

Since $\mathbb{E}_{\pi'} \varphi = \mathbb{E}_\pi \varphi(e^{-\delta} \cdot)$, we also have $\int \pi'(x) (\beta \|x\|^2 + \beta R \|x\| + d) dx \lesssim \beta m^2 + \beta Rm + d$. We thus conclude that

$$\text{TV}(\pi, \pi') \lesssim \delta(\beta m^2 + \beta Rm + d).$$

Next,

$$\begin{aligned} \text{TV}(\pi', \bar{\pi}_\delta) &= \frac{1}{2} \int \left| \int (\pi'(x + \sigma u) - \pi'(x)) \phi(u) du \right| dx \\ &\lesssim \iint |\pi'(x + \sigma u) - \pi'(x)| \phi(u) du dx \\ &\lesssim \iint (\pi'(x + \sigma u) + \pi'(x)) |V(e^\delta(x + \sigma u)) - V(e^\delta x)| \phi(u) du dx. \end{aligned}$$

Again, by smoothness,

$$\begin{aligned} V(e^\delta(x + \sigma u)) - V(e^\delta x) &\leq \|\nabla V(e^\delta x)\|e^\delta \sigma \|u\| + \frac{\beta}{2}e^{2\delta}\sigma^2\|u\|^2 \\ &\lesssim \beta(e^\delta\|x\| + R)e^\delta \sigma \|u\| + \beta e^{2\delta}\sigma^2\|u\|^2 \\ &\lesssim \beta(\|x\| + R)\delta^{\frac{1}{2}}\|u\| + \beta\delta\|u\|^2. \end{aligned}$$

Therefore,

$$\text{TV}(\pi', \bar{\pi}_\delta) \lesssim \beta\delta^{\frac{1}{2}} \iint (\pi'(x + \sigma u) + \pi'(x))(\|u\|\|x\| + \|u\|R + \delta^{\frac{1}{2}}\|u\|^2)\phi(u)dudx.$$

Note that

$$\begin{aligned} &\iint \pi'(x)(\|u\|\|x\| + \|u\|R + \delta^{\frac{1}{2}}\|u\|^2)\phi(u)dudx \\ &\lesssim \mathbb{E}_{\pi'} \|\cdot\|d^{\frac{1}{2}} + Rd^{\frac{1}{2}} + \delta^{\frac{1}{2}}d \leq md^{\frac{1}{2}} + Rd^{\frac{1}{2}} + \delta^{\frac{1}{2}}; \\ &\iint \pi'(x + \sigma u)(\|u\|\|x\| + \|u\|R + \delta^{\frac{1}{2}}\|u\|^2)\phi(u)dudx \\ &= \iint \pi'(y)(\|u\|\|y - \sigma u\| + \|u\|R + \delta^{\frac{1}{2}}\|u\|^2)\phi(u)dudy \\ &\lesssim \iint \pi'(y)(\|u\|\|y\| + \|u\|R + \delta^{\frac{1}{2}}\|u\|^2)\phi(u)dudy \\ &\lesssim md^{\frac{1}{2}} + Rd^{\frac{1}{2}} + \delta^{\frac{1}{2}}. \end{aligned}$$

Therefore, $\text{TV}(\pi', \bar{\pi}_\delta) \lesssim \beta\delta^{\frac{1}{2}}(md^{\frac{1}{2}} + Rd^{\frac{1}{2}} + \delta^{\frac{1}{2}})$. The proof is complete. \square

E.5 DISCUSSION ON THE OVERALL COMPLEXITY OF RDS

In RDS, an accurate score estimate $s_t \approx \nabla \log \bar{\pi}_t$ is critical for the algorithmic efficiency. Existing methods estimate scores through different ways. Here, we review the existing methods and their complexity guarantees for sampling, and leverage Thm. 5 to derive the complexity of normalizing constant estimation. Throughout this section, we always assume that the target distribution $\pi \propto e^{-V}$ satisfies $m^2 := \mathbb{E}_\pi \|\cdot\|^2 < \infty$ and that V is β -smooth.

(I) Reverse diffusion Monte Carlo. The seminal work directly leveraged the following Tweedie's formula (Robbins, 1992) to estimate the score: Huang et al. (2024a)

$$\nabla \log \bar{\pi}_t(x) = \mathbb{E}_{\bar{\pi}_{0|t}(x_0|x)} \frac{e^{-t}x_0 - x}{1 - e^{-2t}}, \quad (34)$$

where

$$\bar{\pi}_{0|t}(x_0|x) \propto_{x_0} \exp\left(-V(x_0) - \frac{\|x_0 - e^t x\|^2}{2(e^{2t} - 1)}\right) \quad (35)$$

is the posterior distribution of Y_0 conditional on $Y_t = x$ in the OU process (Eq. (3)). The paper proposed to sample from $\bar{\pi}_{0|t}(\cdot|x)$ by LMC and estimate the score via empirical mean, which has a provably better LSI constant than the target distribution π (see Huang et al. (2024a, Lem. 2)). They show that if the scores $\nabla \log \bar{\pi}_t$ are uniformly β -Lipschitz, and assume that there exists some $c > 0$ and $n > 0$ such that for any $r > 0$, $V + r\|\cdot\|^2$ is convex for $\|x\| \geq \frac{c}{r^n}$, then w.p. $\geq 1 - \zeta$, the overall complexity for guaranteeing $\text{KL}(\mathbb{Q}\|\mathbb{Q}^\dagger) \lesssim \varepsilon^2$ is

$$O\left(\text{poly}\left(d, \frac{1}{\zeta}\right) \exp\left(\frac{1}{\varepsilon}\right)^{O(n)}\right),$$

which is also the complexity of obtaining a \hat{Z} satisfying Eq. (5).

(II) Recursive score diffusion-based Monte Carlo. A second work Huang et al. (2024b) proposed to estimate the scores in a recursive scheme. Assuming the scores $\nabla \log \bar{\pi}_t$ are uniformly β -Lipschitz, they established a complexity

$$\exp\left(\beta^3 \log^3\left(\beta, d, m^2, \frac{1}{\zeta}\right)\right)$$

in order to guarantee $\text{KL}(\mathbb{Q}||\mathbb{Q}^\dagger) \lesssim \varepsilon^2$ w.p. $\geq 1 - \zeta$.

(III) Zeroth-order diffusion Monte Carlo. The following work He et al. (2024) proposed a zeroth-order method that leverages rejection sampling to sample from $\bar{\pi}_{0|t}(\cdot|x)$. When V is β -smooth, they showed that with a small early stopping time δ , the overall complexity for guaranteeing $\text{KL}(\mathbb{Q}||\mathbb{Q}^\dagger) \lesssim \varepsilon^2$ is

$$\exp\left(\tilde{O}(d) \log \beta \log \frac{1}{\varepsilon}\right).$$

(IV) Self-normalized estimator. Finally, a recent work Vacher et al. (2025) proposed to estimate the scores in a different approach:

$$\nabla \log \bar{\pi}_t(x) = -\frac{1}{1 - e^{-2t}} \frac{\mathbb{E}[\xi e^{-V(e^t(x-\xi))}]}{\mathbb{E}[e^{-V(e^t(x-\xi))}]}, \quad \text{where } \xi \sim \mathcal{N}(0, (1 - e^{-2t})I).$$

Assume that V is β -smooth, and the distributions along the OU process starting from $\pi \propto e^{-V}$ and $\pi' \propto e^{-2V}$ have potentials whose Hessians are uniformly $\succeq -\beta I$, then the complexity for guaranteeing $\mathbb{E} \text{KL}(\mathbb{Q}||\mathbb{Q}^\dagger) \lesssim \varepsilon^2$ is

$$O\left(\left(\frac{\beta(m^2 \vee d)}{\varepsilon}\right)^{O(d)}\right).$$

F SUPPLEMENTARY LEMMAS

Lemma 9. For $x > 0$ and $\varepsilon \in (0, \frac{1}{2})$, define $x_0 := |\log x|$ and $x_1 := |x - 1|$. Then $x_i \geq \varepsilon$ implies $x_{1-i} \geq \frac{\varepsilon}{2}$, and $x_i \leq \varepsilon$ implies $x_{1-i} \leq 2\varepsilon$, for both $i = 0, 1$.

This follows from the standard calculus approximation $\log x \approx x - 1$ when $x \approx 1$. The proof is straightforward and is left as an exercise for the reader.

Lemma 10. For any $0 \leq a \leq b \leq 1$ and $r \geq 1$, $b^r - a^r \leq r(b - a)$.

Proof. This is immediate from the decreasing property of the function $\varphi(x) := x^r - rx$, $x \in [0, 1]$, since $\varphi'(x) = r(x^{r-1} - 1) \leq 0$. \square

Lemma 11 (The median trick (Jerrum et al., 1986)). Let $\widehat{Z}_1, \dots, \widehat{Z}_N$ be $N(\geq 3)$ i.i.d. random variables satisfying

$$\Pr\left(\left|\frac{\widehat{Z}_n}{Z} - 1\right| \leq \varepsilon\right) \geq \frac{3}{4}, \quad \forall n \in [1, N],$$

and let \widehat{Z}_* be the median of $\widehat{Z}_1, \dots, \widehat{Z}_N$. Then

$$\Pr\left(\left|\frac{\widehat{Z}_*}{Z} - 1\right| \leq \varepsilon\right) \geq 1 - e^{-\frac{N}{72}}.$$

In particular, for any $\zeta \in (0, \frac{1}{4})$, choosing $N = \lceil 72 \log \frac{1}{\zeta} \rceil$, the probability is at least $1 - \zeta$.

Proof. Let $A_n := \left\{ \left| \frac{\widehat{Z}_n}{Z} - 1 \right| > \varepsilon \right\}$, which are i.i.d. events happening w.p. $p \leq \frac{1}{4}$. If $\left| \frac{\widehat{Z}_*}{Z} - 1 \right| > \varepsilon$, then there are at least $\lfloor \frac{N}{2} \rfloor$ A_n 's happening, i.e., $S_N := \sum_{n=1}^N 1_{A_n} \geq \lfloor \frac{N}{2} \rfloor$. Then,

$$\begin{aligned} \Pr \left(\left| \frac{\widehat{Z}_*}{Z} - 1 \right| > \varepsilon \right) &\leq \Pr \left(S_N \geq \left\lfloor \frac{N}{2} \right\rfloor \right) = \Pr \left(S_N - \mathbb{E} S_N \geq \left\lfloor \frac{N}{2} \right\rfloor - pN \right) \\ &\leq \Pr \left(S_N - \mathbb{E} S_N \geq \frac{N}{12} \right) \leq e^{-\frac{N}{72}}, \end{aligned}$$

where the first inequality on the second line follows from the fact that $\lfloor \frac{N}{2} \rfloor \geq \frac{N-1}{2} \geq \frac{N}{3}$ for all $N \geq 3$, and the last inequality is due to the Hoeffding's inequality. \square

Lemma 12. *The update rule of AIS (Eq. (14)) is:*

$$X_{T_\ell} = e^{-\Lambda(T_\ell)} X_0 - \left(\int_0^{T_\ell} e^{-(\Lambda(T_\ell) - \Lambda(t))} dt \right) \nabla V(X_0) + \left(2 \int_0^{T_\ell} e^{-2(\Lambda(T_\ell) - \Lambda(t))} dt \right)^{\frac{1}{2}} \xi,$$

where $\Lambda(t) := \int_0^t \lambda \left(\theta_{\ell-1} + \frac{\tau}{T_\ell} (\theta_\ell - \theta_{\ell-1}) \right) d\tau$, and $\xi \sim \mathcal{N}(0, I)$ is independent of X_0 .

Proof. By Itô's formula, we have

$$d \left(e^{\Lambda(t)} X_t \right) = e^{\Lambda(t)} \left(\Lambda'(t) X_t dt + dX_t \right) = e^{\Lambda(t)} \left(-\nabla V(X_0) dt + \sqrt{2} dB_t \right).$$

Integrating over $t \in [0, T_\ell]$, we obtain

$$\begin{aligned} e^{\Lambda(T_\ell)} X_{T_\ell} - X_0 &= - \left(\int_0^{T_\ell} e^{\Lambda(t)} dt \right) \nabla V(X_0) + \sqrt{2} \int_0^{T_\ell} e^{\Lambda(t)} dB_t, \\ \implies X_{T_\ell} &= e^{-\Lambda(T_\ell)} X_0 - \left(\int_0^{T_\ell} e^{-(\Lambda(T_\ell) - \Lambda(t))} dt \right) \nabla V(X_0) + \sqrt{2} \int_0^{T_\ell} e^{-(\Lambda(T_\ell) - \Lambda(t))} dB_t, \end{aligned}$$

and $\sqrt{2} \int_0^{T_\ell} e^{-(\Lambda(T_\ell) - \Lambda(t))} dB_t \sim \mathcal{N} \left(0, \left(2 \int_0^{T_\ell} e^{-2(\Lambda(T_\ell) - \Lambda(t))} dt \right) I \right)$ by Itô isometry. \square

Lemma 13. *The update rule of the RDS (Eq. (16)) is*

$$X_{t_{k+1}} = e^{t_{k+1} - t_k} X_{t_k} + 2(e^{t_{k+1} - t_k} - 1) s_{T-t_k}(X_{t_k}) + \Xi_k,$$

where

$$\Xi_k := \int_{t_k}^{t_{k+1}} \sqrt{2} e^{-(t-t_{k+1})} dB_t \sim \mathcal{N} \left(0, (e^{2(t_{k+1} - t_k)} - 1) I \right),$$

and the correlation matrix between Ξ_k and $B_{t_{k+1}} - B_{t_k}$ is

$$\text{Corr}(\Xi_k, B_{t_{k+1}} - B_{t_k}) = \frac{\sqrt{2}(e^{t_{k+1} - t_k} - 1)}{\sqrt{(e^{2(t_{k+1} - t_k)} - 1)(t_{k+1} - t_k)}} I.$$

Proof. By applying Itô's formula to Eq. (16) for $t \in [t_k, t_{k+1}]$, we have

$$\begin{aligned} d(e^{-t} X_t) &= e^{-t} (-X_t dt + dX_t) = e^{-t} (2s_{T-t_k}(X_{t_k}) dt + \sqrt{2} dB_t) \\ \implies e^{-t_{k+1}} X_{t_{k+1}} - e^{-t_k} X_{t_k} &= 2(e^{-t_k} - e^{-t_{k+1}}) s_{T-t_k}(X_{t_k}) + \int_{t_k}^{t_{k+1}} \sqrt{2} e^{-t} dB_t. \end{aligned}$$

The covariance between two zero-mean Gaussian random variables Ξ_k and $B_{t_{k+1}} - B_{t_k}$ is

$$\begin{aligned} \text{Cov}(\Xi_k, B_{t_{k+1}} - B_{t_k}) &= \mathbb{E} \left[\Xi_k (B_{t_{k+1}} - B_{t_k})^\top \right] \\ &= \mathbb{E} \left[\left(\int_{t_k}^{t_{k+1}} \sqrt{2} e^{-(t-t_{k+1})} dB_t \right) \left(\int_{t_k}^{t_{k+1}} dB_t \right)^\top \right] \\ &= \int_{t_k}^{t_{k+1}} \sqrt{2} e^{-(t-t_{k+1})} dt \cdot I = \sqrt{2}(e^{t_{k+1} - t_k} - 1) I. \end{aligned}$$

Finally, $\text{Corr}(u, v) = \text{diag}(\text{Cov } u)^{-\frac{1}{2}} \text{Cov}(u, v) \text{diag}(\text{Cov } v)^{-\frac{1}{2}}$ yields the correlation. \square

Lemma 14 (Chewi (2022, Lemma 4.E.1)). *Consider a probability measure $\mu \propto e^{-U}$ on \mathbb{R}^d .*

1. *If $\nabla^2 U \succeq \alpha I$ for some $\alpha > 0$ and x_* is the global minimizer of U , then $\mathbb{E}_\mu \|\cdot - x_*\|^2 \leq \frac{d}{\alpha}$.*

2. *If $\nabla^2 U \preceq \beta I$ for some $\beta > 0$, then $\mathbb{E}_\mu \|\nabla U\|^2 \leq \beta d$.*

Lemma 15. *Define $\hat{\pi}_\lambda \propto \exp(-V - \frac{\lambda}{2} \|\cdot\|^2)$, $\lambda \geq 0$. Then under Assump. 2, $\mathbb{E}_{\hat{\pi}_\lambda} \|\cdot\|^2 \leq m^2$ for all $\lambda \geq 0$.*

Proof. Let $V_\lambda := V + \frac{\lambda}{2} \|\cdot\|^2$, and $Z_\lambda = \int e^{-V_\lambda} dx$, so $\hat{\pi}_\lambda = e^{-V_\lambda - \log Z_\lambda}$. We have

$$\begin{aligned} \frac{d}{d\lambda} \log Z_\lambda &= \frac{Z'_\lambda}{Z_\lambda} = -\frac{1}{Z_\lambda} \int e^{-V_\lambda} V'_\lambda dx = -\frac{1}{2} \mathbb{E}_{\hat{\pi}_\lambda} \|\cdot\|^2, \\ \implies \frac{d}{d\lambda} \log \hat{\pi}_\lambda &= -V'_\lambda - \frac{d}{d\lambda} \log Z_\lambda = \frac{1}{2} (\mathbb{E}_{\hat{\pi}_\lambda} \|\cdot\|^2 - \|\cdot\|^2), \\ \implies \frac{d}{d\lambda} \mathbb{E}_{\hat{\pi}_\lambda} \|\cdot\|^2 &= \int \|\cdot\|^2 \left(\frac{d}{d\lambda} \log \hat{\pi}_\lambda \right) d\hat{\pi}_\lambda = \frac{1}{2} \left((\mathbb{E}_{\hat{\pi}_\lambda} \|\cdot\|^2)^2 - \mathbb{E}_{\hat{\pi}_\lambda} \|\cdot\|^4 \right) \leq 0. \end{aligned}$$

□

Lemma 16. *If a function U on \mathbb{R}^d satisfies $0 \prec \nabla^2 U \preceq \beta I$ for some $\beta > 0$, and for any $t \geq 0$, let x_t be the global minimizer of $U + \frac{t}{2} \|\cdot\|^2$. We have $\|x_t\| \leq \frac{\|x_0\|}{1 + \frac{t}{\beta}}$.*

Proof. Since $\nabla U(x_t) + tx_t = 0$, taking time derivative yields $\nabla^2 U(x_t) \dot{x}_t + x_t + t \dot{x}_t = 0$. Due to convexity, $\dot{x}_t = -(\nabla^2 U(x_t) + tI)^{-1} x_t$. Therefore,

$$\frac{1}{2} \frac{d}{dt} \|x_t\|^2 = x_t^\top \dot{x}_t = -x_t^\top (\nabla^2 U(x_t) + tI)^{-1} x_t \leq -\frac{\|x_t\|^2}{\beta + t},$$

which implies $\frac{d}{dt} \left(\left(1 + \frac{t}{\beta}\right)^2 \|x_t\|^2 \right) \leq 0$, and thus the proof is complete. □

G REVIEW AND DISCUSSION ON THE ERROR GUARANTEE (EQ. (5))

G.1 LITERATURE REVIEW OF EXISTING BOUNDS

Estimation of Z . Traditionally, the statistical properties of an estimator are typically analyzed through its bias and variance. However, deriving closed-form expressions of the variance of \hat{Z} and \hat{F} in JE remains challenging. Recall that the estimator $\hat{Z} = Z_0 e^{-W(X)}$, $X \sim \mathbb{P}^\rightarrow$ for $Z = Z_0 e^{-\Delta F}$, and that JE implies Bias $\hat{Z} = 0$. For general (sub-optimally) controlled SDEs, Hartmann & Richter (2024) established both upper and lower bounds of the relative error of the importance sampling estimator, yet bounds tailored for JE are not well-studied. Inspired by this, we establish an upper bound on the *normalized variance* $\text{Var} \frac{\hat{Z}}{Z}$ in Prop. 3 at the end of this section using techniques in Rényi divergence. However, we remark that connecting this upper bound to the properties of the curve (e.g., action) is non-trivial, which we leave for future work.

Estimation of F . Turning to the estimator $\hat{F} = -\log \hat{Z}$ for $F = -\log Z$, we have

$$\text{Bias } \hat{F} = \mathbb{E}_{\mathbb{P}^\rightarrow} W - \Delta F = \mathcal{W} - \Delta F = \mathcal{W}_{\text{diss}}.$$

Bounding the average dissipated work $\mathcal{W}_{\text{diss}} = \text{KL}(\mathbb{P}^\rightarrow \|\mathbb{P}^\leftarrow) = -\mathbb{E}_{\mathbb{P}^\rightarrow} \int_0^T (\partial_t \log \tilde{\pi}_t)(X_t) dt$ remains challenging as well, as the law of X_t under \mathbb{P}^\rightarrow is unknown, thus complicating the bounding of the expectation. To the best of our knowledge, Chen et al. (2020) established a lower bound in terms of $W_2(\pi_0, \pi_1)$ via the Wasserstein gradient flow, but an upper bound remains elusive. Furthermore, $\mathbb{E} \hat{F}^2 = \mathbb{E}_{\mathbb{P}^\rightarrow(X)} (\log Z_0 - W(X))^2$ is similarly intractable to analyze.

For multiple estimators, i.e., $\hat{F}_K := -\log \left(Z_0 \frac{1}{M} \sum_{k=1}^K e^{-W(X^{(k)})} \right)$ where $X^{(1)}, \dots, X^{(K)}$ i.i.d. \mathbb{P}^\rightarrow , Zuckerman & Woolf (2002; 2004) (see also Lelièvre et al. (2010, Sec. 4.1.5)) derived approximate asymptotic bounds on Bias \hat{F}_K and Var \hat{F}_K via the delta method (or equivalently, the central limit theorem and Taylor expansions). Precise and non-asymptotic bounds remain elusive to date.

G.2 EQUIVALENCE IN COMPLEXITIES FOR ESTIMATING Z AND F

We prove the claim in Sec. 3 that estimating Z with $O(\varepsilon)$ relative error and estimating F with $O(\varepsilon)$ absolute error share the same complexity up to absolute constants. This follows directly from Lem. 9: for any $\varepsilon \in (0, \frac{1}{2})$,

$$\text{Eq. (5)} \implies \Pr\left(|\hat{F} - F| \leq 2\varepsilon\right) \geq \frac{3}{4}, \quad \text{and} \quad \text{Eq. (5)} \iff \Pr\left(|\hat{F} - F| \leq \frac{\varepsilon}{2}\right) \geq \frac{3}{4}.$$

G.3 EQ. (5) IS WEAKER THAN BIAS AND VARIANCE

We demonstrate that Eq. (5) is a weaker criterion than controlling bias and variance, which is an immediate result from the Chebyshev inequality:

$$\begin{aligned} \Pr\left(\left|\frac{\hat{Z}}{Z} - 1\right| \geq \varepsilon\right) &\leq \frac{1}{\varepsilon^2} \mathbb{E}\left(\frac{\hat{Z}}{Z} - 1\right)^2 = \frac{\text{Bias}^2 \hat{Z} + \text{Var} \hat{Z}}{\varepsilon^2 Z^2}, \\ \Pr\left(|\hat{F} - F| \geq \varepsilon\right) &\leq \frac{\mathbb{E}(\hat{F} - F)^2}{\varepsilon^2} = \frac{\text{Bias}^2 \hat{F} + \text{Var} \hat{F}}{\varepsilon^2}. \end{aligned}$$

On the other hand, suppose one has established a bound in the following form:

$$\Pr\left(\left|\frac{\hat{Z}}{Z} - 1\right| \geq \varepsilon\right) \leq p(\varepsilon), \quad \text{for some } p : [0, \infty) \rightarrow [0, 1],$$

and assume that \hat{Z} is unbiased. Then this implies

$$\text{Var} \frac{\hat{Z}}{Z} = \mathbb{E}\left(\frac{\hat{Z}}{Z} - 1\right)^2 = \int_0^\infty \Pr\left(\left(\frac{\hat{Z}}{Z} - 1\right)^2 \geq \varepsilon\right) d\varepsilon \leq \int_0^\infty p(\sqrt{\varepsilon}) d\varepsilon.$$

G.4 AN UPPER BOUND ON THE NORMALIZED VARIANCE OF \hat{Z} IN JARZYNSKI EQUALITY

Proposition 3. *Under the setting of JE (Thm. 1), let $(v_t)_{t \in [0, T]}$ be any vector field that generates $(\tilde{\pi}_t)_{t \in [0, T]}$, and define \mathbb{P} as the path measure of Eq. (19). Then,*

$$\text{Var} \frac{\hat{Z}}{Z} \leq \left[\mathbb{E}_{\mathbb{P}} \exp\left(14 \int_0^T \|v_t(X_t)\|^2 dt\right) \right]^{\frac{1}{2}} - 1.$$

Proof. The proof is inspired by Chewi et al. (2022). Note that

$$\text{Var} \frac{\hat{Z}}{Z} = \mathbb{E}\left(\frac{\hat{Z}}{Z}\right)^2 - 1 = \mathbb{E}_{\mathbb{P}^\leftarrow} \left(e^{-W(X) + \Delta F}\right)^2 - 1 = \mathbb{E}_{\mathbb{P}^\rightarrow} \left(\frac{d\mathbb{P}^\leftarrow}{d\mathbb{P}^\rightarrow}\right)^2 - 1,$$

which is the χ^2 divergence from \mathbb{P}^\leftarrow to \mathbb{P}^\rightarrow . Recall the $q(> 1)$ -Rényi divergence defined as $R_q(\mu \parallel \nu) = \frac{1}{q-1} \log \mathbb{E}_\nu \left(\frac{d\mu}{d\nu}\right)^q$, and that $\chi^2(\mathbb{P}^\leftarrow \parallel \mathbb{P}^\rightarrow) = e^{R_2(\mathbb{P}^\leftarrow \parallel \mathbb{P}^\rightarrow)} - 1$. By the weak triangle inequality of Rényi divergence (Chewi, 2022, Lem. 6.2.5):

$$R_2(\mathbb{P}^\leftarrow \parallel \mathbb{P}^\rightarrow) \leq \frac{3}{2} R_4(\mathbb{P}^\leftarrow \parallel \mathbb{P}) + R_3(\mathbb{P} \parallel \mathbb{P}^\rightarrow).$$

We now bound $\mathbb{E}_{\mathbb{P}} \left(\frac{d\mathbb{P}^\rightarrow}{d\mathbb{P}}\right)^q$ for any $q \in \mathbb{R}$. By Girsanov theorem (Lem. 1),

$$\log \frac{d\mathbb{P}^\rightarrow}{d\mathbb{P}}(X) = \int_0^T \left(-\frac{1}{\sqrt{2}} \langle v_t(X_t), dB_t \rangle - \frac{1}{4} \|v_t(X_t)\|^2 dt\right), \quad \text{a.s. } X \sim \mathbb{P}.$$

Therefore,

$$\begin{aligned}
& \mathbb{E}_{\mathbb{P}} \left(\frac{d\mathbb{P}^{\rightarrow}}{d\mathbb{P}} \right)^q \\
&= \mathbb{E}_{\mathbb{P}} \exp \int_0^T \left(-\frac{q}{\sqrt{2}} \langle v_t(X_t), dB_t \rangle - \frac{q}{4} \|v_t(X_t)\|^2 dt \right) \\
&= \mathbb{E}_{\mathbb{P}} \exp \left[\int_0^T \left(-\frac{q}{\sqrt{2}} \langle v_t(X_t), dB_t \rangle - \frac{q^2}{2} \|v_t(X_t)\|^2 dt \right) + \int_0^T \left(\frac{q^2}{2} - \frac{q}{4} \right) \|v_t(X_t)\|^2 dt \right] \\
&\leq \left(\mathbb{E}_{\mathbb{P}} \exp \left[\int_0^T \left(-\sqrt{2}q \langle v_t(X_t), dB_t \rangle - q^2 \|v_t(X_t)\|^2 dt \right) \right] \right)^{\frac{1}{2}} \\
&\cdot \left(\mathbb{E}_{\mathbb{P}} \exp \left[\left(q^2 - \frac{q}{2} \right) \int_0^T \|v_t(X_t)\|^2 dt \right] \right)^{\frac{1}{2}},
\end{aligned}$$

where the last line is by the Cauchy-Schwarz inequality. Let $M_t := -\sqrt{2}q \int_0^t \langle v_r(X_r), dB_r \rangle$, $X \sim \mathbb{P}$ be a continuous martingale with quadratic variation $[M]_t = \int_0^t 2q^2 \|v_r(X_r)\|^2 dr$. By Karatzas & Shreve (1991, Chap. 3.5.D), the process $t \mapsto e^{M_t - \frac{1}{2}[M]_t}$ is a super martingale, and hence $\mathbb{E} e^{M_T - \frac{1}{2}[M]_T} \leq 1$. Thus, we have

$$\mathbb{E}_{\mathbb{P}} \left(\frac{d\mathbb{P}^{\rightarrow}}{d\mathbb{P}} \right)^q \leq \left(\mathbb{E}_{\mathbb{P}} \exp \left[\left(q^2 - \frac{q}{2} \right) \int_0^T \|v_t(X_t)\|^2 dt \right] \right)^{\frac{1}{2}}$$

From Girsanov theorem (Lem. 4), we can similarly obtain the following RN derivative:

$$\log \frac{d\mathbb{P}^{\leftarrow}}{d\mathbb{P}}(X) = \int_0^T \left(-\frac{1}{\sqrt{2}} \langle v_t(X_t), *dB_t^{\leftarrow} \rangle - \frac{1}{4} \|v_t(X_t)\|^2 dt \right), \text{ a.s. } X \sim \mathbb{P}.$$

and use the same argument to show that $\mathbb{E}_{\mathbb{P}} \left(\frac{d\mathbb{P}^{\leftarrow}}{d\mathbb{P}} \right)^q$ has exactly the same upper bound as $\mathbb{E}_{\mathbb{P}} \left(\frac{d\mathbb{P}^{\rightarrow}}{d\mathbb{P}} \right)^q$. In particular, we can use the same martingale argument, whereas now the *backward* continuous martingale is defined as $M'_t := -\sqrt{2}q \int_t^T \langle v_r(X_r), *dB_r^{\leftarrow} \rangle$, $X \sim \mathbb{P}$, with quadratic variation $[M']_t = \int_t^T 2q^2 \|v_r(X_r)\|^2 dr$. Therefore, we conclude that

$$\begin{aligned}
R_2(\mathbb{P}^{\leftarrow} \|\mathbb{P}^{\rightarrow}) &\leq \frac{1}{4} \log \mathbb{E}_{\mathbb{P}} \exp \left(14 \int_0^T \|v_t(X_t)\|^2 dt \right) + \frac{1}{4} \log \mathbb{E}_{\mathbb{P}} \exp \left(5 \int_0^T \|v_t(X_t)\|^2 dt \right) \\
&\leq \frac{1}{2} \log \mathbb{E}_{\mathbb{P}} \exp \left(14 \int_0^T \|v_t(X_t)\|^2 dt \right).
\end{aligned}$$

□

H FURTHER DISCUSSION ON RELATED WORKS

Related works. We briefly review some related works, and defer detailed discussion to App. H.

- **Methods for normalizing constant estimation.** We mainly discuss two classes of methods here. First, the *equilibrium* methods, such as TI (Kirkwood, 1935) and its variants (Brosse et al., 2018; Ge et al., 2020; Chehab et al., 2023; Kook & Vempala, 2024), which involve sampling sequentially from a series of equilibrium Markov transition kernels. Second, the *non-equilibrium* methods, such as AIS (Neal, 2001), which samples from a non-equilibrium SDE that gradually evolves from a prior distribution to the target distributions. In App. H.1, we show that TI is a special case of AIS using the “perfect” transition kernels. Recent years have also witnessed the emergence of *learning-based* non-equilibrium methods for normalizing constant estimation, which are typically byproducts of sampling algorithms (Zhang & Chen, 2022; Nüsken & Richter, 2021; Richter & Berner, 2024; Sun et al., 2024; Vargas et al., 2024; Albergo & Vanden-Eijnden, 2024; Blessing et al., 2025; Chen et al., 2025). Additionally, there are also several methods based on particle filtering (e.g., Kostov & Whiteley (2017); Jasra et al. (2018); Ruzayqat et al. (2022)).

- **Variance reduction in JE and AIS.** Our poof methodology focuses on the discrepancy between the sampling path measure and the reference path measure, which is related to the variance reduction technique in applying JE and AIS. For example, Vaikuntanathan & Jarzynski (2008) introduced the idea of escorted simulation, Hartmann et al. (2017) proposed a method for learning the optimal control protocol in JE through the variational characterization of free energy, and Doucet et al. (2022) leveraged score-based generative model to learn the optimal backward kernel. Quantifying the discrepancy between path measures is the core of our analysis.
- **Complexity analysis for normalizing constant estimation.** Chehab et al. (2023) studied the asymptotic statistical efficiency of the curve for TI measured by the asymptotic mean-squared error, and highlighted the advantage of the geometric interpolation. In terms of non-asymptotic analysis, existing works mainly rely on the isoperimetry of the target distribution. For instance, Andrieu et al. (2016) derived bounds of bias and variance for TI under Poincaré inequality, Brosse et al. (2018) provided complexity guarantees for TI under both strong and weak log-concavity conditions, while Ge et al. (2020) improved the complexity under strong log-concavity using multilevel Monte Carlo.

H.1 THERMODYNAMIC INTEGRATION

(I) Review of TI. We first briefly review the thermodynamic integration (TI) algorithm. Its essence is to write the free-energy difference as an integral of the derivative of free energy. Consider the general curve of probability measures $(\pi_\theta)_{\theta \in [0,1]}$ defined in Eq. (11). Then,

$$\frac{d}{d\theta} \log Z_\theta = -\frac{1}{Z_\theta} \int e^{-V_\theta(x)} \partial_\theta V_\theta(x) dx = -\mathbb{E}_{\pi_\theta} \partial_\theta V_\theta \implies \log \frac{Z}{Z_0} = -\int_0^1 \mathbb{E}_{\pi_\theta} \partial_\theta V_\theta d\theta. \quad (36)$$

One may choose time points $0 = \theta_0 < \dots < \theta_M = 1$ and approximate Eq. (36) by a Riemann sum:

$$\log \frac{Z}{Z_0} \approx -\sum_{\ell=0}^{M-1} (\theta_{\ell+1} - \theta_\ell) \mathbb{E}_{\pi_{\theta_\ell}} \partial_\theta |_{\theta=\theta_\ell} V_\theta, \quad (37)$$

where the expectation under each π_{θ_ℓ} can be estimated by sampling from π_{θ_ℓ} . Nevertheless, there is a way of writing the exact equality instead of the approximation in Eq. (37): since

$$\log \frac{Z_{\theta_{\ell+1}}}{Z_{\theta_\ell}} = \log \int \frac{1}{Z_{\theta_\ell}} e^{-V_{\theta_\ell}(x)} e^{-(V_{\theta_{\ell+1}}(x) - V_{\theta_\ell}(x))} dx = \log \mathbb{E}_{\pi_{\theta_\ell}} e^{-(V_{\theta_{\ell+1}} - V_{\theta_\ell})},$$

by summing over $\ell = 0, \dots, M-1$, we have

$$\log \frac{Z}{Z_0} = \sum_{\ell=0}^{M-1} \log \mathbb{E}_{\pi_{\theta_\ell}} e^{-(V_{\theta_{\ell+1}} - V_{\theta_\ell})}, \quad (38)$$

which constitutes the estimation framework used in Brosse et al. (2018); Ge et al. (2020); Chehab et al. (2023); Kook & Vempala (2024). Hence, we also use TI to name this algorithm.

(II) TI as a special case of AIS. We follow the notations used in Thm. 3 to demonstrate the following claim: *TI (Eq. (38)) is a special case of AIS with every transition kernel $F_\ell(x, \cdot)$ chosen as the perfect proposal π_{θ_ℓ} .*

Proof. In AIS, with $F_\ell(x, \cdot) = \pi_{\theta_\ell}$ in the forward path \mathbb{P}^\rightarrow , we have $\mathbb{P}^\rightarrow(x_{0:M}) = \prod_{\ell=0}^{M-1} \pi_{\theta_\ell}(x_\ell)$. In this special case,

$$W(x_{0:M}) = \log \prod_{\ell=0}^{M-1} \frac{e^{-V_{\theta_\ell}(x_\ell)}}{e^{-V_{\theta_{\ell+1}}(x_\ell)}},$$

and hence the AIS equality becomes the following identity, exactly the same as Eq. (36):

$$\frac{Z}{Z_0} = e^{-\Delta \mathcal{F}} = \mathbb{E}_{\mathbb{P}^\rightarrow} e^{-W} = \prod_{\ell=0}^{M-1} \mathbb{E}_{\pi_{\theta_\ell}} e^{-(V_{\theta_{\ell+1}} - V_{\theta_\ell})}, \quad (39)$$

□

(III) Distinction between equilibrium and non-equilibrium methods. In our AIS framework, the distinction lies in the choice of the transition kernels $F_\ell(x, \cdot)$ within the AIS framework.

In equilibrium methods, the transition kernels are ideally set to the perfect proposal π_{θ_ℓ} . However, in practice, exact sampling from π_{θ_ℓ} is generally infeasible. Instead, one can apply multiple MCMC iterations targeting π_{θ_ℓ} , leveraging the mixing properties of MCMC algorithms to gradually approach the desired distribution π_{θ_ℓ} . Nonetheless, unless using exact sampling methods such as rejection sampling – which is exponentially expensive in high dimensions – the resulting sample distribution inevitably remains biased with a finite number of MCMC iterations.

In contrast, non-equilibrium methods employ transition kernels specifically designed to transport $\pi_{\ell-1}$ toward π_ℓ , often following a curve of probability measures. This distinguishes them as inherently non-equilibrium. A key advantage of this approach over the equilibrium one is its ability to provide unbiased estimates, as demonstrated in JE and AIS.

H.2 PROOF OF THE SECOND PART OF LEM. 6

Recall that our goal is to estimate π_0 's normalizing constant $Z_0 = \int e^{-V_0} dx$, where V_0 is β -strongly convex and 3β -smooth, with global minimizer x^* satisfying $\|x^*\| \leq R \lesssim \frac{1}{\sqrt{\beta}}$. The aim is to obtain an estimator $\widehat{Z}_0 \approx Z_0$ such that

$$\Pr(\mathcal{F}) \leq \frac{1}{8}, \text{ where } \mathcal{F} := \left\{ \left| \frac{\widehat{Z}_0}{Z_0} - 1 \right| \geq \frac{\varepsilon}{8} \right\}. \quad (40)$$

Following the discussion above, the TI algorithm goes as follows. Consider a sequence of non-negative numbers $\lambda_0 > \lambda_1 > \dots > \lambda_K = 0$, where there exists a common $\gamma_0 > 0$ such that $\lambda_k = (1 + \gamma_0) \lambda_{k+1}$, for all $k \in \llbracket 0, K-2 \rrbracket$. Let $\rho_k := \frac{1}{\zeta_k} e^{-f_k}$, where $f_k := V_0 + \frac{\lambda_k}{2} \|\cdot\|^2$ is $(\beta + \lambda_k)$ -strongly-convex and $(3\beta + \lambda_k)$ -smooth. One can write

$$Z_0 = \zeta_K = \zeta_0 \prod_{k=0}^{K-1} \underbrace{\frac{\zeta_{k+1}}{\zeta_k}}_{=: G_k}, \text{ where } G_k = \mathbb{E}_{\rho_k} \underbrace{\exp\left(\frac{\lambda_k - \lambda_{k+1}}{2} \|\cdot\|^2\right)}_{=: g_k},$$

and estimate each G_k by

$$\widehat{G}_k := \frac{1}{N} \sum_{n=1}^N g_k(\widehat{X}_n^{(k)}), \quad \widehat{X}_n^{(k)} \stackrel{\text{i.i.d.}}{\sim} \widehat{\rho}_k \approx \rho_k,$$

so the final estimator is $\widehat{Z}_0 := \widehat{\zeta}_0 \prod_{k=0}^{K-1} \widehat{G}_k$, in which $\widehat{\zeta}_0 \approx \zeta_0$. To proceed, we first prove the following lemma.

Lemma 17. *If*

1. $\text{TV}(\widehat{\rho}_k, \rho_k) \leq \delta \asymp \frac{1}{NK}$, for all $k \in \llbracket 0, K-1 \rrbracket$.
2. The estimate $\widehat{\zeta}_0$ satisfies $\left| \log \frac{\widehat{\zeta}_0}{\zeta_0} \right| \lesssim \varepsilon$.
3. For all $k \in \llbracket 0, K-1 \rrbracket$, the following equation holds:

$$\frac{\mathbb{E}_{\rho_k} g_k^2}{(\mathbb{E}_{\rho_k} g_k)^2} \leq 1 + O(1). \quad (41)$$

Then with $N \asymp \frac{K}{\varepsilon^2}$, Eq. (40) holds.

Proof. By definition of TV distance, for each pair of (n, k) one can construct a random variable $X_n^{(k)} \sim \rho_k$ that only depends on $\widehat{X}_n^{(k)}$ and satisfies $\Pr(\widehat{X}_n^{(k)} \neq X_n^{(k)}) \leq \delta$. Define the event

$$\mathcal{E} = \left\{ \widehat{X}_n^{(k)} = X_n^{(k)} : \forall n \in \llbracket 1, N \rrbracket, k \in \llbracket 0, K-1 \rrbracket \right\}.$$

By independence, $\Pr(\mathcal{E}) \geq (1 - \delta)^{NK} \geq 1 - \delta NK \gtrsim 1$. If $\Pr(\mathcal{F}|\mathcal{E}) \leq \frac{1}{16}$ and $\Pr(\mathcal{E}^c) \leq \frac{1}{16}$, then

$$\Pr(\mathcal{F}) = \Pr(\mathcal{F}|\mathcal{E})\Pr(\mathcal{E}) + \Pr(\mathcal{F}|\mathcal{E}^c)\Pr(\mathcal{E}^c) \leq \Pr(\mathcal{F}|\mathcal{E}) + \Pr(\mathcal{E}^c) \leq \frac{1}{8},$$

as desired.

To obtain $\Pr(\mathcal{F}|\mathcal{E}) \leq \frac{1}{16}$, from now on we *always* assume that \mathcal{E} happens, and omit the conditional notation $(\cdot|\mathcal{E})$ in probability and expectation for simplicity. Note that in this case, $\widehat{G}_k = \frac{1}{N} \sum_{n=1}^N g_k(X_n^{(k)})$, $X_n^{(k)} \stackrel{\text{i.i.d.}}{\sim} \rho_k$, so $\mathbb{E}\widehat{G}_k = G_k$. One can upper bound the probability of large relative error as follows, leveraging Lem. 9:

$$\begin{aligned} \Pr\left(\left|\frac{\widehat{Z}_0}{Z_0} - 1\right| \gtrsim \varepsilon\right) &\leq \Pr\left(\left|\log \frac{\widehat{Z}_0}{Z_0}\right| \gtrsim \varepsilon\right) = \Pr\left(\left|\log \frac{\widehat{\zeta}_0}{\zeta_0} + \log \prod_{k=0}^{K-1} \frac{\widehat{G}_k}{G_k}\right| \gtrsim \varepsilon\right) \\ &\leq \Pr\left(\left|\log \frac{\widehat{\zeta}_0}{\zeta_0} + \log \prod_{k=0}^{K-1} \frac{\widehat{G}_k}{G_k}\right| \gtrsim \varepsilon\right) \\ &\leq \Pr\left(\left|\log \prod_{k=0}^{K-1} \frac{\widehat{G}_k}{G_k}\right| \gtrsim \varepsilon\right) \leq \Pr\left(\left|\prod_{k=0}^{K-1} \frac{\widehat{G}_k}{G_k} - 1\right| \gtrsim \varepsilon\right) \\ &\lesssim \frac{1}{\varepsilon^2} \mathbb{E}\left(\prod_{k=0}^{K-1} \frac{\widehat{G}_k}{G_k} - 1\right)^2 = \frac{1}{\varepsilon^2} \left(\prod_{k=0}^{K-1} \frac{\mathbb{E}\widehat{G}_k^2}{G_k^2} - 1\right), \end{aligned}$$

where the last line is due to Markov inequality. Choosing $N \asymp \frac{K}{\varepsilon^2}$ yields

$$\frac{\mathbb{E}\widehat{G}_k^2}{G_k^2} - 1 = \frac{\text{Var}\widehat{G}_k^2}{G_k^2} = \frac{\mathbb{E}\rho_k g_k^2 - (\mathbb{E}\rho_k g_k)^2}{N(\mathbb{E}\rho_k g_k)^2} \lesssim \frac{1}{N},$$

which implies

$$\Pr\left(\left|\frac{\widehat{Z}_0}{Z_0} - 1\right| \gtrsim \varepsilon\right) \lesssim \frac{1}{\varepsilon^2} \left(\prod_{k=0}^{K-1} \frac{\mathbb{E}\widehat{G}_k^2}{G_k^2} - 1\right) \leq \frac{1}{\varepsilon^2} \left(\left(1 + \frac{1}{N}\right)^K - 1\right) \lesssim \frac{K}{N\varepsilon^2} \lesssim 1.$$

□

The following lemmas show how to accurately estimate ζ_0 and how to satisfy Eq. (41).

Lemma 18. *With $\lambda_0 \asymp \frac{d\beta}{\varepsilon}$, $\widehat{\zeta}_0 := \exp\left(-V_0(0) + \frac{\|\nabla V_0(0)\|^2}{2(3\beta + \lambda_0)}\right) \left(\frac{2\pi}{3\beta + \lambda_0}\right)^{\frac{d}{2}}$ satisfies $\left|\log \frac{\widehat{\zeta}_0}{\zeta_0}\right| \lesssim \varepsilon$.*

Proof. By assumption, f_0 is $(\beta + \lambda_0)$ -strongly-convex and $(3\beta + \lambda_0)$ -smooth. Using quadratic upper and lower bounds on f_0 ,

$$\exp\left(-f_0(0) + \frac{\|\nabla f_0(0)\|^2}{2(3\beta + \lambda_0)}\right) \left(\frac{2\pi}{3\beta + \lambda_0}\right)^{\frac{d}{2}} \leq \zeta_0 \leq \exp\left(-f_0(0) + \frac{\|\nabla f_0(0)\|^2}{2(\beta + \lambda_0)}\right) \left(\frac{2\pi}{\beta + \lambda_0}\right)^{\frac{d}{2}}.$$

Since $f_0(0) = V_0(0)$, $\|\nabla f_0(0)\| = \|\nabla V_0(0)\| = \|\nabla V_0(0) - \nabla V_0(x')\| \leq 3\beta\|x'\| \lesssim \sqrt{\beta}$,

$$1 \leq \frac{\zeta_0}{\widehat{\zeta}_0} \leq \exp\left(\frac{\beta\|\nabla V_0(0)\|^2}{(\beta + \lambda_0)(3\beta + \lambda_0)}\right) \left(1 + \frac{2\beta}{\beta + \lambda_0}\right)^{\frac{d}{2}} \leq \exp\left(\frac{\beta^2}{(\beta + \lambda_0)(3\beta + \lambda_0)} + \frac{d\beta}{\beta + \lambda_0}\right).$$

So $\lambda_0 \asymp \frac{d\beta}{\varepsilon}$ implies $\frac{\beta^2}{(\beta + \lambda_0)(3\beta + \lambda_0)} + \frac{d\beta}{\beta + \lambda_0} \lesssim \varepsilon$. □

Lemma 19. *For $k = K - 1$, $\lambda_k \asymp \frac{\beta}{\sqrt{d}}$ implies Eq. (41).*

Proof. When $k = K - 1$, $g_k = \exp\left(\frac{\lambda_k}{2}\|\cdot\|^2\right)$. We have

$$\frac{\mathbb{E}\rho_k g_k^2}{(\mathbb{E}\rho_k g_k)^2} = \mathbb{E}_{\pi_0} \exp\left(\frac{\lambda_k}{2}\|\cdot\|^2\right) \mathbb{E}_{\pi_0} \exp\left(-\frac{\lambda_k}{2}\|\cdot\|^2\right).$$

Define

$$h_1(\lambda) := \mathbb{E}_{\pi_0} \exp(\lambda \|\cdot\|^2) \mathbb{E}_{\pi_0} \exp(-\lambda \|\cdot\|^2), \quad \lambda \in \left[0, \frac{\beta}{4}\right].$$

One can take derivative to obtain

$$\frac{d}{d\lambda} \log h_1(\lambda) = \int_{-\lambda}^{\lambda} \text{Var}_{\bar{\rho}_s} \|\cdot\|^2 ds,$$

where $\bar{\rho}_s \propto \exp(-V_0 + s \|\cdot\|^2)$ is $\frac{\beta}{2}$ -strongly-log-concave and thus satisfies $\frac{2}{\beta}$ -LSI. Hence,

$$\text{Var}_{\bar{\rho}_s} \|\cdot\|^2 \leq \frac{8}{\beta} \mathbb{E}_{\bar{\rho}_s} \|\cdot\|^2.$$

Let x'_s be the global minimizer of $V_0 - s \|\cdot\|^2$. By Lem. 16, $\|x'_s\| \leq R$. Leveraging Lem. 14, we have

$$\text{Var}_{\bar{\rho}_s} \|\cdot\|^2 \lesssim \frac{1}{\beta} (\mathbb{E}_{\bar{\rho}_s} \|\cdot - x'_s\|^2 + \|x'_s\|^2) \leq \frac{1}{\beta} \left(\frac{2d}{\beta} + R^2 \right) \lesssim \frac{d}{\beta^2}.$$

So $\frac{d}{d\lambda} \log h_1(\lambda) \lesssim \frac{\lambda d}{\beta^2}$, and thus

$$\frac{\mathbb{E}_{\rho_k} g_k^2}{(\mathbb{E}_{\rho_k} g_k)^2} = h_1\left(\frac{\lambda_k}{2}\right) = \exp\left(O\left(\frac{\lambda_k^2 d}{\beta^2}\right)\right) = 1 + O\left(\frac{\lambda_k^2 d}{\beta^2}\right) = 1 + O(1).$$

□

Lemma 20. For $k \in \llbracket 0, K-2 \rrbracket$, Eq. (41) holds with $\gamma_0 = \frac{1}{\sqrt{d}}$.

Proof. One can write $g_k = \exp\left(\frac{\gamma_0 \lambda_{k+1}}{2} \|\cdot\|^2\right)$. Simple calculation yields

$$\frac{\mathbb{E}_{\rho_k} g_k^2}{(\mathbb{E}_{\rho_k} g_k)^2} = \frac{\mathbb{E}_{\pi_0} \exp\left(-\frac{(1+\gamma_0)\lambda_{k+1}}{2} \|\cdot\|^2\right) \mathbb{E}_{\pi_0} \exp\left(-\frac{(1-\gamma_0)\lambda_{k+1}}{2} \|\cdot\|^2\right)}{\mathbb{E}_{\pi_0} \exp\left(-\frac{\lambda_{k+1}}{2} \|\cdot\|^2\right)^2}.$$

Define

$$h_2(\gamma) := \mathbb{E}_{\pi_0} \exp\left(-\frac{(1+\gamma)\lambda}{2} \|\cdot\|^2\right) \mathbb{E}_{\pi_0} \exp\left(-\frac{(1-\gamma)\lambda}{2} \|\cdot\|^2\right), \quad \gamma \in \left[0, \frac{1}{2}\right].$$

One can similarly show

$$\frac{d}{d\gamma} \log h_2(\gamma) = \frac{\lambda^2}{4} \int_{1-\gamma}^{1+\gamma} \text{Var}_{\tilde{\rho}_t} \|\cdot\|^2 dt,$$

where $\tilde{\rho}_t \propto \exp(-V_0 - \frac{t\lambda}{2} \|\cdot\|^2)$ is $(\beta + t\lambda)$ -strongly-log-concave and thus satisfies $\frac{1}{\beta+t\lambda}$ -LSI. Hence, $\text{Var}_{\tilde{\rho}_t} \|\cdot\|^2 \leq \frac{8}{\beta+t\lambda} \mathbb{E}_{\tilde{\rho}_t} \|\cdot\|^2$.

Let x''_t be the global minimizer of $V_0 + \frac{t\lambda}{2} \|\cdot\|^2$. By Lem. 16, $\|x''_t\| \leq \frac{R}{1+\frac{t\lambda}{3\beta}}$. Therefore,

$$\text{Var}_{\tilde{\rho}_t} \|\cdot\|^2 \lesssim \frac{1}{\beta + t\lambda} (\mathbb{E}_{\tilde{\rho}_t} \|\cdot - x''_t\|^2 + \|x''_t\|^2) \lesssim \frac{1}{\beta + t\lambda} \left(\frac{d}{\beta + t\lambda} + \frac{\beta^2 R^2}{(\beta + t\lambda)^2} \right).$$

As a result,

$$\begin{aligned} \frac{d}{d\gamma} \log h_2(\gamma) &\lesssim \lambda^2 \int_{1-\gamma}^{1+\gamma} \frac{1}{\beta + t\lambda} \left(\frac{d}{\beta + t\lambda} + \frac{\beta^2 R^2}{(\beta + t\lambda)^2} \right) dt \\ &\leq \lambda^2 \int_{1-\gamma}^{1+\gamma} \frac{1}{t\lambda} \left(\frac{d}{t\lambda} + \frac{\beta^2 R^2}{t^2 \lambda^2} \right) dt \\ &\lesssim \lambda^2 \gamma \cdot \frac{1}{\lambda} \left(\frac{d}{\lambda} + \frac{\beta^2 R^2}{\lambda^2} \right) = \gamma \left(d + \frac{\beta^2 R^2}{\lambda} \right) \\ \implies \log \frac{h_2(\gamma_0)}{h_2(0)} &\lesssim \gamma_0^2 \left(d + \frac{\beta^2 R^2}{\lambda} \right) = 1 + \frac{\beta^2 R^2}{d\lambda}. \end{aligned}$$

Since $\lambda_{k+1} \geq \lambda_{K-1} \asymp \frac{\beta}{\sqrt{d}}$ and $R \lesssim \frac{1}{\sqrt{\beta}}$, $\frac{\beta^2 R^2}{d\lambda_{k+1}} \lesssim 1$, so $\frac{\mathbb{E}_{\rho_k} g_k^2}{(\mathbb{E}_{\rho_k} g_k)^2} \leq 1 + O(1)$. □

Finally, one can compute the total complexity as follows. The choice $\lambda_0 \asymp \frac{d\beta}{\varepsilon}$, $\lambda_{K-1} \asymp \frac{\beta}{\sqrt{d}}$, and $\lambda_k = \left(1 + \frac{1}{\sqrt{d}}\right) \lambda_{k+1}$ implies $K = \tilde{\Theta}(\sqrt{d})$, and thus $N \asymp \frac{K}{\varepsilon^2} = \tilde{\Theta}\left(\frac{\sqrt{d}}{\varepsilon^2}\right)$. For each k , it is necessary to obtain N i.i.d. approximate samples from ρ_k that are $\delta \asymp \frac{1}{NK} = \tilde{\Theta}\left(\frac{\varepsilon^2}{d}\right)$ -close in TV distance. Using proximal sampler (Fan et al., 2023), the complexity for obtaining one sample is $\tilde{O}(\sqrt{d})$ (note that the condition numbers of f_k 's are uniformly bounded by 3), so the total oracle complexity is $NK \cdot \tilde{O}(\sqrt{d}) = \tilde{O}\left(\frac{d^{\frac{3}{2}}}{\varepsilon^2}\right)$. \square

H.3 PATH INTEGRAL SAMPLER AND CONTROLLED MONTE CARLO DIFFUSION

In this section, we briefly discuss two learning-based samplers used for normalizing constant estimation and refer readers to the original papers for detailed derivations. The path integral sampler (PIS) shares structural similarities with the RDS framework discussed in Thm. 5, using the time-reversal of a universal noising process that transforms any distribution into a prior – such as the OU process in RDS that converges to the standard normal or the Brownian bridge in PIS that converges to the delta distribution at zero. In contrast, the controlled Monte Carlo diffusion (CMCD) extends the JE framework from Sec. 4, focusing on learning the compensatory drift term along an arbitrary interpolating curve $(\pi_\theta)_{\theta \in [0,1]}$, as long as the density of each intermediate distribution π_θ is known up to a constant.

Path integral sampler (PIS, Zhang & Chen (2022)). The PIS learns the drift term of a reference SDE that interpolates the delta distribution at 0 and the target distribution π , which is closely connected with the Brownian bridge and the Föllmer drift (Chewi, 2022).

Fix a time horizon $T > 0$. For any drift term $(u_t)_{t \in [0,T]}$, let \mathcal{Q}^u be the path measure of the following SDE:

$$dX_t = u_t(X_t)dt + dB_t, \quad t \in [0, T]; \quad X_0 \stackrel{\text{a.s.}}{=} 0.$$

In particular, when $u \equiv 0$, the marginal distribution of X_T under \mathcal{Q}^0 is $\mathcal{N}(0, TI) =: \phi_T$. Define another path measure \mathcal{Q}^* by

$$\mathcal{Q}^*(d\xi_{[0,T]}) := \mathcal{Q}^0(d\xi_{[0,T]}|\xi_T)\pi(d\xi_T) = \mathcal{Q}^0(d\xi_{[0,T]})\frac{d\pi}{d\phi_T}(\xi_T), \quad \forall \xi \in C([0, T]; \mathbb{R}^d)$$

and consider the problem

$$u^* = \underset{u}{\operatorname{argmin}} \operatorname{KL}(\mathcal{Q}^u \| \mathcal{Q}^*) \implies \mathcal{Q}^{u^*} = \mathcal{Q}^*.$$

One can calculate the KL divergence between these path measures via Girsanov theorem (Lem. 1):

$$\log \frac{d\mathcal{Q}^u}{d\mathcal{Q}^*}(X) = W^u(X) + \log Z, \quad \text{a.s. } X \sim \mathcal{Q}^u, \quad \text{where}$$

$$W^u(X) = \int_0^T \langle u_t(X_t), dB_t \rangle + \frac{1}{2} \int_0^T \|u_t(X_t)\|^2 dt - \frac{\|X_T\|^2}{2T} + V(X_T) - \frac{d}{2} \log 2\pi T,$$

which implies $Z = \mathbb{E}_{\mathcal{Q}^u} e^{-W^u}$, and $\operatorname{KL}(\mathcal{Q}^u \| \mathcal{Q}^*) = \mathbb{E}_{\mathcal{Q}^u} W^u + \log Z$. On the other hand, directly applying Lem. 1 gives

$$\operatorname{KL}(\mathcal{Q}^u \| \mathcal{Q}^*) = \frac{1}{2} \int_0^T \mathbb{E}_{\mathcal{Q}^u} \|u_t(X_t) - u_t^*(X_t)\|^2 dt.$$

In Zhang & Chen (2022, Theorem 3), the authors considered the effective sample size (ESS) defined by $\operatorname{ESS}^{-1} = \mathbb{E}_{\mathcal{Q}^u} \left(\frac{d\mathcal{Q}^*}{d\mathcal{Q}^u}\right)^2$ as the convergence criterion, and stated that $\operatorname{ESS} \geq 1 - \varepsilon$ as long as $\sup_{t \in [0, T]} \|u_t - u_t^*\|_{L^\infty}^2 \leq \frac{\varepsilon}{T}$. However, this condition is generally hard to verify since the closed-form expression of u^* is unknown, and the L^∞ bound might be too strong. Using the criterion (Eq. (5)) and the same methodology in proving the convergence of JE (Thm. 2), we can establish

an improved result on the convergence guarantee of this estimator, relating the relative error to the training loss of u , which is defined as

$$\min_u L(u) := \mathbb{E}_{\mathcal{Q}^u} \left[\frac{1}{2} \int_0^T \|u_t(X_t)\|^2 dt - \frac{\|X_T\|^2}{2T} + V(X_T) \right] = \text{KL}(\mathcal{Q}^u \parallel \mathcal{Q}^*) - \log Z + \frac{d}{2} \log 2\pi T$$

Proposition 4. Consider the estimator $\hat{Z} := e^{-W^u(X)}$, $X \sim \mathcal{Q}^u$ for Z . To achieve both $\text{KL}(\mathcal{Q}_T^u \parallel \pi) \lesssim \varepsilon^2$ (with \mathcal{Q}_T^u representing the law of X_T in the sampled trajectory $X \sim \mathcal{Q}^u$) and $\Pr\left(\left|\frac{\hat{Z}}{Z} - 1\right| \leq \varepsilon\right) \geq \frac{3}{4}$, it suffices to choose u that satisfies

$$L(u) = -\log Z + \frac{d}{2} \log 2\pi T + O(\varepsilon^2).$$

Proof.

$$\Pr\left(\left|\frac{\hat{Z}}{Z} - 1\right| \geq \varepsilon\right) = \mathcal{Q}^u\left(\left|\frac{d\mathcal{Q}^*}{d\mathcal{Q}^u} - 1\right| \geq \varepsilon\right) \lesssim \frac{\text{TV}(\mathcal{Q}^u, \mathcal{Q}^*)}{\varepsilon} \lesssim \frac{\sqrt{\text{KL}(\mathcal{Q}^u \parallel \mathcal{Q}^*)}}{\varepsilon}.$$

Therefore, ensuring $\text{KL}(\mathcal{Q}^u \parallel \mathcal{Q}^*) \lesssim \varepsilon^2$ up to some sufficiently small constant guarantees that the above probability remains bounded by $\frac{1}{4}$. Furthermore, by the data-processing inequality, $\text{KL}(\mathcal{Q}_T^u \parallel \pi) \leq \text{KL}(\mathcal{Q}^u \parallel \mathcal{Q}^*) \lesssim \varepsilon^2$. \square

Controlled Monte Carlo Diffusion (CMCD, Vargas et al. (2024)). We borrow the notations from Sec. 4 due to its similarity with JE.

Given $(\tilde{\pi}_t)_{t \in [0, T]}$ and the ALD (Eq. (6)), we know from the proof of Thm. 1 that to make $X_t \sim \tilde{\pi}_t$ for all t , the compensatory drift term $(v_t)_{t \in [0, T]}$ must generate $(\tilde{\pi}_t)_{t \in [0, T]}$. Now, consider the task of learning such a vector field $(u_t)_{t \in [0, T]}$ by matching the following forward and backward SDEs:

$$\begin{aligned} \mathcal{P}^\rightarrow : dX_t &= (\nabla \log \tilde{\pi}_t + u_t)(X_t)dt + \sqrt{2}dB_t, \quad X_0 \sim \tilde{\pi}_0, \\ \mathcal{P}^\leftarrow : dX_t &= (-\nabla \log \tilde{\pi}_t + u_t)(X_t)dt + \sqrt{2}dB_t^\leftarrow, \quad X_T \sim \tilde{\pi}_T, \end{aligned}$$

where the loss is $\text{KL}(\mathcal{P}^\rightarrow \parallel \mathcal{P}^\leftarrow)$, discretized in training. Obviously, when trained to optimality, both \mathcal{P}^\rightarrow and \mathcal{P}^\leftarrow share the marginal distribution $\tilde{\pi}_t$ at every time t . By Girsanov theorem (Lem. 4), one can prove the following identity for a.s. $X \sim \mathcal{P}^\rightarrow$: $\log \frac{d\mathcal{P}^\rightarrow}{d\mathcal{P}^\leftarrow}(X) = W(X) + C^u(X) - \Delta F$, where ΔF and $W(X)$ are defined as in Thm. 1, and

$$C^u(X) := - \int_0^T (\langle u_t(X_t), \nabla \log \tilde{\pi}_t(X_t) \rangle + \nabla \cdot u_t(X_t)) dt.$$

We refer readers to Vargas et al. (2024, Prop. 3.3) for the detailed derivation. By $\mathbb{E}_{\mathcal{P}^\rightarrow} \frac{d\mathcal{P}^\leftarrow}{d\mathcal{P}^\rightarrow} = 1$, we know that $\mathbb{E}_{\mathcal{P}^\rightarrow} e^{-W(X) - C^u(X)} = e^{-\Delta F}$. As the paper has not established inference-time performance guarantee given the training loss, we prove the following result characterizing the relationship between the training loss and the accuracy of the sampled distribution as well as the estimated normalizing constant.

Proposition 5. Let $\hat{Z} = Z_0 e^{-W(X) - C^u(X)}$, $X \sim \mathcal{P}^\rightarrow$ be an unbiased estimator of $Z = Z_0 e^{-\Delta F}$. Then, to achieve both $\text{KL}(\mathcal{P}_T^\rightarrow \parallel \pi) \lesssim \varepsilon^2$ (where $\mathcal{P}_T^\rightarrow$ is the law of X_T in the sampled trajectory $X \sim \mathcal{P}^\rightarrow$) and $\Pr\left(\left|\frac{\hat{Z}}{Z} - 1\right| \leq \varepsilon\right) \geq \frac{3}{4}$, it suffices to choose u that satisfies $\text{KL}(\mathcal{P}^\rightarrow \parallel \mathcal{P}^\leftarrow) \lesssim \varepsilon^2$.

Proof. The proof of this theorem follows the same reasoning as that of Prop. 4. For normalizing constant estimation,

$$\Pr\left(\left|\frac{\hat{Z}}{Z} - 1\right| \geq \varepsilon\right) = \mathcal{P}^\rightarrow\left(\left|\frac{d\mathcal{P}^\leftarrow}{d\mathcal{P}^\rightarrow} - 1\right| \geq \varepsilon\right) \lesssim \frac{\text{TV}(\mathcal{P}^\rightarrow, \mathcal{P}^\leftarrow)}{\varepsilon} \lesssim \frac{\sqrt{\text{KL}(\mathcal{P}^\rightarrow \parallel \mathcal{P}^\leftarrow)}}{\varepsilon} \lesssim 1.$$

For sampling, the result is an immediate corollary of the data-processing inequality. \square

I CONCLUSION AND FUTURE WORK

In this paper, we analyzed the complexity of normalizing constant estimation using JE, AIS, and RDS, establishing non-asymptotic convergence guarantees based on insights from continuous-time analysis. Our analysis of JE (Thm. 2) applies to general interpolation curves without requiring explicit isoperimetric assumptions, which significantly extends prior work limited to log-concave distributions. While our main results (Thms. 2 and 4) provide upper complexity bounds, their tightness remains an open question. Deriving general lower bounds would further clarify whether curves with large action inherently require more oracle calls for both sampling and normalizing constant estimation, thereby rigorously validating the arguments in Sec. 6. We also conjecture that our proof techniques can be further extended to samplers beyond overdamped LD (e.g., Hamiltonian or underdamped LD (Sohl-Dickstein & Culpepper, 2012)), and may be applied to estimating normalizing constants of compactly supported distributions on \mathbb{R}^d (e.g., convex bodies volume estimation (Cousins & Vempala, 2018)) and discrete distributions (e.g., Ising model and restricted Boltzmann machines (Huber, 2015; Krause et al., 2020)) via the Poisson stochastic integral framework (Ren et al., 2025a;b). We leave these directions for future research.