



Estimation of Normalizing Constant (Partition Function, Free Energy)

Task: given an unnormalized probability density $\pi \propto e^{-V}$, estimate its normalizing constant (a.k.a. partition function) $Z = \int_{\mathbb{R}^d} e^{-V(x)} dx$ or free energy $F = -\log Z$.

It is a crucial problem in Bayesian statistics, statistical mechanics, and machine learning, and is challenging in high dimensions or when π is multimodal (i.e., V has a complex landscape):

Example 1. [Free energy difference.] Given two energy functions U_0, U_1 , estimate the free-energy difference $\Delta F := -\frac{1}{\beta} \log(\int e^{-\beta U_1} dx / \int e^{-\beta U_0} dx)$.

Example 2. [Marginal likelihood in latent variable models.] In latent variable models such as VAE, the marginal likelihood of a data point x is $p_\theta(x) = \int p_\theta(x|z)p(z)dz$, which is equivalent to the normalizing constant of the posterior distribution $p_\theta(z|x) \propto_z p_\theta(x|z)p(z)$.

Example 3. [Volume of convex bodies.] Estimate the volume of a convex body $\mathcal{K} \iff$ the normalizing constant of the uniform distribution on \mathcal{K} , $\pi \propto 1_{\mathcal{K}}$.

Importance sampling: with a prior distribution $\mu = \frac{1}{Z_\mu} e^{-U}$, we have $\frac{Z_\pi}{Z_\mu} = \frac{1}{Z_\mu} \int e^{-V} dx = \mathbb{E}_\mu \frac{e^{-V}}{e^{-U}}$. However, this estimator suffers from high variance due to the mismatch between μ and π .

Annealing for Addressing Multimodality

Annealing: construct a sequence of intermediate distributions that bridge the target and the prior distributions. This idea motivates several popular methods:

- In statistics: path sampling, **annealed importance sampling**, sequential Monte Carlo, etc.
- In statistical mechanics: thermodynamic integration, **Jarzynski equality**, etc.

Contributions: we aim to establish a rigorous **non-asymptotic** analysis of estimators based on JE and AIS, while introducing minimal assumptions on the target distribution. We also propose a new algorithm based on reverse diffusion samplers (RDS) to tackle a potential shortcoming of AIS.

Problem Setting

Criterion: given an accuracy threshold ε , study the complexity (measured by the number of calls to the oracles V and ∇V) required to obtain an estimator \hat{Z} of Z such that $\Pr\left(\left|\frac{\hat{Z}}{Z} - 1\right| \leq \varepsilon\right) \geq \frac{3}{4}$. Note that the probability can be boosted to any $1 - \zeta$ using the median trick.

Annealing curve: we define a curve of probability measures $(\pi_\theta = \frac{1}{Z_\theta} e^{-V_\theta})_{\theta \in [0,1]}$ from a prior distribution to the target distribution. $Z_1 = Z$ is what we need to estimate.

Assumptions

- Assump. 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 $\mathcal{A} := \int_0^1 |\dot{\pi}|_\theta^2 d\theta$.
- Assump. 2:** V is β -smooth, $\nabla V(0) = 0$, and $m := \sqrt{\mathbb{E}_\pi \|\cdot\|^2} < +\infty$.

Here, the **metric derivative** of $(\pi_\theta)_{\theta \in [0,1]}$ at θ is defined as $|\dot{\pi}|_\theta := \lim_{\delta \rightarrow 0} \frac{W_2(\pi_{\theta+\delta}, \pi_\theta)}{|\delta|}$, which can be interpreted as the “speed” of this curve. If $|\dot{\pi}|_\theta$ exists and is finite for a.e. $\theta \in [0, 1]$, we say that $(\pi_\theta)_{\theta \in [0,1]}$ is **absolutely continuous (AC)**. Its **action** is defined as $\int_0^1 |\dot{\pi}|_\theta^2 d\theta$, which is a key property characterizing the effectiveness of a curve in annealed sampling.

Analysis of the Jarzynski Equality (JE)

We introduce a reparameterized curve $(\tilde{\pi}_t = \pi_{t/T})_{t \in [0, T]}$ for some large T to be determined later

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

► Jarzynski Equality (JE) [Jar97]

Let \mathbb{P}^\rightarrow be the path measure of ALD. Then the following relation between the work functional W and free energy difference ΔF holds:

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

■ Theorem (Convergence Guarantee of JE)

$\hat{Z} := Z_0 e^{-W(X)}$ with $X \sim \mathbb{P}^\rightarrow$ is an unbiased estimator of $Z = Z_0 e^{-\Delta F}$. Under Assump. 1, it suffices to choose $T = \frac{32\mathcal{A}}{\varepsilon^2}$ to obtain $\Pr\left(\left|\frac{\hat{Z}}{Z} - 1\right| \leq \varepsilon\right) \geq \frac{3}{4}$.

■ Core Lemma (Relationship between Metric Derivative and Continuity Equation [AGS08])

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

Analysis of the Annealed Importance Sampling (AIS)

► Annealed Importance Sampling (AIS) Equality [Nea01]

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 $\mathbb{P}^\rightarrow(x_{0:M}) = \pi_0(x_0) \prod_{\ell=1}^M F_\ell(x_{\ell-1}, x_\ell)$. Then the same relation holds:

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

To study non-asymptotic complexity guarantees, we focus on the **geometric interpolation**:

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

where the annealing schedule $\lambda(\cdot)$ is a decreasing function with $\lambda(0) = 2\beta$ and $\lambda(1) = 0$.

Introduce discrete time points $0 = \theta_0 < \theta_1 < \dots < \theta_M = 1$, and define F_ℓ as running LD targeting π_{θ_ℓ} for time T_ℓ . In practice, we approximate this by running *one step* of **annealed Langevin Monte Carlo (ALMC)** using the exponential integrator discretization scheme with step size T_ℓ .

■ Theorem (Convergence Guarantee of AIS)

Under Assumps. 1 and 2, consider the annealing schedule $\lambda(\theta) = 2\beta(1 - \theta)^r$ for some $1 \leq r \lesssim 1$. We use \mathcal{A}_r to denote the action of $(\pi_\theta)_{\theta \in [0,1]}$. Then the oracle complexity for obtaining an estimate \hat{Z} that satisfies the criterion $\Pr\left(\left|\frac{\hat{Z}}{Z} - 1\right| \leq \varepsilon\right) \geq \frac{3}{4}$ is $\tilde{O}\left(\frac{d^3}{\varepsilon^2} \sqrt{\frac{m\beta\mathcal{A}_r^{\frac{1}{2}}}{\varepsilon^2}} \sqrt{\frac{d\beta^2\mathcal{A}_r^2}{\varepsilon^4}}\right)$.

Normalizing Constant Estimation via Reverse Diffusion Sampler (RDS)

The choice of the interpolation curve $(\pi_\theta)_{\theta \in [0,1]}$ is crucial for the complexity of JE and AIS.

The geometric interpolation is widely used in practice due to the availability of the scores of the distributions π_θ . However, for general target distributions, the action of the curve can be large:

■ Lemma (Exp. Lower Bound on the Action of Geometric Annealing from Gaussian Mixture)

Consider $\pi = \frac{1}{2} \mathcal{N}(0, 1) + \frac{1}{2} \mathcal{N}(m, 1)$ on \mathbb{R} for some large $m \gtrsim 1$, whose potential is $\frac{m^2}{2}$ -smooth. Under the setting in AIS, 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 $r \geq 1$. Then, the action of the curve $(\pi_\theta)_{\theta \in [0,1]}$ is $\mathcal{A}_r \gtrsim m^4 e^{\frac{m^2}{40}}$.

Reverse diffusion samplers (RDS): a series of multimodal samplers inspired by diffusion models. The OU process $dY_t = -Y_t dt + \sqrt{2} dB_t$, $t \in [0, T]$; $Y_0 \sim \pi$ transforms any target distribution π into $\phi := \mathcal{N}(0, I)$ as $T \rightarrow \infty$. Let $Y_t \sim \bar{\pi}_t$. The time-reversal $(Y_t^{\leftarrow} := Y_{T-t} \sim \bar{\pi}_{T-t})_{t \in [0, T]}$ satisfies the SDE $dY_t^{\leftarrow} = (Y_t^{\leftarrow} + 2\nabla \log \bar{\pi}_{T-t}(Y_t^{\leftarrow})) dt + \sqrt{2} dW_t$, $t \in [0, T]$; $Y_0^{\leftarrow} \sim \bar{\pi}_T (\approx \phi)$.

We propose leveraging the curve along the OU process for normalizing constant estimation. To support this idea, we first present the following proposition:

■ Proposition (Polynomial Upper Bound of the Action of the Curve along the OU Process)

Let $\bar{\pi}_t$ be the law of Y_t in the OU process 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, $\int_0^\infty |\dot{\bar{\pi}}|_t^2 dt \leq d\beta + m^2$.

■ Theorem (a Framework for Normalizing Constant Estimation via RDS)

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:

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

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 $\Pr\left(\left|\frac{\hat{Z}}{Z} - 1\right| \leq \varepsilon\right) \geq \frac{3}{4}$, it suffices that $\text{KL}(\mathbb{Q} \|\mathbb{Q}^\dagger) \lesssim \varepsilon^2$, $\text{TV}(\pi, \bar{\pi}_\delta) \lesssim \varepsilon$.

We can use results in [HDH⁺24, HZD⁺24, HRT24, VCK25] to derive the total complexity.

References

- [AGS08] 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.
- [HDH⁺24] Xunpeng Huang, Hanze Dong, Yifan Hao, Yian Ma, and Tong Zhang. Reverse diffusion Monte Carlo. In *The Twelfth International Conference on Learning Representations*, 2024.
- [HRT24] 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.
- [HZD⁺24] 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, editors, *Proceedings of Thirty Seventh Conference on Learning Theory*, volume 247 of *Proceedings of Machine Learning Research*, pages 2438–2493. PMLR, 30 Jun–03 Jul 2024.
- [Jar97] Christopher Jarzynski. Nonequilibrium equality for free energy differences. *Phys. Rev. Lett.*, 78:2690–2693, Apr 1997.
- [Nea01] Radford M. Neal. Annealed importance sampling. *Statistics and Computing*, 11(2):125–139, April 2001.
- [VCK25] Adrien Vacher, Omar Chehab, and Anna Korba. Sampling from multi-modal distributions with polynomial query complexity in fixed dimension via reverse diffusion. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.