

Estimating Normalizing Constants for Log-Concave Distributions: Algorithms and Lower Bounds

Rong Ge*, Holden Lee†, Jianfeng Lu‡

November 4, 2019

Abstract

Estimating the normalizing constant of an unnormalized probability distribution has important applications in computer science, statistical physics, machine learning, and statistics. In this work, we consider the problem of estimating the normalizing constant $Z = \int_{\mathbb{R}^d} e^{-f(x)} dx$ to within a multiplication factor of $1 \pm \varepsilon$ for a μ -strongly convex and L -smooth function f , given query access to $f(x)$ and $\nabla f(x)$. We give both algorithms and lowerbounds for this problem. Using an annealing algorithm combined with a multilevel Monte Carlo method based on underdamped Langevin dynamics, we show that $\tilde{\mathcal{O}}\left(\frac{d^{4/3}\kappa + d^{7/6}\kappa^{7/6}}{\varepsilon^2}\right)$ queries to ∇f are sufficient, where $\kappa = L/\mu$ is the condition number. Moreover, we provide an information theoretic lowerbound, showing that at least $\frac{d^{1-o(1)}}{\varepsilon^{2-o(1)}}$ queries are necessary. This provides a first nontrivial lowerbound for the problem.

*Duke University, Computer Science Department rongge@cs.duke.edu

†Duke University, Mathematics Department holee@math.duke.edu

‡Duke University, Mathematics Department jianfeng@math.duke.edu

Contents

1	Introduction	3
1.1	Notation and assumptions	4
1.2	Roadmap	5
2	Related works	5
3	Annealing for Estimating the Normalizing Constant	6
4	Estimating the Normalizing Constant using Multilevel ULD	8
4.1	Insufficiency of ULD	9
4.2	Multilevel ULD	9
4.3	Technical issues	11
5	Lowerbound on Number of Queries	11
6	Conclusion and Future Work	13
A	Proofs for Annealing Strategy	16
B	Estimating the Normalizing Constant using MALA and Annealing	18
C	Estimating the Normalizing Constant using Multilevel Langevin	21
C.1	The multilevel estimate	21
C.2	Multilevel ULD	26
C.3	Multilevel ULD-RMM	29
C.4	Truncation error and bias	33
C.5	Estimating the normalizing constant	37
D	Proof of Lowerbound	40
E	Quadrature Method for Estimating the Normalizing Constant	43
F	Tools and Auxiliary Lemmas	44

1 Introduction

Given a distribution ρ on a space Ω with base measure dx , defined by $\rho(dx) \propto e^{-f(x)} dx$, its normalizing constant is the integral $Z := \int_{\Omega} e^{-f(x)} dx$. Estimating the normalizing constant is a fundamental problem in theoretical computer science, statistical physics (where it is called the partition function [Bal07, SR⁺10]), and Bayesian statistics [GM98]. In high dimensional settings, even when the function $f(x)$ is convex (and the distribution ρ is log-concave), computing the exact normalizing constant is $\#P$ -hard [DF88]. Hence the goal is to approximate the normalizing constant up to a $1 \pm \varepsilon$ multiplicative accuracy. Approximating the normalizing constant is closely related to the problem of sampling from the distribution ρ [JV86, SJ89, DFK91].

Many polynomial time algorithms, starting from the seminal work of [DFK91], were known for estimating normalizing constants in various settings when $f(x)$ is convex. In this paper, we consider the special case where $\Omega = \mathbb{R}^d$ and $f(x)$ is a L -smooth and μ -strongly convex function (see equation (2)). Given query access to $f(x)$ and/or $\nabla f(x)$, our goal is to estimate the normalizing constant

$$Z = \int_{\mathbb{R}^d} e^{-f(x)} dx \quad (1)$$

within a multiplicative factor of $1 \pm \varepsilon$ with probability more than $3/4$ ¹.

This is a classical setting with applications to Bayesian statistics and machine learning. It is simpler than some of the settings considered before (such as volume estimation) because of strong convexity. Indeed, many faster sampling algorithms are known when f is strongly convex. However, there are very few results for estimating the normalizing constant and they give suboptimal dependencies. On the lowerbound side, although lowerbounds were considered in different settings (e.g., [RV08]), there are no non-trivial lowerbounds when f is strongly convex. In this paper, we give a new algorithm that only requires $\tilde{\mathcal{O}}\left(\frac{d^{\frac{4}{3}}\kappa+d^{\frac{7}{6}}\kappa^{\frac{7}{6}}}{\varepsilon^2}\right)$ queries to $\nabla f(x)$, as well as a lowerbound that shows no algorithm can succeed with $\frac{d^{1-o(1)}}{\varepsilon^{2-o(1)}}$ queries.

In high dimensions, most existing works rely on combining sampling algorithms for log-concave distributions and an annealing procedure. Our algorithm follows a similar recipe. We can use several sampling algorithms including Metropolis-Adjusted Langevin Algorithm (MALA), Underdamped Langevin Diffusion (ULD) and randomized midpoint method for ULD (ULD-RMM). However, a naïve combination of ULD and ULD-RMM with standard annealing procedure results in high query complexity. We use an approach called multilevel Monte Carlo [Gil08, GNS⁺16] to improve the query complexity and running time of the algorithm.

Theorem 1.1 (Upper bound). *Suppose $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is μ -strongly convex and L -smooth, and let $\kappa = \frac{L}{\mu}$. Consider the problem of estimating $\int_{\mathbb{R}^d} e^{-f(x)} dx$ within $1 \pm \varepsilon$ with success probability $3/4$,*

1. *Algorithm 2 (annealing with MALA), solves the problem with $\tilde{\mathcal{O}}\left(\frac{d^2\kappa}{\varepsilon^2} \max\{1, \sqrt{\kappa/d}\}\right)$ queries (Theorem B.3).*
2. *Algorithm 6 (annealing with multilevel Monte Carlo) run using Algorithm 4 (ULD), solves the problem with $\tilde{\mathcal{O}}\left(\frac{d^{\frac{3}{2}}\kappa^2}{\varepsilon^2}\right)$ queries (Theorem C.12).*
3. *Algorithm 6 (annealing with multilevel Monte Carlo) run using Algorithm 5 (ULD-RMM), solves the problem with $\tilde{\mathcal{O}}\left(\frac{d^{\frac{4}{3}}\kappa+d^{\frac{7}{6}}\kappa^{\frac{7}{6}}}{\varepsilon^2}\right)$ queries (Theorem C.13).*

¹For an algorithm, the probability can be easily amplified to $1 - \zeta$ by repeating the algorithm $\mathcal{O}(\log(1/\zeta))$ times and finding the median.

Note that these algorithms are also computationally efficient: for all of these algorithms, the runtime (in terms of number of vector operations in \mathbb{R}^d) is comparable to the number of queries. On the way to proving this theorem, we establish improved rates for estimating an expected value of a function using multilevel ULD. This result may be of independent interest.

Theorem 1.2 (Multilevel ULD). *Let $\rho(dx) \propto e^{-f(x)} dx$, where $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is μ -strongly convex and L -smooth. Let $g : \mathbb{R}^d \rightarrow \mathbb{R}$ be L_g -Lipschitz. Suppose $0 < \varepsilon < \frac{L_g}{\sqrt{\mu}}$. Consider the problem of outputting \hat{R} such that $|\hat{R} - \mathbb{E}_\rho[g(x)]| \leq \varepsilon$. With probability at least $\frac{3}{4}$, Algorithm 3 (Multilevel Monte Carlo) has the following guarantees:*

1. When run using Algorithm 4 (ULD), it succeeds using $\tilde{\mathcal{O}}\left(\frac{L_g^2 d^{\frac{1}{2}} \kappa^2}{\mu \varepsilon^2}\right)$ queries (Theorem C.4).
2. Using Algorithm 5 (ULD-RMM), it succeeds using $\tilde{\mathcal{O}}\left(\frac{L_g^2 (d^{\frac{1}{3}} \kappa + d^{\frac{1}{6}} \kappa^{\frac{7}{6}})}{\mu \varepsilon^2}\right)$ queries (Theorem C.6).

Intuitively, the multi-level Monte Carlo method is a way to reduce the variance of the final sample by coupling several different Markov chains at different step sizes, which reduces the number of queries when the running time of the sampling algorithm depends polynomially on the desired accuracy (see Section 4 for more details).

We also give the first lowerbound for the complexity of estimating the normalizing constant:

Theorem 1.3 (Lower bound). *Even for an L -smooth and μ -strongly convex function $f(x)$ with $\kappa = L/\mu$ being a constant, any algorithm that uses $\frac{d^{1-o(1)}}{\varepsilon^{2-o(1)}}$ queries cannot estimate the normalizing constant of $f(x)$ with accuracy $(1 \pm \varepsilon)$ with probability more than $3/4$.*

Our lowerbound matches the dependency on ε in high dimensions (note that this is impossible in low dimensions due to deterministic quadrature methods; see Appendix E). The lowerbound also shows that there is an inherent dependency on dimension d even when the condition number is a constant, which makes the problem of estimating the normalizing constant different from optimization. The lowerbound is information theoretic. We construct a function with many independent cells with two types. The final normalizing constant depends on the relative fraction of the cells of type 2. Making one query to function f can reveal the type of at most one cell; therefore a standard argument shows estimating the frequencies of cell-types requires a large number of queries.

1.1 Notation and assumptions

For any function f , we let $\mathcal{O}(f)$ and $\Omega(f)$ denote the class of functions that are $\leq Cf$ and $\geq Cf$, respectively, for some constant $C > 0$. Let $\tilde{\mathcal{O}}(f)$ denote the class $\mathcal{O}(f) \cdot \log^{\mathcal{O}(1)}(f)$, and $\tilde{\Omega}(f)$ denote the class $\Omega(f) \cdot \log^{-\mathcal{O}(1)}(f)$. Let $\Theta(f)$ denote the class of functions that are both $\mathcal{O}(f)$ and $\Omega(f)$, and $\tilde{\Theta}(f)$ denote the class of functions that are both $\tilde{\mathcal{O}}(f)$ and $\tilde{\Omega}(f)$.

For a vector $v \in \mathbb{R}^d$, let $\|v\|$ denote its Euclidean norm; and for a matrix $A \in \mathbb{R}^{d \times d}$, $\|A\|$ denotes its spectral norm. For $x, y \in \mathbb{R}$, let $x \wedge y = \min\{x, y\}$ and $x \vee y = \max\{x, y\}$.

The p th Wasserstein distance between two probability measures μ and ν is defined as

$$W_p(\mu, \nu) = \left(\inf_{(X, Y) \in \mathcal{C}(\mu, \nu)} \mathbb{E}[\|X - Y\|^p] \right)^{\frac{1}{p}}$$

where $\mathcal{C}(\mu, \nu)$ denotes the set of couplings between μ and ν . The TV-distance is defined as $d_{\text{TV}}(\mu, \nu) = \sup_A |\mu(A) - \nu(A)|$, where the sup is over all measurable subsets.

Throughout this work, we consider a log-concave distribution $\rho(dx) = \frac{1}{Z} e^{-f(x)} dx$. We assume that the negative log-density function $f(x)$ is twice continuously differentiable, μ -strongly convex and L -smooth: For all $x, y \in \mathbb{R}^d$,

$$\frac{\mu}{2} \|x - y\|^2 \leq f(y) - f(x) - \nabla f(x)^\top (y - x) \leq \frac{L}{2} \|x - y\|^2. \quad (2)$$

As we are concerned about the relative error for estimating the normalizing constant Z , it does not matter if f is shifted by a constant, and hence for simplicity of the presentation, we will assume that f achieves its global minimum at x^* with $f(x^*) = 0$ and only consider the (most challenging) regime that $\mu \ll 1 \ll L$. In fact, to further simplify the presentation, we will also assume $x^* = 0$, *i.e.*, f achieves the minimum at the origin. In practice, we do not know x^* a priori, however, using a first-order optimization method like gradient descent, we can obtain an approximate of x^* within error η using $\kappa \log(1/\eta)$ gradient evaluations. Such cost is negligible compared with other parts of the algorithm.

1.2 Roadmap

First in Section 2 we review existing works on sampling and estimating normalizing constant; in particular we recall guarantees for the sampling algorithms that we use in this paper. In Section 3 we describe the annealing strategy that we use, which is similar to but has different parameters with existing work. We describe the main idea of our algorithm (especially the idea of using the multilevel Monte Carlo method) in Section 4. Then we give the main ideas for the lowerbound in Section 5. Detailed algorithms and proofs are deferred to the appendices.

2 Related works

Many methods have been developed over the years for estimating the normalizing constant (also known as the partition function), see e.g., [GM98, SR⁺10] and references therein. However, not many works have given non-asymptotic rates for algorithms to estimate the normalizing constant of a strongly log-concave distribution. The closest work to ours is the recent work [BDM18], which gives a $\tilde{\mathcal{O}}(\kappa^3 d^3 \varepsilon^{-4})$ upperbound. An upperbound with a rather high power dependence on d is also established in [ARW16] for a different algorithm. Our algorithm and analysis yield much better bounds in comparison.

The estimation of the normalizing constant for a log-concave distribution is closely related to volume computation of a convex set K [DFK91, LS93, LV06] (which can be thought of as the special case where $f(x) = 0$ on K and $f(x) = \infty$ outside of K). This can be done in $\tilde{\mathcal{O}}(\frac{d^3}{\varepsilon^2})$ time [CV18] using an annealing algorithm combined with the Metropolis ball walk. While our setup is quite different, the overall annealing algorithm follows the same spirit, albeit with different parameter choices.

To the best of our knowledge, no lowerbound is known for the problem under consideration. For volume computation of convex set, the best known query lowerbound is $\tilde{\Omega}(d^2)$ given by [RV08] when $\varepsilon = \Theta(1)$. The results are not comparable as the volume of convex body corresponds to a function f that is not strongly convex, and the query is of membership rather than gradient type.

Non-asymptotic error analysis for Monte Carlo sampling algorithms has received a lot of research focus in recent years. One popular type of sampling algorithm is based on the Langevin dynamics,

either the underdamped Langevin dynamics

$$\begin{aligned} dx_t &= v_t dt; \\ dv_t &= -\nabla f(x_t) dt - \gamma v_t dt + \sqrt{2\gamma} dB_t, \end{aligned}$$

where $\gamma > 0$ is a friction parameter and each component of $B_t \in \mathbb{R}^d$ is independent standard Brownian motion, or the overdamped version (which can be obtained by taking $\gamma \rightarrow \infty$ of the underdamped Langevin while rescaling time $t \mapsto t/\gamma$):

$$dx_t = -\nabla f(x_t) dt + \sqrt{2} dB_t.$$

After discretization of the SDE by a numerical integration scheme, the overdamped Langevin dynamics leads to the unadjusted Langevin algorithms, whose explicit non-asymptotic error bounds have been established by recent works [Dal17, DM⁺17, DK17, DMM19, VW19], with complexity $\tilde{\mathcal{O}}\left(\frac{\kappa d}{\mu \varepsilon^2}\right)$ to achieve Wasserstein-2 error ε [DMM19]. The dependence on d and ε can be improved by sampling algorithms based on discretizing the underdamped Langevin dynamics, which has been recently pursued by [CCBJ17, DRD18, MCC⁺19, MFWB19, SL19]. In particular, the very recent work [SL19] gives an upperbound of query complexity $\tilde{\mathcal{O}}\left(\max\left\{\frac{d^{1/3}\kappa}{\mu^{1/3}\varepsilon^{2/3}}, \frac{d^{1/6}\kappa^{7/6}}{\mu^{1/6}\varepsilon^{1/3}}\right\}\right)$ for the ULD-RMM algorithm, upon which we will base our algorithm for the normalizing constant.

Metropolis-Hastings acceptance/rejection can be applied on top of the unadjusted Langevin algorithm. The resulting algorithm is known as Metropolis-Adjusted Langevin algorithm (MALA) [RT96], which was in fact first developed in the chemistry literature known as the smart Monte Carlo algorithm [RDF78]. The non-asymptotic error bound for MALA for log-concave probability distribution was recently studied by [DCWY18, CDWY19]. The result indicates that $\mathcal{O}(\kappa d \log(1/\varepsilon))$ queries to f and ∇f are needed to achieve error ε measured in total variation (TV) distance. Thus using Metropolis-Hastings acceptance/rejection improves the sampling efficiency exponentially in terms of the error ε , but suffers a worse dependence on d .

Besides the Langevin dynamics, sampling algorithms based on the deterministic Hamiltonian dynamics have been also quite popular, known as the Hamiltonian Monte Carlo (HMC) algorithms or hybrid Monte Carlo algorithms originally proposed in [DKPR87]; see also the review [BRSS18]. The non-asymptotic error analysis has been considered recently in [MS17, LSV18, LV18, CV19] for log-concave case and in [BREZ18] for more general cases using coupling arguments.

3 Annealing for Estimating the Normalizing Constant

For estimating the normalizing constant Z , we consider an annealing algorithm similar to previous algorithms for normalization constant estimation (see e.g., [LV06, CV18, BDM18]), but with a choice of annealing sequence different from that in the literature. Similar annealing strategies are widely used in calculation of normalizing constants, such as the annealed importance sampling [Nea01] in the statistic literature and thermodynamic integration [Jar97] in the statistical physics literature.

We define a sequence of auxiliary distributions, given by adding a quadratic function to f , for $i = 1, 2, \dots, M$

$$f_i(x) = \frac{1}{2} \frac{\|x\|^2}{\sigma_i^2} + f(x), \quad (3)$$

where $\sigma_1 \leq \sigma_2 \leq \dots \leq \sigma_M$; for convenience of notation, we also define $\sigma_{M+1} = \infty$ so that $f_{M+1} = f$. Correspondingly, we consider the sequence of distributions

$$\rho_i(dx) = Z_i^{-1} e^{-f_i(x)} dx, \quad (4)$$

where Z_i is the normalizing constant

$$Z_i = \int_{\mathbb{R}^d} e^{-f_i(x)} dx. \quad (5)$$

The estimation of Z is based on the identity

$$Z = Z_{M+1} = Z_1 \prod_{i=1}^M \frac{Z_{i+1}}{Z_i}. \quad (6)$$

In (6), we will approximate Z_1 by the normalizing factor of the Gaussian distribution with variance σ_1^2 . The ratio $\frac{Z_{i+1}}{Z_i}$ for $i = 1, \dots, M$ can be estimated using sampling algorithms for the distribution ρ_i , since

$$\frac{Z_{i+1}}{Z_i} = \int \exp\left(\frac{1}{2}\left(\frac{1}{\sigma_i^2} - \frac{1}{\sigma_{i+1}^2}\right)\|x\|^2\right) \rho_i(dx) = \mathbb{E}_{\rho_i}(g_i) \quad (7)$$

where

$$g_i := \exp\left(\frac{1}{2}\left(\frac{1}{\sigma_i^2} - \frac{1}{\sigma_{i+1}^2}\right)\|x\|^2\right). \quad (8)$$

Thus, if $X_i^{(1)}, \dots, X_i^{(K)}$ are iid sample points generated according to the distribution ρ_i (or its approximation), we can estimate

$$\frac{Z_{i+1}}{Z_i} \approx \frac{1}{K} \sum_{k=1}^K g_i(X_i^{(k)}). \quad (9)$$

For the sequence of σ_i^2 , we choose the following annealing strategy: We start with $\sigma_1^2 = \frac{\varepsilon}{2dL}$ and increase as

$$\sigma_{i+1}^2 = \sigma_i^2 \left(1 + \frac{1}{\sqrt{d}}\right) \quad (10)$$

until σ_M^2 is large enough, as specified below. We remark that a slower annealing procedure of $\sigma_{i+1}^2 = \sigma_i^2(1 + 1/d)$ was previously considered in [CV18], which gives a smaller relative variance of g_i for each stage (on the order of d^{-1}). We take a faster annealing procedure to take advantage of variance reduction by the multilevel Monte Carlo method, cf. Section 4.

In the above sketch of the algorithm, the approximation of Z_1 is guaranteed by the following lemma. Proofs of this and other lemmas in this section are postponed to Appendix A.

Lemma 3.1 (Starting distribution). *Letting $\sigma_1^2 = \frac{\varepsilon}{2dL}$, we have*

$$\left(1 - \frac{\varepsilon}{2}\right) \int_{\mathbb{R}^d} e^{-\frac{1}{2}\frac{\|x\|^2}{\sigma_1^2}} dx \leq Z_1 \leq \int_{\mathbb{R}^d} e^{-\frac{1}{2}\frac{\|x\|^2}{\sigma_1^2}} dx. \quad (11)$$

Next we consider the ratio $\frac{Z_{M+1}}{Z_M}$ in (6). We have

$$\frac{Z_{M+1}}{Z_M} = \int_{\mathbb{R}^d} \exp\left(\frac{\|x\|^2}{2\sigma_M^2}\right) \rho_M(dx) = \mathbb{E}_{\rho_M}(g_M) \quad (12)$$

with $g_M = \exp\left(\frac{\|x\|^2}{2\sigma_M^2}\right)$. To control the accuracy of Monte Carlo estimation of $\mathbb{E}_{\rho_M}(g_M)$, we bound the relative variance in the following lemma. The idea of the proof (deferred to Appendix A) comes from [CV18, Section 7.1], in particular the proof of [CV18, Lemma 7.6].

Lemma 3.2. For any $\sigma_M^2 \geq \frac{2}{\mu}$, we have

$$\frac{\mathbb{E}_{\rho_M}(g_M^2)}{\mathbb{E}_{\rho_M}(g_M)^2} = \mathbb{E}_\rho \exp\left(-\frac{1}{2} \frac{\|x\|^2}{\sigma_M^2}\right) \mathbb{E}_\rho \exp\left(\frac{1}{2} \frac{\|x\|^2}{\sigma_M^2}\right) \leq \exp\left(\frac{4d}{\mu\sigma_M^4}\right).$$

Let us now consider the estimate for $\frac{Z_{i+1}}{Z_i} = \mathbb{E}_{\rho_i}(g_i)$ in (6). To bound the variance of $g_i = \exp\left(\frac{1}{2}(\sigma_i^{-2} - \sigma_{i+1}^{-2})\|x\|^2\right)$ under the distribution ρ_i , let $\sigma^2 = \sigma_{i+1}^2$ and $\sigma_i^2 = \sigma^2/(1+\alpha)$, and calculate

$$\frac{\mathbb{E}_{\rho_i}(g_i^2)}{\mathbb{E}_{\rho_i}(g_i)^2} = \frac{\mathbb{E}_\rho \exp\left(-\frac{1+\alpha}{2} \frac{\|x\|^2}{\sigma^2}\right) \mathbb{E}_\rho \exp\left(-\frac{1-\alpha}{2} \frac{\|x\|^2}{\sigma^2}\right)}{\left(\mathbb{E}_\rho \exp\left(-\frac{1}{2} \frac{\|x\|^2}{\sigma^2}\right)\right)^2}. \quad (13)$$

The next lemma gives an upper bound for the right hand side as $\exp(4\alpha^2 d)$. This suggests the choice $\alpha = \frac{1}{\sqrt{d}}$ used in our annealing strategy to give an $\mathcal{O}(1)$ relative variance. The proof follows along similar lines as the previous lemma.

Lemma 3.3. Let ρ be a logconcave distribution, for $\alpha \leq \frac{1}{2}$, we have

$$\frac{\mathbb{E}_\rho \exp\left(-\frac{1+\alpha}{2} \frac{\|x\|^2}{\sigma^2}\right) \mathbb{E}_\rho \exp\left(-\frac{1-\alpha}{2} \frac{\|x\|^2}{\sigma^2}\right)}{\left(\mathbb{E}_\rho \exp\left(-\frac{1}{2} \frac{\|x\|^2}{\sigma^2}\right)\right)^2} \leq \exp(4\alpha^2 d) \quad (14)$$

With these lemmas, it remains to choose a suitable sampling scheme to estimate $\mathbb{E}_{\rho_i} g_i$ for each i . One possible approach is to use the Metropolis-Adjusted Langevin Algorithm (MALA) to generate independent samples with respect to ρ_i . Using the theoretical guarantees of MALA for strongly log-concave distributions recently established in [DCWY18, CDWY19], and the choice of $\sigma_M^2 = \Theta(\frac{\sqrt{d}}{\mu})$, we arrive at an algorithm with total query complexity $\tilde{\mathcal{O}}\left(\frac{d^2\kappa}{\varepsilon^2} \max\{1, \sqrt{\kappa/d}\}\right)$. This follows from the fact that MALA needs $\tilde{\mathcal{O}}(d\kappa \max\{1, \sqrt{\kappa/d}\})$ queries to achieve ε error in TV distance, and we need $\frac{\sqrt{d}}{\varepsilon^2}$ samples at each annealing stage to achieve relative variance $\tilde{\mathcal{O}}\left(\frac{\varepsilon^2}{\sqrt{d}}\right)$, which leads to relative variance $\mathcal{O}(\varepsilon^2)$ for the product, and thus $\mathcal{O}(\varepsilon)$ relative error. See Appendix B for details. The dimension dependence can however be improved by exploiting the multilevel Monte Carlo algorithm, as we discuss in the following section.

4 Estimating the Normalizing Constant using Multilevel ULD

Without making additional smoothness assumptions, for guarantees in KL or TV error, the best dependence on d known is the $\tilde{\mathcal{O}}(d)$ dependence given by MALA. However, for guarantees in Wasserstein (W_2) error, algorithms based on underdamped Langevin diffusion are known to give better dependence: [CCBJ17] show that to achieve W_2 error ε , underdamped Langevin dynamics (ULD) has query complexity $\tilde{\mathcal{O}}\left(\frac{d^{\frac{1}{2}}\kappa^2}{\mu^{\frac{1}{2}}\varepsilon}\right)$, and [DRD18] improves the dependence on κ to $\kappa^{\frac{3}{2}}$. [SL19] propose the Randomized Midpoint Method (RMM) to estimate the integral in ULD, and obtain query complexity $\tilde{\mathcal{O}}\left(\frac{d^{\frac{1}{2}}\kappa}{\varepsilon^{\frac{2}{3}}\mu^{\frac{1}{3}}} + \frac{d^{\frac{1}{6}}\kappa^{\frac{7}{6}}}{\mu^{\frac{1}{6}}\varepsilon^{\frac{1}{3}}}\right)$.

Focusing on the dependence on d and ε , one may hope that a method which obtains W_2 error using $\mathcal{O}\left(\frac{d^{\gamma}}{\varepsilon^\delta}\right)$ queries can be used to compute the normalizing constant in time $\mathcal{O}\left(\frac{d^{1+\gamma}}{\varepsilon^{2+\delta}}\right)$. However,

we show below that a naive substitution of the algorithm in the annealing procedure described in Section 3 fails. The key ingredient we need to obtain this $d^{1+\gamma}$ dependence is multilevel Monte Carlo, which additionally achieves $\frac{1}{\varepsilon^2}$ dependence in ε . This allows us to obtain the $\tilde{\mathcal{O}}\left(\frac{d^{\frac{3}{2}}\kappa^2}{\varepsilon^2}\right)$ and $\tilde{\mathcal{O}}\left(\frac{d^{\frac{4}{3}}\kappa+d^{\frac{7}{6}}\kappa^{\frac{7}{6}}}{\varepsilon^2}\right)$ rates in Theorem 1.1.

For simplicity, in the proof sketch below we assume the condition number and strong convexity are order 1 ($\kappa = \mathcal{O}(1)$, $\mu = \Theta(1)$), and focus on just the dependence on d and ε . In our main theorem we do work out the dependence on κ . We describe the guarantees that we would obtain by using ULD, but the same story holds for ULD-RMM with improved rates. For details, see Appendix C.

4.1 Insufficiency of ULD

Underdamped Langevin dynamics has the following error guarantee: to estimate the distribution up to W_2 -error ε , we can take step size $\eta = \mathcal{O}\left(\frac{\varepsilon}{\sqrt{d}}\right)$ and number of steps $\frac{T}{\eta} = \tilde{\mathcal{O}}\left(\frac{\sqrt{d}}{\varepsilon}\right)$.

Suppose we use $\tilde{\mathcal{O}}(\sqrt{d})$ temperatures, differing by factors of $1 + \frac{1}{\sqrt{d}}$. We chose the fewest number of temperatures such that the variance of $g_i(x)$ over ρ_i is $\mathcal{O}(1)$. (Using more temperatures, we need improved accuracy for estimating $R_i := \mathbb{E}_{x \sim \rho_i} g_i(x)$ for each temperature, which results in the same running time per temperature.) Then to estimate the normalizing constant within $1 \pm \mathcal{O}(\varepsilon)$, we need to estimate the ratio R_i at each step with relative accuracy $\frac{\varepsilon}{\sqrt{d}}$. We can check that $g_i(x) = \exp\left(\frac{\|x\|^2}{\sigma_i^2(1+\sqrt{d})}\right)$ is $\mathcal{O}\left(\frac{R_i}{\sigma_i}\right)$ -Lipschitz around where ρ_i is concentrated, that is, for x such that $\|x\| = \mathcal{O}(\sigma_i\sqrt{d})$. To estimate the product with ε relative accuracy, we need to estimate each R_i with $\mathcal{O}\left(\frac{\varepsilon R_i}{\sqrt{d}}\right)$ accuracy, so we need to sample from $\tilde{\rho}_i$ with $W_2(\tilde{\rho}_i, \rho_i) \leq \mathcal{O}\left(\frac{\varepsilon\sigma_i}{\sqrt{d}}\right)$. This requires us to choose a step size of $\eta = \mathcal{O}\left(\frac{\varepsilon/\sqrt{d}}{\sqrt{d}}\right) = \mathcal{O}\left(\frac{\varepsilon}{d}\right)$, so each sample takes $\tilde{\mathcal{O}}\left(\frac{d}{\varepsilon}\right)$ queries to obtain. In order to reduce the variance to $\frac{\varepsilon^2}{\sqrt{d}}$, we need $\frac{\sqrt{d}}{\varepsilon^2}$ samples at each temperature, for a total of $\tilde{\mathcal{O}}\left(\sqrt{d} \cdot \frac{d}{\varepsilon} \cdot \frac{\sqrt{d}}{\varepsilon^2}\right) = \mathcal{O}\left(\frac{d^2}{\varepsilon^3}\right)$ steps.

4.2 Multilevel ULD

Multilevel Monte Carlo [Gil08] is a generic way to improve rates for estimating $\mathbb{E}Y$ for a random variable Y , when there are biased estimators Y^η such that (1) as $\eta \rightarrow 0$, $\mathbb{E}Y^\eta \rightarrow \mathbb{E}Y$ and the cost to evaluate Y^η increases, and (2) there is a way to couple Y^η and $Y^{\eta'}$ when $\eta' < \eta$ that significantly reduces the variance, $\text{Var}(Y^\eta - Y^{\eta'}) \ll \text{Var}(Y^\eta)$.

This is the case when we wish to estimate $\mathbb{E}_{x \sim \rho} g(x)$, when ρ can be (approximately) obtained from simulating a stochastic differential equation (SDE) for some time T . In this setting, $Y^\eta = g(X^\eta)$ and $X^\eta = x_T^\eta$, where $x_T^\eta \sim \rho^\eta$ is the point obtained by simulating the SDE with some discretization algorithm \mathcal{A} for time T and step size η . Using the same Brownian motion for simulating x_t^η and $x_t^{\eta'}$ naturally defines a coupling. If g is L_g -Lipschitz, $\text{Var}(g(X^\eta) - g(X^{\eta'})) \leq L_g^2 \mathbb{E}[\|X^\eta - X^{\eta'}\|^2]$. The average distance $\mathbb{E}[\|X^\eta - X^{\eta'}\|^2]$ will be comparable to the Wasserstein error $W_2(\rho^\eta, \rho)$. This is much smaller than the variance of X^η , which is comparable to the variance of $X \sim \rho$.

The idea of Multilevel Monte Carlo (Algorithm 3) is to choose decreasing step sizes η_0, \dots, η_k

(e.g. with $\eta_j = \frac{\eta_0}{2^j}$), and write $g(X^{\eta_k})$ as

$$g(X^{\eta_k}) = g(X^{\eta_0}) + \sum_{j=1}^k [g(X^{\eta_j}) - g(X^{\eta_{j-1}})] \quad (15)$$

We estimate each of these terms by taking N_0 samples at the highest level $X_i^{\eta_0}$, and N_j coupled samples $(X_i^{\eta_j-}, X_i^{\eta_{j-1}+})$, to obtain the estimate

$$\hat{R} := \frac{1}{N_0} \sum_{i=1}^{N_0} g(X_i^{\eta_0}) + \sum_{j=1}^k \frac{1}{N_j} \sum_{i=1}^{N_j} [g(X_i^{\eta_j-}) - g(X_i^{\eta_{j-1}+})]. \quad (16)$$

Suppose we would like to give an estimate with bias ε_b and variance ε_σ^2 . The expected value of \hat{R} is simply $\mathbb{E}_{X^{\eta_k} \sim \rho^{\eta_k}} g(X^{\eta_k})$, so to ensure bias $\leq \varepsilon_b$, it suffices to choose η_k small enough. Supposing the variance of $g(X_i^{\eta_j-}) - g(X_i^{\eta_{j-1}+})$ is $F(\eta_j)$, the total variance is $\frac{\text{Var}(g(X^{\eta_0}))}{N_0} + \sum_{j=1}^k \frac{F(\eta_j)}{N_j}$. For smaller step size, because the variance $F(\eta_j)$ is smaller, it suffices to choose a smaller number of samples N_j , which offsets the increased number of steps $\frac{T}{\eta_j}$. Optimally choosing N_j to balance this with the total time necessary, $\sum_{j=1}^k \frac{TN_j}{\eta_j}$, gives the following.

Lemma (Lemma C.2 with $L_g = \sqrt{\mu}$, $F(\eta) = C\eta^\beta$). *Suppose that $\rho(dx) \propto e^{-f(x)}$, f is μ -strongly convex and $g : \mathbb{R}^d \rightarrow \mathbb{R}$ is $\sqrt{\mu}$ -Lipschitz. Suppose algorithm \mathcal{A} with step size η takes $\frac{T}{\eta}$ gradient queries to generate the random variable X^η . Let X^0 denote the corresponding continuous process. Suppose there is a coupling between X^η and X^0 such that $\mathbb{E}[|X^\eta - X^0|^2] \leq F(\eta) := C\eta^\beta$ (for some $\beta > 1$), and $T(\cdot)$ is a function such that $W_2(\rho^\eta, \rho)^2 \leq F(\eta) \wedge \varepsilon^2$ whenever $T \geq T(\varepsilon)$. Let η_0 be such that $F(\eta_0) = \frac{1}{\mu}$ and $F(\eta_k) \leq \frac{\varepsilon_b^2}{\mu}$. For $T \geq T\left(\frac{\varepsilon}{\sqrt{\mu}}\right)$ and appropriate number of samples N_j , multilevel Monte Carlo (Algorithm 3) run using \mathcal{A} returns an estimate \hat{R} of $\mathbb{E}_\rho g$ satisfying $|\mathbb{E}\hat{R} - \mathbb{E}_\rho g| \leq \varepsilon_b$ and $\text{Var}(\hat{R}) \leq \varepsilon_\sigma^2$ using $\mathcal{O}\left(T\left(\frac{1}{\varepsilon_\sigma^2 \eta_0} + \frac{1}{\eta_k}\right)\right)$ gradient queries.*

Note the scaling above is so that the variance of g over ρ is at most 1. Without multilevel Monte Carlo, the number of gradient queries would be significantly worse: $\mathcal{O}\left(\frac{T}{\eta_k} \cdot \frac{1}{\varepsilon_\sigma^2}\right)$, because we need to take a step size of η_k , and the number of samples to reduce the variance from 1 to ε_σ^2 is $\frac{1}{\varepsilon_\sigma^2}$. Using multilevel MC, we only need to pay $\mathcal{O}\left(\frac{1}{\varepsilon_\sigma^2}\right)$ samples at the highest level $k = 0$, and we only need to take η_0 small enough so that $F(\eta_0) = \frac{1}{\mu}$ (which makes $\text{Var}(g(X^{\eta_0})) \leq 1$).

We use this result to give a non-asymptotic analysis of the rate for multilevel ULD (Theorem C.4) and ULD-RMM (Theorem C.6). The results of [CCBJ17, DRD18] show that for underdamped Langevin dynamics, the hypotheses of the lemma hold with $F(\eta) = \mathcal{O}\left(\frac{d}{\mu}\eta^2\right)$, which suggests we take the largest step size to be $\eta_0 = \mathcal{O}(d^{-\frac{1}{2}})$. For ULD with the randomized midpoint method, [SL19] show that the hypotheses hold with $F(\eta) = \mathcal{O}\left(\frac{d}{\mu}\eta^3\right)$, which suggests we take $\eta_0 = \mathcal{O}\left(d^{-\frac{1}{3}}\right)$.

For the problem of estimating the normalizing constant, for each temperature i we apply Lemma C.2 with $g \leftarrow \frac{g_i}{R_i}$, which has Lipschitz constant $\mathcal{O}\left(\frac{1}{\sigma_i}\right) = \mathcal{O}(\sqrt{\mu_i})$ around where it is concentrated, where μ_i is the strong convexity constant of f_i . Then, to obtain bias $\varepsilon_b = \mathcal{O}\left(\frac{\varepsilon}{\sqrt{d}}\right)$ and variance $\varepsilon_\sigma^2 = \mathcal{O}\left(\frac{\varepsilon^2}{\sqrt{d}}\right)$, we need $\eta_k = \mathcal{O}\left(\frac{\varepsilon_b}{\sqrt{d}}\right)$ and so $\mathcal{O}\left(T\left(\frac{\sqrt{d}\sqrt{d}}{\varepsilon^2} + \frac{\sqrt{d}\sqrt{d}}{\varepsilon}\right)\right) = \tilde{\mathcal{O}}\left(\frac{d}{\varepsilon^2}\right)$ queries.

Since there are $\tilde{\mathcal{O}}(\sqrt{d})$ temperatures, the total number of queries over all temperatures is $\tilde{\mathcal{O}}\left(\frac{d^{3/2}}{\varepsilon^2}\right)$. Similarly for ULD-RMM, we need $\eta_k = \mathcal{O}\left(\frac{\varepsilon_b^{2/3}}{d^{1/3}}\right)$ and so $\mathcal{O}\left(T\left(\frac{d^{1/3}\sqrt{d}}{\varepsilon^2} + \frac{d^{1/3}\sqrt{d}}{\varepsilon^2}\right)\right)$ queries per temperature, and $\tilde{\mathcal{O}}\left(\frac{d^{4/3}}{\varepsilon^2}\right)$ queries in total.

Note that it is important to keep track of ε_σ and ε_b separately when computing the rates for multilevel MC. In our application, we can tolerate a larger ε_σ than ε_b at each temperature. This is because when there are M temperatures, when adding up the contributions from the different temperatures, the standard deviation will only be multiplied by \sqrt{M} , while the bias will be multiplied by M . This allowed us to take $\varepsilon_\sigma = \Theta\left(\frac{\varepsilon}{d^{1/4}}\right) \gg \varepsilon_b = \Theta\left(\frac{\varepsilon}{\sqrt{d}}\right)$. If we lowered ε_σ to make it equal to ε_b , then we need a factor of \sqrt{d} more samples for each temperature.

4.3 Technical issues

We glossed over several technical issues in the above proof sketch. First, we wish to estimate $\mathbb{E}_{x \sim \rho_i} g_i(x)$ where ρ_i is the distribution at the i th temperature and g_i is the ratio, but $g_i(x) = \exp\left(\frac{\|x\|^2}{\sigma_i^2(1+\alpha^{-1})}\right)$ is not Lipschitz. Instead, we truncate it for large x , and using concentration of $\|x\|$ on the log-concave distribution ρ_{i+1} to show that the bias introduced is small (Section C.4, Lemmas C.7 and C.8). More precisely, let $h_i(x) = g_i(x) \wedge \exp\left(\frac{r_i^{+2}}{\sigma_i^2(1+\alpha^{-1})}\right)$. We show that for some choice of $\alpha = \mathcal{O}\left(\frac{1}{\sqrt{d}\log(\frac{1}{\varepsilon})}\right)$, $r_i^+ = \mathbb{E}_{\rho_{i+1}} \|x\| + \Theta\left(\sigma_i \sqrt{(1+\alpha)\log\left(\frac{1}{\varepsilon}\right)}\right)$, we have (1) $\frac{h_i}{\mathbb{E}_{\rho_i} g_i}$ is $\mathcal{O}\left(\frac{1}{\sigma_i}\right)$ -Lipschitz, and (2) the bias introduced is small, $|\mathbb{E}_{\rho_i}(h - g)| \leq \varepsilon$.

We need to know at what radius r_i^+ to truncate g_i ; we can do this by estimating $\mathbb{E}_{x \sim \rho_{i+1}} \|x\|$ using samples and then adding a suitable multiple of σ_i (Lemma C.11). Finally, we put all the bounds together to prove the main Theorem C.12 for ULD and Theorem C.13 for ULD-RMM.

5 Lowerbound on Number of Queries

In this section, we give a lowerbound on the number of queries required to estimate the normalizing constant $\int e^{-f(x)} dx$. More precisely, we prove the following theorem:

Theorem 5.1. *For any fixed constant $\gamma > 0$, for large enough d , given query access to gradient or function value of a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ that is 1.5-smooth and 0.5-strongly convex, any algorithm that makes $o(d^{1-\gamma}\varepsilon^{-(2-\gamma)})$ queries cannot estimate the normalizing constant $Z = \int_{\mathbb{R}^d} e^{-f(x)} dx$ within a multiplicative factor of $1 \pm \varepsilon$ with probability more than 3/4.*

In fact, even if the algorithm is allowed to query any local information (such as the Hessian of f at x), our lowerbound still holds. Our construction also satisfies the Hessian Lipschitz property, which was used in some of the sampling results, see e.g., [DRD18, BDM18, MV18, LWME19]. Note that the bound hides constants that depend on γ , and d needs to be at least as large as $\Omega(1/\gamma)$. One might hope that $d\varepsilon^{-2}$ can be a lowerbound for every dimension d . However, this is impossible as when $d \leq 3$ quadrature methods give better dependency in terms of ε (see Appendix E).

To prove Theorem 5.1, we first construct a k -dimensional function (where $k = \Theta(1/\gamma)$), and show that any algorithm that estimates its normalizing constant requires at least $\Omega(\varepsilon^{-(2-\gamma)})$ queries. Then we construct the function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ in Theorem 5.1 by partitioning the d dimensions into d/k groups of size k , and use a product distribution whose marginal on each group corresponds to the function that we construct for the low-dimensional regime.

Lowerbound for low dimensions In low dimensions, our goal is to give a lowerbound that depends on the accuracy ε :

Theorem 5.2. *For any fixed integer $k > 0$, given query access to gradient or function value of a function $f : \mathbb{R}^k \rightarrow \mathbb{R}$ that is 1.5-smooth and 0.5-strongly convex, any algorithm that makes $o(\varepsilon^{-\frac{2}{1+4/k}})$ queries cannot estimate the normalizing constant $Z = \int_{\mathbb{R}^k} e^{-f(x)} dx$ within a multiplicative factor of $1 \pm \varepsilon$ with probability more than 3/4.*

Note that if we would like to get guarantee in terms of ε similar to Theorem 5.1 we only need to choose k such that $-\frac{2}{1+4/k} = -(2 - \gamma)$. It suffices to choose $k = \Theta(1/\gamma)$.

The main idea of proving this theorem is that we will construct a large number of independent “cells” in the space \mathbb{R}^k , where each cell can be one of two types. The final normalizing constant will depend on how many cells are of type 1. We will then pick a value δ (closely related to the accuracy ε) and consider two distributions of functions: in the first distribution, each cell is of type 1 with probability $1/2 + \delta$; in the second distribution, each cell is of type 1 with probability $1/2 - \delta$. When the number of cells is large enough (much more than $1/\delta^2$), the functions from these two distributions will have different normalizing constants (with large constant probability). However, making one query to the function at best gives information about a single cell. By a standard argument (see Claim D.1) we know in order to distinguish between two Bernoulli random variables with bias δ with better than $1/2$ probability, one needs at least $\Omega(1/\delta^2)$ queries. Any algorithm that uses fewer queries will not be able to distinguish the two distributions, and thus cannot estimate the normalizing constant accurately.

To construct these two distributions, we will start from a basic function $f_0(x) = \frac{\|x\|^2}{2}$. The normalizing constant for this function is well-known:

$$\int_{\mathbb{R}^k} e^{-f_0(x)} dx = (2\pi)^{k/2}.$$

To construct n cells, let $l = 1/(\sqrt{k}n^{1/k})$ (wlog we assume $n^{1/k}$ is an integer), and partition $[-1/\sqrt{k}, 1/\sqrt{k}]$ into $n^{1/k}$ intervals each of length $2l$. Let $I_i (i = 1, 2, \dots, n^{1/k})$ be the i -th interval. Each cell τ will be indicated by a k -tuple $(i_1, i_2, \dots, i_k) \in \{1, 2, \dots, n^{1/k}\}^k$, and the cell τ corresponds to $I_{i_1} \times I_{i_2} \times \dots \times I_{i_k}$ in \mathbb{R}^k .

Next we will discuss how to modify the function within the cells. For cell τ , We will modify the function to be $f_0(x) + cq(\frac{1}{l}(x - v_\tau))$ for x in the cell, where v_τ is the center of cell τ . Note that here the input $\frac{1}{l}(x - v_\tau)$ of q ranges in $[-1, 1]^k$. There are two major constraints for designing the function q : (1) it is possible to modify adjacent cells independently without violating the smoothness and strongly convex constraints; (2) it is possible to choose a large enough c such that $\int_{x \in \tau} \exp(-(f_0(x) + cq(\frac{1}{l}(x - v_\tau)))) dx$ is significantly smaller. The exact property of the q function and the construction is deferred to Lemma D.1 in Section D.

Now, we modify the functions within each cell by adding in a scaled version of q , as in the following lemma:

Lemma 5.3. *For any n where $n^{1/k}$ is an integer, and $l = 1/n^{1/k}$. For each cell $\tau = (i_1, \dots, i_k)$, let v_τ be its center. Construct the function $f(x)$ as*

$$f(x) = \begin{cases} f_0(x), & \text{cell } \tau \text{ is of type 1} \\ f_0(x) + c_\tau q\left(\frac{1}{l}(x - v_\tau)\right), & \text{cell } \tau \text{ is of type 2.} \end{cases}$$

Here q is the function constructed in Lemma D.1. There exists a way to choose c_τ ’s such that no matter what types each cell has, the family of functions satisfies the following properties:

1. $f(x)$ is 1.5-smooth and 0.5-strongly convex.
2. The normalizing constant $Z_f = \int_{\mathbb{R}^k} e^{-f(x)} dx = (2\pi)^{k/2} - C \frac{n_2}{n}$, where n_2 is the number of type-2 cells, and C is at least $\Omega(l^2)$.

With this lemma, one can construct two distributions of functions as follows: choose δ such that $\varepsilon = \Theta(\delta^{1+4/k})$, $n \approx 1/\delta^2$, and let each cell be of type 1 with probability $1/2 \pm \delta$ for the two classes. Claim D.1 shows that any algorithm that makes fewer than $o(1/\delta^2)$ queries cannot distinguish the two distributions, while Lemma 5.3 shows that the normalizing constant for two distributions differ by at least $1 + \Omega(l^2\delta)$ factor where $l = \Theta(n^{1/k}) = \Theta(\delta^{-2/k})$. This gives the desired trade-off in Theorem 5.2. A more detailed proof is given in Appendix D.

Lowerbound for high dimensions To generalize Theorem 5.2, as we mentioned earlier, we partition the d dimensions into d/k groups of size k , and use a product distribution. If we use S_i to denote the set of coordinates for the i -th group, we can write $f(x) = \sum_{i=1}^{d/k} f_i(x_{S_i})$. In particular, for the two distributions of functions that the algorithm is trying to distinguish, the $f_i(x_{S_i})$ are sampled from the two distributions of functions we defined for Theorem 5.2. Since the normalizing constant of $f(x)$ is equal to the product of normalizing constants for f_i 's, the gap between the two distributions is amplified by a power of $d/k = \Omega(d)$. Therefore, in order to achieve accuracy $1 \pm \varepsilon$ for function f , one would need to achieve an accuracy of $1 \pm \varepsilon k/d$ for functions f_i . On the other hand, one query in f can simultaneously give information on d/k of the functions f_i 's. Intuitively, if the lowerbound for the k dimensional case is $L(\varepsilon)$, the new lowerbound should be $L(\varepsilon k/d)/(d/k)$. Together with Theorem 5.2 and the choice $k = \Theta(1/\gamma)$, this gives the guarantee in Theorem 5.1. The detailed proof is given in Appendix D.

6 Conclusion and Future Work

In this paper, using multilevel Monte Carlo method we give a better algorithm for estimating the normalizing constant that only uses $\tilde{\mathcal{O}}\left(\frac{d^{4/3}\kappa+d^{7/6}\kappa^{7/6}}{\varepsilon^2}\right)$ queries to the gradient. We also give the first lowerbound that no algorithm can estimate the normalizing constant up to $1 \pm \varepsilon$ accuracy with $\frac{d^{1-o(1)}}{\varepsilon^{2-o(1)}}$ queries. For well-conditioned functions, the two bounds differ by $\mathcal{O}(d^{1/3+o(1)}\varepsilon^{-o(1)})$. Closing the gap is an immediate open problem, however we are not sure which side (if any) is tight. Any better rate for Langevin dynamics or related methods can give a better running time when combined with the multilevel Monte Carlo framework. On the other hand, improving our lowerbound might involve giving a lowerbound for sampling problems that depends on the dimension d .

There are many other settings where the idea of multilevel Monte Carlo may help improving the upperbound. This includes when only stochastic gradient queries are available (or when f is a sum of simpler functions). We note that multilevel methods can work with stochastic gradients as well [GNS⁺16], and variance reduction techniques are available [CFM⁺18]. It is an interesting question whether multilevel Langevin dynamics or multilevel hybrid Monte Carlo can improve running times for volume estimation of convex sets (like polytopes) [LV17, LV18], or smooth log-concave distributions restricted to convex sets.

Acknowledgements

HL would like to thank Oren Mangoubi for introducing multilevel methods. RG and HL would like to thank Ruoqi Shen for explaining the paper [SL19]. RG acknowledges funding from NSF CCF-1704656, NSF CCF-1845171 (CAREER), Sloan Fellowship and Google Faculty Research Award.

Part of the work was done while RG was visiting the Institute for Advanced Study. The work of JL is supported in part by National Science Foundation via grants DMS-1454939 and CCF-1934964.

References

- [ARW16] C. Andrieu, J. Ridgway, and N. Whiteley, *Sampling normalizing constants in high dimensions using inhomogeneous diffusions*, 2016. preprint, arXiv:1612.07583.
- [Bal07] Roger Balian, *From microphysics to macrophysics: methods and applications of statistical physics*, Vol. 2, Springer Science & Business Media, 2007.
- [BDM18] N. Brosse, A. Durmus, and E. Moulines, *Normalizing constants of log-concave densities*, Electron. J. Statist. **12** (2018), 851–889.
- [BÉ85] Dominique Bakry and Michel Émery, *Diffusions hypercontractives*, Séminaire de probabilités xix 1983/84, 1985, pp. 177–206.
- [BGL13] Dominique Bakry, Ivan Gentil, and Michel Ledoux, *Analysis and geometry of Markov diffusion operators*, Vol. 348, Springer Science & Business Media, 2013.
- [BREZ18] N. Bou-Rabee, A. Eberle, and R. Zimmer, *Coupling and convergence for Hamiltonian Monte Carlo*, 2018. preprint, arXiv:1805.00452.
- [BRSS18] N. Bou-Rabee and J. M. Sanz-Serna, *Geometric integrators and the Hamiltonian Monte Carlo method*, Acta Numer. **27** (2018), 113–206.
- [CCBJ17] Xiang Cheng, Niladri S Chatterji, Peter L Bartlett, and Michael I Jordan, *Underdamped langevin mcmc: A non-asymptotic analysis*, arXiv preprint arXiv:1707.03663 (2017).
- [CDWY19] Y. Chen, R. Dwivedi, M. J. Wainwright, and B. Yu, *Fast mixing of Metropolized Hamiltonian Monte Carlo: Benefits of multi-step gradients*, 2019. preprint, arXiv:1905.12247.
- [CFM⁺18] Niladri Chatterji, Nicolas Flammarion, Yian Ma, Peter Bartlett, and Michael Jordan, *On the theory of variance reduction for stochastic gradient Monte Carlo*, Proceedings of the 35th international conference on machine learning, 2018, pp. 764–773.
- [CV18] B. Cousins and S. Vempala, *Gaussian cooling and $O^*(n^3)$ algorithm for volume and Gaussian volume*, SIAM J. Comput. **47** (2018), 1237–1273. preprint, arXiv:1409.6011v3.
- [CV19] Zongchen Chen and Santosh S Vempala, *Optimal convergence rate of hamiltonian monte carlo for strongly logconcave distributions*, arXiv preprint arXiv:1905.02313 (2019).
- [Dal17] Arnak S Dalalyan, *Theoretical guarantees for approximate sampling from smooth and log-concave densities*, Journal of the Royal Statistical Society: Series B (Statistical Methodology) **79** (2017), no. 3, 651–676.
- [DCWY18] R. Dwivedi, Y. Chen, M. J. Wainwright, and B. Yu, *Log-concave sampling: Metropolis-Hastings algorithms are fast*, 2018. preprint, arXiv:1801.02309.
- [DF88] Martin E. Dyer and Alan M. Frieze, *On the complexity of computing the volume of a polyhedron*, SIAM Journal on Computing **17** (1988), no. 5, 967–974.
- [DFK91] Martin Dyer, Alan Frieze, and Ravi Kannan, *A random polynomial-time algorithm for approximating the volume of convex bodies*, Journal of the ACM (JACM) **38** (1991), no. 1, 1–17.
- [DK17] Arnak S Dalalyan and Avetik G Karagulyan, *User-friendly guarantees for the langevin monte carlo with inaccurate gradient*, arXiv preprint arXiv:1710.00095 (2017).
- [DKPR87] S. Duane, A. D. Kennedy, B. J. Pendleton, and D. Roweth, *Hybrid Monte Carlo*, Phys. Lett. B **195** (1987), 216–222.
- [DMM19] Alain Durmus, Szymon Majewski, and Błażej Miasojedow, *Analysis of langevin monte carlo via convex optimization*, Journal of Machine Learning Research **20** (2019), 1–46.
- [DM⁺17] Alain Durmus, Eric Moulines, et al., *Nonasymptotic convergence analysis for the unadjusted langevin algorithm*, The Annals of Applied Probability **27** (2017), no. 3, 1551–1587.
- [DRD18] A. S. Dalalyan and L. Riou-Durand, *On sampling from a log-concave density using kinetic Langevin diffusions*, 2018. preprint, arXiv:1807.09382.
- [Gil08] Michael B Giles, *Multilevel monte carlo path simulation*, Operations Research **56** (2008), no. 3, 607–617.

[GM98] A. Gelman and X.-L. Meng, *Simulating normalizing constants: From importance sampling to bridge sampling to path sampling*, Stat. Sci. **13** (1998), 163–185.

[GNS⁺16] Mike Giles, Tigran Nagapetyan, Lukasz Szpruch, Sebastian Vollmer, and Konstantinos Zygalakis, *Multi-level monte carlo for scalable bayesian computations*, arXiv preprint arXiv:1609.06144 (2016).

[Jar97] C. Jarzynski, *Nonequilibrium equality for free energy differences*, Phys. Rev. Lett. **78** (1997), 2690–2693.

[JVV86] Mark R Jerrum, Leslie G Valiant, and Vijay V Vazirani, *Random generation of combinatorial structures from a uniform distribution*, Theoretical Computer Science **43** (1986), 169–188.

[LM00] Beatrice Laurent and Pascal Massart, *Adaptive estimation of a quadratic functional by model selection*, Annals of Statistics (2000), 1302–1338.

[LS93] László Lovász and Miklós Simonovits, *Random walks in a convex body and an improved volume algorithm*, Random structures & algorithms **4** (1993), no. 4, 359–412.

[LSV18] Yin Tat Lee, Zhao Song, and Santosh S Vempala, *Algorithmic theory of odes and sampling from well-conditioned logconcave densities*, arXiv preprint arXiv:1812.06243 (2018).

[LV06] László Lovász and Santosh Vempala, *Simulated annealing in convex bodies and an $o^*(n^4)$ volume algorithm*, Journal of Computer and System Sciences **72** (2006), no. 2, 392–417.

[LV17] Yin Tat Lee and Santosh S Vempala, *Geodesic walks in polytopes*, Proceedings of the 49th annual ACM SIGACT symposium on theory of computing, 2017, pp. 927–940.

[LV18] ———, *Convergence rate of riemannian hamiltonian monte carlo and faster polytope volume computation*, Proceedings of the 50th annual ACM SIGACT symposium on theory of computing, 2018, pp. 1115–1121.

[LWME19] Xuechen Li, Denny Wu, Lester Mackey, and Murat A. Erdogdu, *Stochastic Runge-Kutta accelerates Langevin Monte Carlo and beyond*, 2019. preprint, arXiv:1906.07868.

[MCC⁺19] Yi-An Ma, Niladri Chatterji, Xiang Cheng, Nicolas Flammarion, Peter Bartlett, and Michael I Jordan, *Is there an analog of nesterov acceleration for mcmc?*, arXiv preprint arXiv:1902.00996 (2019).

[MFWB19] Wenlong Mou, Nicolas Flammarion, Martin J Wainwright, and Peter L Bartlett, *Improved bounds for discretization of langevin diffusions: Near-optimal rates without convexity*, arXiv preprint arXiv:1907.11331 (2019).

[MS17] Oren Mangoubi and Aaron Smith, *Rapid mixing of Hamiltonian Monte Carlo on strongly log-concave distributions*, arXiv preprint arXiv:1708.07114 (2017).

[MV18] Oren Mangoubi and Nisheeth Vishnoi, *Dimensionally tight bounds for second-order hamiltonian monte carlo*, Advances in neural information processing systems, 2018, pp. 6028–6038.

[Nea01] R. M. Neal, *Annealed importance sampling*, Statist. Comput. **11** (2001), 125–139.

[RDF78] P. J. Rossky, J. D. Doll, and H. L. Friedman, *Brownian dynamics as smart Monte Carlo simulation*, J. Chem. Phys. **69** (1978), 4628.

[RT96] G. O. Roberts and R. L. Tweedie, *Exponential convergence of Langevin distributions and their discrete approximations*, Bernoulli **2** (1996), no. 4, 341–363.

[RV08] L. Rademacher and S. Vempala, *Dispersion of mass and the complexity of randomized geometric algorithms*, Adv. Math. **219** (2008), 1037–1069.

[SJ89] Alistair Sinclair and Mark Jerrum, *Approximate counting, uniform generation and rapidly mixing markov chains*, Information and Computation **82** (1989), no. 1, 93–133.

[SL19] Ruoqi Shen and Yin Tat Lee, *The randomized midpoint method for log-concave sampling*, arXiv preprint arXiv:1909.05503 (2019).

[SR⁺10] Gabriel Stoltz, Mathias Rousset, et al., *Free energy computations: A mathematical perspective*, World Scientific, 2010.

[VW19] Santosh S Vempala and Andre Wibisono, *Rapid convergence of the unadjusted langevin algorithm: Log-sobolev suffices*, arXiv preprint arXiv:1903.08568 (2019).

A Proofs for Annealing Strategy

We provide proofs here for Lemmas in Section 3.

Proof of Lemma 3.1. Without loss of generality, we assume $x^* = 0$ (as it amounts to a change of variable $x \rightarrow x - x^*$ which does not affect the normalizing constant). The upper bound is obvious since $f(x) \geq 0$ by our assumption (recall that we only concern about the relative error for normalizing constant, so that shifting f by a constant has no impact). For the lower bound of Z_1 , note that $f(x) \leq \frac{1}{2}L\|x\|^2$, we have

$$\begin{aligned} Z_1 &= \int_{\mathbb{R}^d} e^{-f(x)-\frac{1}{2}\frac{\|x\|^2}{\sigma_1^2}} dx \\ &\geq \int_{\mathbb{R}^d} e^{-\frac{1}{2}(L+\sigma_1^{-2})\|x\|^2} dx = \left(2\pi(L+\sigma_1^{-2})^{-1}\right)^{d/2} \end{aligned} \quad (17)$$

Thus,

$$\frac{Z_1}{(2\pi\sigma_1^2)^{d/2}} = \left(1 + \sigma_1^2 L\right)^{-d/2} \geq e^{-d\sigma_1^2 L/2} \quad (18)$$

which is larger than $1 - \frac{\varepsilon}{2}$ for $\sigma_1^2 = \frac{\varepsilon}{2dL}$. \square

Proof of Lemma 3.2. Define

$$h(t) := \mathbb{E}_\rho e^{-t\|x\|^2} \mathbb{E}_\rho e^{t\|x\|^2}.$$

We have

$$\begin{aligned} \frac{h'(t)}{h(t)} &= \frac{\mathbb{E}_\rho(\|x\|^2 e^{t\|x\|^2})}{\mathbb{E}_\rho e^{t\|x\|^2}} - \frac{\mathbb{E}_\rho(\|x\|^2 e^{-t\|x\|^2})}{\mathbb{E}_\rho e^{-t\|x\|^2}} \\ &= \int_{-t}^t v'(s) ds, \end{aligned}$$

where

$$v(s) := \frac{\mathbb{E}_\rho(\|x\|^2 e^{s\|x\|^2})}{\mathbb{E}_\rho e^{s\|x\|^2}},.$$

Thus

$$\begin{aligned} v'(s) &= \frac{\mathbb{E}_\rho(\|x\|^4 e^{s\|x\|^2}) \mathbb{E}_\rho e^{s\|x\|^2} - \left(\mathbb{E}_\rho(\|x\|^2 e^{s\|x\|^2})\right)^2}{\left(\mathbb{E}_\rho e^{s\|x\|^2}\right)^2} \\ &= \text{Var}_{\rho_s}(\|x\|^2), \end{aligned}$$

where ρ_s is a distribution with $\frac{d\rho_s}{d\rho} \propto e^{s\|x\|^2}$. Since ρ is strongly log-concave with convexity parameter μ , ρ_s satisfies the Poincaré inequality with constant $1/(\mu - 2s) \leq 2/\mu$ for $s \leq \frac{1}{4}\mu$, thus

$$\text{Var}_{\rho_s}(\|x\|^2) \leq \frac{8}{\mu} \mathbb{E}_{\rho_s}(\|x\|^2) \leq \frac{16}{\mu} d, \quad (19)$$

where the last inequality follows from the concentration property of log-concave distribution.

Therefore,

$$\begin{aligned}
\ln h\left(\frac{1}{2\sigma_M^2}\right) &= \ln h(0) + \int_0^{1/(2\sigma_M^2)} \frac{h'(t)}{h(t)} dt \\
&\leq \int_0^{1/(2\sigma_M^2)} \int_{-t}^t \text{Var}_{\rho_s}(\|x\|^2) ds dt \\
&\leq \int_0^{1/(2\sigma_M^2)} \int_{-t}^t \frac{16d}{\mu} dt \\
&= \frac{4d}{\mu\sigma_M^4}
\end{aligned}$$

where in the last inequality we have used (19) that $s \leq 1/(2\sigma_M^2) \leq \frac{\mu}{4}$ by our assumption on σ_M^2 . Thus we arrive at

$$\mathbb{E}_\rho \exp\left(-\frac{1}{2} \frac{\|x\|^2}{\sigma_M^2}\right) \mathbb{E}_\rho \exp\left(\frac{1}{2} \frac{\|x\|^2}{\sigma_M^2}\right) \leq \exp\left(\frac{4d}{\mu\sigma_M^4}\right). \quad \square$$

Proof of Lemma 3.3. Define

$$h(\alpha) := \mathbb{E}_\rho \exp\left(-\frac{1+\alpha}{2} \frac{\|x\|^2}{\sigma^2}\right) \mathbb{E}_\rho \exp\left(-\frac{1-\alpha}{2} \frac{\|x\|^2}{\sigma^2}\right).$$

It follows then

$$\begin{aligned}
\frac{h'(\alpha)}{h(\alpha)} &= -\frac{1}{2\sigma^2} \left(\frac{\mathbb{E}_\rho \left(\|x\|^2 \exp\left(-\frac{1+\alpha}{2} \frac{\|x\|^2}{\sigma^2}\right) \right)}{\mathbb{E}_\rho \exp\left(-\frac{1+\alpha}{2} \frac{\|x\|^2}{\sigma^2}\right)} - \frac{\mathbb{E}_\rho \left(\|x\|^2 \exp\left(-\frac{1-\alpha}{2} \frac{\|x\|^2}{\sigma^2}\right) \right)}{\mathbb{E}_\rho \exp\left(-\frac{1-\alpha}{2} \frac{\|x\|^2}{\sigma^2}\right)} \right) \\
&= -\frac{1}{2\sigma^2} \int_{1-\alpha}^{1+\alpha} v'(t) dt,
\end{aligned}$$

where $v(t)$ is defined as

$$v(t) := \frac{\mathbb{E}_\rho \left(\|x\|^2 \exp\left(-\frac{t}{2} \frac{\|x\|^2}{\sigma^2}\right) \right)}{\mathbb{E}_\rho \exp\left(-\frac{t}{2} \frac{\|x\|^2}{\sigma^2}\right)}.$$

Explicit calculation gives

$$\begin{aligned}
v'(t) &= -\frac{1}{2\sigma^2} \frac{\mathbb{E}_\rho \left(\|x\|^4 \exp\left(-\frac{t}{2} \frac{\|x\|^2}{\sigma^2}\right) \right) \mathbb{E}_\rho \exp\left(-\frac{t}{2} \frac{\|x\|^2}{\sigma^2}\right) - \left(\mathbb{E}_\rho \left(\|x\|^2 \exp\left(-\frac{t}{2} \frac{\|x\|^2}{\sigma^2}\right) \right) \right)^2}{\left(\mathbb{E}_\rho \exp\left(-\frac{t}{2} \frac{\|x\|^2}{\sigma^2}\right) \right)^2} \\
&= -\frac{1}{2\sigma^2} \text{Var}_{\rho_t}(\|x\|^2).
\end{aligned}$$

Here ρ_t is the distribution given by

$$\frac{d\rho_t}{d\rho} \propto \exp\left(-\frac{t}{2} \frac{\|x\|^2}{\sigma^2}\right).$$

By the Poincarè inequality and concentration property of strongly log-concave measure

$$\text{Var}_{\rho_t}(\|x\|^2) \leq \frac{4\sigma^2}{t} \mathbb{E}_{\rho_t} \|x\|^2 \leq 8 \frac{\sigma^4}{t^2} d.$$

Therefore, we arrive at the inequality

$$\begin{aligned}
\frac{h'(\alpha)}{h(\alpha)} &= \frac{1}{4\sigma^4} \int_{1-\alpha}^{1+\alpha} \text{Var}_{\rho_t}(\|x\|^2) dt \\
&\leq 2d \int_{1-\alpha}^{1+\alpha} \frac{1}{t^2} dt \\
&= 2d \left(\frac{1}{1-\alpha} - \frac{1}{1+\alpha} \right) \\
&\leq 8d\alpha.
\end{aligned}$$

This gives

$$\ln h(\alpha) - \ln h(0) = \int_0^\alpha \frac{h'(\alpha)}{h(\alpha)} d\alpha \leq 4d\alpha^2.$$

Thus, we arrive that

$$\frac{h(\alpha)}{h(0)} \leq e^{4d\alpha^2},$$

which is the desired inequality by the definition of h . \square

B Estimating the Normalizing Constant using MALA and Annealing

Let us first recall the Metropolis adjusted Langevin algorithm (MALA) [RT96], Algorithm 1, which is a Metropolis-Hastings algorithm with the proposal step given by discretized overdamped Langevin diffusion.

Algorithm 1 Metropolis adjusted Langevin algorithm (MALA)

Input: Step size h and a sample x_0 from a starting distribution μ_0

Output: Sequence x_1, x_2, \dots

- 1: **for** $i = 0, 1, \dots$ **do**
- 2: Draw $z_{i+1} \sim \mathcal{N}(x_i - h\nabla f(x_i), 2hI)$
- 3: Compute $\alpha_{i+1} \leftarrow \min \left\{ 1, \frac{\exp(-f(z_{i+1}) - |x_i - z_{i+1} + h\nabla f(z_{i+1})|^2_2/(4h))}{\exp(-f(x_i) - |z_{i+1} - x_i + h\nabla f(x_i)|^2_2/(4h))} \right\}$
- 4: With probability α_{i+1} accept the proposal $x_{i+1} \leftarrow z_{i+1}$
- 5: With probability $1 - \alpha_{i+1}$ reject the proposal $x_{i+1} \leftarrow x_i$
- 6: **end for**

Following the recent theoretical analysis for MALA [DCWY18, CDWY19], we consider the $\frac{1}{2}$ -lazy version of MALA, namely, for each step, for probability $\frac{1}{2}$ one stays at the previous iterate and for probability $\frac{1}{2}$ one takes a MALA step. The laziness guarantees that the Markov chain is aperiodic and hence has a unique invariant measure, given by the target distribution thanks to the Metropolis acceptance-rejection step. The convergence of the empirical measure to the target measure has been established in [DCWY18, CDWY19], which we recall here:

Theorem B.1 ([CDWY19, Theorem 2]). *Assume the target distribution ρ is strongly log-concave with L -smooth and μ -strongly convex negative log-density. Then given the initial distribution $\rho_0 = \mathcal{N}(x^*, \frac{1}{L}I)$, the $\frac{1}{2}$ -lazy version of MALA with step size $h = c(Ld \max\{1, \sqrt{\kappa/d}\})^{-1}$ achieves*

$$d_{\text{TV}}(\rho_n, \rho) \leq \delta$$

for steps

$$n \geq Cd\kappa \log \frac{d}{\delta} \max\{1, \sqrt{\kappa/d}\},$$

where c and C above are universal constants.

The above Theorem assumes x^* , the minimum of $f(x)$. In practice, we do not know x^* a priori, however, using a first-order method like gradient descent, we can obtain an η -approximate mode \tilde{x} using $\kappa \log(1/\eta)$ gradient evaluations. If we instead take the initial distribution $\tilde{\rho}_0 = \mathcal{N}(\tilde{x}, \frac{1}{2L}I)$, the warmness parameter with respect to the target distribution becomes $\exp(\frac{d}{2} \log(2\kappa) + L\eta^2)$ instead of $\kappa^{d/2}$ for $\rho_0 = \mathcal{N}(x^*, \frac{1}{L}I)$. As discussed in [DCWY18, Section 3.2], with a slightly modified step size, the MALA sampling then requires

$$n \geq Cd\kappa \log \frac{d}{\delta} \max\{1, \sqrt{\kappa/d}\} \left(2 + \frac{2L\eta^2}{d \log \kappa}\right)$$

steps to achieve TV error less than δ . Thus with a negligible amount of increased cost for finding \tilde{x} that is $1/\sqrt{L}$ accurate: $\|\tilde{x} - x^*\| \leq 1/\sqrt{L}$, we have the number of steps of MALA for achieving δ error in TV norm remains $\mathcal{O}(d\kappa \log \frac{d}{\delta} \max\{1, \sqrt{\kappa/d}\})$.

Come back to the problem of estimating the normalizing constant. We will estimate the normalizing constant based on the annealing algorithm. The Lemma 3.2 suggests the choice of σ_M^2 to be larger than $\frac{2\sqrt{d}}{\mu}$ so that it satisfies the assumption of the Lemma the last stage has the same $\mathcal{O}(1)$ relative variance as the previous steps, guaranteed by Lemma 3.3. This implies that the number of stages

$$M \leq C\sqrt{d} \left(\log \frac{\kappa d}{\varepsilon} + 1 \right) = \tilde{\mathcal{O}}(\sqrt{d}). \quad (20)$$

Given the annealing sequence, we approximate Z_1 by the normalizing constant of Gaussian with variance σ_1^2 . Lemma 3.1 guarantees that this would only introduce at most $\varepsilon/2$ relative error. Thus the task remains to estimate the ratio Z_{i+1}/Z_i for $i = 1, \dots, M$, or equivalently to estimate the expectation of

$$g_i = \exp\left(\frac{1}{2}\left(\frac{1}{\sigma_i^2} - \frac{1}{\sigma_{i+1}^2}\right)\|x\|^2\right) \quad (21)$$

under the distribution ρ_i , proportional to $\exp(-\frac{1}{2}\frac{\|x\|^2}{\sigma_{i+1}^2} - f(x)) dx$. Suppose we generate K iid samples $X_i^{(1)}, \dots, X_i^{(K)}$ according to ρ_i , we estimate the ratio Z_{i+1}/Z_i by

$$\hat{g}_i = \frac{1}{K} \sum_{k=1}^K g_i(X_i^{(k)}). \quad (22)$$

Denote the short hand $\bar{g}_i = \mathbb{E}_{\rho_i} g_i$, we use the relative variance bounds shown in Lemma 3.3 and Lemma 3.2 to upper bound

$$\begin{aligned} \mathbb{E}(\hat{g}_i^2) &= \frac{1}{K^2} \left(\sum_{k=1}^K \mathbb{E}(g_i(X_i^{(k)}))^2 + K(K-1)\bar{g}_i^2 \right) \\ &\leq \frac{1}{K^2} (e^4 K + K(K-1))\bar{g}_i^2 \\ &\leq \left(1 + \frac{60}{K}\right)\bar{g}_i^2 \end{aligned} \quad (23)$$

Lemma B.2. Let Y_i , $i = 1, \dots, M$ be independent variables and let $\bar{Y}_i = \mathbb{E}Y_i$. Assume there exists $\eta > 0$ such that $\eta M \leq \frac{1}{5}$ and

$$\mathbb{E}Y_i^2 \leq (1 + \eta)\bar{Y}_i^2,$$

then for any $\varepsilon > 0$

$$\mathbb{P}\left(\frac{|Y_1 \cdots Y_M - \bar{Y}_1 \cdots \bar{Y}_M|}{\bar{Y}_1 \cdots \bar{Y}_M} \geq \frac{\varepsilon}{2}\right) \leq \frac{5\eta M}{\varepsilon^2}.$$

Proof. The proof follows the Chebyshev's inequality:

$$\begin{aligned} \mathbb{P}\left(\frac{|Y_1 \cdots Y_M - \bar{Y}_1 \cdots \bar{Y}_M|}{\bar{Y}_1 \cdots \bar{Y}_M} \geq \frac{\varepsilon}{2}\right) &\leq \frac{4}{\varepsilon^2} \frac{\text{Var}(Y_1 \cdots Y_M)}{\bar{Y}_1^2 \cdots \bar{Y}_M^2} \\ &= \frac{4}{\varepsilon^2} \left(\frac{\mathbb{E}(Y_1^2 \cdots Y_M^2)}{\bar{Y}_1^2 \cdots \bar{Y}_M^2} - 1 \right) \\ &\leq \frac{4}{\varepsilon^2} ((1 + \eta)^M - 1) \\ &\leq \frac{4}{\varepsilon^2} (e^{\eta M} - 1) \\ &\leq \frac{5\eta M}{\varepsilon^2}, \end{aligned}$$

where the last inequality follows from $e^{\eta M} - 1 \leq \frac{5}{4}\eta M$ for $\eta M \leq \frac{1}{5}$. \square

Applying Lemma B.2 by taking $Y_i = \hat{g}_i$ and $\eta = \frac{60}{K}$, we obtain

$$\mathbb{P}\left(\frac{|\hat{g}_1 \cdots \hat{g}_M - \bar{g}_1 \cdots \bar{g}_M|}{\bar{g}_1 \cdots \bar{g}_M} \geq \frac{\varepsilon}{2}\right) \leq \frac{300M}{\varepsilon^2 K}. \quad (24)$$

This suggests us to take the number of samples $K = \frac{1200M}{\varepsilon^2}$, so that the right hand side of above is bounded by $\frac{1}{4}$. Since we have M stages in total, the total number of samples we need in the whole algorithm is

$$N_{\text{tot}} = MK = \mathcal{O}\left(\frac{M^2}{\varepsilon^2}\right) = \tilde{\mathcal{O}}\left(\frac{d}{\varepsilon^2}\right). \quad (25)$$

To generate the iid samples $X_i^{(1)}, \dots, X_i^{(K)}$, $i = 1, \dots, M$, we will use the $\frac{1}{2}$ -lazy version of MALA algorithm, and choose parameter $\delta = \frac{1}{4} \frac{1}{N_{\text{tot}}}$, so that for probability at least $\frac{3}{4}$, every sample in our algorithm is guaranteed to follow the desired distribution, since we have in total N_{tot} samples.

Note that we have a uniform bound over the condition number of ρ_i , $i = 1, \dots, M$ by $\kappa = L/\mu$ thanks to the strongly log-concave assumption on ρ . Thus, for each sample, the number of steps it takes is bounded by $\mathcal{O}(d\kappa \log(dN_{\text{tot}}) \max\{1, \sqrt{\kappa/d}\})$ by Theorem B.1.

We summarize the procedure of estimating the normalizing constant based on the MALA sampling below.

Algorithm 2 Annealing algorithm for normalizing constant based on MALA

Input: μ -strongly convex and L -smooth function f , error threshold ε

Output: An estimate \widehat{Z} for the normalizing constant $Z = \int e^{-f(x)} dx$ within relative error $\mathcal{O}(\varepsilon)$

```

1:  $\sigma_1^2 \leftarrow \frac{\varepsilon}{2dL}$ 
2:  $M \leftarrow \left\lceil \log\left(\frac{2d^{3/2}\kappa}{\varepsilon}\right) / \log\left(1 + \frac{1}{\sqrt{d}}\right) \right\rceil$ 
3:  $K \leftarrow \frac{1200M}{\varepsilon^2}$ 
4:  $\widehat{Z} \leftarrow (2\pi\sigma_1^2)^{d/2}$ 
5: for  $i = 1, 2, \dots, M$  do
6:   if  $i < M$  then
7:      $\sigma_{i+1}^2 \leftarrow \sigma_i^2 \left(1 + \frac{1}{\sqrt{d}}\right)$ 
8:   else
9:      $\sigma_{i+1}^2 \leftarrow \infty$ 
10:  end if
11:  Use  $\frac{1}{2}$ -lazy MALA to generate random variables  $X_i^{(1)}, \dots, X_i^{(K)}$  iid wrt  $\rho_i$  with TV error
    guarantee  $\delta = \frac{1}{4MK}$ .
12:   $\widehat{g}_i \leftarrow \frac{1}{K} \sum_{k=1}^K \exp\left(\frac{1}{2}\left(\frac{1}{\sigma_i^2} - \frac{1}{\sigma_{i+1}^2}\right)\|X_i^{(k)}\|^2\right)$ 
13:   $\widehat{Z} \leftarrow \widehat{Z} \widehat{g}_i$ 
14: end for
15: return  $\widehat{Z}$ 

```

Putting together all the above estimates, we arrive at the following guarantee for the Algorithm 2.

Theorem B.3. *Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be a μ -strongly convex and L -smooth function. With probability of success at least $\frac{3}{4}$, Algorithm 2 gives an estimate \widehat{Z} of the normalizing constant $Z = \int e^{-f(x)} dx$ with relative error ε with query complexity*

$$\mathcal{O}\left(MK \frac{d^2}{\varepsilon^2} \log(dMK) \kappa \max\{1, \sqrt{\kappa/d}\}\right) = \widetilde{\mathcal{O}}\left(\frac{d^2}{\varepsilon^2} \kappa \max\{1, \sqrt{\kappa/d}\}\right).$$

C Estimating the Normalizing Constant using Multilevel Langevin

In Section C.1 we introduce multilevel Monte Carlo, a generic way to obtain a faster rate for estimating an expected value. Multilevel Monte Carlo reduces the variance in the estimate by simulating a SDE with multiple step sizes in a coupled fashion. We give guarantees for multilevel Monte Carlo for a general setting, assuming properties of the SDE and the coupling. In Section C.2 and C.3 we apply the multilevel Monte Carlo to ULD and ULD with RMM, respectively. These two sections prove the two parts of Theorem 1.2. In Section C.4, we introduce a truncation procedure to solve the technical issues mentioned in Section 4.3, namely that the function we are estimating is not Lipschitz. Finally in Section C.5 we apply multilevel ULD and ULD-RMM to normalizing constant estimation.

C.1 The multilevel estimate

We consider multilevel Monte Carlo for the following setting: We wish to estimate $\mathbb{E}_{X \sim \rho} g(X)$, where ρ cannot be sampled from exactly, but can be (approximately) sampled from by simulating a SDE for some time T . Suppose we have a discretization algorithm \mathcal{A} that given time T and

step size η , simulates the SDE with step size η , making $\mathcal{O}(T/\eta)$ queries (i.e., a constant number of queries per iteration), and returns a sample $X^\eta = x_T^\eta \sim \rho^\eta$. Smaller η naturally gives more accurate samples, but it also requires more queries and takes longer time. Naively, we would just run \mathcal{A} at a step size η small enough so that $|\mathbb{E}_{X \sim \rho^\eta} g(X^\eta) - \mathbb{E}_{X \sim \rho} g(X)| \leq \frac{\varepsilon}{2}$, and take enough samples. If we need to take $\eta = \varepsilon^{-\gamma}$, then this gives a rate of $\mathcal{O}\left(\frac{1}{\varepsilon^{2+\gamma}}\right)$.

Multilevel Monte Carlo method takes advantage of coupling of \mathcal{A} at two step sizes to reduce the variance. Assume that we can run \mathcal{A} coupled between two step sizes, to generate $(X^\eta, X^{\eta/2})$ such that $\text{Var}(g(X^\eta) - g(X^{\eta/2})) \ll \text{Var}(g(X^\eta))$ decays sufficiently fast, multilevel Monte Carlo leads to a faster rate $\tilde{\mathcal{O}}\left(\frac{1}{\varepsilon^2}\right)$ for estimating $\mathbb{E}_{X \sim \rho} g(X)$. The dependence on other parameters will also be improved.

To achieve this, multilevel Monte Carlo uses the estimator

$$\widehat{R} := \frac{1}{N_0} \sum_{i=1}^{N_0} g(X_i^{\eta_0}) + \sum_{j=1}^k \frac{1}{N_j} \sum_{i=1}^{N_j} [g(X_i^{\eta_0/2^j+}) - g(X_i^{\eta_0/2^{j-1}-})] \quad (26)$$

where $X_i^{\eta_0}$ are samples at the highest level (step size), and $(X_i^{\eta_0/2^j+}, X_i^{\eta_0/2^j-})$ are coupled samples at level j . For larger j , the variance $\text{Var}(g(X_i^{\eta_0/2^j+}) - g(X_i^{\eta_0/2^{j-1}-}))$ is smaller, so fewer samples are needed, offsetting the increased query complexity. We note that $\mathbb{E}\widehat{R} = \mathbb{E}[g(X^{\eta_0/2^j})]$, so the bias is determined by the smallest step size. On the other hand, minimizing the variance requires optimizing the sample sizes N_j .

We work out non-asymptotic rates for multilevel Monte Carlo, given the guarantees on \mathcal{A} (the rate of decay of the variance and bias of individual estimates in the step size η). The result is similar to [Gil08, Theorem 3.1], which works out the asymptotic rates when the variance and bias follow a power law in η . However, we will need to work out the rates when the desired bias ε_b and variance ε_σ^2 are different, because for our application of estimating normalizing constant, we can tolerate a larger ε_σ than ε_b at each temperature.

Note also the complication that in our setting, the bias depends not just on the step size, but also the time T . We simulate a SDE where ρ is the stationary distribution, so running the algorithm for a finite time T introduces some bias ε , even as the step size $\eta \rightarrow 0$. Hence, we assume that the bias is bounded by $G(\eta) \vee \varepsilon$, whenever $T \geq T(\varepsilon)$, and need to set T large enough. In our setting, the Markov processes will converge exponentially, so this only introduces a $\log\left(\frac{1}{\varepsilon}\right)$ factor.

Lemma C.1. *Let \mathcal{A} be an algorithm that given a parameter T (e.g. time) and $\eta > 0$ (e.g., discretizations with step size η), returns X^η . Let ρ^η be the distribution of X^η . Suppose also that $X \sim \rho$ (the distribution we are trying to approximate) and there are couplings between any two of the random variables. Suppose the following hold for any $\eta \leq \eta_{\max}$:*

1. *If X^η and $X^{\eta'}$ are coupled, the variance satisfies $\text{Var}[g(X^\eta) - g(X^{\eta'})] \leq F(\eta)$ whenever $\eta' \leq \frac{\eta}{2}$, where F is a non-decreasing, non-negative function satisfying $\sum_{j=0}^{\infty} \left(\frac{F(\eta/2^j)}{\eta/2^j}\right)^a \leq C_F \left(\frac{F(\eta)}{\eta}\right)^a$ for some universal constant C_F and any $a \in \{\frac{1}{2}, 1\}$.*
2. *The bias satisfies $|\mathbb{E}g(X^\eta) - \mathbb{E}g(X)| \leq G(\eta) \vee \varepsilon$, for non-decreasing function G , whenever $T \geq T(\varepsilon)$.*
3. *The variance satisfies $\text{Var}[g(X)] \leq c$.*
4. *Algorithm \mathcal{A} takes $\frac{T}{\eta}$ queries (e.g., to $\nabla \log(\rho)$) to compute a sample X^η .*

Suppose $\eta_j = \frac{\eta_0}{2^j}$ and η_0, η_k, N_j , and T are chosen so that the following hold:

Algorithm 3 Multilevel Monte Carlo

Input: Initial point x_0 , time T , largest step size η_0 , number of levels k , number of samples N_0, \dots, N_k , function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, function $g : \mathbb{R}^d \rightarrow \mathbb{R}$.

Input: Sampling algorithm $\mathcal{A}(x_0, f, \eta, T)$ which can give coupled samples $(x^\eta, x^{\eta/2})$ (or individual samples x^η).

Output: Estimate of $\mathbb{E}_{x \sim \rho} g(x)$ where $\rho(dx) \propto e^{-f(x)} dx$

```

1: for  $1 \leq i \leq N_0$  do
2:   Run  $\mathcal{A}$  with initial point  $x_0$ , function  $f$ , step size  $\eta_0$ , and time  $T$  to obtain  $X_i^{\eta_0}$ .
3: end for
4: for  $1 \leq j \leq k$  do
5:   for  $1 \leq i \leq N_j$  do
6:     Run coupled  $\mathcal{A}$  with initial point  $x_0$ , function  $f$ , step size  $\eta = \eta_0/2^{j-1}$ , and time  $T$ , to
       obtain  $(X_i^{\eta^-}, X_i^{\eta/2^+})$ .
7:   end for
8: end for
9: return  $\frac{1}{N_0} \sum_{i=1}^{N_0} g(X_i^{\eta_0}) + \sum_{j=1}^k \frac{1}{N_j} \sum_{i=1}^{N_j} [g(X_i^{\eta_0/2^j-}) - g(X_i^{\eta_0/2^{j-1}+})]$ .

```

- $F(\eta_0) = c$ and $\eta_0 \leq \eta_{\max}$.
- $G(\eta_k) \leq \varepsilon_b$.
- $N_j \geq \frac{1}{\varepsilon_\sigma^2} \sqrt{\frac{F(\eta_0)\eta_j F(\eta_j)}{\eta_0}}$.
- $T \geq T(\varepsilon_b)$.

Then the estimate (26) satisfies $|\mathbb{E}\widehat{R} - \mathbb{E}_\rho g| \leq \varepsilon_b$ and $\text{Var}(\widehat{R}) \leq \varepsilon_\sigma^2$. Taking N_j to be the minimum possible, the number of queries needed is

$$Q = T \left(\frac{4C_F^2 c}{\varepsilon_\sigma^2 \eta_0} + \frac{2}{\eta_k} \right) = \mathcal{O} \left(T \left(\frac{c}{\varepsilon_\sigma^2 \eta_0} + \frac{1}{\eta_k} \right) \right).$$

Note for example that the decay condition on F is satisfied when $F(\eta) = C\eta^\beta$ for some $\beta > 1$. This is the most favorable case in [Gil08, Theorem 3.1]; reduced speedups are still available in the regime $\beta \leq 1$.

Proof. Let $T = T(\varepsilon_b)$. The number of queries needed is $\sum_{j=1}^k \frac{T N_j}{\eta_j}$.

We claim that the total variance is $\text{Var}(\widehat{R}) \leq 4 \sum_{j=1}^k \frac{F(\eta_j)}{N_j}$, and the bias is $|\mathbb{E}\widehat{R} - \mathbb{E}_\rho g| \leq G(\eta_k) \vee \varepsilon_b$.

To see the expression for the variance, write

$$\widehat{R} = \frac{1}{N_0} \sum_{i=1}^{N_0} g(X_i) + \frac{1}{N_0} \sum_{i=1}^{N_0} [g(X_i^{\eta_0}) - g(X_i)] + \sum_{j=1}^k \frac{1}{N_j} \sum_{i=1}^{N_j} [g(X_i^{\eta_0/2^j+}) - g(X_i^{\eta_0/2^{j-1}-})]$$

so that the total variance is (the first two terms are not independent, but the others are)

$$\text{Var}(\widehat{R}) \leq \frac{2c}{N_0} + \frac{2F(\eta_0)}{N_0} + \sum_{j=1}^k \frac{F(\eta_j)}{N_j} \leq 4 \sum_{j=0}^k \frac{F(\eta_j)}{N_j}$$

since $\frac{c}{N_0} \leq \frac{F(\eta_0)}{N_0}$ by assumption on η_0 .

For the bias, note that $|\mathbb{E}\hat{R} - \mathbb{E}_\rho g| = |\mathbb{E}g(X^{\eta_k}) - \mathbb{E}_\rho g| \leq G(\eta_k) \vee \varepsilon_b$ by assumption.

To justify our choice of N_j , note that by Cauchy-Schwarz,

$$\underbrace{\left(\sum_{j=0}^k \frac{TN_j}{\eta_j} \right)}_{\text{(number of time steps)}} \underbrace{\left(\sum_{j=0}^k \frac{F(\eta_j)}{N_j} \right)}_{\text{(upper bound on variance)}} \geq T \left(\sum_{j=0}^k \sqrt{\frac{F(\eta_j)}{\eta_j}} \right)^2.$$

If the bound on variance is kept constant, because the RHS does not depend on N_j , then the number of steps is minimized when equality happens above. Equality happens when $N_j = K\sqrt{\eta_j F(\eta_j)}$ for some constant K . When $N_j \geq K\sqrt{\eta_j F(\eta_j)}$ the variance is bounded by

$$\text{Var}(\hat{R}) = 4 \sum_{j=0}^n \frac{F(\eta_j)}{N_j} = 4 \sum_{j=0}^n \frac{F(\eta_j)}{K\sqrt{\eta_j F(\eta_j)}} \leq 4 \sum_{j=0}^n \frac{1}{K} \sqrt{\frac{F(\eta_j)}{\eta_j}} = \frac{4C_F}{K} \sqrt{\frac{F(\eta_0)}{\eta_0}}$$

by assumption on the decay of F . By choosing $K = \frac{4C_F}{\varepsilon_\sigma^2} \sqrt{\frac{F(\eta_0)}{\eta_0}}$, the variance is bounded by ε_σ^2 . Then the requirement on N_j is $N_j \geq \frac{4C_F}{\varepsilon_\sigma^2} \sqrt{\frac{F(\eta_0)\eta_j F(\eta_j)}{\eta_0}}$.

It remains to compute the number of time steps. With the minimum choice of N_j , the number of time steps is

$$\begin{aligned} \sum_{j=0}^k \frac{TN_j}{\eta_j} &\leq \sum_{j=0}^k \frac{T}{\eta_j} \left(\frac{4C_F}{\varepsilon_\sigma^2} \sqrt{\frac{F(\eta_0)\eta_j F(\eta_j)}{\eta_0}} + 1 \right) \\ &= \sum_{j=0}^k \frac{4TC_F}{\varepsilon_\sigma^2} \sqrt{\frac{F(\eta_0)F(\eta_j)}{\eta_0\eta_j}} + \sum_{j=0}^k \frac{T}{\eta_j} \\ &\stackrel{(i)}{\leq} \frac{4TC_F^2}{\varepsilon_\sigma^2} \frac{F(\eta_0)}{\eta_0} + \frac{2T}{\eta_k} \\ &\leq T \left(\frac{4C_F^2 c}{\varepsilon_\sigma^2 \eta_0} + \frac{2}{\eta_k} \right). \end{aligned}$$

where (i) uses the assumption on decay of F and the fact that $\frac{T}{\eta_k}$ is a decaying geometric series with largest term $\frac{T}{\eta_k}$. \square

We put the lemma in a more convenient form for our applications.

Lemma C.2. *Suppose $g : \mathbb{R}^d \rightarrow \mathbb{R}$ is L_g -Lipschitz. Let \mathcal{A} be an algorithm that given a parameter T and $\eta > 0$, returns X^η . Let ρ^η be the distribution of X^η . Suppose also that $X^0 \sim \rho^0$ (e.g., the continuous process with the same initial distribution) and $X \sim \rho$ (the distribution we are trying to approximate) and there are couplings between any two of the random variables. Suppose the following hold for any $\eta \leq \eta_{\max}$:*

1. *If X^η and X^0 are coupled, then $\mathbb{E}[\|X^\eta - X^0\|^2] \leq F(\eta)$, where F is a non-decreasing, non-negative function satisfying $\sum_{j=0}^\infty \left(\frac{F(\eta/2^j)}{\eta/2^j}\right)^a \leq C_F \left(\frac{F(\eta)}{\eta}\right)^a$ for some universal constant C_F and any $a \in \{\frac{1}{2}, 1\}$.*
2. *If $T \geq T(\varepsilon)$, then $W_2(\rho^\eta, \rho)^2 \leq F(\eta) \vee \varepsilon^2$.*

3. ρ satisfies a Poincaré inequality with constant c . (In particular, this is satisfied for $c = \frac{1}{\mu}$ if $\rho(dx) \propto e^{-f(x)} dx$ and f is μ -strongly convex.)
4. Algorithm \mathcal{A} takes $\frac{T}{\eta}$ queries (e.g., to $\nabla \log(\rho)$) to compute a sample X^η .

Suppose $\eta_j = \frac{\eta_0}{2^j}$ and η_0, η_k, N_j , and T are chosen so that the following hold:

- $F(\eta_0) = \frac{c}{4}$, $\eta_0 \leq \eta_{\max}$.

- $F(\eta_k) \leq \frac{\varepsilon_b^2}{4L_g^2}$.

- $N_j \geq \frac{4L_g^2}{\varepsilon_\sigma^2} \sqrt{\frac{F(\eta_0)\eta_j F(\eta_j)}{\eta_0}}$.

- $T \geq T\left(\frac{\varepsilon_b}{L_g}\right)$

Then the estimate (26) satisfies $|\mathbb{E}\hat{R} - \mathbb{E}_\rho g| \leq \varepsilon_b$ and $\text{Var}(\hat{R}) \leq \varepsilon_\sigma^2$. Taking N_j to be the minimum possible, the number of queries needed is

$$Q = \mathcal{O}\left(T\left(\frac{cL_g^2}{\varepsilon_\sigma^2\eta_0} + \frac{1}{\eta_k}\right)\right).$$

Moreover, we have $W_2(\rho^{\eta_k}, \rho) \leq \frac{\varepsilon_b}{L_g}$.

Proof. We check that the conditions of Lemma C.1 are satisfied with $F(\eta) \leftarrow 4L_g^2 F(\eta)$, $G(\eta) \leftarrow 2L_g \sqrt{F(\eta)}$, $c \leftarrow cL_g^2$, and $T(\varepsilon) \leftarrow T\left(\frac{\varepsilon}{L_g}\right)$. Substituting then gives the parameters.

1. Using the fact that g is L_g Lipschitz, Cauchy-Schwarz, and the Minkowski inequality,

$$\begin{aligned} \text{Var}[g(X^\eta) - g(X^{\eta'})] &\leq L_g^2 \mathbb{E} \left[\left\| X^\eta - X^{\eta'} \right\|^2 \right] \\ &\leq L_g^2 \left(\mathbb{E} \left[\left\| X^\eta - X^0 \right\|^2 \right]^{\frac{1}{2}} + \mathbb{E} \left[\left\| X^0 - X^{\eta'} \right\|^2 \right]^{\frac{1}{2}} \right)^2 \leq 4L_g^2 F(\eta). \end{aligned}$$

2. Using the fact that g is L_g Lipschitz, for $T \geq T\left(\frac{\varepsilon}{L_g}\right)$,

$$|\mathbb{E}g(X^\eta) - \mathbb{E}g(X)| \leq L_g W_2(X^\eta, X) \leq L_g \left(\sqrt{f(\eta)} \vee \frac{\varepsilon}{L_g} \right) = L_g \sqrt{F(\eta)} \vee \varepsilon.$$

3. Since $\|\nabla g(x)\| \leq L_g$, the Poincaré inequality implies that $\text{Var}_\rho(g) \leq \frac{1}{\mu} \int_{\mathbb{R}^d} \|\nabla g(x)\|^2 dx \leq cL_g^2$. When $\rho(dx) \propto e^{-f(x)}$ and f is μ -strongly convex, it satisfies a Poincaré inequality by Bakry-Émery, Theorem F.3.

4. This follows directly.

Finally, note that by choice of η_k , $W_2(\rho^{\eta_k}, \rho) \leq F(\eta_k) \wedge \frac{\varepsilon_b}{L_g} = \frac{\varepsilon_b}{L_g}$. \square

C.2 Multilevel ULD

Underdamped Langevin diffusion with parameters γ, u is given by the following SDE:

$$\begin{aligned} dv_t &= -\gamma v_t dt - u \nabla f(x_t) dt + \sqrt{2\gamma u} dB_t \\ dx_t &= v_t dt \end{aligned}$$

where $x_t, v_t \in \mathbb{R}^d$ and B_t is standard Brownian motion. Under mild conditions, the SDE is ergodic with stationary distribution proportional to $e^{-(f(x) + \frac{1}{2u}\|v\|^2)}$. Compared to overdamped Langevin dynamics on log-concave distributions, it is known to enjoy an improved rate of convergence in W_2 distance. Here, v_t is thought of as velocity, and $-\gamma v_t$ is a drag term. ULD is closely related to Hamiltonian Monte Carlo.

The discrete dynamics with step size η can be described by

$$\begin{aligned} dv_t^\eta &= -\gamma v_t^\eta dt - u \nabla f(x_{\lfloor t/\eta \rfloor \eta}^\eta) dt + \sqrt{2\gamma u} dB_t \\ dx_t^\eta &= v_t^\eta dt. \end{aligned}$$

We will take $\gamma = 2$ and $u = \frac{1}{L}$. By integration, we can derive the explicit discrete-time update rule [CCBJ17, Lemma 10]:

$$v_{t+\eta}^\eta = e^{-2\eta} v_t^\eta + \frac{1}{L} \int_0^\eta e^{-2(\eta-s)} \nabla f(x_t^\eta) ds + \frac{2}{\sqrt{L}} \int_0^\eta e^{2(s-\eta)} dB_{t+s} \quad (27)$$

$$= e^{-2\eta} v_t^\eta + \frac{1}{2L} (1 - e^{-2\eta}) \nabla f(x_t^\eta) + \underbrace{\frac{2}{\sqrt{L}} \int_0^\eta e^{2(s-\eta)} dB_{t+s}}_{=: W_{1,t}^\eta} \quad (28)$$

$$x_{t+\eta}^\eta = x_t^\eta + \frac{1}{2} (1 - e^{-2\eta}) v_t^\eta + \frac{1}{2L} \int_0^\eta (1 - e^{-2(\eta-s)}) \nabla f(x_t^\eta) ds + \frac{1}{\sqrt{L}} \int_0^\eta (1 - e^{2(s-\eta)}) dB_{t+s} \quad (29)$$

$$= x_t^\eta + \frac{1}{2} (1 - e^{-2\eta}) v_t^\eta + \frac{1}{2L} \left(\eta - \frac{1}{2} (1 - e^{-2\eta}) \right) \nabla f(x_t^\eta) ds + \underbrace{\frac{1}{\sqrt{L}} \int_0^\eta (1 - e^{2(s-\eta)}) dB_{t+s}}_{=: W_{2,t}^\eta} \quad (30)$$

where all the instances of Brownian motion are the same. Let $G_t^\eta = \int_0^\eta e^{2s} dB_{t+s}$ and $H_t^\eta = \int_0^\eta dB_{t+s}$. As calculated in [SL19, Lemma 5],

$$\begin{aligned} \begin{pmatrix} G_t^\eta \\ H_t^\eta \end{pmatrix} &\sim N \left(\mathbf{0}, \begin{pmatrix} \frac{1}{4}(e^{4\eta} - 1) & \frac{1}{2}(e^{2\eta} - 1) \\ \frac{1}{2}(e^{2\eta} - 1) & \eta \end{pmatrix} \otimes I_d \right) \\ W_{1,t}^\eta &= e^{-2\eta} G_t^\eta \\ W_{2,t}^\eta &= H_t^\eta - e^{-2\eta} G_t^\eta. \end{aligned}$$

Define $S_{G_t^\eta, H_t^\eta}^\eta$ to be the map sending (x_t^η, v_t^η) to $(x_{t+\eta}^\eta, v_{t+\eta}^\eta)$ as defined above. As shorthand, because the η can be inferred, we write this as $S_{(G,H)_t}^\eta$.

We define a coupling between the continuous and discrete dynamics, or between discrete dynamics with different step sizes, by having the processes share the same Brownian motion. We refer

to this as synchronous coupling. When coupling the dynamics with step sizes η and $\eta/2$, we have

$$\begin{aligned} G_t^\eta &= \int_0^\eta e^{2s} dB_{t+s} = \int_0^{\eta/2} e^{2s} dB_{t+s} + \int_{\eta/2}^\eta e^{2s} dB_{t+s} \\ &= \int_0^{\eta/2} e^{2s} dB_{t+s} + e^\eta \int_0^{\eta/2} e^{2s} dB_{t+\frac{\eta}{2}+s} = G_t^{\eta/2} + e^\eta G_{t+\eta/2}^{\eta/2} \\ H_t^\eta &= \int_0^\eta dB_{t+s} = \int_0^{\eta/2} dB_{t+s} + \int_0^{\eta/2} dB_{t+\eta/2+s} = H_t^{\eta/2} + H_{t+\eta/2}^{\eta/2}. \end{aligned}$$

This leads to the update in Algorithm 4.

Algorithm 4 Coupled Underdamped Langevin Dynamics (ULD)

Input: Initial point $x_0 \in \mathbb{R}^d$, function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ (with gradient access).

Input: Time T and step size η OR bound on strong convexity μ , condition number κ , and desired accuracy ε .

Output: Coupled samples $(X^{\eta-}, X^{\eta/2+})$.

- 1: **if** ε is given **then**
- 2: Let $\eta = \frac{\varepsilon}{208\kappa} \sqrt{\frac{\mu}{d}}$.
- 3: Let $T = \frac{\kappa}{2} \log\left(\frac{48(d/\mu)}{\varepsilon}\right)$.
- 4: **end if**
- 5: Let $t = 0$.
- 6: Let $(x_0^\eta, v_0^\eta) = (x_0^{\frac{\eta}{2}}, v_0^{\frac{\eta}{2}}) = (x_0, 0)$.
- 7: **while** $t < T$ **do**
- 8: Draw $\begin{pmatrix} G_t^{\eta/2} \\ H_t^{\eta/2} \end{pmatrix}, \begin{pmatrix} G_{t+\eta/2}^{\eta/2} \\ H_{t+\eta/2}^{\eta/2} \end{pmatrix} \sim N\left(\mathbf{0}, \begin{pmatrix} \frac{1}{4}(e^{2\eta} - 1) & \frac{1}{2}(e^\eta - 1) \\ \frac{1}{2}(e^\eta - 1) & \eta/2 \end{pmatrix} \otimes I_d\right)$.
- 9: Let

$$\begin{aligned} G_t^\eta &= G_t^{\eta/2} + e^\eta G_{t+\eta/2}^{\eta/2} \\ H_t^\eta &= H_t^{\eta/2} + H_{t+\eta/2}^{\eta/2}. \end{aligned}$$

- 10: Let $(x_{t+\eta}^{\eta/2}, v_{t+\eta}^{\eta/2}) = S_{(G,H)_{t+\eta/2}}^{\eta/2} \circ S_{(G,H)_t}^{\eta/2}(x_t^{\eta/2}, v_t^{\eta/2})$.
- 11: Let $(x_{t+\eta}^\eta, v_{t+\eta}^\eta) = S_{(G,H)_t}^\eta(x_t^\eta, v_t^\eta)$.
- 12: Set $t \leftarrow t + \eta$.
- 13: **end while**
- 14: Output $(x_t^\eta, x_t^{\eta/2})$.

The main result on underdamped Langevin we will use is the following.

Theorem C.3 (Convergence of ULD, [CCBJ17, Theorem 1]). *Suppose f is twice continuously differentiable, μ -strongly convex, and L -smooth, and let $\kappa = \frac{L}{\mu}$. Let $\rho(dx) \propto e^{-f(x)} dx$.*

Let ρ^η be the distribution of discretized underdamped Langevin with step size η after time T , under the initial distribution $\delta_{(x,v)=(x_0,0)}$. Let the initial distance to optimum $x^ = \operatorname{argmin} f$ satisfy $\|x_0 - x^*\| \leq D$.*

1. *Let x^η, x^0 be synchronously coupled trajectories from the discrete and continuous processes.*

Let $T \geq \frac{\kappa}{2} \log \left(\frac{24\sqrt{\frac{d}{\mu} + D^2}}{\varepsilon} \right)$, and ρ^η be the distribution of x_T^η . Then for $\eta = O\left(\frac{1}{\kappa}\right)$ dividing into T ,

$$\begin{aligned}\mathbb{E}[||x_T^\eta - x_T^0||^2] &\leq \mathcal{O}\left(\frac{d}{\mu} \cdot \kappa^2 \eta^2\right) \\ W_2(\rho^\eta, \rho)^2 &\leq \mathcal{O}\left(\frac{d}{\mu} \cdot \kappa^2 \eta^2\right) \vee \varepsilon^2\end{aligned}$$

2. For step size $\eta \leq \frac{\varepsilon}{104\kappa} \sqrt{\frac{1}{d/\mu + D^2}}$ and $T \geq \frac{\kappa}{2} \log \left(\frac{24\sqrt{\frac{d}{\mu} + D^2}}{\varepsilon} \right)$, we have $W_2(\rho^\eta, \rho) \leq \varepsilon$. The algorithm makes $\frac{T}{\eta}$ queries to ∇f .

Proof. The second part is [CCBJ17, Theorem 1].² Their proof essentially establishes the first part of the theorem: In their notation, $W_2(\rho^\eta, \rho)$ is $W_2(p^{(n)}, p^*)$ where $n = \frac{T}{\eta}$. They show that

$W_2(p^{(n)}, p^*) \leq T_1 + T_2$, where $T_1 \leq \frac{\varepsilon}{2}$ with the choice of T , and $T_2 \leq 4\kappa\eta\sqrt{\frac{32 \cdot 26 \left(\frac{d}{\mu} + D^2\right)}{5}}$. This establishes the bound on $W_2(\rho^\eta, \rho)$.

For the bound on $\mathbb{E}[||x_T^\eta - x_T^0||^2]$, note that their bounds on Wasserstein distance come from synchronously coupling the continuous and discrete processes. In their notation, $q^{(n)}$ is the distribution of (x^η, v^η) at the n th step, $p^{(n)}$ is the distribution of x at the n th step, $\tilde{\Phi}_\eta$ is one step of the discrete process, and Φ_η is the exact underdamped Langevin process for the same amount of time. The same induction in (9)–(10) of [CCBJ17] shows that $W_2(\tilde{\Phi}_\eta^{T/\eta} q^{(0)}, \Phi_\eta^{T/\eta} q^{(0)}) \leq \frac{1}{1-e^{-\eta/2\kappa}} \eta^2 \sqrt{\frac{8\mathcal{E}_K}{5}}$ and $W_2(p^{(n)}, p'^{(n)}) \leq T_2 := \frac{1}{1-e^{-\eta/2\kappa}} \eta^2 \sqrt{\frac{32\mathcal{E}_K}{5}}$, where $p'^{(n)}$ is the distribution of the continuous process after n steps. The bound on Wasserstein distance is attained by synchronous coupling of the two processes. Their bound $T_2 \leq 4\kappa\eta\sqrt{\frac{32 \cdot 26 \left(\frac{d}{\mu} + D^2\right)}{5}}$ then establishes the bound on $\mathbb{E}[||x_T^\eta - x_T^0||^2]$. \square

The number of steps $\frac{T}{\eta}$ has a $\frac{\kappa^2}{\varepsilon}$ dependence on κ and ε . We note that [DRD18] has a better dependence, $\kappa^{\frac{3}{2}} \left(\kappa^{\frac{1}{2}} \wedge \frac{d^{\frac{1}{2}}}{\mu^{\frac{1}{2}} \varepsilon^2} \right)$, and can be used to give better bounds in Theorem C.12. However, as ULD-RMM has faster running time (Theorem C.5), we will work with the simpler bound in [CCBJ17]. The next Theorem gives the first part of Theorem 1.2.

Theorem C.4 (Rate of Multilevel ULD). *Let $\rho(dx) \propto e^{-f(x)} dx$, where $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is μ -strongly convex and L -smooth. Let $g : \mathbb{R}^d \rightarrow \mathbb{R}$ be L_g -Lipschitz. Suppose that x_0 satisfies $\|x_0 - x^*\| \leq D = O\left(\sqrt{\frac{d}{\mu}}\right)$. Then Algorithm 3 run using Algorithm 4 (ULD) started at x_0 with parameters*

$$\begin{aligned}\eta_0 &= \Theta\left(\frac{1}{d^{\frac{1}{2}} \kappa}\right) & T &\geq \frac{\kappa}{2} \log \left(\frac{24L_g\sqrt{\frac{d}{\mu} + D^2}}{\varepsilon_b} \right) \\ \eta_k &= \Theta\left(\frac{\varepsilon_b}{L_g \kappa} \sqrt{\frac{\mu}{d}}\right) & N_j &\geq \Omega\left(\frac{L_g^2 d^{\frac{3}{4}} \kappa^{\frac{3}{2}} \eta_j^{\frac{3}{2}}}{\mu \varepsilon_\sigma^2}\right)\end{aligned}$$

²Note that they actually show the theorem with the $\sqrt{\frac{d}{\mu} + D^2}$ inside the log, although this is not reflected in their theorem statement.

outputs \widehat{R} such that $|\mathbb{E}\widehat{R} - \mathbb{E}_\rho g| \leq \varepsilon_b$, and $\text{Var}(\widehat{R}) \leq \varepsilon_\sigma^2$. This takes $\mathcal{O}\left(\kappa^2\sqrt{d}\log\left(\frac{L_g}{\varepsilon_b}\cdot\sqrt{\frac{d}{\mu}}\right)\left(\frac{L_g^2}{\mu\varepsilon_\sigma^2} + \frac{L_g}{\sqrt{\mu\varepsilon_b}}\right)\right)$ gradient evaluations. Moreover, letting ρ^η be the distribution of x_T^η , we have $W_2(\rho, \rho^{\eta_k}) \leq \frac{\varepsilon_b}{L_g}$.

In particular, for $\varepsilon \leq \frac{L_g}{\sqrt{\mu}}$, taking $\varepsilon_b = \varepsilon_\sigma = \frac{\varepsilon}{3}$, $\mathbb{P}\left(|\widehat{R} - \mathbb{E}_\rho g| > \varepsilon\right) \leq \frac{1}{4}$, and the algorithm uses $\mathcal{O}\left(\frac{L_g^2\kappa^2\sqrt{d}}{\mu\varepsilon^2}\log\left(\frac{L_g}{\varepsilon_b}\cdot\sqrt{\frac{d}{\mu}}\right)\right)$ gradient evaluations.

Proof. We check that the conditions of Lemma C.2 hold with $F(\eta) = \frac{Cd\kappa^2\eta^2}{\mu}$ (for some C), $c = \frac{1}{\mu}$, and $T(\varepsilon) = \frac{\kappa}{2}\log\left(\frac{24\sqrt{\frac{d}{\mu}+D^2}}{\varepsilon^2}\right)$. Conditions 1, 2, and 4 follow from Theorem C.3(1), and condition 3 follows since f is μ -strongly convex.

We choose η_0 so that $\frac{Cd\kappa^2\eta_0^2}{\mu} = F(\eta_0) = \frac{1}{4\mu}$, leading to $\eta_0 = \Theta\left(\frac{1}{d^{\frac{1}{2}}\kappa}\right)$. Note that we do have $\eta_0 \leq \eta_{\max} = \Theta\left(\frac{1}{\kappa}\right)$. We choose η_k so that $\frac{Cd\kappa^2\eta_k^2}{\mu} \leq \frac{\varepsilon_b^2}{L_g^2}$, leading to $\eta_k = \mathcal{O}\left(\frac{\varepsilon_b}{\kappa L_g}\sqrt{\frac{\mu}{d}}\right)$. We choose N_j so that

$$N_j \geq \frac{4L_g^2}{\varepsilon_\sigma^2}\sqrt{\frac{f(\eta_0)\eta_j f(\eta_j)}{\eta_0}} = \mathcal{O}\left(\frac{L_g^2}{\varepsilon_\sigma^2}\sqrt{\frac{\left(\frac{1}{\mu}\right)\left(\frac{d\kappa^2\eta_j^3}{\mu}\right)}{\left(\frac{1}{d^{\frac{1}{2}}\kappa}\right)}}\right) = \mathcal{O}\left(\frac{L_g^2 d^{\frac{3}{4}} \kappa^{\frac{3}{2}} \eta_j^{\frac{3}{2}}}{\mu\varepsilon_\sigma^2}\right).$$

We choose $T \geq T\left(\frac{\varepsilon_b}{L_g}\right)$. Finally, the number of queries is

$$\begin{aligned} Q &= \mathcal{O}\left(T\left(\frac{cL_g^2}{\varepsilon_\sigma^2\eta_0} + \frac{1}{\eta_k}\right)\right) = \mathcal{O}\left(\kappa\log\left(\frac{L_g\sqrt{d/\mu}}{\varepsilon_b}\right)\left(\frac{L_g^2/\mu}{\varepsilon_\sigma^2\left(\frac{1}{d^{\frac{1}{2}}\kappa}\right)} + \frac{\kappa L_g}{\varepsilon_b}\sqrt{\frac{d}{\mu}}\right)\right) \\ &= \mathcal{O}\left(\kappa^2\sqrt{d}\log\left(\frac{L_g}{\varepsilon_b}\cdot\sqrt{\frac{d}{\mu}}\right)\left(\frac{L_g^2}{\mu\varepsilon_\sigma^2} + \frac{L_g}{\sqrt{\mu\varepsilon_b}}\right)\right). \end{aligned}$$

The last part follows since

$$\mathbb{P}\left(|\widehat{R} - \mathbb{E}_\rho g| > \varepsilon\right) \leq \mathbb{P}\left(|\widehat{R} - \mathbb{E}\widehat{R}| > \frac{2\varepsilon}{3}\right) \leq \frac{1}{4}. \quad \square$$

C.3 Multilevel ULD-RMM

In the integral formulation of the dynamics (27) and (29), the difference between the continuous and discrete dynamics is that in the continuous dynamics, we have the current gradient $\nabla f(x_{t+s})$ instead of the gradient at the last time step $\nabla f(x_t^\eta)$. The idea of the randomized midpoint method (RMM) [SL19] is to estimate the integrals by their value at $s = \alpha\eta$ for a uniformly random $\alpha \in [0, 1]$, instead of at $s = 0$. This reduces the bias caused by the one-step numerical quadrature with the price of increasing the standard deviation, which accumulates much slower than the bias in the numerical integration. This is in fact similar to our choice of $\varepsilon_b \ll \varepsilon_\sigma$ later for using the multilevel Monte Carlo method combined with annealing. The estimate of $x_{t+\alpha\eta}$, which we denote by y_t^η , is

obtained using the discretization with step size $\alpha\eta$. The update is given by

$$\begin{aligned} y_t^\eta &= x_t^\eta + \frac{1}{2}(1 - e^{-2\alpha\eta})v_t^\eta - \frac{1}{2}u \left(\alpha\eta - \frac{1}{2}(1 - e^{-2\alpha\eta}) \right) \nabla f(x_t^\eta) + \frac{1}{\sqrt{L}}W_{1,t}^\eta \\ x_{t+\eta}^\eta &= x_t^\eta + \frac{1}{2}(1 - e^{-2\eta})v_t^\eta - \frac{\eta}{2L}(1 - e^{-2(1-\alpha)\eta})\nabla f(y_t^\eta) + \frac{1}{\sqrt{L}}W_{2,t}^\eta \\ v_{t+\eta} &= v_t^\eta e^{-2\eta} - u\eta e^{-2(1-\alpha)\eta}\nabla f(y_t^\eta) + \frac{2}{\sqrt{L}}W_{3,t}^\eta \end{aligned}$$

where

$$W_{1,t}^\eta = H_1 - e^{-2\alpha\eta}G_1 \quad W_{2,t}^\eta = (H_1 + H_2) - e^{-2\eta}(G_1 + G_2) \quad W_{3,t}^\eta = e^{-2\eta}(G_1 + G_2)$$

and

$$\begin{aligned} G_{1,t}^\eta &= \int_0^{\alpha\eta} e^{2s} dB_{t+s} & H_{1,t}^\eta &= \int_0^{\alpha\eta} dB_{t+s} \\ G_{2,t}^\eta &= \int_{\alpha\eta}^\eta e^{2s} dB_{t+s} & H_{2,t}^\eta &= \int_{\alpha\eta}^\eta dB_{t+s}. \end{aligned}$$

Writing $G_i = G_{i,t}^\eta$ and $H_i = H_{i,t}^\eta$, define $R_{\alpha,G_1,H_1,G_2,H_2}^\eta$ to be the map sending (x_t^η, v_t^η) to $(x_{t+\eta}^\eta, v_{t+\eta}^\eta)$ as defined above.

To define the coupled dynamics, note that once we have selected α_1 and α_2 for step size $\eta/2$ for time steps t and $t + \eta/2$ respectively, one way to define a uniformly random $\alpha \in [0, 1]$ is to take $\frac{\alpha_1}{2}$ or $\frac{1+\alpha_2}{2}$ each with probability $\frac{1}{2}$. This coupling has the advantage that we have $t + \alpha\eta = t + \alpha_1\frac{\eta}{2}$ or $t + \frac{\eta}{2} + \alpha_2\frac{\eta}{2}$, so we can calculate the $W_{i,t}^\eta$ in terms of quantities already computed. (This coupling is out of convenience only; it is the fact that we use the same Brownian motion that reduces the variance, not the fact that α is coupled to α_1 and α_2 .) A straightforward calculation gives the updates for coupled ULD-RMM, Algorithm 5. For ease of notation we drop the subscripts and superscripts for G and H .

This gives the following improved rates.

Theorem C.5 (Convergence of ULD-RMM, [SL19, Theorem 3]). *Suppose f is twice continuously differentiable, μ -strongly convex, and L -smooth, and let $\kappa = \frac{L}{\mu}$. Let $\rho(dx) \propto e^{-f(x)} dx$.*

Let ρ^η be the distribution of the Randomized Midpoint Method for ULD with step size η after time T , under the initial distribution $\delta_{(x,v)=(x^,0)}$.*

1. *Let x^η, x^0 be synchronously coupled points from the discrete and continuous processes. Let $T \geq 2\kappa \log\left(\frac{20(d/\mu)}{\varepsilon^2}\right)$, and ρ^η be the distribution of x_T^η . For η smaller than some constant,*

$$\begin{aligned} \mathbb{E}[|x_T^\eta - x_T|^2] &\leq \mathcal{O}\left(\left(\frac{d\kappa\eta^6}{\mu} + \frac{d\eta^3}{\mu}\right) \log\left(\frac{\sqrt{d/\mu}}{\varepsilon}\right)\right) \\ W_2(\rho^\eta, \rho)^2 &\leq \mathcal{O}\left(\left(\frac{d\kappa\eta^6}{\mu} + \frac{d\eta^3}{\mu}\right) \log\left(\frac{\sqrt{d/\mu}}{\varepsilon}\right)\right) \vee \varepsilon^2 \end{aligned}$$

2. *Let $c > 0$ be a small enough constant. For step size $\eta \leq c \min\left\{\frac{\varepsilon^{\frac{1}{3}}}{\kappa^{\frac{1}{6}} \log^{\frac{1}{6}}\left(\frac{\sqrt{d/\mu}}{\varepsilon}\right)}, \left(\frac{\mu}{d}\right)^{\frac{1}{6}}, \frac{\varepsilon^{\frac{2}{3}}}{\log^{\frac{1}{3}}\left(\frac{\sqrt{d/\mu}}{\varepsilon}\right)} \left(\frac{\mu}{d}\right)^{\frac{1}{3}}\right\}$, and time $T \geq 2\kappa \log\left(\frac{20(d/\mu)}{\varepsilon^2}\right)$, $W_2(\rho^\eta, \rho) \leq \varepsilon$. The algorithm makes $\frac{2T}{\eta}$ queries to ∇f .*

Algorithm 5 Coupled Underdamped Langevin Dynamics with Randomized Midpoint Method (ULD-RMM)

Input: Initial point $x_0 \in \mathbb{R}^d$, function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ (with gradient access).

Input: Time T and step size η OR bound on strong convexity μ , condition number κ , and desired accuracy ε .

Output: Coupled samples $(X^{\eta-}, X^{\eta/2+})$.

```

1: if  $\varepsilon$  is given then
2:   Let  $\eta = c \max \left\{ \frac{\varepsilon^{\frac{1}{3}}}{\kappa^{\frac{1}{6}}} \left( \frac{\mu}{d} \right)^{\frac{1}{6}} \log^{-\frac{1}{6}} \left( \frac{1}{\varepsilon} \right), \varepsilon^{\frac{2}{3}} \left( \frac{\mu}{d} \right)^{\frac{1}{3}} \log^{-\frac{1}{3}} \left( \frac{1}{\varepsilon} \right) \right\}$ , where  $c$  is a small enough universal
   constant.
3:   Let  $T = 2\kappa \log \left( \frac{20}{\varepsilon^2} \right)$ .
4: end if
5: Let  $t = 0$ .
6: Let  $(x_0^\eta, v_0^\eta) = (x_0^{\frac{\eta}{2}}, v_0^{\frac{\eta}{2}}) = (x_0, 0)$ .
7: while  $t < T$  do
8:   Let  $\alpha_1, \alpha_2$  be random numbers in  $[0, 1]$ :  $\alpha_1, \alpha_2 \sim U([0, 1])$ .
9:   Draw  $\begin{pmatrix} G_1^{(i)} \\ H_1^{(i)} \end{pmatrix} \sim N \left( \mathbf{0}, \begin{pmatrix} \frac{1}{4}(e^{2\alpha\eta} - 1) & \frac{1}{2}(e^{\alpha\eta} - 1) \\ \frac{1}{2}(e^{\alpha\eta} - 1) & \alpha\eta/2 \end{pmatrix} \otimes I_d \right)$  for  $i = 1, 2$ .
10:  Draw  $\begin{pmatrix} G_2^{(i)} \\ H_2^{(i)} \end{pmatrix} \sim N \left( \mathbf{0}, \begin{pmatrix} \frac{1}{4}(e^{2\eta} - e^{2\alpha\eta}) & \frac{1}{2}(e^\eta - e^{\alpha\eta}) \\ \frac{1}{2}(e^\eta - e^{\alpha\eta}) & (1 - \alpha)\eta/2 \end{pmatrix} \otimes I_d \right)$  for  $i = 1, 2$ .
11:  if random coin flip = heads then
12:    Set  $\alpha = \frac{\alpha_1}{2}$ 

$$G_1 = G_1^{(1)} \quad H_1 = H_1^{(1)}$$


$$G_2 = G_2^{(1)} + e^\eta(G_1^{(2)} + G_2^{(2)}) \quad H_2 = H_2^{(1)} + H_1^{(2)} + H_2^{(2)}$$

13:  else
14:    Set  $\alpha = \frac{1+\alpha_2}{2}$  and

$$G_1 = G_1^{(1)} + G_2^{(1)} + e^\eta G_1^{(2)} \quad H_1 = H_1^{(1)} + H_2^{(1)} + H_1^{(2)}$$


$$G_2 = e^\eta G_2^{(2)} \quad H_2 = H_2^{(2)}.$$

15:  end if
16:  Let  $(x_{t+\eta}^{\eta/2}, v_{t+\eta}^{\eta/2}) = R_{\alpha_2, G_1^{(2)}, H_1^{(2)}, G_2^{(2)}, H_2^{(2)}}^{\eta/2} \circ R_{\alpha_1, G_1^{(1)}, H_1^{(1)}, G_2^{(1)}, H_2^{(1)}}^{\eta/2}(x_t^{\eta/2}, v_t^{\eta/2})$ .
17:  Let  $(x_{t+\eta}^\eta, v_{t+\eta}^\eta) = R_{\alpha, G_1, H_1, G_2, H_2}^\eta(x_t^\eta, v_t^\eta)$ .
18:  Set  $t \leftarrow t + \eta$ .
19: end while
20: Output  $(x_t^\eta, x_t^{\eta/2})$ .

```

Proof. The second part is exactly [SL19, Theorem 3].

In their notation, (x_n, v_n) is the n th iterate of their algorithm, and (y_n, w_n) is the n th step of the exact ULD, started from a random point from the stationary distribution. Examining their proof, they show that

$$\mathbb{E}[||x_N - y_N||^2 + ||(x_N + v_N) - (y_N + w_N)||^2]$$

$$\leq e^{-\frac{N\eta}{2\kappa}} \underbrace{\mathbb{E}[||x_0 - y_0||^2 + ||(x_0 + v_0) - (y_0 + w_0)||^2]}_{\leq \frac{5d}{\mu}} + \mathcal{O}\left(\left(\frac{\kappa d\eta^6}{\mu} + \frac{d\eta^3}{\mu}\right) \log\left(\frac{d/\mu}{\varepsilon^2}\right)\right) \\ + \mathcal{O}(\kappa\eta^7 + \eta^3)\mathbb{E}[||x_N - y_N||^2 + ||(x_N + v_N) - (y_N + w_N)||^2]$$

For $\eta = O(\kappa^{-\frac{1}{7}})$, we have that the last term is $\leq \frac{1}{2}\mathbb{E}[||x_N - y_N||^2 + ||(x_N + v_N) - (y_N + w_N)||^2]$, so

$$\mathbb{E}[||x_N - y_N||^2 + ||(x_N + v_N) - (y_N + w_N)||^2] \leq e^{-\frac{N\eta}{2\kappa}} \frac{5d}{\mu} + \mathcal{O}\left(\left(\frac{\kappa d\eta^6}{\mu} + \frac{d\eta^3}{\mu}\right) \log\left(\frac{d/\mu}{\varepsilon^2}\right)\right).$$

By choice of T (or N), this term is $\leq \frac{\varepsilon}{2}$. This establishes the bound on $W_2(\rho^\eta, \rho)$.

Finally, note that we can replace (y_n, w_n) by the exact ULD started with the same initial condition. Then the same derivation holds, except that the first term is 0. This shows the bound on $\mathbb{E}[||x_T^\eta - x_T||^2]$. \square

Combining Theorem C.5 with Lemma C.2, we can prove the second part of Theorem 1.2.

Theorem C.6 (Rate of Multilevel ULD-RMM). *Let $\rho(dx) \propto e^{-f(x)} dx$, where $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is μ -strongly convex and L -smooth. Let $g : \mathbb{R}^d \rightarrow \mathbb{R}$ be L_g -Lipschitz. Then Algorithm 3 run using Algorithm 5 (ULD) started at x^* with parameters*

$$\eta_0 = \Theta\left(\frac{1}{d^{\frac{1}{6}}\kappa^{\frac{1}{6}}\log\left(\frac{L_g^2}{\varepsilon_b^2} \cdot \frac{d}{\mu}\right)^{\frac{1}{6}}} \wedge \frac{1}{d^{\frac{1}{3}}\log\left(\frac{L_g^2}{\varepsilon_b^2} \cdot \frac{d}{\mu}\right)^{\frac{1}{3}}}\right) \quad T = \mathcal{O}\left(\kappa\log\left(\frac{L_g^2}{\varepsilon_b^2} \cdot \frac{d}{\mu}\right)\right) \\ \eta_k = \Theta\left(\frac{\varepsilon_b^{\frac{1}{3}}\mu^{\frac{1}{6}}}{d^{\frac{1}{6}}\kappa^{\frac{1}{6}}L_g^{\frac{1}{3}}\log\left(\frac{L_g^2}{\varepsilon_b^2} \cdot \frac{d}{\mu}\right)^{\frac{1}{6}}} \wedge \frac{\varepsilon_b^{\frac{2}{3}}\mu^{\frac{1}{3}}}{d^{\frac{1}{3}}L_g^{\frac{2}{3}}\log\left(\frac{L_g^2}{\varepsilon_b^2} \cdot \frac{d}{\mu}\right)^{\frac{1}{3}}}\right) \quad N_j \geq \Omega\left(\frac{L_g^2}{\varepsilon_\sigma^2}\sqrt{\frac{f(\eta_0)\eta_j f(\eta_j)}{\eta_0}}\right)$$

outputs \hat{R} such that $|\mathbb{E}\hat{R} - \mathbb{E}_\rho g| \leq \varepsilon_b$, and $\text{Var}(\hat{R}) \leq \varepsilon_\sigma^2$. This takes

$$\mathcal{O}\left(\left(\kappa^{\frac{7}{6}}d^{\frac{1}{6}}\log\left(\frac{L_g^2}{\varepsilon_b^2} \cdot \frac{d}{\mu}\right)^{\frac{7}{6}} + \kappa d^{\frac{1}{3}}\log\left(\frac{L_g^2}{\varepsilon_b^2} \cdot \frac{d}{\mu}\right)^{\frac{4}{3}}\right)\frac{L_g^2}{\mu\varepsilon_\sigma^2} + \frac{\kappa^{\frac{7}{6}}d^{\frac{1}{6}}L_g^{\frac{1}{3}}}{\varepsilon_b^{\frac{1}{3}}}\log\left(\frac{L_g^2}{\varepsilon_b^2} \cdot \frac{d}{\mu}\right)^{\frac{7}{6}} + \frac{\kappa d^{\frac{1}{3}}L_g^{\frac{2}{3}}}{\varepsilon_b^{\frac{2}{3}}}\log\left(\frac{L_g^2}{\varepsilon_b^2} \cdot \frac{d}{\mu}\right)^{\frac{4}{3}}\right)$$

gradient evaluations. Moreover, letting ρ^η be the distribution of x_T^η , we have $W_2(\rho, \rho^{\eta_k}) \leq \frac{\varepsilon_b}{L_g}$.

In particular, for $\varepsilon \leq \frac{L_g}{\sqrt{\mu}}$, taking $\varepsilon_b = \varepsilon_\sigma = \frac{\varepsilon}{2}$, $\mathbb{P}(|\hat{R} - \mathbb{E}_\rho g| > \varepsilon) \leq \frac{1}{4}$, and the algorithm uses $\mathcal{O}\left(\frac{L_g^2\kappa^{\frac{7}{6}}d^{\frac{1}{6}}\log\left(\frac{L_g^2}{\varepsilon_b^2} \cdot \frac{d}{\mu}\right)^{\frac{7}{6}} + \kappa d^{\frac{1}{3}}\log\left(\frac{L_g^2}{\varepsilon_b^2} \cdot \frac{d}{\mu}\right)^{\frac{4}{3}}}{\mu\varepsilon^2}\right)$ gradient evaluations.

Proof. We check that the conditions of Lemma C.2 hold with $F(\eta) = C\left(\frac{d\kappa\eta^6}{\mu} + \frac{d\eta^3}{\mu}\right) \log\left(\frac{L_g^2}{\varepsilon_b^2} \cdot \frac{d}{\mu}\right)$ (for some C), $c = \frac{1}{\mu}$, and $T(\varepsilon) = 2\kappa\log\left(\frac{20(d/\mu)}{\varepsilon^2}\right)$. Conditions 1, 2, and 4 follow from Theorem C.5(1), and condition 3 follows since f is μ -strongly convex.

We choose η_0 so that

$$C \left(\frac{d\kappa\eta_0^6}{\mu} + \frac{d\eta_0^3}{\mu} \right) \log \left(\frac{L_g^2}{\varepsilon_b^2} \cdot \frac{d}{\mu} \right) = F(\eta_0) = \frac{1}{4\mu}$$

and η_k so that

$$\begin{aligned} \frac{\varepsilon_b^2}{4L_g^2} &\geq C \left(\frac{d\kappa\eta_k^6}{\mu} + \frac{d\eta_k^3}{\mu} \right) \log \left(\frac{L_g^2}{\varepsilon_b^2} \cdot \frac{d}{\mu} \right) = F(\eta_k) \\ &\Leftarrow \frac{d\kappa\eta_k^6}{\mu} \log \left(\frac{L_g^2}{\varepsilon_b^2} \cdot \frac{d}{\mu} \right) \leq \mathcal{O} \left(\frac{\varepsilon_b^2}{L_g^2} \right) \\ \text{and } \frac{d\eta_k^3}{\mu} \log \left(\frac{L_g^2}{\varepsilon_b^2} \cdot \frac{d}{\mu} \right) &\leq \mathcal{O} \left(\frac{\varepsilon_b^2}{L_g^2} \right), \end{aligned}$$

leading to the given bounds on η_0 and η_k . We choose T so that $T \geq T \left(\frac{\varepsilon_b}{L_g} \right)$. We do have that $\eta_0 \leq \eta_{\max} = \Theta(1)$. Substituting the bounds on η_0 and η_k into $Q = \mathcal{O} \left(T \left(\frac{L_g^2/\mu}{\varepsilon_b^2 \eta_0} + \frac{1}{\eta_k} \right) \right)$ gives the bound on the number of queries. \square

C.4 Truncation error and bias

There is a technical point that the ratio g is not Lipschitz, as it grows exponentially for large $\|x\|$; however, because large x 's are very unlikely under ρ , the expected value of g changes very little if we replace it by a “clamped” version of g (Lemma C.7).

Lemma C.7 (Truncation error). *Suppose that $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is a μ -strongly convex function, ρ is a probability measure on \mathbb{R}^d with $\rho(dx) = \frac{1}{Z} e^{-\left(\frac{1}{2}\frac{\|x\|^2}{\sigma^2} + f(x)\right)}$ dx , and $g(x) = \exp \left(\frac{\|x\|^2}{2\sigma^2(1+\alpha^{-1})} \right)$ for $\alpha \in (0, \infty]$. Let ρ' be the probability distribution with $\frac{d\rho'}{d\rho} \propto g(x)$, and $\bar{r} = \mathbb{E}_{x \sim \rho'} \|x\|$. For any $r \geq \bar{r}$,*

$$\frac{\int_{\|x\| \geq r} g(x) \rho(dx)}{\mathbb{E}_{x \sim \rho} g(x)} \leq \exp \left(-\frac{1}{2} \left(\frac{1}{\sigma^2(1+\alpha)} + \mu \right) (r - \bar{r})^2 \right).$$

Note that we allow $\alpha = \infty$, in which case $\alpha^{-1} = 0$, $g(x) = \exp \left(\frac{\|x\|^2}{2\sigma^2} \right)$, and the bound is $\exp \left(-\frac{\mu(r - \bar{r})^2}{2} \right)$.

Proof. Note that $\frac{1}{\sigma^2} - \frac{1}{\sigma^2(1+\alpha^{-1})} = \frac{1}{\sigma^2(\alpha+1)}$, so ρ' is $\left(\frac{1}{\sigma^2(\alpha+1)} + \mu \right)$ -strongly convex. Then for any r ,

$$\frac{\int_{\|x\| \geq r} g(x) \rho(dx)}{\int_{\mathbb{R}^d} g(x) \rho(dx)} = \mathbb{P}_{x \sim \rho'} (\|x\| \geq r).$$

By Theorem F.3 and F.4 on the 1-Lipschitz function $\|x\|$, we have the concentration bound

$$\mathbb{P}_{x \sim \rho'} (\|x\| \geq r) \leq \exp \left(-\frac{1}{2} \left(\frac{1}{\sigma^2(1+\alpha)} + \mu \right) (r - \bar{r})^2 \right).$$

\square

Lemma C.8 (Bias calculation). *Suppose that $f : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a convex function, ρ is a probability measure on \mathbb{R}^d with $\rho(dx) = \frac{1}{Z} e^{-\left(\frac{1}{2} \frac{\|x\|^2}{\sigma^2} + f(x)\right)} dx$, and $g(x) = g_r(\|x\|)$, where $g_r(r) := e^{\frac{r^2}{2\sigma^2(1+\alpha^{-1})}}$. Let \bar{r} be as in Lemma C.7. Suppose one of the following hold.*

1. Suppose $r^+ = \bar{r} + c\sigma$ for $c \geq 0$. Let $L_h = \frac{\bar{r} + c\sigma}{\sigma^2(1+\alpha^{-1})} \exp\left(4\alpha^2 d + \frac{\bar{r}c}{\sigma(1+\alpha^{-1})} + \frac{c^2}{2(1+\alpha^{-1})}\right)$ and $\varepsilon_1 = \exp\left(-\frac{c^2}{2(1+\alpha)}\right)$.
2. Suppose f is μ -strongly convex, $\sigma^2 \geq \frac{2}{\mu}$, $\alpha = \infty$, and $r^+ = \bar{r} + \frac{c}{\sqrt{\mu}}$ for $c \geq 0$. Let $L_h = \frac{\bar{r} + \frac{c}{\sqrt{\mu}}}{\sigma^2} \exp\left(\frac{4d}{\mu\sigma^4} + \frac{\bar{r}c}{\sigma^2\sqrt{\mu}} + \frac{c^2}{2\sigma^2\mu}\right)$ and $\varepsilon_1 = \exp\left(-\frac{c^2}{2}\right)$.

Define $h(y) = g(y) \wedge g_r(r^+)$. Then h is $L_h \cdot \mathbb{E}_\rho g$ -Lipschitz and

$$\frac{|\mathbb{E}_{x \sim \tilde{\rho}} h(x) - \mathbb{E}_{x \sim \rho} g(x)|}{\mathbb{E}_{x \sim \rho} g(x)} \leq \varepsilon_1 + L_h W_2(\rho, \tilde{\rho}).$$

Proof. By the triangle inequality,

$$\frac{|\mathbb{E}_{x \sim \tilde{\rho}} h(x) - \mathbb{E}_{x \sim \rho} g(x)|}{\mathbb{E}_{x \sim \rho} g(x)} \leq \frac{|\mathbb{E}_{x \sim \rho} [h(x) - g(x)]|}{\mathbb{E}_{x \sim \rho} g(x)} + \frac{|\mathbb{E}_{x \sim \tilde{\rho}} h(x) - \mathbb{E}_{x \sim \rho} h(x)|}{\mathbb{E}_{x \sim \rho} g(x)}.$$

In either case, the first expression is bounded by Lemma C.7:

$$\frac{|\mathbb{E}_{x \sim \rho} [h(x) - g(x)]|}{\mathbb{E}_{x \sim \rho} g(x)} \leq \varepsilon_1$$

To bound the second expression, we note that h is Lipschitz with constant $\max_{\|x\| \leq r^+} \|\nabla g(x)\|$. Thus by Kantorovich-Rubinstein duality,

$$\frac{|\mathbb{E}_{x \sim \tilde{\rho}} h(x) - \mathbb{E}_{x \sim \rho} h(x)|}{\mathbb{E}_{x \sim \rho} g(x)} \leq \frac{W_1(\rho, \tilde{\rho}) \max_{\|x\| \leq r^+} \|\nabla g(x)\|}{\mathbb{E}_{x \sim \rho} g(x)} \leq W_2(\rho, \tilde{\rho}) \frac{\max_{\|x\| \leq r^+} \|\nabla g(x)\|}{\mathbb{E}_{x \sim \rho} g(x)}$$

It remains to show the bound $\frac{\max_{\|x\| \leq r^+} \|\nabla g(x)\|}{\mathbb{E}_{x \sim \rho} g(x)} \leq L_h$. We consider the two cases separately.

Case 1. First, we compare the numerator to $g_r(\bar{r})$. Let ρ' be the probability density with $\frac{d\rho'}{d\rho} = g(x)$. Then $\mathbb{E}_{x \sim \rho'} g(x) = \frac{\mathbb{E}_\rho(g^2)}{\mathbb{E}_\rho g}$, so

$$\mathbb{E}_\rho g = \mathbb{E}_\rho g \frac{E_\rho g}{\mathbb{E}_\rho(g^2)} \mathbb{E}_{\rho'} g = \frac{(\mathbb{E}_\rho g)^2}{\mathbb{E}_\rho(g^2)} \mathbb{E}_{\rho'} g \geq \exp(-4\alpha^2 d) \mathbb{E}_{\rho'} g$$

using (13) and Lemma 3.3. Now, by definition of \bar{r} in Lemma C.7 because g_r is convex, $\mathbb{E}_{x \sim \rho'} g(x) = \mathbb{E}_{x \sim \rho'} g_r(\|x\|) \geq g_r(\mathbb{E}_{x \sim \rho'} \|x\|) = g_r(\bar{r})$. Hence

$$\begin{aligned} \frac{\max_{\|x\| \leq r^+} \|\nabla g(x)\|}{\mathbb{E}_{x \sim \rho} g(x)} &= \frac{\frac{d}{dx} g_r(x)|_{x=r^+}}{\mathbb{E}_{x \sim \rho} g(x)} \leq \frac{\frac{r^+}{\sigma^2(1+\alpha^{-1})} g_r(r^+)}{\exp(-4\alpha^2 d) g_r(\bar{r})} \\ &= \frac{r^+}{\sigma^2(1+\alpha^{-1})} \exp\left(4\alpha^2 d + \frac{r^{+2} - \bar{r}^2}{2\sigma^2(1+\alpha^{-1})}\right) \\ &= \frac{\bar{r} + c\sigma}{\sigma^2(1+\alpha^{-1})} \exp\left(4\alpha^2 d + \frac{\bar{r}c}{\sigma(1+\alpha^{-1})} + \frac{c^2}{2(1+\alpha^{-1})}\right) = L_h. \end{aligned}$$

Case 2. Similar to the first case,

$$\mathbb{E}_\rho g = \frac{(\mathbb{E}_\rho g)^2}{\mathbb{E}_\rho g^2} \mathbb{E}_{\rho'} g \geq \exp\left(-\frac{4d}{\mu\sigma_M^2}\right) \mathbb{E}_{\rho'} g \geq \exp\left(-\frac{4d}{\mu\sigma_M^2}\right) g_r(\bar{r})$$

using Lemma 3.2, noting that the condition on σ^2 is satisfied. We now have

$$\begin{aligned} \frac{\max_{\|x\| \leq r^+} \|\nabla g(x)\|}{\mathbb{E}_{x \sim \rho} g(x)} &= \frac{\frac{d}{dx} g_r(x)|_{x=r^+}}{\mathbb{E}_{x \sim \rho} g(x)} \leq \frac{\frac{r^+}{\sigma^2} g_r(r^+)}{\exp\left(-\frac{4d}{\mu\sigma^4}\right) g_r(\bar{r})} \\ &= \frac{r^+}{\sigma^2} \exp\left(\frac{4d}{\mu\sigma^4} + \frac{r^{+2} - \bar{r}^2}{2\sigma^2}\right) \\ &= \frac{\bar{r} + c/\sqrt{\mu}}{\sigma^2} \exp\left(\frac{4d}{\mu\sigma^4} + \frac{\bar{r}c}{\sigma^2\sqrt{\mu}} + \frac{c^2}{2\sigma^2\mu}\right) = L_h. \end{aligned} \quad \square$$

Corollary C.9. *Keep the setup of Lemma C.8. Then*

1. *If $\alpha \leq \frac{\log 2}{2\sqrt{d} \log(\frac{1}{\varepsilon})}$ and $r^+ \in \bar{r} + \sigma \sqrt{(1+\alpha) \log(\frac{1}{\varepsilon})} [\sqrt{2}, 2]$, then in Lemma C.8(1), $L_h \leq \frac{112e}{\sigma}$.*
2. *If $\sigma^2 \geq 4 \left(\frac{\sqrt{d} \vee \sqrt{\log(\frac{1}{\varepsilon})}}{\mu} \right) \left(1 \vee \frac{1}{\sqrt{\mu}} \right)$, $r^+ \in \bar{r} + \frac{1}{\sqrt{\mu}} \sqrt{\log(\frac{1}{\varepsilon})} [\sqrt{2}, 2]$, and $\varepsilon \leq \frac{1}{4}$, then in Lemma C.8(2), $L_h \leq 2e^2 \sqrt{\mu}$.*

In either case, $\frac{|\mathbb{E}_{\tilde{\rho}} h - \mathbb{E}_\rho g|}{|\mathbb{E}_\rho g|} \leq \varepsilon_1 + L_h W_2(\rho, \tilde{\rho})$.

Proof. To show (1), write $r^+ = \bar{r} + c\sigma$. Then $c \geq \sqrt{2(1+\alpha) \log(\frac{1}{\varepsilon_1})}$, so we have $e^{-\frac{c^2}{2(1+\alpha)}} \leq \varepsilon_1$. By Lemma F.2,

$$\bar{r} \leq \sigma \sqrt{1+\alpha} (\sqrt{d} + 2\sqrt{2 \log 2}) \leq 5\sigma\sqrt{d}.$$

Noting that $c \leq 2\sqrt{(1+\alpha) \log(\frac{1}{\varepsilon})}$,

$$\begin{aligned} \frac{\bar{r} + c\sigma}{\sigma^2(1+\alpha^{-1})} &\leq \frac{5\sqrt{d}}{\sigma(1+\alpha^{-1})} + \frac{2\alpha\sqrt{\log(\frac{1}{\varepsilon})}}{\sigma} \\ \frac{\bar{r}c}{\sigma(1+\alpha^{-1})} &\leq 10\alpha\sqrt{d \log\left(\frac{1}{\varepsilon}\right)} \\ \frac{c^2}{2(1+\alpha^{-1})} &\leq 2\alpha \log\left(\frac{1}{\varepsilon}\right). \end{aligned}$$

Substituting into the definition of L_h in Lemma C.8(1),

$$\begin{aligned}
L_h &= \frac{\bar{r} + c\sigma}{\sigma^2(1 + \alpha^{-1})} \exp \left(4\alpha^2 d + \frac{\bar{r}c}{\sigma(1 + \alpha^{-1})} + \frac{c^2}{2(1 + \alpha^{-1})} \right) \\
&\stackrel{(i)}{\leq} \left(\frac{5\sqrt{d}}{\sigma(1 + \alpha^{-1})} + \frac{2\alpha\sqrt{\log\left(\frac{1}{\varepsilon_1}\right)}}{\sigma} \right) \exp \left(4\alpha^2 d + 10\alpha\sqrt{d\log\left(\frac{1}{\varepsilon}\right)} + 2\alpha\log\left(\frac{1}{\varepsilon}\right) \right) \\
&\stackrel{(ii)}{\leq} \frac{1}{\sigma} \left(\frac{5}{2} + 2\alpha\sqrt{\log\left(\frac{1}{\varepsilon}\right)} \right) \exp \left(1 + 10\alpha\sqrt{d\log\left(\frac{1}{\varepsilon}\right)} \right) \left(\frac{1}{\varepsilon} \right)^{2\alpha} \\
&\stackrel{(iii)}{\leq} \frac{7}{2\sigma} \exp(1 + 5\log 2) 2 = \frac{112e}{\sigma},
\end{aligned}$$

where (i) follows from substitution, (ii) follows from $\alpha \leq \frac{1}{2\sqrt{d}}$, and (iii) follows from $\alpha \leq \frac{\log 2}{2\sqrt{d}\log\left(\frac{1}{\varepsilon}\right)}$.

To show (2), write $r^+ = \bar{r} + \frac{c}{\sqrt{\mu}}$. Then $c \geq \sqrt{2\log\left(\frac{1}{\varepsilon}\right)}$, so we have $\exp\left(-\frac{\mu(r^+ - \bar{r})^2}{2}\right) = \exp\left(-\frac{c^2}{2}\right) \leq \varepsilon$. By Lemma F.2,

$$\bar{r} \leq \frac{1}{\sqrt{\mu}}(\sqrt{d} + 2\sqrt{2\log 2}) \leq 5\sqrt{\frac{d}{\mu}}$$

Noting that $c \leq 2\sqrt{\log\left(\frac{1}{\varepsilon}\right)}$,

$$\begin{aligned}
\frac{\bar{r} + \frac{c}{\sqrt{\mu}}}{\sigma^2} &\leq \frac{5\sqrt{d} + 2\sqrt{\log\left(\frac{1}{\varepsilon}\right)}}{\sigma^2\sqrt{\mu}} \stackrel{(i)}{\leq} \frac{7}{4} \\
\frac{4d}{\mu\sigma^4} &\stackrel{(ii)}{\leq} \frac{1}{4} \\
\frac{\bar{r}c}{\sigma^2\sqrt{\mu}} &\leq \frac{5\sqrt{d\log\left(\frac{1}{\varepsilon}\right)}}{\sigma^2\mu} \stackrel{(iii)}{\leq} \frac{5}{4} \\
\frac{c^2}{2\sigma^2\mu} &\stackrel{(iv)}{\leq} \frac{2\log\left(\frac{1}{\varepsilon}\right)}{\sigma^2\mu} \leq \frac{2}{4}.
\end{aligned}$$

where in (i) we use $\sigma^2 \geq 4\frac{\sqrt{d}}{\sqrt{\mu}}$ and $\sigma^2 \geq 4\frac{\sqrt{\log\left(\frac{1}{\varepsilon}\right)}}{\sqrt{\mu}}$, in (ii) we use $\sigma^2 \geq 4\frac{\sqrt{d}}{\sqrt{\mu}}$, in (iii) we use $\sigma^2 \geq 4 \cdot \frac{\sqrt{d}}{\sqrt{\mu}} \cdot \frac{\sqrt{\log\left(\frac{1}{\varepsilon}\right)}}{\sqrt{\mu}}$, and in (iv) we use $\sigma^2 \geq 4 \cdot \frac{\sqrt{\log\left(\frac{1}{\varepsilon}\right)}}{\sqrt{\mu}} \cdot \frac{\sqrt{\log\left(\frac{1}{\varepsilon}\right)}}{\sqrt{\mu}}$. Substituting into the definition of L_h in Lemma C.8(2),

$$\begin{aligned}
L_h &= \frac{\bar{r} + \frac{c}{\sqrt{\mu}}}{\sigma^2} \exp \left(\frac{4d}{\mu\sigma^4} + \frac{\bar{r}c}{\sigma^2\sqrt{\mu}} + \frac{c^2}{2\sigma^2\sqrt{\mu}} \right) \\
&\leq \frac{7}{4} \exp \left(\frac{1}{4} + \frac{5}{4} + \frac{2}{4} \right) \leq 2e^2.
\end{aligned}$$

In either case, by Lemma C.8, $\frac{|\mathbb{E}_{\bar{\rho}}h - \mathbb{E}_{\rho}g|}{|\mathbb{E}_{\rho}g|} \leq \varepsilon_1 + \varepsilon_2$. \square

C.5 Estimating the normalizing constant

Before stating the main algorithm, let us first understand how errors in individual annealing steps can be composed to give the final error for estimating the normalizing constant.

Lemma C.10. *Suppose the following hold.*

1. (Estimate of partition function at highest temperature) $\frac{\hat{Z}_1}{Z_1} \in [e^{-\varepsilon_1}, e^{\varepsilon_1}]$.
2. (Bias of ratio) For $1 \leq i \leq M$, letting $\tilde{R}_i = \mathbb{E}\hat{R}_i$, $|\tilde{R}_i - R_i| \leq \frac{\varepsilon_2 R_i}{2M}$.
3. (Variance of ratio) For $1 \leq i \leq M$, \hat{R}_i is independent with $\text{Var}(\hat{R}_i) \leq \frac{\varepsilon_3^2 \tilde{R}_i^2}{40M}$.

Then $\mathbb{P}\left(\frac{\hat{Z}_1 \prod_{i=1}^M \hat{R}_i}{Z_1 \prod_{i=1}^M R_i} \notin [e^{-(\varepsilon_1 + \varepsilon_2 + \varepsilon_3)}, e^{\varepsilon_1 + \varepsilon_2 + \varepsilon_3}]\right) \leq \frac{1}{8}$.

Proof. From (2) we get that $\frac{\tilde{R}_i}{R_i} \in [1 - \frac{\varepsilon_2}{2M}, 1 + \frac{\varepsilon_2}{2M}] \subseteq [e^{-\frac{\varepsilon_2}{M}}, e^{\frac{\varepsilon_2}{M}}]$.

From (3) and Lemma B.2, $\mathbb{P}\left(\prod_{i=1}^M \frac{\hat{R}_i}{R_i} \notin [e^{-\varepsilon_3}, e^{\varepsilon_3}]\right) \leq \frac{5\varepsilon_3^2 M}{40\varepsilon_3^2 M} = \frac{1}{8}$.

Factoring $\frac{\hat{Z}_1 \prod_{i=1}^M \hat{R}_i}{Z_1 \prod_{i=1}^M R_i} = \frac{\hat{Z}_1}{Z_1} \cdot \prod_{i=1}^M \frac{\hat{R}_i}{R_i} \cdot \prod_{i=1}^M \frac{\hat{R}_i}{R_i}$, the result now follows. \square

We are now ready to introduce the main algorithm for estimating the normalizing constant. Algorithm 6 first estimates the thresholds r_i^+ to cut off g_i in Lemma C.8 so that the resulting estimate has bias $\mathcal{O}(\varepsilon)$. Then it calls the Multilevel Monte Carlo algorithm at each temperature with the truncated functions h_i . We can choose which Monte Carlo algorithm to use; we will consider both the ULD and ULD-RMM algorithms. Note that an alternative to estimating $\bar{r}_i^+ = \mathbb{E}_{X \sim \rho_{i+1}} \|X\|$ separately is to use the samples obtained from the multilevel procedure; we only estimate it separately to make the proof simpler.

To prove the correctness of Algorithm 6, we rely on guarantees proved in Theorem C.4 and Theorem C.6, as well as the truncation in Section C.4. The final ingredient is to show that Algorithm 6 estimates r_i^+ and r_M^+ correctly for the truncation in Section C.4 to work.

Lemma C.11. *Suppose $\alpha \leq \frac{1}{4}$ and $\sigma_M^2 \geq \frac{1}{\mu}$. In Algorithm 6, with probability $\geq \frac{7}{8}$ both the following hold:*

1. For $1 \leq i < M$, $r_i^+ \in \bar{r}_i + \sigma_i \sqrt{(1 + \alpha) \log(\frac{1}{\varepsilon})} [\sqrt{2}, 2]$.

2. $r_M^+ \in \bar{r}_M + \frac{1}{\sqrt{\mu}} \sqrt{\log(\frac{1}{\varepsilon})} [\sqrt{2}, 2]$.

Proof. Let μ_i be the strong convexity constant of f_i . Let $\tilde{\rho}_{i+1}$ be the distribution of the output of $\mathcal{A}(x_0, f_{i+1}, \varepsilon = \frac{\sigma_i}{8})$. By guarantee of algorithm \mathcal{A} and the fact that $\|\cdot\|$ is 1-Lipschitz, $|\mathbb{E}_{x \sim \tilde{\rho}_{i+1}} \|x\| - \mathbb{E}_{x \sim \rho_{i+1}} \|x\|| \leq \frac{\sigma_i}{8}$.

Now

$$\begin{aligned}
\text{Var}_{x \sim \tilde{\rho}_{i+1}} (\|x\|) &\leq \mathbb{E}_{x \sim \tilde{\rho}_{i+1}} \left[(\|x\| - \mathbb{E}_{y \sim \rho_{i+1}} \|y\|)^2 \right] \\
&\stackrel{(i)}{\leq} \inf_{(x,y) \in \mathcal{C}(\tilde{\rho}_{i+1}, \rho_{i+1})} \left[\mathbb{E}_{(x,y)} [(\|x\| - \|y\|)^2]^{\frac{1}{2}} + \mathbb{E}_{y \sim \rho_{i+1}} [(\|y\| - \mathbb{E}_{y \sim \rho_{i+1}} \|y\|)^2]^{\frac{1}{2}} \right]^2 \\
&= \left(W_2(\tilde{\rho}_{i+1}, \rho_{i+1}) + \text{Var}_{y \sim \rho_{i+1}} (\|y\|)^{\frac{1}{2}} \right)^2 \\
&\stackrel{(ii)}{\leq} \left(\frac{\sigma_i}{8} + \frac{5}{4} \sigma_i \right)^2 \leq \left(\frac{11}{8} \sigma_i \right)^2 \leq 2\sigma_i^2
\end{aligned}$$

Algorithm 6 Multilevel Monte Carlo for normalizing constant estimation

Input: Initial point x_0 , function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, initial temperature σ_1 , final temperature σ_{\max} , multiplier α , desired accuracy ε .

Input: Algorithm $\mathcal{A}(x_0, f)$ which: (1) given (η, T) , returns coupled samples $(X^\eta, X^{\eta/2})$, (2) given ε , returns samples $X \sim \tilde{\rho}$ such that $W_2(\rho, \tilde{\rho}) \leq \varepsilon$, where $\rho(dx) \propto e^{-f(x)} dx$.

Input: Algorithm $\mathcal{B}(L, \mu, L_h, \varepsilon_b, \varepsilon_\sigma)$ to set parameters $(T, \eta_0, k, N_0, \dots, N_k)$ for the multilevel Monte Carlo.

Output: Estimate of $Z = \int_{\mathbb{R}^d} e^{-f(x)} dx$.

- 1: Let $\widehat{Z} = \widehat{Z}_1 = (2\pi\sigma_1^2)^d$.
- 2: Let the number of levels be $M = \lceil \log_{1+\alpha} \left(\frac{\sigma_{\max}^2}{\sigma_1^2} \right) \rceil + 1$.
- 3: For each $1 \leq i \leq M$ let Let $\sigma_i = \sigma_1(1 + \alpha)^{(i-1)/2}$.
- 4: **for** $1 \leq i \leq M - 1$ **do**
- 5: Let $f_{i+1}(x) = \frac{\|x\|^2}{2\sigma_{i+1}^2} + f(x)$.
- 6: Run algorithm $\mathcal{A}(x_0, f_{i+1}, \varepsilon = \frac{\sigma_i}{8})$ to obtain $S = 2^{10}M$ samples x_i , and let $\widehat{r}_i = \frac{1}{S} \sum_{j=1}^S \|x_j\|$.
- 7: Let $r_i^+ = \widehat{r}_i + \sigma_i \sqrt{2(1 + \alpha) \log \left(\frac{8}{\varepsilon} \right) + \frac{1}{4}}$.
- 8: **end for**
- 9: Run algorithm $\mathcal{A}(x_0, f, \varepsilon = \frac{1}{8\sqrt{\mu}})$ to obtain $S = 2^{10}M$ samples x_j , and let $\widehat{r}_M = \frac{1}{S} \sum_{j=1}^S \|x_j\|$.
- 10: Let $r_M^+ = \widehat{r}_M + \frac{1}{\sqrt{\mu}} \sqrt{2 \log \left(\frac{8}{\varepsilon} \right) + \frac{1}{4}}$.
- 11: **for** $1 \leq i \leq M$ **do**
- 12: If $i = M$ set $\alpha \leftarrow \infty$.
- 13: Let $g_i(x) = \exp \left(\frac{\|x\|^2}{2\sigma_i^2(1+\alpha^{-1})} \right)$.
- 14: Let $h_i(x) = g_i(x) \wedge \exp \left(\frac{r_i^{+2}}{2\sigma_i^2(1+\alpha^{-1})} \right)$.
- 15: Run Algorithm 3 (Multilevel Monte Carlo) on functions f_i, h_i with sampling algorithm \mathcal{A} and with parameters set by $\mathcal{B} \left(L + \frac{1}{\sigma_i^2}, \mu + \frac{1}{\sigma_i^2}, L_h, \frac{\varepsilon}{16M}, \frac{\varepsilon}{128\sqrt{M}} \right)$, to obtain estimate \widehat{R}_i of $\mathbb{E}_{\rho_i} h_i(x)$.
- 16: Set $\widehat{Z} \leftarrow \widehat{Z} \widehat{R}_i$.
- 17: **end for**
- 18: **return** \widehat{Z}

where in (i) we use Minkowski's inequality and in (ii) we use the fact that $\text{Var}_{x \sim \rho_{i+1}}(\|x\|) \leq \frac{1}{\mu_{i+1}}$ by Theorem F.3, and for $\alpha \leq \frac{1}{4}$, $\text{Var}_{x \sim \rho_{i+1}}(\|x\|) \leq \frac{1}{\mu_{i+1}} \leq \sigma_{i+1}^2 \leq \sigma_i^2(1 + \alpha)^2 \leq \left(\frac{5}{4}\right)^2 \sigma_i^2$. Then since $S = 2^{10}M$,

$$\text{Var}_{x_j \sim \tilde{\rho}_{i+1}} \left(\frac{1}{S} \sum_{j=1}^S \|x_j\| \right) \leq \frac{2\sigma_i^2}{2^{10}M} = \frac{\sigma_i^2}{2^9 M}.$$

Thus by the triangle inequality and the bound on the bias,

$$\mathbb{P} \left(\left| \widehat{r}_i - \mathbb{E}_{x \sim \rho_{i+1}} \|x\| \right| \geq \frac{\sigma_i}{4} \right) \leq \mathbb{P} \left(\left| \widehat{r}_i - \mathbb{E}\widehat{r}_i \right| \geq \frac{\sigma_i}{8} \right) \leq \frac{\sigma_i^2 / (2^9 M)}{\sigma_i^2 / 2^6} \leq \frac{1}{8M}.$$

The analogous statement for $i = M$ follows similarly with $f_{M+1} = f$ and $\rho_{M+1} = \rho$ by noting $\text{Var}_{x \sim \rho}(\|x\|) \leq \frac{1}{\mu} \leq \sigma_M^2$, using the assumption on σ_M^2 . By the union bound, letting $s_i =$

$\begin{cases} \sigma_i, & 1 \leq i < M \\ \frac{1}{\sqrt{\mu}}, & i = M, \\ 1 \leq i \leq M-1, \end{cases}$ we have $\mathbb{P}(\forall i \in [1, M], |\hat{r}_i - \mathbb{E}_{x \sim \rho_{i+1}} \|x\|| \leq \frac{s_i}{4}) \geq \frac{7}{8}$. Under this event, for

$$r_i^+ \in \bar{r}_i + \sigma_i \left(\sqrt{2(1+\alpha) \log \left(\frac{1}{\varepsilon_1} \right)} + \left[0, \frac{1}{2} \right] \right) \subseteq \bar{r}_i + \sigma_i \sqrt{(1+\alpha) \log \left(\frac{1}{\varepsilon} \right)} [\sqrt{2}, 2],$$

and for $i = M$,

$$r_M^+ \in \bar{r}_M + \frac{1}{\sqrt{\mu}} \left(\sqrt{2 \log \left(\frac{1}{\varepsilon_1} \right)} + \left[0, \frac{1}{2} \right] \right) \subseteq \bar{r}_M + \frac{1}{\sqrt{\mu}} \sqrt{\log \left(\frac{1}{\varepsilon} \right)} [\sqrt{2}, 2] \quad \square$$

Finally we are ready to state and prove the main theorems.

Theorem C.12 (Multilevel ULD for estimating the normalizing constant). *Let $f(x)$ be μ -strongly convex and L -smooth. Let $Z = \int_{\mathbb{R}^d} e^{-f(x)} dx$. Let $\alpha = \frac{\log 2}{2\sqrt{d} \log(\frac{8}{\varepsilon})} \wedge \frac{1}{4}$, $\sigma_1 = \frac{\varepsilon}{8dL}$, and $\sigma_{\max} = 4 \left(\frac{\sqrt{d} \vee \sqrt{\log(\frac{8}{\varepsilon})}}{\mu} \right) \left(1 \vee \frac{1}{\sqrt{\mu}} \right)$. Algorithm 6 with Algorithm 4 as the sampling algorithm \mathcal{A} , with parameters set by Theorem C.4 computes \hat{Z} such that with probability $\frac{3}{4}$, $\frac{\hat{Z}}{Z} \in [1 - \varepsilon, 1 + \varepsilon]$. The number of queries to $\nabla f(x)$ is $\tilde{\mathcal{O}}\left(\frac{d^{\frac{3}{2}} \kappa^2}{\varepsilon^2}\right)$.*

Proof. Let μ_i, L_i, κ_i be the strong convexity constant, smoothness constant, and condition number of f_i . Note that $\kappa_i = \frac{L_i}{\mu_i} \leq \frac{L_i + \frac{1}{\sigma_i^2}}{\mu_i + \frac{1}{\sigma_i^2}} \leq \kappa$, so we can always bound the dependence on κ_i by κ ; we will use this fact implicitly. Let $\tilde{\rho}_i$ be the distribution of $x_T^{\eta_k}$, where η_k, T are the smallest step size and time for the i th temperature. Let \hat{R}_i be the estimate at the i th temperature, $\hat{R}_i := \mathbb{E}\hat{R}_i = \mathbb{E}_{\tilde{\rho}_i} h_i$, and $R_i = \mathbb{E}_{\rho_i} g_i$. For ease of computation, let ε_1 and $\varepsilon_2 \leq \frac{1}{4}$ be such that $\sigma_1 = \frac{\varepsilon_1}{dL}$ (our assumption has $\varepsilon_1 = \frac{\varepsilon}{8}$) and $\sigma_{\max} = 4 \left(\frac{\sqrt{d} \vee \sqrt{\log(\frac{1}{\varepsilon_2})}}{\mu} \right) \left(1 \vee \frac{1}{\sqrt{\mu}} \right)$ and $\alpha = \frac{\log 2}{2\sqrt{d} \log(\frac{1}{\varepsilon_2})} \wedge \frac{1}{4}$ (our assumption has $\varepsilon_2 = \frac{\varepsilon}{8}$).

By assumption on α and σ_{\max} , by Lemma C.11, with probability $\geq \frac{7}{8}$, Corollary C.9(1) is satisfied for $(r^+, \bar{r}) = (r_i^+, \bar{r}_i)$ for $1 \leq i \leq M-1$ and (2) is satisfied for $(r^+, \bar{r}) = (r_M^+, \bar{r}_M)$. Then $L_{h_i} = \mathcal{O}\left(\frac{1}{\sigma_i}\right)$ and $L_{h_M} = \mathcal{O}(\sqrt{\mu})$. In either case, $L_{h_i} = \mathcal{O}(\sqrt{\mu_i})$ and h_i/R_i is L_{h_i} -Lipschitz. For the rest of the proof, we will condition on the event that the hypothesis of Corollary C.9 are satisfied.

By Corollary C.9, $\frac{|\tilde{R}_i - R_i|}{R_i} \leq \varepsilon_2 + L_h W_2(\rho, \tilde{\rho}_i)$. In order to make $\text{Var}(\hat{R}_i) \leq \frac{\varepsilon_2^2 R_i^2}{256M} =: \varepsilon_\sigma^2 R_i^2$ and $|\tilde{R}_i - R_i| \leq \frac{\varepsilon_2 R_i}{2M} =: \varepsilon_b R_i$, by Theorem C.4, the number of queries required is

$$Q = \mathcal{O} \left(\kappa_i^2 \sqrt{d} \log \left(\frac{L_{h_i}}{\varepsilon_b} \cdot \sqrt{\frac{d}{\mu_i}} \right) \left(\frac{L_{h_i}^2}{\mu_i \varepsilon_\sigma^2} + \frac{L_{h_i}}{\sqrt{\mu_i} \varepsilon_b} \right) \right) = \mathcal{O} \left(\kappa^2 \sqrt{d} \log \left(\frac{\sqrt{d} M}{\varepsilon_2} \right) \left(\frac{M}{\varepsilon_2^2} + \frac{M}{\varepsilon_2} \right) \right)$$

where we substitute ε_b and ε_σ and use $L_{h_i} = \mathcal{O}(\sqrt{\mu_i})$.

Also by Theorem C.4, $W_2(\rho_i, \tilde{\rho}_i) \leq \frac{\varepsilon_b}{L_{h_i}}$, so $\frac{|\tilde{R}_i - R_i|}{R_i} \leq \varepsilon_2 + L_{h_i} W_2(\rho_i, \tilde{\rho}_i) \leq \varepsilon_2 + \frac{\varepsilon_2}{2M} \leq \frac{1}{2}$, where in the last step we use $\varepsilon_2 \leq \frac{1}{4}$. Hence $\tilde{R}_i \geq \frac{1}{2} R_i$ and $\text{Var}(\hat{R}_i) \leq \frac{\varepsilon^2 R_i^2}{256M} \leq \frac{\varepsilon^2 \tilde{R}_i^2}{64M}$.

By choice of σ_1 , by Lemma 3.1, $1 \leq \frac{\hat{Z}_1}{Z_1} \leq \frac{1}{1-\varepsilon_1} \leq e^{2\varepsilon_1}$. We also have $|\tilde{R}_i - R_i| \leq \frac{\varepsilon_2 R_i}{2M}$ and $\text{Var}(\hat{R}_i) \leq \frac{\varepsilon_2^2 \tilde{R}_i^2}{64M}$. By Lemma C.10, $\mathbb{P}\left(\frac{\hat{Z}}{Z} \notin [e^{-(2\varepsilon_1+2\varepsilon_2)}, e^{(2\varepsilon_1+2\varepsilon_2)}]\right) \leq \frac{1}{8}$. Taking $\varepsilon_1 = \varepsilon_2 = \frac{\varepsilon}{8}$ as in our assumptions, and recalling that we conditioned on an event of probability $\geq \frac{7}{8}$, we have that $\frac{\hat{Z}}{Z} \in [1 - \varepsilon, 1 + \varepsilon]$ with probability $\geq \frac{3}{4}$.

The total number of levels is $M = \log_{1+\alpha} \left(\frac{\sigma_M^2}{\sigma_1^2} \right) = \tilde{\mathcal{O}}(\sqrt{d})$. The total query complexity is $QM = \tilde{\mathcal{O}}\left(\kappa^2 \sqrt{d} M^2\right) = \tilde{\mathcal{O}}\left(\frac{\kappa^2 d^{\frac{3}{2}}}{\varepsilon^2}\right)$. \square

Theorem C.13 (Multilevel ULD-RMM for estimating the normalizing constant). *Let $f(x)$ be μ -strongly convex and L -smooth. Let $Z = \int_{\mathbb{R}^d} e^{-f(x)} dx$. Define α , σ_i , and σ_{\max} as in Theorem C.12, and let $x_0 = x^* = 0$. Algorithm 6 with Algorithm 5 as the sampling algorithm \mathcal{A} , with parameters set by Theorem C.6 computes \hat{Z} such that with probability $\frac{3}{4}$, $\frac{\hat{Z}}{Z} \in [1 - \varepsilon, 1 + \varepsilon]$. The number of queries to $\nabla f(x)$ is $\tilde{\mathcal{O}}\left(\frac{d^{\frac{4}{3}} \kappa + d^{\frac{7}{6}} \kappa^{\frac{7}{6}}}{\varepsilon^2}\right)$.*

Note that we assume $x_0 = x^*$ as Theorem C.5 makes that assumption; however, we note that we can use gradient descent to approximately find x^* , and that the analysis of [SL19] can tolerate a warm start.

Proof. The proof is the same as Theorem C.12. The only difference is that the number of queries at a level is given by Theorem C.6 instead:

$$\begin{aligned} Q &= \mathcal{O}\left(\left(\kappa^{\frac{7}{6}} d^{\frac{1}{6}} \log\left(\frac{L_g^2}{\varepsilon_b^2} \cdot \frac{d}{\mu}\right)^{\frac{7}{6}} + \kappa d^{\frac{1}{3}} \log\left(\frac{L_g^2}{\varepsilon_b^2} \cdot \frac{d}{\mu}\right)^{\frac{4}{3}}\right) \frac{L_g^2}{\mu \varepsilon_\sigma^2} \right. \\ &\quad \left. + \frac{\kappa^{\frac{7}{6}} d^{\frac{1}{6}} L_g^{\frac{1}{3}}}{\varepsilon_b^{\frac{1}{3}}} \log\left(\frac{L_g^2}{\varepsilon_b^2} \cdot \frac{d}{\mu}\right)^{\frac{7}{6}} + \frac{\kappa d^{\frac{1}{3}} L_g^{\frac{2}{3}}}{\varepsilon_b^{\frac{2}{3}}} \log\left(\frac{L_g^2}{\varepsilon_b^2} \cdot \frac{d}{\mu}\right)^{\frac{4}{3}}\right) \\ &= \mathcal{O}\left(\kappa^{\frac{7}{6}} d^{\frac{1}{6}} \left(\frac{M}{\varepsilon^2} + \frac{M^{\frac{1}{3}}}{\varepsilon^{\frac{1}{3}}}\right) + \kappa d^{\frac{1}{3}} \left(\frac{M}{\varepsilon^2} + \frac{M^{\frac{2}{3}}}{\varepsilon^{\frac{2}{3}}}\right)\right) \end{aligned}$$

The total query complexity is $QM = \tilde{\mathcal{O}}((d^{\frac{1}{3}} \kappa + d^{\frac{1}{6}} \kappa^{\frac{7}{6}}) M^2) = \tilde{\mathcal{O}}\left(\frac{d^{\frac{4}{3}} \kappa + d^{\frac{7}{6}} \kappa^{\frac{7}{6}}}{\varepsilon^2}\right)$. \square

D Proof of Lowerbound

In this section we provide the missing details for the lowerbound in Section 5.

Result on biased coin

Claim D.1. *Given independent samples of a random variable X , where X is drawn from Bernoulli distribution with either $p = 1/2 + \delta$ or $p = 1/2 - \delta$, any algorithm that looks at $o(1/\delta^2)$ samples of X cannot decide which distribution X is drawn from without probability better than $1/2 + c$ for any constant $c > 0$.*

This is very standard and we give a proof here just for completeness.

Proof. Let Y and Z be two Bernoulli random variables with $p_Y = 1/2 + \delta$ and $p_Z = 1/2 - \delta$ of being 1 respectively. Then the KL-divergence between these two distributions is $KL(Y\|Z) \leq \mathcal{O}(\delta^2)$. Let Y^n and Z^n be n independent samples of Y and Z , by property of KL divergence we know $KL(Y^n\|Z^n) = nKL(Y\|Z) \leq \mathcal{O}(n\delta^2)$. When $n = o(1/\delta^2)$, $KL(Y^n\|Z^n) = o(1)$. Finally by Pinsker's inequality we know the TV-distance between Y^n and Z^n is at most $\sqrt{KL(Y^n\|Z^n)/2} = o(1)$. Therefore it is impossible to distinguish between Y^n and Z^n with any probability $1/2 + c$ for constant $c > 0$. \square

Construction of function q First, the lowerbound construction needs a function q which we use to modify the initial function f_0 . We construct such a q function in the following lemma:

Lemma D.1. *There exists a function $q : [-1, 1]^k \rightarrow \mathbb{R}$ that satisfies*

1. *For any $x \in [-1, 1]^k$ with at least one coordinate $x_i = \pm 1$, $q(x) = 0$, $\nabla q(x) = 0$ and $\nabla^2 q(x) = 0$.*
2. *For any $x \in [-1, 1]^k$, $0 \leq q(x) \leq 1$, $\|\nabla^2 q(x)\| \leq 36k$.*
3. *For any $x \in [-1/2, 1/2]^k$, $q(x) \geq 3^{-k}$.*

Proof. We construct q as a product of individual coordinates. Let $p : [-1, 1] \rightarrow \mathbb{R}$ be the function $p(x) = (1+x)^3(1-x)^3$. It is easy to verify that $p(-1) = p'(-1) = p''(-1) = p(1) = p'(1) = p''(1) = 0$, $0 \leq p(x) \leq 1$ and $p(x) \geq 1/3$ when $x \in [-1/2, 1/2]$.

Now we define $q(x) = p(x_1)p(x_2) \cdots p(x_k)$. If any coordinate $x_i (i = 1, 2, \dots, k)$ is 1 or -1, we have $q(x) = 0$ because $p(x_i) = 0$. The gradients $\frac{\partial q}{\partial x_i} = p'(x_i)p(x_1)p(x_2) \cdots p(x_{i-1})p(x_{i+1}) \cdots p(x_k) = 0$; for any $j \neq i$, $\frac{\partial q}{\partial x_j}$ has a factor of $p(x_i)$ so it is also 0. Similarly, all the second order partial derivatives will have a factor of $p(x_i)$, $p'(x_i)$ or $p''(x_i)$, so the Hessian is also 0. Therefore we have verified Property 1.

For Property 2, we observe that for $i \neq j$, $\frac{\partial^2 q}{\partial x_i \partial x_j}(x) = p'(x_i)p'(x_j) \prod_{t \neq i, j} p(x_t)$. It is easy to verify that $|p'(x_i)| \leq 6$ for any value of $x_i \in [-1, 1]$, therefore $|\frac{\partial^2 q}{\partial x_i \partial x_j}(x)| \leq 36$. Similarly, we also know for any i , $|\frac{\partial^2 q}{(\partial x_i)^2}(x)| \leq 36$. Therefore, the Hessian matrix $\nabla^2 q(x)$ is a $k \times k$ matrix with entries no larger than 36, so we have $\|\nabla^2 q(x)\| \leq \|\nabla^2 q(x)\|_F \leq 36k$.

Property 3 follows immediately from $p(x) \geq 1/3$ when $x \in [-1/2, 1/2]$. \square

Proof of Lemma 5.3 Using the construction of q , one can select a type for each of the cell and construct a corresponding function as in Lemma 5.3. We give the proof of the lemma here:

Proof. First, by Lemma D.1, the q function has 0 value, gradient and Hessian at the boundary. Therefore the function value, gradient and Hessian of $f(x)$ agrees with $f_0(x)$ on the boundary. As a result, the function we construct is still twice differentiable on every point.

For any cell τ , by Lemma D.1 the function $q\left(\frac{1}{l}(x - v_\tau)\right)$ for $x \in \tau$ has Hessian bounded by $\frac{36k}{l^2}$. We will make sure that every c_τ is bounded by $\frac{l^2}{72k}$, so the function $c_\tau q\left(\frac{1}{l}(x - v_\tau)\right)$ has a Hessian with spectral norm at most 1/2. Since $\nabla^2 f(x) = \nabla^2 f_0(x) + c_\tau \nabla^2 q\left(\frac{1}{l}(x - v_\tau)\right)$, by standard matrix perturbation bounds, the Hessian of f always satisfies $0.5I \preceq \nabla^2 f \preceq 1.5I$, which implies $f(x)$ is 1.5-smooth and 0.5-strongly convex.

For the second property, note that $f(x) \geq f_0(x)$ as both c_τ and q are positive. Therefore $\int e^{-f(x)} dx$ is always smaller than $\int e^{-f_0(x)} dx$. For each cell τ , let

$$C_\tau = \int_{x \in \tau} \left[\exp(-f_0(x)) - \exp\left(-f_0(x) + \frac{l^2}{72k} q\left(\frac{1}{l}(x - v_\tau)\right)\right) \right] dx.$$

Therefore C_τ is the amount of decrease in normalizing constant if we choose $c_\tau = \frac{l^2}{72k}$ (the maximum allowed value). We first show a lowerbound on C_τ :

$$\begin{aligned}
C_\tau &= \int_{x \in \tau} \left[\exp(-f_0(x)) - \exp\left(-f_0(x) + \frac{l^2}{72k} q\left(\frac{1}{l}(x - v_\tau)\right)\right) \right] dx \\
&= \int_{x \in \tau} e^{-f_0(x)} \left[1 - \exp\left(-\frac{l^2}{72k} q\left(\frac{1}{l}(x - v_\tau)\right)\right) \right] dx \\
&\geq \int_{\|x - v_\tau\|_\infty \leq l/2} e^{-f_0(x)} \left[1 - \exp\left(-\frac{l^2}{72k} q\left(\frac{1}{l}(x - v_\tau)\right)\right) \right] dx \\
&\geq \int_{\|x - v_\tau\|_\infty \leq l/2} e^{-f_0(x)} (1 - \exp(-\frac{l^2}{72k3^k})) dx \\
&\geq \int_{\|x - v_\tau\|_\infty \leq l/2} e^{-1} (1 - \exp(-\frac{l^2}{72k3^k})) dx \\
&= \Omega(l^{2+k}) = \Omega\left(\frac{l^2}{n}\right).
\end{aligned}$$

Let τ^* be the cell with the smallest C_{τ^*} , set $c_{\tau^*} = \frac{l^2}{72k}$. Set all the c_τ 's carefully in $[0, \frac{l^2}{72k}]$ so that the decrease in every cell is equal to C_{τ^*} (this is always possible because the amount of decrease is continuous and monotonically increasing with respect to c_τ), and we have the second property. \square

Proof of Theorem 5.2 Now we are ready to prove the lowerbound Theorem 5.2 for a constant number of dimensions.

Proof. Fix an desired accuracy δ small enough, choose $n \geq 100/\delta^2$ and make sure $n^{1/k}$ is an integer (when $\delta < 1$ we still have $n = \mathcal{O}(1/\delta^2)$).

Consider two distributions of functions \mathcal{F}_1 and \mathcal{F}_2 . In \mathcal{F}_1 , each cell is of type 1 with probability $1/2 + \delta$ independently, in \mathcal{F}_2 , each cell is of type 1 with probability $1/2 - \delta$ independently. After the types of cells are decided, function f is constructed according to Lemma 5.3.

Clearly, querying any point of $f(x)$ (whether the query is on function value or gradient) can give information about at most one cell. Therefore by Claim D.1, any algorithm that makes $o(1/\delta^2)$ queries will not be able to distinguish whether the function comes from \mathcal{F}_1 or \mathcal{F}_2 with probability better than 0.6.

On the other hand, by standard concentration bounds and the fact that $n \geq 100/\delta^2$, we know with at least 0.99 probability functions in \mathcal{F}_1 has at most $n(1 - \delta)/2$ type 2 cells, and functions in \mathcal{F}_2 has at least $n(1 + \delta)/2$ type 2 cells. By Lemma 5.3, we know with probability at least 0.99, the normalizing constant $Z \geq (2\pi)^{k/2} - C(1 - \delta)/2 =: \theta_1$ for $f \sim \mathcal{F}_1$, and $Z \leq (2\pi)^{k/2} - C(1 + \delta)/2 =: \theta_2$ for $f \sim \mathcal{F}_2$. Therefore, if an algorithm can estimate the normalizing constant with accuracy better than $\sqrt{\theta_1/\theta_2} - 1$ with probability 3/4, it is going to be able to distinguish \mathcal{F}_1 and \mathcal{F}_2 with probability better than 0.6, which is impossible.

Now, by Lemma 5.3, we know $C = \Omega(l^2)$, therefore $\theta_1/\theta_2 = 1 + \Omega(C\delta) = 1 + \Omega(l^2\delta)$. Using the fact that $l = 1/(\sqrt{k}n^{1/k})$ and $n = \Theta(1/\delta^2)$, we know $\sqrt{\theta_1/\theta_2} - 1 = \Omega(\delta^{1+4/k})$. The Theorem follows by choosing δ such that $\varepsilon = \Theta(\delta^{1+4/k})$. \square

Proof of Theorem 5.1 Finally we extend Theorem 5.2 to Theorem 5.1.

Proof of Theorem 5.1. The proof is very similar to Theorem 5.2. Fix a constant k depending only on γ that we will determine later. We will break the d coordinates of input x into $d' = \lfloor d/k \rfloor$ groups of size k each (ignoring the remainder). Let x_{S_i} be the input x restrict to the i -th group of coordinates. The function we construct will be a sum of functions $f(x) = \sum_{i=1}^{d'} f_i(x_{S_i})$.

Fix an desired accuracy δ small enough, choose $n \geq 100/\delta^2$ and make sure $n^{1/k}$ is a constant (when δ is small enough we still have $n = \mathcal{O}(1/\delta^2)$).

Consider two distributions of functions \mathcal{F}_1 and \mathcal{F}_2 same as in the proof of Theorem 5.2. When $f \sim \mathcal{F}_1$, construct $f_1, f_2, \dots, f_{d'}$ independently using Lemma 5.3, where each cell is of type 1 with probability $1/2 + \delta$; when $f \sim \mathcal{F}_2$, construct $f_1, f_2, \dots, f_{d'}$ independently using Lemma 5.3, where each cell is of type 1 with probability $1/2 - \delta$.

It is easy to see that the normalizing constant for f is the product of normalizing constant of $f_1, f_2, \dots, f_{d'}$. By construction in Lemma 5.3 and calculations in Theorem 5.2, there exists a constant Z such that the normalizing constant for f_i is $Z(1 + \Omega(\delta^{1+4/k}))$ with probability at least 0.99 when $f \sim \mathcal{F}_1$, and $Z(1 - \Omega(\delta^{1+4/k}))$ with probability at least 0.99 when $f \sim \mathcal{F}_2$. When $\delta^{1+4/k}d \leq 1/5$, by Lemma B.2 we know with probability at least 0.99, the normalizing constant for $f \sim \mathcal{F}_1$ is at least $Z^{d'}(1 + \Omega(\delta^{1+4/k}d)) =: \theta_1$, and the normalizing constant for $f \sim \mathcal{F}_2$ is at most $Z^{d'}(1 - \Theta(\delta^{1+4/k}d)) =: \theta_2$. When the number of queries is $o(1/d\delta^2)$, no algorithm can distinguish between these two distributions, which means no algorithm can estimate the normalizing constant with accuracy better than $\sqrt{\theta_1/\theta_2} - 1 = \Theta(\delta^{1+4/k}d)$.

If we set $\varepsilon = \Theta(\delta^{1+4/k}d)$, then any algorithm that uses $o\left(d^{\frac{1-4/k}{1+4/k}} \varepsilon^{-\frac{2}{1+4/k}}\right)$ queries cannot estimate the normalizing constant with multiplicative error $1 \pm \varepsilon$ with probability better than 3/4. Finally, we choose $k = \lceil 8/\gamma \rceil$, so $\frac{2}{1+4/k} \geq 2 - \gamma$ and $\frac{1-4/k}{1+4/k} \geq 1 - \gamma$, which gives the guarantee in the theorem. \square

E Quadrature Method for Estimating the Normalizing Constant

Alternative to the Monte Carlo strategy as discussed, for lower dimensions, a deterministic quadrature scheme for $Z = \int e^{-f(x)} dx$ might be computationally less expensive.

First, we recall that for $X \in \mathbb{R}^d$ a random variable distributed according to a logconcave distribution with $\mathbb{E}(\|X\|^2) \leq R^2$. Restricted the support of X to a ball with radius $2R\log(1/\varepsilon)$ captures at least $1 - \varepsilon/2$ fraction of the mass. Thus it suffices to integrate $e^{-f(x)}$ inside a square Q_{R_0} centered at the origin of radius $R_0 = 2\sqrt{d/\mu\log(1/\varepsilon)}$.

Inside the square Q_{R_0} , we use a trapezoidal quadrature rule with grid spacing h to integrate $e^{-f(x)}$. Denote the estimate from quadrature as S_h , the error is bounded from above by

$$\left| \int_{Q_{R_0}} e^{-f(x)} dx - S_h \right| \leq C \text{vol}(Q_{R_0}) h^2 d^2 \max_{x \in Q_{R_0}} \|\nabla^2(\exp(-f(x)))\|. \quad (31)$$

The Hessian of $e^{-f(x)}$ can be bounded from above by

$$\begin{aligned} \max_{x \in Q_{R_0}} \|\nabla^2(\exp(-f(x)))\| &= \max_{x \in Q_{R_0}} (\|\nabla^2 f(x)\| + \|\nabla f(x)\|^2) \exp(-f(x)) \\ &\leq \max_{x \in Q_{R_0}} (L + L^2 \|x\|^2) \exp(-\frac{\mu}{2} \|x\|^2) \\ &\leq L + \frac{2L^2}{e\mu} \\ &\leq CL(1 + \kappa). \end{aligned}$$

Thus, to make the right hand side of (31) error ε , we need

$$h \leq Cd^{-1}L^{-1/2}(1+\kappa)^{-1/2}\text{vol}(Q_{R_0})^{-1/2}\varepsilon^{-1/2}.$$

The number of quadrature points is given by

$$\begin{aligned} N &= \mathcal{O}(\text{vol}(Q_{R_0})^{1+d/2}L^{d/2}(1+\kappa)^{d/2}\varepsilon^{d/2}) \\ &= \tilde{\mathcal{O}}\left(\left(\frac{d}{\mu}\right)^{d/2+d^2/4}d^dL^{d/2}(1+\kappa)^{d/2}\varepsilon^{d/2}\right). \end{aligned}$$

While this complexity has a better dependence in ε for low dimension ($d \leq 3$), the dependence in dimension is much worse than that of the Monte Carlo method.

F Tools and Auxiliary Lemmas

We note some concentration results and functional inequalities for log-concave distributions.

Lemma F.1 (Concentration around mode for log-concave distributions). *Suppose $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is a convex $\frac{1}{\sigma^2}$ -strongly convex function with minimum at 0, and let ρ be a probability measure on \mathbb{R}^d with $\rho(dx) \propto e^{-f(x)} dx$. Then for any r , $\mathbb{P}_{x \sim \rho}(\|x\| \geq r) \leq \mathbb{P}_{x \sim N(0, \sigma^2)}(\|x\| \geq r)$.*

Proof. Without loss of generality, $f(0) = 0$. Using spherical coordinates, we have

$$\mathbb{P}_{x \sim \rho}(\|x\| \geq r) = \frac{\int_{\mathbb{S}^{d-1}} \int_r^\infty s^{d-1} e^{-f(sv)} ds d\mathbb{S}^{d-1}(v)}{\int_{\mathbb{S}^{d-1}} \int_0^\infty s^{d-1} e^{-f(sv)} ds d\mathbb{S}^{d-1}(v)} \quad (32)$$

Let

$$A(v) = \int_r^\infty s^{d-1} e^{-f(sv)} ds d\mathbb{S}^{d-1}(v) \quad C = \int_r^\infty s^{d-1} e^{-\frac{s^2}{2\sigma^2}} ds d\mathbb{S}^{d-1}(v) \quad (33)$$

$$B(v) = \int_0^r s^{d-1} f(sv) ds d\mathbb{S}^{d-1}(v) \quad D = \int_0^r s^{d-1} e^{-\frac{s^2}{2\sigma^2}} ds d\mathbb{S}^{d-1}(v). \quad (34)$$

We will show that $\frac{A(v)}{B(v)} \leq \frac{C}{D}$. Then $\frac{A(v)}{A(v) + B(v)} \leq \frac{C}{C+D}$, so

$$\frac{\int_{\mathbb{S}^{d-1}} \int_r^\infty s^{d-1} e^{-f(sv)} ds d\mathbb{S}^{d-1}(v)}{\int_{\mathbb{S}^{d-1}} \int_0^\infty s^{d-1} e^{-f(sv)} ds d\mathbb{S}^{d-1}(v)} = \frac{\int_{\mathbb{S}^{d-1}} A(v) d\mathbb{S}^{d-1}(v)}{\int_{\mathbb{S}^{d-1}} A(v) + B(v) d\mathbb{S}^{d-1}(v)} \quad (35)$$

$$\leq \sup_{v \in \mathbb{S}^{d-1}} \frac{A(v)}{A(v) + B(v)} \leq \frac{C}{C+D} = \mathbb{P}_{x \sim N(0, \sigma^2 I_d)}(\|x\| \geq r). \quad (36)$$

It suffices to show $\frac{A(v)}{B(v)} \leq \frac{C}{D}$. For this, we first prove the following claim: If a, c are positive functions on Ω_1 , and b, d are nonnegative functions on Ω_2 , then $\frac{\int_{\Omega_1} a(x) dx}{\int_{\Omega_2} b(x) dx} \leq \frac{\int_{\Omega_1} c(x) dx}{\int_{\Omega_2} d(x) dx}$.

To see the claim, note that

$$\frac{\int_{\Omega_1} c(x) dx}{\int_{\Omega_2} d(x) dx} = \frac{\int_{\Omega_1} a(x) \cdot \frac{c(x)}{a(x)} dx}{\int_{\Omega_2} b(x) \cdot \frac{d(x)}{b(x)} dx} \geq \frac{\int_{\Omega_1} a(x) dx \cdot \inf_{\Omega_1} \frac{c}{a}}{\int_{\Omega_2} b(x) dx \cdot \inf_{\Omega_1} \frac{d}{b}} \geq \frac{\int_{\Omega_1} a(x) dx}{\int_{\Omega_2} b(x) dx}.$$

Now we show that the claim implies $\frac{A(v)}{B(v)} \leq \frac{C}{D}$. We have

$$\inf_{s \in [r, \infty)} e^{-\frac{s^2}{2\sigma^2} + f(s)} \leq \sup_{s \in [0, r]} e^{-\frac{s^2}{2\sigma^2} + f(s)}$$

because $f(s) - \frac{s^2}{2\sigma^2}$ is an increasing function. Thus the claim implies that $\frac{A(v)}{B(v)} \leq \frac{C}{D}$. \square

Lemma F.2. Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be a m -strongly convex function and let $\rho(dx) \propto e^{-f(x)} dx$. Let $x^* = \operatorname{argmin}_x f(x)$ be the mode. Then

$$\mathbb{E}_{x \sim \rho} \|x - x^*\| \leq \frac{1}{\sqrt{m}} (\sqrt{d} + 2\sqrt{2 \log 2}).$$

Proof. By Lemma F.1 and the χ^2 tail bound from [LM00],

$$\begin{aligned} & \mathbb{P} \left(\|x - x^*\|^2 \geq \frac{1}{m} \left(d + 2(\sqrt{d \log 2} + \log 2) \right) \right) \\ & \leq \mathbb{P}_{x \sim N(0, \frac{\sigma^2}{m} I_d)} \left(\|x - x^*\|^2 \geq \frac{1}{m} \left(d + 2(\sqrt{d \log 2} + \log 2) \right) \right) \leq \frac{1}{2} \end{aligned}$$

$$\text{so } \mathbb{P} \left(\|x - x^*\| \geq \frac{1}{\sqrt{m}} \left(\sqrt{d} + \sqrt{2 \log 2} \right) \right) \leq \frac{1}{2}.$$

By Theorem F.3 and Theorem F.4,

$$\mathbb{P} \left(\|x - x^*\| \leq \mathbb{E} \|x - x^*\| - \frac{c}{\sqrt{m}} \right) \leq \exp \left(-\frac{c^2}{2} \right). \quad (37)$$

Taking $c = \sqrt{2 \log 2}$, we get this is $\leq \frac{1}{2}$.

Hence the sets $\left\{ x : \|x - x^*\| \leq \frac{1}{\sqrt{m}} \left(\sqrt{d} + \sqrt{2 \log 2} \right) \right\}$ and $\left\{ x : \|x - x^*\| \geq \mathbb{E} \|x - x^*\| - \frac{\sqrt{2 \log 2}}{\sqrt{m}} \right\}$ must intersect, so

$$\mathbb{E} \|x - x^*\| - \frac{\sqrt{2 \log 2}}{\sqrt{m}} \leq \frac{1}{\sqrt{m}} \left(\sqrt{d} + \sqrt{2 \log 2} \right),$$

as needed. \square

Theorem F.3 (Bakry-Émery [BÉ85, BGL13]). Suppose f is μ -strongly convex. Then $\rho(dx) \propto e^{-f(x)} dx$ satisfies a Poincaré inequality with constant $\frac{1}{\mu}$ ($\frac{1}{\mu} \int_{\mathbb{R}^d} \|\nabla g(x)\|^2 dx \geq \operatorname{Var}_{\rho}(g)$ for all g where the integral is defined) and a log-Sobolev inequality with constant $\frac{1}{\mu}$.

Theorem F.4 (Log-Sobolev inequality implies Gaussian measure concentration, [BGL13, (5.4.2)]). Suppose $\rho(dx)$ is a distribution on \mathbb{R}^d that satisfies a log-Sobolev inequality with constant C . Let $g : \mathbb{R}^d \rightarrow \mathbb{R}$ be L -Lipschitz. Then

$$\mathbb{P} (|g - \mathbb{E}_{\rho} g| \geq r) \leq 2 \exp \left(-\frac{r^2}{2CL^2} \right).$$