
Sampling from Log-Concave Distributions with Infinity-Distance Guarantees

Oren Mangoubi
Worcester Polytechnic Institute

Nisheeth K. Vishnoi
Yale University

Abstract

For a d -dimensional log-concave distribution $\pi(\theta) \propto e^{-f(\theta)}$ constrained to a convex body K , the problem of outputting samples from a distribution ν which is ε -close in infinity-distance $\sup_{\theta \in K} |\log \frac{\nu(\theta)}{\pi(\theta)}|$ to π arises in differentially private optimization. While sampling within total-variation distance ε of π can be done by algorithms whose runtime depends polylogarithmically on $\frac{1}{\varepsilon}$, prior algorithms for sampling in ε infinity distance have runtime bounds that depend polynomially on $\frac{1}{\varepsilon}$. We bridge this gap by presenting an algorithm that outputs a point ε -close to π in infinity distance that requires at most $\text{poly}(\log \frac{1}{\varepsilon}, d)$ calls to a membership oracle for K and evaluation oracle for f , when f is Lipschitz. Our approach departs from prior works that construct Markov chains on a $\frac{1}{\varepsilon^2}$ -discretization of K to achieve a sample with ε infinity-distance error, and present a method to directly convert continuous samples from K with total-variation bounds to samples with infinity bounds. This approach also allows us to obtain an improvement on the dimension d in the running time for the problem of sampling from a log-concave distribution on polytopes K with infinity distance ε , by plugging in TV-distance running time bounds for the Dikin Walk Markov chain.

1 Introduction

The problem of sampling from a log-concave distribution is as follows: For a convex body $K \subseteq \mathbb{R}^d$ and a convex function $f : K \rightarrow \mathbb{R}$, output a sample θ from the distribution $\pi(\theta) \propto e^{-f(\theta)}$. This is a basic problem in computer science, statistics, and machine learning, with applications to optimization and integration [1, 29], Bayesian statistics [39], reinforcement learning [6], and differential privacy [32, 20, 2, 24]. Sampling exactly from π is known to be computationally hard for most interesting cases of K and f [16] and, hence, the goal is to output samples from a distribution ν that is at a small (specified) “distance” to π . For applications such as computing the integral of π , bounds in the total variation (TV) distance [1] or KL divergence (which implies a TV bound) are sufficient. In applications such as computing the expectation of a Lipschitz function with respect to π , Wasserstein distance may also be sufficient. In differentially private optimization [32, 20, 2, 19, 24], one requires bounds on the stronger infinity-distance –

$$d_{\infty}(\nu, \pi) := \sup_{\theta \in K} \left| \log \frac{\nu(\theta)}{\pi(\theta)} \right|$$

– to guarantee pure differential privacy, and TV, KL, or Wasserstein bounds are insufficient; see [14].

Pure differential privacy (DP) is the strongest notion of DP and has been extensively studied (see e.g. the survey [14]). It has advantages over weaker notions of differential privacy. E.g., when privacy of “groups” of individuals (rather than just single individuals) must be preserved, any mechanism which is (pure) ε -DP (with respect to single individuals), is also $k\varepsilon$ -DP with respect to subsets of k individuals. Motivated by applications to differential privacy, we study the problem of designing efficient algorithms to output samples from a distribution ν which is ε -close in d_{∞} to π .

Related works. Several lines of work have designed Markov chains that generate samples from distributions that are close to a given log-concave distribution. These results differ in both their assumptions on the log-density and its support, as well as the distance used to measure closeness. One line of work includes bounds for sampling from a log-concave distribution on a compactly supported convex body within TV distance $O(\delta)$, including results with running time that is polylogarithmic in $\frac{1}{\delta}$ [1, 30, 29, 34] (as well as other results which give a running time bound that is polynomial in $\frac{1}{\delta}$ [18, 17, 5, 4]). In addition to assuming access to a value oracle for f , some Markov chains just need access to a membership oracle for K [1, 30, 29], while others assume that K is a given polytope: $\{\theta \in \mathbb{R}^d : A\theta \leq b\}$ [23, 33, 35, 34, 26]. They often also assume that K is contained in a ball of radius R and contains a ball of smaller radius r . Many of these results assume that the target log-concave distribution satisfies a “well-rounded” condition which says that the variance of the target distribution is $\Theta(d)$ [30, 29], or that it is in isotropic position (all eigenvalues of its covariance matrix are $\Theta(1)$) [25]; when applied to log-concave distributions that are not well-rounded or isotropic, these results require a “rounding” pre-processing procedure to find a linear transformation which makes the target distribution well-rounded or isotropic. Finally, it is often assumed that the function f is such that f is L -Lipschitz or β -smooth [34], including works handling the widely-studied special case when f is uniform on K where $L = \beta = 0$ (see e.g. [28, 23, 33, 35, 26, 8, 21]).

Another line of work gives sampling algorithms with bounds on the distance to the target density π in terms of Wasserstein distance [13, 11], KL divergence [40, 12], and Renyi divergence [36]. In contrast to works which assume access to an oracle for the value of f , many of these results instead assume access to an oracle for the gradient of f and require the log-density to be L -Lipschitz or β -smooth on all of \mathbb{R}^d (or on, e.g., a cube containing K) for some $L, \beta > 0$. However, as noted earlier, bounds in the Wasserstein distance, KL divergence, and α -Renyi divergence (for $\alpha < \infty$) also do not imply bounds on the infinity distance, and the running time bounds provided by these works are polynomial in $\frac{1}{\varepsilon}$. (See also Appendix A for additional discussion and challenges.)

Among prior works that give algorithms with bounds on d_∞ , [20] applies the grid walk Markov chain of [1] to sample from a uniform distribution on a convex body. [2] extends the approach of [20] to log-Lipschitz log-concave distributions. Unlike the TV-distance case where algorithms whose running time depends *logarithmically* on the error are known (e.g., [1, 29, 34]), the best available bounds for sampling within $O(\varepsilon)$ infinity-distance [20, 2] have runtime that is polynomial in $\frac{1}{\varepsilon}$ and a relatively large polynomial in d .

Our contributions. We present a new approach to output samples, which come with d_∞ bounds, from a log-concave and log-Lipschitz distribution constrained to a convex body. Specifically, when $K := \{\theta : A\theta \leq b\}$ is a polytope (where one is given A and b) our main result (Theorem 2.1) guarantees samples from a distribution that is within $O(\varepsilon)$ error in d_∞ and whose runtime depends logarithmically on $\frac{1}{\varepsilon}$ compared to the polynomial dependence of [2]. Our approach departs from prior works that construct Markov chains on a $\frac{1}{\varepsilon^2}$ -discretization of K to convert a sample with ε infinity-distance error, and we present a method (Algorithm 1) to directly convert continuous samples from K with total-variation bounds to samples with infinity bounds (Theorem 2.2). This continuous-space approach also allows us to obtain an improvement on the dimension d in the running time when K is a polytope by plugging in TV-distance running time bounds for the Dikin Walk Markov chain of [34]. As immediate applications, we obtain faster algorithms for differentially private empirical risk minimization (Corollary 2.4) and low rank approximation (Corollary 2.5).

2 Results

Let $B(v, s) := \{z \in \mathbb{R}^d : \|z - v\|_2 \leq s\}$ and ω denote the matrix-multiplication constant.

Theorem 2.1 (Main result) *There exists an algorithm which, given $\varepsilon, L, r, R > 0$, $A \in \mathbb{R}^{m \times d}$, $b \in \mathbb{R}^m$ that define a polytope $K := \{\theta \in \mathbb{R}^d : A\theta \leq b\}$ contained in a ball of radius R , a point $a \in \mathbb{R}^d$ such that K contains a ball $B(a, r)$ of smaller radius r , and an oracle for the value of a convex function $f : K \rightarrow \mathbb{R}^d$, where f is L -Lipschitz, and defining π to be the distribution $\pi \propto e^{-f}$, outputs a point from a distribution ν such that $d_\infty(\nu, \pi) < \varepsilon$. Moreover, with very high probability¹ this algorithm takes $O(T)$ function evaluations and $O(T \times md^{\omega-1})$ arithmetic operations, where $T = O((m^2d^3 + m^2dL^2R^2) \times [LR + d \log(\frac{Rd+LRd}{r\varepsilon})])$.*

¹The number of steps is $O(\tau \times T)$, where $\mathbb{E}[\tau] \leq 3$, $\mathbb{P}(\tau \geq t) \leq (\frac{2}{3})^t$ for $t \geq 0$, and $\tau \leq O(d \log(\frac{R}{r}) + LR)$ w.p. 1.

In comparison to the polynomial in $\frac{1}{\varepsilon}$ runtime bounds of [2], Theorem 2.1 guarantees a runtime that is *logarithmic* in $\frac{1}{\varepsilon}$, and also improves the dependence on the dimension d , in the setting where K is a polytope. Specifically, [2] show that the number of steps of the grid walk to sample from π with infinity-distance error at most ε is

$$O\left(\frac{1}{\varepsilon^2}(d^{10} + d^6 L^4 R^4) \times \text{polylog}\left(\frac{1}{\varepsilon}, \frac{1}{r}, R, L, d\right)\right)$$

(Lemma 6.5 in the Arxiv version of [2]). When applying their algorithm to the setting where f is constrained to a polytope $K = \{x \in \mathbb{R}^d : Ax \leq b\}$, each step of their grid walk Markov chain requires computing a membership oracle for K and the value of the function f . The membership oracle can be computed in $O(md)$ arithmetic operations. Thus the bound on the number of arithmetic operations for each step of their grid walk is $O(md)$ (provided that each function evaluation takes at most $O(md)$ arithmetic operations). Thus the bound on the number of arithmetic operations to obtain a sample from π is $O(\frac{1}{\varepsilon^2}(md^{11} + md^7 L^4 R^4) \times \text{polylog}(\frac{1}{\varepsilon}, \frac{1}{r}, R, L, d))$. Thus, Theorem 2.1 improves on this bound by a factor of roughly $\frac{1}{\varepsilon^2 m^3} d^{8-\omega}$. For example, when $m = O(d)$, as may be the case in differentially private applications, the improvement is $\frac{1}{\varepsilon^2} d^{5-\omega}$.

We note that the bounds of [2] also apply in the more general setting where K is a convex body with membership oracle. One can extend our bounds to achieve a runtime that is logarithmic in $\frac{1}{\varepsilon}$ (and polynomial in d, L, R) in the more general setting where K is a convex body with membership oracle; we omit the details (see Remark 2.3).

Moreover, we also note that while there are several results which achieve $O(\delta)$ TV bounds in time logarithmic in $\frac{1}{\delta}$, TV bounds do not in general imply $O(\varepsilon)$ bounds on the KL or Renyi divergence, or on the infinity-distance, for any $\delta > 0$.² On the other hand, an ε -infinity-distance bound does immediately imply a bound of ε on the KL divergence D_{KL} , and α -Renyi divergence D_α , since $D_{\text{KL}}(\mu, \pi) \leq d_\infty(\mu, \pi)$ and $D_\alpha(\mu, \pi) \leq d_\infty(\mu, \pi)$ for any $\alpha > 0$ and any pair of distributions μ, π . Thus, under the same assumptions on K and f , Theorem 2.1 implies a method of sampling from a Lipschitz concave log-density on K with ε KL and Renyi divergence error in a number of arithmetic operations that is *logarithmic* in $\frac{1}{\varepsilon}$, with the same bound on the number of arithmetic operations.

The polynomial dependence on $\frac{1}{\varepsilon}$ in [2] is due to the fact that they rely on a discrete-space Markov chain [1], on a grid with cells of width $w = O(\frac{\varepsilon}{L\sqrt{d}})$, to sample from π within $O(\varepsilon)$ infinity-distance.

Since their Markov chain's runtime bound is polynomial in w^{-1} , they get a runtime bound for sampling within $O(\varepsilon)$ infinity-distance that is polynomial in $\frac{1}{\varepsilon}$. The proof of Theorem 2.1 bypasses the use of discrete grid-based Markov chains by introducing Algorithm 1 which transforms any sample within $\delta = O(\varepsilon e^{-d-nLR})$ TV distance of the distribution $\pi \propto e^{-f}$ on the *continuous* set K (as opposed to a discretization of K), into a sample within $O(\varepsilon)$ infinity-distance from π . This allows us to make use of a continuous-space Markov chain, whose step size is not restricted to a grid of width $O(\frac{\varepsilon}{L\sqrt{d}})$ and is instead independent of ε , to obtain a sample within $O(\varepsilon)$ infinity-distance from π in time that is logarithmic in $\frac{1}{\varepsilon}$.

Theorem 2.2 (Main technical contribution: Converting TV bounds to infinity-distance bounds)

There exists an algorithm (Algorithm 1) which, given $\varepsilon, r, R, L > 0$, a membership oracle for a convex body K contained in a ball of radius R and containing a ball $B(0, r)$, and an oracle which outputs a point from a distribution μ which has TV distance

$$\delta \leq O\left(\varepsilon \times \left(\frac{R(d \log(R/r) + LR)^2}{\varepsilon r}\right)^{-d} e^{-LR}\right)$$

from a distribution $\pi \propto e^{-f}$ where $f : K \rightarrow \mathbb{R}$ is an L -Lipschitz function (see Appendix B for the exact values of δ and related hyper-parameters), outputs a point $\hat{\theta} \in K$ such that the distribution ν

²For instance, if $\pi(\theta) = 1$ with support on $[0, 1]$, for every $\delta > 0$ there is a distribution ν where $\|\nu - \pi\|_{\text{TV}} \leq 2\delta$ and yet $d_\infty(\nu, \pi) \geq D_{\text{KL}}(\nu, \pi) \geq \frac{1}{2}$. ($\nu(\theta) = e^{\frac{1}{\delta}}$ on $\theta \in [0, \delta e^{-\frac{1}{\delta}}]$, $\nu(\theta) = \frac{1-\delta}{1-\delta e^{-\frac{1}{\delta}}}$ on $(\delta e^{-\frac{1}{\delta}}, 1]$ and $\nu(\theta) = 0$ otherwise)

of $\hat{\theta}$ satisfies $d_\infty(\nu, \pi) \leq \varepsilon$. Moreover, with very high probability³ this algorithm finishes in $O(1)$ calls to the sampling and membership oracles, plus $O(d)$ arithmetic operations.

To the best of our knowledge Theorem 2.2 is the first result which for any $\varepsilon, L, r, R > 0$, when provided as input a sample from a continuous-space distribution on a convex body K within some TV distance $\delta = \delta(\varepsilon, r, R, L) > 0$ from a given L -log-Lipschitz distribution π on K , where K is contained in a ball of radius R and containing a ball of smaller radius r , outputs a sample with distribution within infinity-distance $O(\varepsilon)$ from π . This is in contrast to previous works [2] (see also [20] which applies only to the special case where π is the uniform distribution on K) which instead require as input a sample with bounded TV distance from the restriction of π on a discrete grid on K , and then convert this discrete-space sample into a sample within infinity-distance $O(\varepsilon)$ from the continuous-space distribution $\pi : K \rightarrow \mathbb{R}$.

Algorithm 1: Interior point TV to infinity-distance converter

Input: $d \in \mathbb{N}$

Input: A membership oracle for a convex body $K \in \mathbb{R}^d$ and an $r > 0$ such that $B(0, r) \subseteq K$.

Input: A sampling oracle which outputs a point from a distribution $\mu : K \rightarrow \mathbb{R}$

Output: A point $\hat{\theta} \in K$.

```

1 Hyperparameters:  $\Delta > 0, \tau_{\max} \in \mathbb{N}$  (set in Appendix B)
2 for  $i = 1, \dots, \tau_{\max}$  do
3   Sample a point  $\theta \sim \mu$ 
4   Sample a point  $\xi \sim \text{Unif}(B(0, 1))$ 
5   Set  $Z \leftarrow \theta + \Delta r \xi$ 
6   Set  $\hat{\theta} \leftarrow \frac{1}{1-\Delta} Z$ 
7   If  $\hat{\theta} \in K$ , output  $\hat{\theta}$  with probability  $\frac{1}{2}$  and halt. Otherwise, continue.
8 end
9 Sample a point  $\hat{\theta} \sim \text{Unif}(B(0, r))$ 
10 Output  $\hat{\theta}$ 

```

Remark 2.3 (Extension to convex bodies with membership oracles) We note that Theorem 2.1 can be extended to the general setting where K is an arbitrary convex body in a ball of radius R and containing a ball of smaller radius r , and we only have membership oracle access to K . Namely, one can plug in the results of [29] to our Theorem 2.2 to generate a sample from a L -Lipschitz concave log-density on an arbitrary convex body K in a number of operations that is (poly)-logarithmic in $\frac{1}{\varepsilon}, \frac{1}{r}$ and polynomial on d, L, R . We omit the details.

Applications to differentially private optimization. Sampling from distributions with $O(\varepsilon)$ infinity-distance error has many applications to differential privacy. Here, the goal is to find a randomized mechanism $h : \mathcal{D}^n \rightarrow \mathcal{R}$ which, given a dataset $x \in \mathcal{D}^n$ consisting of n datapoints, outputs model parameters $\hat{\theta} \in \mathcal{R}$ in some parameter space \mathcal{R} , which minimize a given (negative) utility function $f(\theta, x)$, under the constraint that the output $\hat{\theta}$ preserves the pure ε -differential privacy of the data points x . A randomized mechanism $h : \mathcal{D}^n \rightarrow \mathcal{R}$ is said to be ε -differentially private if for any datasets $x, x' \in \mathcal{D}$ which differ by a single datapoint, and any $S \subseteq \mathcal{R}$, we have that

$$\mathbb{P}(h(x) \in S) \leq e^\varepsilon \mathbb{P}(h(x') \in S);$$

see [14].

As one application of Theorem 2.1, we consider the problem of finding an (approximate) minimum $\hat{\theta}$ of an empirical risk function $f : K \times \mathcal{D}^n \rightarrow \mathbb{R}$ under the constraint that the output $\hat{\theta}$ is ε -differentially private, where $f(\theta, x) := \sum_{i=1}^n \ell_i(\theta, x_i)$. Following [2], we assume that the $\ell_i(\cdot, x)$ are L -Lipschitz for all $x \in \mathcal{D}^n, i \in \mathbb{N}$, for some given $L > 0$. In this setting [2] show that the minimum ERM utility bound under the constraint that $\hat{\theta}$ is pure ε -differentially private, $\mathbb{E}_{\hat{\theta}}[f(\hat{\theta}, x)] - \min_{\theta \in K} f(\theta, x) = \Theta(\frac{dLR}{\varepsilon})$, is achieved if one samples $\hat{\theta}$ from the exponential mechanism $\pi \propto e^{-\frac{\varepsilon}{2LR} f}$ with infinity-distance error at most $O(\varepsilon)$. Plugging Theorem 2.1 into the framework of

³Algorithm 1 finishes in τ calls to the sampling and membership oracles, plus $O(\tau d)$ arithmetic operations, where $\mathbb{E}[\tau] \leq 3$ and $\mathbb{P}(\tau \geq t) \leq (\frac{2}{3})^t$ for all $t \geq 0$ and $\tau \leq 5d \log(\frac{R}{r}) + 5LR + 2$ w.p. 1.

the exponential mechanism, we obtain a pure ε -differentially private mechanism which achieves the minimum expected risk (Corollary 2.4, see Section C.1 for a proof).

Corollary 2.4 (Differentially private empirical risk minimization) *There exists an algorithm which, given $\varepsilon, L, r, R > 0, A \in \mathbb{R}^{m \times d}, b \in \mathbb{R}^m$ that define a polytope $K := \{\theta \in \mathbb{R}^d : A\theta \leq b\}$ contained in a ball of radius R and containing a ball $B(0, r)$ of smaller radius r , and a convex function $f(\theta, x) := \sum_{i=1}^n \ell_i(\theta, x_i)$, where each $\ell_i : K \rightarrow \mathbb{R}$ is L -Lipschitz, outputs a random point $\hat{\theta} \in K$ which is pure ε -differentially private and satisfies*

$$\mathbb{E}_{\hat{\theta}}[f(\hat{\theta}, x)] - \min_{\theta \in K} f(\theta, x) \leq O\left(\frac{dLR}{\varepsilon}\right).$$

Moreover, this algorithm takes at most $T \times md^{\omega-1}$ arithmetic operations plus T evaluations of the function f , where $T = O\left((m^2d^3 + m^2dn^2\varepsilon^2) \times (\varepsilon n + d) \log^2\left(\frac{nRd}{r\varepsilon}\right)\right)$.

Corollary 2.4 improves on the previous bound [2] of $O\left(\left(\frac{1}{\varepsilon^2}(m+n)d^{11} + \varepsilon^2n^4(m+n)d^7\right) \times \text{polylog}\left(\frac{nRd}{r\varepsilon}\right)\right)$ arithmetic operations by a factor of roughly $\max\left(\frac{d^8-\omega}{\varepsilon^2m^2}, \frac{1}{\varepsilon m^2}nd^5\right)$, in the setting where the ℓ_i are L -Lipschitz on a polytope K and each ℓ_i can be evaluated in $O(d)$ operations. See Appendix C.1 for a proof of this corollary.

As another application, we consider the problem of finding a low-rank approximation of a sample covariance matrix $\Sigma = \sum_{i=1}^n u_i u_i^\top$ where the datapoints $u_i \in \mathbb{R}^d$ satisfy $\|u_i\| \leq 1$, in a differentially private manner. Given any $k > 0$, the goal is to find a (random) rank- k projection matrix P which maximizes the average variance $\mathbb{E}_P[\langle \Sigma, P \rangle]$ of the matrix Σ (also referred to as the utility of P), under the constraint that the mechanism which outputs the matrix P is ε -differentially private. This problem has many applications to statistics and machine learning, including differentially private principal component analysis (PCA) [7, 3, 15, 24].

When privacy is not a concern, the solution P which maximizes the variance is just a projection matrix onto the subspace spanned by top- k eigenvectors of Σ , and the maximum variance satisfies $\langle \Sigma, P \rangle = \sum_{i=1}^k \lambda_i$, where $\lambda_1 \geq \dots \geq \lambda_d > 0$ denote the eigenvalues of Σ . However, when privacy is a concern, there is a tradeoff between the desired privacy level ε and the utility $\mathbb{E}_P[\langle \Sigma, P \rangle]$, and the maximum utility $\mathbb{E}_P[\langle \Sigma, P \rangle]$ one can achieve decreases with the privacy parameter ε . The best current utility bound for an ε -differentially private low rank approximation algorithm was achieved in [24], who show that one can find a pure ε -differentially private random rank- k projection matrix P such that $\mathbb{E}_P[\langle \Sigma, P \rangle] \geq (1 - \delta) \sum_{i=1}^k \lambda_i$ whenever $\sum_{i=1}^k \lambda_i \geq \Omega\left(\frac{dk}{\varepsilon\delta} \log \frac{1}{\delta}\right)$ for any $\delta > 0$. To generate the matrix P , their algorithm generates a sample, with infinity-distance error $O(\varepsilon)$, from a Lipschitz concave log-density on a polytope, and transforms this sample into a projection matrix P . The sampling algorithm used in [24] has a bound of $\text{poly}\left(\frac{1}{\varepsilon}, d, \lambda_1 - \lambda_d\right)$ arithmetic operations and they leave as an open problem whether this can be improved from a polynomial dependence on $\frac{1}{\varepsilon}$ to a logarithmic dependence on $\frac{1}{\varepsilon}$. Corollary 2.5 shows that a direct application of Theorem 2.1 resolves this problem. (See Section C.2 for a proof.)

Corollary 2.5 (Differentially private low rank approximation) *There exists an algorithm which, given a sample covariance matrix $\Sigma = \sum_{i=1}^n u_i u_i^\top$ for datapoints $u_i \in \mathbb{R}^d$ satisfying $\|u_i\| \leq 1$, its eigenvalues $\lambda_1 \geq \dots \geq \lambda_d > 0$, an integer k , and $\varepsilon, \delta > 0$, outputs a random rank- k symmetric projection matrix P such that P is ε -differentially private and satisfies the utility bound*

$$\mathbb{E}_P[\langle \Sigma, P \rangle] \geq (1 - \delta) \sum_{i=1}^k \lambda_i(\Sigma)$$

whenever $\sum_{i=1}^k \lambda_i(\Sigma) \geq C \frac{dk}{\varepsilon\delta} \log \frac{1}{\delta}$ for some universal constant $C > 0$. Moreover the number of arithmetic operations is logarithmic in $\frac{1}{\varepsilon}$ and polynomial in d and $\lambda_1 - \lambda_d$.

3 Proof Overviews

Given any ε , and a function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, the goal is to sample from a distribution $\pi(\theta) \propto e^{-f(\theta)}$, constrained to a d -dimensional convex body K with infinity-distance error at most $O(\varepsilon)$ in a number

of arithmetic operations that is logarithmic in $\frac{1}{\varepsilon}$. We assume that K is contained in a ball of some radius $R > 0$ and contains a ball of some radius $r > 0$, and f is L -Lipschitz. In addition we would also like our bounds to be polylogarithmic in $\frac{1}{\varepsilon}$, and polynomial in d, L, R with a lower-order dependence on the dimension d than currently available bounds for sampling from Lipschitz concave log-densities on a polytope in infinity-distance [2]. We note that since whenever K is contained in a ball of radius R and contains a ball $B(0, r)$ of smaller radius r , we also have that $B(0, r) \subseteq K \subseteq B(0, 2R)$, without loss of generality, we may assume that $B(0, r) \subseteq K \subseteq B(0, R)$ as this would only change the bounds provided in our main theorems by a constant factor.

The main ingredient in the proof of Theorem 2.1 is Theorem 2.2 that uses Algorithm 1 to transform a TV-bounded sample into a sample from π with error bounded in d_∞ . Subsequently, we invoke Theorem 2.1 when K is given as a polytope $K := \{x \in \mathbb{R}^d : Ax \leq b\}$ and plug in the Dikin Walk Markov chain of [34] which generates independent samples from π with bounded TV error. We first present an overview of the proof of Theorem 2.2. (The full proof has been omitted to space restrictions and presented in Appendix B.) The proof of Theorem 2.1 is presented in Section 3.2.

3.1 Converting samples with TV bounds to infinity-distance bounds; proof of Theorem 2.2

Impossibility of obtaining log-dependence on infinity-distance via grid walk. One approach is to observe that if e^{-f} has support on a discrete space S with at most $|S|$ points, then any ν such that $\|\nu - \pi\|_{\text{TV}} \leq \varepsilon$ also satisfies

$$d_\infty(\nu, \pi) \leq 2|S| \frac{\max_{z \in S} e^{-f(z)}}{\min_{z \in S} e^{-f(z)}} \times \varepsilon$$

for any $\varepsilon \leq \min_{z \in S} \pi(z)$. This suggests forming a grid G over K , then using a discrete Markov chain to generate a sample θ within $O(\varepsilon)$ TV distance of the discrete distribution $\pi_G \propto e^{-f}$ with support on the grid G , and then designing an algorithm which takes as input θ and outputs a point with bounded infinity-distance to the continuous distribution π . This approach was used in [20] in the special case when π is uniform on K , and then extended by [2] to log-Lipschitz log-concave distributions. In their approach, [2] first run a “grid-walk” Markov chain on a discrete grid in a cube containing K . They then apply the bound from [1] which says that the grid walk obtains a sample Z within TV distance $O(\delta)$ from the distribution $\propto e^{-f}$ (restricted to the grid) in time that is polylogarithmic in $\frac{1}{\delta}$ and quadratic in a^{-1} , where a is the distance between neighboring grid points. Since their grid has size $|S| = \Theta\left(\left(\frac{R}{a}\right)^d\right)$, a TV distance of $O(\delta)$ automatically implies an infinity-distance of $O(\delta c |S|)$, where c is the ratio of the maximum to the minimum probability mass satisfying $c \leq e^{LR}$ since f is L -Lipschitz on $K \subseteq B(0, R)$. Thus, by using the grid walk to sample within TV-distance $\delta = O\left(\frac{\varepsilon}{|S|c}\right)$ from the discrete distribution π_G , they obtain a sample Z which also has infinity-distance $O(\varepsilon)$ from π_G . Finally, to obtain a sample from the distribution $\pi \propto e^{-f}$ on the continuous space K , they sample a point uniformly from the “grid cell” $[Z - a, Z + a]^d$ centered at Z . Since f is L -Lipschitz, the ratio $\frac{e^{-f(z)}}{e^{-f(w)}}$ is bounded by $O(\varepsilon)$ for all w in the grid cell $[Z - a, Z + a]^d$ as long as $a = O\left(\frac{\varepsilon}{L\sqrt{d}}\right)$, implying that the sample is an infinity-distance of $O(\varepsilon)$ from $\pi \propto e^{-f}$. However, since the running time bound of the grid walk is quadratic in a^{-1} , the grid coarseness $a = O\left(\frac{\varepsilon}{L\sqrt{d}}\right)$ needed to achieve $O(\varepsilon)$ infinity-distance from π leads to a running time bound which is *quadratic* in $\frac{1}{\varepsilon}$.

To get around this problem, rather than relying on the use of a discrete-space Markov chain such as the grid walk to sample within $O(\varepsilon)$ infinity-distance from π , we introduce an algorithm (Algorithm 1) which transforms any sample within $\delta = O(\varepsilon e^{-d-LR})$ TV distance from the distribution $\pi \propto e^{-f}$ on the *continuous* space K (as opposed to a grid-based discretization of K), into a sample within $O(\varepsilon)$ infinity-distance from π . This allows us to make use of a continuous-space Markov chain, such as the Dikin walk of [34], whose step-size is not restricted by a grid of coarseness $w = O\left(\frac{\varepsilon}{L\sqrt{d}}\right)$ and instead is independent of ε , in order to generate a sample within $O(\varepsilon)$ infinity-distance from π in runtime that is *logarithmic* in $\frac{1}{\varepsilon}$.

Converting continuous space TV-bounded samples to infinity-distance bounded samples. As discussed in Section 1, there are many Markov chain results which allow one to sample from a log-

the distribution π on K with infinity-distance error ε it is sufficient sample a point in the Δr -interior of K and to then apply a “stretching” operation to K .

We still need a method of sampling within $O(\varepsilon)$ infinity-distance error of π on all of K , including in the region $K \setminus \text{int}_{\Delta r}(K)$ near the boundary of K . Towards this end, after Algorithm 1 generates a point $Z = \theta + \xi$ from the above-mentioned distribution $\hat{\mu}$, it then multiplies Z by $\frac{1}{1-\Delta}$ and returns the resulting point $\hat{\theta} := \frac{1}{1-\Delta}Z$ if it is K , in other words, if $Z \in (1-\Delta)K$. If we can show that $(1-\Delta)K \subseteq \text{int}_{\Delta r}(K)$, then this would imply that $\left| \log \frac{\hat{\mu}(z)}{\pi(z)} \right| \leq \varepsilon$ for all $z \in (1-\Delta)K$, and hence that the distribution of $\hat{\nu}$ of the returned point $\hat{\theta}$ satisfies

$$\left| \log \frac{\hat{\nu}(z)}{\pi((1-\Delta)z)} \right| \leq \left| \log \frac{\hat{\mu}((1-\Delta)z)}{\pi((1-\Delta)z)} \right| + \log \frac{1}{(1-\Delta)^d} \leq O(\varepsilon)$$

for all $z \in K$. Since f is L -Lipschitz we have $\left| \log \frac{\pi(\theta)}{\pi((1-\Delta)\theta)} \right| = O(\varepsilon)$ for all $\theta \in K$, and hence we would then have that the distribution $\hat{\nu}$ of the point $\hat{\theta}$ returned by Algorithm 1 satisfies

$$\left| \log \frac{\hat{\nu}(z)}{\pi(z)} \right| = \left| \log \frac{\hat{\nu}(z)}{\pi((1-\Delta)z)} \right| + O(\varepsilon) \leq O(\varepsilon) \quad \forall z \in K. \quad (1)$$

However, for (1) to hold, we still need to show that $(1-\Delta)K \subseteq \text{int}_{\Delta r}(K)$ (proved in Lemma B.1). In other words, we would like to show that for any point $Z \in (1-\Delta)K$, there is a ball $B(Z, \Delta r)$ centered at Z of radius Δr contained in K . To show this fact, it is sufficient to consider the convex hull \mathcal{C} of $B(0, r) \cup \left\{ \frac{1}{1-\Delta}Z \right\} \subseteq K$, and show that it contains the ball $B(Z, \Delta r)$. Towards this end, we make the following geometric construction: we let p be a point such that the line $p\hat{\theta}$ is tangent to $B(0, r)$, and q the point on $p\hat{\theta}$ which minimizes the distance $\|q - Z\|_2$ (see Figure 1). Since $\angle op\hat{\theta}$ and $\angle Zq\hat{\theta}$ are both right angles, we have that the triangles $Op\hat{\theta}$ and $Zq\hat{\theta}$ are similar triangles, and hence that $\frac{\|Z-q\|_2}{r} = \frac{\|Z-\hat{\theta}\|_2}{\|\hat{\theta}-0\|_2}$. In other words,

$$\|Z - q\|_2 = \frac{\|Z - \hat{\theta}\|_2}{\|\hat{\theta} - 0\|_2} \times r = \frac{\left\| Z - \frac{1}{1-\Delta}Z \right\|_2}{\left\| \frac{1}{1-\Delta}Z \right\|_2} \times r = \Delta \times r,$$

implying a ball of radius Δr centered at Z is contained in $\mathcal{C} \subseteq K$, and hence that

$$(1-\Delta)K \subseteq \text{int}_{\Delta r}(K).$$

Bounding the infinity distance error. To complete the bound on the infinity-distance of the distribution $\hat{\nu}$ of the point returned by Algorithm 1 to the target distribution π , we must show both a lower bound (Lemma B.5) and an upper bound (Lemma B.6) on the ratio $\frac{\hat{\nu}(\theta)}{\pi(\theta)}$ at every point $\theta \in K$. Both the upper and lower bounds are necessary to bound the infinity-distance $d_\infty(\hat{\nu}(\theta), \pi(\theta)) = \sup_{\theta \in K} \left| \log \frac{\hat{\nu}(\theta)}{\pi(\theta)} \right|$.

Both Lemmas B.5 and B.6 require the input point to have TV error $\delta < \varepsilon \left(\frac{R}{\Delta r}\right)^{-d} e^{-LR}$. The term $\left(\frac{R}{\Delta r}\right)^{-d}$ is a lower bound on the ratio of the volume of K to the volume of the smoothing ball $B(0, \Delta r)$; this bound holds since K is contained in a ball of radius R . The term e^{-LR} is a lower bound on the ratio $\frac{\min_{w \in K} \pi(w)}{\max_{w \in K} \pi(w)}$ of the minimum value of the density π to the maximum value of π at any two points in K ; this bound holds since f is L -Lipschitz.

The above choice of δ ensures that in any ball $B(z, \Delta r)$ with center z in the Δr -interior of K , the distribution μ of the input point, which satisfies $\|\mu - \pi\|_{\text{TV}} \leq \delta$, will have between $e^{-\varepsilon}$ and e^ε times the probability mass which the target distribution π has inside the ball $B(z, \Delta r)$. Thus, when the distribution μ is smoothed by adding noise uniformly distributed on a ball of radius Δr , the smoothed distribution $\hat{\nu}(\theta)$ is within $e^{-\varepsilon}$ and e^ε times the target probability density $\pi(\theta)$ at any point θ in the Δr -interior of K , allowing us to bound the infinity distance error of the smoothed distribution $\hat{\nu}$ at any point θ in the Δr -interior of K . We then apply this fact, together with Lemma

B.1 which says that for any point $\theta \in K$ the point $(1 - \Delta)\theta$ is in the Δr -interior of K , to bound the distribution ν of the output point (after the stretching operation) as follows,

$$\nu(\theta) \geq (1 - \Delta)^d \hat{\nu}((1 - \Delta)\theta) \stackrel{\text{Lemma B.1}}{\geq} (1 - \Delta)^d \pi((1 - \Delta)\theta) \times e^{-\varepsilon} \geq \pi(\theta) e^{-\frac{\varepsilon}{2}}. \quad (2)$$

Here our choice of hyperparameter $\Delta \leq \frac{\varepsilon}{\max(d, LR)}$ ensures that $(1 - \Delta)^d = \Omega(1)$ and, since f is L -Lipschitz, that $\pi((1 - \Delta)\theta) \geq e^{-\varepsilon} \pi(\theta)$. This proves the lower bound (Lemma **B.5**). The proof of the upper bound (Lemma **B.6**) follows in a similar way as equation (2) but with the inequalities going in the opposite direction.

Bounding the number of iterations and concluding the proof of Theorem 2.2 We still need to deal with the problem that the point $\hat{\theta}$ may not be accepted. If this occurs, roughly speaking, we repeat the above procedure until a point $\hat{\theta}$ with distribution $\hat{\nu}$ is accepted. To bound the number of iterations, we show that $\hat{\theta}$ is in K with high probability. Towards this end, we first use the facts that f is L -Lipschitz and $K \subseteq B(0, R)$, to show that the probability a point sampled from $\pi \propto e^{-f}$ lies inside $(1 - \Delta)K$ is at least $(1 - \Delta)^d e^{-L\Delta R} \geq \frac{9}{10}$ (Lemma **B.2**). Lemma **B.2** says that if you stretch the polytope by a factor of $\frac{1}{1 - \Delta}$, then most of the volume of the stretched polytope $(\frac{1}{1 - \Delta})K$ remains inside the original polytope K . The term $(1 - \Delta)^d$ is just the ratio of the volume of $(1 - \Delta)K$ to the volume of K . And, since f is L -Lipschitz, the term $e^{-L\Delta R}$ bounds the ratio $\frac{\pi(\theta)}{\pi(\frac{1}{1 - \Delta}\theta)}$ of the target density at any point $\theta \in K$ to the value of π at the point $\frac{1}{1 - \Delta}\theta$ to which the stretching operation transports θ , whenever $\frac{1}{1 - \Delta}\theta \in K$. The choice of hyperparameter $\Delta \leq \frac{\varepsilon}{\max(d, LR)}$ ensures that the acceptance probability $(1 - \Delta)^d e^{-L\Delta R}$ guaranteed by Lemma **B.2** is at least $\frac{9}{10}$. Since the convex body $(1 - \Delta)K$ contains the ball $B(0, \frac{r}{2})$, applying Lemma **B.1** a second time (this time to the convex body $(1 - \Delta)K$) we get that

$$(1 - 3\Delta)K \subseteq \text{int}_{\Delta r}((1 - \Delta)K).$$

Thus, by Lemma **B.2** we have that θ lies inside $\text{int}_{\Delta r}((1 - \Delta)K)$ with probability at least $\frac{9}{10} - \delta \geq \frac{8}{10}$ (as θ is sampled from π with TV error $\leq \delta$). Therefore, since $\xi \sim B(0, \Delta r)$, we must also have that the probability that the point $\hat{\theta} = \frac{1}{1 - \Delta}(\theta + \xi)$ is in K (and is therefore not rejected) is greater than $\frac{8}{10}$ ⁴. This implies that the number of iterations until our algorithm returns a point $\hat{\theta}$ is less than $k > 0$ with probability at least $1 - 2^{-k}$, and the expected number of iterations is at most 2 (proved in Corollary **B.4**).

Since each iteration requires one random sample θ from the distribution μ , and one call to a membership oracle for K (to determine if $\frac{1}{1 - \Delta}Z \in K$), the number of sampling oracle and membership oracle calls required by Algorithm **1** is $O(1)$ with very high probability. Therefore, with high probability, Algorithm **1** returns a point $\hat{\theta}$ from a distribution with infinity-distance at most ε from π after $O(1)$ calls to the sampling and membership oracles.

Since Algorithm **1** succeeds with probability $1 - 2^{-k}$ after k iterations, after

$$\tau_{\max} = 5d \log\left(\frac{R}{r}\right) + 5LR + \varepsilon$$

iterations Algorithm **1** will have succeeded with probability roughly $1 - \varepsilon \left(\frac{R}{r}\right)^{-5d} e^{-5LR}$. In the very unlikely event that Algorithm **1** still has not succeeded after τ_{\max} iterations, Algorithm **1** simply outputs a point sampled from the uniform distribution on the ball $B(0, r)$ of radius r contained in K . The probability mass of the target distribution π inside this ball is at least as large as $\left(\frac{R}{r}\right)^{-d} e^{-LR}$; thus, since f is L -Lipschitz, we show in Corollary **B.4** that outputting a sample from the uniform distribution on this ball with probability $\varepsilon \left(\frac{R}{r}\right)^{-5d} e^{-5LR}$ does not change the ∞ -distance error of the sample returned by the algorithm by more than ε .

⁴In Algorithm **1** we reject $\hat{\theta}$ with a slightly higher probability to ensure that, in differential privacy applications, in addition to the privacy of the point returned by the algorithm, the runtime is also ε -differentially private.

3.2 Completing the proof of Theorem 2.1

Proof: By Theorem 1 the output of Algorithm 1 has infinity-distance error bounded by ε as long as the input samples have TV distance error bounded by

$$\delta \leq O\left(\varepsilon \times \left(\frac{R(d \log(R/r) + LR)^2}{\varepsilon r}\right)^{-d} e^{-LR}\right),$$

and, with high probability, Algorithm 1 requires $O(1)$ such independent samples. To generate a sample from π with TV error $O(\delta)$ when $K = \{\theta \in \mathbb{R}^d : A\theta \leq b\}$ is a polytope defined by m inequalities, we use the Dikin Walk Markov chain of [34]. This Markov chain requires an initial point from some distribution μ_0 which is w -warm with respect to the stationary distribution π , that is, $\sup_{z \in K} \frac{\mu_0(z)}{\pi(z)} \leq w$. To obtain a warm start, we let μ_0 be the uniform distribution on the ball with radius r contained in K and sample from μ_0 . Since f is L -Lipschitz, and K is contained in a ball of radius R , μ_0 is w -warm with

$$w \leq \frac{1}{\text{Vol}(B(0, r))} \times \left(\frac{\max_{z \in K} \pi(\theta)}{\min_{z \in K} \pi(\theta)} \times \text{Vol}(B(0, R))\right) \leq \left(\frac{R}{r}\right)^d \times e^{RL}.$$

From [34], we have that from this w -warm start the Dikin Walk Markov chain requires at most $O((m^2 d^4 + m^2 d^2 L^2 R^2) \log(\frac{w}{\delta}))$ steps to generate a sample with TV distance at most δ from π , where each step makes one function evaluation and $O(md^{\omega-1})$ arithmetic operations. Plugging in the above values of δ, w the number of Markov chain steps is

$$T = O((m^2 d^3 + m^2 d L^2 R^2) \times [LR + d \log(\frac{Rd + LRd}{r\varepsilon})])$$

to generate each independent sample with the required TV error $O(\delta)$. Since the number of independent samples required as input for Algorithm 1 is $O(1)$ w.h.p., the number of arithmetic operations for Algorithm 1 to output a point with at most ε infinity-distance error is $O(T \times md^{\omega-1})$.

Finally, we note that in the more general setting where K is a convex body with membership oracle (but not necessarily) a polytope, we can instead use, for instance, the hit-and-run Markov chain of [29] to generate samples from π with TV error $O(\delta)$ in a number of membership and function evaluation oracle calls that is polynomial in d and poly-logarithmic in $\frac{1}{\delta}, R, r$. We can then plug this sample into our Algorithm 1 to obtain a sample from π with infinity-distance error $O(\varepsilon)$ in a number of oracle calls that is (poly)-logarithmic in $\frac{1}{\varepsilon}, \frac{1}{r}$ and polynomial on d, L, R . (see Remark 2.3). ■

4 Conclusion, Limitations, and Future Work

To the best of our knowledge, this is the first work that presents an algorithm for sampling from log-concave distributions on convex bodies that comes with infinity-distance bounds and whose running time depends logarithmically on $1/\varepsilon$. Towards this, the main technical contribution is Algorithm 1 (and Theorem 2.2) which achieves this improved dependence on ε by taking as input continuous samples from a convex body with TV bounds and converting them to samples with infinity-distance bounds.

On the other hand, our bounds are polynomial in LR , yet there are algorithms for sampling from logconcave distributions $\pi \propto e^{-f}$ on a convex body in the total variation distance that are poly-logarithmic in R and do not assume f to be Lipschitz [29]. Thus, the main open problem that remains is whether one can also obtain running time bounds for sampling in the infinity-distance which are poly-logarithmic in R and do not require f to be Lipschitz.

Our main result also has direct applications to differentially private optimization (Corollaries 2.4 and 2.5). Differential privacy is a notion which has been embraced in many technologies in societal contexts where privacy of individuals is a concern. Hence, we see our work to have a potential of positive societal impact and do not foresee any potential negative societal impacts.

Acknowledgments and Disclosure of Funding

This research was supported in part by NSF CCF-2104528, CCF-1908347, and CCF-2112665 awards.

References

- [1] David Applegate and Ravi Kannan. Sampling and integration of near log-concave functions. In *Proceedings of the twenty-third annual ACM symposium on Theory of computing*, pages 156–163, 1991.
- [2] Raef Bassily, Adam Smith, and Abhradeep Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 464–473. IEEE, 2014.
- [3] Avrim Blum, Cynthia Dwork, Frank McSherry, and Kobbi Nissim. Practical privacy: the SuLQ framework. In *Proceedings of the twenty-fourth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 128–138, 2005.
- [4] Nicolas Brosse, Alain Durmus, Éric Moulines, and Marcelo Pereyra. Sampling from a log-concave distribution with compact support with proximal Langevin Monte Carlo. In *Conference on Learning Theory*, pages 319–342. PMLR, 2017.
- [5] Sebastien Bubeck, Ronen Eldan, and Joseph Lehec. Finite-time analysis of projected Langevin Monte Carlo. In *Advances in Neural Information Processing Systems*, pages 1243–1251. Cite-seer, 2015.
- [6] Olivier Chapelle and Lihong Li. An empirical evaluation of Thompson sampling. *Advances in neural information processing systems*, 24:2249–2257, 2011.
- [7] Kamalika Chaudhuri, Anand Sarwate, and Kaushik Sinha. Near-optimal differentially private principal components. *Advances in neural information processing systems*, 25:989–997, 2012.
- [8] Yuansi Chen, Raaz Dwivedi, Martin J Wainwright, and Bin Yu. Vaidya walk: A sampling algorithm based on the volumetric barrier. In *2017 55th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1220–1227. IEEE, 2017.
- [9] Xiang Cheng and Peter Bartlett. Convergence of Langevin MCMC in KL-divergence. In *Algorithmic Learning Theory*, pages 186–211. PMLR, 2018.
- [10] 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(3):651–676, 2017.
- [11] Arnak S Dalalyan and Lionel Riou-Durand. On sampling from a log-concave density using kinetic Langevin diffusions. *Bernoulli*, 26(3):1956–1988, 2020.
- [12] Alain Durmus, Szymon Majewski, and Blazej Miasojedow. Analysis of Langevin Monte Carlo via convex optimization. *J. Mach. Learn. Res.*, 20:73–1, 2019.
- [13] Alain Durmus and Eric Moulines. Nonasymptotic convergence analysis for the unadjusted Langevin algorithm. *Annals of Applied Probability*, 27(3):1551–1587, 2017.
- [14] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–407, 2014.
- [15] Cynthia Dwork, Kunal Talwar, Abhradeep Thakurta, and Li Zhang. Analyze gauss: optimal bounds for privacy-preserving principal component analysis. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 11–20, 2014.
- [16] M. E. Dyer and A. M. Frieze. On the complexity of computing the volume of a polyhedron. *SIAM Journal on Computing*, 17(5):967–974, 1988.
- [17] Alan Frieze and Ravi Kannan. Log-Sobolev inequalities and sampling from log-concave distributions. *The Annals of Applied Probability*, 9(1):14–26, 1999.
- [18] Alan Frieze, Ravi Kannan, and Nick Polson. Sampling from log-concave distributions. *The Annals of Applied Probability*, pages 812–837, 1994.

- [19] Arun Ganesh and Kunal Talwar. Faster differentially private samplers via Rényi divergence analysis of discretized Langevin MCMC. *Advances in Neural Information Processing Systems*, 33, 2020.
- [20] Moritz Hardt and Kunal Talwar. On the geometry of differential privacy. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 705–714, 2010.
- [21] He Jia, Aditi Laddha, Yin Tat Lee, and Santosh S. Vempala. Reducing isotropy and volume to KLS: an $O^*(n^3\psi^2)$ volume algorithm. In *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 961–974. ACM, 2021.
- [22] Richard Jordan, David Kinderlehrer, and Felix Otto. The variational formulation of the Fokker–Planck equation. *SIAM journal on mathematical analysis*, 29(1):1–17, 1998.
- [23] Ravi Kannan and Hariharan Narayanan. Random walks on polytopes and an affine interior point method for linear programming. *Mathematics of Operations Research*, 37(1):1–20, 2012.
- [24] Jonathan Leake, Colin S McSwiggen, and Nisheeth K Vishnoi. A polynomial-time algorithm and applications for matrix sampling from Harish-Chandra–Itzykson–Zuber densities. In *ACM symposium on theory of computing STOC*, 2021.
- [25] Yin Tat Lee and Santosh Vempala. Eldan’s stochastic localization and the KLS hyperplane conjecture: An improved lower bound for expansion. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 998–1007. IEEE, 2017.
- [26] Yin Tat Lee and Santosh S Vempala. Convergence rate of Riemannian Hamiltonian Monte Carlo and faster polytope volume computation. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1115–1121, 2018.
- [27] László Lovász and Miklós Simonovits. Random walks in a convex body and an improved volume algorithm. *Random structures & algorithms*, 4(4):359–412, 1993.
- [28] László Lovász and Santosh Vempala. Hit-and-run is fast and fun. *preprint, Microsoft Research*, 2003.
- [29] László Lovász and Santosh Vempala. Fast algorithms for logconcave functions: Sampling, rounding, integration and optimization. In *Foundations of Computer Science, 2006. FOCS'06. 47th Annual IEEE Symposium on*, pages 57–68. IEEE, 2006.
- [30] László Lovász and Santosh Vempala. The geometry of logconcave functions and sampling algorithms. *Random Structures & Algorithms*, 30(3):307–358, 2007.
- [31] Oren Mangoubi and Aaron Smith. Mixing of Hamiltonian Monte Carlo on strongly log-concave distributions: Continuous dynamics. *The Annals of Applied Probability*, in press.
- [32] Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*, pages 94–103. IEEE, 2007.
- [33] Hariharan Narayanan. Randomized interior point methods for sampling and optimization. *The Annals of Applied Probability*, 26(1):597–641, 2016.
- [34] Hariharan Narayanan and Alexander Rakhlin. Efficient sampling from time-varying log-concave distributions. *The Journal of Machine Learning Research*, 18(1):4017–4045, 2017.
- [35] Sushant Sachdeva and Nisheeth K Vishnoi. The mixing time of the Dikin walk in a polytope—a simple proof. *Operations Research Letters*, 44(5):630–634, 2016.
- [36] Santosh Vempala and Andre Wibisono. Rapid convergence of the unadjusted Langevin algorithm: Isoperimetry suffices. In *Advances in Neural Information Processing Systems*, 2019.
- [37] Cédric Villani. *Optimal transport: old and new*, volume 338. Springer, 2009.
- [38] Nisheeth K. Vishnoi. An introduction to Hamiltonian Monte Carlo method for sampling. *CoRR*, abs/2108.12107, 2021.

- [39] Max Welling and Yee W Teh. Bayesian learning via stochastic gradient Langevin dynamics. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 681–688. Citeseer, 2011.
- [40] Andre Wibisono. Sampling as optimization in the space of measures: The Langevin dynamics as a composite optimization problem. In *Conference on Learning Theory*, pages 2093–3027. PMLR, 2018.

Checklist

The checklist follows the references. Please read the checklist guidelines carefully for information on how to answer these questions. For each question, change the default **[TODO]** to **[Yes]**, **[No]**, or **[N/A]**. You are strongly encouraged to include a **justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license to the code and datasets? [N/A]
- Did you include the license to the code and datasets? [N/A]
- Did you include the license to the code and datasets? [N/A]

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes]
 - (c) Did you discuss any potential negative societal impacts of your work? [Yes] (See Section 4)
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
 - (b) Did you include complete proofs of all theoretical results? [Yes] (See Section 3.2, Appendix B and C)
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See supplementary material.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [N/A]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]