

Escaping the Local Minima via Simulated Annealing: Optimization of Approximately Convex Functions

Alexandre Belloni

The Fuqua School of Business, Duke University

Tengyuan Liang

Department of Statistics, The Wharton School, University of Pennsylvania

Hariharan Narayanan

Department of Statistics and Department of Mathematics, University of Washington

Alexander Rakhlin

Department of Statistics, The Wharton School, University of Pennsylvania

Abstract

We consider the problem of optimizing an approximately convex function over a bounded convex set in \mathbb{R}^n using only function evaluations. The problem is reduced to sampling from an *approximately* log-concave distribution using the Hit-and-Run method, which is shown to have the same \mathcal{O}^* complexity as sampling from log-concave distributions. In addition to extend the analysis for log-concave distributions to approximate log-concave distributions, the implementation of the 1-dimensional sampler of the Hit-and-Run walk requires new methods and analysis. The algorithm then is based on simulated annealing which does not relies on first order conditions which makes it essentially immune to local minima.

We then apply the method to different motivating problems. In the context of zeroth order stochastic convex optimization, the proposed method produces an ϵ -minimizer after $\mathcal{O}^*(n^{7.5}\epsilon^{-2})$ noisy function evaluations by inducing a $\mathcal{O}(\epsilon/n)$ -approximately log concave distribution. We also consider in detail the case when the “amount of non-convexity” decays towards the optimum of the function. Other applications of the method discussed in this work include private computation of empirical risk minimizers, two-stage stochastic programming, and approximate dynamic programming for online learning.

1. Introduction and Problem Formulation

Let $\mathcal{K} \subset \mathbb{R}^n$ be a convex set, and let $F : \mathbb{R}^n \rightarrow \mathbb{R}$ be an approximately convex function over \mathcal{K} in the sense that

$$\sup_{x \in \mathcal{K}} |F(x) - f(x)| \leq \epsilon/n \tag{1}$$

for some convex function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ and $\epsilon > 0$. In particular, F may be discontinuous. We seek to find $x \in \mathcal{K}$ such that

$$F(x) - \min_{y \in \mathcal{K}} F(y) \leq \epsilon \tag{2}$$

using only function evaluations of F . This paper presents a randomized method based on simulated annealing that satisfies (2) in expectation (or with high probability). Moreover, the number of

required function evaluations of F is at most $\mathcal{O}^*(n^{4.5})$ (see Corollary 18), where \mathcal{O}^* hides poly-logarithmic factors in n and ϵ^{-1} . Our method requires only a membership oracle for the set \mathcal{K} . In Section 7, we consider the case when the amount of non-convexity in (1) can be much larger than ϵ/n for points away from the optimum.

In the oracle model of computation, access to function values at queried points is referred to as the zeroth-order information. Exact function evaluation of F may be equivalently viewed as approximate function evaluation of the convex function f , with the oracle returning a value

$$F(x) \in [f(x) - \epsilon/n, f(x) + \epsilon/n]. \quad (3)$$

A closely related problem is that of convex optimization with a *stochastic* zeroth order oracle. Here, the oracle returns a noisy function value $f(x) + \eta$. If η is zero-mean and subgaussian, the function values can be averaged to emulate, with high probability, the approximate oracle (3). The randomized method we propose has an $\mathcal{O}^*(n^{7.5}\epsilon^{-2})$ oracle complexity for convex stochastic zeroth order optimization, which, to the best of our knowledge, is the best that is known for this problem. We refer to Section 6 for more details.

The motivation for studying zeroth-order optimization is plentiful, and we refer the reader to Conn et al. (2009) for a discussion of problems where derivative-free methods are essential. In Section 8 we sketch three areas where the algorithm of this paper can be readily applied: private computation with distributed data, two-stage stochastic programming, and online learning algorithms.

2. Prior Work

The present paper rests firmly on the long string of work by Kannan, Lovász, Vempala, and others (Lovász and Simonovits, 1993; Kannan et al., 1997; Kalai and Vempala, 2006; Lovász and Vempala, 2006a,b, 2007). In particular, we invoke the key lower bound on conductance of Hit-and-Run from Lovász and Vempala (2006a) and use the simulated annealing technique of Kalai and Vempala (2006). Our analysis extends Hit-and-Run to approximately log-concave distributions which required new theoretical results and implementation adjustments. In particular, we propose a unidimensional sampling scheme that mixes fast to a truncated approximately log-concave distribution on the line.

Sampling from β -log-concave distributions was already studied in the early work of Applegate and Kannan (1991) with a discrete random walk based on a discretization of the space. In the case of non-smooth densities and unrestricted support, sampling from approximate log-concave distributions has also been studied in Belloni and Chernozhukov (2009) where the hidden convex function f is quadratic. This additional structure was motivated by the central limit theorem in statistical applications and leads to faster mixing rates. Both works used ball walk-like strategies. Neither work considered random walks that allow for long steps like Hit-and-Run.

The present work was motivated by the question of information-based complexity of zeroth-order stochastic optimization. The paper of Agarwal et al. (2013) studies a somewhat harder problem of regret minimization with zeroth-order feedback. Their method is based on the pyramid construction of Nemirovskii and Yudin (1983) and requires $\mathcal{O}(n^{33}\epsilon^{-2})$ noisy function evaluations to achieve a regret (and, hence, an optimization guarantee) of ϵ . The method of Liang et al. (2014) improved the dependence on the dimension to $\mathcal{O}^*(n^{14})$ using a Ball Walk on the epigraph of the function in the spirit of Bertsimas and Vempala (2004). The present paper further reduces this

dependence to $\mathcal{O}^*(n^{7.5})$ and still achieves the optimal ϵ^{-2} dependence on the accuracy. The best known lower bound for the problem is $\Omega(n^2\epsilon^{-2})$ (see [Shamir \(2012\)](#)).

Other relevant work includes the recent paper of [Dyer et al. \(2013\)](#) where the authors proposed a simple random walk method that requires only approximate function evaluations. As the authors mention, their algorithm only works for smooth functions and sets \mathcal{K} with smooth boundaries — assumptions that we would like to avoid. Furthermore, the effective dependence of [Dyer et al. \(2013\)](#) on accuracy is worse than ϵ^{-2} .

3. Preliminaries

Throughout the paper, the functions F and f satisfy (1) and f is convex. The Lipschitz constant of f with respect to ℓ_∞ norm will be denoted by L , defined as the smallest number such that $|f(x) - f(y)| \leq L\|x - y\|_\infty$ for $x, y \in \mathcal{K}$. Assume the convex body $\mathcal{K} \subseteq \mathbb{R}^n$ to be well-rounded in the sense that there exist $r, R > 0$ such that $\mathcal{B}_2^n(r) \subseteq \mathcal{K} \subseteq \mathcal{B}_2^n(R)$ and $R/r \leq \mathcal{O}(\sqrt{n})$.¹ For a non-negative function g , denote by π_g the normalized probability measure induced by g and supported on \mathcal{K} .

Definition 1 A function $h : \mathcal{K} \rightarrow \mathbb{R}_+$ is *log-concave* if

$$h(\alpha x + (1 - \alpha)y) \geq h(x)^\alpha h(y)^{1-\alpha}$$

for all $x, y \in \mathcal{K}$ and $\alpha \in [0, 1]$. A function is called *β -log-concave* for some $\beta \geq 0$ if

$$h(\alpha x + (1 - \alpha)y) \geq e^{-\beta} h(x)^\alpha h(y)^{1-\alpha}$$

for all $x, y \in \mathcal{K}$ and $\alpha \in [0, 1]$.

Definition 2 A function $g : \mathcal{K} \rightarrow \mathbb{R}_+$ is *ξ -approximately log-concave* if there is a log-concave function $h : \mathcal{K} \rightarrow \mathbb{R}_+$ such that

$$\sup_{x \in \mathcal{K}} |\log h(x) - \log g(x)| \leq \xi.$$

Lemma 3 If the function g is $\beta/2$ -approximately log-concave, then g is β -log-concave.

For one-dimensional functions, the above lemma can be reversed:

Lemma 4 (Belloni and Chernozhukov (2009), Lemma 9) If g is a unidimensional β -log-concave function, then there exists a log-concave function h such that

$$e^{-\beta} h(x) \leq g(x) \leq h(x) \text{ for all } x \in \mathbb{R}.$$

Remark 5 (Gap Between β -Log-Concave Functions and ξ -Approximate Log-Concave Functions)

A consequence of Lemma 4 is that β -log-concave functions are equivalent to β -approximately log-concave functions when the domain is unidimensional. However, such equivalence no longer holds in higher dimensions. In the case the domain is \mathbb{R}^n , [Green et al. \(1952\)](#); [Cholewa \(1984\)](#) established that β -log-concave functions are $\frac{\beta}{2} \log_2(2n)$ -approximately log-concave. [Laczkovich \(1999\)](#) showed that there are functions such that the factor that relates these approximations cannot be less than $\frac{1}{4} \log_2(n/2)$.

1. This condition can be relaxed by applying a pencil construction as in [Lovász and Vempala \(2007\)](#).

4. Sampling from Approximate Log-Concave Distributions via Hit-and-Run

In this section we analyze the Hit-and-Run procedure to simulate random variables from a distribution induced by an approximate log-concave function. The Hit-and-Run algorithm is as follows in Algorithm 1.

Algorithm 1 Hit-and-Run

Input: a target distribution π_g on \mathcal{K} induced by a nonnegative function g ; $x \in \text{dom}(g)$; linear transformation Σ ; number of steps m

Output: a point $x' \in \text{dom}(g)$ generated by one-step Hit-and-Run walk

initialization: a starting point $x \in \text{dom}(g)$

for $i = 1, \dots, m$ **do**

1. Choose a random line ℓ that passes through x . The direction is uniform from the surface of ellipse given by Σ acting on sphere
2. On the line ℓ run the unidimensional rejection sampler with π_g restricted to the line (and supported on \mathcal{K}) to propose a successful next step x'

end

In order to handle approximate log-concave functions we need to address implementation issues and address the theoretical difficulties caused by deviations from log-concavity which can include discontinuities. The main implementation difference lies in the unidimensional sampler. No longer a binary search yields the maximum over the line and its end points since β -log-concave functions can be discontinuous and multimodal. We now turn to these questions.

4.1. Unidimensional sampling scheme

As a building block of the randomized method for solving the optimization problem (2), we introduce a one-dimensional sampling procedure. Let g be a unidimensional β -log-concave function on a bounded line segment ℓ , and let π_g be the induced normalized measure. The following guarantee will be proved in this section.

Lemma 6 *Let g be a β -log-concave function and let ℓ be a bounded line segment ℓ on \mathcal{K} . Given a target accuracy $\tilde{\epsilon} \in (0, e^{-2\beta}/2)$, Algorithm 2 produces a point $X \in \ell$ with a distribution $\hat{\pi}_{g,\ell}$ such that*

$$d_{\text{tv}}(\pi_{g,\ell}, \hat{\pi}_{g,\ell}) \leq 3e^{2\beta}\tilde{\epsilon}.$$

Moreover, the method requires $\mathcal{O}^(1)$ evaluations of the unidimensional β -log-concave function g if β is $\mathcal{O}(1)$.*

The proposed method for sampling from the β -log-concave function g is a *rejection* sampler that requires two initialization steps. We first show how to implement step (a).

For the β -log-concave function g , let h be a log-concave function in Lemma 4 and let \tilde{L} denote the Lipschitz constant of the convex function $-\log h$. In the following two results, the \mathcal{O}^* notation hides a $\log(\tilde{L})$ factor.

Lemma 7 (Initialization Step (a)) *Algorithm 3 finds a point $p \in \ell$ that satisfies $g(p) \geq e^{-3\beta} \max_{z \in \ell} g(z)$. Moreover, this step requires $\mathcal{O}^*(1)$ function evaluations.*

Algorithm 2 Unidimensional rejection sampler

Input: 1-dim β -log-concave function g defined on a bounded segment $\ell = [\underline{x}, \bar{x}]$; accuracy $\tilde{\epsilon} > 0$

Output: A sample x with distribution $\hat{\pi}_{g,\ell}$ close to $\pi_{g,\ell}$

Initialization: **(a)** compute a point $p \in \ell$ s.t. $g(p) \geq e^{-3\beta} \max_{z \in \ell} g(z)$ **(b)** given target accuracy $\tilde{\epsilon}$, find two points e_{-1}, e_1 on two sides of p s.t.

$$\begin{aligned} e_{-1} &= \underline{x} \text{ if } g(\underline{x}) \geq \frac{1}{2}e^{-\beta}\tilde{\epsilon}g(p), \quad \frac{1}{2}e^{-\beta}\tilde{\epsilon}g(p) \leq g(e_{-1}) \leq \tilde{\epsilon}g(p) \text{ otherwise} \\ e_1 &= \bar{x} \text{ if } g(\bar{x}) \geq \frac{1}{2}e^{-\beta}\tilde{\epsilon}g(p), \quad \frac{1}{2}e^{-\beta}\tilde{\epsilon}g(p) \leq g(e_1) \leq \tilde{\epsilon}g(p) \text{ otherwise;} \end{aligned} \quad (4)$$

while *sample rejected* **do**

pick $x \sim \text{unif}([e_{-1}, e_1])$ and pick $r \sim \text{unif}([0, 1])$ independently. If $r \leq g(x)/\{g(p)e^{3\beta}\}$ then accept x and stop. Otherwise, reject x

end

Algorithm 3 Initialization Step (a)

Input: unidimensional β -log-concave function g defined on a bounded interval $\ell = [\underline{x}, \bar{x}]$

Output: a point $p \in \ell$ s.t. $g(p) \geq e^{-3\beta} \max_{z \in \ell} g(z)$

while *did not stop* **do**

set $x_l = \frac{3}{4}\underline{x} + \frac{1}{4}\bar{x}$, $x_c = \frac{1}{2}\underline{x} + \frac{1}{2}\bar{x}$ and $x_r = \frac{1}{4}\underline{x} + \frac{3}{4}\bar{x}$
 If $|\log g(x_l) - \log g(x_r)| > \beta$, set $[\underline{x}, \bar{x}]$ as either $[x_l, \bar{x}]$ or $[\underline{x}, x_r]$ accordingly
 else if $|\log g(x_l) - \log g(x_c)| > \beta$, set $[\underline{x}, \bar{x}]$ as either $[x_l, \bar{x}]$ or $[\underline{x}, x_c]$ accordingly
 else if $|\log g(x_r) - \log g(x_c)| > \beta$, set $[\underline{x}, \bar{x}]$ as either $[\underline{x}, x_r]$ or $[x_c, \bar{x}]$ accordingly
 else output $p = \arg \max_{x \in \{x_l, x_c, x_r\}} f(x)$ and stop

end

Lemma 8 (Initialization Step (b)) *Let $\ell = [\underline{x}, \bar{x}]$ and $p \in \ell$. The binary search algorithm finds $e_{-1} \in [\underline{x}, p]$ and $e_1 \in [p, \bar{x}]$ such that (4) holds. Moreover, this step requires $\mathcal{O}^*(1)$ function evaluations.*

According to Lemmas 6, 7, 8, the unidimensional sampling method produces a sample from a distribution that is close to the desired β -log-concave distribution. Furthermore, the method requires a number of queries that is logarithmic in all the parameters.

4.2. Mixing time

In this section, we will analyze mixing time of the Hit-and-Run algorithm with a $\beta/2$ -approximate log-concave function g , namely there exists log-concave h such that $\sup_{\mathcal{X}} |\log g - \log h| \leq \beta/2$. In particular, this implies that g is β -log-concave, according to Lemma 3. In this section, we provide the analysis of Hit-and-Run with the linear transformation $\Sigma = I$ and remark that the results extend to other linear transformations employed to round the log-concave distributions.

The mixing time of a geometric random walk can be bounded through the spectral gap of the induced Markov chain. In turn, the spectral gap relates to the so called conductance which has been a key quantity in the literature. Consider the transition probability of Hit-and-Run with a density g ,

namely

$$P_u^g(A) = \frac{2}{n\pi_n} \int_A \frac{g(x)dx}{\mu_g(u, x)|x - u|^{n-1}}$$

where $\mu_g(u, x) = \int_{\ell(u, x) \cap \mathcal{K}} g(y)dy$. Let $\pi_g(x) = \frac{g(x)}{\int_{y \in \mathcal{K}} g(y)dy}$ be the probability measure induced by the function g . The conductance for a set $S \subset \mathcal{K}$ with $0 < \pi_g(S) < 1$ is defined as

$$\phi^g(S) = \frac{\int_{x \in S} P_x^g(\mathcal{K} \setminus S) d\pi_g}{\min\{\pi_g(S), \pi_g(\mathcal{K} \setminus S)\}},$$

and ϕ^g is the minimum conductance over all measurable sets. The s -conductance is, in turn, defined as

$$\phi_s^g = \inf_{S \subset \mathcal{K}, s < \pi_g(S) \leq 1/2} \frac{\int_{x \in S} P_x^g(\mathcal{K} \setminus S) d\pi_g}{\pi_g(S) - s}.$$

By definition we have $\phi^g \leq \phi_s^g$ for all $s > 0$.

The following theorem provides us an upper bound on the mixing time of the Markov chain based on conductance. Let $\sigma^{(0)}$ be the initial distribution and $\sigma^{(m)}$ the distribution of the m -th step of the random walk of Hit-and-Run with exact sampling from the distribution π_g restricted to the line.

Theorem 9 (Lovász and Simonovits (1993); Lovász and Vempala (2007), Lemma 9.1) *Let $0 < s \leq 1/2$ and let $g : \mathcal{K} \rightarrow \mathbb{R}_+$ be arbitrary. Then for every $m \geq 0$,*

$$d_{\text{tv}}(\pi_g, \sigma^{(m)}) \leq H_0 \left(1 - \frac{(\phi^g)^2}{2}\right)^m \quad \text{and} \quad d_{\text{tv}}(\pi_g, \sigma^{(m)}) \leq H_s + \frac{H_s}{s} \left(1 - \frac{(\phi_s^g)^2}{2}\right)^m$$

where $H_0 = \sup_{x \in \mathcal{K}} \pi_g(x)/\sigma^0(x)$ and $H_s = \sup\{|\pi_g(A) - \sigma^{(0)}(A)| : \pi_g(A) \leq s\}$.

Building on Lovász and Vempala (2006a), we prove the following result that provides us with a lower bound on the conductance for Hit-and-Run induced by a log-concave h . The proof of the result below follows the proof of Theorem 3.7 in Lovász and Vempala (2006a) with modifications to allow unbounded sets \mathcal{K} without truncating the random walk.

Theorem 10 (Conductance Lower Bound for Log-concave Measures with Unbounded Support)

Let h be a log-concave function in \mathbb{R}^n such that the level set of measure $\frac{1}{8}$ contains a ball of radius r . Define $R = (\mathbb{E}_h \|X - z_h\|^2)^{1/2}$, where $z_h = \mathbb{E}_h X$ and X is sampled from the log concave measure induce by h . Then for any subset S , with $\pi_h(S) = p \leq \frac{1}{2}$, the conductance of Hit-and-Run satisfies

$$\phi^h(S) \geq \frac{1}{C_1 n \frac{R}{r} \log^2 \frac{nR}{rp}}$$

where $C_1 > 0$ is a universal constants.

Although Theorem 10 is new, very similar conductance bounds allowing for unbounded sets were establish before. Indeed in Section 3.3 of Lovász and Vempala (2006a) the authors discuss the case of unbounded \mathcal{K} and propose to truncate the set to its effective diameter and use the fact that this distribution would be close to the distribution of the unrestricted set. Such truncation needs

to be enforced which requires to change the implementation of the algorithm and lead to another (small) layer of approximation errors. Theorem 10 avoids this explicit truncation and truncation is done implicitly in the proof only. We note that when applying the simulated annealing technique, even if we start with a bounded set, by diminishing the temperature, we are effectively stretching the sets which would essentially require to handle unbounded sets.

We now argue that conductance of Hit-and-Run with β -approximate log-concave measures can be related to the conductance with log-concave measures.

Theorem 11 *Let g be a $\beta/2$ -approximate log-concave measure and h be any log-concave function with the property in Definition 2. Then the conductance and s -conductance of the random walk induced by g are lower bounded as*

$$\phi^g \geq e^{-3\beta} \phi^h \quad \text{and} \quad \phi_s^g \geq e^{-3\beta} \phi_{s/e^\beta}^h.$$

We apply Theorem 9 to show contraction of $\sigma^{(m)}$ to π_g in terms of the total variation distance.

Theorem 12 (Mixing Time for Approximately-log-concave Measure) *Let π_g is the stationary measure associated with the Hit-and-Run walk based on a $\beta/2$ -approximately log-concave function g , and $M = \|\sigma^{(0)}/\pi_g\| = \int (d\sigma^{(0)}/d\pi_g)d\sigma^{(0)}$. There is a universal constant $C < \infty$ such that for any $\gamma \in (0, 1/2)$, if*

$$m \geq Cn^2 \frac{e^{6\beta} R^2}{r^2} \log^4 \frac{e^\beta MnR}{r\gamma^2} \log \frac{M}{\gamma}$$

then m steps of the Hit-and-Run random walk based on g yield $d_{\text{tv}}(\pi_g, \sigma^{(m)}) \leq \gamma$.

Remark 13 *The value M in Theorem 12 bounds the impact of the initial distribution $\sigma^{(0)}$ which can be potentially far from the stationary distribution. In the Simulated Annealing application of next section, we will show in Lemma 15 that we can “warm start” the chain by carefully picking an initial distribution such that $M = \mathcal{O}(1)$.*

Theorem 12 shows γ -closeness between the distribution $\sigma^{(m)}$ and the corresponding stationary distribution. However, the stationary distribution is not exactly g since the unidimensional sampling procedure described earlier truncates the distribution to improve mixing time. The following theorem shows that these concerns are overtaken by the geometric mixing of the random walk. Let $\hat{\pi}_{g,\ell}$ denote the distribution of the unidimensional sampling scheme (Algorithm 2) along the line ℓ and $\pi_{g,\ell}$ denote the distribution of the unidimensional sampling scheme proportional to g along line ℓ .

Theorem 14 *Let $\hat{\sigma}^{(m)}$ denote the distribution of the Hit-and-Run with the unidimensional sampling scheme (Algorithm 2) after m steps. For any $0 < s < 1/2$, the algorithm maintains that*

$$d_{\text{tv}}(\hat{\sigma}^{(m)}, \pi_g) \leq 2d_{\text{tv}}(\hat{\sigma}^{(0)}, \sigma^{(0)}) + m \sup_{\ell \subset \mathcal{K}} d_{\text{tv}}(\hat{\pi}_{g,\ell}, \pi_{g,\ell}) + \left\{ H_s + \frac{H_s}{s} \left(1 - \frac{(\phi_s^g)^2}{2} \right)^m \right\}$$

where the supremum is taken over all lines ℓ in \mathcal{K} . In particular, for a target accuracy $\gamma \in (0, 1/e)$, if $d_{\text{tv}}(\hat{\sigma}^{(0)}, \sigma^{(0)}) \leq \gamma/8$, s such that $H_s \leq \gamma/4$, $m \geq \{2/(\phi_s^g)^2\} \log(\{H_s/s\}\{4/\gamma\})$, and the precision of the unidimensional sampling scheme to be $\tilde{\epsilon} = \gamma e^{-2\beta}/\{12m\}$, we have

$$d_{\text{tv}}(\hat{\sigma}^{(m)}, \pi_g) \leq \gamma.$$

5. Optimization via Simulated Annealing

We now turn to the main goal of the paper: to exhibit a method that produces an ϵ -minimizer of the nearly convex function F in expectation. Fix the pair f, F with the property (1), and define a series of functions

$$h_i(x) = \exp(-f/T_i), \quad g_i(x) = \exp(-F/T_i)$$

for a chain of temperatures $\{T_i, i = 1, \dots, K\}$ to be specified later. It is immediate that h_i 's are log-concave. Lemma 3, in turn, implies that g_i 's are $\frac{2\epsilon}{nT_i}$ -log-concave.

We now introduce the simulated annealing method that proceeds in epochs and employs the Hit-and-Run procedure with the unidimensional sampler introduced in the previous section. The overall simulated annealing procedure is identical to the algorithm of Kalai and Vempala (2006), with differences in the analysis arising from F being only approximately convex.

Algorithm 4 Simulated annealing

Input: A series of temperatures $\{T_i, 1 \leq i \leq K\}$, K =number of epochs, $x \in \text{int}\mathcal{K}$

Output: a candidate point x for which $F(x) \leq \min_{y \in \mathcal{K}} F(y) + \epsilon$ holds

initialization: well-rounded convex body \mathcal{K} and $\{X_0^j, 1 \leq j \leq N\}$ i.i.d. samples from uniform measure on \mathcal{K} , N -number of strands, set $\mathcal{K}_0 = \mathcal{K}$, and $\Sigma_0 = I$

while i -th epoch, $1 \leq i \leq K$ **do**

1. calculate the i -th rounding linear transformation \mathcal{T}_i based on $\{X_{i-1}^j, 1 \leq j \leq N\}$ and let $\Sigma_i = \mathcal{T}_i \circ \Sigma_{i-1}$
2. draw N i.i.d. samples $\{X_i^j, 1 \leq j \leq N\}$ from measure π_{g_i} using Hit-and-Run with linear transformation Σ_i and with N warm-starting points $\{X_{i-1}^j, 1 \leq j \leq N\}$

end

output $x = \underset{1 \leq j \leq N, 1 \leq i \leq K}{\text{argmin}} F(X_i^j)$.

Before stating the optimization guarantee of the above simulated annealing procedure, we prove that the warm-start property of the distributions between successive epochs and the rounding guarantee given by N samples.

5.1. Warm start and mixing

We need to prove that the measures between successive temperatures are not too far away in the ℓ_2 sense, so that the samples from the previous epoch can be treated as a warm start for the next epoch. The following result is an extension of Lemma 6.1 in (Kalai and Vempala, 2006) to β -log-concave functions.

Lemma 15 *Let $g(x) = \exp(-F(x))$ be a β -log-concave function. Let μ_i be a distribution with density proportional to $\exp\{-F(x)/T_i\}$, supported on \mathcal{K} . Let $T_i = T_{i-1} \left(1 - \frac{1}{\sqrt{n}}\right)$. Then*

$$\|\mu_i/\mu_{i+1}\| \leq C_\gamma = 5 \exp(2\beta/T_i)$$

Next we account for the impact of using the final distribution from the previous epoch $\sigma^{(0)}$ as a “warm-start.”

Theorem 16 Fix a target accuracy $\gamma \in (0, 1/e)$ and let g be an $\beta/2$ -approximately log-concave function in \mathbb{R}^n . Suppose the simulated annealing algorithm (Algorithm 4) is run for $K = \sqrt{n} \log(1/\rho)$ epochs with temperature parameters $T_i = (1 - 1/\sqrt{n})^i$, $0 \leq i \leq K$. If the Hit-and-Run with the unidimensional sampling scheme (Algorithm 2) is run for $m = \mathcal{O}^*(n^3)$ number of steps prescribed in Theorem 12, the algorithm maintains that

$$d_{\text{tv}}(\hat{\sigma}_i^{(m)}, \pi_{g_i}) \leq e\gamma \quad (5)$$

at every epoch i , where $\hat{\sigma}_i^{(m)}$ is the distribution of the m -th step of Hit-and-Run. Here, m depends polylogarithmically on ρ^{-1} .

5.2. Optimization guarantee

We prove an extension of Lemma 4.1 in (Kalai and Vempala, 2006):

Theorem 17 Let f be a convex function. Let X be chosen according to a distribution with density proportional to $\exp\{-f(x)/T\}$. Then

$$\mathbb{E}_f f(X) - \min_{x \in \mathcal{K}} f(x) \leq (n+1)T$$

Furthermore, if F is such that $|F - f|_\infty \leq \rho$, for X chosen from a distribution with density proportional to $\exp\{-F(x)/T\}$, we have

$$\mathbb{E}_F f(X) - \min_{x \in \mathcal{K}} f(x) \leq (n+1)T \cdot \exp(2\rho/T)$$

The above Theorem implies that the final temperature T_K in the simulated annealing procedure needs to be set as $T_K = \epsilon/n$. This, in turn, leads to $K = \sqrt{n} \log(n/\epsilon)$ epochs. The oracle complexity of optimizing F is then, informally,

$$\mathcal{O}^*(n^3) \text{ queries per sample} \times \mathcal{O}^*(n) \text{ parallel strands} \times \mathcal{O}^*(\sqrt{n}) \text{ epochs} = \mathcal{O}^*(n^{4.5})$$

The following corollary summarizes the computational complexity result:

Corollary 18 Suppose F is approximately convex and $|F - f| \leq \epsilon/n$ as in (1). The simulated annealing method with $K = \sqrt{n} \log(n/\epsilon)$ epochs produces a random point such that

$$\mathbb{E}f(X) - \min_{x \in \mathcal{K}} f(x) \leq \epsilon, \quad \mathbb{E}F(X) - \min_{x \in \mathcal{K}} F(x) \leq 2\epsilon.$$

Furthermore the number of oracle queries required by the method is $\mathcal{O}^*(n^{4.5})$.

6. Stochastic Convex Zeroth Order Optimization

Let $f : \mathcal{K} \rightarrow \mathbb{R}$ be the unknown convex L -Lipschitz function we aim to minimize. Within the model of convex optimization with stochastic zeroth-order oracle \mathbf{O} , the information returned upon a query $x \in \mathcal{K}$ is $f(x) + \epsilon_x$ where ϵ_x is the zero mean noise. We shall assume that the noise is sub-Gaussian with parameter σ . That is, $\mathbb{E} \exp(\lambda \epsilon_x) \leq \exp(\sigma^2 \lambda^2 / 2)$. It is easy to see from Chernoff's bound that for any $t \geq 0$ $\mathbb{P}(|\epsilon_x| \geq \sigma t) \leq 2 \exp(-t^2/2)$. We can decrease the noise level by repeatedly querying at x . Fix $\tau > 0$, to be determined later. The average $\bar{\epsilon}_x$ of τ observations is concentrated

as $\mathbb{P}(|\bar{\epsilon}_x| \geq \sigma t/\sqrt{\tau}) \leq 2\exp(-t^2/2)$. To use the randomized optimization method developed in this paper, we view $f(x) + \bar{\epsilon}_x$ as the value of $F(x)$ returned upon a single query at x . Since the randomized method does not re-visit x with probability 1, the function F is “well-defined”.

Let us make the above discussion more precise by describing three oracles. Oracle \mathbf{O}' draws noise ϵ_x for each $x \in \mathcal{K}$ prior to optimization. Upon querying $x \in \mathcal{K}$, the oracle deterministically returns $f(x) + \epsilon_x$, even if the same point is queried twice. Given that the optimization method does not query the same point (with probability one), this oracle is equivalent to an *oblivious* version of oracle \mathbf{O} of the original zeroth order stochastic optimization problem.

To define \mathbf{O}_α , let \mathcal{N}_α be an α -net in ℓ_∞ which can be taken as a box grid of \mathcal{K} . If $\mathcal{K} \subseteq RB_\infty$, the size of the net is at most $(R/\alpha)^n$. The oracle draws ϵ_x for each element $x \in \mathcal{N}_\alpha$, independently. Upon a query $x' \in \mathcal{K}$, the oracle deterministically returns $f(x) + \epsilon_x$ for $x \in \mathcal{N}_\alpha$ which is closest to x' . Note that \mathbf{O}_α is no more powerful than \mathbf{O}' , since the learner only obtains the information on the α -net. Oracle \mathbf{O}_α^τ is a small modification of \mathbf{O}_α . This modification models a repeated query at the same point, as described earlier. Parametrized by τ (the number of queries at the same point), oracle \mathbf{O}_α^τ draws random variables ϵ_x for each $x \in \mathcal{N}_\alpha$, but sub-Gaussian parameter of ϵ_x is $\sigma/\sqrt{\tau}$. The optimization algorithm pays for τ oracle calls upon a single call to \mathbf{O}_α^τ .

We argued that \mathbf{O}_α^τ is no more powerful than the original zeroth order oracle given that the algorithm does not revisit the point. In the rest of the section, we will work with \mathbf{O}_α^τ as the oracle model. For any x , denote the projection to the \mathcal{N}_α to be $\mathcal{P}_{\mathcal{N}_\alpha}(x)$. Define $F : \mathcal{K} \mapsto \mathbb{R}$ as

$$F(x) = f(\mathcal{P}_{\mathcal{N}_\alpha}(x)) + \epsilon_{\mathcal{P}_{\mathcal{N}_\alpha}(x)}$$

where $\mathcal{P}_{\mathcal{N}_\alpha}(x)$ is the closest to x point of \mathcal{N}_α in the ℓ_∞ sense. Clearly, $|F - f|_\infty \leq \max_{x \in \mathcal{N}_\alpha} |\epsilon_x| + \alpha L$ where L is the (ℓ_∞) Lipschitz constant. Since $(\epsilon_x)_{x \in \mathcal{N}_\alpha}$ define a finite collection of sub-Gaussian random variables with sub-Gaussian parameter σ , we have that with probability at least $1 - \delta$

$$\max_{x \in \mathcal{N}_\alpha} |\epsilon_x| \leq \sigma \sqrt{\frac{2n \log(R/\alpha) + 2 \log(1/\delta)}{\tau}}$$

From now on, we condition on this event, which we call \mathcal{E} . To guarantee (1), we set the above upper bound equal to $\frac{\epsilon}{2n} = \alpha L$ where τ is the parameter from oracle \mathbf{O}_α^τ . Solving for τ and α :

$$\tau = \frac{\sigma^2 n^2 (8n \log(R/\alpha) + 8 \log(1/\delta))}{\epsilon^2} = \frac{\sigma^2 n^2 (8n \log(2LRn/\epsilon) + 8 \log(1/\delta))}{\epsilon^2} = \mathcal{O}^*(n^3/\epsilon^2)$$

and $\alpha = \epsilon/(2Ln)$. Note here L affects τ only logarithmically, and, in particular, we could have defined the Lipschitz constant with respect to ℓ_2 . We also observe that the oracle model depends on α and, hence, on the target accuracy ϵ . However, because the dependence on α is only logarithmic, we can take α to be much smaller than ϵ .

Together with the $\mathcal{O}^*(n^{4.5})$ oracle complexity proved in the previous section for optimizing F , the choice of $\tau = \mathcal{O}^*(n^3\epsilon^{-2})$ evaluations per time step yields a total oracle complexity of

$$\mathcal{O}^*(n^{7.5}\epsilon^{-2})$$

for the problem of stochastic convex optimization with zeroth order information. We observe that a factor of n^2 in oracle complexity comes from the union bound over the exponential-sized discretization of the set. This (somewhat artificial) factor can be reduced or removed under additional assumptions on the noise, such as a draw from a Gaussian process with spatial dependence over \mathcal{K} .

Alternatively, this n^2 factor could be removed completely if we could take a union bound over the polynomial number of points visited by the algorithm. Such an argument, however, appears to be tricky.

7. Optimization of Non-convex Functions with Decreasing Fluctuations

Assume the non-convex function $F(x)$ has the property that the “amount of non-convexity” is decreasing as x gets close to its global minimum x^* . If one has some control on the rate of this decrease, it is possible to break the optimization problem into stages, where at each stage one optimizes to the current level of non-convexity, redefines the optimization region to guarantee a smaller amount of “non-convexity,” and proceeds to the next stage. We are not aware of optimization methods for such a problem, and it is rather surprising that one may obtain provable guarantees through simulated annealing.

As one example, consider the problem of stochastic zeroth order optimization where the noise level decreases as one approaches the optimal point. Then, one would expect to obtain a range of oracle complexities between $\log(1/\epsilon)$ and $1/\epsilon^2$ in terms of the rate of the noise decrease.

Let us formalize the above discussion. Suppose there exists a 1-Lipschitz α -strongly convex function $f(x)$ with minimum achieved at $x^* \in \mathcal{K}$:

$$f(x) - f(x^*) \geq \langle \nabla f(x^*), x - x^* \rangle + \frac{\alpha}{2} \|x - x^*\|^2 \geq \frac{\alpha}{2} \|x - x^*\|^2.$$

Define a measure of “non-convexity” of F with respect to f in a ball of radius r around x^* :

$$\Delta(r) := \sup_{x \in \mathcal{B}_2^n(x^*, r)} |F(x) - f(x)|.$$

We have in mind the situation where $\Delta(r)$ decreases as r decreases to 0. At the t th stage of the optimization problem, suppose we start with an ℓ_2 ball $\mathcal{B}_2^n(x_{t-1}, 2r_t)$ of radius $2r_t$ with the property $\mathcal{B}_2^n(x^*, 3r_t) \supset \mathcal{B}_2^n(x_{t-1}, 2r_t) \supset \mathcal{B}_2^n(x^*, r_t)$. Next, we run the simulated annealing procedure for the approximately log-concave function defined over this ball. After $\mathcal{O}^*(n^{4.5})$ queries, we are provided with a point x_t such that in expectation (or high probability)

$$f(x_t) - f(x^*) \leq Cn \cdot \Delta(3r_t)$$

with some universal constant $C > 0$. Thanks to strong convexity, $\frac{\alpha}{2} \|x_t - x^*\|^2 \leq f(x_t) - f(x^*) \leq Cn \cdot \Delta(3r_t)$ which suggests the recursive definition of r_{t+1} :

$$\frac{\alpha}{2Cn} r_{t+1}^2 := \frac{\alpha}{2Cn} \|x_t - x^*\|_{\ell_2}^2 \leq \Delta(3r_t).$$

At stage $t + 1$ we restrict the region to be $\mathcal{B}_2^n(x_t, 2r_{t+1}) \supset \mathcal{B}_2^n(x^*, r_{t+1})$ and run the optimization algorithm again with the new parameter of approximate convexity. The recursion formula for the radius from r_t to r_{t+1} satisfies $\frac{\alpha}{2Cn} r_{t+1}^2 \leq \Delta(3r_t)$. The recursion formula yields a fixed point — a “critical radius” r^* where no further improved can be achieved, with $\frac{\alpha}{2Cn} (r^*)^2 = \Delta(3r^*)$. Let us explore two examples:

$$\text{Polynomial : } \Delta(r) = cr^p, \quad 0 < p < 2, \quad \text{Logarithmic : } \Delta(r) = c \log(1 + dr),$$

where $c, d > 0$ are constants. For the polynomial case, the critical radius is $r^* = \left(\frac{2 \cdot 3^p c C n}{\alpha}\right)^{\frac{1}{2-p}}$, and the required number of epochs is at most $\frac{\log \log(r_0/r^*) + \log(1/\epsilon)}{\log(2/p)}$ if we want to get $r_t = (1 + \epsilon)r$. For the logarithmic case, the critical radius is the unique non-zero solution to $\frac{2cCn}{\alpha} \log(1 + 3dr) = r^2$. We conclude that at an $\mathcal{O}^*(1)$ multiplicative overhead on the number of oracle calls, we can optimize to any level of precision above the fixed point r^* of the non-convexity decay function.

8. Further Applications

In this section, we sketch several applications of the zeroth-order optimization method we introduced. Our treatment is cursory, meant only to give a sense of the range of possible domains.

8.1. Private computation with distributed data

Suppose $i = 1, \dots, n$ are entities—say, hospitals—that each possess private data in the form of m covariate-response pairs $\{(x_{i,j}, y_{i,j})\}_{j=1}^m$. A natural approach to analyzing the aggregate data is to compute a minimizer w^* of

$$f(w) = \frac{1}{mn} \sum_{i,j} \ell(x_{i,j}, y_{i,j}; w) + R(w) \quad (6)$$

for some convex regularization function R and a convex (in w) loss function ℓ . For instance, $\ell(x_{i,j}, y_{i,j}; w) = (y_{i,j} - x_{i,j} \cdot w)^2$ and $R(w) = 0$ would correspond to the problem of linear regression. Given that the hospitals are not willing to release the data to a central authority that would perform the computation, how can the objective (6) be minimized? We propose to use the simulated annealing method of this paper. To this end, we need to specify what happens when the value $f(w_t)$ at the current point w_t is requested. Consider the following idea. The current w_t is passed to a randomly chosen hospital $I_t \sim \text{unif}(1, \dots, n)$. The hospital, in turn, privately chooses an index $J_t \sim \text{unif}(1, \dots, m)$, computes the loss $\ell(x_{I_t, J_t}, y_{I_t, J_t}; w_t)$, adds zero-mean noise $\eta_t \sim N(0, 1)$, and passes the resulting value $v_t = \ell(x_{I_t, J_t}, y_{I_t, J_t}; w_t) + \eta_t$ back to the central authority. Since the computation is done privately by the hospital, the only value released to the outside world is the noisy residual. It is easy to check that v_t is an unbiased estimate of $\mathbb{E}[v_t] = f(w_t)$ with respect to the random variables (I_t, J_t) and η_t . Moreover, the noise level with respect to each source of randomness is of constant order. By repeatedly querying for the noisy value at w_t , the algorithm can reduce the noise variance, as in Section 6, yet—importantly—the returned value is for a potentially different random choice of the hospital and the data point. This latter fact means that repeated querying does not allow the central authority to learn a specific data point. Interestingly, the additional layer of privacy given by the zero-mean noise η_t presents no added difficulty to the minimization procedure, except for slightly changing a constant in the number of required queries.

8.2. Two-stage stochastic programming

Dyer et al. (2013) discuss the following mathematical programming formulation:

$$\max \quad px + \mathbb{E}[\max\{qy | Wy \leq Tx - \xi, y \in \mathbb{R}^{n_1}\}] \quad \text{subject to } Ax \leq b, \quad (7)$$

where $q \in \mathbb{R}^{n_1}$, $W \in \mathbb{R}^{d \times n_1}$, and $T \in \mathbb{R}^{d \times n}$. The expectation is taken over the random variable ξ . This problem is concave in x , and can be solved in two stages. If, given x , an approximate value for the inner expected maximum can be computed, the problem falls squarely into the setting of zeroth order optimization with approximate function evaluations. While the method of Dyer et al. (2013) is simpler, its dependence on the target accuracy ϵ is worse. Additionally, the method of this paper can deal directly with constraint sets with non-smooth boundaries; the method can also handle more general functions in (7) that are not smooth.

References

- Adamczak, R., Litvak, A. E., Pajor, A., Tomczak-Jaegermann, N., et al. (2010). Quantitative estimates of the convergence of the empirical covariance matrix in log-concave ensembles. *J. Amer. Math. Soc.*, 23(2):535–561.
- Agarwal, A., Foster, D. P., Hsu, D., Kakade, S. M., and Rakhlin, A. (2013). Stochastic convex optimization with bandit feedback. *SIAM Journal on Optimization*, 23(1):213–240.
- Applegate, D. and Kannan, R. (1991). Sampling and integration of near log-concave functions. In *Proceedings of the twenty-third annual ACM symposium on Theory of computing*, pages 156–163. ACM.
- Belloni, A. and Chernozhukov, V. (2009). On the computational complexity of mcmc-based estimators in large samples. *The Annals of Statistics*, pages 2011–2055.
- Bertsimas, D. and Vempala, S. (2004). Solving convex programs by random walks. *Journal of the ACM (JACM)*, 51(4):540–556.
- Bourgain, J. (1996). Random points in isotropic convex sets. *Convex geometric analysis*, 34:53–58.
- Cholewa, P. W. (1984). Remarks on the stability of functional equations. *Aequationes Mathematicae*, 27(1):76–86.
- Conn, A. R., Scheinberg, K., and Vicente, L. N. (2009). *Introduction to derivative-free optimization*, volume 8. Siam.
- Dyer, M., Kannan, R., and Stougie, L. (2013). A simple randomised algorithm for convex optimisation. *Mathematical Programming*, pages 1–23.
- Green, J. W. et al. (1952). Approximately convex functions. *Duke Mathematical Journal*, 19(3):499–504.
- Guédon, O. and Rudelson, M. (2007). L_p -moments of random vectors via majorizing measures. *Advances in Mathematics*, 208(2):798–823.
- Kalai, A. and Vempala, S. (2006). Simulated annealing for convex optimization. *Mathematics of Operations Research*, 31(2):253–266.
- Kannan, R., Lovász, L., and Simonovits, M. (1997). Random walks and an $o^*(n^5)$ volume algorithm for convex bodies. *Random structures and algorithms*, 11(1):1–50.
- Laczkovich, M. (1999). The local stability of convexity, affinity and of the jensen equation. *aequationes mathematicae*, 58(1-2):135–142.
- Liang, T., Narayanan, H., and Rakhlin, A. (2014). On zeroth-order stochastic convex optimization via random walks. *arXiv preprint arXiv:1402.2667*.
- Lovász, L. and Simonovits, M. (1993). Random walks in a convex body and an improved volume algorithm. *Random structures & algorithms*, 4(4):359–412.

- Lovász, L. and Vempala, S. (2006a). 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.
- Lovász, L. and Vempala, S. (2006b). Hit-and-run from a corner. *SIAM Journal on Computing*, 35(4):985–1005.
- Lovász, L. and Vempala, S. (2007). The geometry of logconcave functions and sampling algorithms. *Random Structures & Algorithms*, 30(3):307–358.
- Nemirovskii, A. S. and Yudin, D. B. (1983). *Problem Complexity and Method Efficiency in Optimization*. John Wiley.
- Rakhlín, A., Shamir, O., and Sridharan, K. (2012). Relax and randomize: From value to algorithms. In *Advances in Neural Information Processing Systems*, pages 2141–2149.
- Robert, C. P. and Casella, G. (2004). *Monte Carlo statistical methods*, volume 319. Citeseer.
- Rudelson, M. (1999). Random vectors in the isotropic position. *Journal of Functional Analysis*, 164(1):60–72.
- Shamir, O. (2012). On the complexity of bandit and derivative-free stochastic convex optimization. *CoRR*, abs/1209.2388.
- Tropp, J. A. (2012). User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 12(4):389–434.
- Vershynin, R. (2010). Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*.

Appendix A. Rounding for β -log-concave functions

The simulated annealing procedure runs $N = \mathcal{O}^*(n)$ strands of random walk to round the log-concave distribution into near-isotropic position (say 1/2-near-isotropic) at each temperature. The N strands do not interact and thus the computation within each epoch can be parallelized, further reducing the time complexity of the algorithm. For N i.i.d isotropic random vectors $X_i \in \mathbb{R}^n, 1 \leq i \leq N$ sampled from a log-concave measure, the following concentration holds when N is large enough:

$$\sup_{\|v\|_{\ell_2}=1} \left| v^T \left(\frac{1}{N} \sum_{i=1}^N X_i X_i^T \right) v - 1 \right| \leq \frac{1}{2}$$

or, equivalently,

$$\frac{1}{2} \leq \sigma_{\min} \left(\frac{1}{N} \sum_{i=1}^N X_i X_i^T \right) \leq \sigma_{\max} \left(\frac{1}{N} \sum_{i=1}^N X_i X_i^T \right) \leq \frac{3}{2}. \quad (8)$$

Theorems of this type have been first achieved for uniform measures on the convex body \mathcal{K} (measures with bounded Orlicz ψ_2 norm). [Bourgain \(1996\)](#) proved this holds as long as $N \geq Cn \log^3 n$. [Rudelson \(1999\)](#) improved this bound to $N \geq Cn \log^2 n$. For log-concave measures (with bounded Orlicz ψ_1 norm through Borell's lemma), [Guédon and Rudelson \(2007\)](#) proved a stronger version where $N \geq Cn \log n$. See also ([Adamczak et al., 2010](#)) for further improvements. For bounded (almost surely) vectors, we can instead appeal to the the following literature. Theorem 5.41 in ([Vershynin, 2010](#)) yields a spectral concentration bound for heavy tail random matrices with isotropic independent rows. (See also [Tropp \(2012\)](#) Theorem 4.1 for matrix Bernstein's type inequalities.)

For our problem, we need to prove (8) for independent near isotropic rows with β -log-concave measures. There are two ways to achieve this goal. The first is to invoke the [Guédon and Rudelson \(2007\)](#)'s result. Random vectors sampled from β -log-concave still belong to the Orlicz ψ_1 family, thus $N \geq \mathcal{O}(n \log n)$ is enough to achieve the goal with high probability. The second way is through the following lemma:

Lemma 19 ([Vershynin \(2010\)](#), [Theorem 5.41](#)) *Let X be an $N \times n$ matrix whose rows X_i are independent isotropic random vectors in \mathbb{R}^n . Let R be a number such that $\|X_i\|_{\ell_2} \leq R$ almost surely for all i . Then for every $t \geq 0$, one has*

$$\sqrt{N} - tR \leq \sigma_{\min}(X) \leq \sigma_{\max}(X) \leq \sqrt{N} + tR$$

with probability at least $1 - 2n \exp(-ct^2)$, where $c > 0$ is an absolute constant.

Clearly if we take $N \geq CR^2 \log n$, we have

$$\frac{1}{2} \leq \sigma_{\min} \left(\frac{1}{N} \sum_{i=1}^N X_i X_i^T \right) \leq \sigma_{\max} \left(\frac{1}{N} \sum_{i=1}^N X_i X_i^T \right) \leq \frac{3}{2}$$

with probability at least $1 - n^{-C}$, since \mathcal{K} is uniformly bounded within $R = \mathcal{O}(\sqrt{n})r$ and isotropic condition implies $r = \mathcal{O}(1)$ (which translates into $\|X_i\|_{\ell_2} \leq \mathcal{O}(\sqrt{n})$). Thus we conclude that $N = \times(n \log n)$ is enough for bringing a β -log concave measure into isotropic position.

Appendix B. Proofs of Section 3

We start with two useful lemmas that can be found in [Lovász and Vempala \(2007\)](#).

Lemma 20 ([Lovász and Vempala \(2007\)](#), [Lemma 5.19](#)) *Let $h : \mathbb{R}^n \rightarrow \mathbb{R}$ be a log-concave function. Define $M_h := \max h$ and $L_h(t) = \{x \in \mathbb{R}^n : h(x) \geq t\}$. Then for $0 < s < t < M_h$*

$$\frac{\text{vol}(L_h(s))}{\text{vol}(L_h(t))} \leq \left(\frac{\log(M_h/s)}{\log(M_h/t)} \right)^n.$$

Lemma 21 ([Lovász and Vempala \(2007\)](#), [Lemma 5.6\(a\)](#)) *Let X be a random point drawn from a log-concave distribution $h : \mathbb{R} \rightarrow \mathbb{R}_+$ and let $M_h := \max_{x \in \mathbb{R}} h(x)$. Then for every $t \geq 0$*

$$\mathbb{P}(h(X) \leq t) < \frac{t}{M_h}.$$

Proof [Proof of [Lemma 3](#)] The proof is straightforward:

$$\begin{aligned} g(\alpha x + (1 - \alpha)y) &\geq e^{-\beta/2} h(\alpha x + (1 - \alpha)y) \geq e^{-\beta/2} h(x)^\alpha h(y)^{1-\alpha} \\ &\geq e^{-\beta/2} (e^{-\beta/2} g(x))^\alpha (e^{-\beta/2} g(y))^{1-\alpha} \geq e^{-\beta} g(x)^\alpha g(y)^{1-\alpha}. \end{aligned}$$

■

Appendix C. Proofs of Section 4.1

Proof [Proof of [Lemma 7](#)]

Consider a unidimensional β -log-concave function $g : \mathbb{R} \rightarrow \mathbb{R}$. In view of [Lemma 4](#), g can be “sandwiched” by a log-concave function h such that $e^{-\beta} h(x) \leq g(x) \leq h(x)$.

Given ℓ , we want to find $p \in \ell$ such that $g(p) \geq e^{-3\beta} \max_{z \in \ell} g(z)$. We use the following 3-point method, inspired by [Agarwal et al. \(2013\)](#), to provide such a point. Let us work with the convex function

$$\tilde{h} = -\log h$$

and a nearly-convex function

$$\tilde{g} = -\log g.$$

The sandwiching guarantee can be written as

$$\tilde{h}(x) \leq \tilde{g}(x) \leq \tilde{h}(x) + \beta.$$

We now claim that each iteration of the “while” loop of [Algorithm 3](#) maintains the following property: either the length of the interval is reduced by at least $3/4$ while still containing the optimal point, or we have the output point p that satisfies

$$\tilde{g}(p) \leq \min \tilde{g} + 3\beta$$

In the latter case, $g(p) \geq e^{-3\beta} \max_{z \in \ell} g(z)$ as desired.

There are essentially two cases. First, if $\tilde{g}(x_l) - \tilde{g}(x_r) > \beta$ (or similarly we can argue for $|\tilde{g}(x_l) - \tilde{g}(x_c)| > \beta$ and $|\tilde{g}(x_r) - \tilde{g}(x_c)| > \beta$), we have

$$\tilde{h}(x_l) + \beta \geq \tilde{g}(x_l) > \tilde{g}(x_r) + \beta \geq \tilde{h}(x_r) + \beta$$

and thus $\tilde{h}(x_l) > \tilde{h}(x_r)$. Because of convexity of \tilde{h} we can safely remove $[\underline{x}, x_l]$ with the remaining interval still containing the point we are looking for. Second case is when

$$|\tilde{g}(x_l) - \tilde{g}(x_r)| \leq \beta, \quad |\tilde{g}(x_l) - \tilde{g}(x_c)| \leq \beta, \quad |\tilde{g}(x_r) - \tilde{g}(x_c)| \leq \beta$$

Here, we can show the function $g(x)$ is flat enough for $[\underline{x}, \bar{x}]$ and thus the best of x_l, x_c, x_r are good enough. It is not hard to see that

$$|\tilde{h}(x_l) - \tilde{h}(x_r)| \leq 2\beta, \quad |\tilde{h}(x_l) - \tilde{h}(x_c)| \leq 2\beta, \quad |\tilde{h}(x_r) - \tilde{h}(x_c)| \leq 2\beta.$$

Consider the point x_l . By convexity of \tilde{h} , there must be a supporting line $k_l(x)$ that is below the convex function \tilde{h} and such that $k_l(x_l) = \tilde{h}(x_l)$. Thus

$$\min_{[\underline{x}, x_l]} \tilde{h}(x) \geq \min_{[\underline{x}, x_l]} k_l(x) \geq k_l(x_l) - 2\beta = \tilde{h}(x_l) - 2\beta$$

using the fact that $|\underline{x} - x_l| = |x_l - x_c|$. Similarly we can prove

$$\min_{[x_l, x_r]} \tilde{h}(x) \geq \tilde{h}(x_c) - 2\beta, \quad \min_{[x_r, \bar{x}]} \tilde{h}(x) \geq \tilde{h}(x_r) - 2\beta.$$

Thus

$$\min_{[\underline{x}, \bar{x}]} \tilde{h}(x) \geq \min(\tilde{h}(x_l), \tilde{h}(x_c), \tilde{h}(x_r)) - 2\beta.$$

By sandwiching

$$\min_{[\underline{x}, \bar{x}]} \tilde{g}(x) \geq \min(\tilde{g}(x_l), \tilde{g}(x_c), \tilde{g}(x_r)) - 3\beta$$

and, hence,

$$\tilde{g}(p) \leq \min_{[\underline{x}, \bar{x}]} \tilde{g}(x) + 3\beta.$$

It remains to show that the algorithm will terminate in an $\mathcal{O}^*(1)$ number of steps. Let \tilde{L} be the Lipschitz constant of \tilde{h} . By the time the interval is shrunk to $|\underline{x} - \bar{x}| \leq \beta/\tilde{L}$, the algorithm must have entered the second case above and terminated. \blacksquare

Proof [Proof of Lemma 8] Consider a unidimensional β -log-concave function $g : \mathbb{R} \rightarrow \mathbb{R}$. In view of Lemma 4, g can be “sandwiched” by a log-concave function h such that $e^{-\beta}h(x) \leq g(x) \leq h(x)$.

We consider the interval $[x_l, x_r] = [p, \bar{x}]$ (the other case follows similarly). If $g(\bar{x}) \geq \frac{1}{2}e^{-\beta}\tilde{\epsilon}g(p)$, set $e_1 = \bar{x}$. Otherwise we have $g(\bar{x}) < \frac{1}{2}e^{-\beta}\tilde{\epsilon}g(p)$ and we proceed. The procedure always query the midpoint x_m of current interval $[x_l, x_r]$. If $g(x_m) > \tilde{\epsilon}g(p)$ set $x_l = x_m$, or if $g(x_m) < \frac{1}{2}e^{-\beta}\tilde{\epsilon}g(p)$ set $x_r = x_m$, and continue the search. Either operation halves the interval. If the midpoint x_m is

such that $\frac{1}{2}e^{-\beta}\tilde{\epsilon}g(p) \leq g(x_m) \leq \tilde{\epsilon}g(p)$, stop the process and return $e_1 = x_m$. At every iteration, the interval $[x_l, x_r]$ is such that $g(x_l) > \tilde{\epsilon}g(p)$ and $g(x_r) < \frac{1}{2}e^{-\beta}\tilde{\epsilon}g(p)$. We now claim that the algorithm must terminate in an $\mathcal{O}^*(1)$ number of steps. Let $\tilde{h} = -\log h$, and let \tilde{L} be the Lipschitz constant of \tilde{h} . As soon as the length of the current interval $|x_l - x_r| < 1/(2\tilde{L})$, we have $|\tilde{h}(x_l) - \tilde{h}(x_r)| < 1/2$. Thus $h(x_l)/h(x_r) < e^{1/2}$ and $g(x_l)/g(x_r) < e^{1/2+\beta}$, implying that both $g(x_l) > \tilde{\epsilon}g(p)$ and $g(x_r) < \frac{1}{2}e^{-\beta}\tilde{\epsilon}g(p)$ cannot be true at the same time as $2e^\beta > e^{1/2+\beta}$. Hence, the algorithm terminates in a number of steps that is logarithmic in \tilde{L} . \blacksquare

Proof [Proof of Lemma 6] Let h be the log-concave function associated with the β -log-concave function g in the sense of Lemma 4, so that $e^{-\beta}h(x) \leq g(x) \leq h(x)$ for all $x \in \ell$, and $L_f(t)$ denote the (upper) level set of a function f at level t .

We note that since $L_g(t) \subset L_h(t)$ and $e^{-\beta}h(x) \leq g(x) \leq h(x)$, (4) implies that

$$L_h(e^\beta\tilde{\epsilon}g(p)) \subseteq [e_{-1}, e_1] \subseteq L_h(\frac{1}{2}e^{-\beta}\tilde{\epsilon}g(p)). \quad (9)$$

Moreover, either $e_{-1} = \underline{x}$ or $g(e_{-1}) \leq \tilde{\epsilon}g(p)$ which implies $h(e_{-1}) \leq e^\beta\tilde{\epsilon}g(p) < \frac{1}{2}e^{-\beta}g(p)$ if $\tilde{\epsilon} < \frac{1}{2}e^{-2\beta}$.

The stationary distribution for this sampling scheme is a truncated distribution according to the β -log-concave distribution g restricted to $[e_{-1}, e_1]$. (Indeed, this correspond to the classic Accept-Reject method to simulate g based on the uniform distribution with constant $M := g(p)e^{3\beta}$, see Robert and Casella (2004) page 49.) Therefore

$$\begin{aligned} d_{\text{tv}}(\pi_{g,\ell}, \hat{\pi}_{g,\ell}) &= d_{\text{tv}}(\pi_{g,\ell}1\{\ell \setminus [e_{-1}, e_1]\}, \hat{\pi}_{g,\ell}1\{\ell \setminus [e_{-1}, e_1]\}) + d_{\text{tv}}(\pi_{g,\ell}1\{[e_{-1}, e_1]\}, \hat{\pi}_{g,\ell}1\{[e_{-1}, e_1]\}) \\ &\leq \mathbb{P}_{z \sim g}(z \notin [e_{-1}, e_1]) + \frac{\mathbb{P}_{z \sim g}(z \notin [e_{-1}, e_1])}{1 - \mathbb{P}_{z \sim g}(z \notin [e_{-1}, e_1])}. \end{aligned}$$

Next we verify that the truncation error (in the total variation norm) of restricting g to $[e_{-1}, e_1]$ instead of ℓ is of the desired order. By Lemma 21 which quantifies the tail decay of unidimensional log-concave measures, we have

$$\begin{aligned} \mathbb{P}_{z \sim g}(z \notin [e_{-1}, e_1]) &\leq e^\beta \cdot \mathbb{P}_{z \sim h}(z \notin [e_{-1}, e_1]) \leq e^\beta \cdot \mathbb{P}_{z \sim h}(z \notin L_h(e^\beta\tilde{\epsilon}g(p))) \\ &\leq e^\beta \cdot \mathbb{P}_{z \sim h}(z \notin L_h(e^\beta\tilde{\epsilon}M_h)) \leq e^\beta \cdot \mathbb{P}_{z \sim h}(h(z) \leq e^\beta\tilde{\epsilon}M_h) \leq e^{2\beta} \cdot \tilde{\epsilon} \end{aligned}$$

where we used (9).

Thus, provided that $e^{2\beta}\tilde{\epsilon} \leq 1/2$, the total variation distance between the truncated measure $\hat{\pi}_{g,\ell}$ supported on $[e_{-1}, e_1]$ satisfies

$$d_{\text{tv}}(\pi_{g,\ell}, \hat{\pi}_{g,\ell}) \leq 3e^{2\beta}\tilde{\epsilon}.$$

In order to bound the number of evaluations we first bound the probability of the event the event $\{g(x) > \frac{1}{2}e^{-2\beta}g(p)\}$. Indeed, by (9), if $X \sim U([e_{-1}, e_1])$ and $Z \sim U(L_h(\frac{1}{2}\tilde{\epsilon}e^{-\beta}g(p)))$, we have

$$\begin{aligned} \mathbb{P}_X\left(g(X) \geq \frac{1}{2}e^{-2\beta}g(p)\right) &\geq \mathbb{P}_X\left(h(X) \geq \frac{1}{2}e^{-\beta}g(p)\right) \\ &\geq \mathbb{P}_Z\left(h(Z) \geq \frac{1}{2}e^{-\beta}h(p)\right) \\ &= \frac{\text{vol}(L_h(\frac{1}{2}e^{-\beta}h(p)))}{\text{vol}(L_h(\frac{1}{2}\tilde{\epsilon}e^{-\beta}h(p)))} \end{aligned}$$

By Lemma 20 with $s = \frac{1}{2}\tilde{\epsilon}e^{-\beta}h(p)$ and $t = \frac{1}{2}e^{-\beta}h(p)$, it follows that

$$\frac{\text{vol}(L_h(\frac{1}{2}e^{-\beta}h(p)))}{\text{vol}(L_h(\frac{1}{2}\tilde{\epsilon}e^{-\beta}h(p)))} \geq \frac{\log \frac{\max h}{\frac{1}{2}e^{-\beta}h(p)}}{\log \frac{\max h}{\frac{1}{2}\tilde{\epsilon}e^{-\beta}h(p)}} = \frac{\log \frac{\max h}{h(p)} + \log 2 + \beta}{\log \frac{\max h}{h(p)} + \log 2 + \log 1/\tilde{\epsilon} + \beta} \geq \frac{\log 2 + \beta}{\log 2 + \beta + \log 1/\tilde{\epsilon}} \geq \frac{\log 2}{\log(2/\tilde{\epsilon})}.$$

Then, since $r \sim U([0, 1])$ we have

$$\mathbb{P}\left(r \leq \frac{e^{-3\beta}g(x)}{g(p)} \mid g(x) \geq \frac{1}{2}e^{-2\beta}g(p)\right) \geq \frac{1}{2}e^{-5\beta}$$

and thus

$$\mathbb{P}\left(r \leq \frac{e^{-3\beta}g(x)}{g(p)}\right) \geq \frac{e^{-5\beta} \log 2}{2 \log(2/\tilde{\epsilon})}.$$

Since we have a lower bound on the acceptance probability on each sampling step, the number of iterations we need to sample is of the order $\frac{e^{-5\beta} \log 2}{2 \log(2/\tilde{\epsilon})}$. This quantity is $\mathcal{O}(\log(1/\tilde{\epsilon}))$ if β is $\mathcal{O}(1)$. ■

Appendix D. Proofs of Section 4.2

Proof [Proof of Theorem 10] The proof follows closely the arguments in the proof of Theorem 3.7 in Lovász and Vempala (2006a) for bounded sets with modifications to avoid the truncation device discussed in Section 3.3 of Lovász and Vempala (2006a). Define the step-size $F(x)$ by $P(\|x - y\| \leq F(x)) = 1/8$ where y is a random step from x . Next define $\lambda(x, t) = \text{vol}((x + tB) \cap L(\frac{3}{4}f(x)))/\text{vol}(tB)$ and $s(x) = \sup\{t > 0 : \lambda(x, t) \geq 63/64\}$. Finally, $\alpha(x) = \inf\{t \geq 3 : P(f(y) \geq tf(x)) \leq 1/16\}$ where y is a hit-and-run step from x .

Let $\mathcal{K} = S_1 \cup S_2$ be a partition into measurable sets, where $S_1 = S$ and $p = \pi_h(S_1) \leq \pi_h(S_2)$. For for $D = R \log(C'n^2/p)$ we will prove that

$$\int_{S_1} P_x(S_2) dx \geq \frac{1}{CnD \log \frac{nD}{p}} \pi_h(S_1). \quad (10)$$

Consider the points that are deep inside these sets with respect to 1-step distribution

$$S'_1 = \{x \in S_1 : P_x(S_2) < 1/1000\} \text{ and } S'_2 = \{x \in S_2 : P_x(S_1) < 1/1000\},$$

and the complement $S'_3 = \mathcal{K} \setminus S'_1 \cup S'_2$.

Suppose $\pi_h(S'_1) < \pi_h(S_1)/2$. Then

$$\int_{S_1} P_x(S_2) dx \geq \frac{1}{1000} \pi_h(S_1 \setminus S'_1) \geq \frac{1}{2000} \pi_h(S_1)$$

which proves (10). Thus we can assume

$$\pi_h(S'_1) \geq \pi_h(S_1)/2 \text{ and } \pi_h(S'_2) \geq \pi_h(S_2)/2. \quad (11)$$

Define the exceptional subset W as set of points u for which $\alpha(u)$ is very large

$W = W_1 \cup W_2$, where $W_1 := \{u \in S : \alpha(u) \geq 2^{27}nD/p\}$ and $W_2 := \{u \in \mathcal{K} : \|x - z_h\| \geq D\}$.

By Lemma 6.10 in [Lovász and Vempala \(2007\)](#), $\pi_h(W_1) \leq p/\{2^{23}nD\}$ and, by Lemma 5.17 in [Lovász and Vempala \(2007\)](#), $\pi_h(W_2) \leq \frac{p}{Cn^2}$. Now for any $u \in S'_1 \setminus W$ and $v \in S'_2 \setminus W$

$$d_{tv}(P_u, P_v) \geq P_u(S_1) - P_v(S_1) = 1 - P_u(S_2) - P_v(S_1) > 1 - \frac{1}{500}$$

from the definition of S'_1 and S'_2 . Thus by Lemma 6.8 of [Lovász and Vempala \(2006b\)](#), we have

$$d_h(u, v) \geq \frac{1}{128 \log(3 + \alpha(u))} \geq \frac{1}{2^{12} \log \frac{nD}{p}} \text{ or } |u - v| \geq \frac{1}{4\sqrt{n}} \max\{F(u), F(v)\}.$$

By Lemma 3.2 of [Lovász and Vempala \(2006b\)](#), the latter implies that

$$|u - v| \geq \frac{1}{2^8 \sqrt{n}} \max\{s(u), s(v)\}$$

In either case, by Lemma 3.5 in [Lovász and Vempala \(2006a\)](#), for any point $x \in [u, v]$, we have

$$\begin{aligned} s(x) &\leq 2^{14} \log \frac{nD}{p} |u - v| \sqrt{n} \\ &\leq 2^{14} \log \frac{nD}{p} d_{\mathcal{K} \setminus W_2}(u, v) D \sqrt{n} \end{aligned}$$

where the second inequality follows from $u, v \in \mathcal{K} \setminus W_2$.

Recall the original partition S'_1, S'_2 and $S'_3 = \mathcal{K} \setminus \{S'_1 \cup S'_2\}$ of \mathcal{K} . We will apply Theorem 2.1 of [Lovász and Vempala \(2006b\)](#) with a different partition. Consider the partition of $\mathcal{K} \setminus W_2$ defined as $\bar{S}_1 = S'_1 \setminus W$, $\bar{S}_2 = S'_2 \setminus W$ and $\bar{S}_3 = \mathcal{K} \setminus \{W_2 \cup S'_1 \cup S'_2\}$. These definitions imply that $\bar{S}_3 \subset S'_3 \cup W$ so that

$$\pi_h(\bar{S}_3) \leq \pi_h(S'_3) + \pi_h(W). \quad (12)$$

Define for $x \in \mathcal{K} \setminus W_2$

$$\bar{h}(x) = \frac{s(x)}{2^{16} D \sqrt{n} \log \frac{nD}{p}}$$

It follows that for any $u \in S'_1 \setminus W$ and $v \in S'_2 \setminus W$ and $x \in [u, v]$, we have $\bar{h}(x) \leq d_{\mathcal{K} \setminus W_2}(u, v)/3$, and since $\mathcal{K} \setminus W_2$ is a convex body, we have by Theorem 2.1 of [Lovász and Vempala \(2006b\)](#) that

$$\int_{\bar{S}_3} h(x) dx \geq \frac{\int_{\mathcal{K} \setminus W_2} \bar{h}(x) h(x) dx}{\int_{\mathcal{K} \setminus W_2} h(x) dx} \min \left\{ \int_{\bar{S}_1} h(x) dx, \int_{\bar{S}_2} h(x) dx \right\} \quad (13)$$

Although $\mathbb{E}_h(s(x))$ is large, we need a lower bound on $\int_{\mathcal{K} \setminus W_2} s(x) h(x) dx / \int_{\mathcal{K}} h(x) dx$. Since $s(x)$ can be large if \mathcal{K} is unbounded we modify the standard bound next. Because the level set of

measure $1/8$ contains a ball of radius r , we have

$$\begin{aligned}
 \int_{\mathcal{K} \setminus W_2} s(x)h(x)dx / \int_{\mathcal{K}} h(x)dx &= \int_{\mathcal{K} \setminus W_2} \int_0^{s(x)} dt h(x)dx / \int_{\mathcal{K}} h(x)dx \\
 &= \int_0^\infty \int_{\{x \in \mathcal{K} \setminus W_2 : \lambda(x,t) \geq 63/64\}} h(x)dx dt / \int_{\mathcal{K}} h(x)dx \\
 &\geq \frac{1}{16} \int_0^\infty \int_{\{x \in \mathcal{K} \setminus W_2 : \lambda(x,t) \geq 3/4\}} h(x)dx dt / \int_{\mathcal{K}} h(x)dx \\
 &\geq \frac{1}{16} \int_0^\infty \left(\int_{\{x \in \mathcal{K} : \lambda(x,t) \geq 3/4\}} h(x)dx - \int_{W_2} h(x)dx \right)_+ dt / \int_{\mathcal{K}} h(x)dx \\
 &\geq \frac{1}{16} \int_0^\infty \left(\frac{1}{2} - \frac{12t\sqrt{n}}{r} - \frac{p}{Cn^2} \right)_+ dt \\
 &\geq \frac{r}{2^{12}\sqrt{n}}
 \end{aligned}$$

where the first and third inequality follows from page 998 in [Lovász and Vempala \(2006b\)](#), the second by definition of W_2 where we take C and n large enough. Therefore, dividing both sides of (13) by $\int_{\mathcal{K}} h(x)dx$, (12), $\pi_h(S'_i) \geq \pi_h(S_i)/2$, $i = 1, 2$, and $p = \pi_h(S_1) \leq \pi_h(S_2)$, we have

$$\begin{aligned}
 \pi_h(S'_3) + \pi_h(W) &\geq \pi_h(\bar{S}_3) \geq \frac{\int_{\mathcal{K} \setminus W_2} \bar{h}(x)h(x)dx}{\int_{\mathcal{K} \setminus W_2} h(x)dx} \min\{\pi_h(\bar{S}_1), \pi_h(\bar{S}_2)\} \\
 &= \frac{\int_{\mathcal{K} \setminus W_2} s(x)h(x)dx / \int_{\mathcal{K}} h(x)dx}{2^{16}D\sqrt{n} \log \frac{nD}{p}} \min\{\pi_h(S'_1 \setminus W), \pi_h(S'_2 \setminus W)\} \\
 &\geq \frac{1}{2^{28}n(D/r) \log \frac{nD}{p}} \left\{ \min\{\pi_h(S'_1), \pi_h(S'_2)\} - \pi_h(W) \right\} \\
 &\geq \frac{1}{2^{28}n(D/r) \log \frac{nD}{p}} \left\{ \frac{1}{2} \min\{\pi_h(S_1), \pi_h(S_2)\} - \pi_h(W) \right\} \\
 &\geq \frac{1}{2^{28}n(D/r) \log \frac{nD}{p}} \left\{ \frac{\pi_h(S_1)}{2} - p/4 \right\} \geq \frac{\pi_h(S_1)}{2^{32}n(D/r) \log \frac{nD}{p}}
 \end{aligned}$$

where we used that $\pi_h(W) \leq p/4$. Therefore,

$$\int_{S_1} P_x(S_2)dx \geq \frac{\pi_h(S'_3)}{2000} \geq \frac{\pi_h(S_1)}{Cn(D/r) \log(nD/p)}$$

■

Proof [Proof of Theorem 11] Define the shorthand $\rho = e^{\beta/2}$. By sandwiching,

$$\rho^{-1}h(x) \leq g(x) \leq \rho h(x)$$

Then

$$\begin{aligned}
 \rho^{-2}\pi_h(x) &\leq \pi_g(x) \leq \rho^2\pi_h(x) \\
 \rho^{-2}P_u^h(A) &\leq P_u^g(A) \leq \rho^2P_u^h(A)
 \end{aligned}$$

for any $x, u \in \mathcal{K}$ and $A \subset \mathcal{K}$. Thus we have

$$\phi^g(S) \geq \rho^{-6}\phi^g(S).$$

The s -conductance bound can be derived as follows.

$$\begin{aligned}
 \phi_s^g &= \inf_{A \subset \mathcal{K}, s \leq \pi_g(A) \leq 1/2} \frac{\int_{x \in A} P_x^g(K \setminus A) d\pi_g}{\pi_g(A) - s} \\
 &\geq \rho^{-6} \inf_{A \subset \mathcal{K}, s \leq \pi_g(A) \leq 1/2} \frac{\int_{x \in A} P_x^h(K \setminus A) d\pi_h}{\pi_h(A) - s/\rho^2} \\
 &\geq \rho^{-6} \inf_{A \subset \mathcal{K}, s/\rho^2 \leq \pi_h(A) \leq 1/2} \frac{\int_{x \in A} P_x^h(K \setminus A) d\pi_h}{\pi_h(A) - s/\rho^2} = \rho^{-6} \phi_{s/\rho^2}^h
 \end{aligned}$$

■

Proof [Proof of Theorem 12] The H_s defined in Theorem 9 can be upper bounded by

$$\begin{aligned}
 H_s &= \sup_{A: \pi_g(A) \leq s} \int_A \left| \frac{d\sigma^{(0)}}{d\pi_g} - 1 \right| d\pi_g \\
 &\leq \sup_{A: \pi_g(A) \leq s} \left\{ \int_A \left(\frac{d\sigma^{(0)}}{d\pi_g} - 1 \right)^2 d\pi_g \cdot \int_A d\pi_g \right\}^{1/2} \\
 &\leq s^{1/2} \sup_{A: \pi_g(A) \leq s} \left\{ \int_A \left(\frac{d\sigma^{(0)}}{d\pi_g} - 1 \right)^2 d\pi_g \right\}^{1/2} \leq s^{1/2} \|\sigma^{(0)}/\pi_g\|^{1/2} = s^{1/2} M^{1/2}.
 \end{aligned}$$

Let us now use upper bound of Theorem 9 with $s = \left(\frac{\gamma}{2M^{1/2}}\right)^2$ and $D = \frac{2R}{r}$, as well as Theorem 11. We obtain

$$\begin{aligned}
 d_{tv}(\pi_g, \sigma^{(m)}) &\leq \frac{\gamma}{2} + \frac{2M}{\gamma} \left(1 - \frac{(\phi_s^g)^2}{2} \right)^m \\
 &\leq \frac{\gamma}{2} + \frac{2M}{\gamma} \left(1 - \frac{(\rho^{-6} \phi_{s/\rho^2}^h)^2}{2} \right)^m \\
 &\leq \frac{\gamma}{2} + \frac{2M}{\gamma} \exp\left(-\frac{m(\rho^{-6} \phi_{s/\rho^2}^h)^2}{2}\right)
 \end{aligned}$$

where $\rho = e^{\beta/2}$. In view of Theorem 10,

$$\phi_{s/\rho^2}^h \geq \frac{cr}{nR \log^2 \left(\frac{\rho^2 nR}{r(\gamma/2M^{1/2})^2} \right)}$$

we arrive at

$$d_{tv}(\pi_g, \sigma^{(m)}) \leq \frac{\gamma}{2} + \frac{2M}{\gamma} \exp\left(-\frac{m}{2} \left(\frac{\rho^{-6} cr}{nR \log^2 \frac{\rho^2 nRM}{r\gamma^2}} \right)^2\right)$$

Hence, if

$$m \geq Cn^2 \frac{\rho^{12} R^2}{r^2} \log^4 \frac{\rho^2 MnR}{r\gamma^2} \log \frac{M}{\gamma}$$

then

$$d_{\text{tv}}(\pi_g, \sigma^{(m)}) \leq \gamma. \quad \blacksquare$$

Proof [Proof of Theorem 14] Step 1. (Main Step) By the triangle inequality we have

$$d_{\text{tv}}(\hat{\sigma}^{(m)}, \pi_g) \leq d_{\text{tv}}(\hat{\sigma}^{(m)}, \sigma^{(m)}) + d_{\text{tv}}(\sigma^{(m)}, \pi_g) \quad (14)$$

The last term in (14) converges to zero at a geometric rate in m by Theorem 12. Specifically we have

$$d_{\text{tv}}(\sigma^{(m)}, \pi_g) \leq H_s + \frac{H_s}{s} \left(1 - \frac{(\phi_s^g)^2}{2}\right)^m.$$

To bound $d_{\text{tv}}(\hat{\sigma}^{(m)}, \sigma^{(m)})$, the total variation distance after m steps between the two random walks from their corresponding starting distributions $\hat{\sigma}^{(0)}$ and $\sigma^{(0)}$, write for any measurable set A

$$\hat{\sigma}^{(m)}(A) = \int (P_x^{\hat{g}})^{(m)}(A) \hat{\sigma}^{(0)} dx \quad \text{and} \quad \sigma^{(m)}(A) = \int (P_x^g)^{(m)}(A) \sigma^{(0)} dx$$

so that

$$d_{\text{tv}}(\sigma^{(m)}, \hat{\sigma}^{(m)}) \leq \sup_{u \in \mathcal{K}} d_{\text{tv}}\left((P_u^{\hat{g}})^{(m)}, (P_u^g)^{(m)}\right) + 2d_{\text{tv}}(\hat{\sigma}^{(0)}, \sigma^{(0)}).$$

The result follows from Step 2 that shows $\sup_{u \in \mathcal{K}} d_{\text{tv}}\left((P_u^{\hat{g}})^{(m)}, (P_u^g)^{(m)}\right) \leq m \sup_{\ell \subset \mathcal{K}} d_{\text{tv}}(\hat{\pi}_{g,\ell}, \pi_{g,\ell})$.

Step 2. (Error Propagation Bound in m Steps) The unidimensional sampling scheme produces a sample from a truncated distribution (see Lemma 6). That is, at each step of the Hit-and-Run algorithm, we are sampling from a truncated measure according to a truncated function \hat{g} along each line ℓ of g (approximately-log-concave function in \mathbb{R}^n). Let us denote the transition probability kernel starting from u for this truncated function to be $P_u^{\hat{g}}$ and the kernel for the original function is P_u^g . Let us bound the total variation distance between these two kernels through the spherical (elliptical) coordinate system, and with $p(\cdot)$ being the density corresponding to the measure $P(\cdot)$.

Suppose that $\sup_{\ell \subset \mathcal{K}} d_{\text{tv}}(\hat{\pi}_{g,\ell}, \pi_{g,\ell}) \leq \tilde{\epsilon}$. Since $p_u^{\hat{g}}(\theta) = p_u^g(\theta) = p(\theta)$, it holds that

$$\begin{aligned} 2d_{\text{tv}}(P_u^{\hat{g}}, P_u^g) &= \int |p_u^{\hat{g}}(r|\theta)p_u^{\hat{g}}(\theta) - p_u^g(r|\theta)p_u^g(\theta)| dr d\theta \\ &\leq \int \left\{ \int |p_u^{\hat{g}}(r|\theta) - p_u^g(r|\theta)| dr \right\} p(\theta) d\theta \\ &\leq 2\tilde{\epsilon} \int p(\theta) d\theta = 2\tilde{\epsilon} \end{aligned}$$

where on each line (over all θ according to the measure given by the linear transformation composed with uniform direction) the truncated distribution is an $\tilde{\epsilon}$ approximation to π_g . We now claim that

the m -fold iterate of the Hit-and-Run kernel satisfies $d_{\text{tv}} \left((P_u^{\hat{g}})^{(m)}, (P_u^g)^{(m)} \right) \leq m\tilde{\epsilon}$. Let us prove this by induction. Suppose it holds for $m - 1$ steps. Then

$$\begin{aligned}
 2d_{\text{tv}} \left((P_u^{\hat{g}})^{(m)}, (P_u^g)^{(m)} \right) &= \int \left| \int (p_u^{\hat{g}})^{(m-1)}(y) p_y^{\hat{g}}(x) - (p_u^g)^{(m-1)}(y) p_y^g(x) dy \right| dx \\
 &\leq \int \left| \int (p_u^{\hat{g}})^{(m-1)}(y) p_y^{\hat{g}}(x) - (p_u^g)^{(m-1)}(y) p_y^{\hat{g}}(x) dy \right| dx \\
 &\quad + \int \left| \int (p_u^g)^{(m-1)}(y) p_y^{\hat{g}}(x) - (p_u^g)^{(m-1)}(y) p_y^g(x) dy \right| dx \\
 &\leq \int \left| (p_u^{\hat{g}})^{(m-1)}(y) - (p_u^g)^{(m-1)}(y) \right| \left(\int p_y^{\hat{g}}(x) dx \right) dy \\
 &\quad + \int (p_u^g)^{(m-1)}(y) \left(\int |p_y^{\hat{g}}(x) - p_y^g(x)| dx \right) dy \\
 &\leq 2d_{\text{tv}} \left((P_u^{\hat{g}})^{(m-1)}, (P_u^g)^{(m-1)} \right) + 2 \max_y d_{\text{tv}}(P_y^{\hat{g}}, P_y^g) \\
 &\leq 2(m-1)\tilde{\epsilon} + 2\tilde{\epsilon} = 2m\tilde{\epsilon}
 \end{aligned}$$

■

Appendix E. Proofs of Section 5

Proof [Proof of Lemma 15] Define

$$Y(a) = \int_{\mathcal{K}} e^{-F(x)a} dx$$

With this notation, we have that

$$\left\| \frac{\mu_i}{\mu_{i+1}} \right\| = \frac{Y(2/T_i - 1/T_{i+1})Y(1/T_{i+1})}{Y(1/T_i)^2}$$

Define

$$G(x, t) = g\left(\frac{x}{t}\right)^t.$$

Then it holds that

$$\begin{aligned}
 G(\lambda(x, t) + (1 - \lambda)(x', t')) &= g\left(\frac{\lambda x + (1 - \lambda)x'}{\lambda t + (1 - \lambda)t'}\right)^{\lambda t + (1 - \lambda)t'} \\
 &= g\left(\frac{\lambda t}{\lambda t + (1 - \lambda)t'} \frac{x}{t} + \frac{(1 - \lambda)t'}{\lambda t + (1 - \lambda)t'} \frac{x'}{t'}\right)^{\lambda t + (1 - \lambda)t'} \\
 &\geq \exp(-\beta(\lambda t + (1 - \lambda)t')) g\left(\frac{x}{t}\right)^{\lambda t} g\left(\frac{x'}{t'}\right)^{(1 - \lambda)t'} \\
 &= \exp(-\beta(\lambda t + (1 - \lambda)t')) G(x, t)^\lambda G(x', t')^{(1 - \lambda)}
 \end{aligned}$$

Because

$$\int_{\mathcal{K}} G(x, t) dx = \int_{\mathcal{K}} g\left(\frac{x}{t}\right)^t dx = t^n \int_{\mathcal{K}} g(x)^t dx = t^n Y(t)$$

through Prékopa-Leindler inequality, we have

$$\left(\frac{a+b}{2}\right)^{2n} Y\left(\frac{a+b}{2}\right)^2 \geq \exp(-\beta(a+b)) a^n Y(a) b^n Y(b)$$

Take $a = 2/T_i - 1/T_{i+1}$ and $b = 1/T_{i+1}$. Then we have

$$\begin{aligned} \left\| \frac{\mu_i}{\mu_{i+1}} \right\| &= \frac{Y(2/T_i - 1/T_{i+1})Y(1/T_{i+1})}{Y(1/T_i)^2} \\ &\leq \left(\frac{1/T_i^2}{(2/T_i - 1/T_{i+1})(1/T_{i+1})} \right)^n \exp(2\beta/T_i) \\ &= \left(1 + \frac{1}{n - 2\sqrt{n}} \right)^n \exp(2\beta/T_i) \leq e^{n/(n-2\sqrt{n})} \exp(2\beta/T_i) \leq 5 \exp(2\beta/T_i) \end{aligned}$$

■

Proof [Proof of Theorem 16] Let us prove (5) by induction. Suppose at the end of epoch i , we have

$$d_{\text{tv}}(\hat{\sigma}_i^{(m)}, \pi_{g_i}) \leq \left(1 + \frac{1}{n \log 1/\rho} \right)^i \gamma$$

We identify $\hat{\sigma}_i^{(m)} = \hat{\sigma}_{i+1}^{(0)}$, $\sigma_{i+1}^{(0)} = \pi_{g_i}$. Hence, by Step 2 in the proof of Theorem 14, we have

$$\begin{aligned} d_{\text{tv}}(\hat{\sigma}_{i+1}^{(m)}, \pi_{g_{i+1}}) &\leq d_{\text{tv}}(\sigma_{i+1}^{(m)}, \pi_{g_{i+1}}) + d_{\text{tv}}(\hat{\sigma}_{i+1}^{(m)}, \sigma_{i+1}^{(m)}) \\ &\leq d_{\text{tv}}(\sigma_{i+1}^{(m)}, \pi_{g_{i+1}}) + m\tilde{\epsilon} + d_{\text{tv}}(\hat{\sigma}_{i+1}^{(0)}, \sigma_{i+1}^{(0)}) \\ &= d_{\text{tv}}(\sigma_{i+1}^{(m)}, \pi_{g_{i+1}}) + m\tilde{\epsilon} + 2d_{\text{tv}}(\hat{\sigma}_i^{(m)}, \pi_{g_i}) \\ &\leq \left(1 + \frac{1}{n \log 1/\rho} \right)^i \gamma \cdot \frac{1}{4n \log 1/\rho} + \left(1 + \frac{1}{n \log 1/\rho} \right)^i \gamma \cdot \frac{1}{4n \log 1/\rho} + \left(1 + \frac{1}{n \log 1/\rho} \right)^i \gamma \\ &\leq \left(1 + \frac{1}{n \log 1/\rho} \right)^{i+1} \gamma \end{aligned}$$

by choosing $m = \mathcal{O}^*(n^3)$ and

$$\tilde{\epsilon} = \frac{1}{m} \cdot \left(1 + \frac{1}{n \log 1/\rho} \right)^i \gamma \cdot \frac{1}{4n \log 1/\rho}.$$

Thus the final epoch, $i = \sqrt{n} \log 1/\rho$, the error is at most $\left(1 + \frac{1}{n \log 1/\rho} \right)^{\sqrt{n} \log 1/\rho} \gamma \leq e\gamma$. ■

Proof [Proof of Theorem 17] In (Kalai and Vempala, 2006), the authors proved the theorem for the case $f(x) = c \cdot x$ and claimed it can be extended to arbitrary convex functions. Here we

give a proof for the sake of completeness. [Kalai and Vempala \(2006\)](#) proved above inequality with arbitrary convex set K . Let's see how to relate an arbitrary convex function $f(x)$ to a linear function by increasing the dimension by 1. Consider a convex set $\mathcal{K} \in \mathbb{R}^n$ and a continuous convex function $f : \mathbb{R}^n \rightarrow \mathbb{R}$. Consider the epigraph

$$\tilde{\mathcal{K}} := \{(x, y) : x \in \mathcal{K}, y \geq f(x)\}.$$

Define the linear function $\tilde{f}(x, y) = c \cdot (x, y)$, where $(x, y) \in \tilde{\mathcal{K}}$ and $c = (\mathbf{0}, 1) \in \mathbb{R}^{n+1}$. We have

$$\mathbb{E}\tilde{f}(X, Y) \geq \mathbb{E}f(X) \quad \text{the first one increase mass on large values}$$

$$\min_{(x,y) \in \tilde{\mathcal{K}}} \tilde{f}(x, y) = \min_{x \in \mathcal{K}} f(x)$$

Thus

$$\mathbb{E}f(X) - \min_{x \in \mathcal{K}} f(x) \leq \mathbb{E}\tilde{f}(X, Y) - \min_{(x,y) \in \tilde{\mathcal{K}}} \tilde{f}(x, y) \leq (n+1)T$$

proof completed.

For the second claim, it is not hard to see that

$$\mathbb{E}_F f(X) - \min_{x \in \mathcal{K}} f(x) = \mathbb{E}_F \left[f(X) - \min_{x \in \mathcal{K}} f(x) \right].$$

Since adding a constant to the function does not have an effect on the density, we can assume without loss of generality that $\min_{x \in \mathcal{K}} f(x) = 0$. Thus we have

$$\begin{aligned} \mathbb{E}_F f(X) &= \frac{\int_{\mathcal{K}} f(x) \exp\{-F(x)/T\} dx}{\int_{\mathcal{K}} \exp\{-F(x)/T\} dx} \leq \frac{\int_{\mathcal{K}} f(x) \exp\{-f(x)/T\} dx \cdot \exp(\rho/T)}{\int_{\mathcal{K}} \exp\{-f(x)/T\} dx \cdot \exp(-\rho/T)} \\ &\leq \exp(2\rho/T) \cdot \mathbb{E}f(X) \leq (n+1)T \cdot \exp(2\rho/T) \end{aligned}$$

and

$$\mathbb{E}_F f(X) - \min_{x \in \mathcal{K}} f(x) \leq (n+1)T \cdot \exp(2\rho/T). \quad \blacksquare$$

Proof [Proof of Corollary 18] We choose $\rho = \epsilon/n$ and $T_K = \rho$. Given the final temperature, $K = \sqrt{n} \log(n/\epsilon)$. The optimization guarantee follows from Theorem 17. The number of queries is $\mathcal{O}^*(n^3)$ for one sample in one phase (Theorem 12) times $\mathcal{O}^*(n)$ samples per phase for rounding (Section A) times $K = \mathcal{O}^*(\sqrt{n})$ phases. The resulting distribution, however, is only $e\gamma$ -close to the distribution with density proportional to $\exp\{-nF/\epsilon\}$ (by Theorem 16). The guarantee of Theorem 17 holds for the latter distribution, and we need to upper bound the effect of having a sample from an almost-desired distribution. Thankfully, γ enters logarithmically in oracle complexity. Since f is L -Lipschitz and domain is bounded, the range of function values over \mathcal{K} is bounded by $B = \mathcal{O}(nLR)$. Then γ can be chosen as ϵ/B , which again only impacts oracle complexity by terms logarithmic in n, L, R . \blacksquare