

Faster Polytope Rounding, Sampling, and Volume Computation via a Sub-linear Ball Walk

Oren Mangoubi
*Worcester Polytechnic Institute
 Worcester, MA*

Nisheeth K. Vishnoi
*Yale University
 New Haven, CT*

Abstract—This paper studies the problem of “isotropically rounding” a polytope $K \subseteq \mathbb{R}^n$, that is, computing a linear transformation which makes the uniform distribution on the polytope have roughly identity covariance matrix. It is assumed that $K \subseteq \mathbb{R}^n$ is defined by m linear inequalities. We introduce a new variant of the ball walk Markov chain and show that, roughly, the expected number of arithmetic operations per-step of this Markov chain is $O(m)$ that is sub-linear in the input size mn – the per-step time of all prior Markov chains. Subsequently, we apply this new variant of the ball walk to obtain a rounding algorithm that gives a factor of \sqrt{n} improvement on the number of arithmetic operations over the previous bound which uses the hit-and-run algorithm. Since the cost of the rounding pre-processing step is in many cases the bottleneck in improving sampling or volume computation running time bounds, our results imply improved bounds for these tasks. Our algorithm achieves this improvement by a novel method of computing polytope membership, where one avoids checking inequalities which are estimated to have a very low probability of being violated. We believe that this method is likely to be of independent interest for constrained sampling and optimization problems.

Keywords-polytope rounding; volume computation; sampling; random walks

I. INTRODUCTION

The task of bringing a polytope into near-isotropic position is an important problem in mathematics and theoretical computer science (TCS). In TCS, this problem is closely linked with the widely studied problem of computing a polytope’s volume [6], [15], [8], [19], [4], and often serves as an important pre-processing step for these algorithms. For $a > 0$, we say that a convex body K is in a -isotropic position if the uniform distribution on K has covariance matrix Σ_K and mean μ_K satisfying

$$\frac{1}{a^2} I_n \preccurlyeq \Sigma_K \preccurlyeq a^2 I_n \quad \text{and} \quad \|\mu_K\|_2 \leq \frac{1}{10}a.$$

Formally, consider the following problem, where B denotes the unit ball centered at the origin:

Problem 1 (Bringing a polytope into isotropic position). *Given a polytope $K := \{x \in \mathbb{R}^n : Ax \leq b\}$, with $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$ such that $rB \subseteq K \subseteq RB$ for some $r, R > 0$, generate a matrix $\tilde{\Omega} \in \mathbb{R}^{n \times n}$ and vector $\tilde{\mu} \in \mathbb{R}^n$ such that $\tilde{K} := \tilde{\Omega}^{-\frac{1}{2}}(K - \tilde{\mu})$ is in 2-isotropic position.*

The problems of sampling from the uniform distribution on a polytope and of bringing a polytope into isotropic position are closely related. On the one hand, bringing a polytope into isotropic position can improve the running time of Markov chain-based sampling algorithms [14]. On the other hand, it is known that $n \log(n)$ independent samples from the uniform distribution on a polytope suffices to bring a polytope into $O(1)$ -isotropic position [22]. All current Markov chains which are used to sample from the uniform distribution on a polytope defined by m inequalities use at least mn arithmetic operations per Markov chain step to implement, and it is currently an open problem how to improve the number of arithmetic operations to fewer than mn [14]. mn is the size of the input and is also the time required to check whether a given point is in K or not.

The main focus of this paper is to develop Markov chains for sampling that allow us to bypass this mn barrier, and obtain faster algorithms for rounding, sampling, and volume computation. In particular, we introduce an implementation of the ball walk Markov chain [1], [8], which improves the expected number of arithmetic operations to roughly $O(m)$ operations per ball walk step. Our improvement in the per-step complexity applies in the special case when the polytope is in near-isotropic position, and we are given an $O(1)$ -warm start in the n^{-3} interior of the polytope.

We then apply a recent result of [14] which says that, starting at any point on the n^{-3} interior of K , the ball walk, together with a rejection sampling post-processing step from [8], can generate a sample from the uniform distribution on K in $O(n^{2.5} \log(\frac{1}{\varepsilon}))$ “proper” ball walk steps (that is, only counting the steps where the ball walk changes position). If X_0 is also $O(1)$ -warm, the expected number of steps (both proper and improper) is also $O(n^{2.5} \log(\frac{1}{\varepsilon}))$ [8]. Multiplying the two, we get that the expected number of operations for our implementation of the ball walk to generate a sample with TV error ε , is, roughly speaking¹, $O(m \times n^{2.5} \log(\frac{1}{\varepsilon}))$.

¹Since the two random variables may be correlated, we cannot simply multiply their expectations. Instead, we treat these two expectations separately until the very end of our proof, and then use Markov’s inequality to bound each with probability 9/10. We then multiply our bounds for the two random variables that we obtained from Markov’s inequality to get a bound for the product of the random variables which holds with probability 8/10.

Therefore, if we re-start the ball walk at the same $O(1)$ -warm initial point X_0 after we generate each sample, we can use our algorithm to generate p samples that are (conditionally on X_0) jointly independent and uniformly distributed on K with TV error ε , after roughly $O(pmn^{2.5} \log(\frac{1}{\varepsilon}))$ operations.

Using our implementation of the ball walk to generate $n \log(n)$ samples, we can use results from [22] to compute a sample mean and sample covariance matrix for the uniform distribution on K which allows us to bring any polytope K that is in 15-isotropic position into 2-isotropic position with probability $1 - \varepsilon$, in roughly $\tilde{O}(mn^{3.5} \log(\frac{1}{\varepsilon}))$ operations, if we are given an $O(1)$ -warm start X_0 , with X_0 in the n^{-3} interior of K . We use this idea in an iterative manner to obtain a rounding algorithm which can bring any polytope $rB \subseteq K \subseteq RB$ into 2-isotropic position with probability $1 - \varepsilon$. Specifically, we show the following:

Theorem I.1 (Main theorem: Bringing a polytope into isotropic position). *There exists an algorithm which, given a polytope $K := \{x \in \mathbb{R}^n : Ax \leq b\}$, with $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$ such that $rB \subseteq K \subseteq RB$ for some $r, R > 0$, and $\varepsilon > 0$, generates a matrix $\tilde{\Omega} \in \mathbb{R}^{n \times n}$ and vector $\tilde{\mu} \in \mathbb{R}^n$ such that the polytope $\tilde{K} := \tilde{\Omega}^{-\frac{1}{2}}(K - \tilde{\mu})$ is in 2-isotropic position with probability at least $1 - \varepsilon$, in $\tilde{O}(mn^{4.5} \text{polylog}(\frac{1}{\varepsilon}, \frac{R}{r}))$ arithmetic operations.*

To prove Theorem I.1 (see Theorem V.6 for the version of Theorem I.1 specific to our algorithm), we consider a sequence of convex bodies

$$K_i := K \cap (1 + 1/n)^i rB,$$

and bring these convex bodies into isotropic position sequentially, starting from $K_0 = rB$. We are able to do this since one can show that the same linear transformation which brings K_i into 2-isotropic position also brings K_{i+1} into 15-isotropic position. Since there are $n \log(\frac{R}{r})$ convex bodies K_i in the sequence, our algorithm brings K into 2-isotropic position in $\tilde{O}(mn^{4.5} \text{polylog}(\frac{1}{\varepsilon}, \frac{R}{r}))$ operations. Our rounding algorithm improves the best previous $\tilde{O}(mn^5 \text{polylog}(\frac{1}{\varepsilon}, \frac{R}{r}))$ bound of [19] for bringing a polytope into isotropic (or just “well-rounded”) position by a factor of \sqrt{n} , making progress on an open problem (see #4 in Section 8 of [4]). We get an improvement of \sqrt{n} rather than n since our bound for the number of operations per ball walk step needs the convex bodies K_i to be kept in isotropic position at each i (requiring us to generate $\tilde{\Theta}(n^2)$ independent samples), while [19] only need to keep their sequence of convex bodies in well-rounded position at each iteration (which they can do using only $\tilde{\Theta}(n)$ independent samples). On the other hand, each of our samples requires $n^{1.5}$ fewer operations per sample in expectation: we get a factor of n fewer operations from our improved bound on the expected number of operations per Markov chain step, and an additional factor of \sqrt{n} fewer operations because

the bound on the number of ball walk steps on isotropic convex bodies is smaller by a factor of \sqrt{n} than the bound for the hit-and-run Markov chain used in [19]. Our bound on the number of operations to put K in isotropic position is therefore smaller by a factor of \sqrt{n} compared to [19].

A. Application to volume computation

Bringing K into isotropic position allows us to then use the volume computation algorithm of [4] to compute the volume of K with error δ , in $\tilde{O}(\frac{mn^4}{\delta^2})$ operations after we pre-process K into isotropic position. Hence, starting with K far from isotropic position, we can compute the volume of K in roughly $\tilde{O}(mn^{4.5} \log(\frac{R}{r}) + \frac{mn^4}{\delta^2})$ operations:

Corollary I.2 (Computing the volume of a polytope). *There exists an algorithm which, given a polytope $K := \{x \in \mathbb{R}^n : Ax \leq b\}$, with $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$ such that $rB \subseteq K \subseteq RB$ for some $r, R > 0$, and $\varepsilon, \delta > 0$, computes with probability at least $1 - \varepsilon$ the volume of K up to a factor of $1 + \delta$ in $\tilde{O}(mn^{4.5} \text{polylog}(\frac{1}{\varepsilon}, \frac{R}{r}) + \frac{mn^4}{\delta^2} \text{polylog}(\frac{1}{\delta}, \frac{1}{\varepsilon}))$ arithmetic operations.*

In the regime where $\delta^{-1} = O(1)$ and $m \geq n^3$, the best current algorithm, which uses [19] for pre-processing and [4] for volume computation, gives a bound of $\tilde{O}(mn^5 \text{polylog}(\frac{1}{\varepsilon}, \frac{R}{r}) + \frac{mn^4}{\delta^2} \text{polylog}(\frac{1}{\delta}, \frac{1}{\varepsilon}))$ operations. Corollary I.2 improves this bound by a factor of \sqrt{n} . Moreover, since our result benefits from recent improvements in the bound on the Cheeger constant of an isotropic convex body, Corollary I.2 makes progress towards the open problem of connecting improved bounds on the Cheeger constant to faster volume computation (See section 2.2.3 of [14]). In Table I we give bounds for different algorithms which can be used to compute the volume.

B. Application to sampling

Pre-processing a polytope into isotropic position is also a bottleneck for the problem of sampling from the uniform distribution on the polytope. If we use our rounding algorithm (Theorem I.1) to bring K into 2-isotropic position, and then use the hit-and-run algorithm to generate a sample from the uniform distribution on K , we obtain a sample from the uniform distribution on K with TV error ε in $\tilde{O}(mn^{4.5} \text{polylog}(\frac{1}{\varepsilon}, \frac{R}{r}))$ arithmetic operations:

Corollary I.3 (Sampling from non-rounded polytope). *There exists an algorithm which, given a polytope $K := \{x \in \mathbb{R}^n : Ax \leq b\}$, with $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$ such that $rB \subseteq K \subseteq RB$ for some $r, R > 0$, and $\varepsilon > 0$, generates a sample uniformly distributed on K with TV error ε , in $\tilde{O}(mn^{4.5} \text{polylog}(\frac{1}{\varepsilon}, \frac{R}{r}))$ arithmetic operations.*

In the regime where $m \geq n^2$, the current best bound is $\tilde{O}(mn^5 \text{polylog}(\frac{1}{\varepsilon}, \frac{R}{r}))$ operations (if one uses a rounding pre-processing step from [19], and then the hit-and-run

Markov chain [18]) or, depending on the matrix multiplication exponent $2 < \omega \leq 3$, $\tilde{O}(mn^{\omega+2.5} \log(\frac{1}{\varepsilon}))$ for the John walk [3]. Corollary I.3 improves on [19] by a factor of \sqrt{n} , and on [3] by a factor of $n^{\omega-2}$. However, for smaller values of m algorithms such as Riemannian HMC [13] can be faster (see Table II).

Algorithm	# of arithmetic operations
Ball walk + rounding [8]	$mn^6\delta^{-2}$
Ball walk + rounding [11] + Gaussian cooling	$mn^{5.5} + mn^4\delta^{-2}$
Hit-and-run [19] + simulated annealing	$mn^5\delta^{-2}$
Riemannian HMC [13]	$m^2n^{\omega-1/3}\delta^{-2}$
Gaussian cooling [5]	$mn \times \max(n^2(R/r)^2, n^3)\delta^{-2}$
Algorithm 1&2 [our paper] + Gaussian cooling	$mn^{4.5} + mn^4\delta^{-2}$

Table I: Bounds on the number of arithmetic operations to compute the volume of a polytope K with $rB \subseteq K \subseteq RB$ (logarithmic factors of r, R, ε, d, m are not shown). Here ω is the matrix multiplication exponent, currently $\omega \approx 2.37$. (note: Gaussian cooling assumes (r, R) -well rounded, which is somewhat weaker than $rB \subseteq K \subseteq RB$)

Algorithm	# of arithmetic operations
Ball walk + rounding [11]	$mn^{5.5}$
Hit-and-run + rounding [19]	mn^5
Dikin walk [9]	$m^2n^{\omega+1}$
John walk [3] + Dikin walk initialization	$mn^{\omega+2.5}$
Geodesic walk* [12]	$(m^2n^{\omega-\frac{1}{4}} \log(\beta) \text{ assuming a } \beta\text{-warm start})$
Riemannian HMC* [13]	$(m^2n^{\omega-1/3} \log(\beta) \text{ assuming a } \beta\text{-warm start})$
Vaidya walk [3] + Dikin walk initialization	$m^{1.5}n^{\omega+3/2}$
Algorithm 1&2 [our paper]	$mn^{4.5}$

Table II: Bounds on the number of arithmetic operations to generate one sample from the uniform distribution on a polytope K with $rB \subseteq K \subseteq RB$ (logarithmic factors of r, R, ε, d, m are not shown). The matrix multiplication exponent ω is currently $\omega \approx 2.37$ for the best known matrix multiplication algorithm. (*Note: RHMC and Geodesic walk assume a warm start. While one can obtain a $(R/r)^n$ -warm start in our setting by initializing from the uniform distribution on rB , this causes the running times of RHMC and Geodesic walk to gain an additional factor of n .)

Note that in Table II the Dikin, Geodesic and Vaidya walks,

as well as Riemannian HMC, have bounds with dependence on m of at least $m^{1.5}$. In particular, in the regime $m > n^2$, they have slower bounds than the algorithms which have linear dependence on m , including our algorithm as well as the hit-and-run algorithm of [19] and the ball walk of [11]. Finally, note that, given the current bound $\omega \approx 2.37$, our algorithm has faster bounds compared to the John walk.

C. Key technical ideas

The algorithmic techniques we use in our variant of the ball walk are inspired from stochastic gradients, where one queries an oracle by subsampling. Instead of subsampling a small subset of component gradients, at each step of the ball walk we check a small subset of the inequalities defining our polytope. The challenge is in determining which inequalities are important at any given time. The main difficulty lies in the fact that the ball walk is much more likely to violate inequalities corresponding to nearby hyperplanes than far-away ones. The reason is that, if the Markov chain's steps are uniformly distributed, by the isoperimetric inequality [20] (and convexity of the polytope) the Markov chain will in expectation spend at least roughly half of its time a distance of $\frac{1}{n}$ from the boundary of the polytope, if the polytope is in near-isotropic position. Hence, in expectation, the Markov chain will be a distance of $\frac{1}{n}$ from one or more faces of the polytope at least half of the time and one cannot simply check a uniform random subset of these inequalities. Instead, our Markov chain estimates the distance to each inequality, and checks only those inequalities which have a non-negligible probability of being violated. This idea turns out to be challenging to execute and we expand on it in Section III. We believe our method of checking membership may be of independent interest to constrained sampling and optimization problems.

II. OUR ALGORITHMS

A. Notation

For any finite-volume subset $S \subseteq \mathbb{R}^n$, let Σ_S denote the covariance matrix of the uniform distribution on S and let μ_S denote the mean of the uniform distribution on S . Let π_K denote the uniform distribution on a convex body K . Let $\hat{\pi}_K \equiv \hat{\pi}_K^\eta$ denote the speedy distribution for step size η , where the speedy distribution is the stationary distribution of the proper steps of the ball walk with step size η . The Markov chain formed by the proper steps of the ball walk is called the speedy walk. A convex body K is said to be a -isotropic if

$$\frac{1}{a^2} I_n \preceq \Sigma_K \preceq a^2 I_n.$$

If, furthermore, $\|\mu_K\|_2 \leq \frac{1}{10}a$, we say that K is in a -isotropic position. A convex body K is said to be (r, \mathfrak{R}) -rounded if

$$\mathbb{E}_{X \sim \pi_K} [\|X\|_2^2] \leq \mathfrak{R}^2$$

and K contains a ball of radius r . A convex body which is $O(1)$ -isotropic is also $(O(1), O(\sqrt{n}))$ -rounded, although the converse is not true [17]. If a convex body is (r, \mathfrak{R}) -rounded with $\frac{\mathfrak{R}}{r} = O(\sqrt{n})$, we say it is *well-rounded*. For any subset $S \subseteq \mathbb{R}^n$ and any point $x \in \mathbb{R}^n$, let

$$\text{dist}(x, S) := \inf_{y \in S} \|x - y\|_2$$

denote the distance from x to the subset S and let ∂S denote the boundary of the subset S . Let $B(x, r)$ denote the Euclidean ball with center x and radius r , and denote the unit ball by the shorthand notation $B := B(0, 1)$. We say that a probability distribution $\mu : \mathbb{R}^n \rightarrow \mathbb{R}$ (or a random variable with distribution μ) is β -warm with respect to a probability distribution $\nu : \mathbb{R}^n \rightarrow \mathbb{R}$ if

$$\frac{\mu(x)}{\nu(x)} \leq \beta \quad \forall x \in \mathbb{R}^n.$$

We denote the probability distribution of a random variable X by $\mathcal{L}(X)$. Finally, we define the ball walk Markov chain $\tilde{X}_0, \tilde{X}_1, \dots$ on K with initial point $\tilde{X}_0 \in K$ and step size $\eta > 0$, by the recursion

$$\tilde{X}_{i+1} = \begin{cases} \tilde{X}_i + \eta \xi_i & \text{if } \tilde{X}_i + \eta \xi_i \in K \\ \tilde{X}_{i+1} = \tilde{X}_i & \text{otherwise,} \end{cases}$$

where ξ_0, ξ_1, \dots are iid uniform on the unit ball.

B. Algorithm 1 (Sampling)

Algorithm 1 generates independent samples approximately uniform on a convex body

$$K := \{x \in \mathbb{R}^n : Ax \leq b\} \cap \rho B,$$

for some $\rho > 0$. It has two main components: a Markov chain obtained as a subsequence of the ball walk which generates samples from the “speedy distribution”, and a rejection sampling method which obtains uniform samples from these “speedy distributed” samples. The “While” loop generates the ball walk X_1, \dots . To determine whether a ball walk proposal is inside K without checking all m inequalities at each step, every time that the algorithm checks an inequality $A_j x \leq b$, it stores in memory the distance h_j to the corresponding hyperplane H_j . The algorithm then waits $\frac{h_j}{\alpha \eta / \sqrt{n}}$ steps until re-computing that inequality, where $\alpha > 0$ is a parameter set by the user. The idea is that, since with high probability the ball walk makes steps of size $O(\eta / \sqrt{n})$ in the direction of the hyperplane H_j , the ball walk is unlikely to propose a step which crosses H_j before taking $\frac{h_j}{\alpha \eta / \sqrt{n}}$ steps. This allows us to ensure that our implementation of the ball walk remains inside K with high probability.

The ball walk is run until a fixed number of proper steps Y_s are made. The Markov chain Y_1, Y_2, \dots formed by the proper steps of the ball walk is called the “speedy walk”. Since fast mixing bounds are available for the speedy walk but not for the ball walk, we generate our samples

from the speedy walk, that is, we run the ball walk for a fixed number \mathcal{I} of proper steps. This gives us a sample $Y_{\mathcal{I}}$ approximately from the stationary distribution of the speedy walk. Unfortunately, the speedy walk does not have uniform stationary distribution; the samples from the speedy walk have a different distribution called the “speedy distribution”. To obtain uniformly distributed samples from our speedy-distributed samples we use a rejection sampling method, Algorithm 4.15 from [8] (reproduced in our Algorithm 1 as Steps 35-40). To obtain p independent samples, we run the speedy walk p times starting at the same initial point Y_0 but using different independent random Gaussian vectors ξ_i each time (this is the outer “For” loop). The parameter “Modified” can be set to “ON” or “OFF”. If it is “ON” we use our new implementation of the ball walk, while if it is “OFF” we use the usual implementation where all m inequalities are checked at each step. The purpose of the parameter “Modified” is only to simplify the exposition of our proofs; in practice we always set “Modified = ON.”

Algorithm 1 Sampling (Modified speedy walk)

input: $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$, $\rho > 0$, with $\frac{1}{10}B \subseteq K$ where $K := \{x \in \mathbb{R}^n : Ax \leq b\} \cap \rho B$
input: step size $\eta > 0$, tolerance $\alpha > 0$, $\mathcal{I} > 0$, maximum number i_{\max} of proper+improper steps
input: Initial point $Y_0 \in K$, Modified $\in \{\text{ON, OFF}\}$

- 1: **for** $j = 1$ to m **do** \triangleright Initialization: compute distance to all m hyperplanes
- 2: Set $h_j = b_j - A_j X_0$
- 3: **end for**
- 4: Sort the h_j ’s in increasing order, and denote the ordered set of h_j ’s by H
- 5: Set $i = 0$, $s = 0$, and $k = 1$
- 6: **while** $k \leq p$ **do**
- 7: Set $X_0 = Y_0$
- 8: **while** $s \leq \mathcal{I}$ and $i \leq i_{\max}$ **do**
- 9: Set $i \leftarrow i + 1$, and set $K = \text{True}$
- 10: Sample $\xi_i \sim \text{unif}(B(0, 1))$
- 11: Set $\hat{X}_{i+1} = X_i + \eta \xi_i$ \triangleright Ball walk proposal
- 12: **if** Modified = ON **then** \triangleright Determining which inequalities to check
- 13: Set j^* to be the largest value of j such that $h_j < \alpha \frac{\eta}{\sqrt{n}} \times i$
- 14: **for** $j = 1$ to j^* **do** \triangleright Compute distance to possibly-nearby hyperplanes
- 15: Set $h_j = b_j - A_j X_i$
- 16: Insert h_j into H such that H remains in increasing order
- 17: **if** $h_j < 0$ and $K = \text{True}$ **then**
- 18: Set $K = \text{False}$
- 19: **end if**
- 20: **end for**

```

21:      if  $\|\hat{X}_{i+1}\|_2 > \rho$  then  $\triangleright$  Check if in ball  $\rho B$ 
22:          Set  $K = \text{False}$ 
23:      end if
24:      else  $\triangleright$  Conventional membership query, if
         modifications “turned off”
25:          Check if  $\hat{X}_{i+1}$  satisfies all  $m$  inequalities and
         is in  $\rho B$ , and if not set  $K \leftarrow \text{False}$ 
26:      end if
27:      if  $K = \text{True}$  then
28:          Set  $Y_s = X_i$   $\triangleright$  Speedy walk step
29:          Set  $X_{i+1} = \hat{X}_{i+1}$ 
30:          Set  $s \leftarrow s + 1$ 
31:      else
32:          Set  $X_{i+1} = X_i$ 
33:      end if
34:      end while
35:      Set  $Z_k = \frac{2n}{2n-1}Y_s$ 
36:      if  $Z_k \in K$  then
37:          “accept”  $Z_k$  and set  $k \leftarrow k + 1$ 
38:      else
39:          “reject”  $Z_k$   $\triangleright$  rejection sampling to get uniform
         distribution
40:      end if
41:  end while
42: output: Samples  $Z_1, \dots, Z_p$  which are approximately
         uniformly distributed on  $K$ . Output these samples if  $i \leq$ 
 $i_{\max}$ . Otherwise, output “Failure”.

```

C. Algorithm 2 (Rounding)

Using Algorithm 1 as a subroutine, we obtain an algorithm (Algorithm 2) for bringing a polytope into isotropic position:

The goal of Algorithm 2 is to inductively bring into isotropic position a sequence of convex bodies

$$K_1 \subseteq K_2 \subseteq \dots, \quad \text{where } K_i := (1 + 1/n)^i r B \cap K,$$

starting with K_1 . At each iteration of the “For” loop, we use Algorithm 1 to generate samples from the uniform distribution on a convex body \hat{K}_i which is an affine transformation of K_i in 15-isotropic position obtained at the previous step of the “For” loop (Step 11). Using these samples, Algorithm 1 computes a sample covariance matrix and mean for \hat{K}_i (Steps 12 and 13), which allows it to compute an affine transformation that puts K_i into 2-isotropic position as “ \tilde{K}_i ” and K_{i+1} into 15-isotropic position as “ \hat{K}_{i+1} ”. Steps 15-21 generate a point X_0 which is $O(1)$ -warm with respect to the uniform distribution and in the n^{-3} -interior of \tilde{K}_i . The point X_0 , which is also $O(1)$ -warm with respect to the uniform distribution on \hat{K}_{i+1} and in its n^{-3} -interior, is then used in the next iteration of the “For” loop as an initial point when Algorithm 1 is used to generate samples from the uniform distribution on \hat{K}_{i+1} .

Algorithm 2 Rounding

input: $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$
input: $r, R > 0$ such that $rB \subseteq K \subseteq RB$, where
 $K := \{x \in \mathbb{R}^n : Ax \leq b\}$
input: $p \in \mathbb{N}$, $\varepsilon > 0$, Modified $\in \{\text{ON}, \text{OFF}\}$

- 1: Set $\hat{\Sigma}_0 = r^2 I_n$ and $\hat{\mu}_0 = 0$
- 2: Set $i^* = n \log_2(\frac{R}{r})$
- 3: **for** $k = 1$ to p **do**
- 4: Sample Z_k from the uniform distribution on B
- 5: **if** $\text{dist}(Z_k, \partial B) \geq n^{-3}$ **then**
- 6: Set $\hat{X}_0 = Z_k$
- 7: **end if**
- 8: **end for**
- 9: **for** $i = 1$ to $i^* - 1$ **do**
- 10: Define $K_i := (1 + 1/n)^i r B \cap K$ (just a definition,
 no computation here)
- 11: Use Algorithm 1 with parameter “Modified” and
 initial point X_0 to generate p points Z_1, \dots, Z_p
 approximately from the uniform distribution on
 $\hat{K}_i \cap 20\sqrt{n} \log(\frac{40n^2}{\varepsilon})B$, where $\hat{K}_i := \hat{\Sigma}_i^{-\frac{1}{2}}(K_i - \hat{\mu}_i)$
 and is represented by the inequalities with matrix
 $A\hat{\Sigma}_i^{\frac{1}{2}} \in \mathbb{R}^{m \times n}$, and vector $b - A\hat{\mu}_i \in \mathbb{R}^m$.
- 12: Set $\hat{\mu}_{i+1} = \frac{1}{p} \sum_{j=1}^p Z_j + \hat{\mu}_i$
- 13: Set $\hat{\Sigma}_{i+1} = [\frac{1}{p} \sum_{j=1}^p (Z_j - \hat{\mu}_i)^\top (Z_j - \hat{\mu}_i)]\hat{\Sigma}_i$
- 14: Set Interior = False
- 15: **while** Interior = False **do** \triangleright generate a starting
 point, uniform on the “ n^{-3} -interior” of \hat{K}_i
- 16: Use Algorithm 1 with parameter “Modified”
 and initial point X_0 to generate a single sample \hat{X}_0
 approximately from the uniform distribution on
 $\hat{K}_i \cap 20\sqrt{n} \log(\frac{40n^2}{\varepsilon})B$.
- 17: Set $\hat{X}'_0 \leftarrow \hat{\Sigma}_{i+1}^{-\frac{1}{2}}\hat{\Sigma}_i^{\frac{1}{2}}(\hat{X}_0 + \hat{\mu}_i - \hat{\mu}_{i+1})$
- 18: **if** $\text{dist}(\hat{X}'_0, \partial[\hat{K}_i \cap 20\sqrt{n} \log(\frac{40n^2}{\varepsilon})B]) \geq n^{-3}$,
 where $\tilde{K}_i := \hat{\Sigma}_{i+1}^{-\frac{1}{2}}(K_i - \hat{\mu}_{i+1})$ **then**
- 19: Set $X_0 = \hat{X}'_0$
- 20: Set Interior = True
- 21: **end if**
- 22: **end while**
- 23: **end for**
- 24: **output:** $\hat{\Sigma}_{i^*}, \hat{\mu}_{i^*}, X_0$

III. TECHNICAL OVERVIEW OF OUR MAIN RESULT: THEOREM I.1

A. Rounding polytopes via sampling

Most algorithms which bring a polytope into isotropic position work by generating independent or near-independent samples which are approximately uniformly distributed in the polytope. These samples allow one to compute the sample mean and sample covariance matrix for the polytope. In [22] it was shown that $n \log(n)$ samples suffice to bring

a polytope into isotropic position. However, one is still left with the problem of generating independent uniform samples from the polytope. Typically this is done by running a Markov chain on the polytope whose stationary distribution is equal (or in some sense close to) the uniform distribution in the polytope. However, the number of steps for which one must run the Markov chain to obtain uniform independent samples in many cases itself depends on the extent to which the polytope is isotropic, or the extent to which it is rounded.²

Because of this, most rounding algorithms start with the ball contained in the polytope (which is very easy to put in well-rounded position), and gradually deform the polytope at each iteration (for instance by considering the intersection of the polytope with a ball of increasing radius). One alternates between sampling steps where one samples from the convex body, and steps where one uses these samples to compute an affine transformation which keeps the convex body well-rounded (for example, this transformation can be achieved by computing the sample covariance matrix). For instance, this is the case for the “ball walk” Markov chain used in [8], and the “hit-and-run” Markov chain used in [19]. In particular, the algorithm of [19] requires only $n \log(n)$ samples to round the convex body. Key to this is that the hit-and-run Markov chain does not require isotropic position but rather only that the polytope be well-rounded. This fact was used in [19], together with a “pencil construction,” to provide a rounding algorithm where one computes a well-rounded polytope at $n \log(\frac{R}{r})$ iterations each using $\log(n)$ samples, and using n samples to bring the polytope into isotropic position only every $\log(n)$ iterations. Since the hit-and-run Markov chain requires n^3 steps to generate a uniform sample, they require roughly $n^4 \log(n) \log(\frac{R}{r})$ Markov chain steps to put a polytope into isotropic position. If the polytope is defined by m inequalities, the hit-and-run algorithm uses mn arithmetic operations to compute polytope membership at each step of the Markov chain, giving a bound of $mn \times n^4 \log(n) \log(\frac{R}{r})$ arithmetic operations to round the polytope; this is currently the fastest running time bound for rounding this class of polytopes in our setting.

This mn cost of computing each step of the Markov chain is a feature of all current Markov chain sampling algorithms on polytopes defined by m inequalities [14]. However, one can imagine that there may be ways of reducing the cost of computing polytope membership. One approach is to use a Markov chain called “coordinate hit-and-run” [7]. This algorithm works in the same way as the usual hit-and-run algorithm, except that it only takes steps in (random) coordinate directions. Hence, checking each inequality takes only $O(1)$ arithmetic operations, meaning that each step of

²While there are some algorithms such as the Dikin walk which do not depend on how isotropic the polytope is, these do not currently provide the fastest methods of bringing a polytope $rB \subseteq K \subseteq RB$ into isotropic position in our setting.

the Markov chain would roughly require only m operations. Unfortunately, since there are as of yet no polynomial-in-dimension mixing time bounds for coordinate hit-and-run, one cannot currently use coordinate hit-and-run to obtain better running time bounds for rounding.

B. A first attempt

As an alternative approach to coordinate hit-and-run one might consider using the stochastic gradient technique to reduce the cost of computing each step of a Markov chain which stays inside a polytope. For instance, one might attempt to apply stochastic gradients to a Markov chain such as the Dikin walk which, instead of computing polytope membership, makes use of the log-barrier function of the polytope to remain inside the polytope. The log-barrier at any point in the Markov chain is given by

$$\phi(x) = - \sum_{j=1}^m \log(A_j x - b_j),$$

with Hessian

$$\nabla^2 \phi(x) = - \sum_{j=1}^m \frac{A_j^\top A_j}{(A_j x - b_j)^2}$$

(here A_j is a row-vector). If m is large, one might try to estimate the gradient, or in the case of the Dikin walk the Hessian, of the log-barrier function by taking a small subset of the polytope’s inequalities and using these to estimate the sum. Unfortunately this “stochastic Hessian” gives a very bad approximation for points which are near a face of the polytope, since the term $\frac{1}{(A_j x - b_j)^2}$ corresponding to the nearest face can be much larger than the combined contributions of the terms corresponding to all the other faces, and a small subsample of the polytope inequalities will most likely not include the single very large term. Since a uniform random sampling is unlikely to include this overwhelmingly large term in the sum, this “stochastic Hessian” version of the Dikin walk is likely to very quickly leave the polytope.

One might instead consider an approach related to stochastic gradients, but for Markov chains such as the ball walk which, instead of computing a barrier function, compute polytope membership at each step. To determine polytope membership, one typically checks all m inequalities at each step of the Markov chain. One may instead consider checking only a small uniformly random subset of $\frac{m}{n}$ of these inequalities at each step. Unfortunately, this approach cannot work if one wishes to sample from the uniform distribution on a polytope. The reason is that, if the Markov chain’s steps are uniformly distributed, by the isoperimetric inequality [20] (and convexity of the polytope) the Markov chain will in expectation spend at least half of its time a distance of $\frac{1}{n}$ from the boundary of the polytope. Hence, in expectation, the Markov chain will be a distance

of $\frac{1}{n}$ from one face of the polytope at least half of the time. For a ball walk with optimal step size of $\eta = \Theta(\frac{1}{\sqrt{n}})$, with high probability the ball walk takes a step of size $\frac{1}{n}$ in the direction of this face. Since the ball walk may have an $\Omega(1)$ probability of proposing a step which violates the inequality corresponding to the closest face, if one only checks a small random subset of size $\frac{m}{n}$ of the inequalities one is likely to miss checking the inequality corresponding to this face, which in many cases would cause the Markov chain to leave the polytope with probability $\Omega(1)$ at each step (this is the case, for instance, if the polytope is a cube). Hence, we cannot limit ourselves to checking a random subsample of the inequalities.

C. Our method of computing polytope membership

Recall that the isoperimetric inequality [20] implies that any Markov chain whose stationary distribution is close to the uniform distribution on the polytope will spend on average at least $\frac{1}{n}$ of its time a distance $\frac{1}{n}$ from the boundary of the polytope. This suggests that if we wish to compute only a small subsample of the polytope inequalities at each step of the Markov chain, we must make sure that our subsample includes all those inequalities whose corresponding face is close to the current Markov chain step. While it is possible for us to compute the distance

$$h_j(i) := A_j X_i - b_j$$

between the current point in the Markov chain X_i to the hyperplane H_j corresponding to each inequality (A_j, b_j) , we do not wish to compute this distance $h_j(i)$ for each j at every step i in the Markov chain, since this takes the same mn operations needed to check each inequality. To get around this problem, rather than computing each distance, we instead only compute the distance to any given hyperplane at a small fraction of the steps. To determine which $h_j(i)$ should be computed at any given step, we estimate a high-probability lower bound $L_j(i)$ for $h_j(i)$ and only compute $h_j(i)$ if the Markov chain is likely to propose a step violating the j th inequality, that is, if this lower bound is

$$L_j(i) = O\left(\frac{1}{n}\right).$$

To estimate the lower bound on $h_j(i)$, we apply concentration inequalities to the steps of the ball walk. Specifically, we use concentration inequalities for spherical caps to show that with probability at least $1 - \frac{\varepsilon}{m}$ the ball walk will never take any steps of size more than $\frac{\eta}{\sqrt{n}} \log(\frac{i_{\max}}{\varepsilon})$ in the direction of H_j , if the ball walk is run for at most i_{\max} steps (i_{\max} is a “cutoff time” parameter which can be set by the user of Algorithm 1; if the algorithm takes more than i_{\max} steps, the algorithm terminates without outputting any samples).

Hence, if we set our high-probability lower bound to be

$$L_j(i) = h_j(i^*(i)) - (i - i^*(i)) \times \frac{\eta}{\sqrt{n}} \log\left(\frac{i_{\max}}{\varepsilon}\right),$$

where $i^*(i)$ is the last time before step i that the distance to H_j was computed, then with probability at least $1 - \varepsilon$ the ball walk will never leave the polytope (Lemma IV.1, and Step 13 in Algorithm 1).

D. Using anti-concentration to prove expected running time bounds for our algorithm

Even though we have shown that our algorithm obtains the exact same samples as the usual implementation of the ball walk with high probability, we still have to show that it reduces the number of inequalities one has to check at each step of the Markov chain. Towards this end, we prove an anti-concentration inequality (Lemma IV.3) for the uniform distribution on a convex body to show that the expected number of inequalities our algorithm checks at any given step is roughly $\frac{m}{n}$. Roughly, this inequality says that a uniform random point on an isotropic convex body has probability at most $O(\Delta)$ of being within a distance Δ of any given codimension-1 hyperplane. The main obstacle in applying the anti-concentration inequality is that, while we are able to guarantee that a random step of the Markov chain is $O(1)$ -warm with respect to the uniform distribution on K , if we are to only check a small fraction $\frac{1}{n}$ of the inequalities (in expectation) at each step, the steps where one checks the inequality cannot be uniformly distributed or even $O(1)$ -warm with respect to the uniform distribution but instead can only be $\Theta(n)$ -warm at best. This is because one has to ensure that, if the Markov chain is within a distance of one ball walk step from a given inequality, this inequality will be checked with very high probability, but only with probability $O(\frac{1}{n})$ at a “typical” step.

To get around this problem we instead bound the waiting time between any given step i of the Markov chain and the most recent time $i^*(i)$ that the distance to the hyperplane was computed by the algorithm, as a function of the distance to this hyperplane at the current step i of the Markov chain, allowing us to apply the anti-concentration inequality. Specifically, we use our bound on the size of the ball walk’s steps in the direction of H_j to show that $i - i^*(i)$ must be greater than $\frac{h_j(i)\sqrt{n}}{\eta \log(\frac{m i_{\max}}{\varepsilon})}$. Since $\eta = \Theta(\frac{1}{\sqrt{n}})$, applying the anti-concentration inequality on $h_j(i)$ gives a lower bound on $\mathbb{E}[i - i^*(i)]$ that is roughly equal to n , if we are given an $O(1)$ -warm start with respect to the uniform distribution. This in turn implies that given an $O(1)$ -warm start, we compute the distance to any given hyperplane at only a small fraction $\frac{1}{n}$ of the ball walk steps in expectation (Lemma IV.5). Hence, rather than taking mn operations, our algorithm is able to compute each step of the ball walk in only m operations in expectation, an improvement by a factor of n .

E. Improved bounds for Rounding a polytope

Unfortunately, we cannot use the rounding meta-algorithm of [19], which requires only $n \log(n)$ samples to round the polytope, with our implementation of the ball walk. The reason is that, while the rounding meta-algorithm in [19] keeps a sequence of convex bodies well-rounded at each of the n iterations, it only keeps the convex bodies in isotropic position at a small number $\log(n)$ of the iterations. To obtain our bounds on the expected frequency at which one needs to compute the distance to each H_j , we must make sure that the convex body is in isotropic position at each iteration; our rounding algorithm (Algorithm 2) uses $n^2 \log(n)$ samples instead of $n \log(n)$ (see the next paragraph for a discussion of the rounding algorithm). Hence, it would seem at first that, despite the fact that we improve the expected number of arithmetic operations at each step of the Markov chain by a factor of n , this improvement would be offset by the fact that we need n times as many samples, which we require to keep the polytope in isotropic position at every iteration. However, there are additional benefits to keeping a polytope in isotropic position. In particular, recent improvements towards weaker versions of the KLS conjecture imply that the best current bound on the mixing time of the proper steps of the ball walk (also called the “speedy walk”) also improves to $n^{2.5}$ by a factor of \sqrt{n} if the polytope is in isotropic position as opposed to the roughly n^3 mixing time bound available for the hit-and-run Markov chain when the polytope is well-rounded but not $O(1)$ -isotropic [11]. Hence, keeping the convex body in isotropic position allows us to combine our factor of n improvement with the \sqrt{n} improvement in the mixing time from [11]. The number of arithmetic operations to round a convex body is then at most roughly $mn^{4.5}$, an improvement of \sqrt{n} over the mn^5 bound of [19]. Note that it is enough to bound the number of arithmetic operations in expectation, since one can always start over if the rounding algorithm takes more than its expected number of steps.

F. Rounding a polytope by sampling from isotropic position (Algorithm 2)

More specifically, in order to efficiently generate the samples needed to bring a polytope into 2-isotropic position one should first ensure that the polytope from which one samples is, say, in 15-isotropic position. Towards this end, one can consider a sequence of nested convex bodies

$$K_i := K \cap (1 + 1/n)^i rB.$$

The initial polytope $K_0 = rB$ is just the ball contained inside K , which can be brought into isotropic position by multiplying this ball by $\frac{1}{r\sqrt{d}}$. Since the diameter of K_i increases by a factor of only $1 + 1/n$, at each step, one can show that the volume of these convex bodies does not increase by more than e at each step i , and that for any transformation

that brings K_{i-1} into 2-isotropic position, applying the same transformation to K_i would bring it into 15-isotropic position. This suggests an iterative algorithm (Algorithm 2), where one samples from a 15-isotropic convex body which is a linear transformation of K_{i-1} , allowing one to bring K_{i-1} into 2-isotropic position. The same transformation brings K_i into 15-isotropic position, allowing one to iteratively bring the sequence of polytopes into 2-isotropic position by alternating between sampling and linear transformation steps (Lemma V.2). This takes $n \log(\frac{R}{r})$ iterations to bring the polytope K into isotropic position, and uses $n \log(n)$ samples at each iteration; the number of samples to round the polytope is then roughly $n^2 \log(n) \log(\frac{R}{r})$. To allow us to apply Lemma IV.5 and bound the expected fraction of the time that our implementation of the ball walk (Algorithm 1) checks a given inequality, we must still show the ball walk has a warm start at each iteration of Algorithm 2. We can obtain a warm start for K_i by using a sample from the $i-1$ iteration which is approximately uniformly distributed on K_{i-1} (Steps 15-21 of Algorithm 2). Since K_{i-1} contains at least $\frac{1}{e}$ of the volume of K_i , a uniformly distributed point on K_{i-1} provides us with an $O(1)$ -warm start for K_i (Lemma V.4).

Remark III.1. *If the KLS conjecture is proved true, the mixing time bound of the speedy walk on convex bodies in isotropic position would decrease by an additional factor of \sqrt{n} to just roughly n^2 , potentially allowing us to improve our running time to roughly mn^4 . On the other hand, as noted in [11], it is not known how to connect improvements in the KLS conjecture to the current-best rounding algorithm which uses hit-and-run from a well-rounded but not isotropic position [19]. We note, however, that since further improvements to KLS would only apply to the ball walk from a warm start, even using our method, one would have to find a way to modify our rounding algorithm to allow it to use approximately-independent samples from a warm start rather than fully-independent samples from a cold start.*

G. Organization of the rest of the paper

In the rest of the paper, we prove the main result (Theorem I.1, proved for our specific algorithm as Theorem V.6), and its corollaries I.2 and I.3.

In Section IV we bound the accuracy of Algorithm 1 and (roughly speaking) the expected number of arithmetic operations it performs when it is used to sample from a polytope in $O(1)$ -isotropic position. In Section IV-A we bound the probability our implementation of the ball walk leaves the polytope. In Section IV-B we prove an anti-concentration bound, which we use in Section IV-C to bound the expected frequency at which our implementation of the ball walk checks a given inequality. In Sections IV-D and IV-E we recall results from [8] and [11] which allow us to then bound the mixing time of the speedy walk (the proper

steps of the ball walk) and the expected number of steps.

In Section V we bound the success probability and expected number of arithmetic operations of the rounding algorithm (Algorithm 2). In Section V-A we bound the success probability of Algorithm 2. In Section V-B we bound the expected number of arithmetic operations made by Algorithm 2 under the assumption that it provides a warm start to the ball walk subroutine (Algorithm 1) at each iteration of Algorithm 2, and in Section V-C we show that this warm start assumption holds. In Section V-D we verify that the running time of steps where one does not check inequalities have only negligible contribution to the running time of Algorithm 2. In Section V-E we combine these results to complete the proof of our main theorem for rounding (Theorem V.6).

In Section VI we prove our results for volume computation (Corollary I.2) and sampling (Corollary I.3) for polytopes which may be far from isotropic position.

IV. SAMPLING FROM AN ISOTROPIC POSITION.

A. Bounding the distance traveled in any direction

In this section we bound the distance traveled by the Markov chain in the direction orthogonal to the plane H_j after i steps.

Lemma IV.1. *Fix $\hat{\varepsilon} > 0$ and suppose that $\alpha \geq 4 \log(\frac{2m i_{\max}}{\hat{\varepsilon}})$ in Algorithm 1. Then with probability at least $1 - \hat{\varepsilon}$ we have that, given the same random vectors ξ_i , the output of Algorithm 1 is the same regardless of whether we set Modified = ON or Modified = OFF.*

Proof: By the concentration inequality for spherical caps [10], for $\xi \sim \text{uniform}(B(0, 1))$, for the j th row A_j of the matrix A we have

$$\mathbb{P}\left(\left|A_j \frac{\xi}{\|\xi\|_2}\right| \geq t\right) \leq 2e^{-(n-2)t^2/2},$$

and hence

$$\mathbb{P}\left(\left|\eta A_j \frac{\xi}{\|\xi\|_2}\right| \geq \frac{\eta t}{\sqrt{n}}\right) \leq 2e^{-t^2/2}.$$

Thus,

$$\mathbb{P}\left(\left|\eta A_j \frac{\xi}{\|\xi\|_2}\right| \geq \frac{\eta}{\sqrt{n}} 2 \log\left(\frac{2}{\delta}\right)\right) \leq \delta,$$

for every $\delta > 0$. Therefore, for every $\hat{\varepsilon} > 0$,

$$\begin{aligned} & \mathbb{P}\left(\sum_{\ell=1}^i \left|\eta A_j \frac{\xi_\ell}{\|\xi\|_2}\right| \geq i \times \frac{\eta}{\sqrt{n}} 2 \log\left(\frac{2i}{\hat{\varepsilon}}\right)\right) \\ & \leq \mathbb{P}\left(\left|\eta A_j \frac{\xi_\ell}{\|\xi\|_2}\right| \geq \frac{\eta}{\sqrt{n}} 2 \log\left(\frac{2i}{\hat{\varepsilon}}\right) \text{ for some } \ell \in [i]\right) \\ & \leq i \times \frac{\hat{\varepsilon}}{i} = \hat{\varepsilon}. \end{aligned}$$

Hence, for our implementation of the ball walk X_1, X_2, \dots in Algorithm 1 we have

$$\begin{aligned} & \mathbb{P}\left(\sup_{i \leq \ell \leq k} |A_j X_\ell - A_j X_i| \geq (k-i) \frac{2\eta}{\sqrt{n}} \log\left(\frac{2i_{\max}}{\hat{\varepsilon}}\right)\right. \\ & \quad \left. \text{for some } 0 \leq i \leq k \leq i_{\max}\right) \\ & \leq \mathbb{P}\left(\left|\eta A_j \frac{\xi_\ell}{\|\xi\|_2}\right| \geq \frac{2\eta}{\sqrt{n}} \log\left(\frac{2i_{\max}}{\hat{\varepsilon}}\right) \text{ for some } \ell \in [i_{\max}]\right) \\ & \leq i_{\max} \times \frac{\hat{\varepsilon}}{i_{\max}} \\ & \leq \hat{\varepsilon}. \end{aligned}$$

Thus, if $\alpha \geq 2 \log\left(\frac{2m i_{\max}}{\hat{\varepsilon}}\right)$, we have

$$\begin{aligned} & \mathbb{P}\left(\sup_{i \leq \ell \leq k} |A_j X_\ell - A_j X_i| \geq (k-i) \frac{\eta}{\sqrt{n}} \alpha\right. \\ & \quad \left. \text{for some } 0 \leq i \leq k \leq i_{\max}\right) \\ & \leq \frac{\hat{\varepsilon}}{m}. \end{aligned} \quad (1)$$

Let Z_i be the usual ball walk Markov chain (where we check every inequality at each step) which evolves according to the following update equations:

$$Z_0 = X_0$$

$$Z_{i+1} = \begin{cases} Z_i + \eta \xi_i & \text{if } Z_i + \eta \xi_i \in K \\ Z_i & \text{otherwise.} \end{cases}$$

Then inequality (1) implies that if we set $\alpha \geq 2 \log\left(\frac{2m i_{\max}}{\hat{\varepsilon}}\right)$ in Algorithm 1, then with probability at least $1 - \hat{\varepsilon}$ we have that $X_i = Z_i$ for every $i \in [i_{\max}]$.

Therefore, with probability at least $1 - \hat{\varepsilon}$, the output of Algorithm 1 is the same regardless of whether we set Modified = ON or Modified = OFF. \blacksquare

Remark IV.2. *We have to bound the sum of the absolute value of the distance, rather than the sum of the variance, since the rejection step could introduce a bias (for instance, if the Markov chain is traveling along the face of a polytope).*

B. Anti-concentration bounds for isotropic convex bodies

Lemma IV.3. *Suppose that the uniform distribution π_K on K has identity covariance matrix (that is, K is 1-isotropic) and that $X \sim \pi_K$ is a random vector uniformly distributed on K . Let H be any codimension-1 hyperplane. Then we have*

$$\mathbb{P}(\text{dist}(X, H) \leq \hat{\varepsilon}) \leq 2\hat{\varepsilon} \quad \forall \hat{\varepsilon} > 0.$$

Proof: Let $A \in \mathbb{R}^{1 \times n}$ be a row vector and $b \in \mathbb{R}$ a real number such that $Ax = b$ is the equation for the Hyperplane H . Let $X \sim \pi_K$ be a random vector uniformly distributed on K . Denote the distribution of AX by π_K^A . Note that π_K^A is a marginal distribution of π_K . First, we note two facts:

- 1) All marginals of a logconcave distribution are logconcave (Theorem 2.2 of [23]).
- 2) If the covariance matrix of any distribution π satisfies

$$\sigma_1 I_n \preceq \Sigma \preceq \sigma_2 I_n,$$

then the variance $\mathbf{A}\Sigma\mathbf{A}^\top$ of its marginal π^A in the subspace defined by \mathbf{A} satisfies

$$\sigma_1 \leq \mathbf{A}\Sigma\mathbf{A}^\top \leq \sigma_2.$$

By the above facts, we have that the distribution of $\mathbf{A}X \in \mathbb{R}$ is isotropic (i.e., it has variance 1) and is logconcave. Let x^* be a maximizer of π_K^A . By Lemma 5.5(a) in [16], we have

$$\pi_K^A(x^*) \leq 1. \quad (2)$$

Hence, for any $\hat{\varepsilon} > 0$ we have

$$\begin{aligned} \mathbb{P}(|\mathbf{A}X - b| \leq \hat{\varepsilon}) &= \int_{b-\hat{\varepsilon}}^{b+\hat{\varepsilon}} \pi_K^A(x) dx \\ &\leq \int_{b-\hat{\varepsilon}}^{b+\hat{\varepsilon}} \pi_K^A(x^*) dx \\ &\stackrel{(Eq. 2)}{\leq} \int_{b-\hat{\varepsilon}}^{b+\hat{\varepsilon}} 1 dx \\ &= 2\hat{\varepsilon}. \end{aligned}$$

■

Remark IV.4. Note that the bounds in [21] are more general than what we need since they apply to hyperplanes of any codimension. We only care about codimension-1 hyperplanes, and can reduce the problem to obtaining anti-concentration bounds for a 1-dimensional isotropic logconcave distribution. This allows us to get a tight bound in Lemma IV.3 without assuming the KLS conjecture. This bound is tight (up to a universal constant) since it is tight for the special case of the unit cube and the regular simplex.

C. Bounding the frequency of constraint checking

To simplify notation, define $\hat{\eta} := \frac{1}{16}\eta\sqrt{n}$, and $\gamma := 10\alpha\hat{\eta}$.

Lemma IV.5. Suppose that K contains a ball of radius $r = \frac{1}{10}$. Fix $\hat{\varepsilon} > 0$ and set the algorithmic parameter $\alpha \geq 4\log(\frac{2i_{\max}}{\hat{\varepsilon}})$. Consider any row A_j of A and entry b_j of b . Suppose that the initial point X_0 is β -warm with respect to the uniform distribution for some $\beta > 0$. Let N_j be the number of steps (excluding the first step) of the Markov chain in Algorithm 1 with Modified = ON at which the algorithm checks inequality (A_j, b_j) and let N be the number of Markov chain steps. Let $F_j := \frac{N_j}{N}$ be the frequency of checking this inequality (excluding the first check). Then

$$\mathbb{E}[F_j] \leq 16n^{-1}\gamma\beta + \frac{32\gamma}{n} \times \beta \log(n/\gamma) + \frac{1}{N}\beta\hat{\varepsilon}.$$

Proof: First, we note that the stationary distribution of the steps of the ball walk (including improper and proper steps) is uniform on K . Let

$$X = X_0, X_1, X_2, \dots, X_N$$

be the Markov chain generated by Algorithm 1 with Modified = ON, initial point $X_0 = Y_0$, and random vectors ξ_1, \dots . Let

$$\tilde{X} = \tilde{X}_1, \tilde{X}_2, \dots, \tilde{X}_N$$

be the first N steps of the Markov chain generated by Algorithm 1 with Modified = OFF, using the same initial point $\tilde{X}_0 = Y_0$ and the same random vectors³. Using the same initial point and random vectors defines a coupling between X and \tilde{X} . Then $X_k = \tilde{X}_k$ for all k if and only if $X_k \in K$ for all $k \leq N$.

Let G be the event that

$$\sup_{i \leq \ell \leq k} |A_j \tilde{X}_\ell - A_j \tilde{X}_i| < (k-i) \times \frac{\eta}{\sqrt{n}}\alpha$$

for all $0 \leq i \leq k \leq i_{\max}$ and all $j \in [m]$. By Equation (1) in the proof of Lemma IV.1, we have that

$$\mathbb{P}(G) \geq 1 - \beta\hat{\varepsilon}.$$

Also by the proof of Lemma IV.1 we have that $X_k \in K$ for all $k \leq N$ if G occurs. Hence,

$$\mathbb{P}(X_k = \tilde{X}_k \forall k \leq i_{\max}) \geq \mathbb{P}(G). \quad (3)$$

Recall that $h_j(i) := b_j - A_j X_i$ is the distance from the Markov chain X_i to the hyperplane corresponding to the inequality (A_j, b_j) at step i . Rather than checking the inequality (A_j, b_j) at each step of the ball walk, Algorithm 1 waits some number of steps $w(i)$ after checking this inequality at some step i . More generally, we define

$$w(i) := \max\left(\left\lfloor \frac{\sqrt{n}}{\alpha\eta} h_j(i) \right\rfloor, 1\right)$$

regardless of whether the inequality is actually checked at step i (we can think of $w(i)$ as the amount of time the algorithm *would* have waited if it had checked the inequality at step i).

Let $i^*(k)$ be the step at which the inequality is checked for the k th time. Let $\mathbb{k}(i)$ be the number of times the inequality has been checked after i Markov chain steps (in particular, we have $\mathbb{k}(i) \leq i$). Let

$$i^*(i) := i(\mathbb{k}(i))$$

be the last time the inequality was checked.

³We consider all steps $\tilde{X}_k = \tilde{X}_{\tilde{N}}$, for all $k \geq \tilde{N}$, where \tilde{N} is the number of Markov chain steps computed by the algorithm. That is, the Markov chain remains stuck forever at the same point after Algorithm 1 halts.

Then, if the Markov chain X does not leave K , the total number N_j of times the inequality is checked is

$$\begin{aligned} N_j &= \sum_{k=1}^{N_j} w(\mathbb{i}(k)) \\ &= \sum_{i=1}^N \frac{1}{w(\mathbb{i}^*(i))} \quad (\text{If } G \text{ does not occur}), \end{aligned} \quad (4)$$

where the first equality holds because $w(\mathbb{i}(k)) = 1$ for all k . Therefore, we have

$$\mathbb{E}[N_j] \leq \sum_{i=1}^N \mathbb{E} \left[\frac{1}{w(\mathbb{i}^*(i))} \right] + \mathbb{P}(G^c). \quad (5)$$

We will show that, if G occurs, then

$$w(i-s) \geq \frac{1}{4}w(i) \quad \forall 0 \leq s \leq \frac{1}{8}w(i).$$

Suppose that G occurs. Without loss of generality we may assume that $w(i) > 2$ (since otherwise we have $w(i-s) > 1 > \frac{1}{4}w(i)$). Then

$$\begin{aligned} h_j(i-s) &\geq h_j(i) - s \times \frac{\eta}{\sqrt{n}}\alpha \\ &\geq h_j(i) - \frac{1}{8}w(i) \times \frac{\eta}{\sqrt{n}}\alpha \\ &= h_j(i) - \frac{2}{8} \frac{\sqrt{n}}{\alpha\eta} h_j(i) \times \frac{\eta}{\sqrt{n}}\alpha \\ &= \frac{3}{4}h_j(i). \end{aligned} \quad (6)$$

Therefore,

$$\begin{aligned} w(i-s) &= \max \left(\left\lfloor \frac{\sqrt{n}}{\alpha\eta} h_j(i-s) \right\rfloor, 1 \right) \\ &\stackrel{(Eq. 6)}{\geq} \max \left(\left\lfloor \frac{\sqrt{n}}{\alpha\eta} \frac{3}{4}h_j(i) \right\rfloor, 1 \right) \\ &\geq \max \left(\frac{\sqrt{n}}{\alpha\eta} \frac{3}{4}h_j(i), 1 \right) - 1 \\ &\geq \frac{3}{4} \max \left(\frac{\sqrt{n}}{\alpha\eta} h_j(i), 1 \right) - 1 \\ &\geq \frac{3}{4}w(i) - 1 \\ &\geq \frac{1}{4}w(i), \end{aligned}$$

where the last inequality holds since we assumed without loss of generality that $w(i) > 2$. Therefore whenever G occurs we have

$$w(i-s) \geq \frac{1}{4}w(i) \quad \forall i, s \in \mathbb{Z}^+, s \in \left[0, \frac{1}{8}w(i)\right]. \quad (7)$$

Suppose (towards a contradiction) that $w(\mathbb{i}^*(i)) < \frac{1}{8}w(i)$. But we always have

$$\mathbb{i}^*(i) + w(\mathbb{i}^*(i)) > i,$$

since $w(\mathbb{i}^*(i))$ is the amount of time we wait to check the inequality after step $\mathbb{i}^*(i)$, and, by definition of $\mathbb{i}^*(i)$ we have not yet re-checked the inequality at step i . Hence we would have

$$i - \mathbb{i}^*(i) < w(\mathbb{i}^*(i)) < \frac{1}{8}w(i).$$

Then by Inequality 7 we would have

$$w(\mathbb{i}^*(i)) \geq \frac{1}{4}w(i),$$

which contradicts our assumption that $w(\mathbb{i}^*(i)) < \frac{1}{8}w(i)$. Therefore, by contradiction we have that

$$w(\mathbb{i}^*(i)) \geq \frac{1}{8}w(i) \quad \forall i \in \mathbb{Z}^+. \quad (8)$$

Hence, combining Equations (5) and (8) we have

$$\begin{aligned} \mathbb{E}[N_j] &\stackrel{(Eq. 5)}{\leq} \sum_{i=1}^N \mathbb{E} \left[\frac{1}{w(\mathbb{i}^*(i))} \right] + \mathbb{P}(G^c) \\ &\stackrel{(Eq. 8)}{\leq} 8 \sum_{i=1}^N \mathbb{E} \left[\frac{1}{w(i)} \right] + \mathbb{P}(G^c). \end{aligned} \quad (9)$$

Therefore it is enough to bound $\mathbb{E}[\frac{1}{w(i)}]$ for each i .

Bounding $\mathbb{E}[\frac{1}{w(i)}]$: Fix any $i \in [N]$. First, we note that without loss of generality we may assume that $X_0 \equiv \tilde{X}_0$ is a 1-warm start, since the bound on $\mathbb{E}[F_j]$ for the β -warm case for general $\beta \geq 1$ will be at most β times as large as the bound for the 1-warm special case.

In the special case where $X_0 \equiv \tilde{X}_0$ is a 1-warm start, $\tilde{X}_i \sim \pi_K$ is uniformly distributed on K . Then by Lemma IV.3 we have

$$\begin{aligned} \mathbb{E} \left[\frac{1}{w(i)} \right] &= \mathbb{E} \left[\frac{1}{\max(\lfloor \frac{\sqrt{n}}{\alpha\eta} h_j(i) \rfloor, 1)} \right] \\ &= \mathbb{E} \left[\frac{1}{\max(\lfloor \frac{n}{10\alpha\hat{\eta}} h_j(i) \rfloor, 1)} \right] \\ &\leq 1 \times \mathbb{P} \left(h_j(i) \leq \frac{10\alpha\hat{\eta}}{n} \right) \\ &\quad + 2\mathbb{E} \left[\frac{10\alpha\hat{\eta}}{nh_j(i)} \times \mathbb{1}\{h_j(i) \geq \frac{10\alpha\hat{\eta}}{n}\} \right] \\ &\stackrel{(Lemma IV.3)}{\leq} 2n^{-1}\gamma \\ &\quad + 2\mathbb{E} \left[\frac{\gamma}{nh_j(i)} \times \mathbb{1}\{h_j(i) \geq \frac{\gamma}{n}\} \right], \end{aligned} \quad (10)$$

where $\gamma := 10\alpha\hat{\eta}$.

But

$$\begin{aligned} &\mathbb{E} \left[\frac{\gamma}{nh_j(i)} \times \mathbb{1}\{h_j(i) \geq \frac{\gamma}{n}\} \right] \\ &= \int_0^\infty \mathbb{P} \left(\frac{\gamma}{nh_j(i)} \times \mathbb{1}\{h_j(i) \geq \frac{\gamma}{n}\} \geq t \right) dt \\ &= \int_0^\infty \mathbb{P} \left(\frac{\gamma}{nh_j(i)} \times \mathbb{1}\{1 \geq \frac{\gamma}{nh_j(i)}\} \geq t \right) dt \end{aligned} \quad (11)$$

$$\begin{aligned}
&= \int_0^1 \mathbb{P} \left(\frac{\gamma}{nh_j(i)} \times \mathbb{1}\{1 \geq \frac{\gamma}{nh_j(i)}\} \geq t \right) dt \\
&\leq \int_0^1 \mathbb{P} \left(\frac{\gamma}{nh_j(i)} \geq t \right) dt \\
&= \int_0^1 \mathbb{P} \left(\frac{\gamma}{tn} \geq h_j(i) \right) dt \\
&= -\frac{\gamma}{n} \int_{\infty}^{\frac{\gamma}{n}} \mathbb{P}(u \geq h_j(i)) u^{-2} du \\
&= \frac{\gamma}{n} \int_{\frac{\gamma}{n}}^{\infty} \mathbb{P}(u \geq h_j(i)) u^{-2} du \\
&= \frac{\gamma}{n} \int_{\frac{\gamma}{n}}^1 \mathbb{P}(u \geq h_j(i)) u^{-2} du \\
&\quad + \frac{\gamma}{n} \int_1^{\infty} \mathbb{P}(u \geq h_j(i)) u^{-2} du \\
&\stackrel{(\text{Lemma IV.3})}{\leq} \frac{\gamma}{n} \int_{\frac{\gamma}{n}}^1 u^1 \times 2 \times u^{-2} du \\
&\quad + \frac{\gamma}{n} \int_1^{\infty} 1 \times u^{-2} du \\
&= \frac{2\gamma}{n} \int_{\frac{\gamma}{n}}^1 u^{-2+1} du + \frac{\gamma}{n} \\
&= \frac{2\gamma}{n} \times \log(n/\gamma) + \frac{\gamma}{n} \\
&\leq \frac{4\gamma}{n} \times \log(n/\gamma).
\end{aligned}$$

Hence, combining Inequalities (10) and (11) we have

$$\mathbb{E} \left[\frac{1}{w(i)} \right] \leq 2n^{-1}\gamma + \frac{4\gamma}{n} \times \log(n/\gamma). \quad (12)$$

Thus, by Inequality (9) we have

$$\begin{aligned}
\mathbb{E}[F_j] &= \mathbb{E} \left[\frac{N_j}{N} \right] \quad (13) \\
&\stackrel{(\text{Eq. 9})}{\leq} 8 \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[\frac{1}{w(i)} \right] + \frac{1}{N} \mathbb{P}(G^c) \\
&\stackrel{(\text{Eq. 12})}{\leq} \frac{8}{N} \sum_{i=1}^N \left[2n^{-1}\gamma + \frac{4\gamma}{n} \times \log(n/\gamma) \right] + \frac{1}{N} \mathbb{P}(G^c) \\
&= \frac{8}{N} \times N \left[2n^{-1}\gamma + \frac{4\gamma}{n} \times \log(n/\gamma) \right] + \frac{1}{N} \mathbb{P}(G^c) \\
&\leq 16n^{-1}\gamma + \frac{32\gamma}{n} \times \log(n/\gamma) + \frac{1}{N} \hat{\varepsilon}.
\end{aligned}$$

Hence, in the general- β case we get:

$$\mathbb{E}[F_j] \leq 16n^{-1}\gamma\beta + \frac{32\gamma}{n} \times \beta \log(n/\gamma) + \frac{1}{N} \beta \hat{\varepsilon}. \quad \blacksquare$$

D. Mixing time of the speedy walk

To bound the mixing time of the ball walk, one can consider the speedy walk. The speedy walk is the same Markov chain as the ball walk except that we leave out the steps where the ball walk does not change position. Since the ball walk ends up staying for more time at certain points than the speedy walk, the speedy walk has a different stationary distribution $\hat{\pi}_K$ called the “speedy distribution”. Denote by the random variable τ_i the stopping time which is equal to the number of proper+improper steps taken until the ball walk has taken i proper steps. Then the random walk Z_1, Z_2, \dots where $Z_i = X_{\tau_i}$ is the “speedy walk”.

We recall the following Theorem⁴ from [8], of which Theorem 18 and the following paragraph in [11] is a corollary:

Lemma IV.6 (Speedy walk (Theorem 18 and following paragraph in [11], Theorem 4.1 in [8])). *Suppose that K is 15-isotropic and fix $\hat{\varepsilon} > 0$. Given an initial point X_0 which is a β -warm start with respect to the speedy distribution, the ball walk on K with step size $\eta \geq \frac{1}{800\sqrt{n \log(n/\hat{\varepsilon})}}$ satisfies*

$$\|\mathcal{L}(X_{\tau_i}) - \hat{\pi}_K\|_{\text{TV}} \leq \hat{\varepsilon}$$

if $i \geq cn^{2.5} \log^3(\frac{\beta}{\hat{\varepsilon}})$ where $c > 0$ is a universal constant.

If instead the ball walk starts from a non-random point which is a distance at least n^{-c_1} for any constant c_1 , then the ball walk on K with step size $\eta \geq \frac{1}{800\sqrt{n \log(n/\hat{\varepsilon})}}$ satisfies

$$\|\mathcal{L}(X_{\tau_i}) - \hat{\pi}_K\|_{\text{TV}} \leq \hat{\varepsilon}$$

if $i \geq c_2 n^2 D (\log \log D) \log^3(n/\hat{\varepsilon})$ where D is the diameter of K and c_2 is a constant that depends only on c_1 .

Let λ be the probability that the ball walk proposes a step inside the convex body K from a point uniformly distributed on K ; we call λ the *average local conductance*. We will use the following results [8] which allow one to obtain improved average-case running time bounds for the ball walk.⁵

Lemma IV.7 (Corollary 4.6 in [8]). *The average local conductance of K satisfies $\lambda \geq 1 - \frac{\eta\sqrt{n}}{2r}$ if K contains a ball of radius r .*

Using average local conductance [8] gives the following Lemma on the expected number of improper steps taken by the ball walk:

Lemma IV.8 (Theorem 4.10b in [8]). *Suppose that X_0 is distributed according to the speedy distribution on K . Fix*

⁴The “M-distance” used in [8] is bounded above by the warmness β , and bounded below by the TV distance. So the result we quote here is in fact weaker than the “M-distance” version of the result.

⁵Since our goal is to put the convex body in isotropic position, which fails with exponentially small probability in the running time if we obtain iid points, it is enough to bound the average-case running time since we can always just start over if the running time ends up being too long.

$t > 0$. Then the expected number of (proper and improper) ball walk steps needed to get t proper steps is at most $\frac{2t}{\lambda}$.

Remark IV.9. Theorem 4.10b [8] was stated for a specific value of t (their bound on the mixing time). However, in the special case when we already start at the stationary distribution of the speedy walk, if the expectation holds for one value of t , it must also hold for every value of t . (Moreover, we note that even for non-stationary starts (which we do not need here) their proof holds for all values of t .)

E. Bounding the accuracy

Lemma IV.10. Assume that K is a 30-isotropic convex body containing $B(0, \frac{1}{10})$, and that $\text{dist}(X_0, \partial K) \geq n^{-3}$. Fix $\hat{\varepsilon} > 0$. Then Algorithm 1 with $\eta \leq \frac{1}{10\sqrt{8n \log(n/\hat{\varepsilon})}}$, Modified = OFF, and $\mathcal{I} = c_2 n^2 \rho (\log \log \rho) \log^3(n/\hat{\varepsilon})$ outputs independent samples Z_1, Z_2, \dots, Z_p , where each Z_i has TV distance $10\hat{\varepsilon}$ to the uniform distribution on K .

Proof: Since the speedy walk is initialized at a point X_0 , which is a distance at least n^{-3} from the boundary of K , by Lemma IV.6 we have that the samples Y_1, Y_2, \dots obtained by running the speedy walk for

$$\mathcal{I} = c_2 n^2 \rho (\log \log \rho) \log^3(n/\hat{\varepsilon})$$

proper steps each satisfy

$$\|\mathcal{L}(Y_s) - \hat{\pi}_K\|_{\text{TV}} \leq \hat{\varepsilon}$$

for all s . Moreover, these points are independent since each run of the speedy walk starts at the same point X_0 .

By Theorem 4.16 in [8] we have that the samples Z_1, Z_2, \dots obtained from Y_1, Y_2, \dots by the rejection sampling step in Algorithm 1 are uniformly distributed on K with TV error $10\hat{\varepsilon}$. Moreover, since Y_1, Y_2, \dots are jointly independent, Z_1, Z_2, \dots, Z_p are also jointly independent. ■

V. ROUNDING A POLYTOPE

In this section we analyze the running time and accuracy of Algorithm 2.

A. Bounding the success probability of Algorithm 2

In this section we bound the success probability of Algorithm 2. We use the following lemma (Corollary 11 in [2]), which is a corollary of the main result in [22].

Lemma V.1 (Corollary 11 in [2]). Let K be a convex set. Let Y_1, \dots, Y_p be iid uniform random points in K and fix $\hat{\varepsilon} > 0$. Let

$$\bar{Y} := \frac{1}{p} \sum_{i=1}^p Y_i,$$

and let

$$\hat{\Sigma}_Y := \frac{1}{p} \sum_{i=1}^p (Y_i - \bar{Y})(Y_i - \bar{Y})^\top.$$

Then there exists an absolute constant c such that if $p \geq n \times c \log^2(\frac{1}{\hat{\varepsilon}}) \log^2(n)$, the convex set

$$K^\ddagger := \hat{\Sigma}_Y^{-\frac{1}{2}}(K - \bar{Y})$$

is 2-isotropic and $\|\mu_{K^\ddagger}\|_2 < \frac{1}{20}$ with probability at least $1 - \hat{\varepsilon}$.

Note that in the proof of Corollary 11 in [2] it is shown that

$$\|\mu_{K^\ddagger}\|_2 < \frac{1}{20},$$

although this is not mentioned explicitly in [2] in the statement of their Corollary.

Fix $\varepsilon > 0$. From now on we fix the parameters p, \mathcal{I}, η in Algorithm 1 as follows:

- $p \geq n \times c \log^2(\frac{1}{\varepsilon}) \log^2(n)$,
- $\mathcal{I} = c_2 n^2 20\sqrt{n} \log\left(\frac{40n^2 p^2}{\varepsilon}\right) \times \left(\log \log 20\sqrt{n} \log\left(\frac{40n^2 p^2}{\varepsilon}\right)\right) \log^3\left(\frac{np^2}{\varepsilon}\right) \log \log\left(\frac{R}{r}\right)$,
- $\eta = \frac{1}{30\sqrt{n \log(n/\varepsilon)}}$.

Lemma V.2. Suppose that we set parameters modified = OFF, $i_{\max} = \infty$. Then for any value of $\alpha > 0$, with probability at least $1 - \varepsilon$ the convex body

$$\tilde{K}_{i^*} := \hat{\Sigma}_{i^*}^{-\frac{1}{2}}(K - \hat{\mu}_{i^*})$$

outputed by Algorithm 2 is in 2-isotropic position. Moreover, the expected number of iterations of each of the “While” loops in Algorithm 2 is bounded above by 2.

Proof: Recall the definitions from Algorithm 2 where

$$\tilde{K}_{i-1} := \begin{cases} B & \text{for } i = 1 \\ \hat{\Sigma}_i^{-\frac{1}{2}}(K_{i-1} - \hat{\mu}_i) & \text{for } i \geq 2, \end{cases}$$

and

$$\hat{K}_i := \hat{\Sigma}_i^{-\frac{1}{2}}(K_i - \hat{\mu}_i) \quad \forall i \in \mathbb{N}.$$

We prove this theorem by induction:

Inductive assumption: Suppose that

$$\tilde{K}_{i-1} = \hat{\Sigma}_i^{-\frac{1}{2}}(K_{i-1} - \hat{\mu}_i)$$

is 2-isotropic with $\|\mu_{\tilde{K}_{i-1}}\|_2 \leq \frac{1}{20}$, and that

$$B(0, \frac{1}{4}) \subseteq \tilde{K}_{i-1}.$$

Base case: Since for $i = 1$ $\tilde{K}_{i-1} = B$ is 2-isotropic, we must have that \hat{K}_i is 4e \leq 15-isotropic (see the inductive case for why this is true). Therefore \hat{K}_i is 15-isotropic and contains the ball $B(0, \frac{1}{4})$, since

$$\hat{K}_i \supseteq \tilde{K}_{i-1} = B \supseteq B(0, \frac{1}{4}).$$

Then by Lemma IV.10 we have that the points Z_1, \dots, Z_p are independent, and are each a TV distance at most $\frac{\varepsilon}{p^2}$ from the uniform distribution on \hat{K}_i .⁶

Therefore, by Lemma V.1 we have that \tilde{K}_i is 2-isotropic with probability at least $1 - \frac{\varepsilon}{p}$ (since we can couple Z_1, \dots, Z_p to independent random vectors $\hat{Z}_1, \dots, \hat{Z}_p$ which are exactly uniformly distributed on \hat{K}_i , such that

$$\mathbb{P}(X_i = \hat{Z}_i) \geq 1 - \frac{\varepsilon}{p^2},$$

and apply Lemma V.1 to these vectors).

Inductive case: Showing that \tilde{K}_{i-1} being in 2-isotropic position implies that \tilde{K}_i is in 2-isotropic position.

Since

$$K_{i-1} \subseteq K_i \quad \text{and} \quad \text{Vol}(K_i) \leq e \text{Vol}(K_{i-1}),$$

we have

$$\Sigma_{K_i} \succcurlyeq \frac{\text{Vol}(K_{i-1})}{\text{Vol}(K_i)} \Sigma_{K_{i-1}} \succcurlyeq \frac{1}{e} \Sigma_{K_{i-1}}.$$

Moreover, since K is convex and $0 \in K$, we have

$$K_i \subseteq (1 + 1/n)K_{i-1},$$

and

$$\text{Vol}(K_i) \geq \frac{1}{e} \text{Vol}((1 + 1/n)K_{i-1}).$$

Hence, we have

$$\begin{aligned} \Sigma_{K_i} &\preccurlyeq \frac{\text{Vol}((1 + 1/n)K_{i-1})}{\text{Vol}(K_i)} \Sigma_{(1+1/n)K_{i-1}} \\ &= \frac{\text{Vol}((1 + 1/n)K_{i-1})}{\text{Vol}(K_i)} (1 + 1/n)^2 \Sigma_{K_{i-1}} \\ &\preccurlyeq 4e \Sigma_{K_{i-1}}. \end{aligned}$$

Therefore, we have that

$$\frac{1}{e} \Sigma_{K_{i-1}} \preccurlyeq \Sigma_{K_i} \preccurlyeq 4e \Sigma_{K_{i-1}},$$

and hence that

$$\frac{1}{e} \frac{1}{4} I_n \preccurlyeq \frac{1}{e} \Sigma_{\tilde{K}_{i-1}} \preccurlyeq \Sigma_{\hat{K}_i} \preccurlyeq 2e \Sigma_{\tilde{K}_{i-1}} \preccurlyeq 16e I_n.$$

Therefore we have shown that the fact that \tilde{K}_{i-1} is 2-isotropic implies that \hat{K}_i is $4\sqrt{e} \leq 15$ -isotropic. Therefore \hat{K}_i is 15-isotropic and contains the ball $B(0, \frac{1}{4})$, since

$$\hat{K}_i \supseteq \tilde{K}_{i-1} \supseteq B(0, \frac{1}{4}).$$

We now show that the centers of mass of \hat{K}_i and \tilde{K}_{i-1} are a distance at most roughly \sqrt{n} apart. Suppose (towards a contradiction) that

$$\|\mu_{\hat{K}_i} - \mu_{\tilde{K}_{i-1}}\|_2 > 10\sqrt{n} \log\left(\frac{40n}{\varepsilon}\right).$$

⁶See the inductive case for why we get a bound for the uniform distribution on \hat{K}_i even though the Markov chain is on $\hat{K}_i \cap \rho B$ for $\rho = 20\sqrt{n} \log\left(\frac{40n^2}{\varepsilon}\right) B$.

By Lemma 24 in [11], $1 - \frac{\varepsilon}{40n}$ of the volume of the convex body \hat{K}_i is inside the ball of radius $2\sqrt{n} \log\left(\frac{40n}{\varepsilon}\right)$ with center at $\mu_{\hat{K}_i}$. Hence, if the assumption

$$\|\mu_{\hat{K}_i} - \mu_{\tilde{K}_{i-1}}\|_2 > 10\sqrt{n} \log\left(\frac{40n}{\varepsilon}\right)$$

were true, we would have (for $\varepsilon < 0.1$) that a nonzero portion of the volume of \hat{K}_i is a distance of at least $40n\sqrt{n}$ from $\mu_{\hat{K}_i}$, since

$$\frac{\text{Vol}(\hat{K}_i)}{\text{Vol}(\tilde{K}_{i-1})} \leq e \quad \text{and} \quad \hat{K}_i \supseteq \tilde{K}_{i-1}.$$

This is a contradiction since the convex body \hat{K}_i is entirely contained in a ball of radius $15n$ because it is 15-isotropic. Hence by contradiction we have that

$$\|\mu_{\hat{K}_i} - \mu_{\tilde{K}_{i-1}}\|_2 \leq 10\sqrt{n} \log\left(\frac{40n}{\varepsilon}\right).$$

By inductive assumption we have that $\|\mu_{\tilde{K}_{i-1}}\|_2 \leq \frac{1}{5}$, and hence that

$$\|\mu_{\hat{K}_i}\|_2 \leq \|\mu_{\hat{K}_i} - \mu_{\tilde{K}_{i-1}}\|_2 + \|\mu_{\tilde{K}_{i-1}}\|_2 \leq 12\sqrt{n} \log\left(\frac{40n}{\varepsilon}\right).$$

By Lemma 24 in [11], $1 - \frac{\varepsilon}{40n^2 p^2}$ of the volume of the convex body \hat{K}_i is contained in a ball of radius $2\sqrt{n} \log\left(\frac{40n^2 p^2}{\varepsilon}\right)$ centered at $\mu_{\hat{K}_i}$. Hence, since \hat{K}_i being 15-isotropic implies that it is contained in a ball of radius $15n$, we have that

$$\hat{K}_i^\dagger := \hat{K}_i \cap 20\sqrt{n} \log\left(\frac{40n^2 p^2}{\varepsilon}\right) B$$

is 30-isotropic. Moreover, by Lemma 24 in [11] and the fact that

$$\|\mu_{\hat{K}_i}\|_2 \leq 12\sqrt{n} \log\left(\frac{40n}{\varepsilon}\right),$$

we have

$$\text{Vol}(\hat{K}_i^\dagger) \geq \left(1 - \frac{\varepsilon}{2p^2}\right) \text{Vol}(\hat{K}_i).$$

Therefore, since the rejection step in the “While” loop of Algorithm 2 ensures that X_0 is in the n^{-3} -interior of

$$\hat{K}_i \cap 20\sqrt{n} \log\left(\frac{40n^2}{\varepsilon}\right) B,$$

by Lemma IV.10 we have that the points Z_1, \dots, Z_p are independent, and are each a TV distance at most $\frac{\varepsilon}{p^2}$ from the uniform distribution on \hat{K}_i .

Therefore, by Lemma V.1 we have that \hat{K}_i is 2-isotropic and

$$\|\mu_{\hat{K}_i}\|_2 \leq \frac{1}{20}$$

with probability at least $1 - \frac{\varepsilon}{p}$ (Since we can couple Z_1, \dots, Z_p to independent random vectors $\hat{Z}_1, \dots, \hat{Z}_p$ which are exactly uniformly distributed on \hat{K}_i , such that

$$\mathbb{P}(Z_k = \hat{Z}_k) \geq 1 - \frac{\varepsilon}{p^2},$$

and apply Lemma V.1 to these vectors. Therefore, we have Z_1, \dots, Z_p independent with

$$\|\mathcal{L}(Z_k) - \pi_{\tilde{K}_i}\|_{\text{TV}} \leq \frac{\varepsilon}{10p^2 \log(\frac{R}{r})}$$

for all $k \in [p]$.

Bounding the number of iterations of the “While” loop. First, we bound $\text{dist}(\hat{X}'_0, \partial[\tilde{K}_i \cap \rho B])$ for

$$\rho = 20\sqrt{n} \log\left(\frac{40n^2}{\varepsilon}\right).$$

Since \tilde{K}_i is 2-isotropic, it contains a ball of radius $\frac{1}{\sqrt{2}}$. Therefore, by a statement in the proof of Corollary 4.6 in [8], we have that

$$\text{Vol}_{n-1}(\partial\tilde{K}_i) \leq \frac{n}{1/\sqrt{2}} \text{Vol}(\partial\tilde{K}_i).$$

Therefore, we have that

$$\begin{aligned} \mathbb{P}_{Y \sim \text{unif}(\tilde{K}_i)}(\text{dist}(Y, \partial\tilde{K}_i) < n^{-3}) &\leq \frac{n^{-3} \text{Vol}_{n-1}(\partial\tilde{K}_i)}{\text{Vol}(\partial\tilde{K}_i)} \\ &\leq n^{-2}\sqrt{2}. \end{aligned}$$

Moreover, by Lemma 24 in [11], $1 - \frac{\varepsilon}{10n^2p^2}$ of the volume of the convex body \tilde{K}_i is contained in $B(\mu_{\tilde{K}_i}, \frac{\rho}{2})$, and since \tilde{K}_i is in 2-isotropic position,

$$\|\mu_{\tilde{K}_i}\|_2 \leq \frac{1}{5}.$$

Thus, $1 - \frac{\varepsilon}{10n^2p^2}$ of the volume of the convex body \tilde{K}_i is contained in the ball $B(0, \rho)$.

Now, since \hat{X}_0 and Z_1 are generated by the ball walk with the same starting point and parameters, we have

$$\begin{aligned} \|\mathcal{L}(\hat{X}'_0) - \pi_{\tilde{K}_i}\|_{\text{TV}} &= \|\mathcal{L}(\hat{X}_0) - \pi_{\tilde{K}_i}\|_{\text{TV}} \\ &= \|\mathcal{L}(Z_1) - \pi_{\tilde{K}_i}\|_{\text{TV}} \\ &\leq \frac{\varepsilon}{p^2}. \end{aligned}$$

Therefore, we have

$$\begin{aligned} \mathbb{P}\left(\text{dist}(\hat{X}'_0, \partial[\tilde{K}_i \cap \rho B]) < n^{-3}\right) &\leq n^{-2}\sqrt{2} + \frac{\varepsilon}{10n^2p^2} + \|\mathcal{L}(Z_k) - \pi_{\tilde{K}_i}\|_{\text{TV}} \\ &\leq n^{-2}\sqrt{2} + \frac{\varepsilon}{10n^2p^2} + \frac{\varepsilon}{p^2}. \end{aligned}$$

But $\hat{X}'_0 \equiv \hat{X}'_0^k$ is generated independently at each iteration k of the while loop, implying that

$$\begin{aligned} \mathbb{P}\left(\min_{j \leq k} \text{dist}(\hat{X}'_0^j, \partial[\tilde{K}_i \cap \rho B]) \leq n^{-3}\right) &\leq \left(n^{-2}\sqrt{2} + \frac{\varepsilon}{10n^2p^2} + \frac{\varepsilon}{p^2}\right)^k \\ &\leq 2^{-k}. \end{aligned}$$

Therefore, the expected number of iterations of the “While” loop is bounded by

$$\sum_{k=1}^{\infty} 2^{-(k-1)} \leq 2. \quad \blacksquare$$

B. Bounding the expected running time of Algorithm 2

In this section we bound the expected frequency at which Algorithm 2 checks any given inequalities. From now on we set the parameter α of Algorithm 1 to be

$$\alpha = 4 \log\left(\frac{2npi_{\max}}{\varepsilon}\right).$$

Lemma V.3. Suppose that, at each step of Algorithm 2, X_0 is a β -warm start with respect to both the uniform distribution and the speedy distribution. Fix $b^* \leq \frac{i_{\max}}{40\mathcal{I}}$. Then with probability at least $\frac{8}{10}$ the total number of inequality checks made during the first b^* times that Algorithm 2 with Modified = ON invokes the ball walk Markov chain is $\leq 40\beta b^* \mathcal{I} \times m \left[16n^{-1}\gamma + \frac{32\gamma}{n} \times \log(n/\gamma) + \frac{\varepsilon}{np} \right]$.

Proof: Let \mathcal{S}^k be the number of proper+improper steps for the k th Markov chain in Algorithm 2, and let \mathcal{F}_j^k be the frequency at which the k th Markov chain checks the inequality (A_j, b_j) .

Let

$$\mathcal{S}^\dagger := \sum_{k=1}^{b^*} \mathcal{S}^k$$

be the sum of the steps for the first b^* runs of the Markov chain in Algorithm 2, and let

$$\mathcal{F}^\dagger := \frac{1}{b^*} \sum_{k=1}^{b^*} \sum_{j=1}^m \mathcal{F}_j^k$$

be the frequency at which Algorithm 2 checks any inequality.

Now, if Y_0 is a β -warm start for the speedy distribution, by Lemmas IV.8 and IV.7 we have

$$\mathbb{E}[\mathcal{S}^\dagger] \leq \beta \times 4b^* \mathcal{I}.$$

Hence, we have, by Markov’s inequality, that

$$\mathcal{S}^\dagger \leq 40b^* \mathcal{I} \leq i_{\max},$$

with probability at least $\frac{9}{10}$.

Thus, by Lemma IV.5 we have

$$\begin{aligned} \mathbb{E}[\mathcal{F}^\dagger \mathbb{1}\{\mathcal{S}^\dagger \leq i_{\max}\}] &\leq m \times \left[16n^{-1}\gamma\beta + \frac{32\gamma}{n} \times \beta \log(n/\gamma) + \frac{\varepsilon\beta}{nb^*} \right]. \end{aligned}$$

Hence, by Markov’s inequality we have with probability at least $\frac{9}{10}$ that

$$\mathcal{F}^\dagger \mathbb{1}\{\mathcal{S}^\dagger \leq i_{\max}\} \leq 160n^{-1}\gamma\beta + \frac{320\gamma}{n}\beta \log(n/\gamma) + 10\frac{\beta\varepsilon}{nb^*}.$$

Therefore, with probability at least $\frac{8}{10}$ we have that both

$$\mathcal{S}^\dagger \leq 40\beta b^* \mathcal{I} \leq i_{\max}$$

and

$$\mathcal{F}^\dagger \leq 160n^{-1}\gamma\beta + \frac{320\gamma}{n} \times \beta \log(n/\gamma) + 10\frac{\beta\varepsilon}{nb^*},$$

since

$$\mathcal{F}^\dagger \mathbb{1}\{\mathcal{S}^\dagger \leq i_{\max}\} = \mathcal{F}^\dagger$$

whenever

$$\mathcal{S}^\dagger \leq 40\beta b^* \mathcal{I} \leq i_{\max}.$$

Therefore, with probability at least $\frac{8}{10}$ we have that the total number of times Algorithm 1 computes an inequality is

$$\mathcal{S}^\dagger \mathcal{F}^\dagger \leq 40\beta b^* \mathcal{I} \times m \left[16n^{-1}\gamma\beta + \frac{32\gamma}{n}\beta \log(n/\gamma) + \frac{\beta\varepsilon}{nb^*} \right]. \quad \blacksquare$$

C. Bounding the warmth of the start in Algorithm 2 with respect to the speedy and uniform distributions

Lemma V.4. *With probability at least $\frac{6}{10} - \varepsilon$ we have that, the total number of inequality checks for the duration of Algorithm 2 with Modified = ON, is at most*

$$220(2+p)i^* \mathcal{I} \times m \left[16n^{-1}\gamma + \frac{32\gamma}{n} \times \log(n/\gamma) + \frac{\varepsilon}{np} \right].$$

Moreover, if $\hat{\Sigma}_{i^*}, \hat{\mu}_{i^*}, X_0$ are the outputs of Algorithm 2, we have that X_0 is $O(1)$ -warm with respect to the uniform distribution on $\hat{\Sigma}_{i^*}^{-\frac{1}{2}}(K - \hat{\mu}_{i^*})$.

Proof: First, we bound the number of times Algorithm 1 runs the ball walk. By Lemma V.2 we have that the expected number of iterations of the “While” loop at each “For” loop iteration of Algorithm 1 is at most 2. Hence, the expected number of times the ball walk is run by Algorithm 1 is at most $(2+p)i^*$. Therefore, by Markov’s inequality, we have that with probability at least $\frac{9}{10}$ the ball walk is run by Algorithm 1 no more than $10(2+p)i^*$ times.

Next, we bound the speedy and uniform warmth. The main observation of this section is that, for our step size, any distribution which is β -warm with respect to the uniform distribution is 2β -warm with respect to the speedy distribution. This is true since, by Remark 4.12 in [8], one can obtain speedy-distributed samples by starting with uniformly distributed samples, if we take one step of the ball walk starting at each sample, and reject the original point if and only if the ball walk step leaves the convex body. By Lemma IV.7, for our choice of step size η the average acceptance probability λ is at most $\frac{1}{2}$. Therefore, any sample which is β -warm with respect to the uniform distribution must also

be 2β -warm with respect to the speedy distribution⁷, since then

$$\hat{\pi}_K(x) = \pi_K(x) \frac{1}{\lambda} \mathbb{P}(\{\text{one step of ball walk starting at } x \text{ is rejected}\}).$$

Therefore, to apply Lemma V.3, it is enough to show that the initial points in Algorithm 1 are β -warm with respect to the uniform distribution.

Denote the value of \hat{X}_0' at the i th iteration of the “For” loop and k ’th iteration of the “While” loop of Algorithm 2 by $\hat{X}_0' \equiv \hat{X}_0'^{i,k}$. Next, we note that by the proof of Lemma V.2 we have

$$\|\mathcal{L}(\hat{X}_0'^{i,k}) - \pi_{\hat{K}_i}\|_{\text{TV}} \leq \frac{\varepsilon}{10p(i^*)^2}$$

for each i, k . Moreover, since we are using a fixed starting point, $\hat{X}_0'^1, \hat{X}_0'^2, \dots$ are independent. Thus, there exists a sequence of random variables

$$\{\tilde{X}_0^{i,k}\}_{i \in [i^*], k \in [10(2+p)i^*]}$$

such that $\tilde{X}_0^{i,k} = \hat{X}_0'^{i,k}$ with probability at least $\frac{9}{10} - \varepsilon$, where $\tilde{X}_0^{i,k} \sim \pi_{\hat{K}_i}$. Hence, $\tilde{X}_0^{i,k}$ are 1-warm with respect to the uniform distribution on \tilde{K}_i .

Recall from the proof of Lemma V.2 that

$$K_i \subseteq K_{i+1} \quad \text{and} \quad \text{Vol}(K_{i+1}) \leq e\text{Vol}(K_i).$$

Thus, we have

$$\tilde{K}_i \subseteq \hat{K}_{i+1} \quad \text{and} \quad \text{Vol}(\hat{K}_{i+1}) \leq e\text{Vol}(\tilde{K}_i),$$

since

$$\tilde{K}_i := \hat{\Sigma}_{i+1}^{-\frac{1}{2}}(K_i - \hat{\mu}_{i+1}) \text{ and } \hat{K}_{i+1} := \hat{\Sigma}_{i+1}^{-\frac{1}{2}}(K_{i+1} - \hat{\mu}_{i+1}).$$

Therefore, the fact that $\tilde{X}_0^{i,k}$ is 1-warm with respect to the uniform distribution on \tilde{K}_i implies that it is e -warm with respect to the uniform distribution on \hat{K}_{i+1} . Hence, we have that $\tilde{X}_0^{i,k}$ is $2e$ -warm with respect to the speedy distribution on \hat{K}_{i+1} .

Hence, by Lemma V.3 we have that with probability at least $\frac{6}{10} - \varepsilon$, the total number of inequality checks for the duration of Algorithm 2, is at most

$$80e(2+p)i^* \mathcal{I} \times m \left[16n^{-1}\gamma + \frac{32\gamma}{n} \times \log(n/\gamma) + \frac{\varepsilon}{np} \right]. \quad \blacksquare$$

⁷Note, however, that the converse is not true: β -warm with respect to the speedy distribution does not imply 2β warm with respect to the uniform distribution.

D. Counting the running time of subroutines with negligible contributions to the running time

Sorting: First, we argue that the cost of sorting the h_j 's can be ignored.

Algorithm 1 must sort the m inequalities into

$$O\left(\frac{n}{\alpha\eta}\right) = O(n^{1.5})$$

bins, since, if K is $O(1)$ -isotropic, the diameter of K is $O(\sqrt{n})$. This takes $O(m \log(n))$ time. This only occurs once, at the start of the algorithm. Since our bound for the entire algorithm is $O(mn^{4.5})$ and $mn^{4.5} \gg m \log(n)$, we can ignore this cost.

At each step of the Markov chain, Algorithm 1 must move all the bins over by 1 to “add” $\frac{n}{\alpha\eta}$ to each bin, and select all h_j 's such that

$$h_j < \alpha \frac{\eta}{\sqrt{n}} \times i.$$

This can be done in $O(1)$ operations by simply moving a “pointer” at the bin for elements in $[\alpha \frac{\eta}{\sqrt{n}} \times (i-1), \alpha \frac{\eta}{\sqrt{n}} \times i]$ one bin to the right. This is only done once every Markov chain step, and is negligible in comparison to the cost of computing a Markov chain step, implying that we can ignore this $O(1)$ cost.

After computing the new values for all the h_j 's that were selected, Algorithm 1 must then sort these h_j 's into the corresponding bins. Since there are $O(n^{1.5})$ bins, this takes $O(\log(n))$ time for each selected h_j . This is negligible in comparison to the cost of recomputing the value of each h_j , which is $O(n)$ (since we have to take an inner product). Therefore we can ignore this cost.

Applying the linear transformation to K_i : Next, we argue that the cost of applying the linear transformation to put K_i into isotropic position at each iteration of Algorithm 2 can be ignored. This linear transformation requires computing $A\hat{\Sigma}_i^{\frac{1}{2}}$, which takes $O(mn^2)$ operations, and computing $b - A\hat{\mu}_i$, which takes $O(mn)$ operations. There are $i^* = n \log(\frac{R}{r})$ iterations in Algorithm 2, so applying the linear transformations contributes at most $O(mn^3 \log(\frac{R}{r}))$ arithmetic operations. This is much smaller than our bound of $O(mn^{4.5} \log(\frac{R}{r}))$ on the number of operations, so we can ignore the cost of applying the linear transformations.

E. Proof of main theorem for rounding

We now consider the following procedure, where we run Algorithm 2 multiple times until it succeeds. This allows us to put the polytope into isotropic position with a bound on the running time that holds with very high probability (as opposed to just holding in expectation).

Algorithm 3 Rounding in bounded time

input: $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$
input: $r, R > 0$ such that $rB \subseteq K \subseteq RB$, where $K := \{x \in \mathbb{R}^n : Ax \leq b\}$
input: $p \in \mathbb{N}$, $\varepsilon > 0$
input: $\mathcal{I}, i_{\max} \in \mathbb{N}$

- 1: Set Success = False
- 2: **while** Success = False **do**
- 3: Run Algorithm 2 with Modified = ON and with the above inputs until either it outputs $\hat{\Sigma}_{i^*}, \hat{\mu}_{i^*}, X_0$ or until it completes \mathcal{I} inequality checks.
- 4: If $\hat{\Sigma}_{i^*}, \hat{\mu}_{i^*}, X_0$ were obtained, output Success = True
- 5: **end while**
- 6: **output:** $\hat{\Sigma}_{i^*}, \hat{\mu}_{i^*}, X_0$

From now on we choose \mathcal{I}, i_{\max} to satisfy

$$\mathcal{I} = 2000(2+p)i^*\mathcal{I} \quad (14)$$

$$\times m \left[16n^{-1}\gamma + \frac{32\gamma}{n} \log(n/\gamma) + \frac{6\varepsilon}{np} \right],$$

$$i_{\max} = \left(2000(2+p)i^*\mathcal{I} \times m \left[\frac{320m\hat{\eta}}{n\varepsilon} + \frac{640m\hat{\eta}}{n\varepsilon} \log\left(\frac{n\varepsilon}{20m\hat{\eta}}\right) + \frac{\varepsilon}{np} \right] \right)^2.$$

In particular, we have $i_{\max} \geq \mathcal{I}$.

Lemma V.5. Fix $\varepsilon > 0$. Then with probability at least $1 - 2\varepsilon \log(\frac{1}{\varepsilon})$ Algorithm 3 outputs, after at most $\log(\frac{1}{\varepsilon})$ calls to Algorithm 2, an affine transformation $(\hat{\Sigma}_{i^*}, \hat{\mu}_{i^*})$ and a point X_0 , for which

$$K^\dagger := \hat{\Sigma}_{i^*}^{-\frac{1}{2}}(K - \hat{\mu}_{i^*})$$

is in 2-isotropic position. Moreover we have that X_0 is $O(\log(\frac{1}{\varepsilon}))$ -warm with respect to the uniform distribution on K^\dagger .

Proof: First, we note that by Lemma V.4 we have that the “While” loop of Algorithm 3 completes at least one run of Algorithm 2 after $\log(\frac{1}{\varepsilon})$ runs of the “While” loop of Algorithm 3 with probability at least

$$1 - 2^{-\log(\frac{1}{\varepsilon})} \geq 1 - \varepsilon.$$

Moreover, by Lemmas V.2 and V.4, if Algorithm 2 were allowed to keep running even after it uses up its allotted number of arithmetic operations and ball walk steps, it would return an affine transformation which puts K into 2-isotropic position, and an $O(1)$ -warm start for the uniform distribution on this affine transformation of K , with probability at least $1 - \varepsilon$ at each run. Therefore, if we do stop Algorithm 2 after its allotted number of arithmetic operations and ball walk steps, after $\log(\frac{1}{\varepsilon})$ steps we obtain, with probability at

least $1 - \varepsilon - \varepsilon \log(\frac{1}{\varepsilon})$, an affine transformation $(\hat{\Sigma}_{i^*}, \hat{\mu}_{i^*})$ for which

$$K^\dagger := \hat{\Sigma}_{i^*}^{-\frac{1}{2}}(K - \hat{\mu}_{i^*})$$

is in 2-isotropic position and an $O(\log(\frac{1}{\varepsilon}))$ -warm start X_0 for the uniform distribution on K^\dagger . \blacksquare

Fix

$$p = n \times c \log^2(\frac{1}{\varepsilon}) \log^2(n).$$

Recall the shorthand notation

$$\hat{\eta} := \frac{1}{10}\eta\sqrt{n} \quad \text{and} \quad \gamma := 10\alpha\hat{\eta},$$

and that we have fixed the parameters $p, \mathcal{I}, \eta, \alpha$ of Algorithm 1 as follows:

- $\mathcal{I} = c_2 n^2 20\sqrt{n} \log\left(\frac{40n^2 p^2}{\varepsilon}\right) \times \left(\log \log 20\sqrt{n} \log\left(\frac{40n^2 p^2}{\varepsilon}\right)\right) \log^3\left(\frac{np^2}{\varepsilon}\right) \log \log(\frac{R}{r})$,
- $\eta = \frac{1}{30\sqrt{n \log(n/\varepsilon)}}$,
- $\alpha = 4 \log\left(\frac{2np i_{\max}}{\varepsilon}\right)$.

Also recall that we have fixed \mathfrak{I} and i_{\max} in Equation 14, and that i^* is set to

$$i^* = n \log_2\left(\frac{R}{r}\right)$$

in Algorithm 2.

Theorem V.6 (Version of Main Theorem (Th. I.1) specific to our algorithm). Fix $\varepsilon > 0$. With probability at least $1 - 2\varepsilon \log(\frac{1}{\varepsilon})$ Algorithm 3 outputs an affine transformation $(\hat{\Sigma}_{i^*}, \hat{\mu}_{i^*})$ and a point X_0 for which

$$K^\dagger := \hat{\Sigma}_{i^*}^{-\frac{1}{2}}(K - \hat{\mu}_{i^*})$$

is in 2-isotropic position, and X_0 is $O(\log(\frac{1}{\varepsilon}))$ -warm with respect to the uniform distribution on K^\dagger , in $\tilde{O}(mn^{4.5} \log^9(\frac{1}{\varepsilon}) \log^9(\frac{R}{r}))$ arithmetic operations.

Proof: By Lemma V.5, with probability at least $1 - 2\varepsilon \log(\frac{1}{\varepsilon})$ Algorithm 3 outputs an affine transformation $(\hat{\Sigma}_{i^*}, \hat{\mu}_{i^*})$ which puts K into 2-isotropic position, in at most $\mathfrak{I} \log(\frac{1}{\varepsilon})$ inequality checks. The number of arithmetic operations for each inequality check is no more than $3n$, implying that the number of arithmetic operations performed by Algorithm 3 in this event is at most

$$\begin{aligned} & 3n \times \mathfrak{I} \log(\frac{1}{\varepsilon}) \\ &= 6000(2+p)i^*\mathcal{I} \times m \left[16\gamma + 32\gamma \log(n/\gamma) + \frac{6\varepsilon}{p} \right] \log(\frac{1}{\varepsilon}) \\ &= \tilde{O}(mn^{4.5} \log^9(\frac{1}{\varepsilon}) \log^9(\frac{R}{r})). \end{aligned} \quad \blacksquare$$

VI. PROOFS OF COROLLARIES FOR VOLUME COMPUTATION AND SAMPLING

In this section we prove Corollaries I.2 and I.3 from the introduction. First, we prove Corollary I.2 on volume estimation.

Proof of Corollary I.2: By Lemma V.6 we have that with probability at least $1 - \varepsilon$ Algorithm 3 obtains an affine transformation $(\hat{\Sigma}_{i^*}, \hat{\mu}_{i^*})$ such that

$$K^\dagger := \hat{\Sigma}_{i^*}^{-\frac{1}{2}}(K - \hat{\mu}_{i^*})$$

is in 2-isotropic position, in $\tilde{O}(mn^{4.5} \text{polylog}(\frac{1}{\varepsilon}, \frac{R}{r}))$ arithmetic operations. Thus, with probability at least $1 - \varepsilon$, K^\dagger contains a ball $B(\mu_{K^\dagger}, \frac{1}{2})$ where

$$\|\mu_{K^\dagger}\|_2 \leq \frac{1}{5},$$

implying that

$$B(0, \frac{1}{4}) \subseteq K^\dagger.$$

Thus, with probability at least $1 - \varepsilon$, $4K^\dagger$ contains the unit ball B and is 8-isotropic. Hence, if we then apply the volume algorithm [5] to $4K^\dagger$, by Theorem 1.1 of [5] with probability at least $1 - \varepsilon$ we can compute the volume of $4K^\dagger$ up to a factor of $1 + \delta$ in

$$\frac{mn^4}{\delta^2} \log^6\left(\frac{n}{\delta}\right) \log\left(\frac{1}{\varepsilon}\right)$$

arithmetic operations. But

$$\text{Vol}(K) = \frac{1}{4^n} \det(\hat{\Sigma}_{i^*}^{\frac{1}{2}}) \text{Vol}(4K^\dagger).$$

Since $\det(\hat{\Sigma}_{i^*}^{\frac{1}{2}})$ can be computed in $O(n^3)$ arithmetic operations, we can compute with probability at least $1 - \varepsilon$ the volume of K up to a factor of $1 + \delta$ in

$$\tilde{O}\left(mn^{4.5} \text{polylog}\left(\frac{1}{\varepsilon}, \frac{R}{r}\right) + \frac{mn^4}{\delta^2} \log^6\left(\frac{n}{\delta}\right) \log\left(\frac{1}{\varepsilon}\right)\right)$$

arithmetic operations. \blacksquare

We then prove Corollary I.3 on sampling:

Proof of Corollary I.3: By Theorem V.6 we have that, with probability at least $1 - \varepsilon$, Algorithm 3 obtains an affine transformation $(\hat{\Sigma}_{i^*}, \hat{\mu}_{i^*})$ and a random vector X_0 such that

$$K^\dagger := \hat{\Sigma}_{i^*}^{-\frac{1}{2}}(K - \hat{\mu}_{i^*})$$

is in 2-isotropic position and X_0 is $O(1)$ -warm with respect to the uniform distribution on K^\dagger in

$$\tilde{O}\left(mn^{4.5} \text{polylog}\left(\frac{1}{\varepsilon}, \frac{R}{r}\right)\right)$$

arithmetic operations.

Since K^\dagger is in 2-isotropic position it is contained in $B(\mu_{K^\dagger}, 4n)$ where

$$\|\mu_{K^\dagger}\|_2 \leq \frac{1}{5}.$$

Moreover, by Lemma 24 in [11], $1 - \varepsilon$ of the volume of the convex body K^\dagger is contained in a ball of radius $2\sqrt{n} \log(\frac{1}{\varepsilon})$ centered at μ_{K^\dagger} .

Hence, by Theorem 1.1 of [18], if we run the hit-and-run Markov chain sampling algorithm for $\tilde{O}(n^3 \log(\frac{1}{\varepsilon}))$ Markov chain steps on the convex body

$$K^\dagger \cap B(0, 4\sqrt{n} \log(\frac{1}{\varepsilon}))$$

with $O(1)$ -warm start X_0 , we obtain a point Z uniformly distributed on

$$K^\dagger \cap B(0, 4\sqrt{n} \log(\frac{1}{\varepsilon})),$$

with TV error ε .

But since

$$\text{Vol}(K^\dagger \cap B(0, 4\sqrt{n} \log(\frac{1}{\varepsilon}))) \geq (1 - \varepsilon) \text{Vol}(K^\dagger),$$

we have that Z is also uniformly distributed on K^\dagger with TV error 2ε . To obtain a point on K , we compute

$$\tilde{Z} = \hat{\Sigma}_{i^*}^{\frac{1}{2}} Z + \hat{\mu}_{i^*}.$$

Since Z is uniformly distributed on K^\dagger with TV error 2ε , \tilde{Z} must be uniformly distributed on K with TV error 2ε as well.

Therefore, we obtain a point Z from the uniform distribution on K with TV error 2ε , in a number

$$\begin{aligned} \tilde{O}(mn^{4.5} \text{polylog}(\frac{1}{\varepsilon}, \frac{R}{r}) + n^3 \log(\frac{1}{\varepsilon})) \\ = \tilde{O}(mn^{4.5} \text{polylog}(\frac{1}{\varepsilon}, \frac{R}{r})) \end{aligned}$$

of arithmetic operations. \blacksquare

ACKNOWLEDGMENT

The authors would like to thank Sushant Sachdeva and the anonymous reviewers for their helpful comments and suggestions. This research was partially supported by NSF CCF-1908347 and SNSF 200021_182527 grants.

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. ACM, 1991.
- [2] Dimitris Bertsimas and Santosh S Vempala. Solving convex programs by random walks. In *Proceedings of the thiry-fourth annual ACM symposium on theory of computing*, pages 109–115. ACM, 2002.
- [3] Yuansi Chen, Raaz Dwivedi, Martin J Wainwright, and Bin Yu. Fast MCMC sampling algorithms on polytopes. *The Journal of Machine Learning Research*, 19(1):2146–2231, 2018.
- [4] Benjamin Cousins. *Efficient high-dimensional sampling and integration*. PhD thesis, Georgia Institute of Technology, 2017.
- [5] Benjamin Cousins and Santosh S Vempala. Bypassing KLS: Gaussian cooling and an $O^*(n^3)$ volume algorithm. In *Proceedings of the forty-seventh annual ACM symposium on theory of computing*, pages 539–548. ACM, 2015.
- [6] Martin Dyer, Alan Frieze, and Ravi Kannan. A random polynomial-time algorithm for approximating the volume of convex bodies. *Journal of the ACM (JACM)*, 38(1):1–17, 1991.
- [7] Hulda S Haraldsdóttir, Ben Cousins, Ines Thiele, Ronan MT Fleming, and Santosh S Vempala. CHRR: coordinate hit-and-run with rounding for uniform sampling of constraint-based models. *Bioinformatics*, 33(11):1741–1743, 2017.
- [8] Ravi Kannan, László Lovász, and Miklós Simonovits. Random walks and an $O^*(n^5)$ volume algorithm for convex bodies. *Random Structures & Algorithms*, 11(1):1–50, 1997.
- [9] Ravindran 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.
- [10] Michel Ledoux. *The concentration of measure phenomenon*. American Mathematical Soc., 2001.
- [11] Yin Tat Lee and Santosh S Vempala. Eldan’s stochastic localization and the KLS hyperplane conjecture: An improved lower bound for expansion. In *Foundations of Computer Science (FOCS), 2017 IEEE 58th Annual Symposium on*, pages 998–1007. IEEE, 2017.
- [12] Yin Tat Lee and Santosh S Vempala. Geodesic walks in polytopes. In *Proceedings of the 49th Annual ACM SIGACT Symposium on theory of Computing*, pages 927–940. ACM, 2017.
- [13] 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. ACM, 2018.
- [14] Yin Tat Lee and Santosh S Vempala. The Kannan-Lovász-Simonovits conjecture. *arXiv preprint arXiv:1807.03465*, 2018.
- [15] László Lovász and Miklós Simonovits. Random walks in a convex body and an improved volume algorithm. *Random structures & algorithms*, 4(4):359–412, 1993.
- [16] László Lovász and Santosh S Vempala. The geometry of logconcave functions and sampling algorithms. *Random Structures & Algorithms*, 30(3):307–358, 2007.
- [17] László Lovász and Santosh S 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.
- [18] László Lovász and Santosh S Vempala. Hit-and-run from a corner. *SIAM Journal on Computing*, 35(4):985–1005, 2006.

- [19] László Lovász and Santosh S Vempala. Simulated annealing in convex bodies and an $O^*(n^4)$ volume algorithm. *Journal of Computer and System Sciences*, 72(2):392–417, 2006.
- [20] Robert Osserman. The isoperimetric inequality. *Bulletin of the American Mathematical Society*, 84(6):1182–1238, 1978.
- [21] Grigorios Paouris. Small ball probability estimates for log-concave measures. *Transactions of the American Mathematical Society*, 364(1):287–308, 2012.
- [22] Mark Rudelson. Random vectors in the isotropic position. *Journal of Functional Analysis*, 164(1):60–72, 1999.
- [23] Santosh S Vempala. Geometric random walks: a survey. *Combinatorial and computational geometry*, 52(573–612):2, 2005.