Near-Optimal Deterministic Algorithms for Volume Computation via M-ellipsoids

Daniel Dadush* Santosh Vempala†

May 2012

Abstract

We give a deterministic $2^{O(n)}$ algorithm for computing an M-ellipsoid of a convex body, matching a known lower bound. This leads to a nearly optimal deterministic algorithm for estimating the volume of a convex body and improved deterministic algorithms for fundamental lattice problems under general norms.

1 Introduction

Ellipsoids have traditionally played an important role in the study of convex bodies. The classical Lowner-John ellipsoid, for instance, is the starting point for many interesting studies. To recall John's theorem, for any convex body K in $\mathbb{R}^n$, there is an ellipsoid E with centroid x_0 such that

$$x_0 + E \subseteq K \subseteq x_0 + nE.$$

This bound is achieved by the *maximum volume* ellipsoid contained in K.

Ellipsoids have also been critical to the design and analysis of efficient algorithms. The most notable example is the ellipsoid algorithm [1, 2] for linear [3] and convex optimization [4], which represents a frontier of polynomial-time solvability. For the basic problems of sampling and integration in high dimension, the *inertial* ellipsoid defined by the covariance matrix of a distribution is an important ingredient of efficient algorithms [5, 6, 7]. This ellipsoid also achieves the bounds of John's theorem for general convex bodies (for centrally-symmetric convex bodies, the maximum volume ellipsoid achieves the best possible sandwiching ratio of $\sqrt{n}$ while the inertial ellipsoid could still have a ratio of n).

Another ellipsoid that has played a critical role in the development of modern convex geometry is the M-ellipsoid (Milman's ellipsoid). This object was introduced by Milman as a tool to prove fundamental inequalities in convex geometry (see e.g., Chapter 7 of [8]). An M-ellipsoid E of a convex body K has small *covering numbers* with respect to K. We let $N(A, B)$ denote the number of translations of a set B required to cover the set A. Then, as shown by Milman, every convex body K in $\mathbb{R}^n$ has an ellipsoid E for which $N(K, E)N(E, K)$ is bounded by $2^{O(n)}$. This is the best possible bound up to a constant in the exponent. In contrast, the John ellipsoid can have this covering bound as high as $n^{\Omega(n)}$. Intuitively, an M-ellipsoid for K is the largest ellipsoid with the property that roughly $1/2^n$ fraction of its volume is inside

* School of Industrial and Systems Engineering, Georgia Tech. dndadush@gmail.com

† School of Computer Science, Georgia Tech. vempala@gatech.edu

K (as opposed to the entire ellipsoid being in K). The existence of M-ellipsoids now has several proofs in the literature: by Milman [9], multiple proofs by Pisier [8] and, more recently, by Klartag [10].

The complexity of computing these ellipsoids is interesting for its own sake, but also due to several important consequences that we will discuss presently. John ellipsoids are hard to compute, but their sandwiching bounds can be approximated deterministically to within $O(\sqrt{n})$ in polynomial time [4]. Inertial ellipsoids can be approximated to arbitrary accuracy by random sampling in polynomial time. Algorithms for M-ellipsoids have been considered more recently. The proof of Klartag [10] gives a randomized polynomial-time algorithm [11]. This approach is based on estimating a covariance matrix from random samples and seems inherently difficult to derandomize. It has been open to give a deterministic algorithm for constructing an M-ellipsoid that achieves optimal covering bounds. The extent to which randomness is essential for efficiency is a very interesting question in general, and specifically for problems on convex bodies where separations between randomized and deterministic complexity are known in the general oracle model [12, 13]. Here we address the question of deterministic M-ellipsoid construction and consider its algorithmic consequences for volume estimation and also for fundamental lattice problems, namely the Shortest Vector Problem (SVP) and the Bounded Distance Decoding (BDD) problem.

The first new result of this paper is a deterministic $2^{O(n)}$ algorithm for computing an M-ellipsoid of a convex body in the oracle model [4]. This is the best possible up to a constant in the exponent as there is a $2^{\Omega(n)}$ lower bound for deterministic algorithms. We state this result formally, then proceed to its extensions and consequences. For all our algorithmic problems with convex bodies, we only need the body to be specified by a standard *well-guaranteed membership oracle*, i.e., the algorithm has access to a membership oracle for the convex body of interest K, a point x_0 in K and numbers r, R s.t. balls of these radii sandwich K, i.e., $x_0 + rB_2^n \subseteq K \subseteq RB_2^n$ [4]. By *time complexity* of an algorithm, we refer to the total number of oracle calls and additional arithmetic operations (we focus on the dependence of the complexity on the dimension and suppress factors that depend polynomially on the size of the input (in particular $\log(R/r)$.)

Theorem 1.1. *There is a deterministic algorithm that, given any convex body $K \subset \mathbb{R}^n$ specified by a well-guaranteed membership oracle, finds an ellipsoid E such that $N(K,E)N(E,K) \leq 2^{O(n)}$. The time complexity of the algorithm is $2^{O(n)}$ and its space complexity is bounded by a polynomial in n.*

In [14], we gave a deterministic algorithm based on computing an approximate minimum mean-width ellipsoid (or ℓ-ellipsoid, see Section 2.1). The resulting covering bound was $N(K,E)N(E,K) = O(\log n)^n$ rather than the optimal $2^{O(n)}$, and the complexity of the algorithm was also $O(\log n)^n$.

Our approach here is to completely algorithmicize Milman's original existence proof and thereby obtain the best possible deterministic complexity of $2^{O(n)}$. By adjusting the parameters in the resulting algorithm to "slow down" Milman's iteration, we get the optimal trade-off between approximation and complexity for volume computation.

1.1 Deterministic volume computation

The first consequence is for estimating the volume of a convex body. This is an ancient problem that has led to many developments in algorithmic techniques, high-dimensional geometry and probability theory. On one hand, the problem can be solved for any convex body presented in the general membership oracle model in randomized polynomial time to arbitrary accuracy [15]. On the other hand, the following lower bound (improving on [16]) shows that deterministic algorithms cannot achieve such approximations.

Theorem 1.2. *[12] Suppose there is a deterministic algorithm that takes as input a symmetric convex body K satisfying $B_1^n \subset K \subset B_2^n$ and outputs $A(K), B(K)$ such that $A(K) \leq \mathrm{vol}(K) \leq B(K)$ and makes at* *most n^a calls to the membership oracle for K. Then there is some convex body K for which*

$$\frac{B(K)}{A(K)} \geq \left(\frac{cn}{a \log n}\right)^{n/2}$$

where c is an absolute constant.

In particular, the theorem implies that achieving even a $2^{O(n)}$ approximation requires $2^{\Omega(n)}$ oracle calls. The volume of an M-ellipsoid E of K is clearly within a factor of $2^{O(n)}$ of the volume of K. Thus Theorem 1.1 gives a $2^{O(n)}$ algorithm that achieves this volume approximation by computing the volume of the M-ellipsoid found. We state this consequence formally.

Theorem 1.3. *There is a deterministic algorithm of time complexity $2^{O(n)}$ and polynomial space complexity that estimates the volume of a convex body given by a well-guaranteed membership oracle to within a factor of $2^{O(n)}$.*

What is the complexity of achieving a smaller approximation factor for the volume? The following result of Barany and Furedi [17] gives a lower bound.

Theorem 1.4. *[17] For any $0 \leq \epsilon \leq 1$, any deterministic algorithm that estimates the volume of any input convex body to within a $(1+\epsilon)^n$ given only a membership oracle to the body, must make at least $\Omega(1/\epsilon)^{n/2}$ queries to the membership oracle.*

We show that our M-ellipsoid algorithm can be modified to obtain an algorithm that essentially matches this best possible complexity-vs-approximation trade-off for centrally symmetric convex bodies.

Theorem 1.5. *For any $0 \leq \epsilon \leq 1$, there is a deterministic algorithm that computes a $(1+\epsilon)^n$ approximation of the volume of a centrally symmetric convex body given by a well-guaranteed membership oracle in $O(1/\epsilon)^{O(n)}$ time and polynomial space.*

1.2 Deterministic lattice algorithms

Efficient M-ellipsoid construction also has consequences for central lattice problems. We define these problems next. For a convex body $K \subseteq \mathbb{R}^n$ containing the origin, the gauge function of K is

$$\|x\|_K = \inf\{s \geq 0 : x \in sK\}$$

for $x \in \mathbb{R}^n$. For symmetric K (i.e. $K = -K$), $\|\cdot\|_K$ is a usual norm on $\mathbb{R}^n$ (we shall refer to $\|\cdot\|_K$ as the norm induced by K and specify asymmetric whenever relevant). We say that K is *well-centered* if $\mathrm{vol}(K \cap -K) \geq 2^{-O(n)}\mathrm{vol}(K)$ (every convex body is well-centered with respect to its centroid or a point sufficiently close to its centroid).

The Shortest Vector Problem (SVP) is stated as follows: given an n-dimensional lattice L represented by a basis, and a convex body K, find a nonzero $v \in L$ such that $\|v\|_K$ is minimized. In the Closest Vector Problem (CVP), in addition to a lattice and a convex body, we are also given a query point x in $\mathbb{R}^n$, and the goal is to find a vector $v \in L$ that minimizes $\|x - v\|_K$. These problems are central to the geometry of numbers and have applications to integer programming, factoring polynomials, cryptography, etc.

The AKS sieve [18, 19] can be used to solve SVP in *randomized* $2^{O(n)}$ time, also using exponential space and randomness. Finding a deterministic algorithm of this complexity has been an important open problem. In a breakthrough paper, Micciancio and Voulgaris [20] gave deterministic $2^{O(n)}$ algorithms for

SVP and *exact* CVP in the Euclidean norm. The focus then shifted to extending these results to general norms as in the AKS-sieve based randomized algorithms.

Subsequently, [11] gave a reduction from general norm SVP to CVP in the Euclidean norm (or more specifically, to enumerating lattice points in ellipsoids), and thereby avail the algorithm of [20]. The reduction uses a $2^{O(n)}$ space and poly(n) randomness, improving on the AKS sieve, and gives an expected running time of $2^{O(n)}$ for general norm SVP. We now state the main part of the reduction precisely as it is useful for deterministic algorithms as well. For a lattice L and convex body K in $\mathbb{R}^n$, let $G(K, L)$ be the largest number of lattice points contained in any translate of K, i.e.,

$$G(K, L) = \max_{x \in \mathbb{R}^n} |(K + x) \cap L|. \tag{1.1}$$

The main result of [11], using the algorithm of [20], can be stated as follows.

Theorem 1.6. *[11] Given any convex body $K \subseteq \mathbb{R}^n$ along with an ellipsoid E of K and any n-dimensional lattice $L \subseteq \mathbb{R}^n$, the set $K \cap L$ can be computed in deterministic time $G(K, L) \cdot N(K, E)N(E, K) \cdot 2^{O(n)}$.*

For an M-ellipsoid E of K, the numbers $N(K, E)$ and $N(E, K)$ are both bounded by $2^{O(n)}$. From Theorem 1.1, we obtain the folllowing corollary.

Corollary 1.7. *Given any convex body $K \subseteq \mathbb{R}^n$ and any n-dimensional lattice $L \subseteq \mathbb{R}^n$, the set $K \cap L$ can be computed deterministically in time $G(K, L) \cdot 2^{O(n)}$.*

For SVP in any norm, a simple packing argument [11] shows that $G(\lambda_1 K, L) = 2^{O(n)}$, where $\lambda_1 = \inf_L \|x\|_K$, the length of the shortest nonzero vector in L, giving us the following result.

Theorem 1.8. *Given a basis for a lattice L and a well-centered convex body K, both in $\mathbb{R}^n$, the shortest vector in L under the norm $\|.\|_K$ can be found deterministically using $2^{O(n)}$ time and space.*

The reduction from [11] can also be used for a special case CVP in any norm called the *bounded distance decoding* problem. Here one assumes that the distance to the lattice of the query point is bounded by some factor γ times the length of the shortest nonzero lattice vector. In this case, $G(\gamma\lambda_1 K, L) = (2 + \gamma)^{O(n)}$ and we obtain a deterministic $(2 + \gamma)^{O(n)}$ algorithm.

Theorem 1.9. *Given a basis for a lattice L, any well-centered n-dimensional convex body K and a query point x in $\mathbb{R}^n$, the closest vector in L to x in the norm $\|.\|_K$ defined by K can be computed deterministically using $(2 + \gamma)^{O(n)}$ time and space, provided that the minimum distance is at most γ times the length of the shortest nonzero vector of L under $\|\cdot\|_K$.*

It remains open to solve CVP deterministically in time $2^{O(n)}$ with no assumptions on the minimum distance. Even the special case of CVP under any norm other than the Euclidean norm is open.

2 Techniques from convex geometry

2.1 The Lewis ellipsoid

Let $\alpha : \mathbb{R}^{n \times n} \to \mathbb{R}_+$ be a norm on $n \times n$ real matrices. We define the dual norm α^* for any $S \in \mathbb{R}^{n \times n}$ as

$$\alpha^*(S) = \sup\{\mathrm{tr}(SA) \, : \, A \in \mathbb{R}^{n \times n}, \alpha(A) \leq 1\}. \tag{2.1}$$

For a matrix $A \in \mathbb{R}^{n \times n}$, we denote its transpose by A^T, its inverse (when it exists) by A^{-1} and $A^{-T} = (A^{-1})^T$ and its trace by $\mathrm{tr}(A) = \sum_{i=1}^n A_{ii}$.

Theorem 2.1. *[21] For any norm α on $\mathbb{R}^{n\times n}$, there is an invertible linear transformation $A \in \mathbb{R}^{n\times n}$ such that*

$$\alpha(A) = 1 \text{ and } \alpha^*(A^{-1}) = n.$$

The ellipsoid AB_2^n corresponding to the optimal matrix A for a norm α is called the *Lewis* ellipsoid for α. The proof of the above theorem is based on examining the properties of the optimal solution to the following optimization problem:

$$\begin{aligned} \max \det(A) \text{ s.t.} \\ A \in \mathbb{R}^{n\times n} \\ \alpha(A) \leq 1 \end{aligned} \tag{2.2}$$

Lewis showed that the optimal A satisfies $\alpha^*(A^{-1}) = n$ by a simple variational argument (which we use later in Lemma 3.2).

We will be interested in norms α of the following form. Let $K \subseteq \mathbb{R}^n$ denote a symmetric convex body with associated norm $\|\cdot\|_K$, and let γ_n denote the canonical Gaussian measure on $\mathbb{R}^n$. We define the ℓ-norm with respect to K for $A \in \mathbb{R}^{n\times n}$ as

$$\ell_K(A) = \left(\int \|Ax\|_K^2 d\gamma_n(x)\right)^{\frac{1}{2}}.$$

The ℓ-norm was first studied and defined by Tomczak-Jaegermann and Figiel [22]. Roughly speaking, one can think of the ℓ-ellipsoid as the largest ellipsoid with the property that half of its volume is contained in K [8]. The ℓ-norm with respect to the polar $K^* = \{x \in \mathbb{R}^n : \langle x, y\rangle \leq 1 \; \forall y \in K\}$ is then:

$$\ell_{K^*}(A) = \left(\int \|Ax\|_{K^*}^2 d\gamma_n(x)\right)^{\frac{1}{2}}.$$

The norm and dual norm with respect to a convex body K are related by the "roundness" of K, as measured by its Banach-Mazur distance to a Euclidean ball. We recall the latter, then state the connection precisely. For two convex bodies $K, L \subseteq \mathbb{R}^n$ the Banach-Mazur distance between K and L is

$$d_{BM}(K, L) = \inf\{s : s \geq 1, \exists x \in \mathbb{R}^n, T \in \mathbb{R}^{n\times n} \text{ invertible }, TK \subseteq L - x \subseteq sTK\}$$

In words, it is the minimum dilation s such that there is some point x and transformation T for which the set $L - x$ is sandwiched between TK and sTK. The next lemma plays an important role.

Lemma 2.2. *[8] For $A \in \mathbb{R}^{n\times n}$, we have that*

$$\ell_K^*(A) \leq \ell_{K^*}(A^T) \leq 4(1 + \log d_{BM}(K, B_2^n))\ell_K^*(A).$$

As shown by Pisier [8], one can think of the ℓ-ellipsoid as the largest ellipsoid with the property that half of its volume is contained in K.

2.2 Covering numbers and volume estimates

Let $B_2^n \subseteq \mathbb{R}^n$ denote the n-dimensional Euclidean ball. Recall that $N(K, D)$ is the number of translates of D required to cover K. The following bounds for convex bodies $K, D \subset \mathbb{R}^n$ are classical. We use c, C to denote absolute constants throughout the paper.

Lemma 2.3. *For any two symmetric convex bodies K, D,*

$$\frac{\mathrm{vol}(K)}{\mathrm{vol}(K \cap D)} \leq N(K, D) \leq 3^n \frac{\mathrm{vol}(K)}{\mathrm{vol}(K \cap D)}.$$

The next lemma is from [23].

Lemma 2.4. *Let $D \subseteq \beta K$, $\beta \geq 1$. Then,*

$$\mathrm{vol}(\mathrm{conv}\,\{K, D\}) \leq 4\beta n N(D, K)\mathrm{vol}(K).$$

The following are the Sudakov and dual Sudakov inequalities (see e.g., Section 6 of [24]).

Lemma 2.5 (Sudakov Inequality)**.** *For any $t > 0$, convex body $K \subset \mathbb{R}^n$ and invertible matrix $A \in \mathbb{R}^{n\times n}$*

$$N(K, tAB_2^n) \leq e^{C\ell_{K^*}(A^{-T})^2/t^2}.$$

Lemma 2.6 (Dual Sudakov Inequality)**.** *For any $t > 0$, and $A \in \mathbb{R}^{n\times n}$*

$$N(AB_2^n, tK) \leq e^{C\ell_K(A)^2/t^2}.$$

The following lemma gives a simple containment relationship.

Lemma 2.7. *For any $A \in \mathbb{R}^{n\times n}$, A invertible, we have that*

$$\frac{1}{\ell_{K^*}(A^{-1})} K \subseteq AB_2^n \subseteq \ell_K(A)K$$

Proof. We first show that $E = AB_2^n \subseteq \ell_K(A)K$. Assume not, then there exists $x \in E$ such that

$$\|x\|_K = \sup_{y\in K^*} |\langle y, x\rangle| > \ell_K(A).$$

Let $y \in K^*$ be such that $|\langle y, x\rangle| = \|x\|_K$. Then we have

$$\ell_K(A) < |\langle x, y\rangle| \leq \sup_{z\in AB_2^n} |\langle z, y\rangle| = \sup_{z\in B_2^n} |\langle z, A^T y\rangle| = \|A^T y\|_2$$

But now note that

$$\ell_K(A) = \mathrm{E}[\|AX\|_K^2]^{\frac{1}{2}} \geq \mathrm{E}[|\langle y, AX\rangle|^2]^{\frac{1}{2}} = \|A^T y\|_2,$$

a contradiction. Therefore $AB_2^n \subseteq \ell_K(A)K$ as needed. Now applying the same argument on $E^* = A^{-1}B_2^n$ and K^*, we get that $E^* \subseteq \ell_{K^*}(A^{-1})K^*$. From here via duality ($K \subseteq L \to L^* \subseteq K^*$), we get that

$$\frac{1}{\ell_{K^*}(A^{-1})} K = (\ell_{K^*}(A^{-1})K^*)^* \subseteq (A^{-1}B_2^n)^* = AB_2^n$$

as needed. □

2.3 Approximating the ℓ-norm

In our algorithm we will need to approximate the integral defining the ℓ_K norm by a finite sum. Our approximation of the ℓ_K norm is defined as follows:

$$\tilde{\ell}_K(A) = \sum_{x \in \{-1,1\}^n} \frac{1}{2^n} \|Ax\|_K.$$

The next lemma is essentially folklore, we give a known proof here.

Lemma 2.8. *For a symmetric convex body K and any $A \in \mathbb{R}^{n \times n}$, we have*

$$\ell_K(A) \leq \sqrt{8\pi}(1 + \log d_{BM}(K, B_2^n))\tilde{\ell}_K(A).$$

Proof. Let $g_1, \dots, g_n$ denote i.i.d. $N(0,1)$ Gaussians, let $u_1, \dots, u_n$ denote i.i.d. uniform $\{-1,1\}$ random variables and let $A_1, \dots, A_n \in \mathbb{R}^n$ denote the columns of A. Then we have that

$$\begin{aligned}
\ell_K(A) &\leq 4(1 + \log d_{BM}(K, B_2^n)) \sup \left\{ \sum_i \langle A_i, y_i \rangle : \mathrm{E}[\| \sum_i g_i y_i \|_{K^*}^2]^{\frac{1}{2}} \leq 1 \right\} \qquad \text{(using Lemma 2.2)} \\
&\leq 4\sqrt{\frac{\pi}{2}}(1 + \log d_{BM}(K, B_2^n)) \ \sup \left\{ \sum_i \langle A_i, y_i \rangle : \mathrm{E}[\| \sum_i u_i y_i \|_{K^*}^2]^{\frac{1}{2}} \leq 1 \right\} \\
&\leq 4\sqrt{\frac{\pi}{2}}(1 + \log d_{BM}(K, B_2^n)) \ \mathrm{E}[\| \sum_i u_i A_i \|_K^2]^{\frac{1}{2}} = 4\sqrt{\frac{\pi}{2}}(1 + \log d_{BM}(K, B_2^n)) \ \tilde{\ell}_K(A)
\end{aligned}$$

The second inequality follows from the classical comparison $\mathrm{E}[f(u_1, \dots, u_n)] \leq \mathrm{E}[f(\sqrt{\frac{\pi}{2}}g_1, \dots, \sqrt{\frac{\pi}{2}}g_n)]$ for any convex function $f : \mathbb{R}^n \to \mathbb{R}$, and setting $f(x_1, \dots, x_n) = \| \sum_i x_i y_i \|_C^2$. The last inequality follows from the following weak duality relation:

$$\begin{aligned}
\sum_i \langle A_i, y_i \rangle &= \mathrm{E}[\langle \sum_i u_i A_i, \sum_j u_j y_j \rangle] \leq \mathrm{E}[\| \sum_i u_i A_i \|_K \| \sum_j u_j y_j \|_{K^*}] \\
&\leq \mathrm{E}[\| \sum_i u_i A_i \|_K^2]^{\frac{1}{2}} \, \mathrm{E}[\| \sum_j y_j u_j \|_{K^*}^2]^{\frac{1}{2}} \leq \ell_K(A).
\end{aligned}$$

□

The next lemma is a strengthening due to Pisier, using Proposition 8 from [25]. While it is not critical for our results (the difference is only in absolute constants), we use this stronger bound in our analysis.

Lemma 2.9. *For a symmetric convex body K and any $A \in \mathbb{R}^{n \times n}$, we have*

$$\sqrt{\frac{2}{\pi}}\tilde{\ell}_K(A) \leq \ell_K(A) \leq c_1\sqrt{1 + \log d_{BM}(K, B_2^n)}\, \tilde{\ell}_K(A)$$

where c_0, c_1 are absolute constants. Furthermore, by duality, we get that

$$\frac{1}{c_1\sqrt{1 + \log d_{BM}(K, B_2^n)}}\tilde{\ell}_K^*(A) \leq \ell_K^*(A) \leq \sqrt{\frac{\pi}{2}}\tilde{\ell}^*(A).$$

3 Algorithm for computing an M-ellipsoid

In this section, we present the algorithm for computing an M-ellipsoid of an arbitrary convex body given in the oracle model. We first observe that it suffices to give an algorithm for centrally symmetric K. For a general convex body K, we may replace K by the difference body $K - K$ (which is symmetric). An M-ellipsoid for $K - K$ remains one for K, as the covering estimates changes by at most a $2^{O(n)}$ factor. To see this, note that for any ellipsoid E we have that $N(K, E) \leq N(K - K, E)$ and that

$$N(E, K) \leq N(E, K - K)N(K - K, K) \leq N(E, K - K)N(K - K, K \cap -K) \leq N(E, K - K)2^{O(n)},$$

where the last inequality follows by using Lemma 2.3 and the Rogers-Shephard inequality [26], i.e. $\mathrm{vol}(K - K) \leq 4^n \mathrm{vol}(K)$.

Our algorithm has two main components: a subroutine to compute an approximate Lewis ellipsoid for a norm given by a convex body, and an implementation of the iteration that makes this ellipsoid converge to an M-ellipsoid of the original convex body. To compute the approximate ℓ-ellipsoid we use the following convex program:

$$\begin{aligned} \max \det(A)^{\frac{1}{n}} s.t.& \\ A &\in \mathbb{R}^{n \times n}, \text{ symmetric} \\ A &\succeq 0 \\ \tilde{\ell}_K(A) &\leq 1 \end{aligned} \tag{3.1}$$

Here $A \succeq 0$ denotes the constraint that the real symmetric matrix A is positive semidefinite, i.e., all its eigenvalues are nonnegative. So, in contrast to Lewis' program (2.2), we optimize over only symmetric positive semidefinite matrices. Another important difference is that we have replaced the ℓ-norm with $\tilde{\ell}_K$. to make the objective function computable.

With these changes, we can solve the convex program to arbitrary accuracy in polynomial time using the ellipsoid algorithm [4]. The main theorem from [4] is that a convex function can be minimized over a convex body given by a well-guaranteed membership oracle to within accuracy ϵ so that the number of calls to the oracle is polynomial in n, the size of the input representation and $\log(1/\epsilon)$. In more detail, the set S over which we optimize is convex and the logarithm of the objective function is concave over this set. In addition, it is not hard to find a feasible starting point in the set and derive bounds on the radii of balls that sandwich the set (see e.g., [14]). A membership oracle for the feasible region is straightforward: given any $n \times n$ real matrix X, we can verify that it is symmetric positive semidefinite and that its $\tilde{\ell}_K$ norm is at most 1 (in fact we can obtain a *separation oracle* for S, i.e., one that provides a hyperlane that separates an infeasible X from S). We will find approximately optimal solutions to this convex program applied to a series of convex bodies, and ensure that we have an efficient oracle for each one, given only the oracle for the initial body K.

Given a centrally-symmetric convex body K, as a pre-processing step, we put it in approximate John position using the Ellipsoid algorithm in polynomial time [4], so that $B_2^n \subseteq K \subseteq nB_2^n$, i.e., $d_{BM}(K, B_2^n) \leq n$. We then use the M-ellipsoid Algorithm described next. By $\log^{(i)}(n)$ we mean the i'th iterated logarithm, i.e., $\log^{(1)}(n) = \log n, \log^{(2)} n = \log\log n$ and so on.

M-ellipsoid Algorithm

```
1. Let $K_1 = K$ and $T = \log^* n$
2. For $i = 1 \dots T - 1$,
    (a) Compute an approximate $\ell$-ellipsoid of $K_i$ using the convex program (3.1) to get an approximately optimal transformation $A_i$ (the corresponding ellipsoid is $A_i B_2^n$).
    (b) Set
        $r_{in} = \frac{\sqrt{n}}{\log^{(i)}(n)\tilde{\ell}_{K_i}(A_i)}$ and $r_{out} = \log^{(i)}(n)\frac{\tilde{\ell}_{K_i^*}(A_i^{-1})}{\sqrt{n}}$.
    (c) Define
        $K_{i+1} = \mathrm{conv}\{K_i \cap r_{out} A_i B_2^n, r_{in} A_i B_2^n\}$.
3. Output $E = \frac{\sqrt{n}}{\tilde{\ell}_{K_{T-1}}(A_{T-1})} A_{T-1} B_2^n$ as the M-ellipsoid.
```

This is essentially an algorithmic version of Milman's proof of the existence of M-ellipsoids. We try to construct a good ellipsoid for the original body K. However, its quality depends on $d_{BM}(K, B_2^n)$, which can be high in the beginning. Each iteration then constructs a more "round" version of K by taking the convex hull of two bodies, the first is the restriction of the current K to a not-too-large ball, and the second is a smaller ball contained in the first. Thus the new Banach-Mazur distance of K to the unit ball is bounded by the ratio of the radii of these balls which we will maintain as at most polylogarithmic in the previous ratio. Finally, given a good ellipsoid for the new body, i.e., one with small covering number, we will see that only needs a relatively small number of copies of it to cover the original body, i.e., we keep the ratio of volumes bounded. Since the "roundness" is dropping so quickly, the total number of iterations is small and the total blow-up in volume ratios is also small. We formally prove all these properties of the algorithm in the next section.

3.1 Analysis

We note that the time complexity of the algorithm is bounded by $\mathrm{poly}(n)2^{O(n)}$ and the space complexity is polynomial in n. In fact, the only step that takes exponential time is the evaluation of the ℓ-norm constraint of the semidefinite program. This evaluation happens a polynomial number of times. The rest of computation involves applying the ellipsoid algorithm and computing oracles for successive bodies (i.e., an oracle for K_{i+1} given an oracle for K_i). Given membership oracles for two convex bodies A, B, we can build membership oracles for their intersection $A \cap B$ and for their convex hull $\mathrm{conv}\{A, B\}$ [4]. These oracles use only a polynomial (in n) number of calls to the oracles for A and B. The complexity of the oracle grows as $n^{O(i)}$ in the i'th iteration, for a maximum of $n^{O(\log^* n)} = 2^{o(n)}$.

A well-guaranteed oracle for a convex body consists of a membership oracle and a bound on the ratio between two balls that sandwich the body. Our analysis below includes the sandwiching ratio, which gets smaller with each iteration.

We begin by showing that Lewis' optimality condition (Theorem 2.1) is robust to approximation and works when restricted to positive semidefinite transformations. This allows us to establish the desired properties for approximate optimizers of the convex program (3.1). Following this, we will show that the algorithm, with the property established for approximately optimal solutions, finds an M-ellipsoid of the original body.

3.1.1 Approximate Lewis ellipsoids

The main statement of this section is the following.

Theorem 3.1. *Let A be a $(1-\epsilon)$-approximate optimizer to the convex program (3.1) for $\epsilon \leq 1/(36n^4)$. Then*

$$\ell_K(A)\ell_{K^*}(A^{-1}) \leq Cn(1+\log d_{BM}(K, B_2^n))^{\frac{3}{2}}.$$

for an absolute constant $C > 0$.

The proof of the theorem is based on the next lemma. For a matrix T, recall that $\|T\|_F = \sqrt{\sum_{i,j} T_{ij}^2}$ is its Frobenius norm, and $\|T\|_2 = \sup_{x \in B_2^n} \|Tx\|_2$ is the operator norm.

Lemma 3.2. *Let K be such that $B_2^n \subseteq K \subseteq nB_2^n$ and $A \in \mathbb{R}^{n\times n}$, be a $(1-\epsilon)$-approximate optimizer for the convex program (3.1), i.e. $\det(A)^{\frac{1}{n}} \geq (1-\epsilon)OPT$. Then for $\epsilon \leq 1/36n^4$, we have that*

$$\tilde{\ell}_K(A)\tilde{\ell}_K^*(A^{-1}) \leq n(1+6n^2\sqrt{\epsilon}) \leq 2n.$$

Proof. For simplicity of notation, we write $\tilde{\ell}_K(T)$ as $\alpha(T)$ for $T \in \mathbb{R}^{n\times n}$. Take $T \in \mathbb{R}^{n\times n}$ (not necessarily positive semidefinite) satisfying $\alpha(T) \leq 1$.

First note that $I_n/\alpha(I_n)$ is a feasible solution to (3.1) satisfying

$$\det(\frac{I_n}{\alpha(I_n)})^{\frac{1}{n}} = \frac{1}{\alpha(I_n)} \geq \frac{1}{\|I_n\|_F} = \frac{1}{\sqrt{n}}.$$

Let $A_{OPT} \succeq 0$ denote the optimal solution to (3.1). Since $\det(A_{OPT}) \geq \frac{1}{\sqrt{n}}$, we clearly have that $A_{OPT} \succ 0$. Therefore for $\delta > 0$ small enough we have that $A_{OPT} + \delta T \succeq 0$. From this, we see that $(A_{OPT} + \delta T)/\alpha(A_{OPT} + \delta T)$ is also feasible for (3.1) as $\alpha((A_{OPT} + \delta T)/\alpha(A_{OPT} + \delta T)) = 1$. Since A_{OPT} is the optimal solution, we have that

$$\det\left(\frac{A_{OPT} + \delta T}{\alpha(A_{OPT} + \delta T)}\right)^{\frac{1}{n}} \leq \det(A_{OPT})^{\frac{1}{n}}.$$

Rewriting this and using the triangle inequality,

$$\begin{aligned}\det(A_{OPT} + \delta T)^{\frac{1}{n}} &\leq \det(A_{OPT})^{\frac{1}{n}}\alpha(A_{OPT} + \delta T) \leq \det(A_{OPT})^{\frac{1}{n}}(\alpha(A_{OPT}) + \delta\alpha(T)) \\ &\leq \det(A_{OPT})^{\frac{1}{n}}(1+\delta).\end{aligned}$$

Dividing by $\det(A_{OPT})^{\frac{1}{n}}$ on both sides, we get that

$$\det(I_n + \delta A_{OPT}^{-1}T)^{\frac{1}{n}} \leq 1 + \delta. \tag{3.2}$$

Since both sides are equal at $\delta = 0$, we must have the same inequality for the derivatives with respect to δ at 0. This yields

$$\frac{1}{n}\mathrm{tr}(A_{OPT}^{-1}T) \leq 1 \Leftrightarrow \mathrm{tr}(A_{OPT}^{-1}T) \leq n \tag{3.3}$$

Up to this point the proof is essentially the same as Lewis' proof of Theorem 2.1. We now depart from that proof to account for approximately optimal solutions. We will use the following three claims.

Claim 1. $\alpha(T) \leq \|T\|_F \leq n\alpha(T)$.

Proof. (of Claim 1.) Let U denote a uniform vector in $\{-1,1\}^n$. Since $\frac{1}{n}\|x\|_2 \leq \|x\|_K$ for any $x \in \mathbb{R}^n$, we have that

$$\alpha(T) = \mathrm{E}[\|UT\|_K^2]^{\frac{1}{2}} \geq \frac{1}{n}\mathrm{E}[\|UT\|_2^2]^{\frac{1}{2}} = \frac{1}{n}\|T\|_F.$$

Now using the inequality $\|x\|_K \leq \|x\|_2$ for $x \in \mathbb{R}^n$, a similar argument yields $\alpha(T) \leq \|T\|_F$. □

Claim 2. $\|A_{OPT}^{-1}\|_2 \leq n$.

Proof. (of Claim 2.) Let σ denote the largest eigenvalue of A_{OPT}^{-1} and $v \in \mathbb{R}^n$ be an associated unit eigenvector. Since $A_{OPT} \succ 0$, we have that $A_{OPT}^{-1} \succ 0$, and hence $\sigma = \|A^{-1}\|_2$. Now note that $A_{OPT} + \delta vv^T \succ 0$ for any $\delta \geq 0$, and that $\alpha(vv^T) \leq \|vv^T\|_F = \|v\|_2^2 = 1$. Therefore by Equation (3.3), we have that

$$n \geq \mathrm{tr}(A^{-1}(vv^T)) = \mathrm{tr}(\sigma vv^T) = \sigma$$

as needed. □

Claim 3. $A^{-1} \preceq (1 + 6\sqrt{n\epsilon})A_{OPT}^{-1}$.

We can now complete the proof of the lemma (we will prove the last claim presently). Take $T \in \mathbb{R}^{n\times n}$ satisfying $\alpha(T) \leq 1$. By Claim 1, we note that $\|T\|_F \leq n\alpha(T) \leq n$. Now by Equation (3.3), we have that

$$\mathrm{tr}(A^{-1}T) = \mathrm{tr}(A_{OPT}^{-1}T) + \mathrm{tr}((A^{-1} - A_{OPT}^{-1})T) \leq n + \|A^{-1} - A_{OPT}^{-1}\|_F\|T\|_F \leq n + n\|A^{-1} - A_{OPT}^{-1}\|_F$$

We bound the second term using the Claim 3. Since $A^{-1} \preceq (1+6\sqrt{n\epsilon})A_{OPT}^{-1}$, we have that $A^{-1} - A_{OPT}^{-1} \preceq 6\sqrt{n\epsilon}A_{OPT}^{-1}$, and hence, using Claim 2,

$$\|A^{-1} - A_{OPT}^{-1}\|_F \leq \sqrt{n}\|A^{-1} - A_{OPT}^{-1}\|_2 \leq 6n\sqrt{\epsilon}\|A_{OPT}^{-1}\|_2 \leq 6n^2\sqrt{\epsilon}$$

Using this bound, we get

$$\mathrm{tr}(A^{-1}T) \leq n + 6n^3\sqrt{\epsilon} = n(1 + 6n^2\sqrt{\epsilon})$$

for any $T \in \mathbb{R}^{n\times n}$ satisfying $\alpha(T) \leq 1$. Thus we get that $\alpha^*(A^{-1}) \leq n\left(1 + 6n^2\sqrt{\epsilon}\right)$. Together with the constraint $\alpha(A) \leq 1$, the conclusion of the lemma follows. It remains to prove Claim 3.

Proof. (of Claim 3.) Since A is a $(1-\epsilon)$-approximate maximizer to (3.1) we have that

$$\det(A)^{\frac{1}{n}} \geq (1-\epsilon)\det(A_{OPT})^{\frac{1}{n}} \Rightarrow \det(A) \geq (1-n\epsilon)\det(A_{OPT})$$

We begin by proving by proving $A \succeq (1 - 3\sqrt{n\epsilon})A_{OPT}$. Now note that

$$A \succeq (1 - 3\sqrt{n\epsilon})A_{OPT} \quad \Leftrightarrow \quad A_{OPT}^{-\frac{1}{2}}AA_{OPT}^{-\frac{1}{2}} \succeq (1 - 3\sqrt{n\epsilon})I_n$$

Hence letting $B = A_{OPT}^{-\frac{1}{2}}AA_{OPT}^{-\frac{1}{2}}$, it suffices to show that $B \succeq (1 - 3\sqrt{n\epsilon})I_n$. From here, we note that $1 \geq \det(B) = \det(A)/\det(A_{OPT}) \geq (1 - n\epsilon)$. Now from Equation (3.3), we have that

$$\mathrm{tr}(B) = \mathrm{tr}(A_{OPT}^{-\frac{1}{2}}AA_{OPT}^{-\frac{1}{2}}) = \mathrm{tr}(A_{OPT}^{-1}A) \leq n$$

Let $\sigma_1, \dots, \sigma_n$ denote the eigenvalues of B in non-increasing order. We first note that $\sigma_n \leq 1$ since otherwise

$$\det(B) = \prod_{i=1}^{n} \sigma_i \geq \sigma_n^n > 1$$

a contradiction. Furthermore, since $B \succ 0$, we have that $0 < \sigma_n \leq 1$. So we may write $\sigma_n = 1 - \epsilon_0$, for $1 > \epsilon_0 \geq 0$. Now since $\sum_{i=1}^n \sigma_i = \operatorname{tr}(B) \leq n$, by the arithmetic mean - geometric mean inequality we have that

$$\det(B) = \sigma_n \prod_{i=1}^{n-1} \sigma_i = (1-\epsilon_0) \prod_{i=1}^{n-1} \sigma_i \leq (1-\epsilon_0) \left(\frac{\sum_{i=1}^{n-1} \sigma_i}{n-1} \right)^{n-1} \leq (1-\epsilon_0)(1 + \frac{\epsilon_0}{n-1})^{n-1}$$

Using the inequality $1 + x \leq e^x \leq 1 + x + \frac{e-1}{2}x^2$ for $x \in [-1, 1]$, we get that

$$\begin{aligned}(1-\epsilon_0)(1+\frac{\epsilon_0}{n-1})^{n-1} &\leq (1-\epsilon_0)e^{\epsilon_0} \leq (1-\epsilon_0)(1+\epsilon_0+\frac{e-1}{2}\epsilon_0^2) \\ &= 1 - \frac{3-e}{2}\epsilon_0^2 - \frac{e-1}{2}\epsilon_0^3 \leq 1 - \frac{3-e}{2}\epsilon_0^2\end{aligned}$$

From this we have

$$1 - \frac{3-e}{2}\epsilon_0^2 \geq \det(B) \geq (1 - n\epsilon) \;\; \Rightarrow \;\; \epsilon_0 \leq \sqrt{\frac{2}{3-e} n\epsilon} \leq 3\sqrt{n\epsilon}$$

Therefore $\sigma_n = 1 - \epsilon_0 \geq 1 - 3\sqrt{n\epsilon} \Rightarrow B \succeq (1 - 3\sqrt{n\epsilon})I_n \Rightarrow A \succeq (1 - 3\sqrt{n\epsilon})A_{OPT}$ as needed. Hence,

$$A^{-1} \preceq \left(\frac{1}{1 - 3\sqrt{n\epsilon}} \right) A_{OPT}^{-1} \preceq (1 + 6\sqrt{n\epsilon})A_{OPT}^{-1}$$

for $\epsilon \leq 1/36n$, proving the claim. □

This completes the proof of the lemma. □

We can now prove Theorem 3.1.

Proof. (of Theorem 3.1.) Using Lemmas 2.2, 2.9 and 3.2 in that order, we have

$$\begin{aligned}\ell_K(A)\ell_{K^*}(A^{-1}) &\leq 4(1 + \log d_{BM}(K, B_2^n))\ell_K(A)\ell_K^*(A^{-1}) \\ &\leq C(1 + \log d_{BM}(K, B_2^n))^{\frac{3}{2}}\tilde{\ell}_K(A)\tilde{\ell}_K^*(A^{-1}) \\ &\leq 2Cn(1 + \log d_{BM}(K, B_2^n))^{\frac{3}{2}}.\end{aligned}$$

□

3.1.2 Convergence to an M-ellipsoid

Next we turn to proving that the algorithm produces an M-ellipsoid. While the analysis follows the existence proof to a large extent, we need to handle the various approximations incurred. To aid in the analysis of the M-ellipsoid Algorithm for input $K \subseteq \mathbb{R}^n$, we make some additional definitions. Let $a_i = \log^{(i)} n$ and

$T = \log^* n$. Let $K_1, \dots, K_T$ and $A_1, \dots, A_T$ denote the sequence of bodies and transformations generated by the algorithm. Set $K_1^{out} = K_1^{in} = K$, and for $1 \leq i \leq T-1$ define

$$K_{i+1}^{in} = \text{conv}\{K_i^{in}, r_{in}^i A_i B_2^n\} \quad K_{i+1}^{out} = K_i^{out} \cap r_{out}^i A_i B_2^n$$

where r_{in}^i, r_{out}^i are defined as r_{in}, r_{out} in the i'th iteration of the main loop of the M-ellipsoid Algorithm. Recall that

$$r_{in}^i = \frac{\sqrt{n}}{\log^{(i)}(n)\tilde{\ell}_{K_i}(A_i)} \text{ and } r_{out}^i = \log^{(i)}(n)\frac{\tilde{\ell}_{K_i^*}(A_i^{-1})}{\sqrt{n}}.$$

so that

$$\frac{r_{out}^i}{r_{in}^i} = \frac{a_i^2}{n}\tilde{\ell}_{K_i}(A_i)\tilde{\ell}_{K_i^*}(A_i^{-1}).$$

Thus, K_i^{in} contains a ball of radius r_{in}^i while K_i^{out} is contained in a ball of radius r_i^{out}. By construction, we have the relations

$$K \subseteq K_1^{in} \subseteq \cdots \subseteq K_T^{in}, \qquad K \supseteq K_1^{out} \supseteq \cdots \supseteq K_T^{out}, \qquad K_i^{out} \subseteq K_i \subseteq K_i^{in} \ \forall i \in [T]$$

The proof of the main theorem will be based on the following inductive lemmas which quantify the properties of the sequences of bodies defined above.

Lemma 3.3. *$\forall i \in [T]$, we have that $d_{BM}(K_i, B_2^n) \leq C(\log^{(i-1)} n)^{\frac{7}{2}}$.*

Proof. For the base case, we have that $d_{BM}(K_1, B_2^n) \leq \sqrt{n} \leq Cn^{\frac{7}{2}}$ for any constant $C \geq 1$.

For the general case, by construction of K_{i+1} we have that

$$r_{in}^i A_i B_2^n \subseteq K_{i+1} \subseteq r_{out}^i A_i B_2^n.$$

Therefore,

$$\begin{aligned} d_{BM}(K_{i+1}, B_2^n) &\leq \frac{r_{out}^i}{r_{in}^i} \\ &= \frac{a_i^2}{n}\tilde{\ell}_{K_i^*}(A_i^{-1})\tilde{\ell}_{K_i}(A_i) \\ &\leq \frac{C_1 a_i^2}{n}\ell_{K_i^*}(A^{-1})\ell_{K_i}(A_i) \quad \text{(by Lemma 2.9)} \\ &\leq C_2(\log^{(i)} n)^2(1+\log d_{BM}(K_i, B_2^n))^{\frac{3}{2}} \quad \text{(by Theorem3.1)}\,. \end{aligned}$$

Using the fact that $\log^{(i)}(n) \geq 1$, $\forall i \in [T-1]$, a direct computation shows that the above recurrence equation implies the existence of a constant $C > 1$ (depending only on C_1) such that the stated bound on $d_{BM}(K_{i+1}, B_2^n)$ holds. □

Lemma 3.4. *For $i \in [T-1]$, we have that*

$$\max\left\{\frac{\text{vol}(K_i^{out})}{\text{vol}(K_{i+1}^{out})}, \frac{\text{vol}(K_{i+1}^{in})}{\text{vol}(K_i^{in})}\right\} \leq e^{Cn/\log^{(i)} n}$$

Proof. By Lemma 2.3, the fact that $K_i^{out} \subseteq K_i$, Lemma 2.5, Lemma 2.9 and Lemma 3.3, we have that

$$\begin{aligned}\frac{\mathrm{vol}(K_i^{out})}{\mathrm{vol}(K_{i+1}^{out})} &\leq N(K_i^{out}, r_{out}^i A_i B_2^n) \leq N(K_i, r_{out}^i A_i B_2^n) \\ &\leq e^{C(\ell_{K_i^*}(A_i^{-1})/r_{out}^i)^2} = e^{Cn\ell_{K_i^*}(A_i^{-1})^2/(a_i\tilde{\ell}_{K_i^*}(A^{-1}))^2} \\ &\leq e^{Cn\log(d_{BM}(K_i^*,B_2^n))/a_i^2} \leq e^{Cn/\log^{(i)} n}\end{aligned}$$

By Lemma 2.7, 2.9 and 3.3, we see that

$$r_{in}^i A_i B_2^n \subseteq r_{in}^i \ell_{K_i^{in}}(A_i) K_i^{in} \subseteq r_{in}^i \ell_{K_i}(A_i) K_i^{in} \subseteq C_1\sqrt{n} K_i^{in}.$$

Next by Lemma 2.4, the fact that $K_i \subseteq K_i^{in}$, Lemma 2.6, Lemma 2.9 and Lemma 3.3, we have that

$$\begin{aligned}\frac{\mathrm{vol}(K_{i+1}^{in})}{\mathrm{vol}(K_i^{in})} &\leq C_1 4n^{\frac{3}{2}} N(r_{in}^i A_i B_2^n, K_i^{in}) \leq C_1 n^{\frac{3}{2}} N(r_{in}^i A_i B_2^n, K_i) \\ &\leq C_1 n^{\frac{3}{2}} e^{C(\ell_{K_i}(A_i) r_{in}^i)^2} = C_1 n^{\frac{3}{2}} e^{Cn\ell_{K_i}(A_i)^2/(a_i\tilde{\ell}_{K_i}(A_i))^2} \\ &\leq C_1 n^{\frac{3}{2}} e^{Cn\log(d_{BM}(K_i,B_2^n))/a_i^2} \leq C_1 n^{\frac{3}{2}} e^{Cn(1/\log^{(i)} n)} \leq e^{Cn/\log^{(i)} n}\end{aligned}$$

□

We are now ready to complete the proof.

Proof. (of Theorem 1.1.) By construction of K_T, we note that

$$r_{in}^{T-1} A_{T-1} B_2^n \subseteq K_T \subseteq r_{out}^{T-1} A_{T-1} B_2^n$$

where by Lemma 3.3 we have that $r_{out}^{T-1}/r_{in}^{T-1} = O(1)$. Therefore the returned ellipsoid $E = \frac{\sqrt{n}}{\tilde{\ell}_{K_{T-1}(A_{T-1})}} A_{T-1} B_2^n$ (last line of the M-ellipsoid Algorithm satisfies that

$$\frac{1}{C} E \subseteq K_T \subseteq CE$$

for an absolute constant $C \geq 1$. Next by Lemma 2.3, we have that

$$N(K, E), N(E, K) \leq 3^n \frac{\max\{\mathrm{vol}(K), \mathrm{vol}(E)\}}{\mathrm{vol}(K \cap E)}$$

Now we see that

$$K \subseteq K_T^{in} \subseteq CK_T^{in} \qquad E \subseteq CK_T \subseteq CK_T^{in},$$

and that

$$K \supseteq \frac{1}{C} K_T^{out} \qquad E \supseteq \frac{1}{C} K_T \supseteq \frac{1}{C} K_T^{out}.$$

Therefore,

$$\frac{\max\{\mathrm{vol}(K), \mathrm{vol}(E)\}}{\mathrm{vol}(K \cap E)} \leq C^{2n} \frac{\mathrm{vol}(K_T^{in})}{\mathrm{vol}(K_T^{out})}.$$

Finally, by Lemma 3.4 we have that

$$\frac{\mathrm{vol}(K_T^{in})}{\mathrm{vol}(K_T^{out})} = \prod_{i=1}^{T-1} \frac{\mathrm{vol}(K_{i+1}^{in})}{\mathrm{vol}(K_i^{in})} \frac{\mathrm{vol}(K_i^{out})}{\mathrm{vol}(K_{i+1}^{out})} \le \prod_{i=1}^{T-1} e^{2Cn/\log^{(i)} n} = 2^{O(n)}.$$

Combining the above inequalities yields the desired guarantee on the algorithm. The time complexity is $2^{O(n)}$, dominated by the time to evaluate the $\tilde{\ell}_K$-norm. The space is polynomial since all we need to maintain are efficient oracles for the successive bodies K_i, which can be done space-efficiently for the operations of intersection and convex hull used in the algorithm [4]. □

4 An asymptotically optimal volume algorithm

In this section, we show how to modify our M-ellipsoid algorithm to prove Theorem 1.5.

In the M-ellipsoid algorithm of the previous section, we construct a series of convex bodies $K_0 = K, K_1, \ldots, K_T$ such that the covering numbers $N(K, K_T)$ and $N(K_T, K)$ are bounded by $2^{O(n)}$ and the final body K_T has $d_{BM}(K_T, B_2^n) < C$ for some constant C. Our modification will construct a similar sequence of bodies, but rather than bounding covering numbers, we will ensure that

$$e^{-C\epsilon n}\mathrm{vol}(K) \le \mathrm{vol}(K_T) \le e^{C\epsilon n}\mathrm{vol}(K)$$

and

$$d_{BM}(K_T, B_2^n) \le C\frac{\ln(1/\epsilon)^{\frac{5}{2}}}{\epsilon^2}.$$

Then we approximate the volume of K_T by finding an approximate ℓ-ellipsoid E for it, and covering it with translations of a maximal parallelopiped that fits in ϵE. Here is the precise algorithm. The reader can see that it is similar to the iteration from the previous section, but applied at a slower rate.

```
**Deterministic Volume($K, \epsilon$).**

  1. Let $K_1 = K$ and $T = \log^* n$

  2. For $i = 1 \ldots T-1$,

      (a) Compute an approximate $\ell$-ellipsoid of $K_i$ using the convex program (3.1) to get an approximately optimal transformation $A_i$ (the corresponding ellipsoid is $A_i B_2^n$).

      (b) Set
          $$r_{in} = \frac{\epsilon\sqrt{n}}{\sqrt{\ln(1/\epsilon)}C\log^{(i)}(n)\tilde{\ell}_{K_i}(A_i)} \text{ and } r_{out} = \frac{C\sqrt{\ln(1/\epsilon)}\log^{(i)}(n)\tilde{\ell}_{K_i^*}(A_i^*)}{\epsilon\sqrt{n}}.$$

      (c) Define
          $$K_{i+1} = \mathrm{conv}\{K_i \cap r_{out}A_iB_2^n, r_{in}A_iB_2^n\}.$$

  3. Compute the ellipsoid $E = r_{in}A_{T-1}B_2^n$ and a maximum volume parallelopiped $P$ inscribed in $E$ (via the principal components of $A_{T-1}$).

  4. Cover $K_T$ with disjoint copies of $\epsilon P$. Output $k\mathrm{vol}(P)$ where $k$ is the number of copies used.
```

Proof of Theorem 1.5. Let $a_i = \log^{(i)} n$. As in Lemma 3.3, we bound the Banach Mazur via the following recurrence

$$d_{BM}(K_{i+1}, B_2^n) \leq \frac{r_{out}^i}{r_{in}^i} \leq C\frac{\ln(1/\epsilon)}{\epsilon^2}(\log^{(i)}(n))^2(1+\log d_{BM}(K_i, B_2^n))^{\frac{3}{2}}.$$

From the above recurrence a direct computation reveals that for $\forall\ i \in [T]$,

$$d_{BM}(K_i, B_2^n) \leq C\frac{\ln(1/\epsilon)^{5/2}}{\epsilon^2}(\log^{(i-1)}(n))^{\frac{7}{2}}$$

We now show that the volumes of the K_i bodies changes very slowly. This will enable us to conclude that the volume of K_T is very close to the volume of K.

By Lemmas 2.7, 2.9 and the above bound on $d_{BM}(K_i, B_2^n)$, we have that

$$r_{in}^i A_i B_2^n \subseteq r_{in}^i \ell_{K_i}(A_i) K_i \subseteq C\frac{\epsilon\sqrt{n \log d_{BM}(K_i, B_2^n)}}{\sqrt{\ln(1/\epsilon)}\log^{(i)}(n)} K_i \subseteq C\epsilon\sqrt{n} K_i$$

and that

$$r_{out}^i A_i B_2^n = C\frac{\sqrt{\ln(1/\epsilon)}\log^{(i)}(n)\tilde{\ell}_{K^*}(A^{-1})}{\epsilon\sqrt{n}} A_i B_2^n \supseteq C\frac{\ell_{K^*}(A^{-1})}{\epsilon\sqrt{n}} A_i B_2^n \supseteq C\frac{1}{\epsilon\sqrt{n}} K_i.$$

Therefore if $\epsilon \leq C/\sqrt{n}$, then $K_{i+1} = \mathrm{conv}\{r_{in}^i A_i B_2^n, K_i \cap r_{out}^i A_i B_2^n\} = K_i$. Since this holds for all $i \in [T-1]$, we get that $K_T = K$ and hence $\mathrm{vol}(K_T) = \mathrm{vol}(K)$.

Now assume that $\epsilon \geq C/\sqrt{n}$. Then for $i \in [T-1]$, using Lemmas 2.3 and 2.5, we have,

$$\begin{aligned}
\mathrm{vol}(K_{i+1}) &\geq \mathrm{vol}(K_i \cap r_{out} B_2^n) \\
&\geq \frac{\mathrm{vol}(K_i)}{N(K_i, r_{out}^i B_2^n)} \\
&\geq e^{-C(\ell_{K_i^*}(A_i^{-1})/r_{out}^i)^2}\mathrm{vol}(K_i) \\
&\geq e^{-C(\epsilon^2/\ln(1/\epsilon))n\log d_{BM}(K_i, B_2^n)/a_i^2}\mathrm{vol}(K_i) \\
&\geq e^{-Cn\epsilon/\log^{(i)}(n)}\mathrm{vol}(K_i).
\end{aligned}$$

From the above, we get that

$$\frac{\mathrm{vol}(K_T)}{\mathrm{vol}(K)} = \prod_{i=1}^{T-1}\frac{\mathrm{vol}(K_{i+1})}{\mathrm{vol}(K_i)} \geq \prod_{i=1}^{T-1} e^{-Cn\epsilon/\log^{(i)}(n)} \geq e^{-Cn\epsilon}$$

Next via Lemma 2.4, the above containment, and Lemma 2.5, we have,

$$\begin{aligned}
\mathrm{vol}(K_{i+1}) &\leq \mathrm{vol}(\mathrm{conv}\,\{K_i, r_{in} B_2^n\}) \\
&\leq C(\epsilon\sqrt{n})nN(r_{in}B_2^n, K_i)\mathrm{vol}(K_i) \\
&\leq C(\epsilon n^{\frac{3}{2}})e^{C(r_{in}^i \ell_K(A_i))^2}\mathrm{vol}(K_i) \\
&\leq C(\epsilon n^{\frac{3}{2}})e^{C(\epsilon^2/\ln(1/\epsilon))n\log d_{BM}(K_i, B_2^n)/a_i^2}\mathrm{vol}(K_i) \\
&\leq C(\epsilon n^{\frac{3}{2}})e^{Cn\epsilon/\log^{(i)}(n)}\mathrm{vol}(K_i).
\end{aligned}$$

From this, we get that

$$\frac{\mathrm{vol}(K_T)}{\mathrm{vol}(K)} = \prod_{i=1}^{T-1} \frac{\mathrm{vol}(K_{i+1})}{\mathrm{vol}(K_i)} \leq (C\epsilon n^{\frac{3}{2}})^{\log^*(n)} \prod_{i=1}^{T-1} e^{Cn\epsilon/\log^{(i)}(n)} \leq e^{Cn\epsilon},$$

where the above holds as long as $\epsilon = \Omega(\frac{\log n \log^* n}{n})$ (which we have by assumption).

Combining the above inequalities, we get

$$e^{-C\epsilon n}\mathrm{vol}(K) \leq \mathrm{vol}(K_T) \leq e^{C\epsilon n}\mathrm{vol}(K).$$

Let E denote the final ellipsoid computed by the algorithm, and let P denote a maximimum volume inscribed parallelipiped of E. By construction of E and K_T, we have that $E \subseteq K_T \subseteq C\frac{\ln(1/\epsilon)^{5/2}}{\epsilon^2}E$. Therefore the covering produced is contained in $K_T + \epsilon P \subseteq K_T + \epsilon E \subseteq (1+\epsilon)K_T$. Hence the estimate found by the algorithm lies between $\mathrm{vol}(K_T)$ and $\mathrm{vol}((1+\epsilon)K_T) = (1+\epsilon)^n\mathrm{vol}(K_T)$. Thus the overall approximation factor is bounded by $e^{Cn\epsilon}$ as desired.

Next we bound the size of the covering found by the algorithm in Step 4. Noting that $\mathrm{vol}(E) = 2^{O(n)}\mathrm{vol}(P)$, the size of the covering is bounded by

$$\frac{\mathrm{vol}(K_T + \epsilon P)}{\mathrm{vol}(P)} \leq (1+\epsilon)^n \frac{\mathrm{vol}(K_T)}{\mathrm{vol}(P)} \leq C^n(1+\epsilon)^n \frac{\mathrm{vol}(K_T)}{\mathrm{vol}(E)} \leq C^n(1+\epsilon)^n(\ln(1/\epsilon)^{5/2}/\epsilon^2)^n = (1/\epsilon)^{O(n)}.$$

Finally, we describe the enumeration procedure that will ensure that the time bound is $(1/\epsilon)^{O(n)}$ and the space used is polynomial in n. The number of parallelopipeds enumerated could be as high as $(1/\epsilon)^{O(n)}$. However, we do not need to store all the copies that intersect K, we only need the number. To do this using polynomial space, we start with a parallelopiped inside K designated as the *root* and fix an order on its axes. For every other parallelopiped in the axis-aligned tiling, designate its *parent* to be an adjacent node closer to the root in Manhattan distance along the axes of the parallelopiped (i.e., the usual L_1 distance for the centers of the parallelopipeds after transforming parallelopipeds to cuboids), breaking ties using the ordering on coordinates. This ensures that a traversal of the tree defined by this structure takes time linear in the number of nodes in the tree and space linear in the dimension. This is a special case of a more general space-efficient traversal technique studied by Avis and Fukuda [27]. □

Acknowledgements. We are very grateful to Assaf Naor and Grigoris Paouris for helpful pointers and discussions, specifically for showing us the proofs of Lemmas 2.8 and 2.9.

References

[1] N. Z. Shor. Cut-off method with space extension in convex programming problems. *Kibernetika*, 13:9495, 1977.

[2] D. B. Yudin and A. S. Nemirovski. Evaluation of the information complexity of mathematical programming problems (in russian). *Ekonomika i Matematicheskie Metody*, 13(2):3–45, 1976.

[3] L. G. Khachiyan. Polynomial algorithms in linear programming. *USSR Computational Mathematics and Mathematical Physics*, 20:53–72, 1980.

[4] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*. Springer, 1988.

[5] R. Kannan, L. Lovász, and M. Simonovits. Random walks and an $O^*(n^5)$ volume algorithm for convex bodies. *Random Structures and Algorithms*, 11:1–50, 1997.

[6] L. Lovász and S. Vempala. Simulated annealing in convex bodies and an $O^*(n^4)$ volume algorithm. *J. Comput. Syst. Sci.*, 72(2):392–417, 2006.

[7] S. Vempala. Recent progress and open problems in algorithmic convex geometry. In *FSTTCS*, pages 42–64, 2010.

[8] G. Pisier. *The Volume of Convex Bodies and Banach Space Geometry*. Cambridge University Press, 1989.

[9] V. Milman. Inegalites de brunn-minkowski inverse et applications at la theorie locales des espaces normes. *C. R. Acad. Sci. Paris*, 302(1):25–28, 1986.

[10] B. Klartag. On convex perturbations with a bounded isotropic constant. *Geom. and Funct. Anal.*, 16(6):1274–1290, 2006.

[11] D. Dadush, C. Peikert, and S. Vempala. Enumerative lattice algorithms in any norm via m-ellipsoid coverings. In *FOCS*, pages 580–589, 2011.

[12] Z. Furedi and I. Barany. Computing the volume is difficult. In *STOC '86: Proceedings of the eighteenth annual ACM symposium on Theory of computing*, pages 442–447, New York, NY, USA, 1986. ACM.

[13] M. E. Dyer, A. M. Frieze, and R. Kannan. A random polynomial time algorithm for approximating the volume of convex bodies. In *STOC*, pages 375–381, 1989.

[14] D. Dadush and S. Vempala. Deterministic construction of an approximate m-ellipsoid and its application to derandomizing lattice algorithms. In *SODA*, pages 1445–1456, 2012.

[15] M.E. Dyer, A.M. Frieze, and R. Kannan. A random polynomial-time algorithm for approximating the volume of convex bodies. *J. ACM*, 38(1):1–17, 1991.

[16] G. Elekes. A geometric inequality and the complexity of computing volume. *Discrete & Computational Geometry*, pages 289–292, 1986.

[17] Z. Furedi and I. Barany. Approximation of the sphere by polytopes having few vertices. *Proceedings of the AMS*, 102(3):651–659, 1988.

[18] M. Ajtai, R. Kumar, and D. Sivakumar. A sieve algorithm for the shortest lattice vector problem. In *STOC*, pages 601–610, 2001.

[19] V. Arvind and P. S. Joglekar. Some sieving algorithms for lattice problems. In *FSTTCS*, pages 25–36, 2008.

[20] D. Micciancio and P. Voulgaris. A deterministic single exponential time algorithm for most lattice problems based on Voronoi cell computations. In *STOC*, pages 351–358, 2010.

[21] D. R. Lewis. Ellipsoids defined by Banach ideal norms. *Mathematika*, 26(1):18–29, 1979.

[22] T. Figiel and N. Tomczak-Jaegermann. Projections onto hilbertian subspaces of banach spaces. *Israel Journal of Mathematics*, 33:155–171, 1979.

[23] V. Milman. Isomorphic symmetrization and geometric inequalities. In Joram Lindenstrauss and Vitali Milman, editors, *Geometric Aspects of Functional Analysis*, volume 1317 of *Lecture Notes in Mathematics*, pages 107–131. Springer Berlin / Heidelberg, 1988.

[24] A. A. Giannopoulos and V. D. Milman. Chapter 17: Euclidean structure in finite dimensional normed spaces. In W.B. Johnson and J. Lindenstrauss, editors, *Handbook of the Geometry of Banach Spaces*, volume 1, pages 707–779. Elsevier Science B.V., 2001.

[25] G. Pisier. Remarques sur un resultat non publie de b. maurey. *(French) Seminaire d'Analyse Functionnelle 1980-81*, Exp. No. V, 13:1–12, 1981.

[26] C. A. Rogers and G. C. Shephard. The difference body of a convex body. *Arch. Math.*, 8:220–233, 1957.

[27] D. Avis and K. Fukuda. Reverse search for enumeration. *Discrete Applied Mathematics*, 65:21–46, 1993.