

On Khachiyan's algorithm for the computation of minimum-volume enclosing ellipsoids

Michael J. Todd^{a,1}, E. Alper Yıldırım^{b,2}

^a*School of Operations Research and Industrial Engineering, Cornell University, Ithaca, NY 14853, USA*

^b*Department of Industrial Engineering, Bilkent University, 06800 Bilkent, Ankara, Turkey*

Received 30 September 2005; received in revised form 13 February 2007; accepted 27 February 2007

Available online 3 April 2007

Dedicated to the memory of Leo Khachiyan

Abstract

Given $\mathcal{A} := \{a^1, \dots, a^m\} \subset \mathbb{R}^d$ whose affine hull is \mathbb{R}^d , we study the problems of computing an approximate rounding of the convex hull of \mathcal{A} and an approximation to the minimum-volume enclosing ellipsoid of \mathcal{A} . In the case of centrally symmetric sets, we first establish that Khachiyan's barycentric coordinate descent (BCD) method is exactly the polar of the deepest cut ellipsoid method using two-sided symmetric cuts. This observation gives further insight into the efficient implementation of the BCD method. We then propose a variant algorithm which computes an approximate rounding of the convex hull of \mathcal{A} , and which can also be used to compute an approximation to the minimum-volume enclosing ellipsoid of \mathcal{A} . Our algorithm is a modification of the algorithm of Kumar and Yıldırım, which combines Khachiyan's BCD method with a simple initialization scheme to achieve a slightly improved polynomial complexity result, and which returns a small "core set." We establish that our algorithm computes an approximate solution to the dual optimization formulation of the minimum-volume enclosing ellipsoid problem that satisfies a more complete set of approximate optimality conditions than either of the two previous algorithms. Furthermore, this added benefit is achieved without any increase in the improved asymptotic complexity bound of the algorithm of Kumar and Yıldırım or any increase in the bound on the size of the computed core set. In addition, the "dropping idea" used in our algorithm has the potential of computing smaller core sets in practice. We also discuss several possible variants of this dropping technique.

© 2007 Elsevier B.V. All rights reserved.

MSC: 90C25; 90C46; 65K05

Keywords: Löwner ellipsoid; Core sets; Rounding of polytopes; Ellipsoid method; Approximation algorithms

1. Introduction

Let $\mathcal{A} := \{a^1, \dots, a^m\} \subset \mathbb{R}^d$ be a finite set of vectors whose affine hull is \mathbb{R}^d . In this paper, we are concerned with the problem of computing an approximate "rounding" of the convex hull of \mathcal{A} as well as the problem of computing the minimum-volume enclosing ellipsoid (MVEE) of \mathcal{A} , which we shall denote by $\text{MVEE}(\mathcal{A})$, also known as the Löwner or Löwner–John ellipsoid of \mathcal{A} .

¹ This author was supported in part by NSF through grant DMS-0513337 and ONR through Grant N00014-02-1-0057.

² This author was supported in part by NSF through CAREER grant DMI-0237415.

E-mail addresses: miketodd@cs.cornell.edu (M.J. Todd), yildirim@bilkent.edu.tr (E.A. Yıldırım).

The MVEE satisfies [22]

$$\frac{1}{d} \text{MVEE}(\mathcal{A}) \subseteq \text{conv}(\mathcal{A}) \subseteq \text{MVEE}(\mathcal{A}), \quad (1)$$

where $\text{conv}(\mathcal{A})$ denotes the convex hull of \mathcal{A} and the ellipsoid on the left-hand side is obtained by scaling $\text{MVEE}(\mathcal{A})$ around its center by a factor of $1/d$. Furthermore, if \mathcal{A} is centrally symmetric (i.e., if $\mathcal{A} = -\mathcal{A}$), the factor on the left-hand side can be improved to $1/\sqrt{d}$. Therefore, $\text{MVEE}(\mathcal{A})$ provides a rounding of the full-dimensional polytope $\text{conv}(\mathcal{A})$.

Given $\varepsilon > 0$, an ellipsoid $\mathcal{E} \subset \mathbb{R}^d$ is said to be a $(1 + \varepsilon)d$ -rounding of $\text{conv}(\mathcal{A})$ if

$$\frac{1}{(1 + \varepsilon)d} \mathcal{E} \subseteq \text{conv}(\mathcal{A}) \subseteq \mathcal{E}. \quad (2)$$

In the case of centrally symmetric point sets, we will be interested in finding $\sqrt{(1 + \varepsilon)d}$ -roundings, where we replace the factor on the left-hand side of (2) by $1/\sqrt{(1 + \varepsilon)d}$. Similarly, for $\eta > 0$, we say that an ellipsoid $\mathcal{E} \subset \mathbb{R}^d$ is a $(1 + \eta)$ -approximation of $\text{MVEE}(\mathcal{A})$ if

$$\mathcal{A} \subseteq \mathcal{E}, \quad \text{Vol}(\mathcal{E}) \leq (1 + \eta) \text{Vol}(\text{MVEE}(\mathcal{A})), \quad (3)$$

where $\text{Vol}(\mathcal{E})$ denotes the volume of the ellipsoid \mathcal{E} .

In general, there is not a strong relationship between (2) and (3). Given an ellipsoid $\mathcal{E} \subset \mathbb{R}^d$ satisfying (2) for some $\varepsilon > 0$, we have

$$\frac{1}{(1 + \varepsilon)^d d^d} \text{Vol}(\mathcal{E}) \leq \text{Vol}(\text{conv}(\mathcal{A})) \leq \text{Vol}(\text{MVEE}(\mathcal{A})) \leq \text{Vol}(\mathcal{E}),$$

which implies that \mathcal{E} is a $(1 + \eta)$ -approximation of $\text{MVEE}(\mathcal{A})$ for $\eta = (1 + \varepsilon)^d d^d - 1$. (However, for the algorithms we discuss, there is a much tighter bound: when we terminate with an ellipsoid giving a $(1 + \varepsilon)d$ -rounding, we have a $(1 + \eta)$ -approximation of $\text{MVEE}(\mathcal{A})$ for $\eta = (1 + \varepsilon)^{(d+1)/2} - 1$.) On the other hand, if an ellipsoid $\mathcal{E} \subset \mathbb{R}^d$ satisfies (3) for some $\eta > 0$, it may happen that there is no finite value of ε such that (2) is satisfied. For instance, if \mathcal{A} consists of sampled points from the boundary of half of an ellipsoid $\tilde{\mathcal{E}}$ cut by a hyperplane through its center, then $\tilde{\mathcal{E}}$ yields a 2-approximation of $\text{MVEE}(\mathcal{A})$ whereas (2) will not be satisfied for any finite value of ε .

MVEEs play an important role in several diverse applications such as optimal design [32,34], computational geometry [6,12,39]), convex optimization [20,29], computer graphics [9,13], pattern recognition [18], and statistics [31].

Several algorithms have been developed for the MVEE problem. These algorithms can be categorized as first-order algorithms [23,26,31,34,35], second-order interior-point algorithms [30,33], or a combination of the two [23]. For small dimensions d , the MVEE problem can be solved in $O(d^{O(d)}m)$ arithmetic operations using randomized [1,27,39] or deterministic [12] algorithms. A fast implementation is also available in the CGAL library¹ for solving the problem in two dimensions [16]. Khachiyan and Todd [24] established a linear-time reduction of the MVEE problem to the problem of computing a maximum-volume inscribed ellipsoid (MVIE) in a polytope described by a finite number of inequalities. Therefore, the MVEE problem can also be solved using the algorithms developed for the (typically, harder) MVIE problem [4,24,28,42,43]. Since the MVEE problem can be formulated as a determinant maximization problem, the more general algorithms of Vandenberghe et al. [38] and Toh [37] can also be applied.

Khachiyan [23] proposed an algorithm to compute a $(1 + \varepsilon)d$ -rounding of $\text{conv}(\mathcal{A})$ in polynomial time for fixed $\varepsilon > 0$; it can also compute a $(1 + \eta)$ -approximation of $\text{MVEE}(\mathcal{A})$ in polynomial time for fixed $\eta > 0$. Khachiyan's algorithm can be viewed as a barycentric coordinate descent (BCD) method [23] or as a Frank–Wolfe [15] or sequential linear programming algorithm [26,33] for the dual optimization formulation of the MVEE problem. Indeed, this algorithm was proposed for the dual of the MVEE problem, which is the optimal design problem in statistics, by Fedorov [14]; a slight variant without line searches, which may have some advantages in the statistical setting, was independently suggested by Wynn [41] in 1970. So we could unambiguously call this the FW-method. However, Khachiyan's analysis

¹ <http://www.cgal.org>.

was dependent on his choice of initialization and termination rules, so we continue to call this version Khachiyan’s algorithm. His method computes a $(1 + \eta)$ -approximation to $\text{MVEE}(\mathcal{A})$ in

$$O(md^2([(1 + \eta)^{2/(d+1)} - 1]^{-1} + \log d + \log \log m)) \tag{4}$$

operations for $\eta > 0$. (Note that $[(1 + \eta)^{2/(d+1)} - 1]^{-1} = O(d/\eta)$ for $\eta \in (0, 1]$.)

More recently, Kumar and Yildirim [26] proposed a modification of Khachiyan’s algorithm (henceforth the KY algorithm) using a simple initialization scheme, which can compute a $(1 + \eta)$ -approximation to $\text{MVEE}(\mathcal{A})$ in

$$O(md^2([(1 + \eta)^{2/(d+1)} - 1]^{-1} + \log d)) \tag{5}$$

operations, which is a slight improvement over (4). In addition, the algorithm of [26] computes an ellipsoid $\mathcal{E} \subset \mathbb{R}^d$ such that $\mathcal{A} \subseteq \mathcal{E}$ and a subset $\mathcal{X} \subseteq \mathcal{A}$ with the property that

$$\text{Vol}(\text{MVEE}(\mathcal{X})) \leq \text{Vol}(\text{MVEE}(\mathcal{A})) \leq \text{Vol}(\mathcal{E}) \leq (1 + \eta) \text{Vol}(\text{MVEE}(\mathcal{X})),$$

which implies that \mathcal{E} is simultaneously a $(1 + \eta)$ -approximation to $\text{MVEE}(\mathcal{X})$ and to $\text{MVEE}(\mathcal{A})$. Such a set \mathcal{X} is called an “ η -core set” (or a core set) of \mathcal{A} to signify that $\text{conv}(\mathcal{X})$ provides a good approximation of $\text{conv}(\mathcal{A})$. Furthermore, \mathcal{X} satisfies

$$|\mathcal{X}| = O(d[(1 + \eta)^{2/(d+1)} - 1]^{-1} + d \log d), \tag{6}$$

which is independent of m , the number of points in \mathcal{A} . (Note that John [22] shows that a 0-core set of cardinality at most $(d + 1)(d + 2)/2$ exists. In Section 5, we establish that it is in fact possible to extract a core set whose size matches that of the 0-core set with some post-processing. However, the computation in such a procedure dominates all the work performed in the algorithm.) “Small” core set results have previously been established for several geometric optimization problems [25,8,7,11,2] and play an important role in developing efficient and practical algorithms for various large-scale problems in moderate dimensions.

Our contributions in this paper are twofold. In the case of centrally symmetric point sets, we establish that Khachiyan’s BCD method is exactly the polar of the deepest cut ellipsoid method (using two-sided symmetric cuts), but with a much improved analysis. Secondly, for arbitrary point sets, we propose a modification of the algorithm of [26], which computes an approximate solution that satisfies a more complete set of near-optimality conditions than either of the algorithms in [23] or in [26]. Furthermore, we establish that our modification does not lead to any increase in the asymptotic complexity of the algorithm of [26] given by (5). Finally, the KY algorithm of [26] starts with a carefully selected, small core set and expands it gradually. In contrast, our algorithm allows for dropping points in the working core set, which has the potential to compute a smaller core set than that given by (6) (but certainly no bigger asymptotically).

After the first version of this paper was written, we discovered that our method was discovered in 1973 by Atwood [5] in the context of the optimal design problem. Atwood’s work was independent of an earlier 1970 paper of Wolfe [40], who suggests the equivalent “away steps” to improve the convergence of the Frank–Wolfe method. Wolfe hinted at a proof of linear convergence, and this was later established rigorously by Guélat and Marcotte [21]. Their analysis does not apply to the current setting, and recently Ahipasaoglu et al. [3] proved the linear convergence of this method for the MVEE (or equivalently, the optimal design) problem.

The paper is organized as follows. We define notation in the remainder of this section. In Section 2, we review formulations of the MVEE problem as an optimization problem and we describe various approximate optimality conditions. Section 3 discusses an interpretation of Khachiyan’s algorithm as the polar of the deepest symmetric cut ellipsoid algorithm. Then in Section 4, we describe and analyze our modification of this and the KY algorithm. We conclude the paper in Section 5. There we also describe some preliminary computational results demonstrating the efficiency of the proposed method (more results can be found in [3]) and its ability to generate small core sets.

1.1. Notation

Vectors are denoted by lower-case Roman letters. For a vector p , p_i denotes its i th component, and P the diagonal matrix whose diagonal entries are given by these components. Inequalities on vectors apply to each component. We reserve e for the vector of ones of appropriate dimension, which will be clear from the context, and e^j for the j th unit vector. Upper-case Roman letters are reserved for matrices. For a finite set of vectors \mathcal{V} , $\text{span}(\mathcal{V})$ denotes the linear

subspace spanned by \mathcal{V} . Functions and operators are denoted by upper-case Greek letters. Scalars except for m, d , and n are represented by lower-case Greek letters unless they represent components of a vector or elements of a sequence of scalars, vectors or matrices. We reserve i, j , and k for indexing purposes. Upper-case script letters are used for all other objects such as sets, polytopes, and ellipsoids.

2. Preliminaries and formulations

A (full-dimensional) ellipsoid $\mathcal{E}_{Q,c}$ in \mathbb{R}^d is specified by a $d \times d$ symmetric positive definite matrix Q and a center $c \in \mathbb{R}^d$ and admits a representation given by

$$\mathcal{E}_{Q,c} = \{x \in \mathbb{R}^d : (x - c)^T Q (x - c) \leq 1\}. \tag{7}$$

The volume of the ellipsoid $\mathcal{E}_{Q,c}$, denoted by $\text{Vol}(\mathcal{E}_{Q,c})$, is given by $\text{Vol}(\mathcal{E}_{Q,c}) = \rho \det Q^{-1/2}$, where ρ is the volume of the unit ball in \mathbb{R}^d [20]. Similarly, we define the scaled volume by

$$\text{vol}(\mathcal{E}_{Q,c}) := \det Q^{-1/2}. \tag{8}$$

Let $\mathcal{A} := \{a^1, \dots, a^m\} \subset \mathbb{R}^d$ be a finite set of vectors whose affine hull is \mathbb{R}^d . If \mathcal{A} is not centrally symmetric, we define a “lifting” of \mathcal{A} to \mathbb{R}^n , where $n := d + 1$, by

$$\mathcal{A}' := \{\pm q^1, \dots, \pm q^m\} \quad \text{where} \quad q^i := \begin{bmatrix} a^i \\ 1 \end{bmatrix}, \quad i = 1, \dots, m, \tag{9}$$

which is centrally symmetric. It turns out that $\text{MVEE}(\mathcal{A})$ and $\text{MVEE}(\mathcal{A}')$ are closely related [24,30]. More specifically,

$$\text{MVEE}(\mathcal{A}) \times \{1\} = \text{MVEE}(\mathcal{A}') \cap \Pi, \tag{10}$$

where

$$\Pi := \{x \in \mathbb{R}^n : x_n = 1\}. \tag{11}$$

In addition, for any $\eta > 0$, if $\mathcal{E}' \subset \mathbb{R}^n$ is a $(1 + \eta)$ -approximation of $\text{MVEE}(\mathcal{A}')$, then \mathcal{E} , where $\mathcal{E} \times \{1\} := \mathcal{E}' \cap \Pi$, is a $(1 + \eta)$ -approximation of $\text{MVEE}(\mathcal{A})$ [24]. Henceforth, we assume that \mathcal{A} is not centrally symmetric; clearly, if it is, our algorithms can be simplified by omitting this lifting step.

Since \mathcal{A}' is centrally symmetric, it is easy to verify that $\text{MVEE}(\mathcal{A}')$ is centered at the origin. If $\mathcal{A}' \subset \mathcal{E}'_{Q,c}$ where $c \neq 0$, then $\mathcal{A}' \subset \mathcal{E}'_{(1/\theta)Q,0}$, where $\theta := 1 - c^T Q c < 1$, which implies that the latter ellipsoid has a smaller volume by (8). Therefore, the problem of computing $\text{MVEE}(\mathcal{A}')$ can be formulated as the following convex optimization problem:

$$\begin{aligned} (\mathbf{P}(\mathcal{A}')) \min_M \quad & -\log \det M \\ \text{s.t.} \quad & (q^i)^T M q^i \leq 1, \quad i = 1, \dots, m, \\ & M \in \mathbb{R}^{n \times n} \text{ is symmetric and positive definite,} \end{aligned} \tag{12a}$$

where $M \in \mathbb{R}^{n \times n}$ is the decision variable.

The Lagrangian dual of $(\mathbf{P}(\mathcal{A}'))$ is equivalent to

$$\begin{aligned} (\mathbf{D}(\mathcal{A}')) \max_p \quad & \Phi(p) := \log \det A(p) \\ \text{s.t.} \quad & e^T p = 1, \\ & p \geq 0, \end{aligned} \tag{12b}$$

where $p \in \mathbb{R}^m$ is the decision variable and $A : \mathbb{R}^m \rightarrow \mathbb{R}^{n \times n}$ is the linear operator given by

$$A(p) := \sum_{i=1}^m p_i q^i (q^i)^T. \tag{13}$$

Using the fact that the gradient of the function $\Theta(X) := \log \det X$ over symmetric matrices is given by $\nabla \Theta(X) = X^{-1}$, we see that the necessary and sufficient optimality conditions for p^* to solve $(\mathbf{D}(\mathcal{A}'))$ are given by

$$w_i(p^*) + s_i^* = \lambda^*, \quad i = 1, \dots, m, \tag{14a}$$

$$e^T p^* = 1, \tag{14b}$$

$$p_i^* s_i^* = 0, \quad i = 1, \dots, m, \tag{14c}$$

together with $p^* \geq 0$ and $s^* \geq 0$, where

$$w_i(p) := (q^i)^T \Lambda(p)^{-1} q^i, \quad i = 1, \dots, m. \tag{15}$$

For any feasible solution $p \in \mathbb{R}^m$ of $(\mathbf{D}(\mathcal{A}'))$ with $\Phi(p) > -\infty$, we have [23, Lemma 1]

$$\sum_{i=1}^m p_i w_i(p) = n. \tag{16}$$

Therefore, multiplying both sides of (14a) by p_i^* and summing up for $i = 1, \dots, m$, we obtain $\lambda^* = n$ by (14b) and (14c). It follows then that the optimality conditions of $(\mathbf{D}(\mathcal{A}'))$ can be equivalently rewritten as

$$w_i(p^*) \leq n, \quad i = 1, \dots, m, \tag{17a}$$

$$e^T p^* = 1, \tag{17b}$$

together with $p^* \geq 0$. By (17) and (16),

$$p_i^* > 0 \text{ implies } w_i(p^*) = n, \quad i = 1, \dots, m, \tag{18}$$

which are simply the complementary slackness conditions (14c).

Khachiyan’s algorithm [23] is driven by computing a feasible solution \tilde{p} of $(\mathbf{D}(\mathcal{A}'))$ that satisfies the so-called ε -relaxed optimality conditions defined by

$$w_i(\tilde{p}) \leq (1 + \varepsilon)n, \quad i = 1, \dots, m. \tag{19}$$

For a solution \tilde{p} to (19), let $j \in \{1, \dots, m\}$ be such that $\tilde{p}_j > 0$. By (16),

$$\begin{aligned} w_j(\tilde{p}) &= \frac{1}{\tilde{p}_j} \left(n - \sum_{i=1, i \neq j}^m \tilde{p}_i w_i(\tilde{p}) \right), \\ &\geq (n[1 - (1 + \varepsilon)(1 - \tilde{p}_j)]) / \tilde{p}_j, \\ &= n(1 + \varepsilon - \varepsilon / \tilde{p}_j), \end{aligned}$$

where we used (19) and the feasibility of \tilde{p} to derive the inequality. Therefore, such a solution \tilde{p} satisfies a very weak approximate form of the complementary slackness conditions (18).

In view of this observation, we define a more complete set of approximate optimality conditions aimed towards a stronger approximation to the complementary slackness conditions (18). Given $\varepsilon \in (0, 1)$, we say that a feasible solution \hat{p} satisfies the ε -approximate optimality conditions if

$$w_i(\hat{p}) \leq (1 + \varepsilon)n, \quad i = 1, \dots, m, \tag{20a}$$

$$\hat{p}_i > 0 \text{ implies } w_i(\hat{p}) \geq (1 - \varepsilon)n, \quad i = 1, \dots, m. \tag{20b}$$

Note that these conditions imply that $(1 - \varepsilon)n \leq w_i(\hat{p}) \leq (1 + \varepsilon)n$ if $\hat{p}_i > 0$, $i = 1, \dots, m$, which is a better approximation of the complementary slackness conditions (18) than the ε -relaxed optimality conditions (19) provide.

Khachiyan’s algorithm [23] starts with a feasible solution $\tilde{p} > 0$ of $(\mathbf{D}(\mathcal{A}'))$ and improves upon the objective function value by increasing only one component of \tilde{p} at each iteration and then rescaling to regain feasibility. On the other hand, the KY algorithm [26] uses a simple initialization scheme to compute a feasible solution \hat{p} of $(\mathbf{D}(\mathcal{A}'))$ with

only at most $\min\{2d, m\}$ positive components and then uses the same improvement idea. Therefore, neither of these algorithms ever reduces the number of positive components of p at any iteration. In contrast, while our algorithm starts off with the same initial feasible solution \hat{p} of $(\mathbf{D}(\mathcal{A}'))$, we will allow the number of positive components of p to be reduced. Therefore, our algorithm can potentially compute an approximate solution with a smaller number of positive components than an approximate solution computed by either of the previous two algorithms, which, in turn, will result in a smaller core set. In addition, the asymptotic complexity of our algorithm remains the same as the improved complexity result (5) of [26].

We close this section by relating $(\mathbf{D}(\mathcal{A}'))$ directly to the minimum-volume ellipsoid problem. First, its optimal solution p^* can be easily related to MVEE(\mathcal{A}) (see, e.g., [26, Lemma 2.1]). Let $A \in \mathbb{R}^{d \times m}$ be the matrix whose i th column is given by $a^i, i = 1, \dots, m$. Then $\text{MVEE}(\mathcal{A}) = \mathcal{E}_{Q^*, c^*} := \{x \in \mathbb{R}^d : (x - c^*)^T Q^* (x - c^*) \leq 1\}$, where

$$Q^* := \frac{1}{d}(AP^*A^T - Ap^*(Ap^*)^T)^{-1}, \quad c^* := Ap^*. \tag{21}$$

Furthermore,

$$\log \text{vol}(\text{MVEE}(\mathcal{A})) = \frac{d}{2} \log d + \frac{1}{2} \log \det A(p^*). \tag{22}$$

Now suppose that, instead of an optimal solution p^* , we have a feasible solution \tilde{p} of $(\mathbf{D}(\mathcal{A}'))$ satisfying (19). Then, by combining the arguments of Khachiyan [23, Lemmas 2 and 5] and of [26, Lemma 2.1], we find that $\mathcal{E}_{\tilde{Q}, \tilde{c}}$ is a $(1 + ((d + 1)/d)\varepsilon)d$ -rounding of \mathcal{A} and a $(1 + \varepsilon)^{(d+1)/2}$ -approximation to $\text{MVEE}(\mathcal{A})$, where

$$\tilde{Q} := \frac{1}{(1 + \varepsilon)d}(A\tilde{P}A^T - A\tilde{p}(A\tilde{p})^T)^{-1}, \quad \tilde{c} := A\tilde{p}. \tag{23}$$

3. A new interpretation of Khachiyan’s algorithm

In this section, we establish that Khachiyan’s algorithm [23] is exactly the polar of the deepest cut ellipsoid method (using two-sided symmetric cuts), but with a much improved analysis.

We describe Khachiyan’s algorithm below.

Algorithm 3.1. Khachiyan’s algorithm to compute a feasible solution of $(\mathbf{D}(\mathcal{A}'))$ satisfying (19).

Require: Input set of points $\mathcal{A} = \{a^1, \dots, a^m\} \subset \mathbb{R}^d, \varepsilon > 0$.

- 1: $k \leftarrow 0, n \leftarrow d + 1, p^0 \leftarrow (1/m)e$, and $q^i \leftarrow ((a^i)^T, 1)^T, i = 1, \dots, m$.
 - 2: While p^k does not satisfy (19), do
 - 3: **loop**
 - 4: $j \leftarrow \arg \max_{i=1, \dots, m} (q^i)^T A(p^k)^{-1} q^i, \kappa \leftarrow (q^j)^T A(p^k)^{-1} q^j$;
 - 5: $\beta \leftarrow \frac{\kappa - n}{n(\kappa - 1)}$;
 - 6: $p^{k+1} \leftarrow (1 - \beta)p^k + \beta e^j, k \leftarrow k + 1$.
 - 7: **end loop**
 - 8: **Output** p^k .
-

Khachiyan’s algorithm seeks a minimum-volume ellipsoid containing

$$\mathcal{Q} := \text{conv}\{\pm q^1, \dots, \pm q^m\},$$

but, since it is a dual algorithm, it constructs a sequence of ellipsoids

$$\mathcal{E}_k := \{y \in \mathbb{R}^n : y^T A(p^k)^{-1} y \leq 1\}$$

satisfying $\mathcal{E}_k \subseteq \mathcal{Q}$, and stops when $\mathcal{Q} \subseteq \sqrt{(1 + \varepsilon)n} \mathcal{E}_k$. Thus, the polar ellipsoids

$$\mathcal{E}_k^\circ = \{z \in \mathbb{R}^n : z^T A(p^k) z \leq 1\} \tag{24}$$

all contain the polar polytope

$$\mathcal{Q}^\circ = \{z \in \mathbb{R}^n : -1 \leq (q^i)^\top z \leq 1, i = 1, \dots, m\},$$

and the algorithm stops when $(1/\sqrt{(1+\varepsilon)n})\mathcal{E}_k^\circ$ is contained in \mathcal{Q}° .

So suppose that, at a particular iteration, we have the ellipsoid \mathcal{E}_k° , and that \mathcal{Q}° does not contain $(1/\sqrt{(1+\varepsilon)n})\mathcal{E}_k^\circ$. This means that one of the pairs of hyperplanes $(q^i)^\top z = \pm 1$ of \mathcal{Q}° intersects \mathcal{E}_k° “close to the origin.” So the condition $(q^j)^\top A(p^k)^{-1}q^j =: \kappa > (1+\varepsilon)n$ in Khachiyan’s algorithm corresponds to

$$\mathcal{Q}^\circ \subseteq \{z \in \mathcal{E}_k^\circ : -\gamma((q^j)^\top A(p^k)^{-1}q^j)^{1/2} \leq (q^j)^\top z \leq \gamma((q^j)^\top A(p^k)^{-1}q^j)^{1/2}\}$$

for $\gamma := 1/\sqrt{\kappa} < 1/\sqrt{(1+\varepsilon)n}$. (We use γ here for β in [36], to avoid confusion with Khachiyan’s β .) This is exactly the set-up of the deepest two-sided symmetric cut ellipsoid method, which chooses as the next ellipsoid the smallest volume ellipsoid containing the right-hand side above. Further, in the version of Burrell and Todd [10], the initial ellipsoid is also of the form in (24) for a suitable p .

Let us examine the next ellipsoid constructed in this ellipsoid method. According to Theorem 2 of [36], this has the form $\{z \in \mathbb{R}^n : z^\top B_+^{-1}z \leq 1\}$, with

$$B_+ := \delta \left(B - \sigma \frac{(Bq)(Bq)^\top}{q^\top Bq} \right), \tag{25}$$

where $B := A(p^k)^{-1}$, $q := q^j$, and

$$\delta := \frac{n(1-\gamma^2)}{n-1}, \quad \sigma := \frac{1-n\gamma^2}{1-\gamma^2}.$$

Then, using the rank-one modification formula, we obtain

$$\begin{aligned} B_+^{-1} &= \delta^{-1} \left(B^{-1} + \frac{\sigma}{(1-\sigma)q^\top Bq} qq^\top \right) \\ &= \delta^{-1} \left(B^{-1} + \frac{\sigma\gamma^2}{(1-\sigma)} qq^\top \right) \\ &= \delta^{-1}(1+\mu) \left[\left(1 - \frac{\mu}{1+\mu} \right) B^{-1} + \frac{\mu}{1+\mu} qq^\top \right], \end{aligned}$$

where

$$\mu := \frac{\sigma\gamma^2}{(1-\sigma)}.$$

Now substituting in the value for σ , we find $\mu = (1-n\gamma^2)/(n-1)$, so that $1+\mu = n(1-\gamma^2)/(n-1)$, and

$$\frac{\mu}{1+\mu} = \frac{1-n\gamma^2}{n(1-\gamma^2)} = \frac{\gamma^{-2}-n}{n(\gamma^{-2}-1)} = \frac{\kappa-n}{n(\kappa-1)} = \beta,$$

using $\gamma := 1/\sqrt{\kappa}$. Next, substituting in the value for δ , we get $\delta^{-1}(1+\mu) = 1$, so that

$$B_+^{-1} = (1-\beta)B^{-1} + \beta qq^\top = (1-\beta)A(p^k) + \beta qq^\top = A(p^{k+1}).$$

This demonstrates the equivalence of the two methods.

Above, we claimed that Khachiyan provides a much improved analysis of his method. Indeed, he showed [23, Lemma 3] that, at every iteration, the (natural) logarithm of the volume of \mathcal{E}_k increases by $(1/2)(\log(1+\varepsilon_k) - \varepsilon_k/(1+\varepsilon_k))$, where $\max_{i=1,\dots,m} (q^i)^\top A(p^k)^{-1}q^i = (1+\varepsilon_k)n$. He then uses this in his Lemma 4 to provide a bound on the total number of iterations, by bounding the number required to reduce ε_k to 1, then to $\frac{1}{2}$, etc. By contrast, analyses of the

ellipsoid method just look at the worst case volume reduction: in this case, one would bound the number of iterations assuming that the logarithm of the volume only increased by $(1/2)(\log(1 + \varepsilon) - \varepsilon/(1 + \varepsilon))$ at each iteration, where ε is the final tolerance.

Remark. Let us briefly discuss implementation of Khachiyan’s algorithm (and our variants). At each iteration, we need access to the inverse of $A(p^k)$, or to some factorization of it, and to the quantities $(q^i)^T A(p^k)^{-1} q^i$ for each i . At each iteration, $A(p^k)$ is updated by adding a rank-one symmetric matrix to it, and then scaling by a positive number. In some applications, the vectors q^i and hence possibly the matrix $A(p^k)$ will be sparse. We therefore recommend maintaining a Cholesky factorization LDL^T of $A(p^k)$, where L is a lower triangular matrix with unit diagonal and D is diagonal with positive diagonal entries. We also maintain values $\kappa_i = w_i(p^k) = (q^i)^T A(p^k)^{-1} q^i$ for each i . From these values we can determine κ and j . We then compute $\hat{q}^j := L^{-1} q^j$ and from this a more accurate value $(\hat{q}^j)^T D^{-1} \hat{q}^j$ for κ and $\bar{q}^j := (LDL^T)^{-1} q^j = L^{-T} D^{-1} \hat{q}^j$. Noting that

$$A(p^{k+1}) = (1 - \beta) \left[A(p^k) + \frac{\beta}{1 - \beta} q^j (q^j)^T \right],$$

we see that it is enough to update the Cholesky factorization after a rank-one change, and then scale the diagonal matrix to account for the multiplicative factor. The rank-one update can be performed in an efficient and numerically stable way using the technique of Gill et al. in [17]. This uses the already computed vector \hat{q}^j , and requires $O(n^2)$ operations.

We also need to update the quantities κ_i . Using the update of the matrix $A(p^k)^{-1}$ in (25), we can easily check that the formulae

$$\kappa_i^\pm = \delta(\kappa_i - \sigma((q^i)^T \bar{q}^j)^2 / \kappa), \quad i = 1, \dots, m$$

(where δ and σ are computed as above from $\gamma := 1/\sqrt{\kappa}$) yield the new quantities. (For $i = j$, we use the formula $\kappa_j^\pm = \delta(1 - \sigma)\kappa$.) Each update of a κ_i requires $O(n)$ operations to perform the inner product, for a total of $O(mn)$. The total complexity of $O(mn)$ operations for each iteration is the same as in Khachiyan [23], but with a more stable implementation. Note that Sun and Freund [33] state that $O(mn^2)$ operations are needed.

In our modified algorithm, we will occasionally subtract a rank-one matrix from $A(p^k)$ instead of adding one. Here care must be taken to preserve numerical stability; Gill et al. [17] propose a stable technique for this also.

Our implementation recommendation is very similar to that proposed by Goldfarb and Todd [19] for the ellipsoid method. The difference is that here we are proposing maintaining a Cholesky factorization of the possibly sparse matrix $A(p^k)$, and in most iterations we add a rank-one matrix, while [19] maintained a factorization of $A(p^k)^{-1}$, and subtracted a rank-one matrix at each iteration.

4. A modification of the KY algorithm

Let $\mathcal{A} := \{a^1, \dots, a^m\} \subset \mathbb{R}^d$ be a finite set of vectors whose affine hull is \mathbb{R}^d . In this section, we describe a modification of the KY algorithm [26], which, in turn, is derived from Khachiyan’s algorithm [23], to compute a feasible solution \hat{p} of $(\mathbf{D}(\mathcal{A}))$ that satisfies the ε -approximate optimality conditions (20) for any given $\varepsilon > 0$.

The main difference between the KY algorithm and Khachiyan’s algorithm is the choice of the initial feasible solution. The former uses a simple initial volume approximation scheme in an attempt to identify a smaller subset of vectors in \mathcal{A} that provides a reasonable approximation of $\text{conv}(\mathcal{A})$. We outline this scheme in the next subsection.

4.1. Initial volume approximation

Given $\mathcal{A} = \{a^1, \dots, a^m\} \subset \mathbb{R}^d$, the following deterministic algorithm identifies a subset $\mathcal{X}_0 \subseteq \mathcal{A}$ of size at most $\min\{2d, m\}$ such that $\text{vol MVEE}(\mathcal{X}_0)$ is a provable approximation to $\text{vol MVEE}(\mathcal{A})$ [26].

Algorithm 4.1. Volume approximation algorithm.

Require: Input set of points $\mathcal{A} = \{a^1, \dots, a^m\} \subset \mathbb{R}^d$.

- 1: If $m \leq 2d$, then $\mathcal{X}_0 \leftarrow \mathcal{A}$. Return.
- 2: $\Psi \leftarrow \{0\}$, $\mathcal{X}_0 \leftarrow \emptyset$, $k \leftarrow 0$.
- 3: While $\mathbb{R}^d \setminus \Psi \neq \emptyset$, **do**
- 4: **loop**
- 5: $k \leftarrow k + 1$; pick an arbitrary direction $b^k \in \mathbb{R}^d$ in the orthogonal complement of Ψ ;
- 6: $\alpha \leftarrow \arg \max_{i=1, \dots, m} (b^k)^T a^i$, $\mathcal{X}_0 \leftarrow \mathcal{X}_0 \cup \{a^\alpha\}$;
- 7: $\beta \leftarrow \arg \min_{i=1, \dots, m} (b^k)^T a^i$, $\mathcal{X}_0 \leftarrow \mathcal{X}_0 \cup \{a^\beta\}$;
- 8: $\Psi \leftarrow \text{span}(\Psi, \{a^\beta - a^\alpha\})$.
- 9: **end loop**
- 10: **Output** \mathcal{X}_0 .

The following lemma provides information about the running time of Algorithm 4.1 and the quality of the resulting approximation.

Lemma 4.1 (Kumar and Yildirim [26]). *Algorithm 4.1 terminates in $O(md^2)$ time with a subset $\mathcal{X}_0 \subseteq \mathcal{A}$ with $|\mathcal{X}_0|$ at most $\min\{2d, m\}$ such that*

$$\text{Vol MVEE}(\mathcal{A}) \leq d^{2d} \text{Vol MVEE}(\mathcal{X}_0). \tag{26}$$

4.2. Our modification

In this section, we present a modification of the KY algorithm for approximating the MVEE of a given set $\mathcal{A} = \{a^1, \dots, a^m\} \subset \mathbb{R}^d$. Given $\varepsilon > 0$, we establish that our modification computes an approximate solution that satisfies the ε -approximate optimality conditions given by (20). Note that each of the algorithms in [23] and in [26] computes an approximate solution that satisfies the weaker ε -relaxed optimality conditions given by (19). In addition, this added benefit does not lead to an increase in the asymptotic complexity result, i.e., the running time of our modified algorithm is asymptotically the same as that of the improved complexity result of the KY algorithm. Finally, we show that our algorithm returns a core set whose asymptotic size has the same bound as that computed by the KY algorithm. In contrast with the KY algorithm which only adds points to the working core set, our modification allows dropping points throughout the algorithm. Therefore, in practice, our modification is likely to compute a smaller core set than that computed by the KY algorithm.

We outline our algorithm below:

Algorithm 4.2. Modified algorithm to compute a feasible solution of $(\mathbf{D}(\mathcal{A}'))$ satisfying (20).

Require: Input set of points $\mathcal{A} = \{a^1, \dots, a^m\} \subset \mathbb{R}^d$, $\varepsilon > 0$.

- 1: Run Algorithm 4.1 on \mathcal{A} to get output \mathcal{X}_0 .
- 2: Let $p^0 \in \mathbb{R}^m$ be such that $p_i^0 = 1/|\mathcal{X}_0|$ for $a^i \in \mathcal{X}_0$ and $p_i^0 = 0$ otherwise.
- 3: $k \leftarrow 0$, $n \leftarrow d + 1$, and $q^i \leftarrow ((a^i)^T, 1)^T$, $i = 1, \dots, m$.
- 4: $\mathcal{E}_0 \leftarrow \{y \in \mathbb{R}^n : y^T A(p^0)^{-1} y \leq 1\}$.
- 5: While p^k does not satisfy (20), **do**
- 6: **loop**
- 7: $j_+ \leftarrow \arg \max\{(q^i)^T A(p^k)^{-1} q^i : i = 1, \dots, m\}$, $\kappa_+ \leftarrow (q^{j_+})^T A(p^k)^{-1} q^{j_+}$;
- 8: $j_- \leftarrow \arg \min\{(q^i)^T A(p^k)^{-1} q^i : i = 1, \dots, m, p_i^k > 0\}$, $\kappa_- \leftarrow (q^{j_-})^T A(p^k)^{-1} q^{j_-}$;
- 9: $\varepsilon_+ \leftarrow (\kappa_+/n) - 1$, $\varepsilon_- \leftarrow 1 - (\kappa_-/n)$;
- 10: $\varepsilon_k \leftarrow \max\{\varepsilon_+, \varepsilon_-\}$;
- 11: **if** $\varepsilon_k = \varepsilon_+$ **then**

```

12:                                      $\beta_k \leftarrow \frac{\kappa_+ - n}{n(\kappa_+ - 1)}$ ;
13:                                      $p^{k+1} \leftarrow (1 - \beta_k)p^k + \beta_k e^{j_+}, k \leftarrow k + 1$ ;
14:     else
15:                                      $\beta_k \leftarrow \min \left\{ \frac{n - \kappa_-}{n(\kappa_- - 1)}, \frac{p_{j_-}^k}{1 - p_{j_-}^k} \right\}$ ;
16:                                      $p^{k+1} \leftarrow (1 + \beta_k)p^k - \beta_k e^{j_-}, k \leftarrow k + 1$ ;
17:     end if
18:      $\mathcal{E}_k \leftarrow \{y \in \mathbb{R}^n : y^T A(p^k)^{-1} y \leq 1\}$ ;
19:      $\mathcal{X}_k \leftarrow \{a^i \in \mathcal{A} : p_i^k > 0\}$ .
20: end loop
21: Output  $p^k, \mathcal{E}_k$  and  $\mathcal{X}_k$ .
    
```

We now describe Algorithm 4.2 in more detail. Given $\mathcal{A} = \{a^1, \dots, a^m\} \subset \mathbb{R}^d$, Algorithm 4.1 is called on \mathcal{A} to get output $\mathcal{X}_0 \subseteq \mathcal{A}$. This subset is used to identify an initial feasible solution p^0 of $(\mathbf{D}(\mathcal{A}'))$. The KY algorithm also uses the same procedure to obtain an initial feasible point. The main difference lies in the execution of each iteration of the main loop. In contrast with the KY algorithm, Algorithm 4.2 computes not only the farthest point $q^{j_+} \in \mathcal{A}'$ from the origin in terms of the ellipsoidal norm induced by \mathcal{E}_k but also the closest point $q^{j_-} \in \mathcal{A}'$ among those with p_j^k positive. At iteration k ,

$$(q^i)^T A(p^k)^{-1} q^i \leq (1 + \varepsilon_+)n, \quad i = 1, \dots, m, \quad \text{and} \quad (q^i)^T A(p^k)^{-1} q^i \geq (1 - \varepsilon_-)n \quad \text{if } p_i^k > 0 \tag{27}$$

by definition of ε_- and ε_+ , which implies that p^k satisfies the ε_k -approximate optimality conditions given by (20). Both Khachiyan’s algorithm and the KY algorithm work towards improving ε_+ . In contrast, Algorithm 4.2 is driven by improving both ε_+ and ε_- .

Let us now describe how p^k gets updated at iteration k . Since $\mathcal{E}_k \subseteq \mathcal{Q} := \text{conv}\{\pm q^1, \dots, \pm q^m\}$, p^k is updated in a way to yield the maximum increase in the volume of \mathcal{E}_{k+1} . In the case that $\varepsilon_k = \varepsilon_+$, p^{k+1} is given by a convex combination of p^k and e^{j_+} . Since $\log \text{vol}(\mathcal{E}_{k+1})$ is exactly half $\log \det A(p^{k+1})$, β_k is given by the solution of

$$\beta_k := \arg \max_{\beta \in [0,1]} \log \det A((1 - \beta)p^k + \beta e^{j_+}) = \frac{\kappa^+ - n}{n(\kappa^+ - 1)}.$$

Both Khachiyan’s algorithm and the KY algorithm use this update. On the other hand, if $\varepsilon = \varepsilon_-$, then p^{k+1} is obtained from p^k by “moving away” from e^{j_-} . (Sun and Freund [33] and Kumar and Yildirim [26] show that e^{j_+} maximizes a linear approximation to the objective function of $(\mathbf{D}(\mathcal{A}'))$ at the current point over the simplex: similarly, it can be seen that e^{j_-} minimizes the same linear approximation when restricted only to the positive components of p^k .) In this case, β_k is given by the solution of

$$\beta_k := \arg \max_{\beta \in \left[0, \frac{p_{j_-}^k}{1 - p_{j_-}^k}\right]} \log \det A((1 + \beta)p^k - \beta e^{j_-}) = \min \left\{ \frac{n - \kappa_-}{n(\kappa_- - 1)}, \frac{p_{j_-}^k}{1 - p_{j_-}^k} \right\}.$$

Note that the range of β is chosen to ensure the feasibility of p^{k+1} .

4.3. Analysis of the modified algorithm

Our analysis is based very heavily on those for Khachiyan’s and the KY algorithm, but we need to distinguish the three kinds of iterations. If $\varepsilon_k = \varepsilon_+$, we call the k th iteration a *plus-iteration*. If $\varepsilon_k = \varepsilon_-$ and $\beta_k = (n - \kappa_-)/(n(\kappa_- - 1))$, we call it a *minus-iteration*. Finally, if $\varepsilon_k = \varepsilon_-$ and $\beta_k = p_{j_-}^k / (1 - p_{j_-}^k) < (n - \kappa_-)/(n(\kappa_- - 1))$ we call it a *drop-iteration*, because then a component of p^k becomes zero and the associated point a^{j_-} is dropped from \mathcal{X}_k .

We consider the quantities $v_k := \log \det A(p^k)$ and how they change at each iteration. Note that

$$A(p^{k+1}) = (1 - \alpha)A(p^k) + \alpha q^j (q^j)^T$$

for $\alpha = \beta_k$ and $j = j_+$ in a plus-iteration, or $\alpha = -\beta_k$ and $j = j_-$ in a minus-iteration or a drop-iteration. It follows that

$$v_{k+1} - v_k = (n - 1) \log(1 - \alpha) + \log(1 + \alpha(\kappa - 1)),$$

where κ is either κ_+ or κ_- respectively. In a drop-iteration, all we can say is that this is nonnegative. (The function above is monotonic until its maximum.) But in a plus- or minus-iteration, we can substitute for the value of α in terms of κ and then κ in terms of ω (either ε_+ or $-\varepsilon_-$, respectively) to get

$$\begin{aligned} v_{k+1} - v_k &= (n - 1) \log\left(\frac{(n - 1)\kappa}{n(\kappa - 1)}\right) + \log\left(\frac{\kappa}{n}\right) \\ &= (n - 1) \log\left(\frac{(n - 1)(1 + \omega)}{n - 1 + \omega n}\right) + \log(1 + \omega) \\ &= \log(1 + \omega) - (n - 1) \log\left(1 + \frac{\omega}{(n - 1)(1 + \omega)}\right) \\ &\geq \log(1 + \omega) - \frac{\omega}{1 + \omega}. \end{aligned} \tag{28}$$

Lemma 4.2. *In a plus- or a minus-iteration,*

$$v_{k+1} - v_k \geq \log(1 + \varepsilon_k) - \frac{\varepsilon_k}{1 + \varepsilon_k}.$$

Proof. This follows directly from (28) in a plus-iteration. To complete the proof, it suffices using (28) again to show that

$$\log(1 - \varepsilon) + \frac{\varepsilon}{1 - \varepsilon} \geq \log(1 + \varepsilon) - \frac{\varepsilon}{1 + \varepsilon}$$

for $\varepsilon = \varepsilon_- = \varepsilon_k > 0$. But if we define the functions $f(\varepsilon) := \log(1 + \varepsilon) - \varepsilon/(1 + \varepsilon)$ and $g(\varepsilon) := f(-\varepsilon)$, we find $f'(\varepsilon) = \varepsilon/(1 + \varepsilon)^2$ and $g'(\varepsilon) = -f'(-\varepsilon) = \varepsilon/(1 - \varepsilon)^2 \geq \varepsilon/(1 + \varepsilon)^2$ for any nonnegative ε . Also, f and g are both zero at zero, and this gives $g(\varepsilon_k) \geq f(\varepsilon_k)$ and hence the result. \square

Let us define

$$\tau_0 := \min\{k : \varepsilon_k \leq 1\}. \tag{29}$$

Khachiyan’s analysis [23] starts by deriving an upper bound on τ_0 . To that end, let v^* denote the optimal value of $(\mathbf{D}(\mathcal{A}'))$. Using Lemma 4.1 and (22), Kumar and Yildirim [26, Theorem 4.2] establish that $v^* - v_0 = O(d \log d)$. (Khachiyan’s algorithm [23] uses a different initial point which satisfies $v^* - v_0 = O(d \log m)$, and applies a slightly different argument than that below. This is the reason why the KY algorithm achieves a slightly improved complexity result over Khachiyan’s algorithm.) By Lemma 4.2, at each plus- or minus-iteration with $\varepsilon_k \geq 1$, we have $v_{k+1} - v_k \geq \log 2 - 1/2 > 0$. At each drop iteration, we can no longer find a positive lower bound on $v_{k+1} - v_k \geq 0$. However, each such iteration can be paired with a previous iteration where $p_{j_-}^k$ was increased from zero, except for those where $p_{j_-}^k$ is decreased to zero for the first time and was positive at the initial iteration. Note that the initial feasible point p^0 in Algorithm 4.2 has only at most $2d$ positive components (in contrast with m positive components in Khachiyan’s algorithm). Therefore, doubling the iteration count in the analysis of [26], we find that, after at most $2d + O(d \log d) = O(d \log d)$ iterations, Algorithm 4.2 computes a solution p^k with $\varepsilon_k \leq 1$, which implies that

$$\tau_0 = O(d \log d). \tag{30}$$

The next stage of Khachiyan’s analysis is aimed at deriving an upper bound on the number of iterations to obtain an iterate p^k with $\varepsilon_k \leq \varepsilon$. Starting with an iterate with $\varepsilon_k \leq 1$, Khachiyan’s clever argument [23, Lemma 4] is based on computing an upper bound on the number of iterations to obtain the first iterate with $\varepsilon_k \leq \frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \dots$. It follows from this argument that Khachiyan’s algorithm (or the KY algorithm) computes an iterate with $\varepsilon_k \leq \varepsilon$ after at most $O(d/\varepsilon)$ iterations. Due to the possibility of drop-iterations, we can again invoke a similar argument based on iteration doubling

and a “fixed charge” (the initial at most $2d$ positive components of p^k) to establish that Algorithm 4.2 terminates after

$$[\tau := \tau_0 + (\min\{k : \varepsilon_k \leq \varepsilon\} - \tau_0) = O(d \log d) + O(d/\varepsilon)] \tag{31}$$

iterations.

Theorem 4.1. *Given $\mathcal{A} = \{a^1, \dots, a^m\} \subset \mathbb{R}^d$ and $\varepsilon > 0$, Algorithm 4.2 computes a feasible solution p of $(\mathbf{D}(\mathcal{A}'))$ that satisfies the ε -approximate optimality conditions given by (20) in $O(d[\log d + 1/\varepsilon])$ iterations.*

Upon termination of Algorithm 4.2, we have

$$\mathcal{E}_k \subseteq \mathcal{Q} \subseteq \sqrt{(1 + \varepsilon)n} \mathcal{E}_k,$$

where $\mathcal{Q} := \text{conv}\{\pm q^1, \dots, \pm q^m\}$. It follows that $\sqrt{(1 + \varepsilon)n} \mathcal{E}_k \subset \mathbb{R}^n$ is a $\sqrt{(1 + \varepsilon)n}$ -rounding of \mathcal{Q} . In order to obtain a $(1 + \varepsilon)d$ -rounding of $\text{conv}(\mathcal{A})$, it follows from the analysis in [23, Section 3] that Algorithm 4.2 can be called with

$$\varepsilon' := \frac{d}{d + 1} \varepsilon = \frac{d}{n} \varepsilon,$$

to obtain \mathcal{E}_k . Then let $\mathcal{E} \subset \mathbb{R}^d$ be defined by

$$\mathcal{E} \times \{1\} := \sqrt{(1 + \varepsilon')n} \mathcal{E}_k \cap \Pi = \sqrt{1 + (1 + \varepsilon)d} \mathcal{E}_k \cap \Pi,$$

where Π is given by (11). By [23, Lemma 5],

$$\frac{1}{(1 + \varepsilon)d} \mathcal{E} \subseteq \text{conv}(\mathcal{A}) \subseteq \mathcal{E},$$

which implies that \mathcal{E} is a $(1 + \varepsilon)d$ -rounding of $\text{conv}(\mathcal{A})$.

Corollary 4.1. *Given $\mathcal{A} = \{a^1, \dots, a^m\} \subset \mathbb{R}^d$ and $\varepsilon > 0$, Algorithm 4.2 computes a $(1 + \varepsilon)d$ -rounding of $\text{conv}(\mathcal{A})$ in $O(d[\log d + 1/\varepsilon])$ iterations.*

So far we have only discussed obtaining $(1 + \varepsilon)d$ -roundings of $\text{conv}(\mathcal{A})$, not approximate MVEEs. But Khachiyan [23, Theorem 3] shows that a $(1 + \eta)$ -approximation of $\text{MVEE}(\mathcal{A})$ can easily be obtained from a solution to the ε -relaxed optimality conditions for $(\mathbf{D}(\mathcal{A}'))$ if $1 + \eta = (1 + \varepsilon)^{n/2}$. Thus to obtain a $(1 + \eta)$ -approximation of $\text{MVEE}(\mathcal{A})$, we merely apply our algorithm with $\varepsilon = (1 + \eta)^{2/(d+1)} - 1$. From this it is easily seen that our modified algorithm achieves the same complexity (5) as the KY algorithm to obtain a $(1 + \eta)$ -approximation of $\text{MVEE}(\mathcal{A})$, and provides an η -core set \mathcal{X} with the same asymptotic size (6). (Note that each iteration requires $O(md)$ arithmetic operations.)

Corollary 4.2. *Given $\mathcal{A} = \{a^1, \dots, a^m\} \subset \mathbb{R}^d$ and $\eta > 0$, Algorithm 4.2 computes a $(1 + \eta)$ -approximation of $\text{MVEE}(\mathcal{A})$ in $O(md^2([(1 + \eta)^{2/(d+1)} - 1]^{-1} + \log d))$ arithmetic operations and returns an η -core set $\mathcal{X} \subseteq \mathcal{A}$ such that $|\mathcal{X}| = O(d[(1 + \eta)^{2/(d+1)} - 1]^{-1} + d \log d)$.*

5. Final remarks

In this paper, we have established a close relationship between Khachiyan’s BCD method that computes an approximate rounding of the convex hull of a finite set of vectors and the ellipsoid method using two-sided deepest symmetric cuts. Motivated by this, we have proposed an efficient way to implement Khachiyan’s BCD method. We have also proposed a modification of the KY algorithm that computes an approximate solution to a more complete set of approximate optimality conditions than used by either the KY or Khachiyan’s algorithm. In addition, our algorithm maintains the same asymptotic complexity result and the same bound on core sets as the KY algorithm.

The “dropping idea” in our algorithm can also be applied in different forms yielding different variants of our algorithm. For instance, if $\varepsilon_k = \varepsilon_-$ at iteration k of Algorithm 4.2, β_k can be set to its upper bound even if the optimal solution of the linesearch problem defining β_k is in the interior of the range as long as the objective function value does not thereby decrease. This more aggressive dropping technique would lead to the same complexity analysis as that of

Algorithm 4.2 and would likely return smaller core sets. We could even consider reducing each positive component of p^k to zero at every iteration, as long as the corresponding $(q^i)^T A(p^k)^{-1} q^i$ is less than n and the objective function value does not increase, but this would likely be much more expensive.

We have implemented Algorithm 4.2 and made a preliminary comparison to the Khachiyan and KY methods. We solved random problems of sizes (d, m) varying from $(10, 200)$ to $(30, 30000)$ generated as in [3], terminating when we obtained a $(1 + \varepsilon)d$ -rounding for $\varepsilon = 10^{-3}$. From the baseline Khachiyan method, the KY algorithm reduced the time required by 10–18% and the number of iterations by 2–7%, while the new method reduced the time required by 93–98% and the number of iterations by 95–98%. The Khachiyan method always has $p^k > 0$, so all core sets have size m . By contrast, the KY and the new method produce much smaller core sets, typically of size at most $5d$ and never bigger than $10d$. The new method always gave a core set smaller than the KY algorithm, with reductions between 17% and 24%. More extensive computational results can be found in [3].

We finally discuss how to achieve an even smaller core set at the expense of considerable standard linear algebra. The idea is to rewrite $A(p^k)$ at the end of Algorithm 4.2 as a linear combination of fewer matrices than the number of positive components of p^k . One can find a basic feasible solution p to the system $\sum_{i=1}^m p_i q^i (q^i)^T = A(p^k)$, $p \geq 0$. (The equation corresponding to the bottom right entry of the matrix ensures that the sum of the components of p is 1.) Such a solution can be computed in $O(md^4)$ operations and would have at most $(d + 1)(d + 2)/2$ positive components, matching the John [22] bound. Indeed, this system has $n(n + 1)/2 = (d + 1)(d + 2)/2$ rows (note that all matrices are symmetric, so we only need to equate the entries on and above the diagonal) and m columns. We need $O(d^6)$ operations to compute an initial $O(d^2) \times O(d^2)$ basis inverse, and then at most m iterations requiring $O(d^4)$ operations each to get to a basic feasible solution. Since we are implicitly assuming $m > n(n + 1)/2$, the total is $O(md^4)$ operations. However, note that, for constant ε and η , this amount of work dominates all the work performed so far, which is $O(md^2(\log d + 1/\varepsilon))$ for a $(1 + \varepsilon)d$ -rounding of \mathcal{A} , or $O(md^2(\log d + d/\eta))$ for a $(1 + \eta)$ -approximation of $MVEE(\mathcal{A})$.

References

- [1] I. Adler, R. Shamir, A randomization scheme for speeding up algorithms for linear and convex quadratic programming problems with a high constraints-to-variables ratio, *Math. Programming* 61 (1993) 39–52.
- [2] P.K. Aggarwal, R. Poreddy, K.R. Varadarajan, H. Yu, Practical methods for shape fitting and kinetic data structures using core sets, in: *Proceedings of the 20th Annual ACM Symposium on Computational Geometry*, 2004, pp. 263–272.
- [3] D.S. Ahipasaoglu, P. Sun, M.J. Todd, Linear convergence of a modified Frank–Wolfe algorithm for computing minimum volume enclosing ellipsoids, Technical Report TR 1452, School of Operations Research and Industrial Engineering, Cornell University, Ithaca, New York, 2006.
- [4] K.M. Anstreicher, Improved complexity for maximum volume inscribed ellipsoids, *SIAM J. Optim.* 13 (2) (2003) 309–320.
- [5] C.L. Atwood, Sequences converging to D-optimal designs of experiments, *Ann. Statist.* 1 (1973) 342–352.
- [6] G. Barequet, S. Har-Peled, Efficiently approximating the minimum-volume bounding box of a point set in three dimensions, *J. Algorithms* 38 (2001) 91–109.
- [7] M. Bădoiu, K.L. Clarkson, Smaller core-sets for balls, in: *Proceedings of the 14th ACM-SIAM Symposium on Discrete Algorithms*, 2003, pp. 801–802.
- [8] M. Bădoiu, S. Har-Peled, P. Indyk, Approximate clustering via core-sets, in: *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, 2002, pp. 250–257.
- [9] C. Bouville, Bounding ellipsoid for ray-fractal intersection, in: *SIGGRAPH*, 1985, pp. 45–52.
- [10] B.P. Burrell, M.J. Todd, The ellipsoid method generates dual variables, *Math. Oper. Res.* 10 (1985) 688–700.
- [11] T.M. Chan, Faster core-set constructions and data stream algorithms in fixed dimensions, in: *Proceedings of the 20th Annual ACM Symposium on Computational Geometry*, 2004, pp. 152–159.
- [12] B. Chazelle, J. Matoušek, On linear-time deterministic algorithms for optimization problems in fixed dimension, in: *Proceedings of the 4th ACM-SIAM Symposium on Discrete Algorithms*, 1993, pp. 281–290.
- [13] D. Eberly, *3D Game Engine Design*, Morgan Kaufmann, Los Altos, CA, 2001.
- [14] V.V. Fedorov, *Theory of Optimal Experiments*, Academic Press, New York, 1972.
- [15] M. Frank, P. Wolfe, An algorithm for quadratic programming, *Nav. Res. Logist. Quart.* 3 (1956) 95–110.
- [16] B. Gärtner, S. Schönherr, Exact primitives for smallest enclosing ellipses, in: *Proceedings of the 13th Annual ACM Symposium on Computational Geometry*, 1997, pp. 430–432.
- [17] P.E. Gill, W. Murray, M.A. Saunders, Methods for computing and modifying the LDV factors of a matrix, *Math. Comput.* 29 (1975) 1051–1077.
- [18] F. Glineur, Pattern separation via ellipsoids and conic programming, Master’s Thesis, Faculté Polytechnique de Mons, Belgium, 1998.
- [19] D. Goldfarb, M.J. Todd, Modifications and implementation of the ellipsoid algorithm for linear programming, *Math. Programming* 23 (1982) 1–19.
- [20] M. Grötschel, L. Lovász, A. Schrijver, *Geometric Algorithms and Combinatorial Optimization*, Springer, New York, 1988.
- [21] J. Guélat, P. Marcotte, Some comments on Wolfe’s ‘away step’, *Math. Programming* 35 (1986) 110–119.

- [22] F. John, Extremum problems with inequalities as subsidiary conditions, in: *Studies and Essays, presented to R. Courant on his 60th birthday January 8, 1948*, pp. 187–204. Interscience, New York, 1948, reprinted in: J. Moser, (ed.), *Fritz John, Collected Papers*, vol. 2, Birkhäuser, Boston, 1985, pp. 543–560.
- [23] L.G. Khachiyan, Rounding of polytopes in the real number model of computation, *Math. Oper. Res.* 21 (1996) 307–320.
- [24] L.G. Khachiyan, M.J. Todd, On the complexity of approximating the maximal inscribed ellipsoid for a polytope, *Math. Programming* 61 (1993) 137–159.
- [25] P. Kumar, J.S.B. Mitchell, E.A. Yıldırım, Approximate minimum enclosing balls in high dimensions using core-sets, *ACM J. Experimental Algorithmics* 8 (1) (2003).
- [26] P. Kumar, E.A. Yıldırım, Minimum volume enclosing ellipsoids and core sets, *J. Optim. Theory Appl.* 126 (1) (2005) 1–21.
- [27] J. Matoušek, M. Sharir, E. Welzl, A subexponential bound for linear programming, in: *Proceedings of the 8th Annual Symposium on Computational Geometry*, 1992, pp. 1–8.
- [28] A.S. Nemirovskii, On self-concordant convex-concave functions, *Optim. Methods and Software* 11/12 (1999) 303–384.
- [29] Yu.E. Nesterov, Rounding of convex sets and efficient gradient methods for linear programming problems, Discussion Paper 2004-4, CORE, Catholic University of Louvain, Louvain-la-Neuve, Belgium, 2004.
- [30] Yu.E. Nesterov, A.S. Nemirovskii, *Interior Point Polynomial Methods in Convex Programming*, SIAM Publications, Philadelphia, PA, 1994.
- [31] B.W. Silverman, D.M. Titterton, Minimum covering ellipses, *SIAM J. Sci. Statist. Comput.* 1 (1980) 401–409.
- [32] S. Silvey, D. Titterton, A geometric approach to optimal design theory, *Biometrika* 62 (1973) 21–32.
- [33] P. Sun, R.M. Freund, Computation of minimum volume covering ellipsoids, *Oper. Res.* 52 (2004) 690–706.
- [34] D.M. Titterton, Optimal design: some geometrical aspects of D -optimality, *Biometrika* 62 (2) (1975) 313–320.
- [35] D.M. Titterton, Estimation of correlation coefficients by ellipsoidal trimming, *Appl. Statist.* 27 (3) (1978) 227–234.
- [36] M.J. Todd, On minimum volume ellipsoids containing part of a given ellipsoid, *Math. Oper. Res.* 7 (1982) 253–261.
- [37] K.C. Toh, Primal-dual path-following algorithms for determinant maximization problems with linear matrix inequalities, *Comput. Optim. Appl.* 14 (1999) 309–330.
- [38] L. Vandenberghe, S. Boyd, S. Wu, Determinant maximization with linear matrix inequality constraints, *SIAM J. Matrix Anal. Appl.* 19 (2) (1998) 499–533.
- [39] E. Welzl, Smallest enclosing disks (balls and ellipsoids), in: H. Maurer (Ed.), *Proceedings of New Results and New Trends in Computer Science*, vol. 555, Berlin, Germany, June 1991, *Lecture Notes in Computer Science*, Springer, Berlin, pp. 359–370.
- [40] P. Wolfe, Convergence theory in nonlinear programming, in: J. Abadie (Ed.), *Integer and Nonlinear Programming*, North-Holland, Amsterdam, 1970, pp. 1–36.
- [41] H.P. Wynn, The sequential generation of D -optimum experimental design, *Ann. Math. Statist.* 41 (1970) 1655–1664.
- [42] Y. Zhang, An interior-point algorithm for the maximum-volume ellipsoid problem, Technical Report TR98-15, Department of Computational and Applied Mathematics, Rice University, 1998.
- [43] Y. Zhang, L. Gao, On numerical solution of the maximum volume ellipsoid problem, *SIAM J. Optim.* 14 (2003) 53–76.