Lecture Notes on Semidefinite Programming

THE ELLIPSOID METHOD

Fernando Mário de Oliveira Filho
Instituto de Matemática e Estatística, Universidade de São Paulo

7 March, 2014

One of the greatest breakthroughs of the last century in the theory of optimization was Khachiyan's result [2] that the ellipsoid method of nonlinear programming, introduced by Shor and Yudin and Nemirovskii (cf. Chapter 13 of Schrijver [4]), can be extended so as to provide a polynomial-time algorithm for linear programming. Before the ellipsoid method, algorithms for linear programming such as the simplex method could be fast in practice, but had exponential worst-case performance. Kachiyan's result settled the long-standing problem of determining the complexity of linear programming.

The ellipsoid method proved to be inefficient in practice, and to this day it is not used as an algorithm to solve practical problems (cf. §13.5 in Schrijver [4]). In the domain of linear programming, it was supplanted by interior-point methods, which are also polynomial-time algorithms that moreover perform well in practice. A fundamental result of Grötschel, Lovász, and Schrijver [1], however, secured its position as one of the main theoretical tools in optimization, used to prove the polynomial-time solvability of many classes of optimization problems, in particular of many combinatorial optimization problems.

In these notes we will study the ellipsoid method and the main result of Grötschel, Lovász, and Schrijver and see what are the implications to conic programming.

1. Introduction

In what follows, the set of $n \times n$ symmetric matrices is denoted by $\mathcal{S}^n$ and $A \succeq 0$ means that $A \in \mathcal{S}^n$ is positive semidefinite. The trace inner product of A, $B \in \mathcal{S}^n$ is $\langle A, B \rangle = \operatorname{tr} AB$.

Let $K \subseteq \mathbb{R}^n$ be a nonempty, compact, and convex set. Following Grötschel, Lovász, and Schrijver, consider the following two problems in relation to K:

> *Strong optimization problem.* Given a vector $c \in \mathbb{R}^n$, find a vector $x^* \in K$ that maximizes $c^\mathsf{T} x$ in K.

> *Strong separation problem.* Given a vector $y \in \mathbb{R}^n$, decide that $y \in K$, or else find a hyperplane separating y from K, i.e., find a vector $a \in \mathbb{R}^n$ such that $a^\mathsf{T} y > \max\{\, a^\mathsf{T} x : x \in K \,\}$.

Example 1 (Linear programming). The convex set $K \subseteq \mathbb{R}^n$ could be a polytope, given as a system of linear inequalities

$$a_i^\mathsf{T} x \leq b_i \quad \text{for } i = 1, \ldots, m.$$

Then the optimization problem is a linear programming problem. The separation problem consists of, given a vector $y \in \mathbb{R}^n$, determining whether y satisfies all the linear inequalities, or else finding a violated inequality. This can be done by checking each inequality; if they are all satisfied, then $y \in K$; if for some i we have $a_i^\mathsf{T} y > b_i$, then we have our separating hyperplane.

Example 2 (Semidefinite programming). Let C, A_1, …, $A_m \in \mathcal{S}^n$ and b_1, …, b_m be real numbers, and let K be the feasible region of the semidefinite programming problem

$$\begin{aligned} \text{maximize } & \langle C, X\rangle \\ & \langle A_i, X\rangle = b_i \quad \text{for } i = 1, \ldots, m, \\ & X \succeq 0. \end{aligned}$$

The optimization problem for K is exactly the above problem. The separation problem consists of, given a symmetric matrix Y, decide whether it is positive semidefinite and satisfies all equalities above, or find a hyperplane separating Y from the feasible region K.

If Y violates one of the linear equalities, then we immediately have a separating hyperplane, just as in the case of linear programming. Suppose Y satisfies all equalities but is not positive semidefinite. Then for some $a \in \mathbb{R}^n$ we have $a^\mathsf{T} Y a < 0$, and a separating hyperplane is $\langle aa^\mathsf{T}, X\rangle \geq 0$.

Grötschel, Lovász, and Schrijver show that, for any given class of "convex bodies" (a concept that will be defined precisely below), the optimization and separation problems are equivalent, in the sense that a polynomial-time algorithm for any of the two implies a polynomial-time algorithm for the other. There is a technical point, however. In the way the problems are defined above, it is not clear for instance that there is a solution to the optimization problem having only rational coordinates, so it is not even clear how a solution can be specified. Therefore, from the point of view of computational complexity the problems above are not well-defined.

Further assumptions made on K would guarantee that there is always an optimal solution that is rational. For instance, this is the case when one considers rational polyhedra, that is, polyhedra given by systems of rational inequalities. Instead of making further assumptions on K, however, Grötschel, Lovász, and Schrijver define weaker versions of the optimization and separation problems. If $d(x, K)$ denotes the Euclidean distance between x and K, then we consider the problems:

> *Weak optimization problem.* Given a vector $c \in \mathbb{Q}^n$ and a number $\varepsilon > 0$, find a vector $\hat{x} \in \mathbb{Q}^n$ that is ε-close to K, i.e., $d(\hat{x}, K) \leq \varepsilon$, and that almost maximizes $c^\mathsf{T} x$ on K, i.e., for every $x \in K$, $c^\mathsf{T} x \leq c^\mathsf{T} \hat{x} + \varepsilon$.
>
> *Weak separation problem.* Given a vector $y \in \mathbb{Q}^n$ and a number $\delta > 0$, do one of the following: (i) conclude that $d(y, K) \leq \delta$ (i.e., y is almost feasible); or (ii) find a vector $a \in \mathbb{Q}^n$ such that $\|a\| \geq 1$ and for every $x \in K$, $a^\mathsf{T} x \leq a^\mathsf{T} y + \delta$.

Let us now state more precisely the result of Grötschel, Lovász, and Schrijver. A *convex body* is a quintuple (K, n, a_0, r, R), where $n \geq 2$, $K \subseteq \mathbb{R}^n$ is a convex set, $a_0 \in K$, and

$$B(a_0, r) \subseteq K \subseteq B(a_0, R),$$

where $B(p, r)$ is the closed ball of radius r and center p. The assumption that K is contained in a ball is quite natural from an optimization perspective. The assumption, however, that K contains a ball is not so natural. It means that K is full-dimensional, or at least that one knows the affine subspace containing K (cf. Exercise 1). It can be proven that such an assumption is essential (cf. §3 in Grötschel, Lovász, and Schrijver [1]).

Let $\Pi \subseteq \{0, 1\}^*$ be a language. Consider a class of convex bodies

$$\mathcal{K} = \{\, K_\sigma : \sigma \in \Pi \text{ and } K_\sigma \text{ is a convex body} \,\}.$$

So each body has an associated encoding in Π that is used to represent it.

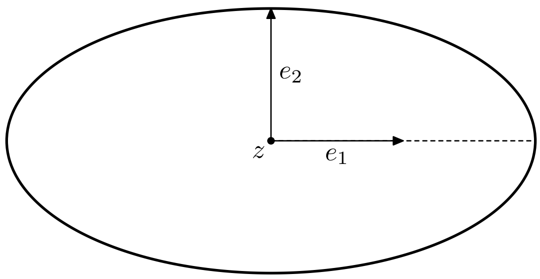

Figure 1. An ellipsoid in $\mathbb{R}^2$ with center z and axes given by the unit vectors $e_1 = (1, 0)$ and $e_2 = (0, 1)$. The longer axis has length 2, whereas the shorter has length 1. This ellipsoid is $\mathrm{ell}(z, A)$ with $A = 4e_1e_1^{\mathsf{T}} + e_2e_2^{\mathsf{T}}$.

The input of the weak optimization problem is then a tuple $(\sigma, n, a_0, r, R, c, \varepsilon)$, where $\sigma \in \Pi$, $K_\sigma = (K, n, a_0, r, R)$, $c \in \mathbb{Q}^n$, and $\varepsilon > 0$. An algorithm for the weak optimization problem for the class $\mathcal{K}$ receives as input such a tuple and solves the weak optimization problem for the convex body K_σ. It runs in polynomial time if its running time is polynomial in the size of the input tuple, which is basically the length of σ plus the lengths of binary representations of n, a_0, r, R, c, and ε.

We may similarly define an algorithm to solve the weak separation problem for $\mathcal{K}$ and what it means for it to be a polynomial-time algorithm. So we have the main result of Grötschel, Lovász, and Schrijver:

Theorem 1. *Let $\mathcal{K}$ be a class of convex bodies. There is a polynomial-time algorithm to solve the weak optimization problem for $\mathcal{K}$ if and only if there is a polynomial-time algorithm to solve the weak separation problem for $\mathcal{K}$.*

Notice that ε is part of the input. So the running time of the algorithm must depend polynomially on the length of a binary representation of ε, which is $O(|\log \varepsilon|)$. This means that by taking a sequence $\varepsilon = 1/2, 1/4, \ldots$ we get a sequence of approximations which converge exponentially fast, while the running time increases polynomially from one value of ε to the next. This exponential convergence rate is what allows Khachiyan to find an optimal solution to a linear programming problem (basically by rounding) instead of just an approximation, and it also allows Grötschel, Lovász, and Schrijver to derive many combinatorial applications of the ellipsoid method.

2. Ellipsoids

Let $z \in \mathbb{R}^n$ and $A \in \mathcal{S}^n$ be a positive definite matrix. An *ellipsoid* is a set

$$\mathrm{ell}(z, A) = \{\, x \in \mathbb{R}^n : (x - z)^{\mathsf{T}} A^{-1} (x - z) \leq 1 \,\}.$$

Vector z is the center of the ellipsoid. Notice that $\mathrm{ell}(z, I)$ is the unit ball with center z.

Since A is positive definite, there is an orthonormal basis $u_1, \ldots, u_n$ of $\mathbb{R}^n$ and positive numbers $\lambda_1, \ldots, \lambda_n$ such that

$$A = \lambda_1 u_1 u_1^{\mathsf{T}} + \cdots + \lambda_n u_n u_n^{\mathsf{T}}.$$

Vectors u_i give the axes of the ellipsoid and numbers λ_i give the squares of the half-lenghts of the axes; see Figure 1.

The fact that A is positive definite also allows us to write $A = U^2$ for some nonsingular matrix $U \in \mathcal{S}^n$. So we see that $\mathrm{ell}(z, A)$ is an affine transformation of the unit ball:

$$\mathrm{ell}(z, A) = \{\, Ux + z : \|x\| \leq 1 \,\}.$$

This makes it easy to compute the volume of an ellipsoid. For $A \in \mathcal{S}^n$ positive definite, let $f_A\colon \mathbb{R}^n \to \{0,1\}$ be such that $f_A(x) = 1$ if and only if $x \in \operatorname{ell}(0, A)$. Then if $A = U^2$ as above we have

$$\begin{aligned}
\operatorname{vol}\operatorname{ell}(z, A) = \operatorname{vol}\operatorname{ell}(0, A) &= \int_{\mathbb{R}^n} f_A(x)\,dx \\
&= \int_{\mathbb{R}^n} f_I(U^{-1}x)\,dx \\
&= |\det U| \int_{\mathbb{R}^n} f_I(x)\,dx \\
&= \sqrt{\det A}\operatorname{vol} B_n,
\end{aligned}$$

where B_n is the unit ball.

3. The method with infinite precision

Let us assume for now that we can compute with real numbers and that we have an oracle to solve the strong separation problem. Let us see how we can find solutions to the optimization problem that are at most ε-distant from the optimal. To this end, the following lemma will be useful (cf. Theorem 13.1 in Schrijver [4]).

Lemma 2. *Let $z \in \mathbb{R}^n$ and let $A \in \mathcal{S}^n$ be positive definite. Given a vector $a \in \mathbb{R}^n$, the ellipsoid* $\operatorname{ell}(z', A')$ *with*

$$z' = z + \frac{1}{n+1} \cdot \frac{Aa}{\sqrt{a^\mathsf{T} Aa}}$$

and

$$A' = \frac{n^2}{n^2 - 1}\left(A - \frac{2}{n+1} \cdot \frac{Aaa^\mathsf{T} A}{a^\mathsf{T} Aa}\right)$$

is the unique minimum-volume ellipsoid containing $\operatorname{ell}(z, A) \cap \{\, x \in \mathbb{R}^n : a^\mathsf{T} x \geq a^\mathsf{T} z \,\}$. *Moreover,*

$$\frac{\operatorname{vol}\operatorname{ell}(z', A')}{\operatorname{vol}\operatorname{ell}(z, A)} < e^{-1/(2n+2)}. \tag{1}$$

Proof. Exercise 2 asks you to prove that when $z = 0$ and $A = I$, the ellipsoid $\operatorname{ell}(z', A')$ given as above indeed is the minimum-volume ellipsoid containing one of the halves of the unit ball. Since ellipsoids are all affine transformations of the unit ball, the parameters for the general case follow from those for the unit ball.

To see (1), we may also assume that $A = I$. Let $u_1, \ldots, u_n$ be an orthonormal basis of $\mathbb{R}^n$, where $u_1 = \|a\|^{-1}a$. Note $u_1, \ldots, u_n$ are all eigenvectors of I each with associated eigenvalue 1. So A' has eigenvalues

$$\frac{n^2}{n^2-1}\left(1 - \frac{2}{n+1}\right) \qquad \text{and} \qquad \frac{n^2}{n^2-1} \text{ with multiplicity } n-1.$$

This implies immediately that

$$\sqrt{\det A'} = \left(\frac{n^2}{n^2-1}\right)^{n/2}\left(1 - \frac{2}{n+1}\right)^{1/2} = \left(\frac{n^2}{n^2-1}\right)^{(n-1)/2}\left(\frac{n}{n+1}\right).$$

From this we have that

$$\frac{\operatorname{vol}\operatorname{ell}(z', A')}{\operatorname{vol}\operatorname{ell}(z, A)} = \sqrt{\det A'} = \left(\frac{n^2}{n^2-1}\right)^{(n-1)/2}\left(\frac{n}{n+1}\right) < e^{-1/(2n+2)},$$

where the last inequality follows from the fact that $1 + x < e^x$ when $x \neq 0$. ∎

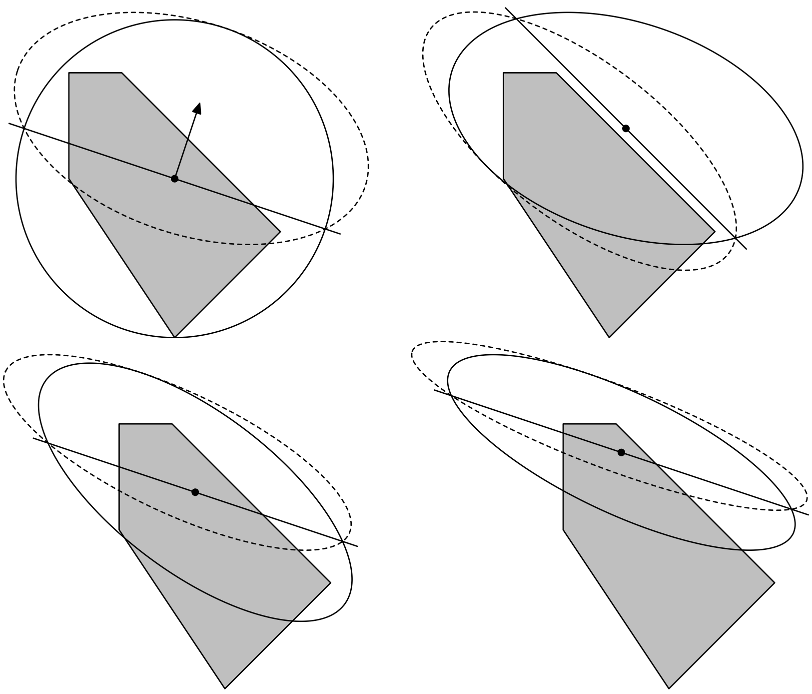

Figure 2. From left-to-right and top-to-bottom, the first four steps of the ellipsoid method with infinite precision. The optimization direction is shown on the first step. On each step the convex body over which we optimize is shown in gray, the current ellipsoid is draw with a full line and its center is marked, the separating hyperplane is shown, and the next ellipsoid to be considered is drawn with a dashed line.

Let (K, n, a_0, r, R) be a convex body. Given $c \in \mathbb{R}^n$ and $\varepsilon > 0$, set

$$N = (2n^2 + 2n)\left\lceil \ln \frac{2R^2\|c\|}{r\varepsilon} \right\rceil.$$

We start be considering a first ellipsoid $\mathrm{ell}(x_0, A_0)$, actually a ball, that contains K, by setting

$$x_0 = a_0 \quad \text{and} \quad A_0 = R^2 I.$$

At a given step $k \geq 0$ of the algorithm we proceed as follows. We run the separation oracle for x_k. If $x_k \in K$, then set $a = c$. If not, then let $d \in \mathbb{R}^n$ give the separating hyperplane, and set $a = -d$. Then we find x_{k+1} and A_{k+1} according to Lemma 2 so that $\mathrm{ell}(x_{k+1}, A_{k+1})$ is the smallest ellipsoid containing $\mathrm{ell}(x_k, A_k) \cap \{\, x \in \mathbb{R}^n : a^\mathsf{T} x \geq a^\mathsf{T} x_k \,\}$ and we proceed to the next step. Figure 2 shows an example of the first few steps of the method applied to a given convex body.

By construction, every ellipsoid computed contains an optimal solution of the original problem. Let j be such that

$$c^\mathsf{T} x_j = \max\{\, c^\mathsf{T} x_k : k = 0, \ldots, N \text{ and } x_k \in K \,\}$$

and set $\hat{x} = x_j$. Let us analyze how close $\hat{x}$ is to an optimal solution.

To this end, let x^* be an optimal solution to the optimization problem. We know that the ball $B(x_0, r)$ is contained in K. Notice that $B(x_0, r) \cap \{\, x \in \mathbb{R}^n : c^\mathsf{T} x = c^\mathsf{T} x_0 \,\}$ is an $(n-1)$-dimensional ball of radius r and center x_0.

Consider then the cone whose base is this ball and whose vertex is x^*. This cone is contained in K because K is convex. The portion of the cone contained in the half-space $\{\, x \in \mathbb{R}^n : c^\mathsf{T} x \geq c^\mathsf{T} \hat{x} \,\}$ is by construction contained in the ellipsoid $\mathrm{ell}(x_N, A_N)$.

The volume of this portion of the cone is

$$\frac{\operatorname{vol} B_{n-1} r^{n-1}(c^{\mathsf T}x^* - c^{\mathsf T}x_0)}{n\|c\|}\left(\frac{c^{\mathsf T}x^* - c^{\mathsf T}\hat{x}}{c^{\mathsf T}x^* - c^{\mathsf T}x_0}\right)^n \le \operatorname{vol}\operatorname{ell}(x_N, A_N)$$
$$\le e^{-N/(2n+2)} \operatorname{vol} B_n R^n,$$

where the last inequality comes from the repeated application of Lemma 2.

From this we get

$$c^{\mathsf T}x^* - c^{\mathsf T}\hat{x} \le e^{-N/(2n^2+2n)} R\left(\frac{n \operatorname{vol} B_n}{\operatorname{vol} B_{n-1}}\right)^{1/n}\left(\frac{c^{\mathsf T}x^* - c^{\mathsf T}x_0}{r}\right)^{\frac{n-1}{n}}\|c\|^{1/n}.$$

Now, notice that

$$c^{\mathsf T}x^* - c^{\mathsf T}x_0 = c^{\mathsf T}(x^* - x_0) \le \|c\|\|x^* - x_0\| \le R\|c\|,$$

and hence

$$c^{\mathsf T}x^* - c^{\mathsf T}\hat{x} \le 2e^{-N/(2n^2+2n)}\frac{R^2}{r}\|c\| \le \varepsilon.$$

So, to achieve a precision of ε, we need a polynomial number of iterations. Of course, in our description we went over many details. The assumption that one can work with real numbers is quite a strong one, and it is *a priori* not clear that one can work only with rational numbers and rational approximations.

4. The ellipsoid method with rational arithmetic

In the previous section we described the ellipsoid method assuming that one could compute with real numbers. Most of the work in proving Theorem 1 is to show that one can work with finite precision and rational arithmetic.

This requires a slight modification of the ellipsoid method as we presented before and careful error estimates to arrive at the result. Here is a precise description of the algorithm.

Consider a convex body (K, n, a_0, r, R), a vector $c \in \mathbb{Q}^n$, and $\varepsilon > 0$. Without loss of generality assume $\varepsilon < r$, $\|c\| \ge 1$, and $n \ge 2$. Assume there is an algorithm to solve the weak separation problem for this convex body.

We start by setting

$$N = 4n^2\left\lceil \ln\frac{2R^2\|c\|}{r\varepsilon}\right\rceil, \quad \delta = \frac{4^{-N}R^2}{300n}, \quad \text{and} \quad p = 5N.$$

As before, the initial ellipsoid $\operatorname{ell}(x_0, A_0)$ is just a ball containing K, that is, we set $x_0 = a_0$ and $A_0 = R^2 I$.

At a given step $k \ge 0$ of the algorithm, we run the weak separation algorithm for x_k. If $d(x_k, K) \le \delta$, then k is called a *feasible index* and $a = c$. If the separation algorithm gives a hyperplane d, then $a = -d$. The ellipsoid $\operatorname{ell}(x_{k+1}, A_{k+1})$ is then defined as follows. We let

$$b_k = \frac{A_k a}{\sqrt{a^{\mathsf T}A_k a}}, \quad x_k^* = x_k + \frac{1}{n+1}b_k, \quad \text{and} \quad A_k^* = \frac{2n^3+3}{2n^2}\left(A_k - \frac{2}{n+1}b_k b_k^{\mathsf T}\right).$$

Then we obtain x_{k+1} from x_k^* and A_{k+1} from A_k^* by rounding every number to p binary digits after the decimal point.

Notice that, with $n^2/(n^2-1)$ in place of $(2n^3+3)/(2n^2)$, $\operatorname{ell}(x_k^*, A_k^*)$ would be the minimum-volume ellipsoid containing one of the halves of $\operatorname{ell}(x_k, A_k)$, cf. Lemma 2, as we used in the description of the method with infinite precision. The enlarged ellipsoid is taken to deal with rounding errors.

Theorem 3. *Let j be such that*

$$c^{\mathsf{T}} x_j = \max\{\, c^{\mathsf{T}} x_k : k = 0, \ldots, N-1 \text{ is feasible} \,\}$$

and set $\hat{x} = x_j$. Then, if x^* is an optimal solution of the optimization problem, we have $c^{\mathsf{T}}\hat{x} \geq c^{\mathsf{T}} x^* - \varepsilon$.

The theorem says that after a polynomial number of iterations of the ellipsoid method we find an almost maximizer, that is, a solution to the weak optimization problem. If the weak separation algorithm runs in polynomial time, since the number of digits used is always polynomial, we see that the whole algorithm runs in polynomial time. The analysis of the algorithm is now more involved because we do not assume infinite precision. Details can be found in the paper by Grötschel, Lovász, and Schrijver [1].

So we have one direction of Theorem 1, namely that a polynomial-time algorithm for separation gives a polynomial-time algorithm for optimization. The other direction is much easier to derive; see Theorem 3.1 in Grötschel, Lovász, and Schrijver [1].

5. Consequences to conic programming

Consider a class of conic programming problems, the k-th problem of which is

$$\begin{aligned} \text{maximize } & c_k^{\mathsf{T}} x \\ & a_{k,i}^{\mathsf{T}} x \leq b_i \quad \text{for } i = 1, \ldots, m_k, \\ & x \in C_k, \end{aligned}$$

where $C_k \subseteq \mathbb{R}^{n_k}$ is a closed and convex cone. Let F_k be the feasible region of this problem

To be even more precise, we let $\Pi = \{\, \sigma_k : k \geq 0 \,\}$ be the language such that σ_k is an encoding of the $a_{k,i}$ and b_i. Assume moreover that for each $\sigma \in \Pi$ there is $a_0 \in \mathbb{R}^{n_k}$ and numbers r and R such that $K_\sigma = (F_k, n_k, a_0, r, R)$ is a convex body. Then $\mathcal{K} = \{\, K_\sigma : \sigma \in \Pi \,\}$ is a class of convex bodies.

Theorem 1 says that there is a polynomial-time algorithm to solve the weak optimization problem for $\mathcal{K}$ if and only if there is a polynomial-time algorithm to solve the weak separation problem for $\mathcal{K}$. To solve the weak separation problem for the k-th conic programming problem we have to check whether a given vector satisfies all linear inequalities and whether it is in the cone C_k. To test whether the vector satisfies all linear inequalities it suffices to test each of them, and this takes polynomial time in the input size. Hence the weak separation problem boils down to the weak separation problem for the cone C_k. In other words: the complexity of conic programming is determined by the complexity of the cone.

This implies that linear programming and semidefinite programming problems can be solved in polynomial time. On the other hand, an example of a class of conic programming problems that are hard to solve is given by copositive programming. Recall that a matrix $A \in \mathcal{S}^n$ is *copositive* if for all $x \in \mathbb{R}^n$, $x \geq 0$, we have $x^{\mathsf{T}} A x \geq 0$. The set of all copositive matrices is a closed and convex cone, but the weak separation problem for it is NP-hard. So Theorem 1 implies that, unless P = NP, there is no polynomial-time algorithm to solve copositive programming problems.

The requirement that we should know an interior point of the feasible region together with balls contained in, and containing the, feasible region can be an obstacle in applying Theorem 1 directly. In many cases however this problem can be avoided. For instance, linear programming problems are very well-behaved, and stronger versions of Theorem 1 hold for them (cf. Schrijver [4]). One can also work-around this issue if the affine subspace spanned by the feasible region is known (see Exercise 1). But, as mentioned before, in general such a requirement cannot be removed (cf. §3 in

Grötschel, Lovász, and Schrijver [1]), and so Theorem 1 does not apply to any class of conic programs.

As a concrete example of how Theorem 1 applies to semidefinite programming, let us consider the problem of computing the Lovász theta number [3] of a graph.

Let $G = (V, E)$ be a graph. We want to find the optimal value $\vartheta(G)$ of the following semidefinite programming problem:

$$\begin{aligned} \text{maximize } & \langle J, X\rangle \\ & \operatorname{tr} X = 1, \\ & X(u, v) = 0 \quad \text{if } uv \in E, \\ & X\colon V \times V \to \mathbb{R} \text{ is positive semidefinite,} \end{aligned}$$

where J is the all-ones matrix. So we have one optimization problem for each graph G, and to apply Theorem 1 we want to describe a class of convex bodies indexed by representations of graphs, so that optimizing a certain objective function over the convex body associated with a given graph is the same as computing the Lovász theta number of the graph.

The feasible region of any problem above is not a full-dimensional subset of $\mathbb{R}^{V\times V}$, so we cannot apply Theorem 1 directly. Notice however that a matrix $X \in \mathbb{R}^{V\times V}$ that is a solution of our problem can be seen as a vector in $\mathbb{R}^N$, where

$$N = |V| + \binom{|V|}{2} - |E|,$$

by looking only at the upper-diagonal entries of X and disregarding those that correspond to edges and therefore are equal to 0.

Even in this space, the feasible region is not full-dimensional, because of the constraint $\operatorname{tr} X = 1$. But if we replace this constraint by $\operatorname{tr} X \leq 1$, the optimal value does not change, and the feasible region becomes full-dimensional. In fact, then the matrix $(2n)^{-1}I$ gives an interior point, and it is easy to find balls around this matrix contained in, and containing the, feasible region.

Exercise 3 shows a way how to solve the weak separation problem for the cone of positive semidefinite matrices in polynomial time. So Theorem 1 implies that the Lovász theta number of a graph can be approximated to any desired specified precision in polynomial time.

6. Exercises

1. Let $\Pi \subseteq \{0,1\}^*$ and for each $\sigma \in \Pi$ let $K_\sigma = (K, n, A, b, a_0, r, R)$, where:

1. $K \subseteq \mathbb{R}^n$ is a convex set;
2. $A \in \mathbb{Q}^{m\times n}$ and $b \in \mathbb{Q}^m$ are such that $H = \{\, x \in \mathbb{R}^n : Ax = b \,\}$ is the affine subspace spanned by K, and moreover the sizes of A and b are bounded by a polynomial on n and the sizes of σ, a_0, r, and R;
3. $a_0 \in K$ and $B(a_0, r) \cap H \subseteq K \subseteq B(a_0, R) \cap H$.

So $\mathcal{K} = \{\, K_\sigma : \sigma \in \Pi \,\}$ is not necessarily a class of convex bodies, because the bodies are not necessarily full-dimensional, but for each body we know the affine subspace it spans.

Fix $\eta > 0$. Given $\sigma \in \Pi$, let $u_1, \ldots, u_k \in \mathbb{Q}^n$ be an orthogonal basis of the subspace H with $1 - \eta \leq \|u_i\| \leq 1$ for $i = 1, \ldots, k$. Such a basis can be computed in polynomial-time on the sizes of A and b.

Now, let $f\colon H \to \mathbb{R}^k$ be such that $f(x) = (\alpha_1, \ldots, \alpha_k)$, where

$$x = \alpha_1 u_1 + \cdots + \alpha_k u_k.$$

(a) Estimate $\tilde{r}$ and $\tilde{R}$ so that $\tilde{K}_\sigma = (f(K), k, f(a_0), \tilde{r}, \tilde{R})$ is a convex body.

(b) Given a polynomial-time algorithm to solve the weak separation problem for the class $\mathcal{K}$, show how to derive a polynomial-time algorithm to solve the weak separation problem for $\tilde{\mathcal{K}} = \{ \tilde{K}_\sigma : \sigma \in \Pi \}$.

(c) Show that, given a polynomial-time algorithm to solve the weak separation problem for $\mathcal{K}$, one can solve the weak optimization problem for $\mathcal{K}$ in polynomial time.

2. In this exercise, we will prove that the ellipsoid described in Lemma 2 is a smallest ellipsoid containing the upper half of the unit ball, though we will not argue that the minimum-volume ellipsoid is unique.

Say we are given points $x_1, \ldots, x_N \in \mathbb{R}^n$ and wish to find the minimum-volume ellipsoid containing these points.

(a) Consider the optimization problem

$$\begin{aligned} \text{maximize } & n + \ln\det A \\ & (x_i - z)^\mathsf{T} A (x_i - z) \le 1 \quad \text{for } i = 1, \ldots, N, \\ & A \in \mathcal{S}^n, \quad A \succeq 0, \end{aligned} \tag{2}$$

where A and z are the variables. Show that if (z, A) is an optimal solution of this problem, then $\operatorname{ell}(z, A^{-1})$ is a minimum-volume ellipsoid containing $x_1, \ldots, x_N$.

(b) Consider the optimization problem

$$\begin{aligned} \text{minimize } & \textstyle\sum_{i=1}^N y_i - \ln\det \sum_{i=1}^N y_i x_i x_i^\mathsf{T} \\ & \textstyle\sum_{i=1}^N y_i x_i = 0, \\ & y_i \ge 0 \quad \text{for } i = 1, \ldots, N, \end{aligned} \tag{3}$$

where the y_i are the variables. Show that, if (z, A) is a feasible solution of (2) and y is a feasible solution of (3), then

$$n + \ln\det A \le \sum_{i=1}^N y_i - \ln\det \sum_{i=1}^N y_i x_i x_i^\mathsf{T}.$$

Hint: The arithmetic-geometric mean inequality says that if $\alpha_1, \ldots, \alpha_n \ge 0$, then

$$\prod_{i=1}^n \alpha_i \le \left(\frac{1}{n} \sum_{i=1}^n \alpha_i \right)^n.$$

Use this to prove that, given a positive semidefinite matrix $A \in \mathcal{S}^n$,

$$\det A \le \left(\frac{\operatorname{tr} A}{n} \right)^n.$$

(c) Use the previous item to prove the following direction of a theorem of John: Let $x_1, \ldots, x_N \in \mathbb{R}^n$ be unit vectors. If there are nonnegative numbers $y_1, \ldots, y_N$ such that

$$\sum_{i=1}^N y_i x_i = 0 \quad \text{and} \quad \sum_{i=1}^N y_i x_i x_i^\mathsf{T} = I,$$

then the unit ball $\operatorname{ell}(0, I)$ is a least-volume ellipsoid containing $x_1, \ldots, x_N$.

(d) Let $e_1, \ldots, e_n$ be the canonical basis of $\mathbb{R}^n$. Prove that the ellipsoid described in Lemma 2 with $z = 0$, $A = I$, and $a = e_1$ is a minimum-volume ellipsoid containing the points $e_1, \pm e_2, \ldots, \pm e_n$.

(e) Let $x_0, \ldots, x_n \in \mathbb{R}^n$ be affinely independent. Show that the minimum-volume ellipsoid containing $x_0, \ldots, x_n$ is $\operatorname{ell}(z, A)$ with

$$z = \frac{1}{n+1}\sum_{i=0}^{n} x_i \quad \text{and} \quad A = \frac{n}{n+1}\sum_{i=0}^{n}(x_i - z)(x_i - z)^{\mathsf{T}}.$$

3. Let $X \in \mathcal{S}^n$.

(a) Let $B \subseteq \{1, \ldots, n\}$ be a set of indices corresponding to a maximal linearly independent set of columns of X. Let $\tilde{X}$ be the principal submatrix of X with rows and columns given by B. Prove that X is positive semidefinite if and only if $\tilde{X}$ is positive definite.

(b) For $k = 1, \ldots, n$, the k-th principal minor of X is the matrix X_k consisting of the first k rows and columns of X. Suppose X is such that $\det X < 0$, but $\det X_1$, $\ldots$, $\det X_{n-1} > 0$. For $i = 1, \ldots, n$ let

$$a_i = (-1)^i M_{in},$$

where M_{in} is the determinant of the (i, n)-minor of X, that is, of the matrix obtained from X by removing row i and column n. Show that

$$a^{\mathsf{T}} X a = \det X_{n-1} \det X.$$

(c) Use the above, together with the fact that a matrix X is positive definite if and only if every principal minor has positive determinant (a fact called Sylvester's criterion), to give a polynomial-time algorithm that either concludes that a rational matrix $X \in \mathcal{S}^n$ is positive semidefinite, or finds $a \in \mathbb{Q}^n$ such that $a^{\mathsf{T}} X a < 0$.

4. Consider a class of semidefinite programming problems, the k-th problem of which being

$$\begin{aligned} &\text{minimize } x_k \\ &\qquad x_0 = 2, \\ &\qquad \begin{pmatrix} 1 & x_{i-1} \\ x_{i-1} & x_i \end{pmatrix} \succeq 0 \quad \text{for } i = 1, \ldots, k. \end{aligned}$$

Let F_k be the feasible region of the k-th problem of this class. We know the affine subspace spanned by F_k and it is easy to find interior points, so the ellipsoid method can be applied to the class of optimization problems above. Does this mean that any problem in the class can be solved in time that is bounded by a polynomial in k? Why?

7. References

[1] M. Grötschel, L. Lovász, and A. Schrijver, The ellipsoid method and its consequences in combinatorial optimization, *Combinatorica* 1 (1981) 169–197.

[2] L.G. Khachiyan, A polynomial algorithm in linear programming, *Soviet Math. Dokl.* 20 (1979) 191–194.

[3] L. Lovász, On the Shannon capacity of a graph, *IEEE Transactions on Information Theory* IT-25 (1979) 1–7.

[4] A. Schrijver, *Theory of Linear and Integer Programming*, John Wiley & Sons, Chichester, 1986.