7 The Ellipsoid Method

Consider a polytope $P = \{x \in \mathbb{R}^n : Ax \geqslant b\}$, given by a matrix $A \in \mathbb{Z}^{m \times n}$ and a vector $b \in \mathbb{Z}^m$. Assume for now that P is bounded and either empty or full-dimensional. Here, P is called *full-dimensional* if $\mathrm{Vol}(P) > 0$. The ellipsoid method takes the following steps to decide whether P is non-empty:

1. Let U be the largest absolute value among the entries of A and b, and define
   $\nu_0 = 0, \quad D_0 = n(nU)^{2n} I, \quad E_0 = E(\nu_0, D_0),$
   $V = (2n)^n (nU)^{n^2}, \quad \nu = n^{-n}(nU)^{-n^2(n+1)},$
   $t^* = \lceil 2(n+1)\log(V/\nu) \rceil.$
2. For $t = 0, \dots, t^*$, do the following:
   (a) If $t = t^*$ then stop; P is empty.
   (b) If $x_t \in P$ then stop; P is non-empty.
   (c) Find a violated constraint, i.e., a row j such that $a_j^T x_t < b_j$.
   (d) Let $E_{t+1} = E(x_{t+1}, D_{t+1})$ with

$$x_{t+1} = x_t + \frac{1}{n+1} \frac{D_t a_j}{\sqrt{a_j^T D_t a_j}},$$
$$D_{t+1} = \frac{n^2}{n^2-1}\left(D_t - \frac{2}{n+1}\frac{D_t a_j a_j^T D_t}{a_j^T D_t a_j}\right).$$

The ellipsoid method is a so-called *interior point method*, because it traverses the interior of the feasible set rather than following its boundary.

7.1 Proof of Correctness

Observe that E_0 is a ball centered at the origin. Given Theorem 6.2, and assuming that (i) $P \subseteq E_0$ and $\mathrm{Vol}(E_0) < V$ and that (ii) P is either empty or $\mathrm{Vol}(P) > \nu$, correctness of the ellipsoid method is easy to see: it either finds a point in P, thereby proving that P is non-empty, or an ellipsoid $E_{t^*} \supseteq P$ with $\mathrm{Vol}(E_{t^*}) < e^{-t^*/2(n+1)}\mathrm{Vol}(E_0) < (\nu/V)\mathrm{Vol}(E_0) < \nu$, in which case P must be empty.

We now show that the above assumptions hold, starting with the inclusion of P in E_0 and the volume of E_0. We use the following lemma.

Lemma 7.1. *Let $A \in \mathbb{Z}^{m\times n}$ and $b \in \mathbb{R}^m$. Let U be the largest absolute value among the entries of A and b. Then every extreme point x of the polytope $P = \{x' \in \mathbb{R}^n : Ax' \geqslant b\}$ satisfies $-(nU)^n \leqslant x_i \leqslant (nU)^n$ for all $i = 1, \dots, n$.*

Proof. Any extreme point x can be written as $x = \hat{A}^{-1}\hat{b}$ for some invertible submatrix $\hat{A} \in \mathbb{Z}^{n\times n}$ of A and subvector $\hat{b} \in \mathbb{R}^n$ of b, corresponding to n linearly independent constraints that are active at x. By Cramer's rule,

$$x_i = \frac{\det \hat{A}^i}{\det \hat{A}},$$

where $\hat{A}^i$ is the matrix obtained by replacing the ith column of $\hat{A}$ by $\hat{b}$. Then, for $i = 1, \dots, n$,

$$\left|\det \hat{A}^i\right| = \left|\sum_{\sigma\in S_n} (-1)^{|\sigma|} \prod_{j=1}^{n} \hat{a}^i_{j,\sigma(j)}\right| \leqslant \sum_{\sigma\in S_n} \prod_{i=1}^{n} |\hat{a}^i_{j,\sigma(j)}| \leqslant n!U^n \leqslant (nU)^n,$$

where $|\sigma|$ is the number of inversions of permutation $\sigma \in S_n$, i.e., the number of pairs i, j such that $i < j$ and $\sigma(i) > \sigma(j)$. Moreover, $\det(\hat{A}) \neq 0$ since $\hat{A}$ is invertible, and $|\det(\hat{A})| \geqslant 1$ since all entries of A are integers. Therefore, $|x_i| \leqslant (nU)^n$ for all $i = 1, \dots, n$. □

If P is bounded, it is therefore contained in a cube with side length $2(nU)^n$. The ball E_0 contains this cube and is itself contained in a cube of volume $V = (2n)^n(nU)^{n^2}$, and thus $P \subseteq E_0$ and $\mathrm{Vol}(E_0) \leqslant V$.

We now turn to the lower bound on the volume of P in the case when it is non-empty.

Lemma 7.2. *Consider a full-dimensional and bounded polytope* $P = \{x \in \mathbb{R}^n : Ax \geqslant b\}$, *where* $A \in \mathbb{Z}^{m\times n}$ *and* $b \in \mathbb{Z}^m$ *and all entries have absolute value at most* U. *Then* $\mathrm{Vol}(P) > n^{-n}(nU)^{-n^2(n+1)}$.

Proof sketch. If P is full-dimensional and bounded and has at least one extreme point, it has $n+1$ extreme points $v^0, \dots, v^n$ that do not lie on a common hyperplane. Let

$$Q = \left\{ x \in \mathbb{R}^n : x = \sum_{k=0}^{n} \lambda_k v^k, \sum_{k=0}^{n} \lambda_k = 1, \lambda_k \geqslant 0 \right\}.$$

Clearly, $Q \subseteq P$ and thus $\mathrm{Vol}(Q) \leqslant \mathrm{Vol}(P)$. It can now be shown that

$$\mathrm{Vol}(Q) = \frac{1}{n!}\left|\det\begin{pmatrix} 1 & \cdots & 1 \\ v^0 & \cdots & v^n \end{pmatrix}\right|.$$

The ith coordinate of v^k is a rational number p^k_i/q^k_i, and by the same argument as in the proof of Lemma 7.1, $|q^k_i| \leqslant (nU)^n$ and $|p^k_i| \geqslant 1$. Therefore,

$$\begin{aligned}\mathrm{Vol}(P) \geqslant \mathrm{Vol}(Q) &\geqslant \frac{1}{n!}\left|\frac{1}{\prod_{i=1}^{n}\prod_{k=0}^{n} q^k_i}\right| \\ &> \frac{1}{n^n}\frac{1}{\prod_{i=1}^{n}\prod_{k=0}^{n}(nU)^n} = n^{-n}(nU)^{-n^2(n+1)}.\end{aligned}$$

□

So far we have assumed that the polytope P is bounded and full-dimensional. We finally lift these assumptions. By Lemma 7.1, all extreme points of P lie in the set $P_B = \{x \in P : |x_i| \leqslant (nU)^n \text{ for all } i = 1, \ldots, n\}$. Moreover, P is non-empty if and only if it has an extreme point. We can therefore test for non-emptiness of P_B instead of P, and P_B is a bounded polytope.

For a polytope P that is not full-dimensional, it is not the case that $\text{Vol}(P) < \nu$ implies $P = \emptyset$, and the ellipsoid method can fail. The following result shows, however, that we can slightly perturb P to obtain a polytope that is either empty or full-dimensional.

LEMMA 7.3. *Let* $P = \{x \in \mathbb{R}^n : Ax \geqslant b\}$*, where* $A \in \mathbb{Z}^{m \times n}$ *and* $b \in \mathbb{Z}^m$ *and all entries have absolute value at most* U*. Let*

$$P_\epsilon = \{\, x \in \mathbb{R}^n : Ax \geqslant b - \epsilon e\}$$

where

$$\epsilon = \frac{1}{2(n+1)} \left((n+1)U\right)^{-(n+1)}$$

and $e^T = (1, \ldots, 1)$*. Then,* $P_\epsilon = \emptyset$ *if and only if* $P = \emptyset$*, and either* $P_\epsilon = \emptyset$ *or* $\text{Vol}(P) > 0$*.*

Proof. We first show that emptiness of P implies emptiness of P_ϵ. If P is empty, then the linear program $\min\{0^T x : Ax \geqslant b\}$ is infeasible and its dual $\max\{\lambda^T b : \lambda^T A = 0^T, \lambda \geqslant 0\}$ is unbounded. There thus has to exist a basic feasible solution λ to the $n+1$ constraints $\lambda^T A = 0^T$, $\lambda^T b = 1$, and $\lambda \geqslant 0$, and, by Lemma 7.1, $\lambda_i \leqslant ((n+1)U)^{n+1}$ for all i. Since λ is a BFS, at most $n+1$ of its components are non-zero, and therefore $\sum_i^m \lambda_i \leqslant (n+1)\left((n+1)U\right)^{n+1}$ and $\lambda^T(b - \epsilon e) = 1 - \epsilon \sum_{i=1}^m \lambda_i \geqslant \frac{1}{2} > 0$. This means that the dual remains unbounded, and the primal infeasible, if we replace b by $b - \epsilon e$, and thus $P_\epsilon = \emptyset$.

It remains to be shown that P_ϵ is full-dimensional if P is non-empty. For this, consider $x \in P$ and let

$$Y = \left\{ y \in \mathbb{R}^n : x_i - \frac{\epsilon}{nU} \leqslant y_i \leqslant x_i + \frac{\epsilon}{nU} \text{ for all } i = 1, \ldots, n \right\}.$$

It is easy to see that Y has volume $(2\epsilon/(nU))^n > 0$ and that $Y \subseteq P_\epsilon$. Thus P_ϵ must be full-dimensional. □

The general case of polytopes P that potentially are unbounded and not full-dimensional can thus be handled by computing the bounded polytope P_B, perturbing it, and then running the ellipsoid method on the resulting polytope.

7.2 The Complexity of the Ellipsoid Method

For a bounded and full-dimensional polytope P given by a matrix A and vector b with integer entries bounded by U, the ellipsoid method decides whether P is empty or not in

$O(n \log(V/\nu)) = O(n^4 \log(nU))$ iterations. It can further be shown that $O(n^6 \log(nU))$ iterations suffice even when P might be unbounded or not full-dimensional.

For the ellipsoid method to have a polynomial running time, however, the number of operations in each iteration also has to be bounded by a polynomial function of n and $\log U$. A potential problem is that the computation of the new ellipsoid involves taking a square root. This means that in general calculations cannot be done exactly, and intermediate results have to be stored with sufficiently many bits to ensure that errors don't accumulate. It turns out that the algorithm can be made to work, with the same asymptotic number of iterations as above, when only $O(n^3 \log U)$ bits are used for each intermediate value. The proof of this result is very technical.

The ellipsoid method has high theoretical significance, because it provided the first polynomial-time algorithm for linear programming and can also be applied to larger classes of convex optimization problems. In practice, however, both the simplex method and a different interior point method, *Karmarkar's algorithm*, tend to be much faster. It turns out that the latter also has a better worst-case performance than the ellipsoid method.