Lecture 8

The Ellipsoid Algorithm*

Recall from Lecture 6 that the duality theorem places the linear programming feasibility and solvability problems in $\mathsf{NP} \cap \text{co-}\mathsf{NP}$. In this class, we will see the ellipsoid algorithm, which was the first polynomial time algorithm for the LP feasibility problem; this places the LP solvability problem in P. The "Ellipsoid algorithm" was introduced by N. Shor in early 1970's as an iterative method for general convex optimization, and later applied by Khachiyan (1979) for linear programs.

8.1 Ellipsoids

In this section, we define an ellipsoid and note some of its useful properties for future use.

Definition 8.1. A (closed) *ball* $B(c, r)$ (in $\mathbb{R}^n$) centered at $c \in \mathbb{R}^n$ with radius r is the set

$$B(c, r) := \{x \in \mathbb{R}^n \,:\, x^T x \leq r^2\}.$$

The set $B(0, 1)$ is called the *unit ball*.

An ellipsoid is just an affine transformation of a ball.

Definition 8.2. An *ellipsoid* E centered at the origin is the image $L(B(0, 1))$ of the unit ball under an *invertible* linear transformation $L : \mathbb{R}^n \to \mathbb{R}^n$. An ellipsoid centered at a general point $c \in \mathbb{R}^n$ is just the translate $c + E$ of some ellipsoid E centered at 0.

We can write the above definition in a more explicit way as follows:

$$\begin{aligned}
L(B(0,1)) &= \{Lx \,:\, x \in B(0,1)\} \\
&= \{y \,:\, L^{-1}y \in B(0,1)\} \\
&= \{y \,:\, (L^{-1}y)^T L^{-1} y \leq 1\} \\
&= \{y \,:\, y^T (LL^T)^{-1} y \leq 1\} \\
&= \{y \,:\, y^T Q^{-1} y \leq 1\}
\end{aligned}$$

*Lecturer: Ryan O'Donnell. Scribe: Srivatsan Narayanan.

where $Q = LL^T$.

What can we say about the matrix $Q = LL^T$? From basic linear algebra, we recall from basic linear algebra that it is *positive definite*. We record this as fact below.

Fact 8.3. *For a symmetric matrix $Q \in \mathbb{R}^{n\times n}$, the following conditions are equivalent:*

1. *$Q = LL^T$ for some $L \in \mathbb{R}^{n\times n}$.*

2. *All the n eigenvalues of Q are nonnegative.*[1]

We say that Q is positive semi-definite *if any of the above conditions hold.*

We will add many more equivalent characterizations to this list later in the course. We will not prove the whole claim in this class; instead we verify just one of the directions to give a flavor.

Proof. (Of 1. $\implies$ 2.) We are given that $Q = LL^T$ for some $L \in \mathbb{R}^{n\times n}$. Suppose λ is an eigenvalue of Q with eigenvector $x \neq 0$; that is, $Qx = \lambda x$. Then

$$\lambda\|x\|^2 = \lambda x^T x = x^T(\lambda x) = x^T(Qx) = x^T LL^T x = (L^T x)^T(L^T x) = \|L^T x\|^2 \geq 0,$$

which shows that λ is real and nonnegative. $\square$

Fact 8.4. *For a symmetric matrix $Q \in \mathbb{R}^{n\times n}$, the following conditions are equivalent:*

1. *$Q = LL^T$ for some* nonsingular[2] *matrix L.*

2. *All the n eigenvalues of Q are strictly positive.*

We say that Q is positive definite *if any of the above conditions hold.*

From the above claims, it is clear that an ellipsoid can equivalently be represented in terms of a positive definite matrix Q.

Definition 8.5. If $Q \in \mathbb{R}^{n\times n}$ is a positive definite matrix, then the ellipsoid associated with Q and centered at $c \in \mathbb{R}^n$ is

$$E(c, Q) := \{c + y \,:\, y^T Q^{-1} y \leq 1\} = \{y \,:\, (y - c)^T Q^{-1}(y - c) \leq 1\}.$$

Remark 8.6. The standard ball $B(0, r)$ is the ellipsoid $E(0, r^2 I)$. More generally, an "axial ellipsoid" with semiaxes $r_1, \ldots r_n$ is given by the ellipsoid $E\left(0, \begin{pmatrix} r_1^2 & & & \Large 0 \\ & r_2^2 & & \\ & & r_3^2 & \\ \Large 0 & & & \ddots \end{pmatrix}\right)$.

[1] Recall that all eigenvalues of a real symmetric matrix are real.

[2] Nonsingular matrices are also known as *invertible*.

The final ingredient is the following fact about the volume of an ellipsoid. Denote by $\mathrm{vol}(A)$ the volume of a set $A \subseteq \mathbb{R}^n$.

Fact 8.7. *If $A \subseteq \mathbb{R}^n$ and L is a linear transformation, then*

$$\mathrm{vol}(L(A)) = |\det L| \cdot \mathrm{vol}(A).$$

In particular, the volume of an ellipsoid $E(c, Q)$ is given by

$$\mathrm{vol}(E(c,Q)) = |\det L| \cdot \mathrm{vol}(B(0,1)) = \sqrt{\det Q} \cdot \mathrm{vol}(B(0,1)).$$

Thus we have related the volume of any ellipsoid to the volume of the unit ball in n dimension. Fortunately, the exact value of the constant of the proportionality, the volume of the unit n-ball, is irrelevant to us.[3]

8.2 The Ellipsoid Algorithm

The ellipsoid algorithm takes as input a convex set, and returns a point from the set provided it is nonempty. (If the set is empty, then we return "empty".) It is clear that this algorithm is useful for testing LP feasibility. Further, since the LP solvability problem reduces to the LP feasibility problem in polynomial time, this algorithm can also be used to solve linear programs as well.

Formally, the ellipsoid algorithm tests if a given convex set $K \subseteq \mathbb{R}^n$ is empty.

8.2.1 Requirements

Apart from the input set K, we assume that we are provided two additional parameters:

1. A number $R \in \mathbb{Q}$ $(R > 0)$ such that $K \subseteq B(0, R)$.
2. A rational $r > 0$ such that either $K = \emptyset$ or $K \supseteq B(c, r)$ for some point c. (Think of this requirement as basically stating that the feasible solution is not completely contained in some affine hyperplane.)

Now, how do we satisfy these two requirements while using the ellipsoid algorithm for solving LPs?

1. The first condition is easily handled. Given a linear program

   $$K = \{Ax = b \,:\, x \geq 0\},$$

   we can add in constraints of the form $-2^L \leq x \leq 2^L$ for each $i \in [n]$ for some $L = \mathrm{poly}(\langle A \rangle, \langle b \rangle, n)$ (without affecting the feasibility of the LP). Since this region is fully contained inside the ball of radius $\sqrt{n}2^L$, we can provide $R = n2^L$ as the parameter.

[3]For the curious: the volume of the unit n-ball has the exact expression $\frac{\pi^{n/2}}{\Gamma(\frac{n}{2}+1)}$, where $\Gamma(\cdot)$ is the "Gamma function". Asymptotically, this volume is $\tilde{\Theta}\left(\frac{C^n}{n^{n/2}}\right)$ for some absolute constant $C > 0$ (hiding $\mathrm{poly}(n)$ factors).

2. The second requirement is slightly tricky since the feasible region of the given LP is contained in the hyperplane $Ax = b$, and hence, as such, there does *not* exist $r > 0$ satisfying the requirements. In this case, suppose the convex set K is empty. Then the hyperplane $Ax = b$ is separated from the positive "orthant" (i.e., the region $x \geq 0$) by a finite width $r > 0$ (that is expressible in polynomially many bits). Then the idea is to draw a small tube around $Ax = b$, so that the system of equations becomes full dimensional.

   Slightly more precisely, there exists some $\epsilon > 0$ such that the new convex set $K' = \{-\epsilon \leq Ax - b \leq \epsilon : x \geq 0\}$ is empty provided K is empty. (See Figure 8.1 for an illustration.) On the other hand, if K is nonempty and $c \in K$, then the set K' contains a ball $B(c, r)$ for some finite $r > 0$. (See Figures 8.2.) Finally, we can show that the numbers ϵ and r are expressible in polynomially many bits. Thus we may provide r as the parameter.

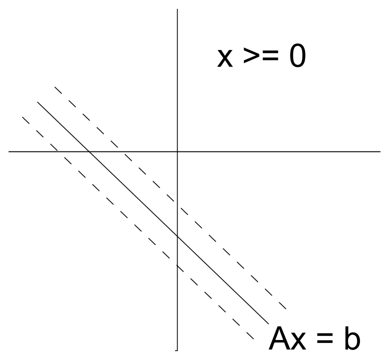

Figure 8.1: Handling the second requirement in the infeasible case. Note that there is a finite width $\epsilon > 0$ such that the region $-\epsilon \leq Ax - b \leq \epsilon$ does not intersect the positive orthant $x \geq 0$.

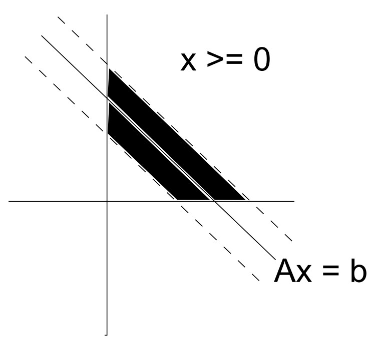

Figure 8.2: Handling the second requirement in the feasible case. Relaxing the LP to $-\epsilon \leq Ax - b \leq \epsilon$ makes the feasible set (the shaded region) full dimensional.

8.2.2 The algorithm

We now describe the ellipsoid algorithm. From now on, it is convenient to assume that the input LP has the form $Ax \leq b$ with $x \geq 0$.

If $n = 1$, then solve the LP directly and terminate; so without loss of generality, assume $n > 1$. The algorithm maintains an ellipsoid E that completely contains K (if K is nonempty). We initialize the algorithm with the ellipsoid $E(0, R^2 I)$ that was promised to satisfy this requirement.

1. Check if the center c is feasible (i.e., if $c \in K$). If so, then we are done.

2. Else, get a "separating hyperplane" through c; i.e., a hyperplane $a^T x = a^T c$ through c such that the set K is completely contained in the half-ellipsoid formed by the intersection of $a^T x \geq a^T c$ with the current ellipsoid. (To implement this step given an LP, we take a violated constraint $a^T x \leq \gamma$; then $a^T x = a^T c$ is a separating hyperplane through c.)

3. Take a smallest volume ellipsoid containing the half-ellipsoid which may contain K. Goto 1.

4. After $N = \text{poly}(n, \langle R \rangle, \langle r \rangle)$ iterations, stop and say $K = \emptyset$.

8.2.3 Analysis of the algorithm

Clearly, Step 1 of the algorithm is correct. Further, at every stage of the algorithm, we maintain the invariant that the convex set K is completely contained inside the current ellipsoid. Thus it only remains to show that if we terminate the algorithm after N steps and find no feasible point, then the set K is indeed empty.

Denote the ellipsoid at the end of iteration k by $E_k = E(c_k, Q_k)$ (the starting ellipsoid is $E_0 = E(0, R^2 I)$). We use the volume of ellipsoids E_k to track the number of iterations of the algorithm. The main claim is the following:

Theorem 8.8 (Volume reduction)**.** *For $k \geq 0$, we have*[4]*:*

$$\frac{\text{Vol}(E(c_{k+1}, Q_{k+1}))}{\text{Vol}(E(c_k, Q_k))} \leq e^{-\frac{1}{2(n+1)}}.$$

We prove this theorem in the next subsection.

Corollary 8.9. *1. After $2(n+1)$ steps, the volume goes down by a factor of $\frac{1}{e}$.*

2. Suppose $N > 2n(n+1)\ln(\frac{R}{r})$. Then after N iterations, the volume becomes smaller than $\text{Vol}(B(0, r))$.

[4] In class, we claimed the slightly stronger upper bound of $e^{-\frac{1}{2n}}$ with no proof. This only causes a constant factor increase in the number of iterations.

Proof. The first part is obvious. To see the second part, the volume of the ellipsoid at the end of N iterations is

$$\mathrm{vol}(E_N) = \mathrm{vol}(E_0)\exp\left(-\frac{N}{2(n+1)}\right) < \mathrm{vol}(B(0,R))\exp\left(-n\ln\left(\frac{R}{r}\right)\right)$$

Rearranging,

$$\mathrm{vol}(E_N) < R^n \mathrm{vol}(B(0,1)) \times \frac{r^n}{R^n} = r^n \mathrm{vol}(B(0,1)) = \mathrm{vol}(B(0,r)) = \mathrm{vol}(B(c,r)).$$

since all balls of the same radius have the same volume. It follows that $K \subseteq E_N$ has a volume strictly less than $\mathrm{vol}(B(c,r))$. But this contradicts the second requirement of the algorithm, unless $K = \emptyset$. $\square$

Thus if we did not find a feasible solution after N iterations, then we can safely terminate. Since N is polynomial in the length of the input, we have shown that this algorithm terminates in polynomial time.

But we are not done yet! Of course, we still need to prove Theorem 8.8. Also, to complete the description of the algorithm, we need to write down the smallest volume ellipsoid containing the half-ellipsoid that may have K. We do both of these in the next subsection.

8.2.4 The description of the half-ellipsoid: a simple case

We first deal with the simple case, where $E_0 = B(0,1)$ and the separating hyperplane is $a_0 = (-1,0,0,\dots,0)$. Our goal is to find an ellipsoid E_1 that contains the region $E_0 \cap \{x : x_1 \geq 0\}$.

Lemma 8.10. *Define $c_1 = (\frac{1}{n+1}, 0, 0, \dots, 0)$, and*

$$Q_1 = \frac{n^2}{n^2-1}\begin{pmatrix} 1-\frac{2}{n+1} & & & & \\ & 1 & & & \mathit{0} \\ & & 1 & & \\ & \mathit{0} & & \ddots & \\ & & & & 1 \end{pmatrix}$$

Then $E_1 = E(c_1, Q_1)$ is the minimum volume ellipsoid containing the half-ball. Moreover,

$$\frac{\mathrm{vol}(E_1)}{\mathrm{vol}(E_0)} \leq e^{-\frac{1}{2(n+1)}}.$$

Proof. In this notes, we only prove that the ellipsoid E_1 contains the desired half-ball and prove the bound on its volume. Although it is true that E_1 is the ellipsoid of *minimum volume*, we do not show that here. Note that this does not affect our algorithm or our analysis in any way.

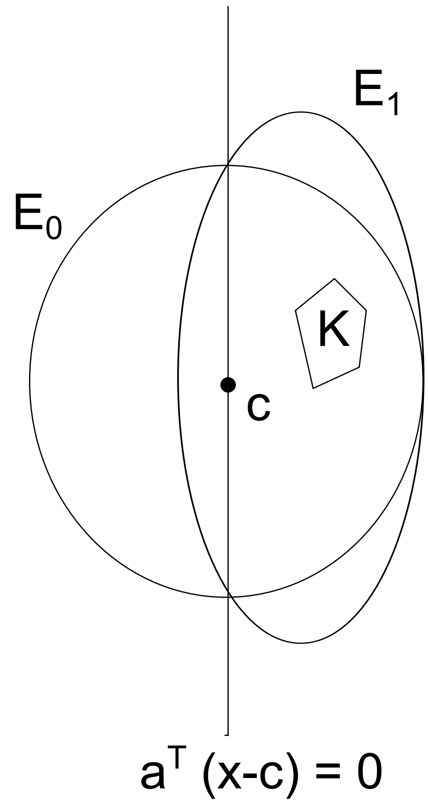

Figure 8.3: Ellipsoid E_1 covering the half-ellipsoid bounded by E_0 and the separating hyperplane $a^T(x-c)=0$.

Take any x in the half-ball; i.e., $\|x\| \leq 1$ and $x_1 \geq 0$. Suppose $x = (x_1, \tilde{x})$ where $\tilde{x} = (x_2, x_3, \ldots, x_n)$. It is easy to verify that

$$Q_1^{-1} = \frac{n^2-1}{n^2} \begin{pmatrix} \frac{n+1}{n-1} & & & \\ & 1 & & \text{\Large 0} \\ \text{\Large 0} & & 1 & \\ & & & \ddots \end{pmatrix}.$$

Consider

$$\begin{aligned}
x^TQ^{-1}x &= \frac{n^2-1}{n^2}\left(x_1 - \frac{1}{n+1}, \tilde{x}\right)^T \begin{bmatrix} \frac{n+1}{n-1} & & & \\ & 1 & & \text{\Large 0} \\ \text{\Large 0} & & 1 & \\ & & & \ddots \end{bmatrix} \left(x_1 - \frac{1}{n+1}, \tilde{x}\right) \\
&= \frac{n^2-1}{n^2}\frac{n+1}{n-1}\left(x_1 - \frac{1}{n+1}\right)^2 + \frac{(n^2-1)}{n^2}\|\tilde{x}\|^2 \\
&= \frac{1}{n^2}((n+1)x_1-1)^2 + \frac{(n^2-1)}{n^2}\|\tilde{x}\|^2 \\
&\overset{(a)}{\le} \frac{1}{n^2}((n+1)x_1-1)^2 + \frac{(n^2-1)}{n^2}(1-x_1^2) \\
&= \frac{(n+1)^2}{n^2}x_1^2 - 2\frac{n+1}{n^2}x_1 + \frac{1}{n^2} + \frac{n^2-1}{n^2} - \frac{n^2-1}{n^2}x_1^2 \\
&= \frac{2(n+1)}{n^2}(x_1^2 - x_1) + 1 \\
&\overset{(b)}{\le} 1,
\end{aligned}$$

where (a) follows from the fact that $x_1^2 + \|\tilde{x}\|^2 = \|x\|^2 \le 1$, and (b) follows from the inequality $0 \le x_1 \le 1$. Therefore, the point x is inside the ellipsoid E_1.

The ratio of the volumes of the ellipsoids E_1 and E_0 is given by

$$\frac{\mathrm{vol}(E_1)}{\mathrm{vol}(E_0)} = \sqrt{\det Q_1} = \sqrt{\left(\frac{n^2}{n^2-1}\right)^n \left(\frac{n-1}{n+1}\right)} = \sqrt{\left(\frac{n^2}{n^2-1}\right)^{n-1} \left(\frac{n}{n+1}\right)^2},$$

after some rearrangement. Using the inequality $1 + z \le e^z$ valid for all real z, we get

$$\frac{\mathrm{vol}(E_1)}{\mathrm{vol}(E_0)} \le \exp\left(\frac{n-1}{2}\cdot\frac{1}{n^2-1} - \frac{1}{n+1}\right) = \exp\left(-\frac{1}{2(n+1)}\right).$$

□

8.2.5 The description of the ellipsoid: the general case

Suppose we have an ellipsoid $E_k = E(c_k, Q_k)$, and we have a separating hyperplane $a_k^Tx = a_k^Tc_k$ through the center c_k. Our goal is to compute the minimum volume ellipsoid E_{k+1} that contains the half-ellipsoid bounded by E_k and $a_k^Tx \ge a_k^Tc_k$.

By the definition of an ellipsoid, there exists some invertible affine transformation L^{-1} that takes E_k to $B(0,1)$ and a_k to $a = (-1, 0, 0, \dots, 0)$. Thus we are back to the simple case in Subsection 8.2.4. Let E' be the ellipsoid just analyzed and take $E_{k+1} = LE'$. This clearly contains the half-ellipsoid. Further,

$$\frac{\mathrm{vol}(E_{k+1})}{\mathrm{vol}(E_k)} = \frac{\mathrm{vol}(L(E'))}{\mathrm{vol}(L(B(0,1)))} = \frac{|\det L|\cdot\mathrm{vol}(E')}{|\det L|\cdot\mathrm{vol}(B(0,1))} = \frac{\mathrm{vol}(E')}{\mathrm{vol}(B(0,1))} \le e^{-\frac{1}{2(n+1)}},$$

by the previous analysis.

For implementation purposes, it is more desirable to describe the ellipsoid E_{k+1} more explicitly by computing the invertible transformation L. We will just state the final result without proof.

Claim 8.11. *The ellipsoid $E_{k+1} = (c_{k+1}, Q_{k+1})$ is given by: $c_{k+1} = c_k - \frac{1}{n+1}h$ and*

$$Q_{k+1} = \frac{n^2}{n^2-1}\left(Q_k - \frac{2}{n+1}hh^T\right).$$

where $h = \sqrt{a_k^T Q_k a_k}$.

Proof. Omitted. □

One final remark is in order about the correctness of our algorithm and its analysis. Note that the description of the half-ellipsoid relies on computing square roots. This makes the preceding analysis valid only for an idealized implementation assuming exact arithmetic. It is possible to handle this precision issues by keeping a good enough approximation of the real quantities using rational numbers. However, the full proof then becomes substantially more complicated. We will deal with this issue (partly) in the next lecture.