

18.409: Problem Set 2

Jonathan Kelner

Overview of Problems 1 and 2

In class, we proved Fritz John's theorem. This showed that any origin symmetric convex body K contains an ellipse E such that

$$E \subseteq K \subseteq \sqrt{n}E, \tag{1}$$

and that any (not necessarily origin symmetric) convex body K contains an ellipse E such that

$$E \subseteq K \subseteq nE. \tag{2}$$

Our proof showed that the maximum volume ellipsoid contained in K meets the conditions described in Equations (1) and (2) and precisely characterized this ellipsoid.

In class, we proved the existence of an ellipsoid satisfying the above conditions, but we did not address how one might compute it algorithmically. In problems 1 and 2, we will investigate how to do this. This will result in algorithms that give ellipsoids meeting slightly weaker constraints, but we will have an efficient algorithm to compute them.

In problem 1, we will study a general (not necessarily symmetric) convex body K , given by a separation oracle (with the usual guarantees about K containing and being contained in balls of certain radii). With just oracle access to K , we will only be able to find an algorithm for a weaker result: we will construct an ellipsoid E such that

$$E \subseteq K \subseteq O(n^{3/2})E. \tag{3}$$

In problem 2, we will allow ourselves a more explicit description of K as the set of points obeying a set of linear inequalities, and we will study the origin symmetric case. (The methods sketched in problems 1 and 2 can be modified to work in both the symmetric and asymmetric cases. I have chosen the versions that require the least algebraic mess.) Given this explicit description, we will produce an ellipsoid E such that

$$E \subseteq K \subseteq \sqrt{n+1}E, \tag{4}$$

thereby coming quite close to the best possible bound.

Problem 2 will thus provide a proof of a slightly weakened version of Fritz John's theorem. Note, however, that it just provides the existence of an ellipsoid meeting Equation (4)—it does not precisely characterize it like the proof in class did. This is sufficient for every algorithmic application of the theorem that I'm aware of.

Problem 1: Approximate John Ellipsoids with a Separation Oracle

Suppose that you are given a separation oracle for a convex body $K \subseteq \mathbb{R}^n$, along with the promise that K contains a small ball around the origin of some radius r and is contained in a ball of radius R . We will show how to compute an ellipsoid E satisfying Equation (3) in time that is polynomial in n and $\log(R/r)$.

To do this, we will actually find a large *simplex* contained in K , which will be a bit easier. You may assume throughout the problem that you have a black box that maximizes linear functionals over K .

(a) It will be more convenient to embed K in \mathbb{R}^{n+1} on the hyperplane $H = \sum_i x_i = 1$. We shall assume that we have done this so that K contains the point $p = (1/(n+1), \dots, 1/(n+1))$. Because we know that K contains a ball of radius r in H , we can change coordinates so that it contains the unit simplex $\Delta = \text{conv}(e_1, \dots, e_{n+1})$, where the e_i are the unit basis vectors in \mathbb{R}^{n+1} .

After changing to this coordinate system, our first goal will be to find a point $p \in K$ that has $|p \cdot e_i| \geq 1 + \epsilon$ for some i and some fixed $\epsilon > 0$, if such a point exists. Describe an algorithm for doing this.

(b) Suppose that you successfully found a point p satisfying the conditions in part (a). Construct a new simplex $\Delta' = \text{conv}(e_1, \dots, e_{i-1}, p, e_{i+1}, \dots, e_{n+1})$. Our algorithm will change coordinates using a linear transformation τ such that $\tau(\Delta') = \Delta$. It will then try to find another p meeting the conditions of part (a) and repeat. Show that this procedure terminates in a number of iterations that is polynomial in n and $\log(R/r)$ with a linear map τ such that τK contains the unit simplex, and in which there is no point in τK meeting the conditions of part (a).

(c) Let $S_\epsilon = \{x \in H \mid |x_i| \leq 1 + \epsilon \text{ for all } i\}$. Part (b) has thus given us a linear map τ such that $\Delta \subseteq \tau K \subseteq S_\epsilon$.

Show that $B(p, \rho_1) \subseteq \Delta$ and $S_\epsilon \subseteq B(p, \rho_2)$, where

- $B(p, \rho_i)$ is the ball around p of a radius ρ_i ,
- $\rho_1 \geq \Theta(1/n)$, and
- $\rho_2 \leq \Theta(\sqrt{n})$.

Note that this gives us an algorithm for finding the ellipsoid guaranteed in Equation (3), along with a proof that such an ellipsoid exists.

Problem 2: A Better Approximation Algorithm When K is Given Explicitly

In this problem, we are going to assume that K is origin symmetric and given to us explicitly as $K = \{x \in \mathbb{R}^n \mid Ax \leq \mathbf{1}\}$, where A is an $m \times n$ matrix $\mathbf{1}$ is the all-ones vector. We will find an ellipsoid meeting the conditions in Equation (4). Note that given just a separation oracle, E obeyed an inequality that was weaker than the optimal one by a factor of \sqrt{n} , whereas here we are getting a result that is asymptotically tight. We will do this by actually working with ellipsoids, not simplices.

In problem 1, we maintained a simplex inside of K at every step. We then found a point that was “too big” with respect to this simplex, and we used it to enlarge the simplex. In this problem, we will proceed in the opposite direction. We will start with an ellipsoid that *contains* K . We’ll then iteratively find proof that it’s too big and find a way to shrink it until we converge to our answer.

Given a well-bounded separation oracle, checking whether a body is contained in a ball and contains a ball are essentially the same, as an oracle for K lets you construct one for its polar K° . However, the explicit description of K breaks the symmetry between the primal and polar. In parts (a) and (b), we’ll explore this asymmetry, which is what allows us to obtain a better approximation in this problem than in problem 1.

(a) Suppose that I tell you that K is contained in the unit ball B_2^n . You cannot algorithmically check this. Instead, provide an algorithm that will do this following:

- If $K \subseteq B_2^n$, it will accept.
- If $K \not\subseteq \sqrt{n}B_2^n$, it will reject.
- Otherwise, in the case where K contains point of norm greater than 1 but does not contain any points of norm greater than \sqrt{n} , the algorithm can do whatever it wants (either accept or reject).

Could you do this, without extensive modification of your algorithm, using just the oracle given in problem 1? (What I'm looking for here is the answer to whether the algorithm you described could easily be made to work in the oracular setting. I am not asking for an elaborate modification if the answer is "yes," or a proof if the answer is "no.")

(b) Now, give an algorithm that determines whether the polytope K contains B_2^n . (You don't get the \sqrt{n} room for error on this one.) Could you do this using just the oracle from problem 1? (Again, I just want a yes/no answer with a short explanation, not anything complicated.)

(c) Let B be the unit ball in \mathbb{R}^2 ,

$$B = \{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 \leq 1\}.$$

Let $E(a, b)$ be the axis-aligned ellipse

$$E(a, b) = \left\{ (x, y) \in \mathbb{R}^2 \mid \frac{x^2}{a} + \frac{y^2}{b} \leq 1 \right\}.$$

Show that for $\gamma > 0$ and $n \geq 2$,

$$B \cap \{(x, y) \in \mathbb{R}^2 \mid |x| \leq \gamma\} \subseteq E(a_\gamma, b_\gamma),$$

where:

- If $\gamma > 1/\sqrt{n}$, $a_\gamma = b_\gamma = 1$.
- If $0 < \gamma \leq 1/\sqrt{n}$,

$$a_\gamma = n\gamma^2$$

and

$$b_\gamma = \frac{n(1 - \gamma^2)}{n - 1}.$$

Note that you are only asked to consider 2-d ellipses in this part— n is just a number, not a dimension.

(d) Use part (c) to prove the following lemma:

Lemma. *Suppose that $K \subseteq \mathbb{R}^n$ is an origin symmetric polytope contained in the unit ball B_2^n . If a point $p \in (1/\sqrt{n+1})B_2^n$ is not contained in K , then K is contained in an ellipse $E = \{x \mid x^T A^{-1} x \leq 1\}$, where A is a matrix with 1 eigenvalue of size $n/(n+1)$ and $n-1$ eigenvalues of size $n^2/((n+1)(n-1))$.*

(e) Show that the ellipse E from the previous part has

$$\frac{\text{Vol}(E)}{\text{Vol}(B_2^n)} \leq e^{-1/\text{poly}(n)}.$$

Use this and the earlier parts of this question to derive a polynomial-time algorithm for finding an ellipse meeting the conditions in Equation (4).

Problem 3: Solving Convex Programs with Random Sampling

Suppose I give you a separation oracle for a convex body K . In class, we showed that this allows us to draw (approximately) uniformly random samples from K . In this problem, we will show how to use this to optimize a linear functional over K . For the duration of this problem, assume that K is contained in an axis-aligned cube of width R around the origin, and that K contains some cube (not necessarily around the origin) of width r . (This algorithm is due to Bertsimas and Vempala.)

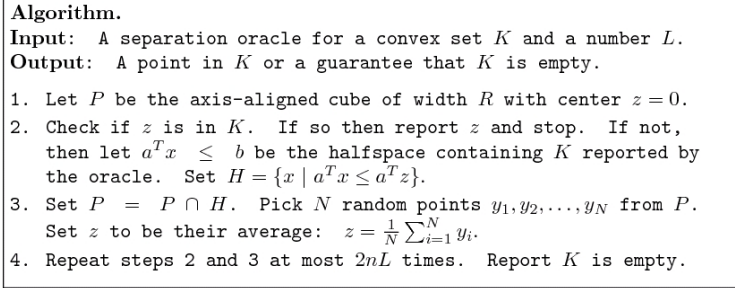


Figure 1: Algorithm for problem 3(d).

(a) In class, we proved Grunbaum’s Theorem:

Theorem. *Let K be a convex body in \mathbb{R}^n . If H is a halfspace that contains the centroid of K , then $\text{Vol}(H \cap K) \geq (1/e)\text{Vol}(K)$.*

In this problem, we will need a slight strengthening of this:

Definition. A convex body K is in *isotropic position* if a point x drawn uniformly at random from K has

- $\mathbf{E}[x] = 0$
- $\mathbf{E}[(v \cdot x)^2] = \|v\|^2$ for any vector $v \in \mathbb{R}^n$.

Theorem. *Let K be a convex body in \mathbb{R}^n in isotropic position, and let p be a point at distance t from the centroid of K . If H is a halfspace that contains p , then $\text{Vol}(H \cap K) \geq (1/e - t)\text{Vol}(K)$.*

Argue that this stronger version of Grunbaum’s Theorem is true. Either of the following will be considered a sufficient answer:

1. Discuss what would need to be changed in our proof of Grunbaum’s Theorem to prove this strengthening. You need not reproduce the whole proof or go into complete detail (unless you find it useful to do so).
2. Show how to deduce this theorem from Grunbaum’s Theorem.

Feel free in either case to weaken the bound to anything of the form $c - O(t)$, where c is any constant strictly between 0 and 1.

(b) Let K be in isotropic position, let p_1, \dots, p_m be samples drawn uniformly at random from K , and let $p = (1/m) \sum_i p_i$. Note that expectation of p is the centroid of K , which is the origin since K is in isotropic position. Prove that

$$\mathbf{E}[\|p\|^2] = \frac{n}{m}.$$

(c) Show how to reduce the question of optimizing a linear functional over K to solving at most logarithmically many instances of the problem, “Find a single point in the polytope $K \cap H$,” for a halfspace H . It is this feasibility problem that our algorithm will solve.

(d) Consider the algorithm shown in Figure 1 (copied from the original paper).

Show that this algorithm succeeds with high probability for some $N = O(n)$ and $L = O(\log(R/r))$. [Hint: Show that the volume of P is likely to shrink by a constant factor in each iteration. Note that you will need to make some argument that handles the case where K is not in isotropic position. This modification should not require anything elaborate, nor should it involve changing the algorithm or strengthening the theorems above.]

MIT OpenCourseWare
<http://ocw.mit.edu>

18.409 Topics in Theoretical Computer Science: An Algorithmist's Toolkit
Fall 2009

For information about citing these materials or our Terms of Use, visit: <http://ocw.mit.edu/terms>.