16. Smallest Enclosing Balls

Lecture on Thursday 27^{th} November, 2008 by Bernd Gärtner <gaertner@inf.ethz.ch>

16.1 Problem Statement and Basics

As usual, we let P be a set of n points, but this time in $\mathbb{R}^d$. We are interested in finding a closed ball of smallest radius that contains all the points in P, see Figure 16.1.

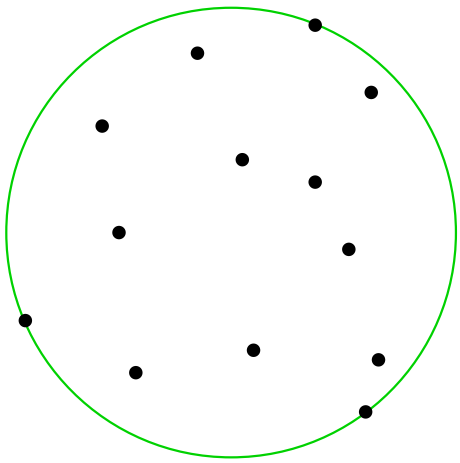

Figure 16.1: *The smallest enclosing ball of a set of points in the plane*

As an "application", imagine a village that wants to build a firehouse. The location of the firehouse should be such that the maximum travel time to any house of the village is as small as possible. If we equate travel time with Euclidean distance, the solution is to place the firehouse in the center of the smallest ball that covers all houses.

16.1.1 Existence

It is not a priori clear that a smallest ball enclosing P exists, but this follows from standard arguments in calculus. As you usually don't find this worked out in papers and textbooks, let us quickly do the argument here.

Fix P and consider the continuous function $\rho : \mathbb{R}^d \to \mathbb{R}$ defined by

$$\rho(c) = \max_{p \in P} \|p - c\|, c \in \mathbb{R}^d$$

Thus, $\rho(c)$ is the radius of the smallest ball centered at c that encloses all points of P. Let q be any point of P, and consider the closed ball

$$B = B(q, \rho(q)) := \{c \in \mathbb{R}^2 \mid \|c - q\| \leq \rho(q)\}.$$

Since B is compact, the function ρ attains its minimum over B at some point c_{opt}, and we claim that c_{opt} is the center of a smallest enclosing ball of P. For this, consider any center $c \in \mathbb{R}^2$. If $c \in B$, we have $\rho(c) \geq \rho(c_{opt})$ by optimality of c_{opt} in B, and if $c \notin B$, we get $\rho(c) \geq \|c - q\| > \rho(q) \geq \rho(c_{opt})$ since $q \in B$. In any case, we get $\rho(c) \geq \rho(c_{opt})$, so c_{opt} is indeed a best possible center.

16.1.2 Uniqueness

Can it be that there are two distinct smallest enclosing balls of P? No, and to rule this out, we use the concept of *convex combinations* of balls. Let $B = B(c, \rho)$ be a closed ball with center c and radius $\rho > 0$. We define the *characteristic function* of B as the function $f_B : \mathbb{R}^2 \to R$ given by

$$f_B(x) = \frac{\|x - c\|^2}{\rho^2}, \quad x \in \mathbb{R}^2.$$

The name characteristic function comes from the following easy

Observation 16.1 *For* $x \in \mathbb{R}^2$*, we have*

$$x \in B \quad \Leftrightarrow \quad f_B(x) \leq 1.$$

Now we are prepared for the convex combination of balls.

Lemma 16.2 *Let* $B_0 = B(c_0, \rho_0)$ *and* $B_1 = (c_1, \rho_1)$ *be two distint balls with characteristic functions* f_{B_0} *and* f_{B_1}*. For* $\lambda \in (0, 1)$*, consider the function* f_λ *defined by*

$$f_\lambda(x) = (1 - \lambda) f_{B_0}(x) + \lambda f_{B_1}(x).$$

Then the following three properties hold.

(i) f_λ *is the characteristic function of a ball* $B_\lambda = (c_\lambda, \rho_\lambda)$*.* B_λ *is called a (proper)* convex combination *of* B_0 *and* B_1*, and we simply write*

$$B_\lambda = (1 - \lambda) B_0 + \lambda B_1.$$

(ii) $B_\lambda \supseteq B_0 \cap B_1$ *and* $\partial B_\lambda \supseteq \partial B_0 \cap \partial B_1$*.*

(iii) $\rho_\lambda < \max(\rho_0, \rho_1)$*.*

A proof of this lemma requires only elementary calculations and can be found for example in the PhD thesis of Kaspar Fischer [?]. Here we will just explain what the lemma means. The family of balls $B_\lambda, \lambda \in (0, 1)$ "interpolates" between the balls B_0 and B_1: while we increase λ from 0 to 1, we continuously transform B_0 into B_1. All intermediate balls B_λ "go through" the intersection of the original ball boundaries (a sphere of dimension $d - 2$). In addition, each intermediate ball contains the intersection of the original balls. This is property (ii). Property (iii) means that all intermediate balls are smaller than the larger of B_0 and B_1. Figure 16.2 illustrates the situation.

Using this lemma, we can easily prove the following

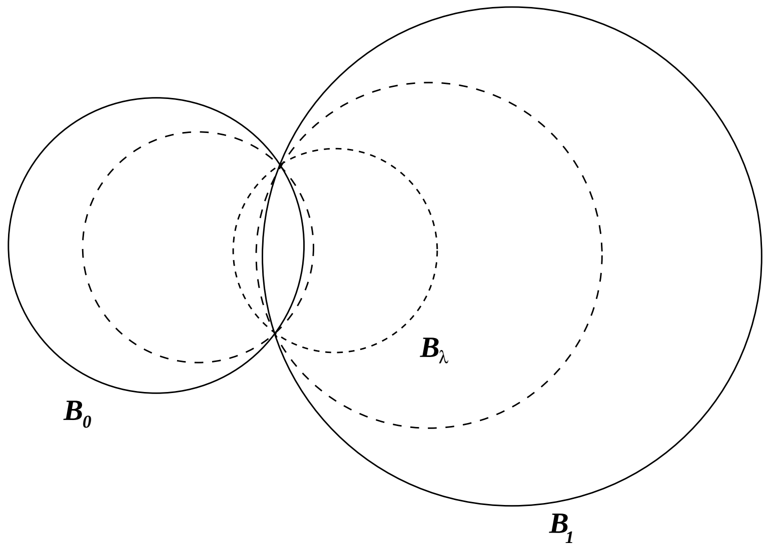

Figure 16.2: *Convex combinations* B_λ *of two balls* B_0, B_1

Theorem 16.3 *Given a finite point set* $P \subseteq \mathbb{R}^d$*, there exists a* unique *ball of smallest radius that contains* P*. We will denote this ball by* $B(P)$.

Proof. If $P = \{p\}$, the unique smallest enclosing ball is $\{p\}$. Otherwise, any smallest enclosing ball of P has positive radius ρ_{opt}. Assume there are two distinct smallest enclosing balls B_0, B_1. By Lemma 16.2, the ball

$$B_{\frac{1}{2}} = \frac{1}{2}B_0 + \frac{1}{2}B_1$$

is also an enclosing ball of P (by property (ii)), but it has smaller radius than ρ_{opt} (by property (iii), a contradiction to B_0, B_1 being smallest enclosing balls. □

16.1.3 Bases

When you look at the example of Figure 16.1, you notice that only three points are essential for the solution, namely the ones on the boundary of the smallest enclosing ball. Removing all other points from P would not change the smallest enclosing ball. Even in cases where more points are on the boundary, it is always possible to find a subset of at most three points (in the $\mathbb{R}^2$ case) with the same smallest enclosing ball. We will next prove a theorem for arbitrary dimensions that implies this; the theorem is based on *Helly's Theorem*, a classical result in convexity.

Theorem 16.4 (Helly's Theorem[?]) *Let* $C_1, \dots C_n$ *be* $n \geq d+1$ *convex subsets of* $\mathbb{R}^d$*. If any* $d+1$ *of the sets have a nonempty common intersection, then the common intersection of* all n *sets is nonempty.*

Even in $\mathbb{R}^1$, this is not entirely obvious. There it says that for any set of intervals with pairwise nonempty overlap there is one point contained in all the intervals. We will not prove Helly's Theorem here but just use it to show the following

Theorem 16.5 *Let* $P \subseteq \mathbb{R}^d$ *be a finite point set. There is a subset* $S \subseteq P, |S| \leq d+1$ *such that* $B(P) = B(S)$.

Proof. If $|P| < d+1$, we may choose $S = P$. Otherwise, let us assume for a contradiction that for all subsets S of size $d+1$, $B(P) \neq B(S)$, meaning that $B(S)$ has smaller radius than $B(P)$. [1]

Let ρ_S denote the radius of $B(S)$ and define

$$\overline{\rho} = \max_{S \subseteq P, |S|=d+1} \rho_S.$$

By assumption, $\overline{\rho} < \rho_{opt}$, the radius of the smallest enclosing ball $B(P)$ of $P = \{p_1, \ldots, p_n\}$. Now define

$$C_i = B(p_i, \overline{\rho}), \quad i = 1, \ldots, n$$

to be the ball around p_i with radius $\overline{\rho}$. We know that the common intersection of all the C_i is empty, since any point in the intersection would be a center of an enclosing ball of P with radius $\overline{\rho} < \rho_{opt}$. Moreover, the C_i are convex, so Helly's Theorem implies that there is a subset S of $d+1$ points whose C_i's also have an empty common intersection. For this set S, we therefore have no enclosing ball of radius $\overline{\rho}$ either, but this is a contradiction, since $\rho_S \leq \overline{\rho}$ proves that there is such a ball. □

The previous theorem motivates the following

Definition 16.6 *Let* $P \subseteq \mathbb{R}^d$ *be a finite point set. A* basis *of* P *is an inlusion-minimal subset* $S \subseteq P$ *such that* $B(P) = B(S)$.

It follows that any basis of P has size at most $d+1$. If the points are in general position (no $k+3$ on a common k-dimensional sphere), then P has a unique basis, and this basis is formed by the set of points on the boundary of $B(P)$.

16.1.4 The trivial algorithm

Definition 16.6 immediately implies the following (rather inefficient) algorithm for computing $B(P)$ and a basis of P, $|P| = n$: for every subset $S \subseteq P, |S| \leq d+1$, compute $B(S)$ (in fixed dimension d, this can be done in constant time). If this is done for all S in order of increasing size, we can easily identify the sets S such that $B(S) \neq B(S \setminus \{p\})$ for all $p \in S$—these are the potential bases of P. For each of these sets S, we check whether $P \subseteq B(S)$. If this is the case, we have $B(S) = B(P)$ (why?), and S is a basis of P.

[1] Because $B(P)$ is an enclosing ball of S, we know that the radius of $B(S)$ is at most the radius of $B(P)$, but they can't be equal, since otherwise $B(S)$ and $B(P)$ would both be smallest enclosing balls of S.

Assuming that d is fixed, the runtime of this algorithm is

$$O\left(n\sum_{i=0}^{d+1}\binom{n}{i}\right) = O(n^{d+2}),$$

because we need time $O(1)$ to compute a single ball $B(S)$ and time $O(n)$ to check for $P \subseteq B(S)$.

If $d = 2$ (the planar case), the trivial algorithm has runtime $O(n^4)$. In the next section, we discuss an algorithm that is substantially better than the trivial one in any dimension.

16.2 The Swiss Algorithm

The name of this algorithm comes from the democratic way in which it works. Let us describe it for the problem of locating the firehouse in a village.

Here is how it is done the Swiss way: a meeting of all n house owners is scheduled, and every house owner is asked to put a slip of paper with his/her name on it into a voting box. Then a constant number c (to be determined later) of slips is drawn at random from the voting box, and the selected house owners have the right to negotiate a location for the firehouse among them. They naturally do this in a selfish way, meaning that they agree on the center of the smallest enclosing ball D of just *their* houses as the proposed location.

The house owners that were not in the selected group now fall into two classes: those that are happy with the proposal, and those that are not. Let's say that a house owner p is happy if and only if his/her house is also covered by D. In other words, p is happy if and only if the proposal would have been the same with p as an additional member of the selected group.

Now, the essence of Swiss democracy is to negotiate until everybody is happy, so as long as there are any unhappy house owners at all, the whole process is repeated. But in order to give the unhappy house owners a higher chance of influencing the outcome of the next round, their slips in the voting box are being doubled before drawing c slips again. Thus, there are now two slips for each unhappy house owner, and one for each happy one.

After round k, a house owner that has been unhappy in i of the k rounds has therefore 2^i slips in the voting box for the next round.

The obvious question is: how many rounds does it take until all house owners are happy? So far, it is not even clear that the meeting ever ends. But Swiss democracy is efficient, and we will see that the meeting actually ends after an expected number of $O(\log n)$ rounds. We will do the analysis for general dimension d (just imagine the village and its houses to lie in $\mathbb{R}^d$).

16.2.1 A lower bound for the total number of slips

Let $S \subseteq P$ be a basis of P. Recall that this means that S is inclusion-minimal with $B(S) = B(P)$.

Observation 16.7 *As long as there are unhappy house owners at all, there is also an unhappy house owner in* S*.*

The proof is simple: let Q be the set of selected house owners in some round. Let us write $B \geq B'$ for two balls if the radius of B is at least the radius of B'.

If all house owners in S are happy, we have

$$B(Q) = B(Q \cup S) \geq B(S) = B(P) \geq B(Q),$$

where the inequalities follow from the corresponding superset relations. The whole chain of inequalities implies that $B(P)$ and $B(Q)$ have the same radius, meaning that they must be equal (we had this argument before).

Since $|S| \leq d + 1$ by Theorem 16.5, we know that after k rounds, some element of S must have doubled its slips at least $k/(d + 1)$ times. This implies the following lower bound on the total number of slips.

Lemma 16.8 *After* k *rounds of the Swiss algorithm, the total number of slips is at least*

$$2^{k/(d+1)}.$$

16.2.2 An upper bound for the total number of slips

First, we want to argue that on average, not too many slips will be doubled in some round.