Appendix G

Smallest Enclosing Balls

This problem is related to the linear programming problem, but in a way it is much simpler, since a unique optimal solution always exists.

We let P be a set of n points in $\mathbb{R}^d$. We are interested in finding a closed ball of smallest radius that contains all the points in P, see Figure G.1.

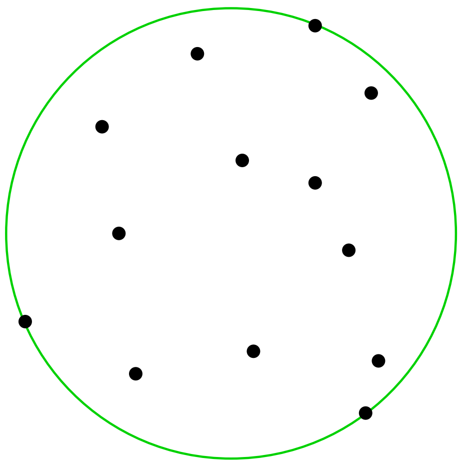

Figure G.1: *The smallest enclosing ball of a set of points in the plane*

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.

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

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

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 $\rho(c) \geqslant \rho(c_{opt})$ by optimality of c_{opt} in B, and if $c \notin B$, we get $\rho(c) \geqslant \|c - q\| > \rho(q) \geqslant \rho(c_{opt})$ since $q \in B$. In any case, we get $\rho(c) \geqslant \rho(c_{opt})$, so c_{opt} is indeed a best possible center.

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 G.1. *For $x \in \mathbb{R}^2$, we have*

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

Now we are prepared for the convex combination of balls.

Lemma G.2. *Let $B_0 = B(c_0, \rho_0)$ and $B_1 = (c_1, \rho_1)$ be two distinct 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 [3]. 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 G.2 illustrates the situation.

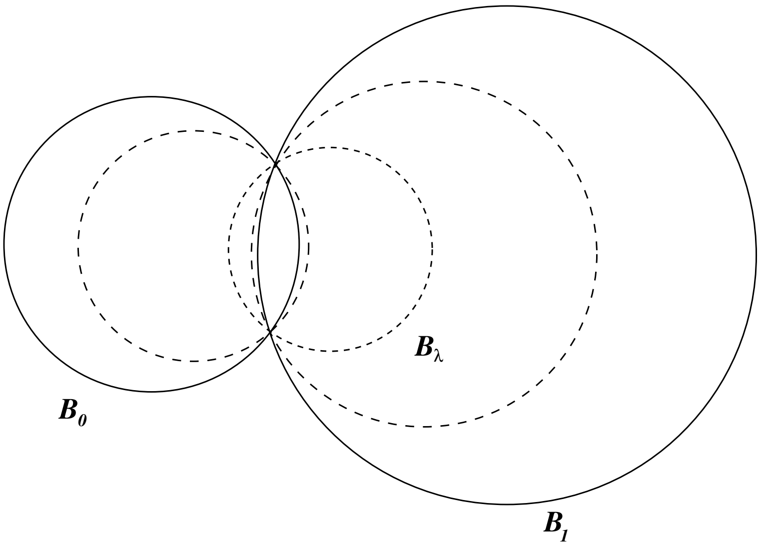

Figure G.2: *Convex combinations* B_λ *of two balls* B_0, B_1

Using this lemma, we can easily prove the following

Theorem G.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 G.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. □

Bases When you look at the example of Figure G.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. This is again a consequence of *Helly's Theorem* (Theorem 4.9).

Theorem G.4. *Let* $P \subseteq \mathbb{R}^d$ *be a finite point set. There is a subset* $S \subseteq P, |S| \leqslant d+1$ *such that* $B(P) = B(S)$*.*

Proof. If $|P| < d+1$, we may choose $S = P$. Otherwise, let ρ_{opt} be the radius of the smallest enclosing ball $B(P)$ of $P = \{p_1, \ldots, p_n\}$. Now define

$$C_i = \{x \in \mathbb{R}^d : \|x - p_i\| < \rho_{opt}\}, \quad i = 1, \ldots, n$$

to be the open ball around p_i with radius ρ_{opt}. 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 smaller than ρ_{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 smaller than ρ_{opt} either. Hence, $B(S)$ has radius at leat ρ_{opt}; but since $S \subseteq P$, the radius of $B(S)$ must also be at most ρ_{opt}, and hence it is equal to ρ_{opt}. But then $B(S) = B(P)$ follows, since otherwise, both $B(P)$ and $B(S)$ would be smallest enclosing balls of S, a contradiction. □

The previous theorem motivates the following

Definition G.5. *Let* $P \subseteq \mathbb{R}^d$ *be a finite point set. A* basis *of* P *is an inclusion-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 $B(P)$.

G.1 The trivial algorithm

Theorem G.4 immediately implies the following (rather inefficient) algorithm for computing $B(P)$: for every subset $S \subseteq P, |S| \leqslant d+1$, compute $B(S)$ (in fixed dimension d, this can be done in constant time), and return the one with largest radius.

Indeed, this works: for all $S \subseteq P$, the radius of $B(S)$ is at most that of $B(P)$, and there must be at least one $S, |S| \leqslant d+1$ (a basis of P) with $B(S) = B(P)$. It follows that the ball $B(T)$ being returned has the same radius as $B(P)$ and is therefore equal by $T \subseteq P$.

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

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

If $d = 2$ (the planar case), the trivial algorithm has runtime $O(n^3)$. In the next section, we discuss an algorithm that is substantially better than the trivial one in any dimension.

In order to adapt Seidel's randomized linear programming algorithm to the problem of computing smallest enclosing balls, we need the following statements.

Exercise G.6. *(i) Let* $P, R \subseteq \mathbb{R}^d, P \cap R = \emptyset$*. If there exists a ball that contains* P *and has* R *on the boundary, then there is also a unique smallest such ball which we denote by* $B(P, R)$*.*

(ii) Let $P, R \subseteq \mathbb{R}^d, P \cap R = \emptyset$*. If* $B(P, R)$ *exists and* $p \in P$ *satisfies* $p \notin B(P \setminus \{p\}, R)$*, then* p *is on the boundary of* $B(P, R)$*, meaning that* $B(P, R) = B(P \setminus \{p\}, R \cup \{p\})$*.*

Prove these two statements!

G.2 Welzl's Algorithm

The idea of this algorithm is the following. Given $P \subseteq \mathbb{R}^d$, the algorithm first recursively computes $B(P \setminus \{p\})$ where $p \in P$ is chosen uniformly at random. Then there are two cases: if $p \in B(P \setminus \{p\})$ we have $B(P) = B(P \setminus \{p\})$ (it's always good to rethink why this holds) and so we are done already. If $p \notin B(P \setminus \{p\})$, we still need to work, but the key fact (that we prove below) is that in this case, p has to be on the *boundary* of $B(P)$. We can therefore recursively compute the smallest enclosing ball of $P \setminus \{p\}$ that has p on its boundary, and this is a simpler problem because it intuitively has one degree of freedom less.

Let us formalize this idea.

Definition G.7. *Let* $P, R \subseteq \mathbb{R}^d$ *be disjoint finite point sets. We define* $B(P, R)$ *as the smallest ball that contains* P *and has the points of* R *on its boundary (*if *this ball exists and is unique).*

It is not hard to see that existence cannot always be guaranteed; for example if R is contained in the convex hull of P, there can be no ball that contains P and has even a single point of R on its boundary. But the following Lemma gives a number of useful properties.

Lemma G.8. *Let* $P, R \subseteq \mathbb{R}^d$ *be disjoint finite point sets, where* R *is affinely independent.*

(i) If there is any ball that contains P *and has* R *on its boundary, then a unique smallest such ball* $B(P, R)$ *exists.*

(ii) If $B(P, R)$ *exists and* $p \notin B(P \setminus \{p\}, R)$*, then* p *is on the boundary of* $B(P, R)$*, meaning that* $B(P, R) = B(P \setminus \{p\}, R \cup \{p\})$*.*

(iii) If $B(P, R)$ *exists, there is a subset* $S \subseteq P$ *of size* $|S| \leqslant d + 1 - |R|$ *such that* $B(P, R) = B(S, R)$*.*

Before we can prove this, we need a helper lemma.

Lemma G.9. *Let* $R \subseteq \mathbb{R}^d, |R| \geqslant 1$ *be an affinely independent point set. Then the set*

$$C(R) := \{c \in \mathbb{R}^d \mid c \textit{ is the center of a ball that has } R \textit{ on its boundary}\} \quad \text{(G.10)}$$

is a linear subspace of dimension $d + 1 - |R|$.

Proof. We have $c \in C(R)$ if and only if there exists a number ρ^2 such that

$$\rho^2 = \|c - p\|^2 = c^\top c - 2c^\top p + p^\top p, \quad p \in R. \quad \text{(G.11)}$$

Defining $\mu = \rho^2 - c^\top c$, this implies

$$\mu = p^\top p - 2c^\top p, \quad p \in R. \quad \text{(G.12)}$$

The set of all (c, μ) satisfying the latter $|R|$ equations is a linear subspace L of R^{d+1}.

Claim. L has dimension $d + 1 - |R|$.

To see this, let us write (G.12) in matrix form as follows.

$$\begin{pmatrix} p_{11} & p_{12} & \cdots & p_{1d} & 1 \\ p_{21} & p_{22} & \cdots & p_{2d} & 1 \\ \vdots & \vdots & & \vdots & \vdots \\ p_{|R|1} & p_{|R|2} & \cdots & p_{|R|d} & 1 \end{pmatrix} \begin{pmatrix} 2c_1 \\ 2c_2 \\ \vdots \\ 2c_d \\ \mu \end{pmatrix} = \begin{pmatrix} p_1^\top p_1 \\ p_2^\top p_2 \\ \cdots \\ p_{|R|}^\top p_{|R|} \end{pmatrix}, \quad \text{(G.13)}$$

where $R = p_1, \dots, p_{|R|}$. We know from linear algebra that the dimension of the solution space L is $d + 1$ minus the rank of the matrix. But this rank is $|R|$ since the rows are linearly independent by affine independence of R (we leave this easy argument to the reader, also as a good exercise to recall the definition of affine independence).

It only remains to show that $C(R)$ is also a linear subspace, and of the same dimension as L. But this holds since the linear function $f : C(R) \to L$ given by

$$f(c) = (c, p_1^\top p_1 - 2c^\top p_1)$$

is a bijection between $C(R)$ and L. The function is clearly injective, but it is also surjective: if $(c, \mu) \in L$, we satisfy (G.11) with $\rho^2 := \mu + c^\top c$, meaning that $c \in C(R)$ and

$$f(c) = (c, p_1^\top p_1 - 2c^\top p_1) = (c, \|c - p_1\|^2 - c^\top c) = (c, \rho^2 - c^\top c) = (c, \mu).$$

□

Now we can proceed with the proof of Lemma G.8.

Proof. The proof of (i) works along the same lines as the one for $B(P)$ in Section G, and it differs only for $R \neq \emptyset$. In this case, choose $s \in R$ arbitrarily.

Let $C(P, R)$ be the set of all $c \in \mathbb{R}^d$ that are centers of balls containing P, and with R on the boundary. Note that $C(P, R)$ is closed, since it results from intersecting the linear space $C(R)$ with the closed set

$$\{c \in \mathbb{R}^d \mid \max_{p \in P} \|c - p\| \leqslant \max_{p \in R} \|p - c\|\}.$$

By assumption, the set $C(P, R)$ is nonempty. We define

$$\rho(c) = \|s - c\|, \quad c \in C(P, R).$$

Thus, $\rho(c)$ is the radius of the *unique* ball centered at c that encloses all points of P and has all points of R on the boundary. To prove that there is some smallest ball containing P and with R on the boundary, we need to show that the continuous function ρ attains a minimum over $C(P, R)$. To be able to restrict attention to a closed and bounded (hence compact) subset of $C(P, R)$, we choose some $c_0 \in C(P, R)$; then $\rho(c_0)$ is certainly an upper bound for the minimum value, meaning that any $c \in C(P, R)$ outside the compact set

$$\{c \in C(P, R) \mid \rho(c) \leqslant \rho(c_0)\}$$

canot be a candidate for the center of a smallest ball. Within this compact set, we do get a minimum c_{opt}, and this is the desired center of a smallest enclosing ball of P that has R on the boundary.

To prove uniquenes, we invoke again the convex combination of balls: Assuming that there are two smallest balls B_0, B_1, then the ball

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

is a smaller ball that still contains P and still has R on its boundary (Lemma G.2 (ii)), a contradiction.

Now for part (ii). Again, convex combinations of balls come to our help. Consider the two balls $B(P, R)$ and $B(P \setminus \{p\}, R)$ (note that existence of the former implies existence of the latter via part (i)). If p is not on the boundary of $B(P, R)$, we have the situation of Figure G.3.

Then there is some small $\varepsilon > 0$ such that the ball

$$(1 - \varepsilon)B(P, R) + \varepsilon B(P \setminus \{p\}, R)$$

(drawn dashed in Figure G.3) still contains P and has R on the boundary, but has smaller volume than $B(P, R)$ by Lemma G.2 (iii), a contradiction.

Now we turn to (iii).

□

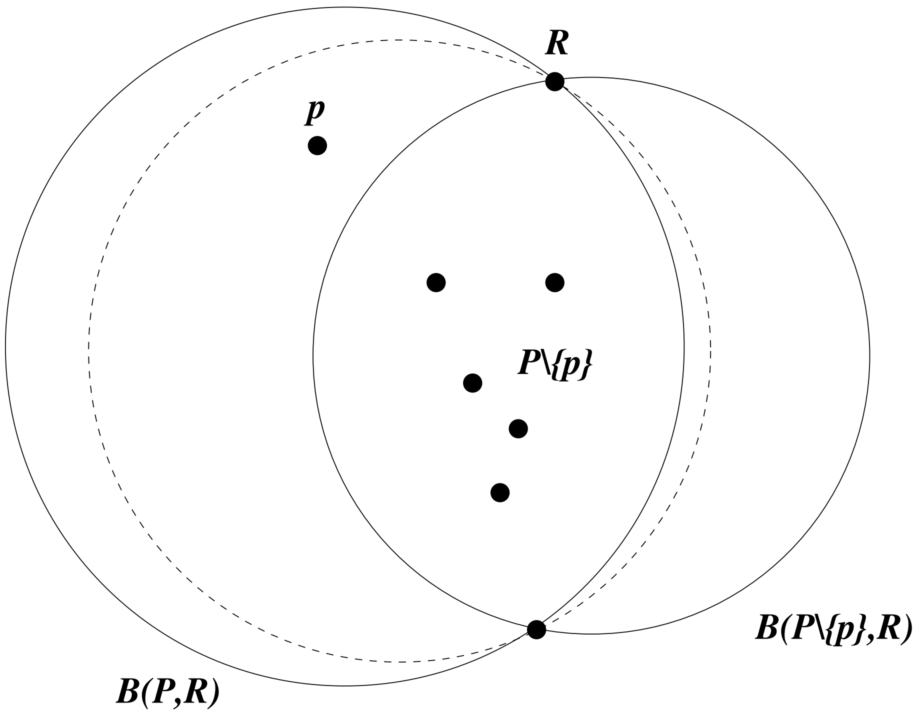

Figure G.3: *Proof of Lemma G.8(ii)*

G.3 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 ballot of paper with his/her name on it into a voting box. Then a constant number r (to be determined later) of ballots 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 ballots in the voting box are being doubled before drawing r ballots again. Thus, there are now two ballots for each unhappy house owner, and one for each happy one.

After round k, a house owner that has been unhappy after exactly i of the k rounds has therefore 2^i ballots 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 $\mathbb{R}^d$).

G.4 The Forever Swiss Algorithm

In the analysis, we want to argue about a fixed round k, but the above algorithm may never get to this round (for large k, we even strongly hope that it never gets there). But for the purpose of the analysis, we formally let the algorithm continue even if everybody is happy after some round (in such a round, no ballots are being doubled).

We call this extension the Forever Swiss Algorithm. A round is called *controversial* if it renders at least one house owner unhappy.

Definition G.14.

(i) Let m_k be the random variable for the total number of ballots after round k of the Forever Swiss Algorithm. We set $m_0 = n$, the initial number of ballots.

(ii) Let C_k be the event that the first k rounds in the Forever Swiss Algorithm are controversial.

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

Observation G.15. *After every controversial round, there is an unhappy house owner in S.*

Proof. Let Q be the set of selected house owners in the round. Let us write $B \geqslant B'$ for two balls if the radius of B is at least the radius of B'.

If all house owners in S were happy with the outcome of the round, we would have

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

where the inequalities follow from the corresponding superset relations. The whole chain of inequalities would then imply that $B(P)$ and $B(Q)$ have the same radius, meaning that they must be equal (we had this argument before). But then, *nobody* would be unhappy with the round, a contradiction to the current round being controversial. □

Since $|S| \leqslant d + 1$ by Theorem G.4, we know that after k rounds, some element of S must have doubled its ballots at least $k/(d + 1)$ times, given that all these rounds were controversial. This implies the following lower bound on the total number m_k of ballots.

Lemma G.16.

$$E(m_k) \geqslant 2^{k/(d+1)}\ \mathrm{prob}(C_k), \quad k \geqslant 0.$$

Proof. By the partition theorem of conditional expectation, we have

$$E(m_k) = E(m_k \mid C_k)\ \mathrm{prob}(C_k) + E(m_k \mid \overline{C_k})\ \mathrm{prob}(\overline{C_k}) \geqslant 2^{k/(d+1)}\ \mathrm{prob}(C_k).$$

□

An upper bound for $E(m_k)$ The main step is to show that the expected increase in the number of ballots from one round to the next is bounded.

Lemma G.17. *For all* $m \in \mathbb{N}$ *and* $k > 0$*,*

$$E(m_k \mid m_{k-1} = m) \leqslant m\left(1 + \frac{d+1}{r}\right).$$

Proof. Since exactly the "unhappy ballots" are being doubled, the expected increase in the total number of ballots equals the expected number of unhappy ballots, and this number is

$$\frac{1}{\binom{m}{r}} \sum_{|R|=r} \sum_{h \notin R} [h \text{ is unhappy with } R] = \frac{1}{\binom{m}{r}} \sum_{|Q|=r+1} \sum_{h \in Q} [h \text{ is unhappy with } Q \setminus \{h\}]. \tag{G.18}$$

Claim: Every $(r+1)$-element subset Q contains at most $d+1$ ballots such that h is unhappy with $Q \setminus \{h\}$.

To see the claim, choose a basis S, $|S| \leqslant d+1$, of the ball resulting from drawing ballots in Q. Only the removal of a ballot h belonging to some house owner $p \in S$ can have the effect that Q and $Q \setminus \{h\}$ lead to different balls. Moreover, in order for this to happen, the ballot h must be the *only* ballot of the owner p. This means that at most one ballot h per owner $p \in S$ can cause h to be unhappy with $Q \setminus \{h\}$.

We thus get

$$\frac{1}{\binom{m}{r}} \sum_{|R|=r} \sum_{h \notin R} [h \text{ is unhappy with } R] \leqslant (d+1)\frac{\binom{m}{r+1}}{\binom{m}{r}} = (d+1)\frac{m-r}{r+1} \leqslant (d+1)\frac{m}{r}. \tag{G.19}$$

By adding m, we get the new expected total number $E(m_k \mid m_{k-1} = m)$ of ballots. □

From this, we easily get our actual upper bound on $E(m_k)$.

Lemma G.20.

$$E(m_k) \leqslant n\left(1 + \frac{d+1}{r}\right)^k, \quad k \geqslant 0.$$

Proof. We use induction, where the case $k = 0$ follows from $m_0 = n$. For $k > 0$, the partition theorem of conditional expectation gives us

$$\begin{aligned} E(m_k) &= \sum_{m \geqslant 0} E(m_k \mid m_{k-1} = m)\, \mathrm{prob}(m_{k-1} = m) \\ &\leqslant \left(1 + \frac{d+1}{r}\right) \sum_{m \geqslant 0} m\, \mathrm{prob}(m_{k-1} = m) \\ &= \left(1 + \frac{d+1}{r}\right) E(m_{k-1}). \end{aligned}$$

Applying the induction hypothesis to $E(m_{k-1})$, the lemma follows. □

Putting it together Combining Lemmas G.16 and G.20, we know that

$$2^{k/(d+1)}\, \mathrm{prob}(C_k) \leqslant n \left(1 + \frac{d+1}{r}\right)^k,$$

where C_k is the event that there are k or more controversial rounds.

This inequality gives us a useful upper bound on $\mathrm{prob}(C_k)$, because the left-hand side power grows faster than the right-hand side power as a function of k, given that r is chosen large enough.

Let us choose $r = c(d+1)^2$ for some constant $c > \log_2 e \approx 1.44$. We obtain

$$\mathrm{prob}(C_k) \leqslant n \left(1 + \frac{1}{c(d+1)}\right)^k / 2^{k/(d+1)} \leqslant n 2^{k \log_2 e/(c(d+1)) - k/(d+1)},$$

using $1 + x \leqslant e^x = 2^{x \log_2 e}$ for all x. This further gives us

$$\mathrm{prob}(C_k) \leqslant n\alpha^k, \tag{G.21}$$

$$\alpha = \alpha(d, c) = 2^{(\log_2 e - c)/c(d+1)} < 1.$$

This implies the following tail estimate.

Lemma G.22. *For any $\beta > 1$, the probability that the Forever Swiss Algorithm performs at least $\lceil \beta \log_{1/\alpha} n \rceil$ controversial rounds is at most*

$$1/n^{\beta-1}.$$

Proof. The probability for at least this many controversial rounds is at most

$$\mathrm{prob}(C_{\lceil \beta \log_{1/\alpha} n \rceil}) \leqslant n\alpha^{\lceil \beta \log_{1/\alpha} n \rceil} \leqslant n\alpha^{\beta \log_{1/\alpha} n} = nn^{-\beta} = 1/n^{\beta-1}.$$

□

In a similar way, we can also bound the expected number of controversial rounds of the Forever Swiss Algorithm. This also bounds the expected number of rounds of the Swiss Algorithm, because the latter terminates upon the first non-controversial round.

Theorem G.23. *For any fixed dimension* d, *and with* $r = \lceil \log_2 e(d+1)^2 \rceil > \log_2 e(d+1)^2$, *the Swiss algorithm terminates after an expected number of* $O(\log n)$ *rounds.*

Proof. By definition of C_k (and using $E(X) = \sum_{m \geqslant 1} \text{prob}(X \geqslant m)$ for a random variable with values in $\mathbb{N}$), the expected number of rounds of the Swiss Algorithm is

$$\sum_{k \geqslant 1} \text{prob}(C_k).$$

For any $\beta > 1$, we can use (G.21) to bound this by

$$\begin{aligned}
\sum_{k=1}^{\lceil \beta \log_{1/\alpha} n \rceil - 1} 1 + n \sum_{k=\lceil \beta \log_{1/\alpha} n \rceil}^{\infty} \alpha^k &= \lceil \beta \log_{1/\alpha} n \rceil - 1 + n \frac{\alpha^{\lceil \beta \log_{1/\alpha} n \rceil}}{1-\alpha} \\
&\leqslant \beta \log_{1/\alpha} n + n \frac{\alpha^{\beta \log_{1/\alpha} n}}{1-\alpha} \\
&= \beta \log_{1/\alpha} n + \frac{n^{-\beta+1}}{1-\alpha} \\
&= \beta \log_{1/\alpha} n + o(1).
\end{aligned}$$

□

What does this mean for $d = 2$? In order to find the location of the firehouse efficiently (meaning in $O(\log n)$ rounds), 13 ballots should be drawn in each round. The resulting constant of proportionality in the $O(\log n)$ bound will be pretty high, though. To reduce the number of rounds, it may be advisable to choose r somewhat larger.

Since a single round can be performed in time $O(n)$ for fixed d, we can summarize our findings as follows.

Theorem G.24. *Using the Swiss Algorithm, the smallest enclosing ball of a set of* n *points in fixed dimension* d *can be computed in expected time* $O(n \log n)$.

The Swiss algorithm is a simplification of an algorithm by Clarkson [1, 2].

The bound of the previous Theorem already compares farovably with the bound of $O(n^{d+1})$ for the trivial algorithm, see Section G.1, but it does not stop here. We can even solve the problem in expected linear time $O(n)$, by using an adaptation of Seidel's linear programming algorithm [4].

Exercise G.25. *Let* H *be a set with* n *elements and* $f : 2^H \to \mathbb{R}$ *a function that maps subsets of* H *to real numbers. We say that* $h \in H$ violates $G \subseteq H$ *if* $f(G \cup \{h\}) \neq f(G)$ *(it follows that* $h \notin G$*). We also say that* $h \in H$ *is* extreme *in* G *if* $f(G \setminus \{h\}) \neq f(G)$ *(it follows that* $h \in G$*).*

Now we define two random variables $V_r, X_r : \binom{H}{r} \to \mathbb{R}$ *where* V_r *maps an* r*-element set* R *to the number of elements that violate* R*, and* X_r *maps an* r*-element set* R *to the number of extreme elements in* R*.*

Prove the following equality for $0 \leqslant r < n$*:*

$$\frac{E(V_r)}{n-r} = \frac{E(X_{r+1})}{r+1}.$$

Exercise G.26. *Imagine instead of doubling the ballots of the unhappy house owners in the Swiss Algorithm, we would multiply their number by some integer* $t \in \mathbb{N}$*. Does the analysis of the algorithm improve (i.e., does one get a better bound on the expected number of rounds, following the same approach)?*

Exercise G.27. *We have shown that for* $d = 2$ *and sample size* $r = 13$*, the Swiss algorithm takes an expected number of* $O(\log n)$ *rounds. Compute the constants, i.e., find numbers* c_1, c_2 *such that the expected number of rounds is always bounded by* $c_1 \log_2 n + c_2$*. Try to make* c_1 *as small as possible.*

G.5 Smallest Enclosing Balls in the Manhattan Distance

We can also compute smallest enclosing balls w.r.t. distances other than the Euclidean distance. In general, if $\delta : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$ is a metric, the smallest enclosing ball problem with respect to δ is the following.

Given $P \subseteq \mathbb{R}^d$, find $c \in \mathbb{R}^d$ and $\rho \in \mathbb{R}$ such that

$$d(c, p) \leqslant \rho, \quad p \in P,$$

and ρ is as small as possible.

For example, if $d(x, y) = \|x - y\|_\infty = \max_{i=1}^d |x_i - y_i|$, the problem is to find a smallest axis-parallel cube that contains all the points. This can be done in time $O(d^2 n)$ by finding the smallest enclosing box. The largest side-length of the box corresponds to the largest extent of the point set in any of the coordinate directions; to obtain a smallest enclosing cube, we simply extend the box along the other directions until all side lengths are equal.

A more interesting case is $d(x, y) = \|x - y\|_1 = \sum_{i=1}^d |x_i - y_i|$. This is the *Manhattan distance.* There, the problem can be written as

$$\begin{array}{ll} \text{minimize} & \rho \\ \text{subject to} & \sum_{i=1}^d |p_{ji} - c_i| \leqslant \rho, \quad j = 1, \ldots, n. \end{array}$$

where p_j is the j-th point and p_{ji} it's i-th coordinate Geometrically, the problem is now that of finding a smallest *cross polytope* (generalized octahedron) that contains the points. Algebraically, we can reduce it to a linear program, as follows.

We replace all $|p_{ji} - c_i|$ by new variables y_{ji} and add the additional constraints $y_{ji} \geqslant p_{ji} - c_i$ and $y_{ji} \geqslant c_i - pji$. The problem now is a linear program.

$$\begin{array}{lll} \text{minimize} & \rho & \\ \text{subject to} & \sum_{i=1}^{d} y_{ji} \leqslant \rho, & j = 1, \ldots, n \\ & y_{ji} \geqslant p_{ji} - c_i, & \forall i, j \\ & y_{ji} \geqslant c_i - p_{ji}, & \forall i, j. \end{array}$$

The claim is that the solution to this linear program also solves the original problem. For this, we need to observe two things: first of all, every optimal solution $(\tilde{c}, \tilde{\rho})$ to the original problem induces a feasible solution to the LP with the same value (simply set $y_{ji} := |p_{ji} - \tilde{c}_i|$), so the LP solution has value equal to $\tilde{\rho}$ or better. The second is that every optimal solution $((\tilde{y}_{ji})_{i,j}, \tilde{\rho})$ to the LP induces a feasible solution to the orginal problem with the same value: by $\sum_{i=1}^{d} \tilde{y}_{ji} \leqslant \tilde{\rho}$ and $\tilde{y}_{ji} \geqslant |p_{ji} - c_i|$, we also have $\sum_{i=1}^{d} |p_{ji} - c_i| \leqslant \tilde{\rho}$. This means, the original problem has value $\tilde{\rho}$ or better. From these two observations it follows that both problems have the same optimal value ρ_{opt}, and an LP solution of this value yields a smallest enclosing ball of P w.r.t. the Manhatten distance.

Questions

88. *Formulate the Swiss Algorithm for computing smallest enclosing balls, and discuss its relation with the Forever Swiss algorithm that we employ for the analysis!*

89. *The analysis of the Forever Swiss algorithm depends on a lower and an upper bound for the expected number of ballots after k controversial rounds. Sketch how these lower and upper bounds can be obtained, and how termination of the algorithm (with high probability) can be derived from them.*

90. *What is the expected runtime of the Swiss Algorithm for computing the smallest enclosing ball of a set of n points in fixed dimension d?*

91. *How can you compute smallest enclising balls in the Manhattan metric?*

References

[1] Kenneth L. Clarkson, A Las Vegas algorithm for linear programming when the dimension is small. In *Proc. 29th Annu. IEEE Sympos. Found. Comput. Sci.*, pp. 452–456, 1988.

[2] Kenneth L. Clarkson, Las Vegas algorithms for linear and integer programming when the dimension is small. *J. ACM*, **42**, (1995), 488–499.

[3] Kaspar Fischer, *Smallest enclosing balls of balls. Combinatorial structure and algorithms*. Ph.D. thesis, ETH Zürich, 2005.

[4] Emo Welzl, Smallest enclosing disks (balls and ellipsoids). In H. Maurer, ed., *New Results and New Trends in Computer Science*, vol. 555 of *Lecture Notes Comput. Sci.*, pp. 359–370, Springer, 1991.