Chapter 1

First Constructions of Epsilon-Nets

Consider the minimum hitting set problem for disks:

> Let P be a set of n points in $\mathbb{R}^2$ and $\mathcal{R}$ a collection of m subsets of P induced by disks in the plane. Then the minimum hitting set problem asks for a hitting set $Q \subseteq P$ for $\mathcal{R}$ of minimum cardinality.

This can be written as an integer program with n variables, say $x_p \in \{0,1\}$ for each $p \in P$, specifying whether p belongs to an optimal solution. Then the size of the optimal solution is simply

$$OPT = \sum_{p} x_p.$$

Relaxing this integer program gives a linear program where $x_p \in [0,1]$, and the goal is to minimize the sum of the x_p's. See the LP on the right.

Minimize $\sum_{p \in P} x_p$

subject to:

$$\sum_{p \in R} x_p \geq 1 \quad \forall R \in \mathcal{R}$$

$$0 \leq x_p \leq 1 \quad \forall p \in P.$$

Let

$$W^* = \sum_{p \in P} x_p$$

denote the value of the linear program. Then the LP constraint implies that the sum of the variables in each set $R \in \mathcal{R}$ is least 1. In other words, each set contains at least $\frac{1}{W^*}$-th of the total weight.

So the initial problem of finding a hitting set for $\mathcal{R}$—which could include sets of small cardinality—now reduces to the problem of finding a hitting set for all sets with weight at least $\frac{1}{W^*}$-th of the total weight. In particular, the LP will assign the elements in a small-sized set of $\mathcal{R}$ relatively large weights, and this guides us in the choice of a near-optimal hitting set.

If we could find a hitting set $Q \subseteq P$ of size at most $C \cdot W^*$ for this problem, for some constant C, then we will have

$$|Q| \leq C \cdot W^* \leq C \cdot OPT,$$

and so Q is a C-approximation to the optimal hitting set.

This task—called *rounding* in optimisation—is precisely the ϵ-net problem we will study.

1.1 Deterministic

An ϵ-net is a hitting set for those sets of $\mathcal{R}$ that contain at least an ϵ-th fraction of the elements of X.

Definition 1.1. *Given a set system $(X, \mathcal{R})$ and a parameter $\epsilon > 0$, a set $N \subseteq X$ is an ϵ-net for $(X, \mathcal{R})$ if for each $R \in \mathcal{R}$ with $|R| \geq \epsilon \cdot |X|$, we have $N \cap R \neq \emptyset$.*

Our goal is to show the existence of an ϵ-net of small size.

We will be interested in the case where $\mathcal{R}$ is derived from configurations of geometric objects. For example, consider the case where the base elements are a set P of n points in the plane, and the set system $\mathcal{R}$ is the primal set system induced on P by disks. See the figure for an example with $n = 16$ and a $\frac{1}{4}$-net consisting of 6 points. In this case, any disk containing at least $\frac{1}{4} \cdot 16 = 4$ points must contain one of the six points of the $\frac{1}{4}$-net.

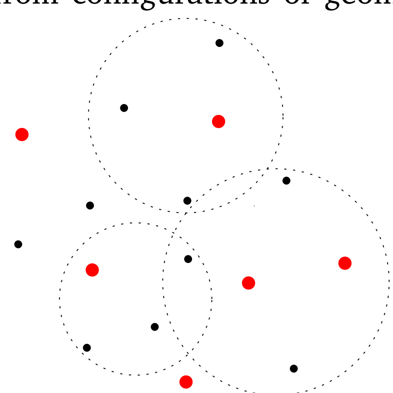

Note that there exist point sets where every ϵ-net must have size $\Omega\left(\frac{1}{\epsilon}\right)$. For example, arrange n points into groups of size ϵn, and place the points in each group within a small circle, and place these circles disjoint from each other. Clearly, N must contain at least one point from each circle, and there are $\lfloor \frac{1}{\epsilon} \rfloor$ disjoint circles.

On the other hand, constructing N by simply arbitrarily picking one point from each circle is not sufficient, as there could exist a disk containing ϵn points of P, but not completely containing any one circle, and so possibly not containing any point of N. See the figure.

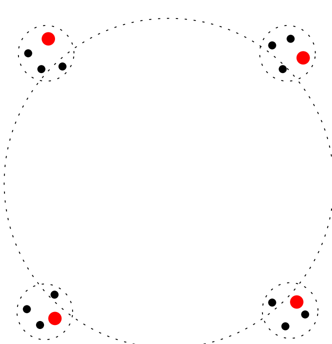

Surprisingly, as we will see later, this lower-bound is asymptotically the right one!

For this section, we show $O\left(\frac{1}{\epsilon}\right)$-sized ϵ-nets for an easier set system.

Theorem 1.1. *Given a set P of n points in the plane, and a parameter $\epsilon > 0$, there exists an ϵ-net of size $O\left(\frac{1}{\epsilon}\right)$ for the primal set system induced by halfplanes in the plane.*

⋆ ⋆ ⋆

Before considering this case, we examine some simpler set systems.

Intervals in $\mathbb{R}$. Given a set P of n points in $\mathbb{R}$, our goal is to pick an $N \subseteq P$ such that any interval that contains at least ϵn points of P contains some point of N. This is easy: sort the points of P by their coordinates and simply pick every ϵn-th point in this order. As each interval must contain a contiguous subset with respect to this ordering, it will be hit by N. The size of N is exactly $\left\lfloor \frac{n}{\epsilon n} \right\rfloor = \left\lfloor \frac{1}{\epsilon} \right\rfloor$.

Anchored rectangles in $\mathbb{R}^2$. Let P be a set of n points in $\mathbb{R}^2$, each with a positive y-coordinate. We seek an ϵ-net N for P with respect to rectangles anchored at the x-axis—in other words, any rectangle intersecting the x-axis and containing at least ϵn points of P should be hit by N.

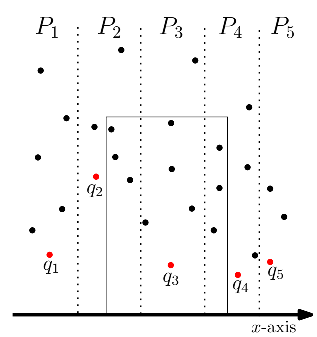

To construct N, assume the points of $P = \{p_1, \ldots, p_n\}$ are sorted by increasing x-coordinates. Partition P into $t = \left\lceil \frac{3}{\epsilon} \right\rceil$ sets $P_1, \ldots, P_t$ of contiguous points, with each P_i containing $\frac{\epsilon n}{3}$ points, except possibly the last set P_t. For each i, add the point with the lowest y-coordinate in P_i, say the point $q_i \in P_i$, to N. This is our ϵ-net, of size at most $\frac{3}{\epsilon}$.

To see why N is an ϵ-net, consider any anchored rectangle R containing at least ϵn points of P. Then R must contain points from at least 3 sets in our partition, say the sets P_i, P_j and P_k, where $i < j < k$. And so R must contain the point q_j with the lowest y-coordinate in P_j.

Half-spaces in $\mathbb{R}^2$. We prove the following theorem.

Theorem 1.1. *Given a set P of n points in the plane, and a parameter $\epsilon > 0$, there exists an ϵ-net of size $O\left(\frac{1}{\epsilon}\right)$ for the primal set system induced by halfplanes in the plane.*

Proof. First consider the easier case where P is in convex position. Then note that any half-space must contain a contiguous subset of P with respect to the order in which the points appear on the convex-hull. As before, picking every ϵn-th point along the convex-hull gives an ϵ-net of size $\lfloor \frac{1}{\epsilon} \rfloor$.

Otherwise, if P is not in convex position, then we'll first map the points in P to a set P' which *is* in convex position, construct an ϵ-net N' for P' as before, and from N' reconstruct an ϵ-net N for P.

Pick any point lying in the convex-hull, $\mathrm{conv}(P)$, of P. Say we pick $o \in P$. For each $p_i \in P$, trace a ray from o through p_i till this ray intersects the boundary of $\mathrm{conv}(P)$ in some point, say the point p_i'. Map p_i to p_i'; if p_i was on $\mathrm{conv}(P)$, then $p_i' = p_i$.

Let P' be the resulting set of $n-1$ points, now in convex position. Pick an ϵ-net N' for P', of size $\frac{1}{\epsilon}$.

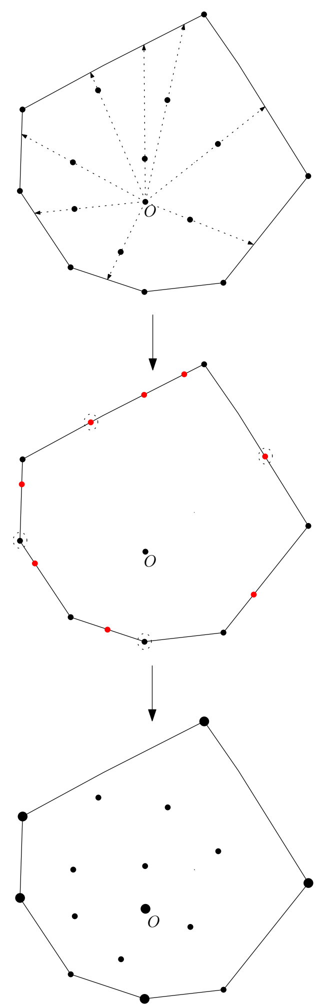

Now we show how to construct an ϵ-net N from N'. Take a point $p_i' \in N'$. If $p_i' = p_i$, i.e., p_i was already on $\mathrm{conv}(P)$, then add p_i to N. Otherwise, p_i' lies on some edge of $\mathrm{conv}(P)$ that is spanned by some two points of P[†]. Add both these points to N. Finally, add the point o to N.

Clearly

$$|N| \leq 1 + 2 \cdot |N'| = 1 + 2 \left\lfloor \frac{1}{\epsilon} \right\rfloor.$$

We claim that N is an ϵ-net. Consider any half-space H containing at least ϵn points of P. If it contains o, we're done as $o \in N$. Otherwise, by construction, we have the property that if $p_i \in H$, then the corresponding point p_i' also lies in H. So H contains at least ϵn points of P'. Therefore H contains a point $p_i' \in N'$, and so it must contain at least one of the two points of the edge of $\mathrm{conv}(P)$ on which p_i' lies. □

[†]Or p_i' lies on a point of P forming a vertex of $\mathrm{conv}(P)$, in which case we replace p_i' with this point instead.

Bibliography and discussion. The proof for halfplanes in $\mathbb{R}^2$ was invented here for didactic purposes; the different original proof is in [1].

[1] J. Komlós, J. Pach, and G. J. Woeginger. Almost tight bounds for epsilon-nets. *Discrete & Computational Geometry*, 7:163–173, 1992.

1.2 Probabilistic

Given a set system $(X, \mathcal{R})$, first observe that, regardless of the structure of $\mathcal{R}$, one can always get the following bound.

Theorem 1.2. *Let $(X, \mathcal{R})$ be a set system with $|X| = n$ and $|\mathcal{R}| = m$. Then given a parameter $\epsilon > 0$, there exists an ϵ-net N for $\mathcal{R}$ of size $O\left(\frac{\ln m}{\epsilon}\right)$.*

There are several (essentially equivalent) ways of viewing its proof. Note that for our purposes, assume that each set in $\mathcal{R}$ has size at least ϵn. Let $\mathcal{R} = \{R_1, \ldots, R_m\}$ be the m sets of $\mathcal{R}$.

Iterative view. Set $\mathcal{R}_1 = \mathcal{R}$. Pick an element of X that hits the maximum number of sets of $\mathcal{R}_1$, say $p_1 \in X$. Add p_1 to N, remove the sets hit by p_1 from $\mathcal{R}_1$ to get $\mathcal{R}_2$. Now re-iterate this procedure on $\mathcal{R}_2$.

Let $\mathcal{R}_i$ be the unhit sets after $i-1$ iterations, and let $p_i \in X$ be the point added to N in the i-th iteration.

At the i-th iteration, by the pigeonhole principle, there exists an element hitting at least

$$\frac{\sum_{R \in \mathcal{R}_i} |R|}{n} \geq \frac{\sum_{R \in \mathcal{R}_i} \epsilon n}{n} = \epsilon \cdot |\mathcal{R}_i|$$

sets of $\mathcal{R}_i$. Remove any $\epsilon |\mathcal{R}_i|$ such sets from $\mathcal{R}_i$ to get $\mathcal{R}_{i+1}$. Thus for any $i \geq 1$, we have

$$|\mathcal{R}_{i+1}| = (1-\epsilon)\,|\mathcal{R}_i| = (1-\epsilon)^2\,|\mathcal{R}_{i-1}| = \cdots = (1-\epsilon)^i \cdot |\mathcal{R}_1|.$$

At the i-th iteration, we have added i elements to N, and there are

$$m\,(1-\epsilon)^i \leq m\, e^{-\epsilon i}$$

unhit sets remaining. For $i = \Theta\left(\frac{\ln m}{\epsilon}\right)$, this becomes some constant, after which one can just add one element from each remaining unhit set to N. Therefore, we can hit all sets with $O\left(\frac{\ln m}{\epsilon}\right)$ elements.

Combinatorial view. Note that the pigeonholing in the proof above is essentially stating that, on average, each element of X hits ϵm sets of $\mathcal{R}$. Of course, this average goes down with the number of iterations, which is what results in the extra logarithmic factor. This can also be viewed succinctly combinatorially, for an integer t we will fix later:

count the number of subsets of X of size t that do not hit all sets of $\mathcal{R}$.

For each set $R \in \mathcal{R}$, there are at most $\binom{n-\epsilon n}{t}$ subsets of size t that do not hit R, and so there at most

$$\sum_{R \in \mathcal{R}} (\text{\# of } t\text{-sized subsets } Q \text{ with } Q \cap R = \emptyset) \;\leq\; m \cdot \binom{n-\epsilon n}{t}$$

subsets of X that do not hit at least one set of $\mathcal{R}$. If this is less than the total number of subsets of size t, then clearly there exists a set of size t that hits all sets of $\mathcal{R}$. A simple calculation shows this for $t = \frac{\ln(m+1)}{\epsilon}$:

$$\begin{aligned}
\frac{m \cdot \binom{n-\epsilon n}{t}}{\binom{n}{t}} &= m \cdot \frac{(n-t)!}{(n-\epsilon n-t)!} \cdot \frac{(n-\epsilon n)!}{n!} = m \cdot \frac{(n-t)(n-t-1)\cdots(n-\epsilon n+1-t)}{n(n-1)\cdots(n-\epsilon n+1)} \\
&= m \cdot \frac{(n-t)}{n} \cdots \frac{(n-\epsilon n+1-t)}{(n-\epsilon n+1)} = m \cdot \left(1 - \frac{t}{n}\right)\left(1 - \frac{t}{n-1}\right) \cdots \left(1 - \frac{t}{n-\epsilon n+1}\right) \\
&\leq m \cdot \left(1 - \frac{t}{n}\right)^{\epsilon n} \leq m \cdot e^{-\frac{t\epsilon n}{n}} = m \cdot e^{-\ln(m+1)} < 1.
\end{aligned}$$

Probabilistic view. Perhaps the simplest view is the probabilistic one: consider picking a random sample

$$S\colon \text{uniform random sample of } X \text{ of size } t.$$

Then

$$\Pr\left[\text{a fixed set } R \in \mathcal{R} \text{ is not hit by } S\right] = \left(1 - \frac{|R|}{n}\right)^t \leq (1-\epsilon)^t \leq e^{-t\epsilon}.$$

So the probability that at least one of the sets of $\mathcal{R}$ is not hit by S is at most $me^{-\epsilon t}$. For $t = \frac{\ln(m+1)}{\epsilon}$, this is less than 1. In particular there is a non-zero probability that S will hit all sets. Therefore, there has to exist at least one such set—in fact, many of them.

If the m sets can be arbitrary, then it is easy to see that asymptotically one cannot do much better than the above bound. For example, take $\mathcal{R}$ to be the power set of X, i.e., $m = 2^n$. Then, for $\epsilon = \frac{1}{2}$, any ϵ-net must have size at least $\frac{n}{2}$, while the above bound gives a $O(n)$-sized net. Obviously, set systems derived from geometry have considerable restrictions. We now show that, by a simple observation, one can actually get smaller ϵ-nets for a wide class of geometric set systems.

⋆ ⋆ ⋆

Axis-parallel rectangles in the plane

Let P be a set of n points in the plane, and $\mathcal{R} = \{P_1, \ldots, P_m\}$ the set system defined by containment by axis-parallel rectangles. In other words, $P_i \in \mathcal{R}$ if and only if there exists an axis-parallel rectangle R such that $P_i = R \cap P$. We now prove the existence of an ϵ-net for $(P, \mathcal{R})$ of size $O\left(\frac{1}{\epsilon}\log\frac{1}{\epsilon}\right)$.

For technical reasons, let $\mathcal{U}$ be a large-enough axis-parallel rectangle containing all points in P.

The key notion is that of a *canonical rectangle* spanned by a set of points.

> A canonical rectangle spanned by $Q \subset \mathcal{U}$ is a rectangle whose each bounding edge either passes through some point of Q or lies on $\partial\mathcal{U}$.

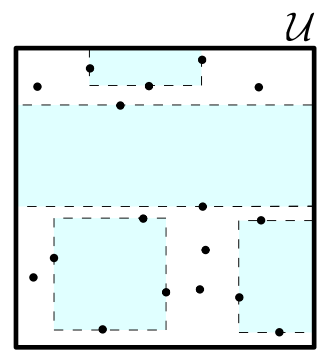

We remark that there are four types of canonical rectangles, namely those 'fixed' by 4, 3, 2 or 1 points. See the figure. For the rest of the proof, we only consider canonical rectangles fixed by 4 points; the other cases are similar.

There are $O\left(n^4\right)$ canonical rectangles spanned by any set of n points in the plane. This implies that $|\mathcal{R}| = O(n^4)$, as an arbitrary rectangle can always be 'shrunk' to get a canonical rectangle, containing precisely the points of the original rectangle.

As before, choose a random sample S by picking each point independently with probability p. The earlier analysis considered the probability that a set $P_i \in \mathcal{R}$ is not hit by S, and then we used the union over *all* sets of $\mathcal{R}$ to upper-bound the probability that there exists some set in $\mathcal{R}$ that is not hit by S. This gave the bound $O\left(\frac{1}{\epsilon}\log|\mathcal{R}|\right)$.

The trick here is to consider the *empty canonical rectangles* spanned by the points of S. Namely those canonical rectangles such that the point on each of its edges belongs to S, and that contain no point of S in their interior. Here is the new observation, using packing properties of Euclidean space.

Claim 1.3. If each empty canonical rectangle spanned by S has less than ϵn points of P, then S is an ϵ-net.

Proof. For contradiction, assume that there exists a rectangle R containing greater than ϵn points that is not hit by S. Then expand R by moving its left, right, top and bottom edges to transform it to an empty canonical rectangle R' spanned by S. But we assumed R' contained less than ϵn points of P, a contradiction. □

Let's do a rough calculation similar to the probabilistic proof of the general case. Say we add each point of P, independently with probability p, to S. Then we have $\mathrm{E}\left[|S|\right] = np$, and so there are an expected $O\left(|S|^4\right) = O\left((np)^4\right)$ canonical rectangles spanned by the points of S. For a fixed such rectangle R, if it contains at least ϵn points of P, then the probability that it does not contain any point of S is $(1-p)^{|R|} \le (1-p)^{\epsilon n}$. By the union bound, the probability that there exists some canonical rectangle spanned by S that does not contain any point of S is at most

$$|S|^4 (1-p)^{\epsilon n} \le (np)^4 \, e^{-p\epsilon n}.$$

To make it less than 1, we set $p = \frac{5}{\epsilon n} \ln \frac{1}{\epsilon}$ and so

$$= n^4 \left(\frac{5 \ln \frac{1}{\epsilon}}{\epsilon n} \right)^4 e^{-5 \ln \frac{1}{\epsilon}} = \frac{1}{\epsilon^4} \left(5 \ln \frac{1}{\epsilon} \right)^4 \epsilon^5 \ll 1.$$

Thus, there exists a set S, of expected size $O\left(\frac{1}{\epsilon} \log \frac{1}{\epsilon}\right)$, such that all canonical rectangles spanned by S either contain less than ϵn points of P or are hit by S. This implies that each *empty* canonical rectangle spanned by S contains less than ϵn points of P. We are done by Claim 1.3: S is an ϵ-net, of expected size $O\left(\frac{1}{\epsilon} \log \frac{1}{\epsilon}\right)$.

Of course, the above proof is technically incorrect for the following reason: we are calculating the probability that a canonical rectangle spanned by the points of S is empty of points of S! Instead, the correct way is to argue that, with non-zero probability, *no* canonical rectangle

(i) spanned by the points of P, and

(ii) containing at least ϵn points of P

ends up as an empty canonical rectangle spanned by the points of S.

Fix a canonical rectangle R spanned by four points of P. It ends up as an empty canonical rectangle in S if and only if

(i) the four points defining R are picked into S, and

(ii) none of the points contained in R are picked into S.

Let $\mathcal{E}_R$ be the event that R ends up as an empty canonical rectangle spanned by S. As each point of P was picked independently into S, we have

$$\Pr\left[\mathcal{E}_R\right] = p^4 \cdot (1-p)^{|R \cap P| - 4}.$$

Using the union bound over all possible $O(n^4)$ canonical rectangles spanned by points of P, the probability that at least one such rectangle containing greater than ϵn points ends up as an empty canonical rectangle in S is

$$\Pr\left[\bigcup_R \mathcal{E}_R\right] \le \sum_R \Pr\left[\mathcal{E}_R\right] = \sum_R p^4 \cdot (1-p)^{|R \cap P| - 4} \le n^4 \cdot p^4 \cdot (1-p)^{\epsilon n - 4} \le \frac{1}{6},$$

for $p = \frac{12}{\epsilon n} \log \frac{1}{\epsilon}$.

Similarly, one can show that the probability that there exists an empty canonical rectangle defined by 3 points is less than $\frac{1}{6}$, and identically for the canonical rectangles defined by 2 points. Therefore, with probability at least $\frac{1}{2}$, each empty canonical rectangle spanned by S has less than ϵn points of P.

By Claim 1.3, S is an ϵ-net, of expected size

$$\mathrm{E}\left[|S|\right] = np = \frac{12}{\epsilon} \log \frac{1}{\epsilon}.$$

By Chernoff bounds, the size of S is very sharply concentrated around its expectation. So the probability that $|S| \geq 2\,\mathrm{E}\left[|S|\right]$ is less than $\frac{1}{2}$†. Then with non-zero probability, S is an ϵ-net of size $O\left(\frac{1}{\epsilon} \log \frac{1}{\epsilon}\right)$.

⋆ ⋆ ⋆

The reader will notice that the only property of axis-aligned rectangles that is used in the proof is the concept of canonical rectangles—that a canonical rectangle R is 'fixed' by 4 points of P (or 3, 2, 1 points), and that R is spanned by S if and only if these points are picked. This is a very general idea, and we next give another example of its use.

Disks in the plane

Let $(P, \mathcal{R})$ be defined by disks in the plane, i.e., $P_i \in \mathcal{R}$ if and only if there exists a disk D in the plane such that $P_i = P \cap D$. We now show that a similar method of analysis also shows the existence of an ϵ-net of size $O\left(\frac{1}{\epsilon} \log \frac{1}{\epsilon}\right)$.

We again have the notion of a canonical disk.

> A canonical disk spanned by $Q \subseteq \mathbb{R}^2$ is either a disk passing through three points of Q, or a halfplane whose bounding line passes through two points of Q (such a halfplane can be seen as disk of infinite radius).

Then observe the following.

Claim 1.4. Let $Q \subset \mathbb{R}^2$ be a finite set of points, and D a disk not containing any point of Q. Then D lies in the union of at most *two* empty canonical disks spanned by Q.

Proof. Assume a disk D does not contain any point of Q, and has center c. Keeping the

†To avoid this very minor technical annoyance, often the sampling is done with a different distribution, by picking S uniformly from P over all subsets of size $t = \Theta\left(\frac{1}{\epsilon} \log \frac{1}{\epsilon}\right)$. Then we avoid having to do this, though then the probability calculation for a set being hit by S is slightly different. We get the same result, though.

center fixed, increase the radius of D until it touches some point, say q_1, of Q. Then, move c in the direction $\vec{q_1 c}$, away from q_1, while increasing the radius so that it still touches q_1. Eventually, the disk will touch another point, say $q_2 \in Q$.

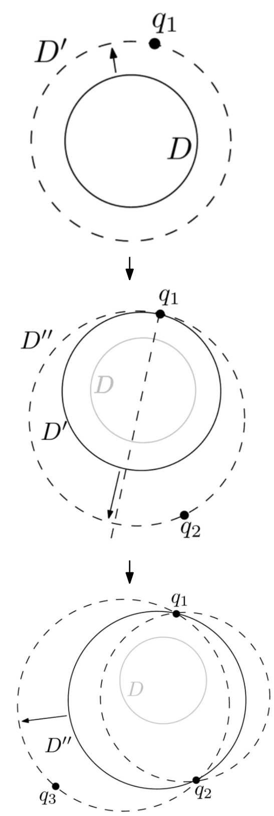

Note that c now lies on the perpendicular bisector of the segment $q_1 q_2$. Now by moving c along this bisector in both directions, one can get two canonical disks, say D_1 and D_2, such that $D \subset D_1 \cup D_2$. $\square$

Pick a random sample by adding each point of P, independently with probability p, to S. The above claim implies that a sample S would be an ϵ-net if one can ensure that each empty canonical disk spanned by S contains less than $\frac{\epsilon n}{2}$ points of P.

A disk D spanned by three points of P is an empty canonical disk spanned by S if and only if the three points are present in S, *and* no point of S lies inside D. The probability of this is

$$p^3 (1-p)^{|D \cap P|-3} .$$

There are $O(n^3)$ canonical disks spanned by P. Thus by the union bound, the probability that a canonical disk containing greater than $\frac{\epsilon n}{2}$ points of P ends up as an empty canonical disk spanned by S is at most

$$n^3 \cdot p^3 (1-p)^{\frac{\epsilon n}{2}-3} .$$

Setting $p = \frac{12}{\epsilon n} \log \frac{1}{\epsilon}$, the above probability becomes less than $\frac{1}{2}$. As before, this implies the existence of an ϵ-net of size $O\left(\frac{1}{\epsilon} \log \frac{1}{\epsilon}\right)$.

Bibliography and discussion. The idea of analyzing random samples using canonical structures is taken from the seminal paper of Clarkson [1].

[1] K. L. Clarkson. New applications of random sampling in computational geometry. *Discrete & Computational Geometry*, 2:195–222, 1987.

2.1 Linear-sized nets for Disks in $\mathbb{R}^2$

We now show that the bound of $O\left(\frac{1}{\epsilon}\log\frac{1}{\epsilon}\right)$ can be further improved in several cases by a fairly general idea. Specifically, for the previously considered range spaces obtained by disks and rectangles in the plane, one can get $o\left(\frac{1}{\epsilon}\log\frac{1}{\epsilon}\right)$ bounds.

The main theorem of this chapter is the following.

Theorem 2.1. *Let P be a set of n points in $\mathbb{R}^2$. Then there exists an ϵ-net N, of size $O\left(\frac{1}{\epsilon}\right)$, for the set system induced on P by disks in the plane.*

Intuitive idea. Recall the probabilistic argument for the existence of an ϵ-net of size $O\left(\frac{1}{\epsilon}\log\frac{1}{\epsilon}\right)$ for disks. We took a uniform random sample $S \subseteq P$ by choosing each point of P independently with probability $p = \Theta\left(\frac{1}{\epsilon n}\log\frac{1}{\epsilon}\right)$. We used this geometric property:

> if every empty canonical disk spanned by S has less than $\frac{\epsilon n}{2}$ points of P, then S is an ϵ-net for disks for P.

Since the total possible number of disks that could end up as empty canonical disks spanned by S is, naively counting, at most $O\left(n^3\right)$, and the probability of each such disk D ending up as an empty canonical disk induced by S is at most $p^3 \cdot (1-p)^{\frac{\epsilon n}{2}}$, the expected number of empty canonical disks in the random sample S can be upper-bounded by

$$O\left(n^3\right) \cdot p^3 \cdot (1-p)^{\frac{\epsilon n}{2}} .$$

To make this less than one, we set $p = \frac{10}{\epsilon n}\log\frac{1}{\epsilon}$. Then $\mathrm{E}[|S|] = np = \Theta\left(\frac{1}{\epsilon}\log\frac{1}{\epsilon}\right)$ and we're done.

The first new idea is to observe that the above analysis is imprecise: the probability of a canonical disk D ending up as an empty canonical disk spanned by S is equal to $p^3\left(1-p\right)^{|D\cap P|}$. Therefore, if $|D \cap P|$ is much larger than $\frac{\epsilon n}{2}$, this probability becomes considerably smaller, in fact decreasing exponentially with $|D \cap P|$. Thus the 'hard case' is considering disks containing $\Theta\left(\epsilon n\right)$ points.

Fortunately, the number of possible canonical disks containing $\Theta(\epsilon n)$ points is considerably smaller than the naive bound of $O\left(n^3\right)$—more generally, the number of canonical disks containing at most k points of P is at most $O\left(nk^2\right)$.

As an example, we compute the expected number of disks, each containing $c \cdot \epsilon n$ points for some fixed constant $c > 1$, that end up as empty canonical disks spanned by a random sample S constructed by picking each point of P independently with probability $\frac{1}{\epsilon n}$.

There are $O\left(n\left(c\epsilon n\right)^2\right)$ such canonical disks, and each ends up in the sample with probability $p^3\left(1-p\right)^{c\epsilon n}$. So the expected number of such disks that will end up as empty canonical disks spanned by S is

$$\Theta\left(n\left(c\epsilon n\right)^2 \cdot p^3\left(1-p\right)^{c\epsilon n}\right) = O\left(n\left(c^2\epsilon^2 n^2\right)\cdot\frac{1}{\epsilon^3 n^3}\cdot e^{-c}\right) = O\left(\frac{1}{\epsilon}\right).$$

This is bad news, since we had hoped to get no such disks. Now unfortunately, S need not be an ϵ-net: an arbitrary disk D' containing at least ϵn points could contain $\frac{\epsilon n}{2}$ points from one such canonical disk D, containing $c\epsilon n$ points, that has ended up as an empty canonical disk spanned by S.

Here is the second new idea: take these $O\left(\frac{1}{\epsilon}\right)$ expected number of disks, each containing roughly $c\epsilon n$ points, that have ended up as empty canonical disks spanned by S. For each such disk D, construct a $\left(\frac{1}{2c}\right)$-net, say S_D, for the set $D \cap P$. If we use the sub-optimal $O(\frac{1}{\epsilon}\log\frac{1}{\epsilon})$ bound from last lecture, we have

$$|S_D| = O(2c\log 2c).$$

Now any disk containing greater than $\frac{1}{2c}\cdot c\epsilon n = \frac{\epsilon n}{2}$ points from $D\cap P$ would be hit by S_D! Therefore, the set $S' = S \cup \bigcup_D S_D$ is an ϵ-net, with expected size

$$\mathrm{E}\left[|S|\right] + \mathrm{E}\left[\#\text{ of empty canonical disks } D,\ |D\cap P| \approx c\epsilon n,\text{ spanned by } S\right]\cdot O(2c\log 2c)$$

$$= np + O\left(\frac{1}{\epsilon}\right)\cdot O\left(2c\log 2c\right) = O\left(\frac{1}{\epsilon}\right) + O\left(\frac{1}{\epsilon}\cdot 2c\log 2c\right) = O\left(\frac{1}{\epsilon}\right).$$

So while it is true that there could be $\Theta(\frac{1}{\epsilon})$ empty canonical disks spanned by S containing $c\epsilon n$ points of P, for each disk, we can add some $O(1)$ additional points such that a disk containing $\frac{\epsilon n}{2}$ points from any such disk would be hit by them.

It remains to take care of all canonical disks containing at least ϵn points. Note that while the number of disks containing at most k points increases polynomially with k—as $O\left(nk^2\right)$, the probability of each ending up as an empty canonical disk decreases exponentially with k, i.e., as $p^3\left(1-p\right)^k$. So essentially the worst case is the above one. We now do the above calculation for all k, and then sum up to get the size of the final ϵ-net.

PROOF OF THEOREM 2.1.

Set $p = \frac{1}{\epsilon n}$, and add each point of P, independently with probability p, to S. We now add additional points to S to get our final net N.

For each empty canonical disk D induced by S, do the following. Let i be the index such that

$$2^{i-1}\cdot\epsilon n < |D\cap P| \le 2^i\cdot\epsilon n.$$

We will add to N an ϵ_i-net S_D for the set system induced by disks on $D \cap P$, where we set ϵ_i such that any disk containing at least $\frac{\epsilon n}{2}$ points from $D \cap P$ would be hit by S_D. In other words, we want

$$\frac{\epsilon n}{2} \geq \epsilon_i \cdot |D \cap P|$$
$$\implies \qquad \epsilon_i \leq \frac{\epsilon n}{2|D \cap P|},$$

which is satisfied if we set $\epsilon_i = \frac{1}{2^{i+1}}$.[†]

Claim 2.2.

$$N = S \quad \cup \bigcup_{\substack{D \text{ empty canonical} \\ \text{disk induced by } S}} S_D \qquad \text{is an } \epsilon\text{-net for } (P, \mathcal{R}).$$

Proof. Let D' be any disk in the plane containing at least ϵn points of P. Then either it is hit by S, or it contains at least $\frac{\epsilon n}{2}$ points from an empty canonical disk induced by S, say disk D. Then D' is hit by S_D. □

It simply remains to do the required calculations and sum up to bound the size of the final ϵ-net N.

For a canonical disk D induced by P, let I_D be the indicator random variable which is 1 if D ends up as an empty canonical disk spanned by S, and 0 otherwise. Then, the expected number of additional points added are

$$\begin{aligned}
\mathrm{E}\left[\sum_{\text{Disks } D} I_D \cdot |S_D|\right] &= \sum_D |S_D| \cdot \mathrm{E}[I_D] = \sum_D |S_D| \cdot \Pr[\text{D is an empty canonical disk}] \\
&= \sum_D |S_D| \cdot p^3 (1-p)^{|D \cap P|} \\
&= \sum_i \sum_{2^i \epsilon n < |D \cap P| \leq 2^{i+1} \epsilon n} |S_D| \cdot p^3 (1-p)^{|D \cap P|} \\
&\leq \sum_i \left|\left\{D \colon 2^i \epsilon n < |D \cap P| \leq 2^{i+1} \epsilon n\right\}\right| \cdot \left(2^{i+1} \log 2^{i+1}\right) \cdot p^3 (1-p)^{2^i \epsilon n} \\
&\leq \sum_i n \left(2^{i+1} \epsilon n\right)^2 \cdot \left(2^{i+1} \log 2^{i+1}\right) \cdot p^3 e^{-p 2^i \epsilon n} \\
&= \sum_i n \left(2^{2i+2} \epsilon^2 n^2\right) \cdot \left(2^{i+1}(i+1)\right) \cdot \frac{1}{\epsilon^3 n^3} e^{-2^i} \\
&= \frac{1}{\epsilon} \sum_i \frac{2^{3i+3}(i+1)}{e^{2^i}}
\end{aligned}$$

[†]As $\frac{1}{2^{i+1}} \leq \frac{\epsilon n}{2(2^i \epsilon n)}$.

$$= O\left(\frac{1}{\epsilon}\right), \text{ since the summation is a decreasing geometric series.}$$

The expected size of S is $np = \frac{1}{\epsilon}$, and so S together with $\bigcup_D S_D$ forms an ϵ-net of expected size $O\left(\frac{1}{\epsilon}\right)$.

Bibliography and discussion. The idea of sampling and refinement was first used to construct optimal-sized cuttings by Chazelle and Friedman [3]. The proof of optimal ϵ-net for disks given above was constructed ex post facto for didactic purposes. A more precise construction and analysis giving ϵ-nets of size at most $\frac{13.4}{\epsilon}$ for the set system induced by disks in the plane was given in [2]. This sampling refinement idea can also be done for dual set systems, and was first done in [4], and improved in [1].

[1] B. Aronov, E. Ezra, and M. Sharir. Small-size ϵ-nets for axis-parallel rectangles and boxes. *SIAM Journal on Computing*, 39(7):3248–3282, 2010.

[2] N. Bus, S. Garg, N. H. Mustafa, and S. Ray. Tighter estimates for ϵ-nets for disks. *Comput. Geom.*, 53:27–35, 2016.

[3] B. Chazelle and J. Friedman. A deterministic view of random sampling and its use in geometry. *Combinatorica*, 10(3):229–249, 1990.

[4] K. Clarkson and K. Varadarajan. Improved approximation algorithms for geometric set cover. *Discrete & Computational Geometry*, 37:43–58, 2007.