1.12 Weak ϵ-nets

Consider the range space $(P, \mathcal{R})$, where P is a set of n points in $\mathbb{R}^d$, and $\mathcal{R}$ consists of all possible subsets induced by convex objects in $\mathbb{R}^d$, i.e., $P' \in \mathcal{R}$ iff there exists a convex object $C \subseteq \mathbb{R}^d$ such that $P' = C \cap P$. A subset $Q \subseteq \mathbb{R}^d$ is a *weak* ϵ-net if $C \cap Q \neq \emptyset$ for all convex objects C containing at least ϵn points of P.

Weak ϵ-nets of size $O(1/\epsilon^{d+1})$. The proof is using the so-called First Selection Lemma. Recall its statement:

Lemma 17 (First Selection Lemma)**.** *Given any set P of n points in $\mathbb{R}^d$, there exists a point $q \in \mathbb{R}^d$ contained in at least $c_d \cdot \binom{n}{d+1}$ d-simplices spanned by P.*

Set $Q_0 = \emptyset, i = 0$ and construct the weak ϵ-net iteratively: If all convex objects C containing at least ϵn points of P are hit by Q_i, we are done. Otherwise, find a point q that lies in the largest number of simplices spanned by $C \cap P$: by the First Selection Lemma, q lies in at least $c_d \cdot \binom{|C \cap P|}{d+1} \geq c_d \cdot \binom{\epsilon n}{d+1}$ such simplices. Note that since C does not contain any point of Q_i, no simplex spanned by $C \cap P$ is hit by any point of Q_i. Set $Q_{i+1} = Q_i \cup \{q\}$, and iterate. At each step, the new point hits $c_d \cdot \binom{\epsilon n}{d+1}$ previously un-hit simplices spanned by P, so the process can go on for at most

$$\frac{\binom{n}{d+1}}{c_d \cdot \binom{\epsilon n}{d+1}} \leq \frac{(en/(d+1))^{d+1}}{c_d \cdot (\epsilon n/(d+1))^{d+1}} = O(\frac{1}{\epsilon^{d+1}})$$

steps, proving the stated result.

A general theme in weak ϵ-nets. Our goal is to improve the above bound to $O(1/\epsilon^d)$. Almost all weak ϵ-net constructions known so far use the following two basic ideas.

First, partition P into t equal-sized sets $P_1, \ldots, P_t$ of n/t points each. Construct a set Q_k such that any convex object intersecting more than $k\epsilon$ sets P_i must contain a point of Q_k. Then Q_t is an ϵ-net: any convex set C containing ϵn points must intersect at least $\epsilon n/(n/t) = t\epsilon$ sets, and so be hit by Q_t. Of course, a convex set intersecting more P_i's is easier to hit as it has stronger structural properties. So the size of Q_k is a decreasing function of k (and an increasing function of $1/\epsilon$).

Second, note that if C intersects only $t\epsilon$ sets, then it contains all points from each P_i that it intersects (to make up the ϵn points it contains). In that case, why not add, recursively, a weak ϵ'-net, for a suitably determined ϵ', for each P_i. The advantage is that if C is not hit by any of the points added recursively, it would have to contain few points (at most $\epsilon' n/t$, by the weak ϵ'-net property) from each P_i, and so intersect considerably more sets – at least $(\epsilon n)/(\epsilon' n/t) = t\epsilon/\epsilon'$ sets. Then $Q_{t/\epsilon'}$ is an ϵ-net, resulting in a lower size of Q. Fixing the parameters for this trade-off then improves the bound from the first idea.

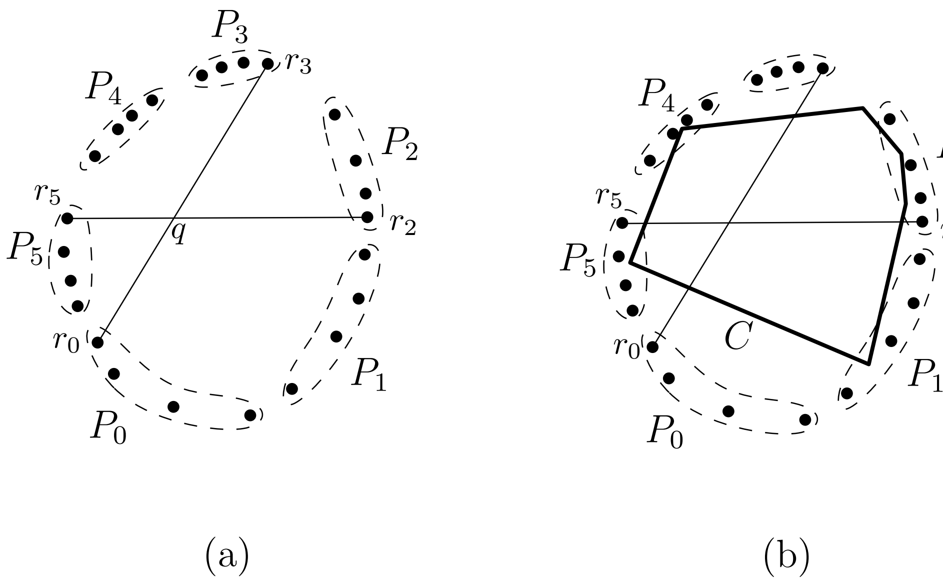

Figure 1.18: (a) Partitioning P along the cyclic order, and the point q for the pair P_2 and P_5. (b) C intersects P_1, P_2, P_4 and P_5, and so the point for the pair P_2 and P_5 must lie in C.

Improved bound in $\mathbb{R}^2$. We first give the improved bound for $\mathbb{R}^2$, following a similar scheme as above. Partition P into equal-sized sets P_1 and P_2 by a vertical line l. Now, a convex object C containing at least $\epsilon n/4$ points from P_1 and P_2 (say $P_1' \subseteq P_1$ and $P_2' \subseteq P_2$) will contain all intersection points of each segment pq, $p \in P_1', q \in P_2'$, with l. So there are at least $(\frac{\epsilon n}{4})^2$ intersections lying within the interval $C \cap l$ on l. Therefore, picking every $\frac{\epsilon^2 n^2}{16}$-th intersection point (when sorted by the y-coordinate) of all $(n/2)^2$ segment intersections with l would hit all such C. We have picked $O(1/\epsilon^2)$ points.

Otherwise, C contains at least $3\epsilon n/4$ points from one partition, say P_1. Then a weak ϵ'-net for P_1, $\epsilon' = (3\epsilon n/4)/(n/2) = 3\epsilon/2$, would hit all such C.

The size, $f(\epsilon)$, of the constructed ϵ-net is:

$$f(\epsilon) = 2f(\frac{3\epsilon}{2}) + O(1/\epsilon^2)$$

which solves to $f(\epsilon) = O(1/\epsilon^2)$.

A further improved bound in $\mathbb{R}^2$ for points in convex position. One can use the structural property of points in convex position to improve the bound above to $\tilde{O}(1/\epsilon)$[9]. Let $P = \{p_0, \ldots, p_{n-1}\}$ be the n points, sorted in the anti-clockwise direction along their order in the convex hull.

Partition P into equal-sized sets $P_0, \ldots, P_{4/\epsilon}$ by their consecutive order in the cyclic sequence, i.e., $P_i = \{p_{i\epsilon n/4}, \ldots, p_{(i+1)\epsilon n/4-1}\}$. Call the first point of each set its *representative*, i.e., P_i

[9] $\tilde{O}$ means ignoring polylogarithmic factors.

has as its representative the point $p_{i\epsilon n/4}$. Let's re-label the representative point of set P_i as the point r_i, and let P' be this set of representative points.

Now, construct the weak ϵ-net Q as follows: for all pairs of sets P_i and P_j, $0 \leq i < j \leq 4/\epsilon$, add to Q the intersection point of the two segments $\overline{r_0 r_{i+1}}$ and $\overline{r_i r_j}$. Note that $|Q| = O(1/\epsilon^2)$. See Figure 1.18(a).

We claim that Q is an ϵ-net: any convex object containing at least ϵn points must intersect at least *four* sets, as each set contains at most $\epsilon n/4$ points. Say C intersects the sets P_i, P_j, P_k and P_l. Then the point added to Q for the pair of sets P_j and P_l will lie inside C. See Figure 1.18(b).

Now the second idea of recursive construction yields the desired result. Say we partition P into t groups $P_1, \ldots, P_t$ as before, and add Q as constructed earlier. By the same argument, any C which intersects at least *four* sets will contain a point of Q. Otherwise, it must contain at least $\epsilon n/3$ points from one of the (at most three) sets that it intersects. Since each set has n/t points, we recursively construct a weak $(\epsilon t/3)$-net for each P_i: then any C containing at least $(\epsilon t/3) \cdot n/t = \epsilon n/3$ from a set will be hit inductively.

The size, $f(\epsilon)$, of the constructed ϵ-net is:

$$f(\epsilon) = t \cdot f(\frac{t\epsilon}{3}) + O(t^2)$$

Setting $t = 3/\sqrt{\epsilon}$, this solves to $f(\epsilon) = \tilde{O}(1/\epsilon)$.

Weak ϵ-nets of size $O(1/\epsilon^d)$. We now present the current-best general bound for weak ϵ-nets in $\mathbb{R}^d$. The idea is an elegant one, which once observed, then works out easily from the general scheme of constructing weak ϵ-nets.

The new idea is to use the following important theorem (which we will cover later on):

Theorem 18 (The Partition Theorem)**.** *Given a set P of n points in $\mathbb{R}^d$ and an integer t, one can partition P into t equal-sized* [10] *sets $P_1, \ldots, P_t$ such that any half-plane intersects the convex-hull of at most $t^{1-1/d}$ sets.*

Given P, partition P into t sets using the Partition Theorem, where the parameter t will be set optimally later on. Pick an arbitrary point of P from each partition as it's representative point, and let P' be this set of t points. Construct a set Q_t with the following property: a centerpoint of any subset $P'' \subseteq P'$ is in Q_t. Now we claim that for $t = 2(d+1)^d/\epsilon^d$, Q_t is a weak ϵ-net.

Why? Well, any convex object containing at least ϵn points of P intersects at least $\epsilon n/(n/t) = t\epsilon$ sets of the partition. Let P'' be the representative points picked from these sets, and $q \in Q_t$ be their centerpoint. If q lies inside C, we are done. Otherwise, there exists a half-plane h separating q from C. And so, by the centerpoint property, the halfspace containing q and

[10] To be precise, within a factor of two of each other.

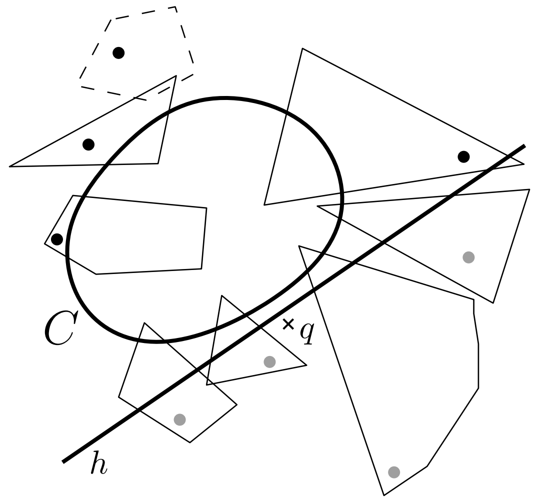

Figure 1.19: The representative points of $P'' \setminus P'''$ are black, while the representative points of P''' are grey.

not containing C must contain at least $t' = |P''|/(d+1)$ points of P'', say P'''. Crucially, the convex-hull of the sets that these points in P''' are representative of must intersect h: each such set intersects C, and also contains a point of P''', and these two lie on different sides of h. See Figure 1.19.

So h intersects the convex-hull of at least $|P''|/(d+1) \geq t\epsilon/(d+1)$ sets of the partition. On the other hand, by the partition theorem, h intersects at most $t^{1-1/d}$ sets. Setting them equal to find t, we get a contradiction if $t > (d+1)^d/\epsilon^d$. Therefore, for $t = 2(d+1)^d/\epsilon^d$, Q_t is a weak ϵ-net.

Note that $|Q_t| = t^{d^2}$: recall that for any subset $P'' \subseteq P$, the centerpoint of P'' is constructed by computing the common intersection of the convex-hulls of all subsets of at least $|P''|/(d+1)$ points of P''. Any vertex of this common intersection is a centerpoint of P'', and is the intersection of d planes, each defined by d points of P''. By adding the intersection points of all d-tuples of planes, where each plane is fixed by d points of P'', one gets the required set Q_t with the stated size.

Now, the standard recursive construction yields our result. Given P, use the partition theorem to get $P_1, \ldots, P_t$. Add Q_t to our weak ϵ-net, and recursively construct a ϵ'-net for each P_i. This forms our constructed set, and we prove now that it is indeed a weak ϵ-net (for appropriate values of t and ϵ').

For a convex object C containing at least ϵn points, if it contains at least $\epsilon' n/t$ points from one of the partitions, the recursive construction guarantees hitting it.

Otherwise, C contains less than $\epsilon' n/t$ points from each set. And therefore C must intersect more than $\epsilon n/(\epsilon' n/t) = t\epsilon/\epsilon'$ sets. Denote these sets by P'', and let $q \in Q_t$ be their center-point. We claim that q lies in C. Otherwise, there exists a half-plane h separating C from q.

Arguing as before, h must then intersect more than $|P''|/(d+1) = \frac{t\epsilon}{\epsilon'(d+1)}$ sets. On the other hand, h intersects at most $t^{1-1/d}$ sets, and so we get a contradiction for $\epsilon' \leq \frac{\epsilon t^{1/d}}{d+1}$. Therefore for this value of ϵ', q must lie in C.

The size, $f(\epsilon)$, of the constructed weak ϵ-net is:

$$f(\epsilon) = t \cdot f(\epsilon') + |Q_t| = t \cdot f(\frac{\epsilon t^{1/d}}{d+1}) + t^{d^2}$$

Setting $t = \tilde{\Theta}(1/\epsilon^{1/d})$, this solves to $f(\epsilon) = \tilde{O}(1/\epsilon^d)$.