CHAPTER 1

A Probabilistic Averaging Technique

This chapter presents a powerful probabilistic proof technique which will be used, in combination with additional ideas, throughout this book. For a simple first application, consider the following problem.

Let $P = \{p_1, \ldots, p_n\}$ be a set of n points in the plane and let $k \geq 0$ be an integer. The k-Delaunay graph of P, denoted by $G_k(P) = (P, E_k)$, is defined as follows:

> $\{p_i, p_j\} \in E_k$ if and only if there exists a closed disk in the plane containing $\{p_i, p_j\}$ and at most k points of $P \setminus \{p_i, p_j\}$.
>
> For each $\{p_i, p_j\} \in E_k$ fix any one such disk and denote it by D_{ij}.

Note that $E_0 \subseteq E_1 \subseteq \cdots \subseteq E_{n-2} = \binom{P}{2}$.

Our goal is to upper bound the number of edges in the k-Delaunay graph of P as a function of n and k. We will prove that $|E_k| = O(n(k+1))$[1].

First, observe that the 0-Delaunay graph of P is simply the Delaunay graph, which is planar and thus E_0 has size at most $3n$. Next, we upper bound $|E_k|$, for any integer $k \geq 1$, by the following argument.

Let S be a random sample constructed by picking each point of P independently with probability $p = \frac{1}{k+1}$ and let $G_0(S)$ be the 0-Delaunay graph of S. We count the expected number of edges in $G_0(S)$ in two ways.

Upper bound: As *any* 0-Delaunay graph on t vertices has at most $3t$ edges, the expected number of edges in $G_0(S)$ is

$$\mathrm{E}\left[3|S|\right] = 3\,\mathrm{E}\left[|S|\right] = 3np = \frac{3n}{k+1}.$$

Lower bound: For any $\{p_i, p_j\} \in E_k$, if both p_i and p_j are picked in S and none of the at most k other points of P lying in D_{ij} are picked in S, then $\{p_i, p_j\}$ is an edge in $G_0(S)$. As each point of P was picked independently, the probability that $\{p_i, p_j\}$ is an edge in $G_0(S)$ is at least

$$p^2 \cdot (1-p)^k = \frac{1}{(k+1)^2} \cdot \left(1 - \frac{1}{k+1}\right)^k \geq \frac{1}{(k+1)^2} \cdot \frac{1}{e}, \tag{1.1}$$

where the last step uses the fact that $\left(1 + \frac{1}{k}\right)^k \leq e$, and thus $\frac{1}{e} \leq \left(\frac{k}{k+1}\right)^k = \left(1 - \frac{1}{k+1}\right)^k$.

[1] The '+1' term is there just to take care of the case $k = 0$.

Using linearity of expectation, Equation (1.1) implies that the expected number of edges of E_k that are present in $G_0(S)$ is at least $|E_k| \cdot \frac{1}{(k+1)^2 e}$.

Combining the upper and lower bounds, we get

$$|E_k| \cdot \frac{1}{(k+1)^2 e} \;\leq\; \text{expected number of edges in } G_0(S) \;=\; \frac{3n}{k+1},$$

implying that $|E_k| = O(n(k+1))$.

Before we move on to other applications of this technique, we make a few remarks.

- The use of random sampling in the above proof is a way to 'implement' a double-counting argument. Essentially we are summing up, over all edges e in $G_k(P)$, the number of subsets $S \subseteq P$ of size $\frac{n}{k+1}$ for which e is an edge in $G_0(S)$[2]. We counted this sum in two ways: iterating over edges of $G_k(P)$ gave a lower bound while iterating over subsets of P gave an upper bound. More precisely, define the set of pairs

$$\mathcal{I} = \big\{\, (e, S) : e \in E_k,\ |S| = \lceil n/(k+1) \rceil,\ e \text{ is in } G_0(S) \big\}.$$

  Then the above double-counting argument gives

$$|E_k| \cdot \binom{n-2-k}{\lceil n/(k+1) \rceil - 2} \;\leq\; |\mathcal{I}| \;\leq\; \binom{n}{\lceil n/(k+1) \rceil} \cdot 3\lceil n/(k+1) \rceil.$$

  Solving this for $|E_k|$ gives $|E_k| = O(n(k+1))$, as before.
- The utility of framing the argument probabilistically is that it beautifully captures the intuition behind the key idea: if there are 'too many' edges in $G_k(P)$, then in expectation more than $3|S|$ of these edges will 'filter through' to $G_0(S)$, for a random sample S. This contradicts the fact that for *any* S, $G_0(S)$ has at most $3|S|$ edges.
- The lower bound follows by considering, for each edge $\{p_i, p_j\}$ of $G_k(P)$, a specific event whose occurrence implies that $\{p_i, p_j\}$ appears as an edge in $G_0(S)$. This need not be the *only* event that could cause $\{p_i, p_j\}$ to be an edge in $G_0(S)$—e.g., there could be a disk other than D_{ij} containing p_i and p_j that happens to not contain any other point of S. Thus our lower bound is not necessarily tight.

  In fact, what we actually want to compute is a lower bound on the probability that there exists *some* disk containing $\{p_i, p_j\}$ and no other point of S. However the events for all possible disks containing $\{p_i, p_j\}$ are not independent, which makes computing this probability difficult. Fortunately, we do not lose much by considering *any one* such disk, and in fact the lower bound is optimal up to constant factors for certain point sets. In particular, the bound $|E_k| = O(n(k+1))$ is tight for the instance of n points lying on a line.
- The calculation, when carried out with probability $p \in (0, 1)$ as a parameter, gives $|E_k| = O\left(\frac{n}{p(1-p)^k}\right)$. The value of p is then set to maximize the denominator. Roughly speaking, as the term $(1-p)^k = e^{-\Theta(pk)}$ decreases exponentially with p, it is best to set p so that $e^{-\Theta(pk)}$ is a constant—that is, $p = \Theta\left(\frac{1}{k}\right)$.

[2] A minor technical difference is that in the probabilistic version, the *expected size* is $\frac{n}{k+1}$.

As for the precise value of p that maximizes the denominator, since the derivative of $p\,(1-p)^k$ with respect to p is $(1-kp-p)\,(1-p)^{k-1}$, it can be verified that $p\,(1-p)^k$ is maximized at $p = \frac{1}{k+1}$.

This also makes sense intuitively: for each edge $\{p_i, p_j\} \in E_k$, the disk D_{ij} contains at most k other points of P and so picking each point with probability less than $\frac{1}{k}$ implies that, in expectation, D_{ij} will not contain any of these points.

- Other applications of this technique follow the same 'template'—pick a random sample and calculate the probability of some event due to it in two ways. The main technical work consists in finding good estimates for certain events; this is typically where a variety of other combinatorial and geometric ideas come into play.

1. Level Sets

Counting pairs is the oldest trick in combinatorics ... every time we count pairs, we learn something from it.

Gil Kalai

Our first application is a variant of the k-Delaunay graph problem. Given a finite set P of points in $\mathbb{R}^d$ and an integer $k \geq 1$, the objective is to upper bound the number of subsets of P of size at most k that are 'realizable' by geometric objects in $\mathbb{R}^d$. We first explain the problem for the case of disks in $\mathbb{R}^2$.
For a set P of n points in $\mathbb{R}^2$, define the set system

$$\mathcal{R}(P) = \left\{ D \cap P : \ D \text{ is a disk in } \mathbb{R}^2 \right\}.$$

We call $\mathcal{R}(P)$ the primal set system induced on P by disks. For any integer $k \geq 1$, let $\mathcal{R}_{=k}(P)$ be the sets of $\mathcal{R}(P)$ of size exactly k and let $\mathcal{R}_{\leq k}(P)$ be the sets of $\mathcal{R}(P)$ of size at most k. That is,

$$\mathcal{R}_{=k}(P) = \left\{ R \in \mathcal{R}(P) : \ |R| = k \right\} \quad \text{and} \quad \mathcal{R}_{\leq k}(P) = \left\{ R \in \mathcal{R}(P) : \ |R| \leq k \right\}.$$

The sets of $\mathcal{R}_{\leq k}(P)$ are called the ($\leq k$)-level sets, or simply ($\leq k$)-sets, of $\mathcal{R}(P)$.

> Observe that $\mathcal{R}_{\leq 2}(P)$—the subsets of P of size at most two that are induced by disks—consists of $O(n)$ sets: the sets of size 1 in $\mathcal{R}_{\leq 2}(P)$ are the points of P and the sets of size 2 are precisely the edges of the Delaunay graph of P. At the other end, $\mathcal{R}_{\leq n}(P)$ is just $\mathcal{R}(P)$, with size $O\left(n^3\right)$.

Our first main result of this section implies both of the above two cases.

Lemma 1.2. *Let P be a set of n points in $\mathbb{R}^2$ and let $\mathcal{R}(P)$ be the primal set system induced on P by disks in the plane. Then for any integer $k \geq 1$,*

$$\left|\mathcal{R}_{\leq k}(P)\right| = O\left(nk^2\right).$$

To simplify the presentation, we will assume that $|P| \geq 3$, and that P is in general position; in particular, no three points lie on a line and no four points lie on a circle.

To prove Lemma 1.2, we will first count a slightly different structure called *canonical disks*, which are disks that are 'fixed' by points of P on their boundary.

Definition 1.3. A canonical disk spanned by $Q \subseteq \mathbb{R}^2$ is a disk whose boundary contains three points of Q.
Furthermore, a canonical disk D spanned by Q is called an empty canonical disk if the interior of D contains no point of Q.

Let $\mathcal{T}(P)$ be the set of all $\binom{n}{3}$ unordered triples of points of P. For a triple $\{p,q,r\} \in \mathcal{T}(P)$, let D_{pqr} be the unique *open* disk whose boundary contains $\{p,q,r\}$; we say that D_{pqr} is spanned by $\{p,q,r\}$. For an integer $k \geq 0$, define the level sets

$$\mathcal{T}_{\leq k}(P) = \left\{ \{p,q,r\} \in \mathcal{T}(P) : \ \left|D_{pqr} \cap P\right| \leq k \right\}.$$

We first observe that the size of $\mathcal{R}_{\leq k}(P)$ is bounded, within a constant factor, by that of $\mathcal{T}_{\leq k}(P)$.

Claim 1.4. For any integer $k \geq 1$, $|\mathcal{R}_{\leq k}(P)| \leq 8 \cdot |\mathcal{T}_{\leq (k-1)}(P)|$.

Proof. Take any $R \in \mathcal{R}_{\leq k}(P)$ and let D be a disk realizing R; that is, $R = D \cap P$.

Now D can be scaled and translated—without any point of P 'crossing' the boundary of D—such that it contains three points of P, say $\{p, q, r\}$, on its boundary. Furthermore at least one of p, q or r belongs to R. See figure.

The interior of D_{pqr} contains at most $k-1$ points of P and so $\{p, q, r\} \in \mathcal{T}_{\leq (k-1)}(P)$. By slightly shifting and scaling D_{pqr}, for each of the 8 possible subsets of $\{p, q, r\}$, one can get a disk containing precisely that subset and all the points of P in the interior of D_{pqr}. One of these subsets is R, implying the claim.

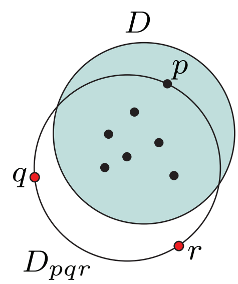

□

We remark here that the constant 8 can be improved with a more careful argument (see discussion).

Now the proof of Lemma 1.2 follows from Claim 1.4 and the following statement.

Lemma 1.5. *Let P be a set of n points in $\mathbb{R}^2$ and let $k \geq 0$ be an integer. Then*

$$|\mathcal{T}_{\leq k}(P)| = O\left(n(k+1)^2\right).$$

Proof. First we establish the case $k = 0$.

Claim 1.6. For any $S \subseteq P$, $|\mathcal{T}_{\leq 0}(S)| \leq 2|S|$.

Proof. $\mathcal{T}_{\leq 0}(S)$ consists of unordered triples of S whose corresponding open disks do not contain any point of S in their interior. If the disk D_{pqr}, spanned by $p, q, r \in S$, contains no point of S in its interior, then by slightly shifting D_{pqr}, it follows that each of the three edges $\{p, q\}$, $\{q, r\}$ and $\{p, r\}$ belong to the Delaunay graph of S. In particular, the triangle with vertices $\{p, q, r\}$ is a face of the Delaunay graph of S. Thus $|\mathcal{T}_{\leq 0}(S)|$ is upper bounded by the number of faces in a planar graph on $|S|$ vertices, which is $2|S| - 4$. □

Now consider the case $k \geq 1$. Construct a random sample S by picking each point of P independently with probability $p = \frac{1}{k+1}$.

We count the expected size of $\mathcal{T}_{\leq 0}(S)$ in two ways.

Upper bound: From Claim 1.6,

$$\mathrm{E}\left[\, |\mathcal{T}_{\leq 0}(S)| \,\right] \leq \mathrm{E}\left[2\,|S|\right] = 2\,np.$$

Lower bound: The key is the following observation:

a triple $\{p, q, r\} \in \mathcal{T}(P)$ is present in $\mathcal{T}_{\leq 0}(S)$ *if and only if* $\{p, q, r\} \subseteq S$ and none of the points in $D_{pqr} \cap P$ are picked in S.

As each point of P was picked independently, for any $\{p, q, r\} \in \mathcal{T}(P)$, we have

$$\Pr\left[\, \{p, q, r\} \in \mathcal{T}_{\leq 0}(S) \,\right] = p^3 \cdot (1-p)^{|D_{pqr} \cap P|}.$$

Therefore, by linearity of expectation,

$$\begin{aligned}
\mathrm{E}\left[\, |\mathcal{T}_{\leq 0}(S)| \,\right] &= \sum_{\{p,q,r\} \in \mathcal{T}(P)} \Pr\left[\, \{p, q, r\} \in \mathcal{T}_{\leq 0}(S) \,\right] \\
&\geq \sum_{\{p,q,r\} \in \mathcal{T}_{\leq k}(P)} \Pr\left[\, \{p, q, r\} \in \mathcal{T}_{\leq 0}(S) \,\right] \\
&= \sum_{\{p,q,r\} \in \mathcal{T}_{\leq k}(P)} p^3 \cdot (1-p)^{|D_{pqr} \cap P|} \\
&\geq \sum_{\{p,q,r\} \in \mathcal{T}_{\leq k}(P)} p^3 \cdot (1-p)^k = |\mathcal{T}_{\leq k}(P)| \cdot p^3 \cdot (1-p)^k.
\end{aligned}$$

Combining the upper and lower bounds,

$$|\mathcal{T}_{\leq k}(P)| \cdot p^3 \cdot (1-p)^k \;\leq\; \mathrm{E}\left[\, |\mathcal{T}_{\leq 0}(S)| \,\right] \;\leq\; 2\,np,$$

and hence $$|\mathcal{T}_{\leq k}(P)| \leq \frac{2\,n}{p^2 \cdot (1-p)^k} = \frac{2\,n\,(k+1)^2}{\left(1 - \frac{1}{k+1}\right)^k} \leq 2e\,n\,(k+1)^2,$$

where the last step follows from the fact that $\left(1 - \frac{1}{k+1}\right)^k \geq \frac{1}{e}$. □

We next prove a similar statement for set systems where the elements are geometric objects in $\mathbb{R}^d$ and the sets are induced by points in $\mathbb{R}^d$. We consider the case of disks in the plane.

Given a set $\mathcal{D} = \{D_1, \ldots, D_n\}$ of n distinct closed disks in $\mathbb{R}^2$, define the set system

$$\mathcal{R}(\mathcal{D}) = \left\{\mathcal{D}_p \;:\; p \in \mathbb{R}^2\right\}, \quad \text{where} \quad \mathcal{D}_p = \left\{D \in \mathcal{D} \;:\; D \ni p\right\}.$$

We call $\mathcal{R}(\mathcal{D})$ the dual set system induced on $\mathcal{D}$ by $\mathbb{R}^2$. Visually, each cell in the arrangement of $\mathcal{D}$ corresponds to a set in $\mathcal{R}(\mathcal{D})$ (note that different cells may correspond to the same subset).

For simplicity we will assume that $\mathcal{D}$ is in general position—in particular, the intersection of the boundaries of every pair of disks of $\mathcal{D}$ is either empty or consists of two distinct points and the intersection of the boundaries of any three disks of $\mathcal{D}$ is empty.

Our goal is to upper bound, for any integer $k \geq 1$, the size of $\mathcal{R}_{\leq k}(\mathcal{D})$. Our main result is the following.

LEMMA 1.7. *Let $\mathcal{D}$ be a set of n closed disks in $\mathbb{R}^2$ and let $\mathcal{R}(\mathcal{D})$ be the dual set system induced on $\mathcal{D}$. Then for any integer $k \geq 1$,*

$$|\mathcal{R}_{\leq k}(\mathcal{D})| = O(nk).$$

As earlier, it suffices to consider canonical sets, defined as follows. Let $\mathcal{V}(\mathcal{D})$ be the set of at most $2\binom{n}{2}$ points in $\mathbb{R}^2$ that are the intersections of boundaries of the disks of $\mathcal{D}$. For any integer $k \geq 0$, define

$$\mathcal{V}_{\leq k}(\mathcal{D}) = \big\{v \in \mathcal{V}(\mathcal{D}) : v \text{ is contained in the } \textit{interior} \text{ of at most } k \text{ disks of } \mathcal{D}\big\}.$$

Similarly one can define $\mathcal{V}_{=k}(\mathcal{D})$. See figure.

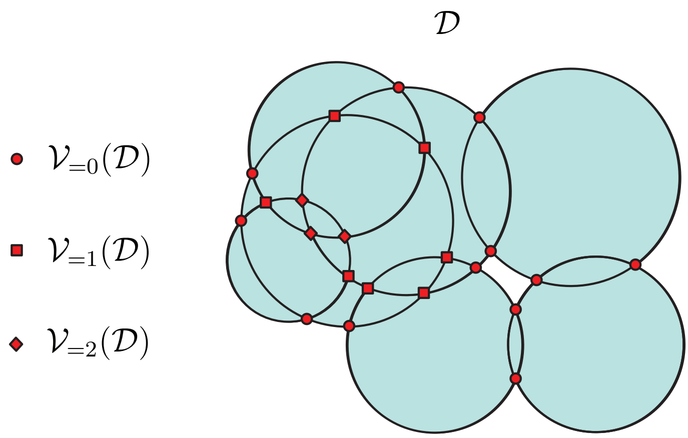

The proof of the following claim is easy and left to the reader.

CLAIM 1.8. For any integer $k \geq 1$, $|\mathcal{R}_{\leq k}(\mathcal{D})| \leq 4 \cdot |\mathcal{V}_{\leq (k-1)}(\mathcal{D})| + |\mathcal{D}|$.

Now the proof of Lemma 1.7 follows from Claim 1.8 and the following statement.

LEMMA 1.9. *For any integer $k \geq 0$, $|\mathcal{V}_{\leq k}(\mathcal{D})| = O(n(k+1))$.*

PROOF. As before, we first upper bound the size of $\mathcal{V}_{\leq 0}(\mathcal{D})$ and then use the averaging technique to upper bound the size of $\mathcal{V}_{\leq k}(\mathcal{D})$ for $k \geq 1$.

CLAIM 1.10. For any $S \subseteq \mathcal{D}$, $|\mathcal{V}_{\leq 0}(S)| \leq 6\,|S|$.

PROOF. Any $v \in \mathcal{V}_{\leq 0}(S)$ is an intersection point between the boundary of two disks of S and is not contained in the interior of any disk of S. Let $G = (S, E)$ be a graph where there is an edge between two disks of S if and only if a common intersection point of their boundaries belongs to $\mathcal{V}_{\leq 0}(S)$. We now show that G is planar, and so $|\mathcal{V}_{\leq 0}(S)| \leq 2|E| \leq 2(3|S| - 6) \leq 6|S|$.
We claim that the following is a plane drawing of G: draw each edge $\{D_i, D_j\} \in E$ as a line segment between the centers of D_i and D_j. Consider any two edges $\{D_i, D_j\}$, $\{D_k, D_l\}$ and let q_{ij}, q_{kl} be the two corresponding points in $\mathcal{V}_{\leq 0}(S)$. Let l be the bisector of q_{ij} and q_{kl}. As both D_i, D_j contain q_{ij} and do not contain q_{kl}, their centers lie on the side of l containing q_{ij}. Similarly the centers of D_k and D_l lie on the side of l containing q_{kl}. Thus the line segments corresponding to the edges $\{D_i, D_j\}$ and $\{D_k, D_l\}$ cannot intersect. $\square$

Now consider the case $k \geq 1$. Construct a random sample S by picking each disk of $\mathcal{D}$ independently with probability $p = \frac{1}{k+1}$.
We will count the expected size of $\mathcal{V}_{\leq 0}(S)$ in two ways.

Upper bound: From Claim 1.10,

$$\mathrm{E}\left[|\mathcal{V}_{\leq 0}(S)|\right] \leq \mathrm{E}[6|S|] = 6\,\mathrm{E}[|S|] = 6\,np.$$

Lower bound: This follows by considering the probability of each vertex in $\mathcal{V}_{\leq k}(\mathcal{D})$ ending up as a vertex of $\mathcal{V}_{\leq 0}(S)$. Let $v \in \mathcal{V}_{\leq k}(\mathcal{D})$ and let $D_i, D_j \in \mathcal{D}$ be the two disks such that v is an intersection point of the boundaries of D_i and D_j. Then

$v \in \mathcal{V}_{\leq 0}(S)$ if and only if $\{D_i, D_j\} \subseteq S$ and every disk of $\mathcal{D}$ containing v in its interior is not present in S.

As there are at most k such disks and each disk of $\mathcal{D}$ was picked independently,

$$\Pr\left[v \in \mathcal{V}_{\leq 0}(S)\right] \geq p^2\,(1-p)^k.$$

Therefore,

$$\begin{aligned}\mathrm{E}\left[|\mathcal{V}_{\leq 0}(S)|\right] &= \sum_{v \in \mathcal{V}(\mathcal{D})} \Pr\left[v \in \mathcal{V}_{\leq 0}(S)\right] \\ &\geq \sum_{v \in \mathcal{V}_{\leq k}(\mathcal{D})} \Pr\left[v \in \mathcal{V}_{\leq 0}(S)\right] \geq |\mathcal{V}_{\leq k}(\mathcal{D})| \cdot p^2 (1-p)^k.\end{aligned}$$

Combining the upper and lower bounds,

$$|\mathcal{V}_{\leq k}(\mathcal{D})| \cdot p^2 (1-p)^k \;\leq\; \mathrm{E}\left[|\mathcal{V}_{\leq 0}(S)|\right] \;\leq\; 6\,np,$$

$$\text{and hence} \quad |\mathcal{V}_{\leq k}(\mathcal{D})| \leq \frac{6\,n}{p\,(1-p)^k} = \frac{6\,n\,(k+1)}{\left(1-\frac{1}{k+1}\right)^k} \leq 6e\,n\,(k+1),$$

where the last step used the fact that $\left(1-\frac{1}{k+1}\right)^k \geq \frac{1}{e}$. □

Primal and dual set systems can be defined more generally:

DEFINITION 1.11. Given a set P of points in $\mathbb{R}^d$ and a (possibly infinite) family $\mathcal{R}$ of geometric objects in $\mathbb{R}^d$, the primal set system induced on P by $\mathcal{R}$ is

$$\left\{O \cap P \,:\, O \in \mathcal{R}\right\}.$$

DEFINITION 1.12. Given a set $\mathcal{R}$ of geometric objects in $\mathbb{R}^d$, the dual set system induced on $\mathcal{R}$ by $\mathbb{R}^d$ is defined as

$$\left\{\mathcal{R}_p \colon p \in \mathbb{R}^d\right\}, \quad \text{where} \quad \mathcal{R}_p = \left\{R \in \mathcal{R} \colon R \ni p\right\}.$$

We now conclude with the case of primal and dual set systems induced by half-spaces in $\mathbb{R}^d$.

Let P be a set of n points in general position in $\mathbb{R}^d$ and $\mathcal{R}(P)$ the primal set system induced on P by downward-facing half-spaces—that is, considering the x_d-axis as vertical, the half-spaces which contain the point that is the 'minus infinity' of the x_d

axis. It can be shown that $|\mathcal{R}(P)| = O(n^d)$. In fact for points in general position there is a precise bound independent of the structure of P (stated without proof):

$$|\mathcal{R}(P)| = \sum_{i=0}^{d} \binom{n}{i}. \tag{1.13}$$

Now let $\mathcal{T}(P)$ be all the $\binom{n}{d}$ subsets of P of size d. For each $e \in \mathcal{T}(P)$, let h_e^+ be the unique downward-facing *open* half-space whose bounding hyperplane contains e. For an integer $k \geq 0$, define the level sets

$$\mathcal{T}_{\leq k}(P) = \left\{ e \in \mathcal{T}(P) : \left| h_e^+ \cap P \right| \leq k \right\}.$$

As earlier, the size of $\mathcal{R}_{\leq k}(P)$ can be upper bounded, within a multiplicative factor, by that of $\mathcal{T}_{\leq k}(P)$.

To bound $|\mathcal{T}_{\leq k}(P)|$ we again first need a bound on $|\mathcal{T}_{\leq 0}(P)|$. Observe that $|\mathcal{T}_{\leq 0}(P)|$ is simply the number of facets on the lower convex-hull of P. It is well-known, to those who know it, that the Upper Bound Theorem for convex polytopes implies that this is at most $2\sum_{i=0}^{\lfloor \frac{d}{2} \rfloor} \binom{n}{i}$ (see discussion). Now the probabilistic averaging technique of this chapter together with this 0-th level bound implies the following (stated without proof).

Theorem 1.14. *Given a set P of n points in $\mathbb{R}^d$ and an integer $k \geq 0$,*

$$|\mathcal{T}_{\leq k}(P)| \leq 2 \left(\frac{e}{\lceil d/2 \rceil} \right)^{\lceil \frac{d}{2} \rceil} \binom{n}{\lfloor \frac{d}{2} \rfloor} \left(k + \left\lceil \frac{d}{2} \right\rceil \right)^{\lceil \frac{d}{2} \rceil}.$$

The above is $O\left(n^{\lfloor d/2 \rfloor} (k+1)^{\lceil d/2 \rceil}\right)$ when the dimension d is considered a constant.

> For $d = 3$, Theorem 1.14 gives a bound of $O(nk^2)$—the same bound, within a multiplicative constant, as the one of Lemma 1.5. This is not a coincidence: there exists a mapping of points in $\mathbb{R}^2$ to $\mathbb{R}^3$, the so-called 'paraboloid lift', with the property that subsets realized by intersection with disks in $\mathbb{R}^2$ can be realized by intersection with half-spaces in $\mathbb{R}^3$. Thus Theorem 1.14 for $d = 3$ implies Lemma 1.5.

For later use, it will be convenient to state Theorem 1.14 in the dual setting.

Definition 1.15. The level of a point $q \in \mathbb{R}^d$ with respect to a set $\mathcal{H}$ of hyperplanes in $\mathbb{R}^d$ is the number of hyperplanes of $\mathcal{H}$ lying strictly below q in the negative x_d direction; that is, the number of hyperplanes intersecting the ray

$$\left\{ q + \lambda (0, \ldots, 0, -1) : \lambda > 0 \right\}.$$

Given a set $\mathcal{H}$ of hyperplanes in $\mathbb{R}^d$ in general position, *a vertex* in the arrangement of $\mathcal{H}$ is a point lying in the intersection of some d hyperplanes of $\mathcal{H}$. Let $\mathcal{V}_{\leq k}(\mathcal{H})$ be the set of vertices of $\mathcal{H}$ of level at most k. Then by duality, Theorem 1.14 is equivalent to the following statement.

Theorem 1.16. *Given a set $\mathcal{H}$ of n hyperplanes in $\mathbb{R}^d$ and an integer $k \geq 0$,*

$$|\mathcal{V}_{\leq k}(\mathcal{H})| \leq 2 \left(\frac{e}{\lceil d/2 \rceil} \right)^{\lceil \frac{d}{2} \rceil} \binom{n}{\lfloor \frac{d}{2} \rfloor} \left(k + \left\lceil \frac{d}{2} \right\rceil \right)^{\lceil \frac{d}{2} \rceil}.$$

Bibliography and discussion. The influential probabilistic technique in this chapter was used in the seminal paper of Clarkson [**Cla87**], and then Clarkson and Shor [**CS89**], exactly for these problems of bounding the combinatorial complexity of configurations. Indeed, it is sometimes called the 'Clarkson-Shor technique' in the discrete and computational geometry literature. See [**APS08**, **Wag08**] for surveys on $(\leq k)$-sets for half-spaces and related set systems, where one can find information related to Theorem 1.14. Details on the Upper Bound Theorem for convex polytopes [**McM701a**] and related topics can be found in [**Zie95**]. See [**Mat99**, Section 3.1] for a discussion around Claim 1.4.

[APS08] P. K. Agarwal, J. Pach, and M. Sharir, *State of the union (of geometric objects)*, Surveys on discrete and computational geometry, Contemp. Math., vol. 453, Amer. Math. Soc., Providence, RI, 2008, pp. 9–48, DOI 10.1090/conm/453/08794. MR2405676

[CS89] K. L. Clarkson and P. W. Shor, *Applications of random sampling in computational geometry. II*, Discrete Comput. Geom. **4** (1989), no. 5, 387–421, DOI 10.1007/BF02187740. MR1014736

[Cla87] K. L. Clarkson, *New applications of random sampling in computational geometry*, Discrete Comput. Geom. **2** (1987), no. 2, 195–222, DOI 10.1007/BF02187879. MR884226

[Mat99] J. Matoušek. *Geometric Discrepancy: An Illustrated Guide*. Springer, 1999.

[McM701a] P. McMullen, *The maximum numbers of faces of a convex polytope*, Mathematika **17** (1970), 179–184, DOI 10.1112/S0025579300002850. MR283691

[Wag08] U. Wagner, *k-sets and k-facets*, Surveys on discrete and computational geometry, Contemp. Math., vol. 453, Amer. Math. Soc., Providence, RI, 2008, pp. 443–513, DOI 10.1090/conm/453/08810. MR2405692

[Zie95] G. M. Ziegler, *Lectures on polytopes*, Graduate Texts in Mathematics, vol. 152, Springer-Verlag, New York, 1995, DOI 10.1007/978-1-4613-8431-1. MR1311028

2. Concentration Bounds for Sums of Bernoulli Variables

I thought it was a rather trivial lemma, but many things are only trivial once you know them.

Herman Chernoff

We present an application of the probabilistic technique to computing tail bounds of some common probability distributions. That is, we would like to upper bound the probability that a random variable gets a value far from its expectation. This is a basic technical ingredient in nearly all the constructions and methods that will be seen later.

The setting is the following.

> Let $I = \{1, 2, \ldots, n\}$ be a set of n elements from which we will pick a random sample. We aim to pick np elements of I, for a given parameter $p \in [0, 1]$.

The 0-1 valued random variable X_i will be used to indicate whether $i \in I$ is picked in our random sample. Our goal is to estimate the probability that the sum of these n variables, $X = \sum_{i=1}^{n} X_i$, falls far from its expectation $\mathrm{E}[X]$. More precisely, for any $\delta \geq 0$, we are interested in bounding $\Pr\big[X \geq (1+\delta)\,\mathrm{E}[X]\big]$ and $\Pr\big[X \leq (1-\delta)\,\mathrm{E}[X]\big]$.

In fact, we consider the more general case where for a fixed set $J \subseteq I$ with $X_J = \sum_{j \in J} X_j$, we are interested in upper bounds on

$$\Pr\big[X_J \geq (1+\delta) \cdot \mathrm{E}[X_J]\big] \qquad \text{and} \qquad \Pr\big[X_J \leq (1-\delta) \cdot \mathrm{E}[X_J]\big].$$

There are several natural ways to pick a random sample from I. Two basic ones, given a parameter p, are the following.

Binomial distribution: Pick each element of I *independently* with probability p. That is, let $X_1, \ldots, X_n$ be n independent 0-1 random variables where

$$X_i = \begin{cases} 1 & \text{with probability } p, \\ 0 & \text{otherwise.} \end{cases}$$

For any $J \subseteq I$, we have

$$\mathrm{E}[X_J] = \mathrm{E}\left[\sum_{j \in J} X_j\right] = \sum_{j \in J} \mathrm{E}[X_j] = \sum_{j \in J} \Pr[X_j = 1] = |J|\,p.$$

One can write the exact equation for the tail bounds using the fact that the value of each X_i was set independently:

$$\text{(1.17)} \qquad \Pr\big[X_J \geq (1+\delta) \cdot |J|\,p\big] = \sum_{i=\lceil (1+\delta) \cdot |J| p \rceil}^{|J|} \Pr[X_J = i]$$
$$= \sum_{i=\lceil (1+\delta) \cdot |J| p \rceil}^{|J|} \binom{|J|}{i} p^i (1-p)^{|J|-i}.$$

As there is no closed-form formula for this, several methods have been proposed to estimate the right-hand side of the above expression (see discussion).

The fact that $\{X_1, \ldots, X_n\}$ are independent has two advantages: first it makes calculations easier and second, for any $J \subseteq I$ the induced probability distribution on X_J remains the same (that is, each element of J is picked independently with probability p). On the other hand, the number of elements X is not fixed and is a random variable with expectation np.

Sampling without replacement: A second natural way to sample is to choose, out of all $\binom{n}{np}$ np-sized subsets of I, one uniformly at random (assume that np is an integer). This then sets the values of $X_1, \ldots, X_n$, with $\sum_i X_i$ being equal to np. Note that for any $J \subseteq I$, $\mathrm{E}[X_J] = |J|\, p$ since for any i,

$$\Pr[X_i = 1] = \frac{\binom{n-1}{np-1}}{\binom{n}{np}} = \frac{(n-1)!}{(np-1)!(n-np)!} \cdot \frac{(np)!(n-np)!}{n!} = \frac{np}{n} = p.$$

More generally, for any $J \subseteq I$, letting $t = |J|$, the probability that $X_j = 1$ for all $j \in J$, can be upper bounded as

$$\Pr\left[\left(\prod_{j \in J} X_j\right) = 1\right] = \frac{\binom{n-t}{np-t}}{\binom{n}{np}} = \frac{(n-t)!}{(np-t)!} \cdot \frac{(np)!}{n!} = \frac{np}{n}\,\frac{np-1}{n-1} \cdots \frac{np-t+1}{n-t+1} \le p^t, \tag{1.18}$$

since each term $\frac{np-i}{n-i} \le p$ for $p \le 1$. Similarly, the probability that $X_j = 0$ for all $j \in J$, can be upper bounded as

$$\Pr\left[\left(\prod_{j \in J} (1 - X_j)\right) = 1\right] = \frac{\binom{n-t}{np}}{\binom{n}{np}} = \frac{(n-t)!}{(n-t-np)!} \cdot \frac{(n-np)!}{n!} = \frac{n-np}{n}\,\frac{n-np-1}{n-1} \cdots \frac{n-np-t+1}{n-t+1} \le (1-p)^t\,, \tag{1.19}$$

since each term $\frac{n-np-i}{n-i} \le (1-p)$ for $i \ge 0$.
We can again write the precise equation for the tail bounds for any $J \subseteq I$:

$$\Pr\left[X_J \ge (1+\delta)\cdot|J|\,p\right] = \sum_{i=\lceil(1+\delta)\cdot|J|p\rceil}^{|J|} \Pr[X_J = i] = \sum_{i=\lceil(1+\delta)\cdot|J|p\rceil}^{|J|} \frac{\binom{|J|}{i}\cdot\binom{n-|J|}{np-i}}{\binom{n}{np}}.$$

The advantage of this distribution is that $X = np$ always; however the variables $\{X_1, \ldots, X_n\}$ are no longer independent. Consequently, for a $J \subseteq I$, the induced probability distribution on X_J is *not* the one where a $(|J|p)$-sized subset of J is chosen uniformly at random from the set of all $(|J|p)$-sized subsets of J.

The variables $X_1, \ldots, X_n$ are an example of *negatively associated* random variables. We note that in this case the tail bounds are even better—that is, more sharply concentrated around the expectation—than for binomial distribution. Intuitively, for any $i, j \in I$, the fact that $X_i = 1$ makes it *less* likely that $X_j = 1$ and the fact that $X_i = 0$ makes it *more* likely that $X_j = 1$.

Formally, if $p \in (0,1)$,

$$\Pr\left[X_j = 1 \mid X_i = 1\right] = \frac{\binom{n-2}{np-2}}{\binom{n-1}{np-1}} = \frac{np-1}{n-1} < p, \qquad \text{and}$$

$$\Pr\left[X_j = 1 \mid X_i = 0\right] = \frac{\binom{n-2}{np-1}}{\binom{n-1}{np}} = \frac{np}{n-1} > p.$$

Our main theorem, a multiplicative version of a tail bound for negatively associated random variables, is the following.

Theorem 1.20. *Let $X_1, \ldots, X_n$ be n indicator random variables and let $\delta > 0$ be a given parameter. Set $X = \sum_{i=1}^{n} X_i$.*

(1) *Let $p_1, \ldots, p_n$ be reals in $[0,1]$, $0 < \sum_{i=1}^{n} p_i < n$, such that*

$$\text{for any } I' \subseteq [n] \qquad \Pr\left[\left(\prod_{i \in I'} X_i\right) = 1\right] \le \prod_{i \in I'} p_i.$$

Let $\tilde{p} = \frac{\sum_i p_i}{n}$. Then

$$\Pr\left[X \ge (1+\delta)\, n\tilde{p}\right] \le \left(\frac{\left(1 - \frac{\tilde{p}\delta}{1-\tilde{p}}\right)^{(1+\delta)\tilde{p}-1}}{(1+\delta)^{(1+\delta)\tilde{p}}}\right)^{n}. \tag{1.21}$$

The above expression can be simplified to give

$$\Pr\left[X \ge (1+\delta)\, n\tilde{p}\right] \le e^{-\frac{\delta^2}{2+\delta} n\tilde{p}}.$$

(2) *Let $r_1, \ldots, r_n$ be reals in $[0,1]$, $0 < \sum_{i=1}^{n} r_i < n$, such that*

$$\text{for any } I' \subseteq [n] \qquad \Pr\left[\left(\prod_{i \in I'} (1 - X_i)\right) = 1\right] \le \prod_{i \in I'} (1 - r_i).$$

Let $\tilde{r} = \frac{\sum_i r_i}{n}$. Then

$$\Pr\left[X \le (1-\delta)\, n\tilde{r}\right] \le \left(\frac{\left(1 + \frac{\delta\tilde{r}}{1-\tilde{r}}\right)^{-1+\tilde{r}-\delta\tilde{r}}}{(1-\delta)^{\tilde{r}(1-\delta)}}\right)^{n}. \tag{1.22}$$

The above expression can be simplified to give

$$\Pr\left[X \le (1-\delta)\, n\tilde{r}\right] \le e^{-\frac{\delta^2}{2} n\tilde{r}}.$$

We remark that the two preconditions of the above theorem imply that for any variable X_i, we have

$$\Pr\left[X_i = 1\right] \le p_i, \quad \text{and}$$

$$\Pr\left[(1 - X_i) = 1\right] \le 1 - r_i \quad \text{or equivalently,} \qquad \Pr\left[X_i = 1\right] \ge r_i.$$

Thus the variables p_i and r_i are upper and lower bounds on the probability that $X_i = 1$, and therefore $n\tilde{r} \le \mathrm{E}\left[\sum_i X_i\right] \le n\tilde{p}$.

The above theorem applies to *both* the two earlier distributions—binomial and sampling without replacement—as it avoids using independence of the X_i's and instead uses an upper bound on $\Pr\left[\left(\prod_{i\in J} X_i\right) = 1\right]$ and $\Pr\left[\left(\prod_{i\in J}(1-X_i)\right) = 1\right]$ for every $J \subseteq I$.

Binomial distribution: Theorem 1.20 and independence implies the following.

COROLLARY 1.23. *Let $I = \{1,\ldots,n\}$ and $p_1,\ldots,p_n \in [0,1]$ be given parameters. Let $R \subseteq I$ be a random sample constructed by picking each $i \in I$ independently with probability p_i. Then for a fixed $J \subseteq I$ and $\delta > 0$,*

$$\begin{aligned}\Pr\Big[\,|J\cap R| \geq (1+\delta)\,\mathrm{E}\left[|J\cap R|\right]\Big] &= \Pr\Big[\,|J\cap R| \geq (1+\delta)\sum_{j\in J} p_j\Big] \\ &\leq e^{-\frac{\delta^2}{2+\delta}\sum_{j\in J} p_j}, \\ \Pr\Big[\,|J\cap R| \leq (1-\delta)\,\mathrm{E}\left[|J\cap R|\right]\Big] &= \Pr\Big[\,|J\cap R| \leq (1-\delta)\sum_{j\in J} p_j\Big] \\ &\leq e^{-\frac{\delta^2}{2}\sum_{j\in J} p_j}.\end{aligned}$$

In particular,

$$\Pr\left[\,|J\cap R| \geq (1+\delta)\sum_{j\in J} p_j \quad \bigcup \quad |J\cap R| \leq (1-\delta)\sum_{j\in J} p_j\right] \leq 2\,e^{-\frac{\delta^2}{2+\delta}\sum_{j\in J} p_j}.$$

Sampling without replacement: Theorem 1.20 together with Equations (1.18) and (1.19) implies the following.

COROLLARY 1.24. *Let $I = \{1,\ldots,n\}$ and $t \in [n]$ be a given parameter. Let $R \subseteq I$ be a random sample of size t chosen uniformly from all $\binom{n}{t}$ t-sized subsets of I. Then for any fixed $J \subseteq I$ and $\delta > 0$,*

$$\begin{aligned}\Pr\left[|J\cap R| \geq (1+\delta)\,|J|\frac{t}{n}\right] &\leq e^{-\frac{\delta^2}{2+\delta}|J|\frac{t}{n}}, \\ \Pr\left[|J\cap R| \leq (1-\delta)\,|J|\frac{t}{n}\right] &\leq e^{-\frac{\delta^2}{2}|J|\frac{t}{n}}.\end{aligned}$$

In particular,

$$\Pr\left[\,|J\cap R| \geq (1+\delta)\,|J|\frac{t}{n} \quad \bigcup \quad |J\cap R| \leq (1-\delta)\,|J|\frac{t}{n}\right] \leq 2\,e^{-\frac{\delta^2}{2+\delta}|J|\frac{t}{n}}.$$

We remark here that these bounds are tight within constant factors in the exponent for certain ranges of δ. Here is one lower bound (stated without proof; see discussion).

THEOREM 1.25. *Let $I = \{1,\ldots,n\}$ and $p \in \left(0,\frac{1}{2}\right]$. Let $R \subseteq I$ be a random sample constructed by picking each $i \in I$ independently with probability p. Then for $\delta \in \left[\sqrt{\frac{3}{np}},\frac{1}{2}\right]$,*

$$\begin{aligned}\Pr\Big[\,|R| \geq (1+\delta)\,np\Big] &\geq e^{-9\delta^2 np}, \\ \Pr\Big[\,|R| \leq (1-\delta)\,np\Big] &\geq e^{-9\delta^2 np}.\end{aligned}$$

Overview of ideas. The proof of Theorem 1.20 will use our probabilistic averaging technique. That is, we will take a random sample of $\{1, 2, \ldots, n\}$ and calculate the probability of a carefully chosen event due to it in two ways.

At first glance, it might seem odd to estimate the probability of a random event—in our case the tail bounds on X—by taking *another* random sample! However it is a mistake to confuse these two separate probability distributions, with very different purposes—one is part of the input problem and the other is part of the averaging proof technique.

Perhaps a more modular way to think about this is to consider the quantity we are bounding—$\Pr[X \geq (1+\delta)\, n\tilde{p}]$—*combinatorially*: the support of the probability distribution consists of 2^n binary strings corresponding to all possible assignments of the 0-1 variables $X_1, \ldots, X_n$. Each string $s \in \{0,1\}^n$ has some probability, say $w(s)$, of being chosen.

> Let $|s|$ denote the number of 1's in s. The precise value of $w(s)$ depends on the probability distribution. For example, when $p_1 = \cdots = p_n = p$ and where np is an integer, $w(s) = p^{|s|}(1-p)^{n-|s|}$ for the binomial distribution. Similarly $w(s) = 1/\binom{n}{np}$ if $|s| = np$, and 0 otherwise, for the sampling without replacement distribution.

Then our goal is to upper bound the combinatorial quantity

$$\sum_{\substack{s \in \{0,1\}^n \\ |s| \geq (1+\delta) n\tilde{p}}} w(s).$$

Seen this way, it is similar to the earlier use of the probabilistic averaging technique to upper bound the sizes of level sets, with one difference being that earlier we were bounding the cardinality instead of a weighted sum.

As a warm-up, we first prove the following weaker bound, called *Markov's inequality*, under the conditions of Theorem 1.20:

$$\Pr\left[X \geq (1+\delta)\, n\tilde{p}\right] \leq \frac{1}{1+\delta}. \tag{1.26}$$

While Markov's inequality has an even simpler direct proof (furthermore, Markov's inequality holds for *any* positive random variable X for which $\mathrm{E}[X]$ exists, with $\mathrm{E}[X]$ replacing $n\tilde{p}$ in the stated bound. That is, X need not be the sum of n indicator variables), the following proof is an easy natural application of the probabilistic averaging technique and gives insight into the proof of Theorem 1.20.

> Let S be a random sample of the index set $I = \{1, 2, \ldots, n\}$ where each index is picked independently with probability q. Note that S is independent of the X_i variables.
>
> We count the following quantity in two ways:
>
> $$\mathrm{E}\left[|S_1|\right], \quad \text{where} \quad S_1 = \{i \in S \colon X_i = 1\}.$$
>
> That is, the expected number of indices $i \in S$ for which $X_i = 1$.

Upper bound: Using linearity of expectation and the fact that we have $\Pr[X_i = 1] \le p_i$,

$$\mathrm{E}[|S_1|] = \sum_{i=1}^{n} \Pr[i \in S \text{ and } X_i = 1] \le \sum_{i=1}^{n} p_i q = n\tilde{p}\, q.$$

The last step used the fact that S and X are independent.

Lower bound: Consider the elements of the event space for the variable $X = X_1 + \cdots + X_n$ for which $X \ge (1+\delta)\, n\tilde{p}$. Note that for each event $\{X_1, \ldots, X_n\}$ with $X = k$, the expected number of indices $i \in S$ with $X_i = 1$ is precisely kq. Thus we have

$$\mathrm{E}\left[\ |S_1| \ \Big| \ X \ge (1+\delta)\, n\tilde{p}\ \right] \ge (1+\delta)\, n\tilde{p}\, q.$$

Summing up over all events,

$$\begin{aligned}\mathrm{E}[|S_1|] = &\ \mathrm{E}\left[|S_1| \ \big|\ X \ge (1+\delta)n\tilde{p}\right] \cdot \Pr[X \ge (1+\delta)n\tilde{p}] \ + \\ &\ \mathrm{E}\left[|S_1| \ \big|\ X < (1+\delta)n\tilde{p}\right] \cdot \Pr[X < (1+\delta)n\tilde{p}] \\ \ge &\ \mathrm{E}\left[|S_1| \ \big|\ X \ge (1+\delta)n\tilde{p}\right] \cdot \Pr[X \ge (1+\delta)n\tilde{p}] \\ \ge &\ (1+\delta)\, n\tilde{p}\, q \cdot \Pr[X \ge (1+\delta)\, n\tilde{p}]\,.\end{aligned}$$

Putting the upper and lower bounds together,

$$(1+\delta)\, n\tilde{p}\, q \cdot \ \Pr[X \ge (1+\delta)\, n\tilde{p}] \ \le \ \mathrm{E}[|S_1|] \ \le \ n\tilde{p}\, q,$$

$$\text{and hence} \qquad \Pr[X \ge (1+\delta)\, n\tilde{p}] \le \frac{1}{1+\delta}.$$

An astute reader will notice that the proof above is needlessly complicated, as the parameter q does not play any role: the dependence on q is linear in both the upper and lower bounds and thus cancels out. Setting $S = I$ (i.e., $q = 1$) gives the standard proof of Markov's inequality. This will not remain the case for the proof of the main theorem, to which we turn to next.

We now prove our main theorem.

Proof of Theorem 1.20. As before, let S be a random sample where each element in $\{1, 2, \ldots, n\}$ is picked independently with probability q.

We count the following quantity in two ways:

$\Pr\left[\prod_{i \in S} X_i = 1\right]$.

That is, the probability that for *each* index $i \in S$, $X_i = 1$.

Note that this probability is over both the choice of S and the choice of X. Furthermore S and X are independent.

Upper bound. It will be instructive to consider it in three, progressively more general, scenarios:

- **each $X_i = 1$ independently with probability p_i:** Then we have

$$\begin{aligned}\Pr\left[\prod_{i\in S} X_i = 1\right] &= \prod_{i=1}^{n}\Big(1-\Pr\left[i\in S \text{ and } X_i = 0\right]\Big)\\ &= \prod_{i=1}^{n}\left(1-q\left(1-p_i\right)\right) \le \left(\frac{\sum_{i=1}^{n}\left(1-q\left(1-p_i\right)\right)}{n}\right)^n\\ &= \left(\frac{n-nq+q\sum_{i=1}^{n}p_i}{n}\right)^n = \left(q\tilde{p}+1-q\right)^n,\end{aligned}$$

where the third step uses the inequality of arithmetic and geometric means, that $\prod_{i=1}^{n} a_i \le \left(\frac{\sum_{i=1}^{n} a_i}{n}\right)^n$ for any non-negative reals $a_1,\ldots,a_n$.

- **$p_1 = \cdots = p_n = p$:** Then $\tilde{p} = p$ and so

$$\begin{aligned}\Pr\left[\prod_{i\in S} X_i = 1\right] &= \sum_{Q\subseteq[n]}\Pr\left[S = Q \text{ and } \prod_{i\in Q} X_i = 1\right]\\ &= \sum_{Q\subseteq[n]}\Pr\left[S=Q\right]\cdot\Pr\left[\prod_{i\in Q}X_i = 1\right] \quad \Big(S,\ X \text{ are independent}\Big)\\ &\le \sum_{i=0}^{n}\binom{n}{i}q^i(1-q)^{n-i}\cdot p^i \quad \Big(\text{by input assumption}\Big)\\ &= (qp+1-q)^n \quad \Big(\text{by the binomial theorem}\Big).\end{aligned}$$

- **the general case:**

$$\begin{aligned}\Pr\left[\prod_{i\in S} X_i = 1\right] &= \sum_{Q\subseteq[n]}\Pr\left[S=Q\right]\cdot\Pr\left[\prod_{i\in Q}X_i = 1\right]\\ &\le \sum_{Q\subseteq[n]} q^{|Q|}(1-q)^{n-|Q|}\cdot\prod_{i\in Q}p_i \quad \Big(\text{by input assumption}\Big)\\ &= (1-q)^n\sum_{Q\subseteq[n]}\prod_{i\in Q}\frac{qp_i}{1-q} = (1-q)^n\prod_{i=1}^{n}\left(1+\frac{qp_i}{1-q}\right),\end{aligned}$$

where the last step uses the fact that $\prod_{i=1}^{n}(1+a_i) = \sum_{Q\subseteq[n]}\prod_{i\in Q}a_i$ (each term in the L.H.S. of this expression, when opened up, corresponds to a choice of either 1 or a from each of the n product terms). Continuing,

$$\begin{aligned}&= (1-q)^n\prod_{i=1}^{n}\left(\frac{qp_i+1-q}{1-q}\right) = \prod_{i=1}^{n}\left(qp_i+1-q\right)\\ &\le \left(q\tilde{p}+1-q\right)^n \quad \Big(\text{as earlier}\Big).\end{aligned}$$

Lower bound. Consider the elements of the event space of $X = X_1+\cdots+X_n$ for which $X \ge (1+\delta)\,n\tilde{p}$. Note that for each instance of $\{X_1,\ldots,X_n\}$ with $X = k$, the probability that for *each* index $i\in S$ we have $X_i = 1$ is exactly $(1-q)^{n-k}$. In our case $k \geq (1+\delta)\, n\tilde{p}$ and since $(1-q)^{n-k}$ is mononotically increasing with k, we have

$$\Pr\left[\prod_{i\in S} X_i = 1 \;\middle|\; X \geq (1+\delta)\, n\tilde{p}\right] \geq (1-q)^{n-(1+\delta)n\tilde{p}}\,.$$

Summing up over all events,

$$\begin{aligned}
\Pr\left[\prod_{i\in S} X_i = 1\right] &= \Pr\left[\prod_{i\in S} X_i = 1 \;\middle|\; X \geq (1+\delta)\, n\tilde{p}\right] \cdot \Pr[X \geq (1+\delta)n\tilde{p}] \;\; + \\
&\quad \Pr\left[\prod_{i\in S} X_i = 1 \;\middle|\; X < (1+\delta)\, n\tilde{p}\right] \cdot \Pr[X < (1+\delta)n\tilde{p}] \\
&\geq \Pr\left[\prod_{i\in S} X_i = 1 \;\middle|\; X \geq (1+\delta)\, n\tilde{p}\right] \cdot \Pr[X \geq (1+\delta)n\tilde{p}] \\
&\geq (1-q)^{n-(1+\delta)n\tilde{p}} \cdot \Pr\left[X \geq (1+\delta)\, n\tilde{p}\right].
\end{aligned}$$

Combining the upper and lower bounds,

$$(1-q)^{n-(1+\delta)n\tilde{p}} \cdot \Pr\left[X \geq (1+\delta)\, n\tilde{p}\right] \;\leq\; \Pr\left[\prod_{i\in S} X_i = 1\right] \;\leq\; \left(1 - q\,(1-\tilde{p})\right)^n$$

$$\Longrightarrow \Pr\left[X \geq (1+\delta)\, n\tilde{p}\right] \leq \left(\frac{1 - q\,(1-\tilde{p})}{(1-q)^{1-(1+\delta)\tilde{p}}}\right)^n.$$

To minimize the R.H.S. of the above expression[3], we set $q = \frac{\delta}{(1-\tilde{p})(1+\delta)}$. Then

$$\begin{aligned}
\Pr\left[X \geq (1+\delta)\, n\tilde{p}\right] &\leq \left(\frac{1 - \frac{\delta}{(1-\tilde{p})(1+\delta)}(1-\tilde{p})}{\left(1 - \frac{\delta}{(1-\tilde{p})(1+\delta)}\right)^{1-(1+\delta)\tilde{p}}}\right)^n \\
&= \left(\frac{\frac{1}{1+\delta}}{\left(\frac{1-\tilde{p}-\tilde{p}\delta}{(1-\tilde{p})(1+\delta)}\right)^{1-(1+\delta)\tilde{p}}}\right)^n = \left(\frac{\left(\frac{1-\tilde{p}-\tilde{p}\delta}{1-\tilde{p}}\right)^{(1+\delta)\tilde{p}-1}}{(1+\delta)^{(1+\delta)\tilde{p}}}\right)^n \\
&= \left(\frac{\left(1 - \frac{\tilde{p}\delta}{1-\tilde{p}}\right)^{(1+\delta)\tilde{p}-1}}{(1+\delta)^{(1+\delta)\tilde{p}}}\right)^n,
\end{aligned}$$

getting the required bound.

The other direction—an upper bound on the probability that the number of 1's in X is at most $(1-\delta)n\tilde{r}$—is equivalent to upper bounding the probability that the number of 0's in X is at least

$$n - (1-\delta)\, n\tilde{r} = \left(\frac{1-(1-\delta)\,\tilde{r}}{(1-\tilde{r})}\right) n\;(1-\tilde{r}) = \left(1 + \frac{\delta\tilde{r}}{1-\tilde{r}}\right) n\,(1-\tilde{r})\,.$$

---

[3]The partial derivative w.r.t. q is $\frac{\left(n\tilde{p}\right)\left((1+\delta)(\tilde{p}-1)q+\delta\right)\left(((\tilde{p}-1)q+1)(1-q)^{\delta\tilde{p}+\tilde{p}-1}\right)^n}{(q-1)((\tilde{p}-1)q+1)}$.

Set $Y_i = 1 - X_i$ for $i = 1, \ldots, n$, and let $Y = \sum_i Y_i$. That is, $Y = n - X$ is a random variable denoting the number of 0's in X. Then

$$\Pr\left[X \le (1-\delta)\, n\tilde{r}\right] = \Pr\left[Y \ge \left(1 + \frac{\delta\tilde{r}}{1-\tilde{r}}\right) n\,(1-\tilde{r})\right].$$

Thus we can apply the previous bound on the variable Y_i's, now with probabilities $(1-r_i)$ instead of p_i. We have $\frac{\sum_{i=1}^{n}(1-r_i)}{n} = (1-\tilde{r})$ and so from Equation (1.21) with $\delta' = \frac{\delta\tilde{r}}{1-\tilde{r}}$,

$$\begin{aligned}
\Pr\left[Y \ge (1+\delta')\, n\,(1-\tilde{r})\right] &\le \left(\frac{\left(1 - \frac{(1-\tilde{r})\delta'}{1-(1-\tilde{r})}\right)^{(1+\delta')(1-\tilde{r})-1}}{(1+\delta')^{(1+\delta')(1-\tilde{r})}}\right)^{n} \\
&= \left(\frac{\left(1 - \frac{\delta\tilde{r}}{\tilde{r}}\right)^{(1+\frac{\delta\tilde{r}}{1-\tilde{r}})(1-\tilde{r})-1}}{\left(1 + \frac{\delta\tilde{r}}{1-\tilde{r}}\right)^{(1+\frac{\delta\tilde{r}}{1-\tilde{r}})(1-\tilde{r})}}\right)^{n} \\
&= \left(\frac{(1-\delta)^{-\tilde{r}(1-\delta)}}{\left(1+\frac{\delta\tilde{r}}{1-\tilde{r}}\right)^{1-\tilde{r}+\delta\tilde{r}}}\right)^{n} = \left(\frac{\left(1+\frac{\delta\tilde{r}}{1-\tilde{r}}\right)^{-1+\tilde{r}-\delta\tilde{r}}}{(1-\delta)^{\tilde{r}(1-\delta)}}\right)^{n},
\end{aligned}$$

getting the required bound.

The simplifications of these expressions are covered in many places and we refer the reader to existing literature on this (see discussion). □

Bibliography and discussion. The proof given of these tail bounds is from [**IK10**], while Theorem 1.25 is from [**KY15**]. A nice exposition of several proofs of tail bounds similar to the one presented here (Chernoff's bound, Bernstein's inequality, Hoeffding's extension) together with the details of simplification of the expressions in Theorem 1.20 can be found in [**Mul18**] (see also [**DP09**]). A discussion on the differences between sampling with and without replacement can be found in [**FK15**, Section 21.5]. A discussion on the asymmetry between the upper and lower tail bounds in Theorem 1.20 can be found in [**AS16**, Appendix A]. Some approximations for sums of binomial coefficients (such as those of Equation (1.17)) can be found in [**GKP94**, Chapter 5] (see also [**Spi19**]). Many other concentration inequalities can be found in the text [**BLM13**].

[AS16] N. Alon and J. H. Spencer, *The probabilistic method*, 4th ed., Wiley Series in Discrete Mathematics and Optimization, John Wiley & Sons, Inc., Hoboken, NJ, 2016. MR3524748

[BLM13] S. Boucheron, G. Lugosi, and P. Massart, *Concentration inequalities*, Oxford University Press, Oxford, 2013. A nonasymptotic theory of independence; With a foreword by Michel Ledoux, DOI 10.1093/acprof:oso/9780199535255.001.0001. MR3185193

[DP09] D. P. Dubhashi and A. Panconesi, *Concentration of measure for the analysis of randomized algorithms*, Cambridge University Press, Cambridge, 2009, DOI 10.1017/CBO9780511581274. MR2547432

[FK15] A. Frieze and M. Karoński, *Introduction to random graphs*, Cambridge University Press, Cambridge, 2016, DOI 10.1017/CBO9781316339831. MR3675279

[GKP94] R. L. Graham, D. E. Knuth, and O. Patashnik, *Concrete mathematics*, 2nd ed., Addison-Wesley Publishing Company, Reading, MA, 1994. A foundation for computer science. MR1397498

[IK10] R. Impagliazzo and V. Kabanets, *Constructive proofs of concentration bounds*, Approximation, randomization, and combinatorial optimization, Lecture Notes in Comput. Sci., vol. 6302, Springer, Berlin, 2010, pp. 617–631, DOI 10.1007/978-3-642-15369-3_46. MR2755867

[KY15] P. Klein and N. E. Young, *On the number of iterations for Dantzig-Wolfe optimization and packing-covering approximation algorithms*, SIAM J. Comput. **44** (2015), no. 4, 1154–1172, DOI 10.1137/12087222X. MR3390154

[Mul18] W. Mulzer, *Five proofs of Chernoff's bound with applications*, Bull. Eur. Assoc. Theor. Comput. Sci. EATCS **124** (2018), 59–76. MR3793013

[Spi19] M. Z. Spivey, *The art of proving binomial identities*, Discrete Mathematics and its Applications (Boca Raton), CRC Press, Boca Raton, FL, 2019, DOI 10.1201/9781351215824. MR3931743

CHAPTER 6

Epsilon-Nets: Combinatorial Bounds

The initial study of ϵ-nets in the field of computational geometry started in the 1980s with the work of Clarkson who showed the existence of ϵ-nets of size $O\left(\frac{1}{\epsilon}\log\frac{1}{\epsilon}\right)$ for specific geometric set systems. He was mainly interested in their algorithmic applications, such as nearest-neighbor queries for a set of points in Euclidean space. Chapter 2 is largely based on his work.

Independently, Haussler and Welzl showed similar bounds in a purely abstract setting, needing just that the given set system has bounded VC-dimension. In fact, what was needed, given a set system $(X, \mathcal{F})$, was the property that there exist an absolute constant d such that

$$\text{for all } Y \subseteq X,\ |\mathcal{F}|_Y| = O\left(|Y|^d\right) \quad \text{(see Lemma 4.3).}$$

They showed, surprisingly, that this is already a sufficient condition for the existence of small ϵ-nets; the following will be the first theorem of this chapter.

Theorem 6.1. *Let $(X, \mathcal{F})$ be a finite set system, $d \geq 1$ an integer such that* VC-dim$(\mathcal{F}) \leq d$, *and $\epsilon \in \left(0, \frac{1}{2}\right)$ a given parameter. Let N be a uniform random sample of X of size $t = \left\lceil \frac{56\,d}{\epsilon} \ln \frac{1}{\epsilon} \right\rceil$. Then N is an ϵ-net of $\mathcal{F}$ with probability at least $\frac{1}{2}$.*

It was later shown that the above bound is optimal within constant factors; that is, for every positive integer n, integer $d \geq 2$ and small-enough parameter $\epsilon > 0$, there is a set system $(X, \mathcal{F})$ with $|X| = n$ and VC-dim$(\mathcal{F}) \leq d$, such that any ϵ-net of $\mathcal{F}$ has size $\Omega\left(\frac{d}{\epsilon}\log\frac{1}{\epsilon}\right)$. This lower bound will be presented in Chapters 10 and 11.

Over the past thirty years, it has been observed that improvements to the Clarkson and Haussler–Welzl bounds are possible for a variety of geometric set systems. We have already seen an example in Chapter 3: $O\left(\frac{1}{\epsilon}\right)$-sized ϵ-nets exist for set systems induced by disks in the plane. Early work towards $o\left(\frac{1}{\epsilon}\log\frac{1}{\epsilon}\right)$ upper bounds was fundamentally geometric, involving spatial partitioning along the ideas seen in Chapter 3.

Over the next twenty years, through the work of Aronov, Chan, Clarkson, Ezra, Ray, Sharir, Varadarajan and others, it was realized that geometry is not really needed. In fact, somewhat surprisingly, an entire suite of optimal bounds can be obtained entirely combinatorially, with the shallow-cell complexity being a key parameter of a set system that dictates the size of ϵ-nets.

> The reason that the shallow-cell complexity of a set system comes into play is the following. Given a set system $(X, \mathcal{F})$, the probability that a set $F \in \mathcal{F}$ is not hit by a uniform random sample $N \subseteq X$ decreases exponentially with the size of F. On the other hand, the number of sets

of $\mathcal{F}$ of size at most k is an increasing function of k whose growth is upper bounded by the shallow-cell complexity of the set system. It turns out that the interplay between these two dictates the sizes of ϵ-nets.

The second main theorem of this chapter will be the following.

THEOREM 6.2. *Let* $(X, \mathcal{F})$ *be a finite set system with shallow-cell complexity* $\varphi_{\mathcal{F}}(\cdot, \cdot)$ *and with* $\text{VC-dim}(\mathcal{F}) \leq d$. *Then for any* $\epsilon \in \left(0, \frac{1}{2}\right)$ *there exists an* ϵ*-net of* $\mathcal{F}$ *of size*

$$O\left(\frac{d}{\epsilon} + \frac{1}{\epsilon} \log \varphi_{\mathcal{F}}\left(\frac{16d}{\epsilon}, 48d\right)\right).$$

Consider the primal set system $\mathcal{R}$ induced by disks in the plane: as $\varphi_{\mathcal{R}}(m, k) = O\left(k^2\right)$ (Lemma 1.2) and $\text{VC-dim}(\mathcal{R}) \leq 3$, Theorem 6.2 implies the existence of ϵ-nets of size $O\left(\frac{1}{\epsilon}\right)$. Thus we recover Theorem 3.3 from just the shallow-cell complexity of $\mathcal{R}$!

1. A First Bound using Ghost Sampling

Mathematics is one of a few fields in which one can do top-level work without a lot of life experience, something that might be key in the arts or humanities. One does not have to have experience raising children through school, dealing with family tragedies, and so forth, to be able to find three numbers whose fourth powers add up to another one.

Noam Elkies

We prove the following.

THEOREM 6.1. *Let $(X, \mathcal{F})$ be a finite set system, $d \geq 1$ an integer such that* VC-dim$(\mathcal{F}) \leq d$, *and $\epsilon \in \left(0, \frac{1}{2}\right)$ a given parameter. Let N be a uniform random sample of X of size $t = \left\lceil \frac{56\,d}{\epsilon} \ln \frac{1}{\epsilon} \right\rceil$. Then N is an ϵ-net of $\mathcal{F}$ with probability at least $\frac{1}{2}$.*

Set $n = |X|$. We can assume that each set in $\mathcal{F}$ has size at least ϵn.

For ease of calculations, we will allow an element to be picked into N multiple times. That is, N will be a *sequence* of size $t = \left\lceil \frac{56d}{\epsilon} \ln \frac{1}{\epsilon} \right\rceil$, where each element in this sequence is chosen uniformly at random from X. In a natural way, a sequence Y is an ϵ-net of $\mathcal{F}$ if Y contains at least one element from each set of $\mathcal{F}$.

Throughout this section we will work with sequences instead of sets. Moreover, the *size* of the intersection of any set $R \in \mathcal{F}$ with a sequence will count multiplicities.

Overview of ideas. Let $\mathcal{T} = X^t$ denote the set of all t-sized sequences of elements of X, and let $\mathcal{T}_b \subseteq \mathcal{T}$ be the sequences which are *not* an ϵ-net of $\mathcal{F}$. Our goal is to upper bound $|\mathcal{T}_b|$. In particular, we will show that $|\mathcal{T}_b| < \frac{|\mathcal{T}|}{2}$, implying that N—constructed by picking t elements uniformly at random, with replacement—is an ϵ-net of $\mathcal{F}$ with probability at least $\frac{1}{2}$. This implies the same property for a uniform random sample of X of size t, proving Theorem 6.1.

Fix an ordering of the sets of $\mathcal{F}$ and

for each $Y \in \mathcal{T}_b$, let $R_Y \in \mathcal{F}$ be the *first* set in the ordering for which Y fails—that is, for which $R_Y \cap Y = \emptyset$.

We will use the probabilistic averaging technique from Chapter 1. That is, we take a uniform random sequence S of X of a certain size s—with replacement, so $S \in X^s$—and examine the relationship between $\mathcal{T}_b$ and S. Specifically,

we count the expected number of sequences $Y \in \mathcal{T}_b$ s.t. $|R_Y \cap S| \geq \frac{\epsilon s}{2}$.

Lower bound: On one hand, for any $Y \in \mathcal{T}_b$,

$$\mathrm{E}\left[\, |R_Y \cap S| \,\right] = \sum_{i=1}^{s} \frac{|R_Y|}{n} \geq \sum_{i=1}^{s} \epsilon = \epsilon s, \tag{6.3}$$

keeping in mind that each element of R_Y is counted with multiplicity in $|R_Y \cap S|$. A tail bound will then imply that the expected number of sets $Y \in \mathcal{T}_b$ for which $|R_Y \cap S| \geq \frac{\epsilon s}{2}$ is at least $\frac{|\mathcal{T}_b|}{2}$.

Upper bound: On the other hand, this expectation can be calculated exactly:

$$\frac{1}{n^s}\sum_{Z\in X^s}\left|\left\{Y\in\mathcal{T}_b\ :\ |R_Y\cap Z|\ \geq\frac{\epsilon s}{2}\right\}\right|$$

$$=\frac{1}{n^s}\ \left|\left\{(Z,Y):\ Z\in X^s,\ Y\in\mathcal{T}_b\subseteq X^t,\ |R_Y\cap Z|\geq\frac{\epsilon s}{2}\right\}\right|. \tag{6.4}$$

Call each of the above (Z,Y) a *satisfying pair*. That is,

$$Z\in X^s,\quad Y\in\mathcal{T}_b\subseteq X^t,\qquad |R_Y\cap Y|=0\quad \text{while}\quad |R_Y\cap Z|\geq\tfrac{\epsilon s}{2}.$$

We will show that these constraints together force an upper bound on the total number of satisfying pairs.

Combining the lower and upper bounds will then give the desired bound on $|\mathcal{T}_b|$.

Proof of Theorem 6.1. For $Z\in X^s$, define

$$\mathcal{T}_Z=\left\{Y\in\mathcal{T}_b\ :\ |R_Y\cap Z|\ \geq\frac{\epsilon s}{2}\right\}.$$

Set $s=t$ and let S be an element chosen uniformly at random from X^s.

We count the expected size of $\mathcal{T}_S$ in two ways.

Lemma 6.5 (Lower bound).

$$\mathrm{E}\left[\,|\mathcal{T}_S|\,\right]=\sum_{Y\in\mathcal{T}_b}\Pr\left[|R_Y\cap S|\geq\frac{\epsilon s}{2}\right]\geq\frac{1}{2}\cdot|\mathcal{T}_b|.$$

Proof. The proof follows from linearity of expectation and the next claim.

Claim 6.6. For any $R\in\mathcal{F}$, $\Pr\left[|R\cap S|\leq\frac{\epsilon s}{2}\right]<\frac{1}{2}$.

Proof. For $i=1,\ldots,s$, let Y_i be an indicator random variable that is 1 if and only if the i-th element of S is in R. Then

$$|R\cap S|=\sum_{i=1}^{s}Y_i,\quad \text{where}\quad \Pr\left[Y_i=1\right]=\frac{|R|}{n}\geq\epsilon.$$

Then Chernoff's bound (Theorem 1.20) applied to the s variables $\{Y_1,\ldots,Y_s\}$ with $\delta=\frac{1}{2}$ implies that

$$\Pr\left[|R\cap S|\leq\left(1-\frac{1}{2}\right)\epsilon s\right]\leq\exp\left(-\frac{\epsilon s}{8}\right)\leq\exp\left(-\frac{56d\ln\frac{1}{\epsilon}}{8}\right)<\frac{1}{2},$$

recalling that $s=t=\left\lceil\frac{56d}{\epsilon}\ln\frac{1}{\epsilon}\right\rceil$. □

The upper bound of $\frac{1}{2}$ in Claim 6.6 can be replaced by $o(1)$ but it doesn't matter for us as this only changes the multiplicative constant factor in the final bound (see discussion). □

The upper bound on $\mathrm{E}\left[|\mathcal{T}_S|\right]$ is implied by the following combinatorial statement.

Lemma 6.7 (Upper bound).

$$\sum_{Z\in X^s}|\mathcal{T}_Z|=\left|\left\{(Z,Y):\quad Z\in X^s,\ Y\in\mathcal{T}_b\subseteq X^t,\ |R_Y\cap Z|\geq\frac{\epsilon s}{2}\right\}\right|<\frac{n^{2t}}{4}. \tag{6.8}$$

PROOF. The trick to showing Equation (6.8) is to use averaging again.

For each $U \in X^{s+t}$ and $R \in \mathcal{F}$, let $\Psi(U, R)$ be the number of ways of partitioning U into two subsequences $Z \in X^s$ and $Y \in X^t$ such that (Z, Y) is a satisfying pair, with $R_Y = R$.

As each satisfying pair (Z, Y) can be combined into a sequence of size $s+t$ in $\binom{s+t}{t}$ ways, we have

$$\binom{s+t}{t} \sum_{Z \in X^s} |\mathcal{T}_Z| = \sum_{U \in X^{s+t}} \sum_{R \in \mathcal{F}} \Psi(U, R). \tag{6.9}$$

We make two observations:

(1) For each fixed $U \in X^{s+t}$, let $\mathcal{F}' \subseteq \mathcal{F}$ be such that all the sets in $\mathcal{F}'$ have the same intersection with U. Then all $R \in \mathcal{F}'$ have the same intersection with any (Z, Y) derived from U, implying that $\Psi(U, R)$ is (possibly) non-zero only for the first set, according to our initial ordering, of $\mathcal{F}'$. As there are $|\mathcal{F}|_U|$ distinct intersections of sets of $\mathcal{F}$ with U, there are at most $|\mathcal{F}|_U|$ sets $R \in \mathcal{F}$ for which $\Psi(U, R)$ is non-zero.

(2) For a fixed $R \in \mathcal{F}$ with $|R \cap U| \geq \frac{\epsilon s}{2}$, there are at most $\binom{s+t-\frac{\epsilon s}{2}}{t}$ ways to select Y from U such that Y does not contain any element of R. This is an upper bound on $\Psi(U, R)$ for any fixed U and R.

The above two observations together with Equation (6.9) imply that $\sum_{Z \in X^s} |\mathcal{T}_Z|$ can be upper bounded by

$$\frac{1}{\binom{s+t}{t}} \cdot \sum_{U \in X^{s+t}} |\mathcal{F}|_U| \cdot \binom{s+t-\frac{\epsilon s}{2}}{t} \leq \sum_{U \in X^{s+t}} \cdot \left(\frac{e(s+t)}{d}\right)^d \cdot \frac{\binom{s+t-\frac{\epsilon s}{2}}{t}}{\binom{s+t}{t}},$$

where the second step follows from Lemma 4.3. It remains to simplify this upper bound:

$$\begin{aligned}
&= n^{2t} \cdot \left(\frac{2et}{d}\right)^d \cdot \frac{2t - \frac{\epsilon t}{2}}{2t} \, \frac{2t - \frac{\epsilon t}{2} - 1}{2t-1} \cdots \frac{t - \frac{\epsilon t}{2} + 1}{t+1} \quad \Big(\text{recalling that } s = t\Big) \\
&= n^{2t} \cdot \left(\frac{2et}{d}\right)^d \cdot \left(1 - \frac{\frac{\epsilon t}{2}}{2t}\right) \cdots \left(1 - \frac{\frac{\epsilon t}{2}}{t+1}\right) \;\leq\; n^{2t} \cdot \left(\frac{2et}{d}\right)^d \cdot \left(1 - \frac{\frac{\epsilon t}{2}}{2t}\right)^t \\
&\leq n^{2t} \cdot \left(\frac{2et}{d}\right)^d \cdot e^{-\frac{\epsilon t}{4}} \leq n^{2t} \cdot \left(\frac{112\, e}{\epsilon} \ln \frac{1}{\epsilon}\right)^d \cdot e^{-14d \ln \frac{1}{\epsilon}} \\
&= n^{2t} \cdot \left(\frac{112\, e}{\epsilon} \ln \frac{1}{\epsilon}\right)^d \cdot \epsilon^{14d} < n^{2t} \cdot \left(\frac{112\, e}{\epsilon^2}\right)^d \cdot \epsilon^{14d} \\
&< n^{2t} \cdot (112\, e)^d \cdot \epsilon^{12d} < \frac{n^{2t}}{4},
\end{aligned}$$

as $(112e)^d \cdot \epsilon^{12d} \leq (112e)^d \cdot \left(\frac{1}{2}\right)^{12d} < \frac{1}{4}$. □

Combining the upper and lower bounds,

$$\frac{1}{2} \cdot |\mathcal{T}_b| \quad \leq \quad \mathrm{E}\left[|\mathcal{T}_S|\right] \quad = \quad \frac{1}{n^s} \sum_{Z \in X^s} |\mathcal{T}_Z| \quad < \quad \frac{n^t}{4},$$

gives $|\mathcal{T}_b| < \frac{n^t}{2} = \frac{|\mathcal{T}|}{2}$, as required. □

A remark: the upper bound used in the proof, that

$$|\mathcal{F}|_U| \leq \left(\frac{e|U|}{d}\right)^d,$$

is a consequence of the fact that $\text{VC-dim}(\mathcal{F}) \leq d$. If instead we only used $|\mathcal{F}|_U| = O(|U|)^d$, the final bound would come out to be $O\left(\frac{d}{\epsilon}\log\frac{d}{\epsilon}\right)$.

We conclude with two remarks.

Ghost sampling: A clever double-counting trick in the proof is to upper bound the number of satisfying pairs (Z, Y)—where $Z \in X^s$ and $Y \in X^t$—by enumerating over all $(s+t)$-sized subsets of X. This trick was also used in the proof of Theorem 5.1, though in that case we had $t = 1$.

In statistics and learning theory literature, this instance of double-counting is called *ghost sampling*, and it is a useful technique to avoid discretization when the base set X is infinite. For the case when the set system $(X, \mathcal{F})$ is finite, there are simpler proofs (see Chapter 12).

In the proof that we presented, we only wanted the probability that the random sample N succeeds to be an ϵ-net of $\mathcal{F}$ to be non-zero, which is sufficient to guarantee the existence of an ϵ-net of the required size. By introducing this probability as a parameter in the sample size and re-working the above proof, we arrive at the following statement (a proof is presented in Chapter 12).

THEOREM 6.10. *Let $(X, \mathcal{F})$ be a finite set system, $d \in \mathbb{N}$ a positive integer such that* $\text{VC-dim}(\mathcal{F}) \leq d$, *and $\epsilon \in \left(0, \frac{1}{2}\right)$ a given parameter. Then there exists an absolute constant $C_6 > 0$ such that a random sample N constructed by picking each point of X independently with probability $\frac{C_6}{\epsilon|X|}\ln\frac{1}{\epsilon^d\gamma}$ is an ϵ-net of $\mathcal{F}$ with probability at least $1-\gamma$.*

Iterative View: As is often the case with double-counting proofs, one can 'unroll' the proof of Theorem 6.1 to an iterative version, as follows[1]. Let $(X_0, \mathcal{F}_0)$ be a set system with $|X_0| = n$ and $\epsilon > 0$ a parameter such that each set of $\mathcal{F}_0$ has size at least ϵn. A straightforward application of Chernoff's bound implies that there exists a $X_1 \subseteq X_0$ such that each $S \in \mathcal{F}_0$ contains at least $\frac{\epsilon n}{2}$ elements of X_1 and further

$$|X_1| \leq |X_0| \cdot \left(\frac{1}{2} + \sqrt{\frac{10\log|\mathcal{F}_0|}{\epsilon n}}\right).$$

Now one can repeat this step for the set system $(X_1, \mathcal{F}_1 = \mathcal{F}|_{X_1})$ to get a set $X_2 \subseteq X_1$ and so on. After the i-th iteration we have a set X_i such that each $S \in \mathcal{F}_0$ contains at least $\frac{\epsilon n}{2^i}$ elements of X_i and furthermore,

$$|X_i| \leq |X_{i-1}| \cdot \left(\frac{1}{2} + \sqrt{\frac{10\log|\mathcal{F}_{i-1}|}{\frac{\epsilon n}{2^{i-1}}}}\right) \leq \quad \cdots \quad \leq n \cdot \prod_{j=0}^{i-1}\left(\frac{1}{2} + \sqrt{\frac{10\log|\mathcal{F}_j|}{\frac{\epsilon n}{2^j}}}\right).$$

[1]Also this can be done by an inductive argument, though that somewhat obscures the ideas.

We continue the iterations as long as iteration i satisfies $\sqrt{\frac{10\log|\mathcal{F}_i|}{\frac{\epsilon n}{2^i}}} \leq \frac{1}{2}$. Say the above procedure runs for t iterations. Now a calculation shows that for all $i \leq t$,

$$|X_i| \leq n \cdot \prod_{j=0}^{i-1}\left(\frac{1}{2} + \sqrt{\frac{10\log|\mathcal{F}_j|}{\frac{\epsilon n}{2^j}}}\right) \leq c_1 \cdot \frac{n}{2^i}, \tag{6.11}$$

where c_1 is a sufficiently large constant. We set $t = \left\lfloor \log \frac{\epsilon n}{c'\, d\, \log\frac{1}{\epsilon}} \right\rfloor$ for a sufficiently large constant c' depending only on c_1. Then it can be verified that for all $i \leq t$,

$$\sqrt{\frac{10\log|\mathcal{F}_i|}{\frac{\epsilon n}{2^i}}} \leq \sqrt{\frac{10\log|\mathcal{F}_t|}{\frac{\epsilon n}{2^t}}} \leq \sqrt{\frac{10\log\left(\frac{ec_1 n}{2^t d}\right)^d}{\frac{\epsilon n}{2^t}}} \leq \sqrt{\frac{10\log\left(\frac{ec_1 c'\log\frac{1}{\epsilon}}{\epsilon}\right)}{c'\log\frac{1}{\epsilon}}} \leq \frac{1}{2},$$

where the second step follows from Equation (6.11) and Lemma 4.3.
Finally, each $S \in \mathcal{F}_0$ contains at least $\frac{\epsilon n}{2^t} \geq c' d \log\frac{1}{\epsilon} \geq 1$ points of X_t. Thus X_t is an ϵ-net of $\mathcal{F}_0$, with

$$|X_t| \leq c_1 \cdot \frac{n}{2^t} = c_1 \cdot \frac{n}{\frac{\epsilon n}{c'\, d\, \log\frac{1}{\epsilon}}} = O\left(\frac{d}{\epsilon}\log\frac{1}{\epsilon}\right).$$

An elegant strengthened form of this idea is the basis of Chapter 8.

Bibliography and discussion. A slightly weaker bound than the main theorem was first shown in [**HW876a**], and built upon the work in [**VC716a**]. The bound of this section is from [**Blu+89**]. The proof is usually stated entirely in probabilistic language while we have chosen a more combinatorial exposition that makes clear its basis in the probabilistic averaging technique of Chapter 1. The proof also follows immediately from the more general notion of ϵ-approximations (see Chapter 13).
The constant '56' in Theorem 6.1 can be replaced, with more precise calculations, by $1 + o(1)$! In particular, it was shown in [**KPW926a**] that a uniform random sample obtained by $\frac{d}{\epsilon}\left(\log\frac{1}{\epsilon} + 2\log\log\frac{1}{\epsilon} + 3\right)$ independent draws is an ϵ-net with probability at least $1 - e^{-d}$.

[Blu+89] A. Blumer, A. Ehrenfeucht, D. Haussler, and M. K. Warmuth, *Learnability and the Vapnik-Chervonenkis dimension*, J. Assoc. Comput. Mach. **36** (1989), no. 4, 929–965, DOI 10.1145/76359.76371. MR1072253

[HW876a] D. Haussler and E. Welzl, *ϵ-nets and simplex range queries*, Discrete Comput. Geom. **2** (1987), no. 2, 127–151, DOI 10.1007/BF02187876. MR884223

[KPW926a] J. Komlós, J. Pach, and G. Woeginger, *Almost tight bounds for ϵ-nets*, Discrete Comput. Geom. **7** (1992), no. 2, 163–173, DOI 10.1007/BF02187833. MR1139078

[VC716a] V. N. Vapnik and A. Ja. Červonenkis, *The uniform convergence of frequencies of the appearance of events to their probabilities* (Russian, with English summary), Teor. Verojatnost. i Primenen. **16** (1971), 264–279. MR0288823

2. Optimal ϵ-Nets using Packings

The question you raise, 'how can such a formulation lead to computations?' doesn't bother me in the least! Throughout my whole life as a mathematician, the possibility of making explicit, elegant computations has always come out by itself, as a byproduct of a thorough conceptual understanding of what was going on. Thus I never bothered about whether what would come out would be suitable for this or that, but just tried to understand—and it always turned out that understanding was all that mattered.

Alexandre Grothendieck

Given a set system $(X, \mathcal{F})$, we now show the existence of small ϵ-nets of $\mathcal{F}$ as a function of its shallow-cell complexity, which we first recall.

> Definition 4.4. A set system $(X, \mathcal{F})$ has shallow-cell complexity $\varphi_{\mathcal{F}}(\cdot,\cdot)$ if for any positive integer k and any finite $Y \subseteq X$, the number of sets in $\mathcal{F}|_Y$ of size at most k is upper bounded by $|Y| \cdot \varphi_{\mathcal{F}}(|Y|, k)$.
>
> For a family $\mathcal{R}$ of geometric objects in $\mathbb{R}^d$—e.g., the family of all half-spaces—the shallow-cell complexity of $\mathcal{R}$ is defined to be the shallow-cell complexity of the primal set system $(\mathbb{R}^d, \mathcal{R})$.

The main theorem we will prove in this section is the following.

Theorem 6.2. *Let $(X, \mathcal{F})$ be a finite set system with shallow-cell complexity $\varphi_{\mathcal{F}}(\cdot,\cdot)$ and with* VC-dim$(\mathcal{F}) \leq d$. *Then for any $\epsilon \in \left(0, \frac{1}{2}\right)$ there exists an ϵ-net of $\mathcal{F}$ of size*

$$O\left(\frac{d}{\epsilon} + \frac{1}{\epsilon}\log \varphi_{\mathcal{F}}\left(\frac{16d}{\epsilon}, 48d\right)\right).$$

Overview of ideas. The proof requires three ideas. For the moment assume that each set in $\mathcal{F}$ has size exactly ϵn.

(1) Fix any *maximal* subset $\mathcal{P} \subseteq \mathcal{F}$ such that every pair of sets of $\mathcal{P}$ have symmetric difference at least $\frac{\epsilon n}{2}$. In other words, $\mathcal{P}$ is a maximal $\left(\epsilon n, \frac{\epsilon n}{2}\right)$-packing of $\mathcal{F}$ (see Definition 5.9). Setting $\delta = \frac{\epsilon n}{2}$, $k = \epsilon n$ and applying Theorem 5.10, we have

$$|\mathcal{P}| \leq \frac{48dn}{\delta} \cdot \varphi_{\mathcal{F}}\left(\frac{8dn}{\delta}, \frac{24dk}{\delta}\right) = O\left(\frac{d}{\epsilon} \cdot \varphi_{\mathcal{F}}\left(O\left(\frac{d}{\epsilon}\right), O(d)\right)\right).$$

The key point here is that for each $F \in \mathcal{F} \setminus \mathcal{P}$, the maximality of $\mathcal{P}$ implies that there exists a set $S \in \mathcal{P}$ such that the size of the symmetric difference between F and S is at most $\frac{\epsilon n}{2}$. In other words, F contains at least $\frac{\epsilon n}{2}$ points from S, where $|S| = \epsilon n$ by assumption. Thus a $\frac{1}{2}$-net N_S for the set system $(S, \mathcal{F}|_S)$ must hit F and consequently $\bigcup_{S \in \mathcal{P}} N_S$ is an ϵ-net of $\mathcal{F}$. Here the sets of $\mathcal{P}$ play the role of canonical objects of Chapter 2.

Simply picking a $\frac{1}{2}$-net N_S *separately* for each $(S, \mathcal{F}|_S)$, $S \in \mathcal{P}$, where each N_S is of constant size by Theorem 6.10, will give an ϵ-net of $\mathcal{F}$ of total size

$$\sum_{S \in \mathcal{P}} |N_S| = O\big(|\mathcal{P}|\big) = O\left(\frac{d}{\epsilon} \cdot \varphi_{\mathcal{F}}\left(O\left(\frac{d}{\epsilon}\right), O(d)\right)\right).$$

This is too big.

(2) The next idea is to 'amortize' the size of the $\frac{1}{2}$-nets by first picking a random sample $R \subseteq X$.

For a fixed $S \in \mathcal{P}$ and $(S, \mathcal{F}|_S)$, Theorem 6.10 states that a random sample of S constructed by picking each point of S independently with probability

$$\begin{aligned} p &= \Theta\left(\frac{1}{(1/2)\,|S|}\log\frac{1}{\gamma} + \frac{d}{(1/2)\,|S|}\log\frac{1}{(1/2)}\right) \\ &= \Theta\left(\frac{1}{\epsilon n}\log\frac{1}{\gamma} + \frac{d}{\epsilon n}\right). \end{aligned}$$

is a $\frac{1}{2}$-net of $\mathcal{F}|_S$ with probability at least $1-\gamma$.

Instead of sampling points separately from each $S \in \mathcal{P}$, we will construct a 'global' random sample R by picking each point of X independently with the above probability p.

For a fixed $S \in \mathcal{P}$, R fails to be a $\frac{1}{2}$-net of $\mathcal{F}|_S$ with probability at most γ. By the union bound, the probability that there exists a $S \in \mathcal{P}$ such that R fails to be a $\frac{1}{2}$-net of $\mathcal{F}|_S$ is at most $|\mathcal{P}| \cdot \gamma$. Setting $\gamma = \frac{1}{|\mathcal{P}|+1}$ implies that, with non-zero probability, R is a $\frac{1}{2}$-net for all $\mathcal{F}|_S$, $S \in \mathcal{P}$. Then

$$\begin{aligned} \mathrm{E}\left[|R|\right] = np &= O\left(\frac{1}{\epsilon}\log\frac{1}{\gamma} + \frac{d}{\epsilon}\right) \\ &= O\left(\frac{1}{\epsilon}\log\left(\frac{d}{\epsilon}\cdot\varphi_{\mathcal{F}}\left(O\left(\frac{d}{\epsilon}\right), O\left(d\right)\right)\right) + \frac{d}{\epsilon}\right) \\ &= O\left(\frac{1}{\epsilon}\log\frac{d}{\epsilon}\right) + O\left(\frac{1}{\epsilon}\log\varphi_{\mathcal{F}}\left(O\left(\frac{d}{\epsilon}\right), O\left(d\right)\right)\right) + O\left(\frac{d}{\epsilon}\right). \end{aligned}$$

The large additional term $O\left(\frac{1}{\epsilon}\log\frac{d}{\epsilon}\right)$ still remains.

(3) Since the expected number of sets in $\mathcal{P}$ for which R fails is at most $|\mathcal{P}| \cdot \gamma$, we had set $\gamma < \frac{1}{|\mathcal{P}|}$ to ensure the existence of a set R with no failures. Instead, we will set γ such that the expected number of $\mathcal{F}|_S$, $S \in \mathcal{P}$, for which R fails to be a $\frac{1}{2}$-net is $O\left(\frac{1}{\epsilon}\right)$. The key point is that for each of these failed sets of $\mathcal{P}$, we can afford to separately add a $O(d)$-sized $\frac{1}{2}$-net for a total of $O\left(\frac{d}{\epsilon}\right)$ additional points. In other words, we set γ so that $|\mathcal{P}| \cdot \gamma = \Theta\left(\frac{1}{\epsilon}\right)$. Then the size of the initial random sample R is

$$\begin{aligned} O\left(\frac{1}{\epsilon}\log\frac{1}{\gamma} + \frac{d}{\epsilon}\right) &= O\left(\frac{1}{\epsilon}\log\left(\epsilon\,|\mathcal{P}|\right) + \frac{d}{\epsilon}\right) \\ &= O\left(\frac{1}{\epsilon}\log\varphi_{\mathcal{F}}\left(O\left(\frac{d}{\epsilon}\right), O\left(d\right)\right) + \frac{d}{\epsilon}\right), \end{aligned}$$

as desired.

We now turn to the formal proof with complete calculations.

PROOF OF THEOREM 6.2. For an integer $j \geq 0$, set

$$\epsilon_j = 2^j \cdot \epsilon \qquad \text{and} \qquad \delta_j = \frac{\epsilon_j n}{2}.$$

For each $j = 1, \ldots, \lceil \log \frac{1}{\epsilon} \rceil$, define

$$\mathcal{F}_j = \left\{ S \in \mathcal{F} \ : \ \epsilon_{j-1} n \ \leq \ |S| \ < \ \epsilon_j n \right\}.$$

Further let $\mathcal{P}_j$ be a *maximal* subset of $\mathcal{F}_j$ satisfying the property that

$$\text{for all } S, S' \in \mathcal{P}_j, \qquad |\Delta(S, S')| \geq \delta_j.$$

As each set in $\mathcal{F}_j$ has size less than $\epsilon_j n$, Theorem 5.10 implies that

$$|\mathcal{P}_j| \leq \frac{48\, dn}{\delta_j} \cdot \varphi_{\mathcal{F}} \left(\frac{8\, dn}{\delta_j}, \frac{24\, d\epsilon_j n}{\delta_j} \right) = \frac{96\, dn}{\epsilon_j n} \cdot \varphi_{\mathcal{F}} \left(\frac{16\, dn}{\epsilon_j n}, \frac{48\, d\epsilon_j n}{\epsilon_j n} \right)$$

$$= \frac{96\, d}{\epsilon_j} \cdot \varphi_{\mathcal{F}} \left(\frac{16\, d}{\epsilon_j}, 48\, d \right). \tag{6.12}$$

CLAIM 6.13. Let $j \in \left\{ 1, \ldots, \lceil \log \frac{1}{\epsilon} \rceil \right\}$. Suppose the set $N_j \subseteq X$ is a $\frac{1}{2}$-net *simultaneously* for all these set systems:

$$\left\{ \left(S, \mathcal{F}_j|_S \right) \ : \ S \in \mathcal{P}_j \right\}.$$

Then N_j hits all the sets of $\mathcal{F}_j$.

PROOF. Let $F \in \mathcal{F}_j$. If $F \in \mathcal{P}_j$, then clearly it is hit by the $\frac{1}{2}$-net of $\mathcal{F}_j|_F$. Otherwise by the maximality of $\mathcal{P}_j$, there exists a $S \in \mathcal{P}_j$ such that

$$|\Delta(F, S)| = |F \setminus S| + |S \setminus F| < \delta_j \qquad \text{and hence} \qquad |F \setminus S| < \delta_j - |S \setminus F|.$$

As $|F| \geq 2^{j-1} \epsilon n = \delta_j$, it follows that

$$|F \cap S| = |F| - |F \setminus S| \geq \delta_j - \left(\delta_j - |S \setminus F| \right) = |S \setminus F|.$$

This implies that $|F \cap S| \geq \frac{|S|}{2}$ and as $F \cap S \in \mathcal{F}_j|_S$, F is hit by the $\frac{1}{2}$-net of $\mathcal{F}_j|_S$. $\square$

Thus it suffices to compute, for each $j = 1, \ldots, \lceil \log \frac{1}{\epsilon} \rceil$, a set N_j such that

$$N_j \ \text{ is a } \tfrac{1}{2}\text{-net of all the } |\mathcal{P}_j| \text{ set systems } (S, \mathcal{F}_j|_S),\ S \in \mathcal{P}_j.$$

We can then return $\bigcup_j N_j$ as an ϵ-net of $\mathcal{F}$. We will construct each N_j separately for each index $j \in \left\{ 1, \ldots, \lceil \log \frac{1}{\epsilon} \rceil \right\}$ by computing the following two sets R_j and M_j, and setting $N_j = R_j \cup M_j$.

Constructing R_j: Let R_j be a sample constructed by picking each point of X independently with probability

$$C_6 \cdot \left(\frac{d}{(1/2) \cdot \epsilon_{j-1} n} \ln \frac{1}{(1/2)} + \frac{1}{(1/2) \cdot \epsilon_{j-1} n} \ln \left(d \cdot \varphi_{\mathcal{F}} \left(\frac{16d}{\epsilon_j}, 48d \right) \right) \right).$$

where C_6 is the constant from Theorem 6.10. For each $S \in \mathcal{P}_j$, we have $|S| \in [\epsilon_{j-1} n, \epsilon_j n)$ and so Theorem 6.10 applied to

$$\text{the set system } \ (S, \mathcal{F}|_S) \quad \text{with} \quad \epsilon = \frac{1}{2}, \quad \gamma = \frac{1}{d \cdot \varphi_{\mathcal{F}} \left(\frac{16d}{\epsilon_j}, 48d \right)},$$

implies that R_j is a $\frac{1}{2}$-net of $\mathcal{F}|_S$ with probability at least $1 - \gamma$. Furthermore,

$$\mathrm{E}\left[|R_j| \right] = O \left(\frac{d}{\epsilon_j} + \frac{1}{\epsilon_j} \log \left(d \cdot \varphi_{\mathcal{F}} \left(\frac{16d}{\epsilon_j}, 48d \right) \right) \right).$$

Constructing M_j: Initialize $M_j = \emptyset$. For each $S \in \mathcal{P}_j$ for which R_j *fails* to be a $\frac{1}{2}$-net, construct a $\frac{1}{2}$-net of $(S, \mathcal{F}_j|_S)$ of size $O(d)$ (by Theorem 6.10) and add it to M_j. Then

$$\begin{aligned}
\mathrm{E}\left[|M_j|\right] &= \sum_{S\in\mathcal{P}_j} \Pr\left[R_j \text{ is not a } \frac{1}{2}\text{-net of } \mathcal{F}_j|_S\right] \cdot \left(\text{size of } \frac{1}{2}\text{-net of } \mathcal{F}_j|_S\right) \\
&\le \sum_{S\in\mathcal{P}_j} \frac{1}{d\cdot\varphi_{\mathcal{F}}\left(\frac{16d}{\epsilon_j}, 48d\right)} \cdot O(d) = |\mathcal{P}_j| \cdot \frac{1}{d\cdot\varphi_{\mathcal{F}}\left(\frac{16d}{\epsilon_j}, 48d\right)} \cdot O(d) \\
&\le \frac{96d}{\epsilon_j} \cdot \varphi_{\mathcal{F}}\left(\frac{16d}{\epsilon_j}, 48d\right) \cdot \frac{1}{d\cdot\varphi_{\mathcal{F}}\left(\frac{16d}{\epsilon_j}, 48d\right)} \cdot O(d) = O\left(\frac{d}{\epsilon_j}\right),
\end{aligned}$$

where the second-to-last step uses the upper bound on $|\mathcal{P}_j|$ given in Equation (6.12).

Thus we can conclude with the expected size of the final ϵ-net N of $\mathcal{F}$:

$$\begin{aligned}
\mathrm{E}[|N|] &= \sum_{j=1}^{\lceil \log \frac{1}{\epsilon}\rceil} \mathrm{E}\left[|R_j| + |M_j|\right] = \sum_{j=1}^{\lceil \log \frac{1}{\epsilon}\rceil} \mathrm{E}\left[|R_j|\right] + \sum_{j=1}^{\lceil \log \frac{1}{\epsilon}\rceil} \mathrm{E}\left[|M_j|\right] \\
&= \sum_{j=1}^{\lceil \log \frac{1}{\epsilon}\rceil} O\left(\frac{d}{\epsilon_j} + \frac{1}{\epsilon_j}\log\left(d\cdot\varphi_{\mathcal{F}}\left(\frac{16d}{\epsilon_j}, 48d\right)\right)\right) + \sum_{j=1}^{\lceil \log \frac{1}{\epsilon}\rceil} O\left(\frac{d}{\epsilon_j}\right) \\
&= \sum_{j=1}^{\lceil \log \frac{1}{\epsilon}\rceil} O\left(\frac{d}{2^j\epsilon} + \frac{1}{2^j\epsilon}\log d + \frac{1}{2^j\epsilon}\log\varphi_{\mathcal{F}}\left(\frac{16d}{2^j\epsilon}, 48d\right)\right) + \sum_{j=1}^{\lceil \log \frac{1}{\epsilon}\rceil} O\left(\frac{d}{2^j\epsilon}\right) \\
&= O\left(\frac{d}{\epsilon}\right) + O\left(\frac{1}{\epsilon}\right)\sum_{j=1}^{\lceil \log \frac{1}{\epsilon}\rceil} \frac{1}{2^j}\cdot\log\varphi_{\mathcal{F}}\left(\frac{16d}{2^j\epsilon}, 48d\right) \\
&\le O\left(\frac{d}{\epsilon}\right) + O\left(\frac{1}{\epsilon}\right)\sum_{j=1}^{\lceil \log \frac{1}{\epsilon}\rceil} \frac{1}{2^j}\cdot\log\varphi_{\mathcal{F}}\left(\frac{16d}{\epsilon}, 48d\right) \\
&= O\left(\frac{d}{\epsilon} + \frac{1}{\epsilon}\log\varphi_{\mathcal{F}}\left(\frac{16d}{\epsilon}, 48d\right)\right).
\end{aligned}$$

□

Bibliography and discussion. The proof in this chapter is from [**MDG18**], building on ideas from several earlier works [**AES10**, **CF90**, **CV07**, **PR08**]. In particular, [**PR08**, **AES10**] are important papers that made progress on ϵ-nets for several basic geometric set systems. It was shown in [**KMP17**] that the bound on sizes of ϵ-nets as a function of their shallow-cell complexity is tight.

[AES10] B. Aronov, E. Ezra, and M. Sharir, *Small-size ϵ-nets for axis-parallel rectangles and boxes*, SIAM J. Comput. **39** (2010), no. 7, 3248–3282, DOI 10.1137/090762968. MR2678074

[CF90] B. Chazelle and J. Friedman, *A deterministic view of random sampling and its use in geometry*, Combinatorica **10** (1990), no. 3, 229–249, DOI 10.1007/BF02122778. MR1092541

[CV07] K. L. Clarkson and K. Varadarajan, *Improved approximation algorithms for geometric set cover*, Discrete Comput. Geom. **37** (2007), no. 1, 43–58, DOI 10.1007/s00454-006-1273-8. MR2279863

[KMP17] A. Kupavskii, N. H. Mustafa, and J. Pach, *Near-optimal lower bounds for ϵ-nets for half-spaces and low complexity set systems*, A journey through discrete mathematics, Springer, Cham, 2017, pp. 527–541. MR3726612

[MDG18] N. H. Mustafa, K. Dutta, and A. Ghosh, *A simple proof of optimal epsilon nets*, Combinatorica **38** (2018), no. 5, 1269–1277, DOI 10.1007/s00493-017-3564-5. MR3884789

[PR08] E. Pyrga and S. Ray, *New existence proofs for ϵ-nets*, Computational geometry (SCG'08), ACM, New York, 2008, pp. 199–207, DOI 10.1145/1377676.1377708. MR2504287

CHAPTER 15

Epsilon-Approximations: Functional Case

Given a set P of n points in $\mathbb{R}^d$ and a parameter $\epsilon > 0$, let A be an ϵ-approximation of the primal set system induced on P by balls in $\mathbb{R}^d$. Then clearly A can be used to approximate P 'combinatorially' with respect to balls:

Let $P_{q,r} = \text{Ball}(q, r) \cap P$ be the set of points of P contained in the ball of radius r centered at q; similarly set $A_{q,r} = \text{Ball}(q, r) \cap A$.

Then for any $q \in \mathbb{R}^d$ and $r > 0$, $|P_{q,r}|$ can be approximated by $|A_{q,r}| \cdot \frac{|P|}{|A|}$, since by the definition of ϵ-approximations,

$$|A_{q,r}| = \frac{|P_{q,r}|\,|A|}{|P|} \pm \epsilon|A| \qquad \text{or equivalently,}$$

$$|P_{q,r}| = |A_{q,r}| \cdot \frac{|P|}{|A|} \pm \epsilon|P|. \tag{15.1}$$

The new idea in this chapter is the observation that since for any $q \in \mathbb{R}^d$, Equation (15.1) holds for *every* radius r, the set A can also be used to approximate the sum of distances from q to the points of P. In particular, here is another property that holds for A: for any $q \in \mathbb{R}^d$ and $r > 0$,

$$\left| \frac{\sum_{p \in P_{q,r}} \text{dist}(p, q)}{n} - \frac{\sum_{p \in A_{q,r}} \text{dist}(p, q)}{|A|} \right| \leq 3\epsilon r. \tag{15.2}$$

To see the intuition for Equation (15.2), we sketch the proof for a weaker bound of $3\sqrt{\epsilon} r$.

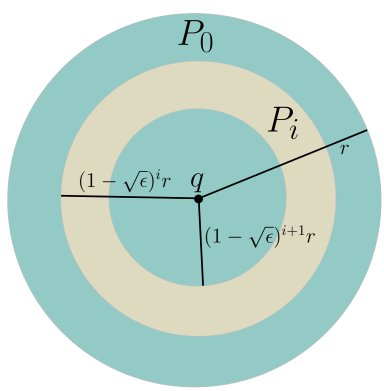

Recall that we say A is an ϵ-approximation of a set $P' \subseteq P$ if $|P' \cap A| = \frac{|P'|\,|A|}{|P|} \pm \epsilon |A|$.

Partition $P_{q,r}$ into disjoint sets $P_0, P_1, \ldots$, where $p \in P_i$ if and only if

$$\mathrm{dist}(p,q) \in \left(\left(1-\sqrt{\epsilon}\right)^{i+1} r, \left(1-\sqrt{\epsilon}\right)^{i} r \right].$$

That is, P_i is the set of points of P lying in the region

$$\mathrm{Ball}\left(q, \left(1-\sqrt{\epsilon}\right)^{i} r\right) \setminus \mathrm{Ball}\left(q, \left(1-\sqrt{\epsilon}\right)^{i+1} r\right).$$

Now the sum of distances of the points of A to q can be approximated by summing up over the P_i's:

$$\left(\sum_{i=0}^{\infty} \left(1-\sqrt{\epsilon}\right)^{i+1} r \cdot |P_i \cap A|\right) < \sum_{p\in A_{q,r}} \mathrm{dist}(p,q) \leq \left(\sum_{i=0}^{\infty} \left(1-\sqrt{\epsilon}\right)^{i} r \cdot |P_i \cap A|\right).$$

We remark that we only need to do the above sum till index $i = \frac{1}{\sqrt{\epsilon}} \ln \frac{1}{\epsilon}$, as after that the average sum of distances is most ϵr in any case. Using the fact that A is a 2ϵ-approximation of each P_i (by Claim 13.6),

$$\left(\sum_{i=0}^{\infty} \left(1-\sqrt{\epsilon}\right)^{i+1} r \left(\frac{|P_i|\,|A|}{n} - 2\epsilon|A|\right)\right) < \sum_{p\in A_{q,r}} \mathrm{dist}(p,q) \leq \left(\sum_{i=0}^{\infty} \left(1-\sqrt{\epsilon}\right)^{i} r \left(\frac{|P_i|\,|A|}{n} + 2\epsilon|A|\right)\right)$$

$$\left(1-\sqrt{\epsilon}\right) \left(\sum_{i=0}^{\infty} \left(1-\sqrt{\epsilon}\right)^{i} \left(\frac{|P_i|}{n} - 2\epsilon\right)\right) < \frac{\sum_{p\in A_{q,r}} \mathrm{dist}(p,q)}{r\,|A|} \leq \left(\sum_{i=0}^{\infty} \left(1-\sqrt{\epsilon}\right)^{i} \left(\frac{|P_i|}{n} + 2\epsilon\right)\right).$$

Using the fact that $2\epsilon \sum_{i=0}^{\infty} \left(1-\sqrt{\epsilon}\right)^{i} = 2\epsilon \cdot \frac{1}{1-\left(1-\sqrt{\epsilon}\right)} = 2\sqrt{\epsilon}$,

$$\left(1-\sqrt{\epsilon}\right) \left(\left(\frac{1}{n}\sum_{i=0}^{\infty} \left(1-\sqrt{\epsilon}\right)^{i} |P_i|\right) - 2\sqrt{\epsilon}\right) < \frac{\sum_{p\in A_{q,r}} \mathrm{dist}(p,q)}{r\,|A|} \leq \left(\frac{1}{n}\sum_{i=0}^{\infty} \left(1-\sqrt{\epsilon}\right)^{i} |P_i|\right) + 2\sqrt{\epsilon}.$$

Similarly approximating $\sum_{p\in P_{q,r}} \mathrm{dist}(p,q)$ over the P_i's gives

$$\left(1-\sqrt{\epsilon}\right) \left(\sum_{i=0}^{\infty} \left(1-\sqrt{\epsilon}\right)^{i} r\,|P_i|\right) < \sum_{p\in P_{q,r}} \mathrm{dist}(p,q) \leq \left(\sum_{i=0}^{\infty} \left(1-\sqrt{\epsilon}\right)^{i} r\,|P_i|\right).$$

Dividing by rn,

$$\left(1-\sqrt{\epsilon}\right) \left(\frac{1}{n}\sum_{i=0}^{\infty} \left(1-\sqrt{\epsilon}\right)^{i} |P_i|\right) < \frac{\sum_{p\in P_{q,r}} \mathrm{dist}(p,q)}{rn} \leq \left(\frac{1}{n}\sum_{i=0}^{\infty} \left(1-\sqrt{\epsilon}\right)^{i} |P_i|\right).$$

These together imply the desired bound:

$$\begin{aligned}\left|\frac{\sum_{p\in P_{q,r}} \mathrm{dist}(p,q)}{r\,n} - \frac{\sum_{p\in A_{q,r}} \mathrm{dist}(p,q)}{r\,|A|}\right| &\leq 2\sqrt{\epsilon} + \sqrt{\epsilon}\left(\frac{1}{n}\sum_{i=0}^{\infty} \left(1-\sqrt{\epsilon}\right)^{i} |P_i|\right) \\ &< 2\sqrt{\epsilon} + \sqrt{\epsilon}\left(\frac{1}{n}\sum_{i=0}^{\infty} |P_i|\right) \\ &\leq 3\sqrt{\epsilon}.\end{aligned}$$

As we will see later, the improvement to $3\epsilon r$ presented later in this chapter follows with more precise calculations.

Consider now another application of the same idea on a slightly more complicated distance function where the point q is replaced by a set of k points. For any set X of k points in $\mathbb{R}^d$ and $p \in P$, define

$$\text{dist}\,(p, X) = \min_{q \in X} \text{dist}\,(p, q)\,. \tag{15.3}$$

Further let

$$P_{X,r} = \{p \in P \;:\; \text{dist}\,(p, X) \leq r\}\,.$$

Observe that $P_{X,r}$ is the set of points of P that lie in the union of the k balls of radius r centered at the points of X. Denote this union by $\text{Ball}\,(X, r)$.

As earlier, our goal is to estimate, for any given $X \in \left(\mathbb{R}^d\right)^k$ and $r \geq 0$, the expression

$$\sum_{p \in P_{X,r}} \text{dist}(p, X).$$

Not surprisingly, if A is an ϵ-approximation of the set system induced on P by the union of k balls, then one can show that

$$\left| \frac{\sum_{p \in P_{X,r}} \text{dist}\,(p, X)}{n} - \frac{\sum_{p \in A_{X,r}} \text{dist}\,(p, X)}{|A|} \right| \leq 3\epsilon r.$$

The first result of this chapter is a more general statement which implies both the above two instances.

The second result is its application to an algorithmic problem central to several domains: the k-median clustering problem, where given a set P of points in $\mathbb{R}^d$ and an integer parameter $k > 0$, the goal is to partition the points of P into k clusters based on certain geometric criteria.

1. A Functional View of Approximations

A common error of judgment among mathematicians is the confusion between telling the truth and giving a logically correct presentation. The two objectives are antithetical and hard to reconcile. Most presentations obeying the current Diktats of linear rigor are a long way from telling the truth; any reader of such a presentation is forced to start writing on the margin, or deciphering on a separate sheet of paper.

The truth of any piece of mathematical writing consists of realizing what the author is "up to"; it is the tradition of mathematics to do whatever it takes to avoid giving away this secret.

Gian-Carlo Rota

Recall the following statement.

Let P be a set of n points in $\mathbb{R}^d$. Each $p \in P$ defines the function

$$\mathrm{dist}\,(p, X) = \min_{q \in X} \mathrm{dist}\,(p, q)\,,$$

where X is a finite set of points in $\mathbb{R}^d$.

Then for any positive integer k and $\epsilon > 0$, there exists an ϵ-approximation $A \subseteq P$ such that for any set X of k points in $\mathbb{R}^d$ and $r \in \mathbb{R}^+$,

$$\left| \frac{\sum_{p \in P_{X,r}} \mathrm{dist}\,(p, X)}{n} - \frac{\sum_{p \in A_{X,r}} \mathrm{dist}\,(p, X)}{|A|} \right| \leq 3\epsilon r, \tag{15.4}$$

where $P_{X,r} = \{p \in P\colon\ \mathrm{dist}\,(p, X) \leq r\}$ and $A_{X,r} = A \cap P_{X,r}$.

Note that the role of each $p \in P$ is captured by the function $\mathrm{dist}\,(p, \cdot)$. We now prove Equation (15.4) in an abstract setting where $\mathrm{dist}\,(p, \cdot)$ is replaced by an arbitrary function $g_p\colon \mathcal{X} \to \mathbb{R}^+$, where $\mathcal{X}$ is a given domain. That is,

$$\begin{aligned} \text{set of all } k\text{-tuples of points in } \mathbb{R}^d \quad &\longrightarrow \quad \text{a domain } \mathcal{X}, \\ \text{set of } n \text{ functions } \mathrm{dist}\,(p, \cdot),\ p \in P \quad &\longrightarrow \quad \text{set } G \text{ of } n \text{ functions from } \mathcal{X} \text{ to } \mathbb{R}^+, \\ P_{X,r} = \{p \in P\colon\ \mathrm{dist}(p, X) \leq r\} \quad &\longrightarrow \quad G_{X,r} = \big\{g \in G\colon g\,(X) \leq r\big\}. \end{aligned}$$

The main theorem of this section states and proves the analog of Equation (15.4) in this abstract setting for a set G of n functions.

THEOREM 15.5. *Let $G = \{g_1, \ldots, g_n\}$ be a set of n functions over a domain $\mathcal{X}$[1], where $g_i\colon \mathcal{X} \to \mathbb{R}^+$. Define the set system $(G, \mathcal{F})$, with*

$$\mathcal{F} = \big\{G_{X,r}\colon X \in \mathcal{X} \text{ and } r \in \mathbb{R}^+\big\}\,, \quad \text{where} \quad G_{X,r} = \big\{g \in G\colon g\,(X) \leq r\big\}.$$

Let $A \subseteq G$ be an ϵ-approximation of $\mathcal{F}$, for a given parameter $\epsilon > 0$. Then for any $X \in \mathcal{X}$ and $r \in \mathbb{R}^+$, setting $A_{X,r} = A \cap G_{X,r}$, we have

$$\left| \frac{\sum_{g \in G_{X,r}} g\,(X)}{|G|} - \frac{\sum_{g \in A_{X,r}} g\,(X)}{|A|} \right| \leq 3\epsilon r.$$

To visualize Theorem 15.5, consider the case when $\mathcal{X} = \mathbb{R}$. The figure illustrates an example of five functions $g_1, \ldots, g_5$. In this example, the set $G_{X,r}$ is simply the set of functions lying below the point (X, r).

[1] $\mathcal{X}$ need not be finite. In the previous example, $\mathcal{X}$ was the set of all k-tuples of points in $\mathbb{R}^d$.

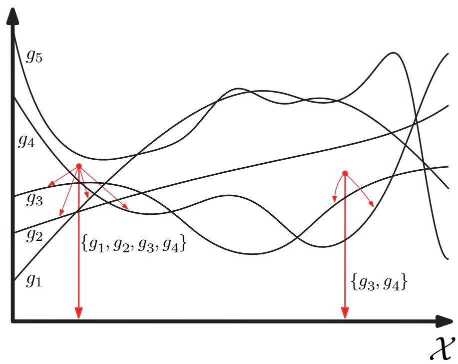

Before we proceed to the proof, we illustrate the versatility of Theorem 15.5 by showing a specific consequence.

Let P be a set of n points in $\mathbb{R}^d$ and $\mathcal{X}$ the set of all k-tuples of points in $\mathbb{R}^d$. Additionally, for each $p \in P$ we are given a function $f_p \colon \mathcal{X} \to \mathbb{R}^+$. Set

$$G = \{f_p \colon p \in P\}.$$

For each $X \in \mathcal{X}$, let r_X be the smallest value for which $G_{X,r} = G$. That is,

$$r_X = \max_{p \in P} f_p(X).$$

Applying Theorem 15.5 to P and G, we arrive at the following.

Corollary 15.6. *Let P be a set of n points in $\mathbb{R}^d$ and k a positive integer. Further each $p \in P$ has an associated function*

$$f_p \colon \left(\mathbb{R}^d\right)^k \to \mathbb{R}^+.$$

These functions define a set system $(P, \mathcal{R})$, with

$$\mathcal{R} = \left\{P_{X,r} \colon X \in \left(\mathbb{R}^d\right)^k \ \text{ and } \ r \in \mathbb{R}^+\right\},$$

$$\text{where } P_{X,r} = \left\{p \in P \colon f_p(X) \leq r\right\}.$$

Let A be an ϵ-approximation of $\mathcal{R}$. Then for any $X \subseteq \mathbb{R}^d$ with $|X| = k$, we have

$$\left| \frac{\sum_{p \in P} f_p(X)}{|P|} - \frac{\sum_{p \in A} f_p(X)}{|A|} \right| \leq 3\epsilon r_X = 3\epsilon \max_{p \in P} f_p(X).$$

Note that for the case where $f_p(X) = \mathrm{dist}(p, X)$, $\mathcal{R}$ is precisely the primal set system induced on P by the union of k equal-radius balls in $\mathbb{R}^d$.

Overview of ideas. For a fixed $X \in \mathcal{X}$ and $r \in \mathbb{R}^+$, we need to relate the quantities

$$\sum_{g \in G_{X,r}} g(X) \qquad \text{and} \qquad \sum_{g \in A_{X,r}} g(X).$$

As earlier, one way to proceed is to partition $G_{X,r}$ into disjoint sets $G_0, G_1, \ldots$, where

$$G_i = \left\{ g \in G \ : \ g(X) \in \left((1-\epsilon)^{i+1} r, (1-\epsilon)^i r \right] \right\}.$$

Then all the functions in G_i have approximately the same value on X, and so one can approximately bound the summation of functions in G_i in terms of $|G_i|$, and the summation of functions in $A \cap G_i$ in terms of $|A \cap G_i|$.

However, we present a different proof based on an elegant trick: we sort the functions in $G_{X,r}$ by increasing $g(X)$ values, and rewrite each of the above two summations as sums of the interval lengths between two consecutive function values in the sorted order. Concretely,

let $G_{X,r} = \{g_1, \ldots, g_t\}$, sorted by increasing $g(X)$ values.

Then any interval

$$C_i = g_i(X) - g_{i-1}(X)$$

is 'contributed' in $\sum_{g \in G_{X,r}} g(X)$ by precisely the functions $\{g_i, \ldots, g_t\}$ (see fig-

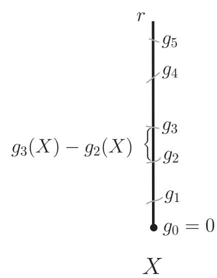

ure). Summing over all consecutive intervals, and using the fact that A is an 2ϵ-approximation of each $\{g_i, \ldots, g_t\}$, we get the required bound.

In the formal proof one has to be a little careful though, as multiple functions might have the same value on X.

Proof of Theorem 15.5. Fix any $X \in \mathcal{X}$ and $r \in \mathbb{R}^+$. Sort the functions in $G_{X,r}$ by increasing $g(X)$ values, and partition $G_{X,r}$ into groups along this order:

$$G_{X,r} = G_1 \cup \cdots \cup G_m,$$

where all the functions in G_i, $i \in [m]$, have the same value on X. Note that $m \leq |G_{X,r}|$. Set

$$G_{\geq i} = G_i \cup \cdots \cup G_m.$$

As A is an ϵ-approximation of the sets $G_1 \cup \cdots \cup G_i$ for all $i \in [m]$, Claim 13.6 (2) implies the following.

Claim 15.7. For each $i \in [m]$, A is a 2ϵ-approximation of $G_{\geq i}$.

Now we sum up the functions in $G_{X,r}$ and $A_{X,r}$ by summing up over the differences between values of adjacent functions. For each $i \in [m]$, fix an arbitrary function $g_i \in G_i$ and let $g_0 = 0$. Then

$$\sum_{g \in G_{X,r}} g(X) = \sum_{i=1}^{m} \big(g_i(X) - g_{i-1}(X)\big) \cdot \big|G_{\geq i}\big|, \qquad \text{and}$$

$$\sum_{g \in A_{X,r}} g(X) = \sum_{i=1}^{m} \big(g_i(X) - g_{i-1}(X)\big) \cdot \big|A \cap G_{\geq i}\big|.$$

Thus the required expression

$$\left| \frac{\sum_{g \in G_{X,r}} g(X)}{n} - \frac{\sum_{g \in A_{X,r}} g(X)}{|A|} \right|$$

is upper bounded by

$$\left| \left(\sum_{i=1}^{m} \big(g_i(X) - g_{i-1}(X)\big) \cdot \frac{\big|G_{\geq i}\big|}{n} \right) - \left(\sum_{i=1}^{m} \big(g_i(X) - g_{i-1}(X)\big) \cdot \frac{\big|A \cap G_{\geq i}\big|}{|A|} \right) \right|$$

$$\leq \sum_{i=1}^{m} \big(g_i(X) - g_{i-1}(X)\big) \cdot \left| \frac{|G_{\geq i}|}{n} - \frac{|A \cap G_{\geq i}|}{|A|} \right|$$

$$\leq \sum_{i=1}^{m} \big(g_i(X) - g_{i-1}(X)\big) \cdot 2\epsilon \leq 2\epsilon r,$$

where the second-to-last step used Claim 15.7. □

Bibliography and discussion. The material in this section is from [**FL11**] (with some simplifications). The size of the ϵ-approximation in Theorem 15.5 depends on a parameter called the *pseudo-dimension*, which is a generalization of the notion of VC-dimension for general functions (see [**HP11**, Chapter 7] for details).

[FL11] D. Feldman and M. Langberg, *A unified framework for approximating and clustering data*, STOC'11—Proceedings of the 43rd ACM Symposium on Theory of Computing, ACM, New York, 2011, pp. 569–578, DOI 10.1145/1993636.1993712. MR2932007

[HP11] S. Har-Peled, *Geometric approximation algorithms*, Mathematical Surveys and Monographs, vol. 173, American Mathematical Society, Providence, RI, 2011, DOI 10.1090/surv/173. MR2760023

2. Application: Sensitivity and Coresets for Clustering

Elegant algorithms are easy to program correctly, as well as being efficient. A clever algorithm that is clean and elegant is much more likely to be used than a messy one. When people understand how an algorithm works, which is much more likely with an elegant algorithm, they are more likely to have confidence in the results it produces.

Also, elegant solutions are much easier to generalize, to extend to other problems. My goal is to find general approaches and solutions, not ad hoc tricks.

Robert Tarjan

Given a set P of n points in $\mathbb{R}^d$, the k-median problem asks to compute a set X of k points that minimizes the cost function[2]

$$\mathrm{Cost}(P,k) = \min_{\substack{X\subseteq\mathbb{R}^d\\ |X|=k}} \mathrm{Cost}(P,X), \quad \text{where} \quad \mathrm{Cost}(P,X) = \sum_{p\in P} \mathrm{dist}(p,X).$$

> One approach towards solving this problem is to first compute a smaller set A that 'approximates' P with respect to $\mathrm{Cost}(P,X)$. That is, for *every* set X of k points in $\mathbb{R}^d$ we would like $\mathrm{Cost}(P,X)$ to be approximately equal to $\mathrm{Cost}(A,X)$ (scaled up appropriately). Then the original problem on P is reduced to finding an X minimizing $\mathrm{Cost}(A,X)$—an easier problem if $|A|$ is much smaller than $|P|$.

This leads to the following definition.

DEFINITION 15.8. Given a set P of n points in $\mathbb{R}^d$ and a parameter $\epsilon > 0$, a set $A \subseteq \mathbb{R}^d$ together with a weight function $w\colon A \to \mathbb{R}$ is an ϵ-coreset for the k-median problem on P if for every $X \subseteq \mathbb{R}^d$ of k points,

$$\sum_{p\in A} \mathrm{dist}\,(p,X)\cdot w\,(p) = (1\pm\epsilon)\cdot \sum_{p\in P}\mathrm{dist}\,(p,X)\,. \tag{15.9}$$

Our goal then is to construct an ϵ-coreset for the k-median problem. We will prove two main theorems, the first of which is the following.

THEOREM 15.10. *Let P be a set of n points in $\mathbb{R}^d$ and $k \in \mathbb{Z}^+$, $\epsilon > 0$ be two given parameters. Define*

$$S = \sum_{p\in P} \sup_{\substack{Y\subseteq\mathbb{R}^d\\ |Y|=k}} \frac{\mathrm{dist}(p,Y)}{\sum_{q\in P}\mathrm{dist}(q,Y)}.$$

Then there exists an ϵ-coreset $A \subseteq P$ of size $O\left(\frac{S^2\, d\, k \log k}{\epsilon^2}\right)$ for the k-median problem on P.

The size of the ϵ-coreset in Theorem 15.10 relies on the seemingly mysterious quantity S; however its proof will demonstrate that S 'falls out' naturally when constructing coresets using ϵ-approximations. There do exist good upper bounds on S but using techniques and ideas outside the scope of this text (see discussion).

[2]Recall that $\mathrm{dist}\,(p,X) = \min_{q\in X}\mathrm{dist}\,(p,X)$.

Our second main result shows that the dependency on S can be removed if one is also given an approximate solution B.

THEOREM 15.11. *Let P be a set of n points in $\mathbb{R}^d$ and $k \in \mathbb{Z}^+, \epsilon > 0$ be two given parameters. Further let $B \subseteq \mathbb{R}^d$ be a set of points and $C \geq 1$ such that*

$$\mathrm{Cost}(P, B) \leq C \cdot \mathrm{Cost}(P, k).$$

Then there exists an ϵ-coreset for the k-median problem on P of size

$$O\left(\frac{C^2\, d\, k \log k}{\epsilon^2} + |B|\right).$$

Note that $|B|$ could be larger than k. Furthermore B is only a C-approximation to $\mathrm{Cost}(P, k)$, where C can be large, even a function of k and n. Theorem 15.11 shows that this approximate solution is already sufficient to get a small ϵ-coreset for P.

Overview of ideas. The proof of Theorems 15.10 and 15.11 rely on the following two insights. Let X be any set of k points in $\mathbb{R}^d$.

Relation to ϵ-approximations: Rewrite Equation (15.9) to get

$$\left|\sum_{p\in P} \mathrm{dist}\,(p, X) - \sum_{p\in A} \mathrm{dist}\,(p, X) \cdot w\,(p)\right| \leq \epsilon \cdot \sum_{p\in P} \mathrm{dist}\,(p, X)\,.$$

This resembles the notion of ϵ-approximations. Indeed, applying Corollary 15.6 to P with functions $f_p\,(X) = \mathrm{dist}\,(p, X)$ for each $p \in P$, an ϵ-approximation $A \subseteq P$ of the set system induced on P by the union of k balls in $\mathbb{R}^d$ satisfies

$$\left|\sum_{p\in P} \mathrm{dist}(p, X) - \sum_{p\in A} \mathrm{dist}(p, X) \frac{|P|}{|A|}\right| \leq 3\,\epsilon \cdot |P| \cdot \max_{p\in P} \mathrm{dist}(p, X).$$

The set A would be an $O\,(\epsilon)$-coreset, with weight function $w(p) = \frac{|P|}{|A|}$, if for each X,

$$|P| \cdot \max_{p\in P} \mathrm{dist}\,(p, X) = O\left(\sum_{q\in P} \mathrm{dist}\,(q, X)\right),$$

or equivalently, if for each p and each X,

$$\mathrm{dist}\,(p, X) = O\left(\frac{\sum_{q\in P} \mathrm{dist}\,(q, X)}{|P|}\right).$$

This is not the case, of course—each distance cannot be upper bounded by the average distance for *all* $X \subseteq \mathbb{R}^d$.

Weighted ϵ-approximations: The condition $\frac{\mathrm{dist}(p,X)}{\sum_{q\in P} \mathrm{dist}(q,X)} = O(\frac{1}{|P|})$ suggests that one should construct an ϵ-approximation according to a weight distribution, which can then be set depending on the relative values of $\mathrm{dist}\,(p, \cdot)$. This idea, sometimes called *importance sampling*, is thematically very similar to the idea in Theorem 8.20. Specifically, consider the following weighted version of Corollary 15.6.

LEMMA 15.12. *Let P be a set of n points in $\mathbb{R}^d$ and k a positive integer. For each $p \in P$ we are given a rational weight m_p and a function $f_p\colon \left(\mathbb{R}^d\right)^k \to \mathbb{R}^+$. These functions define a set system $(P, \mathcal{R})$ with*

$$\mathcal{R} = \left\{P_{X,r}\colon X \subseteq \mathbb{R}^d, |X| = k \text{ and } r \in \mathbb{R}\right\}, \text{ where } P_{X,r} = \left\{p \in P\colon f_p(X) \le r\right\}.$$

Then given $\epsilon > 0$ there exists a multiset $A \subseteq P$ of size $O\left(\frac{\text{VC-dim}(\mathcal{R})}{\epsilon^2}\right)$ such that for any $X \in \left(\mathbb{R}^d\right)^k$,

$$\left|\frac{\sum_{p\in P} f_p(X)}{\sum_{p\in P} m_p} - \frac{\sum_{p\in A} \frac{f_p(X)}{m_p}}{|A|}\right| \le 3\epsilon \left(\max_{p\in P} \frac{f_p(X)}{m_p}\right). \tag{15.13}$$

PROOF. By scaling up, we can assume that each m_p is an integer. Let P' be the set constructed by adding m_p copies of each $p \in P$ to P', where each copy of p is assigned the function $\frac{f_p(X)}{m_p}$. By applying Corollary 15.6 to P', there exists a set A such that for all $X \in \left(\mathbb{R}^d\right)^k$,

$$\left|\frac{\sum_{p'\in P'} f_{p'}(X)}{|P'|} - \frac{\sum_{p'\in A} f_{p'}(X)}{|A|}\right| \le 3\epsilon \left(\max_{p'\in P'} f_{p'(X)}\right). \tag{15.14}$$

Noting that $\sum_{p'\in P'} f_{p'}(X) = \sum_{p\in P} f_p(X)$, Equation (15.14) is equivalent to Equation (15.13).
Finally, as the VC-dimension is unchanged by adding duplicate elements, Theorem 13.2 implies that $|A| = O\left(\frac{\text{VC-dim}(\mathcal{R})}{\epsilon^2}\right)$. □

In the proof of Theorems 15.10 and 15.11 we will set the parameters m_p, $f_p(\cdot)$ and ϵ' such that an ϵ'-approximation A given by Lemma 15.12 can be used to construct the required ϵ-coreset.

Given our preparation, the proof of our first main result is immediate.

THEOREM 15.10. *Let P be a set of n points in $\mathbb{R}^d$ and $k \in \mathbb{Z}^+$, $\epsilon > 0$ be two given parameters. Define*

$$S = \sum_{p\in P} \sup_{\substack{Y\subseteq \mathbb{R}^d\\ |Y|=k}} \frac{\text{dist}(p, Y)}{\sum_{q\in P} \text{dist}(q, Y)}.$$

Then there exists an ϵ-coreset $A \subseteq P$ of size $O\left(\frac{S^2\, d\, k \log k}{\epsilon^2}\right)$ for the k-median problem on P.

PROOF. Set $f_p(X) = \text{dist}(p, X)$ for each $p \in P$ and let m_p be the weight of $p \in P$. These weights will be set later and normalized so that $\sum_{p\in P} m_p = 1$. Further let $\epsilon' > 0$ be a parameter to be set later.
Let A be an ϵ'-approximation given by Lemma 15.12 applied to P with weights m_p and functions $f_p(\cdot)$. That is, for each set X of k points in $\mathbb{R}^d$, A satisfies

$$\left|\sum_{p\in P} \text{dist}(p, X) - \sum_{p\in A} \text{dist}(p, X) \cdot \frac{1}{|A|\, m_p}\right| \le 3\epsilon' \left(\max_{p\in P} \frac{\text{dist}(p, X)}{m_p}\right). \tag{15.15}$$

For A to be an ϵ-coreset, the right-hand side of the above inequality must be upper bounded by $\epsilon \cdot \sum_{p \in P} \operatorname{dist}(p, X)$. The natural choice is to set $m_p = \frac{\operatorname{dist}(p,X)}{\sum_{q \in P} \operatorname{dist}(q,X)}$ for each $p \in P$ and $\epsilon' = \frac{\epsilon}{3}$. However m_p must be independent of X, as A must work for all choices of X! Considering the worst-case bound for m_p over all choices of X leads to the notion of the *sensitivity* of a point.

DEFINITION 15.16. The sensitivity of each $p \in P$ with respect to $\{f_p : p \in P\}$ is defined to be

$$s(p) = \sup_{\substack{Y \subseteq \mathbb{R}^d \\ |Y| = k}} \frac{f_p(Y)}{\sum_{q \in P} f_q(Y)}.$$

Let $S = \sum_{p \in P} s\,(p)$ and set

$$f_p(X) = \operatorname{dist}(p, X) \qquad \text{and} \qquad m_p = \frac{s\,(p)}{S}[3].$$

From Equation (15.15),

$$\left| \sum_{p \in P} \operatorname{dist}(p, X) - \sum_{p \in A} \operatorname{dist}(p, X) \cdot \frac{1}{|A|\, m_p} \right| \leq 3\epsilon' \cdot S \cdot \left(\max_{p \in P} \frac{\operatorname{dist}(p, X)}{s(p)} \right)$$

The R.H.S., after substituting for $s(p)$ and multiplying/dividing by $\sum_{q \in P} \operatorname{dist}(q, X)$:

$$3\epsilon' S \sum_{q \in P} \operatorname{dist}(q, X) \cdot \left(\max_{p \in P} \frac{\frac{\operatorname{dist}(p,X)}{\sum_{q \in P} \operatorname{dist}(q,X)}}{\sup_{\substack{Y \subseteq \mathbb{R}^d \\ |Y| = k}} \frac{\operatorname{dist}(p,Y)}{\sum_{q \in P} \operatorname{dist}(q,Y)}} \right) \leq 3\epsilon' S \left(\sum_{q \in P} \operatorname{dist}(q, X) \right).$$

Setting $\epsilon' = \frac{\epsilon}{3S}$ implies that A is an ϵ-coreset where each $p \in A$ is assigned the weight $\frac{1}{|A|\, m_p}$.

Note that A is an ϵ-approximation of the set system induced on P by the union of k balls in $\mathbb{R}^d$; that is, the set system induced by the k-fold union of balls in $\mathbb{R}^d$. Lemma 10.3 and Theorem 11.6 implies that the VC-dimension of this set system is $\Theta\,(dk \log k)$ and thus $|A| = O\left(\frac{S^2 dk \log k}{\epsilon^2}\right)$ by Lemma 15.12. □

We remark here that to compute A above, we need to compute the weights m_p for each $p \in P$. This, together with the problem of deriving good upper bounds on the total sensitivity S, is a non-trivial algorithmic problem by itself (see discussion).

[3]The division by S is just to get $\sum_p m_p = 1$.

THEOREM 15.11. *Let P be a set of n points in $\mathbb{R}^d$ and $k \in \mathbb{Z}^+, \epsilon > 0$ be two given parameters. Further let $B \subseteq \mathbb{R}^d$ be a set of points and $C \geq 1$ such that*

$$\mathrm{Cost}(P, B) \leq C \cdot \mathrm{Cost}(P, k).$$

Then there exists an ϵ-coreset for the k-median problem on P of size

$$O\left(\frac{C^2\, d\, k \log k}{\epsilon^2} + |B|\right).$$

PROOF. Given B, set the parameters for each $p \in P$ as follows:

$$f_p(X) = \mathrm{dist}(p, X) - \mathrm{dist}(\mathrm{closest}(p, B), X) + \mathrm{dist}\,(p, B)\,,$$
$$m_p = \frac{\mathrm{dist}(p, B)}{\sum_{q \in P} \mathrm{dist}(q, B)},$$

where $\mathrm{closest}\,(p, Q) = \arg\min_{q \in Q} \mathrm{dist}\,(p, q)$ denotes the closest point in Q to p. Note that $f_p\,(X)$ is non-negative[4] due to triangle inequality:

$$\mathrm{dist}(\mathrm{closest}(p, B), X) \leq \mathrm{dist}\,(\mathrm{closest}(p, B), p) + \mathrm{dist}(p, X).$$

Applying Lemma 15.12 with f_p and m_p set above and noting that $\sum_{p \in P} m_p = 1$, we get an ϵ'-approximation A such that for any $X \subseteq \mathbb{R}^d$ of k points,

$$\Bigg|\sum_{p \in P}\Big(\mathrm{dist}(p, X) - \mathrm{dist}(\mathrm{closest}(p, B), X) + \mathrm{dist}\,(p, B)\Big) -$$
$$\sum_{p \in A}\Big(\mathrm{dist}\,(p, X) - \mathrm{dist}\,(\mathrm{closest}\,(p, B)\,, X) + \mathrm{dist}\,(p, B)\Big) \cdot \frac{1}{|A|\, m_p}\Bigg|$$
$$\leq 3\epsilon' \left(\max_{p \in P} \frac{\mathrm{dist}(p, X) - \mathrm{dist}(\mathrm{closest}(p, B), X) + \mathrm{dist}\,(p, B)}{m_p}\right).$$

Using the fact that

$$\sum_{p \in P} \mathrm{dist}(p, B) = \sum_{p \in A}\left(\mathrm{dist}(p, B)\frac{\sum_{q \in P} \mathrm{dist}(q, B)}{|A|\; \mathrm{dist}(p, B)}\right) = \sum_{p \in A} \mathrm{dist}(p, B)\frac{1}{|A|\, m_p},$$

we arrive at

$$\Bigg|\left(\sum_{p \in P} \mathrm{dist}(p, X)\right) - \left(\sum_{p \in P} \mathrm{dist}(\mathrm{closest}(p, B), X)\right)$$
$$- \left(\sum_{p \in A} \mathrm{dist}(p, X)\frac{1}{|A|\, m_p}\right) + \left(\sum_{p \in A} \mathrm{dist}(\mathrm{closest}(p, B), X)\frac{1}{|A|\, m_p}\right)\Bigg|$$
$$\leq 3\epsilon' \left(\max_{p \in P} \frac{\mathrm{dist}(p, X) - \mathrm{dist}(\mathrm{closest}(p, B), X) + \mathrm{dist}\,(p, B)}{m_p}\right). \tag{15.17}$$

R.H.S.: Using a consequence of triangle inequality, that

$$\mathrm{dist}(p, X) - \mathrm{dist}(\mathrm{closest}(p, B), X) \leq \mathrm{dist}(p, \mathrm{closest}(p, B)) = \mathrm{dist}(p, B),$$

[4]Indeed, the additive term $\mathrm{dist}\,(p, B)$ is present in $f_p(X)$ just to make $f_p(X)$ non-negative so that one can apply Lemma 15.12. *Conceptually* we only need $f_p(X) = \mathrm{dist}(p, X) - \mathrm{dist}\,(\mathrm{closest}(p, B), X)$.

as well as substituting the value of m_p, the R.H.S. of Equation (15.17) is at most

$$3\epsilon' \max_{p\in P} \frac{2\,\mathrm{dist}(p,B)}{\frac{\mathrm{dist}(p,B)}{\sum_{q\in P}\mathrm{dist}(q,B)}} = 6\epsilon' \sum_{p\in P} \mathrm{dist}(p,B)$$
$$\leq 6\epsilon' \cdot C \cdot \mathrm{Cost}\,(P,k) \leq 6\epsilon' C \sum_{p\in P} \mathrm{dist}(p,X).$$

L.H.S.: For each $b \in B$, let P_b be the set of points of P whose closest point in B is b. Then the L.H.S. of Equation (15.17) becomes

$$(15.18) \quad \left| \left(\sum_{p\in P} \mathrm{dist}(p,X)\right) - \left(\sum_{b\in B} |P_b| \cdot \mathrm{dist}(b,X)\right) - \left(\sum_{p\in A} \mathrm{dist}(p,X) \frac{1}{|A|\, m_p}\right) \right.$$
$$\left. + \left(\sum_{b\in B} \sum_{p\in P_b\cap A} \mathrm{dist}(b,X) \frac{1}{|A|\, m_p}\right) \right|.$$

We're done—set $\epsilon' = \frac{\epsilon}{6C}$ and return $A \cup B$ as our ϵ-coreset, with weights dictated by Equation (15.18):

$$p \in A \;:\; w(p) = \frac{1}{|A|\, m_p}.$$
$$b \in B \;:\; w(b) = |P_b| - \sum_{p\in A\cap P_b} \frac{1}{|A|\, m_p}.$$

□

We remark that to compute the set A one again needs to compute the weights m_p for all $p \in P$. However this time it is easier as we are also given the set B.

Bibliography and discussion. The beautiful application of this section is from [**FL11**]. See [**VX12**] for upper bounds on sensitivity for a variety of optimization problems. There are many interesting variations, improvements and applications of the basic ideas presented in this section (e.g., see [**HV20**]). We refer the reader to the surveys [**AHPV07**, **Phi18**] for more information on coresets.

[AHPV07] P. K. Agarwal, S. Har-Peled, and K. R. Varadarajan, *Geometric approximation via coresets*, Combinatorial and computational geometry, Math. Sci. Res. Inst. Publ., vol. 52, Cambridge Univ. Press, Cambridge, 2005, pp. 1–30, DOI 10.4171/PRIMS/172. MR2178310

[FL11] D. Feldman and M. Langberg, *A unified framework for approximating and clustering data*, STOC'11—Proceedings of the 43rd ACM Symposium on Theory of Computing, ACM, New York, 2011, pp. 569–578, DOI 10.1145/1993636.1993712. MR2932007

[HV20] L. Huang and N. K. Vishnoi, *Coresets for clustering in euclidean spaces: importance sampling is nearly optimal*, STOC '20—Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, ACM, New York, 2020, pp. 1416–1429. MR4141850

[Phi18] J. M. Phillips. *Coresets and sketches. Handbook of Discrete and Computational Geometry.* CRC Press, 2018. pp. 1269–1286.

[VX12] K. R. Varadarajan and X. Xiao. *On the sensitivity of shape fitting problems. Proceedings of the IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS).* Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2012. pp. 486–497.