Chapter 10

Depth Estimation via Sampling

By Sariel Har-Peled, March 1, 2023[1]

> "Maybe the Nazis told the truth about us. Maybe the Nazis were the truth. We shouldn't forget that: perhaps they were the truth. The rest, just beautiful lies. We've sung many beautiful lies about ourselves. Perhaps that's what I'm trying to do – to sing another beautiful lie."
>
> The roots of heaven, Romain Gary

In this chapter, we introduce a "trivial" yet powerful idea. Given a set S of objects and a point p that is contained in some of the objects, let p's ***weight*** be the number of objects that contain it. We can estimate the depth/weight of p by counting the number of objects that contain it in a random sample of the objects. In fact, by considering points induced by the sample, we can bound the number of "light" vertices induced by S (i.e., vertices that are not contained in many objects). This idea can be extended to bounding the number of "light" configurations induced by a set of objects.

This approach leads to a sequence of short, beautiful, elegant, and correct[2] proofs of several hallmark results in discrete geometry.

While the results in this chapter are not directly related to approximation algorithms, the insights and general approach will be useful for us later, or so one hopes.

10.1. The at most k-levels

Let L be a set of n lines in the plane. A point $p \in \bigcup_{\ell \in \mathsf{L}} \ell$ is of ***level*** k if there are k lines of L strictly below it. The ***k-level*** is the closure of the set of points of level k. Namely, the k-level is an x-monotone curve along the lines of L.

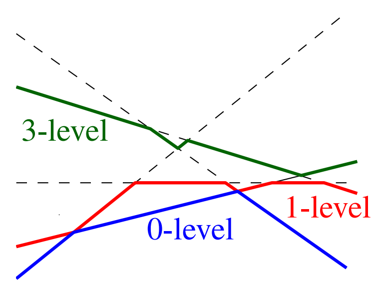

The 0-level is the boundary of the "bottom" face of the arrangement of L (i.e., the face containing the negative y-axis). It is easy to verify that the 0-level has at most $n-1$ vertices, as each line might contribute at most one segment to the 0-level (which is an unbounded convex polygon).

It is natural to ask what the number of vertices at the k-level is (i.e., what the combinatorial complexity of the polygonal chain forming the k-level is). This is a surprisingly hard question, but the same question on the complexity of the at most k-level is considerably easier.

Theorem 10.1.1. *The number of vertices of level at most k in an arrangement of n lines in the plane is $O(nk)$.*

Proof: Pick a random sample R of L, by picking each line to be in the sample with probability $1/k$. Observe that

$$\mathbb{E}\Big[|\mathsf{R}|\Big] = \frac{n}{k}.$$

[2]The saying goes that "hard theorems have short, elegant, and incorrect proofs". This chapter can maybe serve as a counterexample to this claim.

Let $\mathbb{L}_{\leq k} = \mathbb{L}_{\leq k}(\mathsf{L})$ be the set of all vertices of $\mathcal{A}(\mathsf{L})$ of level at most k, for $k > 1$. For a vertex $p \in \mathbb{L}_{\leq k}$, let X_p be an indicator variable which is 1 if p is a vertex of the 0-level of $\mathcal{A}(\mathsf{R})$. The probability that p is in the 0-level of $\mathcal{A}(\mathsf{R})$ is the probability that none of the j lines below it are picked to be in the sample, and the two lines that define it do get selected to be in the sample. Namely,

$$\mathbb{P}[X_p = 1] = \left(1 - \frac{1}{k}\right)^j \left(\frac{1}{k}\right)^2 \geq \left(1 - \frac{1}{k}\right)^k \frac{1}{k^2} \geq \exp\left(-2\frac{k}{k}\right)\frac{1}{k^2} = \frac{1}{e^2k^2}$$

since $j \leq k$ and $1 - x \geq e^{-2x}$, for $0 < x \leq 1/2$.

On the other hand, the number of vertices on the 0-level of R is at most $|\mathsf{R}| - 1$. As such,

$$\sum_{p \in \mathbb{L}_{\leq k}} X_p \leq |\mathsf{R}| - 1.$$

Moreover this, of course, also holds in expectation, implying

$$\mathbb{E}\left[\sum_{p \in \mathbb{L}_{\leq k}} X_p\right] \leq \mathbb{E}\left[|\mathsf{R}| - 1\right] \leq \frac{n}{k}.$$

On the other hand, by linearity of expectation, we have

$$\mathbb{E}\left[\sum_{p \in \mathbb{L}_{\leq k}} X_p\right] = \sum_{p \in \mathbb{L}_{\leq k}} \mathbb{E}[X_p] \geq \frac{|\mathbb{L}_{\leq k}|}{e^2k^2}.$$

Putting these two inequalities together, we get that $\dfrac{|\mathbb{L}_{\leq k}|}{e^2k^2} \leq \dfrac{n}{k}$. Namely, $|\mathbb{L}_{\leq k}| \leq e^2nk$. ■

The connection to depth is simple. Every line defines a halfplane (i.e., the region above the line). A vertex of depth at most k is contained in at most k halfplanes. The above proof (intuitively) first observed that there are at most n/k vertices of the random sample of zero depth (i.e., 0-level of R) and then showed that every such vertex has probability (roughly) $1/k^2$ to have depth zero in the random sample. It thus follows that if the number of vertices of level at most k is μ, then $\mu/k^2 \leq n/k$; namely, $\mu = O(nk)$.

10.2. The crossing lemma

In the following, for a graph G we denote by $n = |V(G)|$ and $m = |E(G)|$ the number of vertices and edges of G, respectively.

A graph G is ***planar*** if it can be drawn in the plane so that none of its edges are crossing.

We state the following well-known fact.

Theorem 10.2.1 (Euler's formula). *For a connected planar graph G, we have $f - m + n = 2$, where f, m, and n are the number of faces, edges, and vertices in a planar drawing of G.*

Lemma 10.2.2. *If G is a planar graph, then $m \leq 3n - 6$.*

Proof: We assume that the number of edges of G is maximal (i.e., no edges can be added without introducing a crossing). If it is not maximal, then add edges till it becomes maximal. This implies that G is a triangulation (i.e., every face is a triangle). Then, every face is adjacent to three edges, and as such $2m = 3f$. By Euler's formula, we have $f - m + n = (2/3)m - m + n = 2$. Namely, $-m + 3n = 6$. Alternatively, $m = 3n - 6$. However, if m is not maximal, this equality deteriorates to the required inequality. ■

For example, the above inequality implies that the complete graph over five vertices (i.e., K_5) is not planar. Indeed, it has $m = \binom{5}{2} = 10$ edges and $n = 5$ vertices, but if it were planar, the above inequality would imply that $10 = m \leq 3n - 6 = 9$, which is of course false. (The reader can amuse himself or herself by trying to prove that $K_{3,3}$, the bipartite complete graph with three vertices on each side, is not planar.)

Kuratowski's celebrated theorem states that a graph is planar if and only if it does not contain either K_5 or $K_{3,3}$ induced inside it (formally, it does not have K_5 or $K_{3,3}$ as a minor).

For a graph G, we define the crossing number of G, denoted as $c(G)$, as the minimal number of edge crossings in any drawing of G in the plane. For a planar graph $c(G)$ is zero, and it is "larger" for "less planar" graphs.

Claim 10.2.3. *For a graph G, we have $c(G) \geq m - 3n + 6$.*

Proof: If $m - 3n + 6 \leq 0 \leq c(G)$, then the claim holds trivially. Otherwise, the graph G is not planar by Lemma 10.2.2. Draw G in such a way that $c(G)$ is realized and assume, for the sake of contradiction, that $c(G) < m - 3n + 6$. Let H be the graph resulting from G, by removing one of the edges from each pair of edges of G that intersects in the drawing. We have $m(H) \geq m(G) - c(G)$. But H is planar (since its drawing has no crossings), and by Lemma 10.2.2, we have $m(H) \leq 3n(H) - 6$, or equivalently, $m(G) - c(G) \leq 3n - 6$. Namely, $m - 3n + 6 \leq c(G)$, which contradicts our assumption. ■

Lemma 10.2.4 (Crossing lemma). *For a graph G, such that $m \geq 6n$, we have $c(G) = \Omega(m^3/n^2)$.*

Proof: We consider a specific drawing D of G in the plane that has $c(G)$ crossings. Next, let U be a random subset of $V(G)$ selected by choosing each vertex to be in the sample with probability $p > 0$.

Let $H = G_U = (U, E')$ be the induced subgraph of G over U. Here, only edges of G with both their endpoints in U "survive" in H; that is, $E' = \left\{ uv \,\middle|\, uv \in E(G) \text{ and } u, v \in U \right\}$.

Thus, the probability of a vertex v to survive in H is p. The probability of an edge of G to survive in H is p^2, and the probability of a crossing (in this specific drawing D) to survive in the induced drawing D_H (of H) is p^4. Let X_v and X_e denote the (random variables which are the) number of vertices and edges surviving in H, respectively. Similarly, let X_c be the number of crossings surviving in D_H. By Claim 10.2.3, we have

$$X_c \geq c(H) \geq X_e - 3X_v + 6.$$

In particular, this holds in the expectation, and by linearity of expectation, we have that

$$\mathbb{E}[X_c] \geq \mathbb{E}[X_e] - 3\,\mathbb{E}[X_v] + 6.$$

Now, by linearity of expectation, we have that $\mathbb{E}[X_c] = c(G)p^4$, $\mathbb{E}[X_e] = mp^2$, and $\mathbb{E}[X_v] = np$, where m and n are the number of edges and vertices of G, respectively. This implies that

$$c(G)p^4 \geq mp^2 - 3np + 6.$$

In particular, $c(G) \geq m/p^2 - 3n/p^3 + 6/p^4 \geq m/p^2 - 3n/p^3$. In particular, setting $p = 6n/m \leq 1$, we have that

$$c(G) \geq \frac{m}{p^2} - \frac{3n}{p^3} = \frac{m^3}{6^2 n^2} - \frac{m^3}{2 \cdot 6^2 n^2} = \frac{m^3}{72n^2}.$$ ■

Surprisingly, despite its simplicity, Lemma 10.2.4 is a very strong tool, as the following results testify.

10.2.1. On the number of incidences

Let P be a set of n distinct points in the plane, and let L be a set of m distinct lines in the plane (note that all the lines might pass through a common point, as we do not assume general position here). Let $I(P, \mathsf{L})$ denote the number of point/line pairs (p, ℓ), where $p \in P$, $\ell \in \mathsf{L}$, and $p \in \ell$. The number $I(P, \mathsf{L})$ is the number of ***incidences*** between lines of L and points of P. Let $I(n, m) = \max_{|P|=n, |\mathsf{L}|=m} I(P, \mathsf{L})$.

The following "easy" result has a long history and required major effort to prove before the following elegant proof was discovered[3].

Lemma 10.2.5. *The maximum number of incidences between n points and m lines is*

$$I(n, m) = O\Big(n^{2/3}m^{2/3} + n\Big).$$

Proof: Let P and L be the set of n points and the set of m lines, respectively, realizing $I(m, n)$. Let G be a graph over the points of P (we assume that P contains an additional point at infinity). We connect two points if they lie consecutively on a common line of L, and we also connect the first and last points on each line with the point at infinity.

Clearly, $m(G) = I + m$ and $n(G) = n + 1$, where $I = I(m, n)$. Now, we can interpret the arrangement of lines $\mathcal{A}(\mathsf{L})$ as a drawing of G, where a crossing of two edges of G is just a vertex of $\mathcal{A}(\mathsf{L})$. We conclude that $c(G) \leq m^2$, since m^2 is a trivial bound on the number of vertices of $\mathcal{A}(\mathsf{L})$. On the other hand, by Lemma 10.2.4, if $I + m = m(G) \geq 6n(G) = 6(n+1)$, we have $c(G) \geq c_1 m(G)^3/n(G)^2$, where c_1 is some constant. Thus,

$$c_1 \frac{(I+m)^3}{(n+1)^2} = c_1 \frac{m(G)^3}{n(G)^2} \leq c(G) \leq m^2.$$

Thus, we have $I = O\Big(m^{2/3}n^{2/3} - m\Big) = O\Big(m^{2/3}n^{2/3}\Big)$.

The other possibility, that $m(G) < 6n(G)$, implies that $I + m \leq 6(n+1)$. That is, $I = O(n)$. Putting the two inequalities together, we have that $I = O\Big(n^{2/3}m^{2/3} + n\Big)$. ■

Lemma 10.2.6. $I(n, n) = \Omega\Big(n^{4/3}\Big)$.

Proof: For a positive integer k, let $[\![k]\!] = \{1, \ldots, k\}$. For the sake of simplicity of exposition, assume that $N = n^{1/3}/2$ is an integer.

Consider the integer point set $P = [\![N]\!] \times [\![8N^2]\!]$ and the set of lines

$$\mathsf{L} = \Big\{ y = ax + b \;\Big|\; a \in [\![2N]\!] \text{ and } b \in [\![4N^2]\!] \Big\}.$$

Clearly, $|P| = n$ and $|\mathsf{L}| = n$.

Now, for any $x \in [\![N]\!]$, $a \in [\![2N]\!]$, and $b \in [\![4N^2]\!]$, we have that

$$y = ax + b \leq 2N \cdot N + 4N^2 \leq 8N^2.$$

Namely, any line in L is incident to a point along each vertical line of the grid of P. As such, every line of L is incident to N points of P. Namely, the total number of incidences between the points of P and L is $|\mathsf{L}|\, N = n \cdot n^{1/3}/2 = n^{4/3}/2$, which implies the claim. ■

[3] Or invented – I have no dog in this fight.

10.2.2. On the number of k-sets

Let P be a set of n points in the plane in general position (i.e., no three points are collinear). A pair of points $p, q \in P$ form a ***k-set*** if there are exactly k points in the (closed) halfplane below the line passing through p and q. Consider the graph $G = (P, E)$ that has an edge for every k-set. We will be interested in bounding the size of E as a function of n. Observe that via duality we have that the number of k-sets is exactly the complexity of the k-level in the dual arrangement $\mathcal{A} = \mathcal{A}(P^\star)$; that is, $P^\star$ is a set of lines, and every k-set of P corresponds to a vertex on the k-level of $\mathcal{A}$.

Lemma 10.2.7 (Antipodality). *Let qp and sp be two k-set edges of G, with q and s to the left of p. Then there exists a point $t \in P$ to the right of p such that pt is a k-set, and* $\mathrm{line}(p, t)$ *lies between* $\mathrm{line}(q, p)$ *and* $\mathrm{line}(s, p)$.

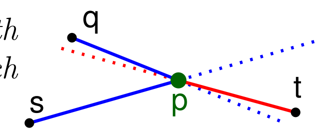

Proof: Let $f(\alpha)$ be the number of points below or on the line passing through p and having slope α, where α is a real number. Rotating this line counterclockwise around p corresponds to increasing α. In the following, let $f_+(\alpha)$ (resp., $f_-(\alpha)$) denote the value of $f(\cdot)$ just to the right (resp., left) of α; formally, $f_-(\alpha) = \lim_{x \to \alpha, x < \alpha} f(x)$ and $f_+(\alpha) = \lim_{x \to \alpha, x > \alpha} f(x)$.

Any point swept over by this line which is to the right of p increases f, and any point swept over to the left of p decreases f by 1.

Let α_q and α_s be the slope of the lines containing qp and sp, respectively. Assume, for the sake of simplicity of exposition, that $\alpha_q < \alpha_s$. Clearly, $f(\alpha_q) = f(\alpha_s) = k$ and $f_+(\alpha_q) = k - 1$. Let y be the smallest value such that $y > \alpha_q$ and $f(y) = k$. Such a y exists since $\alpha_q < \alpha_s$, $f_+(\alpha_q) = k - 1$, $f(\alpha_s) = k$, and there are no three points that are collinear.

We have that $f_-(y) = k - 1$, which implies that the line passing through p with slope $f(y)$ has a point $t \in P$ on it and t is to the right of p. Clearly, if we continue sweeping, the line would sweep over sp, which implies the claim. ■

Lemma 10.2.7 also holds by symmetry in the other direction: Between any two edges to the right of p, there is an antipodal edge on the other side.

Lemma 10.2.8. *Let p be a point of P, and let q be a point to its left, such that $qp \in E(G)$ and it has the largest slope among all such edges. Furthermore, assume that there are $k - 1$ points of P to the right of p. Then, there exists a point $s \in P$, such that $ps \in E(G)$ and ps has larger slope than qp.*

Proof: Let α be the slope of qp, and observe that $f(\alpha) = k$ and $f_+(\alpha) = k - 1$ and $f(\infty) \geq k$. Namely, there exists $y > \alpha$ such that $f(y) = k$. We conclude that there is a k-set adjacent to p on the right, with slope larger than α. ■

So, imagine that we are at an edge $\mathsf{e} = qp \in E(G)$, where q is to the left of p. We rotate a line around p (counterclockwise) till we encounter an edge $\mathsf{e}' = ps \in E(G)$, where s is a point to the right of p. We can now walk from e to e' and continue walking in this way, forming a chain of edges in G. Note that by Lemma 10.2.7, no two such chains can be "merged" into using the same edge. Furthermore, by Lemma 10.2.8, such a chain can end only in the last $k - 1$ points of P (in their ordering along the x-axis). Namely, we decomposed the edges of G into $k - 1$ edge disjoint convex chains (the chains are convex since we rotate counterclockwise as we walk along a chain). See Figure 10.1 for an example.

Lemma 10.2.9. *The edges of G can be decomposed into $k - 1$ convex chains $\mathsf{C}_1, \ldots, \mathsf{C}_{k-1}$.*

Similarly, the edges of G can be decomposed into $m = n - k + 1$ concave chains $\mathsf{D}_1, \ldots, \mathsf{D}_m$.

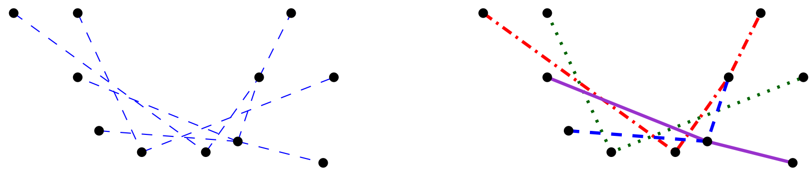

Figure 10.1: An example of 5-sets and their decomposition into four convex chains.

Proof: The first part of the claim is proved above. As for the second claim, rotate the plane by 180°. Every k-set is now an $(n-k+2)$-set, and by the above argumentation, the edges of G can be decomposed into $n-k+1$ convex chains, which are concave in the original orientation. (The figure on the right shows these concave chains for $k = 5$ for the example used above.) ■

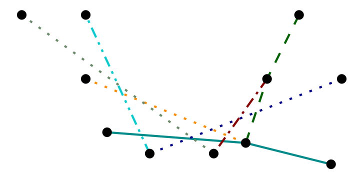

Theorem 10.2.10. *The number of k-sets defined by a set of n points in the plane is $O\left(nk^{1/3}\right)$.*

Proof: The graph G has $n = |P|$ vertices, and let $m = |E(G)|$ be the number of k-sets. By Lemma 10.2.9, any crossing of two edges of G is an intersection point of one convex chain of $\mathsf{C}_1, \ldots, \mathsf{C}_{k-1}$ with a concave chain of $\mathsf{D}_1, \ldots, \mathsf{D}_{n-k+1}$. Since a convex chain and a concave chain can have at most two intersections, we conclude that there are at most $2(k-1)(n-k+1)$ crossings in G. By the crossing lemma (Lemma 10.2.4), there are at least $\Omega(m^3/n^2)$ crossings. Putting these two inequalities together, we conclude $m^3/n^2 = O(nk)$, which implies $m = O\left(nk^{1/3}\right)$. ■

10.3. A general bound for the at most k-weight

We now extend the at most k-level technique to the general moments technique setting. We quickly restate the abstract settings.

Let S be a set of objects. For a subset $\mathsf{R} \subseteq \mathsf{S}$, we define a collection of 'regions' called $\mathcal{F}(\mathsf{R})$. Let $\mathcal{T} = \mathcal{T}(\mathsf{S}) = \bigcup_{\mathsf{R} \subseteq \mathsf{S}} \mathcal{F}(\mathsf{R})$ denote the set of ***all possible regions*** defined by subsets of S. We associate two subsets $D(\sigma), K(\sigma) \subseteq \mathsf{S}$ with each region $\sigma \in \mathcal{T}$. The ***defining set*** $D(\sigma)$ of σ is a subset of S defining the region σ. We assume that for every $\sigma \in \mathcal{T}$, $|D(\sigma)| \leq d$ for a (small) constant d, which is the ***combinatorial dimension***. The ***conflicting set*** $K(\sigma)$ of σ is the set of objects of S such that including any object of $K(\sigma)$ into R prevents σ from appearing in $\mathcal{F}(\mathsf{R})$. The ***weight*** of σ is $\omega(\sigma) = |K(\sigma)|$.

Let $\mathsf{S}, \mathcal{F}(\mathsf{R}), D(\sigma)$, and $K(\sigma)$ be such that for any subset $\mathsf{R} \subseteq \mathsf{S}$, the set $\mathcal{F}(\mathsf{R})$ satisfies the following axioms: (i) For any $\sigma \in \mathcal{F}(\mathsf{R})$, we have $D(\sigma) \subseteq \mathsf{R}$ and $\mathsf{R} \cap K(\sigma) = \emptyset$. (ii) If $D(\sigma) \subseteq \mathsf{R}$ and $K(\sigma) \cap \mathsf{R} = \emptyset$, then $\sigma \in \mathcal{F}(\mathsf{R})$.

Let $\mathcal{T}_{\leq k}(\mathsf{S})$ be the set of regions of $\mathcal{T}$ with weight at most k. Furthermore, assume that the expected number of regions of zero weight of a sample of size r is (or is at most) $\mathbf{f}_0(r)$. Formally, for a sample of size r from S, we denote $\mathbf{f}_0(r) = \mathbb{E}\left[|\mathcal{F}(\mathsf{R})|\right]$. We have the following theorem.

Theorem 10.3.1. *Let S be a set of n objects as above, with combinatorial dimension d, and let k be a parameter. Let R be a random sample created by picking each element of S with probability $1/k$. Then, for some constant c, we have*

$$|\mathcal{T}_{\leq k}(\mathsf{S})| \leq c\, \mathbb{E}\left[k^d \mathbf{f}_0(|\mathsf{R}|)\right].$$

Proof: We reproduce the proof of Theorem 10.1.1. Every region $\sigma \in \mathcal{T}_{\leq k}$ appears in $\mathcal{F}(\mathsf{R})$ with probability $\geq 1/k^d(1-1/k)^k \geq e^{-2}/k^d$. Observe that every sample of size $|\mathsf{R}|$ has equal probability of being picked to be R. As such, we have that $\mathbf{f}_0(r) = \mathbb{E}\Big[|\mathcal{F}(\mathsf{R})| \;\Big|\; |\mathsf{R}| = r\Big]$.

Now, setting $X_\sigma = 1$ if and only if $\sigma \in \mathcal{F}(\mathsf{R})$, we have that

$$\mathbb{E}\Big[\mathbf{f}_0(|\mathsf{R}|)\Big] = \mathbb{E}\Big[\mathbb{E}\Big[|\mathcal{F}(\mathsf{R})| \;\Big|\; |\mathsf{R}| = k\Big]\Big] = \mathbb{E}\Big[|\mathcal{F}(\mathsf{R})|\Big] \geq \mathbb{E}\Big[\sum_{\sigma \in \mathcal{T}_{\leq k}} X_\sigma\Big] = \sum_{\sigma \in \mathcal{T}_{\leq k}} \mathbb{P}\Big[\sigma \in \mathcal{F}(\mathsf{R})\Big] \geq \frac{|\mathcal{T}_{\leq k}|}{k^d e^2}. \qquad \blacksquare$$

Lemma 10.3.2. *Let $\mathbf{f}_0(\cdot)$ be a monotone increasing function which is well behaved; namely, there exists a constant c, such that $\mathbf{f}_0(xr) \leq c\mathbf{f}_0(r)$, for any r and $1 \leq x \leq 2$. Let Y be the number of heads in n coin flips where the probability for head is $1/k$. Then $\mathbb{E}[\mathbf{f}_0(Y)] = O(\mathbf{f}_0(n/k))$.*

Proof: The claim[4] follows easily from Chernoff's inequality. Indeed, we have that

$$\mathbb{E}[Y] = n/k \qquad \text{and} \qquad \mathbb{P}[Y \geq t(n/k)] \leq 2^{-t(n/k)},$$

for $t > 3$, by the simplified form of the Chernoff inequality; see Lemma 10.6.1. Furthermore, by assumption we have that

$$f\Big((t+1)n/k\Big) \leq cf\Big(\tfrac{t+1}{2}(n/k)\Big) \leq c^{\lceil \lg(t+1) \rceil}\mathbf{f}_0(n/k).$$

Putting these two things together, we have that

$$\begin{aligned}\mathbb{E}\Big[\mathbf{f}_0(Y)\Big] &\leq \sum_i \mathbf{f}_0(i)\,\mathbb{P}[Y = i] \leq f\Big(10\frac{n}{k}\Big) + \sum_{t=10}^{k-1} f\Big((t+1)\frac{n}{k}\Big)\,\mathbb{P}[Y \geq t(n/k)] \\ &\leq O\Big(\mathbf{f}_0(n/k)\Big) + \sum_{t=10}^{k} c^{\lceil \lg(t+1) \rceil}\mathbf{f}_0(n/k)2^{-t(n/k)} = O\Big(\mathbf{f}_0(n/k)\Big). \qquad \blacksquare\end{aligned}$$

The following is an immediate consequence of Theorem 10.3.1 and Lemma 10.3.2.

Theorem 10.3.3. *Let S be a set of n objects with combinatorial dimension d, and let k be a parameter. Assume that the number of regions formed by a set of m objects is bounded by a function $\mathbf{f}_0(m)$ and, furthermore, $\mathbf{f}_0(m)$ is well behaved in the sense of Lemma 10.3.2. Then, $|\mathcal{T}_{\leq k}(\mathsf{S})| = O\Big(k^d\mathbf{f}_0(n/k)\Big)$.*

In particular, let $\mathbf{f}_{\leq k}(n) = \max_{|S|=n} |\mathcal{T}_{\leq k}(\mathsf{S})|$ be the maximum number of regions of weight at most k that can be defined by any set of n objects. We have that $\mathbf{f}_{\leq k}(n) = O\Big(k^d\mathbf{f}_0(n/k)\Big)$.

Note that if the function $\mathbf{f}_0(\cdot)$ grows polynomially, then Theorem 10.3.3 applies. It fails if $\mathbf{f}_0(\cdot)$ grows exponentially fast.

10.3.1. Example – k-level in higher dimensions

We need the following fact, which we state without proof.

Theorem 10.3.4 (The upper bound theorem). *The complexity of the convex hull of n points in d dimensions is bounded by $O\Big(n^{\lfloor \mathsf{d}/2 \rfloor}\Big)$.*

[4] Note that this lemma is not implied by Jensen's inequality (see Lemma 10.6.2$_{\text{p9}}$) since if $\mathbf{f}_0(\cdot)$ is convex, then $\mathbb{E}[\mathbf{f}_0(Y)] \geq \mathbf{f}_0(\mathbb{E}[Y]) = \mathbf{f}_0(n/k)$, which goes in the wrong direction.

Example 10.3.5 (At most k-sets). Let P be a set of n points in $\mathbb{R}^{\mathsf{d}}$. A region here is a halfspace with d points on its boundary. The set of regions defined by P is just the faces of the convex hull of P. The complexity of the convex hull of n points in d dimensions is $\mathbf{f}_0(n) = O\left(n^{\lfloor \mathsf{d}/2 \rfloor}\right)$, by Theorem 10.3.4. Two halfspaces h, h' would be considered to be combinatorially different if $P \cap h \neq P \cap h'$. As such, the number of combinatorially different halfspaces containing at most k points of P is at most $O\left(k^{\mathsf{d}} \mathbf{f}_0(n/k)\right) = O\left(k^{\lceil \mathsf{d}/2 \rceil} n^{\lfloor \mathsf{d}/2 \rfloor}\right)$.

10.4. Bibliographical notes

The reader should not mistake the simplicity of the proofs in this chapter with easiness. Almost all the results presented have a long and painful history with earlier proofs being considerably more complicated. In some sense, these results are the limit of mathematical evolution: They are simple, breathtakingly elegant (in some cases), and on their own (without exposure to previous work on the topic), it seems inconceivable that one could come up with them.

At most k-level. The technique for bounding the complexity of the at most k-level (or at most depth k) is generally attributed to Clarkson and Shor [CS89] and more precisely it is from [Cla88]. Previous work on just the two-dimensional variant include [GP84, Wel86, AG86]. Our presentation in Section 10.1 and Section 10.3 follows (more or less) Sharir [Sha03]. The connection of this technique to the crossing lemma is from there.

For a proof of the upper bound theorem (Theorem 10.3.4), see Matoušek [Mat02].

The crossing lemma. The crossing lemma is by Ajtai et al. [ACNS82] and Leighton [Lei84]. The current greatly simplified "proof from the book" is attributed to Sharir. The insight that this lemma has something to do with incidences and similar problems is due to Székely [Szé97]. Elekes [G E97] used the crossing lemma to prove surprising lower bounds on sum and product problems (the proofs of these bounds are surprisingly elegant).

The complexity of k-level and number of k-sets. This is considered to be one of the hardest problems in discrete geometry, and there is still a big gap between the best lower bound [Tót01] and best upper bound currently known [Dey98]. Our presentation in Section 10.2.2 follows suggestions by Micha Sharir and is based on the result of Dey [Dey98] (which was in turn inspired by the work of Agarwal et al. [AACS98]). This problem has a long history, and the reader is referred to Dey [Dey98] for its history.

Incidences. This problem again has a long and painful history. The reader is referred to [PS04, Szé97] for details. The elegant lower bound proof of Lemma 10.2.6 is by Elekes [Ele02].

We only skimmed the surface of some problems in discrete geometry and results known in this field related to incidences and k-sets. Good starting points for learning more are the books by Brass et al. [BMP05] and Matoušek [Mat02].

10.5. Exercises

Exercise 10.1 (Incidence lower bound). Prove that $I(m,n) = \Omega\left(m^{2/3} n^{2/3}\right)$ for any n and m.

Exercise 10.2 (The number of heavy disks). Let P be a set of n points in the plane. A disk d is ***canonical*** if its boundary passes through three points of P. Provide a bound on the number of canonical disks that contain at most k points of P in their interior.

Exercise 10.3 (The number of heavy vertical segments). Let L be a set of n lines in the plane. A vertical segment has weight k if it intersects k segments of L of its interior. Two such vertical segments are distinct if the subsets of segments they intersect are different. Bound the number of distinct vertical segments of weight at most k.

10.6. From previous lectures

Lemma 10.6.1. *Let $X_1, \ldots, X_n$ be n independent Bernoulli trials, where $\mathbb{P}[X_i = 1] = p_i$ and $\mathbb{P}[X_i = 0] = q_i = 1 - p_i$, for $i = 1, \ldots, n$. Let $X = \sum_{i=1}^{n} X_i$ and $\mu = \mathbb{E}[X] = \sum_{i=1}^{n} p_i$. For any $\delta \geq 2e - 1$, we have $\mathbb{P}[X \geq (1+\delta)\mu] \leq 2^{-\mu(1+\delta)}$,*

Lemma 10.6.2 (Jensen's inequality). *Let $f(x)$ be a convex function. Then, for any random variable X, we have $\mathbb{E}[f(X)] \geq f(\mathbb{E}[X])$.*

References

[AACS98] P. K. Agarwal, B. Aronov, T. M. Chan, and M. Sharir. On levels in arrangements of lines, segments, planes, and triangles. *Discrete Comput. Geom.*, 19: 315–331, 1998.

[ACNS82] M. Ajtai, V. Chvátal, M. Newborn, and E. Szemerédi. Crossing-free subgraphs. *Ann. Discrete Math.*, 12: 9–12, 1982.

[AG86] N. Alon and E. Győri. The number of small semispaces of a finite set of points in the plane. *J. Combin. Theory Ser. A*, 41: 154–157, 1986.

[BMP05] P. Brass, W. Moser, and J. Pach. *Research Problems in Discrete Geometry.* Springer, 2005.

[Cla88] K. L. Clarkson. Applications of random sampling in computational geometry, II. *Proc. 4th Annu. Sympos. Comput. Geom.* (SoCG), 1–11, 1988.

[CS89] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4(5): 387–421, 1989.

[Dey98] T. K. Dey. Improved bounds for planar k-sets and related problems. *Discrete Comput. Geom.*, 19(3): 373–382, 1998.

[Ele02] G. Elekes. Sums versus products in number theory, algebra and Erdős geometry — a survey. *Paul Erdős and his Mathematics II.* Bolyai Math. Soc., 2002, pp. 241–290.

[G E97] G. Elekes. On the number of sums and products. *ACTA Arithmetica*: 365–367, 1997.

[GP84] J. E. Goodman and R. Pollack. On the number of k-subsets of a set of n points in the plane. *J. Combin. Theory Ser. A*, 36: 101–104, 1984.

[Lei84] F. T. Leighton. New lower bound techniques for VLSI. *Math. Syst. Theory*, 17: 47–70, 1984.

[Mat02] J. Matoušek. *Lectures on Discrete Geometry*. Vol. 212. Grad. Text in Math. Springer, 2002.

[PS04] J. Pach and M. Sharir. Geometric incidences. *Towards a Theory of Geometric Graphs.* Ed. by J. Pach. Vol. 342. Contemporary Mathematics. Amer. Math. Soc., 2004, pp. 185–223.

[Sha03] M. Sharir. The Clarkson-Shor technique revisited and extended. *Comb., Prob. & Comput.*, 12(2): 191–201, 2003.

[Szé97] L. A. Székely. Crossing numbers and hard Erdős problems in discrete geometry. *Combinatorics, Probability and Computing*, 6: 353–358, 1997.

[Tót01] G. Tóth. Point sets with many k-sets. *Discrete Comput. Geom.*, 26(2): 187–194, 2001.

[Wel86] E. Welzl. More on k-sets of finite sets in the plane. *Discrete Comput. Geom.*, 1: 95–100, 1986.