2 Linearity of expectation

2.1 Setup and basic examples

Often, a random variable X can be written as

$$X = c_1X_1 + c_2X_2 + \cdots + c_nX_n,$$

where c_i are constants and X_i are random variables, not necessarily independent. In these cases, we know that

$$\mathbb{E}[X] = c_1\mathbb{E}[X_1] + \cdots + c_n\mathbb{E}[X_n].$$

However, it is not necessarily true that $\mathbb{E}[XY] = \mathbb{E}[X]\mathbb{E}[Y]$.

Example 2.1

Given a random permutation of $[n]$, how many fixed points do we expect it to have?

Solution. Let A_i be the indicator variable for i being a fixed point: $\sigma(i) = i$. Since i is a fixed point with probability $\frac{1}{n}$, the expected value of A_i is $\frac{1}{n}$, so the expected number of overall fixed points is just $n \cdot \frac{1}{n} = 1$. □

Let's take a look at a basic graph theory problem:

Definition 2.2

A **tournament** is a complete graph with each edge directed (from one endpoint to the other). A **Hamiltonian path** is a directed path through a graph which passes through all vertices.

Theorem 2.3 (Szele, 1943)

For all n, there exists a tournament on n vertices with at least $n!2^{-n+1}$ Hamiltonian paths.

Proof. Start with K_n and randomly orient each edge. Then for each permutation of the edges, the probability that the edges are all directed correctly to form a Hamiltonian cycle in that order is 2^{-n+1} (since there are only two orientations). Thus, by linearity of expectation, the expected number of Hamiltonian paths is $n!2^{-n+1}$, and thus there exists a tournament with at least that many Hamiltonian paths. □

Alon proved in 1990 that the maximum number is asympotically of that magnitude: we can have at most $\frac{n!}{(2-o(1))^n}$ Hamiltonian paths.

Let's now start to look at some more complicated applications.

2.2 Sum-free sets

Definition 2.4

A subset A of an abelian group is **sum-free** if there are no elements $a, b, c \in A$ with $a + b = c$.

An interesting abelian group to consider is the integers:

Theorem 2.5

Every set of n nonzero integers contains a sum-free subset of size at least $\frac{n}{3}$.

Proof. Let A be a set of nonzero integers with $|A| = n$. Pick a real nmber $\theta \in [0, 1]$, and let

$$A_\theta = \left\{ a \in A \mid \{a\theta\} \in \left(\frac{1}{3}, \frac{2}{3} \right) \right\}$$

(in other words, A_θ contains all points with fractional part of $a\theta$ in the middle third). Note that A_θ is always sum-free, since no two elements with fractional part in the middle third can add to a third. Now uniformly pick θ from 0 to 1: since the probability any a is in A_θ is always $\frac{1}{3}$ (since $a\theta$ ranges from 0 to a), the expected number of points in A_θ is $\frac{n}{3}$, and therefore there is some sum-free subset A_θ with size at least $\frac{n}{3}$, as desired. □

The best we can do currently is $\frac{n+2}{3}$, and it's been shown that $\left(\frac{1}{3} + c\right) n$ is not possible asymptotically for any $c > 0$. However, the constant c' in $\frac{1}{3}n + c'$ is still open!

2.3 Cliques

Theorem 2.6 (Ramsey multiplicity)

There exists a 2-coloring of the edges of K_n with a "relatively small number" of t-cliques: there are at most $2^{1-\binom{t}{2}}\binom{n}{t}$ monochromatic copies of K_t.

Proof. Color all the edges randomly. The expected number of monochromatic K_ts is, by linearity of expectation,

$$\binom{n}{t} 2^{1-\binom{t}{2}}$$

since each t vertices we pick has $\binom{t}{2}$ edges and there are only 2 ways to color them to form a monochromatic K_t. Thus, there is a positive probability that the number of monochromatic K_t is at most this number. □

Definition 2.7

Let c_t be the maximum constant such that every 2-edge coloring of K_n has at least $(c_t + o(1))\binom{n}{t}$ monochromatic t-cliques.

In other words, c_t is the best fractional bound on the number of t-cliques, and we've just found that $c_t \leq 2^{1-\binom{t}{2}}$. Can we do better and find a smaller c_t?

It is known that this is tight for $t = 3$: Goodman's theorem implies that we indeed have $c_3 = \frac{1}{4}$. (Proving this is a good exercise in double counting.) We'd initially suspect that equality can also be achieved for $t = 4$, but it was found by Thomason in 1989 that $c_4 < \frac{1}{33} < \frac{1}{2^5}$. Likewise, the bound has been shown to be not tight for all $t > 4$. In fact, the exact value of c_4 is still an open problem.

But can we prove any kind of lower bound for c_t? Specifically, what techniques do we have to proving positive lower bounds? In other words, we're trying to show that there's a lot of t-cliques, and that sounds vaguely like Ramsey's theorem. One thing we could do is find a copy, delete a vertex, and repeat, but this gives a linear number of t-cliques for n^2 edges, which isn't enough for a positive constant. Instead, we'll use the **sampling trick**!

Theorem 2.8

Every 2-coloring of K_n with $n \geq R(t,t)$ contains $\geq \binom{R(t,t)}{t}^{-1} \cdot \binom{n}{t}$ monochromatic K_ts.

Proof. Suppose there are M monochromatic K_ts in our coloring. Let X be any t-clique: then it has a probability of $\frac{M}{\binom{n}{t}}$ of being monochromatic.

But instead, let's pick the same X in a different way. First, pick a random $R(t,t)$ clique, where $R(t,t)$ is the Ramsey number, and then pick a t-vertex subclique of that. (For this trick to work, we need to be able to pick a random $R(t,t)$ clique.) This second procedure has two random steps, but by Ramsey's theorem, there is at least one monochromatic t-clique in this second step! So X is monochromatic with probability at least $\binom{R(t,t)}{t}^{-1}$.

So putting these together,

$$\frac{M}{\binom{n}{t}} \geq \binom{R(t,t)}{t}^{-1}.$$

□

This is likely far from optimal, but at least it gives us a nonzero lower bound on c_t:

Corollary 2.9

For all positive integers t,

$$c_t \geq \binom{R(t,t)}{t}^{-1}.$$

2.4 Independent sets

Let's turn to a new question: what is the maximum number of edges in an n-vertex K_t-free graph? Note that cliques in a graph G are the same as independent sets in $\overline{G}$ (the graph's complement), so this is a very similar idea to what we've been already been discussing.

Theorem 2.10 (Caro-Wei)

Every graph G contains an independent set I of size

$$|I| \geq \sum_{v \in G} \frac{1}{1 + d(v)}.$$

In particular, we should expect large independent sets out of graphs with low degrees, which is convenient for us.

Proof by Alon and Spencer. Consider a random ordering of V, and let I be the set of vertices that appear before all of its neighbors in the graph.

I is an independent set, since no edge can connect two vertices in I (one comes before another). How big is I? By linearity of expectation,

$$\mathbb{E}[|I|] = \sum_{v \in V} \mathbb{P}(v \in I).$$

The probability that a vertex v is in I is $\frac{1}{1+d(v)}$, since there are $d(v) + 1$ total vertices to consider here, v and all of its neighbors, and v must be the one in front. So there's a nonzero probability that an independent set of size at least $\sum_v \frac{1}{1+d(v)}$ exists. □

Now, let's take the complement of Caro-Wei. Independent sets become cliques and vice versa, which yields the following:

Corollary 2.11

Every graph G contains a clique of size

$$S \geq \sum_{v \in G} \frac{1}{(n-1-d(v))+1} = \sum_{v \in G} \frac{1}{n-d(v)}.$$

Note that if we hold the number of degrees fixed, so $\sum d(v) = 2|E|$, the sum is minimized when the $d(v)$s are as close as possible.

So where's the equality case of Caro-Wei (and the corollary after it)? To have maximal independent set size and largest multiplicity, we want something like the following:

Definition 2.12

A **Turán graph** $T_{n,r}$ has n vertices and is an r-partite complete graph, such that each part has either $\left\lfloor \frac{n}{r} \right\rfloor$ or $\left\lfloor \frac{n}{r} \right\rfloor + 1$ vertices.

Note that this graph is K_{r+1}-free, and it turns out this is the extreme example:

Theorem 2.13 (Turán's theorem)

Given a graph G with n vertices that is K_{r+1} free,

$$|E(G)| \leq |E(T_{n,r})| \leq \left(1 - \frac{1}{r}\right) \frac{n^2}{2},$$

where the inequalities are tight if $r|n$.

For simplicity, we'll show a slightly weaker result where we skip the middle part of the inequality.

Proof. Since G is K_{r+1} free, by the complement of Caro-Wei,

$$r \geq \sum_{v \in V} \frac{1}{n-d(v)} \geq \frac{n}{n-\overline{d}}$$

by convexity, where $\overline{d}$ is the average degree of the vertices. Since the average degree is $\frac{2|E|}{n}$, rearranging gives the result. □

We just have to be a bit more careful in the case where r doesn't divide n, but it's not too much more difficult.

2.5 Crossing numbers

The next example may seem a bit less familiar in terms of the techniques it uses. Given a graph G, we can draw it on the plane; it may or may not be planar. A graph is **planar** if we can draw it in a way such that all edges are continuous curves and only intersect at vertices.

Fact 2.14 ("Common folklore knowledge" and Kuratowski's theorem)

K_4 is planar, but K_5 and $K_{3,3}$ are not. It turns out these are the only two minimal examples of nonplanar graphs: any nonplanar graph contains a subgraph that is topologically equivalent to K_5 or $K_{3,3}$.

The idea is that if we see a graph with a lot of edges, it should have a lot of crossings. How many such crossing must K_n or $K_{n,n}$ have? In fact, what's the bound for any G with some large number of edges?

The exact answers to K_n and $K_{n,n}$ are famous open questions, but there are conjectures: they're called Hill's conjecture and the Zarankiewicz conjecture, respectively.

Remark (Historical note). *The problem of drawing the complete bipartite graph with the minimum number of crossings is also called Turán's brick factory problem. During World War II, Turán was forced to work in a brick factory pushing wagons of bricks along rail tracks. The wagons are harder to push when the rail tracks cross. This experience inspired Turán to think about how to design the layout of the tracks in order to minimize the number of crossings.*

The conjecture for $K_{n,n}$ is to either place points antipodal on a sphere and connect geodesics, or put one set on the x-axis and the other on the y-axis. That makes this problem hard: two very different constructions do equally well.

Definition 2.15

The **crossing number** $\text{cr}(G)$ is the minimum number of crossings in a planar drawing of G.

Are there any bounds we can place on this? It seems like we should expect $O(n^4)$ crossings, since any 4 points potentially create a crossing. Is that at least correct up to a constant factor?

We'll start by considering some facts in graph theory:

Proposition 2.16 (Euler's formula)

Given a connected planar graph with V vertices, E edges, and F faces,

$$V - E + F = 2.$$

The next few sentences are easy to get wrong, so we're going to be careful.

Proposition 2.17

Every connected planar graph with at least one cycle (not just a tree) has $3|F| \le 2|E|$.

This is true because every face is surrounded by at least 3 edges, and every edge touches exactly 2 faces.

Plugging this into Euler's formula, we also find that $|E| \le 3|V| - 6$ for all connected planar graphs with at least one cycle. There are some graphs that do not satisfy the conditions above, but that's okay - from similar arguments, we can still deduce that all planar graphs satisfy $|E| \le 3|V|$.

So if there are too many edges, we want to be able to say that there are lots of crossings. Basically, every edge beyond the threshold of $3|V|$ could add a crossing, so if we delete one edge per crossing, we get a planar graph. Thus $|E| - \text{cr}(G) \le 3|V|$, or

$$\text{cr}(G) \ge |E| - 3|V|.$$

But this gives $O(n^2)$ crossings for an n-vertex graph, and we're trying to show that $O(n^4)$ crossings exist. Here's where the probabilistic method comes in: we're going to sample like we did with the Ramsey number to get a better answer.

Theorem 2.18 (Crossing number inequality)

Given a graph G with $|E| \ge 4|V|$,

$$\text{cr}(G) \gtrsim |E|^3/|V|^2.$$

Proof. Let $p \in [0,1]$ be a number that we will decide later, and let G' be obtained from G by randomly picking each vertex with probability p. In other words, randomly delete each vertex (and the edges connected to it) with probability $1-p$.

Our graph G' should satisfy

$$\operatorname{cr}(G') \geq |E'| - 3|V'|,$$

and now take expectations of both sides:

$$\mathbb{E}[\operatorname{cr}(G')] \geq \mathbb{E}[|E'|] - 3\mathbb{E}[|V'|]$$

If we start with a drawing of G, each crossing has 4 vertices that contribute to it. This crossing remains with probability p^4, but note that after we delete some vertices and edges, we can potentially redraw the diagram to have less crossings. So the left hand side has an inequality of the form

$$\mathbb{E}[\operatorname{cr}(G')] \leq p^4 \operatorname{cr}(G).$$

The right hand side is easier:

$$\mathbb{E}[|E'|] = p^2|E|, \mathbb{E}[|V'|] = p|V|.$$

Moving the p^4 to the other side now, we have a new bound:

$$\operatorname{cr}(G) \geq p^{-2}|E| - 3p^{-3}|V|$$

From here, we set p so that we have $4p^{-3}|V| \leq p^{-2}|E|$, but note that this only works if $|E| \geq 4|V|$, since our probability needs to be between 0 and 1. This gives the result that we want: □

Notably, if $|V| = n$ and $|E| \gtrsim n^2$ (is quadratic in n), then $\operatorname{cr}(G) \gtrsim n^4$: the crossing number is quartic in n, as desired!

2.6 Application to incidence geometry

Problem 2.19

Given n points and n lines, what's the maximum number of incidences between them?

Let's formulate this more rigorously:

Definition 2.20

Let $\mathcal{P}$ be a set of points and $\mathcal{L}$ be a set of lines. Define

$$I(\mathcal{P}, \mathcal{L}) = \{(p, \ell) \in \mathcal{P} \times \mathcal{L} : p \in \ell\}$$

to be the set of intersections between a point in $\mathcal{P}$ and a line in $\mathcal{L}$.

We wish to maximize $|I(\mathcal{P}, \mathcal{L})|$.

Example 2.21

Let $\mathcal{P}$ be the lattice grid $[k] \times [2k^2]$, and let $\mathcal{L}$ be the lines with small integer slope: $\mathcal{L} = \{y = mx + b, m \in [k], b \in [k^2]$. Then every line in $\mathcal{L}$ contains k points, so

$$|I(\mathcal{P}, \mathcal{L})| = k^4,$$

which gives $O(n^{4/3})$ incidences.

The natural question to ask is whether this is optimal, and the answer is yes. To prove this, let's start trying to find some upper bounds. Assume temporarily that every line has at least two incidences: clearly, there is a bound

$$I(\mathcal{P}, \mathcal{L}) \le |\mathcal{P}||\mathcal{L}|,$$

which is weak if there are at least 2 points or 2 lines. But let's use the fact that there is at most one line through each pair of points: to do this, we'll double count the number of triples $(p, p', \ell) \in \mathcal{P} \times \mathcal{P} \times \mathcal{L}$ with $p \neq p'$ and $p, p' \in \ell$. On one hand, given two points, we've determined the line, so there are at most $|\mathcal{P}|^2$ such triples. On the other hand, if we count the incidences in terms of lines, the number of triples is

$$\sum_{\ell \in \mathcal{L}} |P \cap \ell|(|P \cap \ell| - 1) \ge \frac{I(\mathcal{P}, \mathcal{L})^2}{|\mathcal{L}|} - I(P, \mathcal{L})$$

where we've done bounding by Cauchy-Schwarz. Putting these together,

$$I(\mathcal{P}, \mathcal{L}) \lesssim |\mathcal{P}||\mathcal{L}|^{1/2} + |\mathcal{L}|.$$

By point-line duality, we can also find an analogous statement if we flip L and P. Either way, for n lines and n points, we're getting $O(n^{3/2})$, which is not as strong as $O(n^{4/3})$.

Remark. *We can make this bound that we found tight in some situations, though: it turns out this is the right number of incidences over a finite field $\mathbb{F}_q^2$ if we take all $\Theta(q^2)$ lines and all q^2 points.*

Back to the Euclidean plane. To make the bound tight, we invoke the topology of Euclidean space and the crossing number theorem. Assume, again, that every line has at least 2 incidences. Draw a graph based on the point-line configuration, where the points are vertices and **consecutive** points on a line form an edge. So each line gets chopped up into some number of segments.

How many edges and vertices are there? The points are vertices, so $|V| = |\mathcal{P}|$. A line with k incidences (and $k \ge 2$) has $k - 1 \ge \frac{k}{2}$ edges, so the number of edges is at least

$$|E| \ge \frac{I(\mathcal{P}, \mathcal{L})}{2}.$$

Two lines can cross at most once, so

$$\text{cr}(G) \le |\mathcal{L}|^2.$$

Provided that the number of incidences is at least 8 times the number of points, we can invoke the crossing number inequality:

$$|\mathcal{L}^2| \ge \text{cr}(G) \gtrsim \frac{|E|^3}{|V|^2} \gtrsim \frac{|I(\mathcal{P}, \mathcal{L})|^3}{|\mathcal{P}|^2}.$$

Rearranging, this gives us

$$I(\mathcal{P}, \mathcal{L}) \lesssim |\mathcal{P}|^{2/3}|\mathcal{L}|^{2/3},$$

but this only works if we have a sufficiently large number of incidences, so we need to add a linear $|\mathcal{P}|$ term. We also need to correct for the fact that we're assuming that there are at least 2 incidences per line, which adds a linear $|\mathcal{L}|$ term:

Theorem 2.22 (Szemerédi-Trotter theorem)
For any set of points and lines,

$$I(\mathcal{P}, \mathcal{L}) \lesssim |\mathcal{P}|^{2/3}|\mathcal{L}|^{2/3} + |\mathcal{P}| + |\mathcal{L}|.$$

This is sharp up to constant factors! As a corollary, n points and n lines always have $O(n^{4/3})$ incidences.

2.7 Derandomization: balancing vectors

We'll start by solving a problem with familiar techniques:

Theorem 2.23
Given $v_1, \cdots, v_n \in \mathbb{R}^n$ unit vectors, there exists $\varepsilon_1, \varepsilon_2, \cdots, \varepsilon_n \in \{-1, 1\}$ such that

$$|\varepsilon_1 v_1 + \cdots + \varepsilon_n v_n| \le \sqrt{n}.$$

This is motivated by considering $v_1, \cdots, v_n$ to be a standard basis: our choices can't get the length of the vector any smaller than $\sqrt{n}$. As a sidenote, we can also show that we can pick the ε_is to make the length at least $\sqrt{n}$.

We want to use linearity of expectation, but we have a small problem: we have an expectation of an absolute value. The easiest way to get around this is to square both sides of our equation!

Proof. Let

$$X = |\varepsilon_1 v_1 + \cdots + \varepsilon_n v_n|^2,$$

and pick each ε_i independently and randomly between $\{-1, 1\}$. X expands out to the sum

$$X = \sum_{i,j=1}^{n} \varepsilon_i \varepsilon_j (v_i \cdot v_j)$$

and now that the absolute values are gone, we can just use linearity of expectation: for $i \ne j$, the expectation is 0, and for $i = j$, we get a contribution of $1 \cdot |v_i|^2 = 1$ from each term. So the expected value of X is n, so with some positive probability $X \le n$ (and also $X \ge n$). □

We can also do this all deterministically: in this case, we don't actually have to use the probabilistic method.

Finding the ε_is algorithmically. We're going to pick our ε_is sequentially and greedily. At each step, we pick the ε_i that minimizes the expected value conditional on the previous choices.

For example, if we pick $\varepsilon_1, \cdots, \varepsilon_{k-1}$, let $w = \varepsilon_1 v_1 + \cdots + \varepsilon_{k-1} v_{k-1}$. Then our conditional probability

$$\mathbb{E}[X \mid \varepsilon_1, \cdots, \varepsilon_k] = \mathbb{E}\left[|w + \varepsilon_k v_k + \varepsilon_n v_n|^2 \mid \varepsilon_1, \cdots, \varepsilon_k\right],$$

and expanding out the square again, this becomes the expected value of

$$|w|^2 + 2\varepsilon_k (w \cdot v_k) + (n - k - 1).$$

To minimize this value, we pick $\varepsilon_k = 1$ if and only if $w \cdot v_k \le 0$. □

Why couldn't we do something like this for the Ramsey number proof, too? The idea is that we can't compute the number of cliques of other subsets easily! (It is "expensive" to do so.) This idea of turning probabilistic proofs into deterministic ones is called **derandomization**.

2.8 Unbalancing lights

Problem 2.24

Consider a grid of $n \times n$ lights, where we only have light switches for each row and column. How can we maximize the number of lightbulbs turned on given some starting configuration?

Represent this as an array of ± 1 numbers. Let $a_{ij} \in \{-1, 1\}$ for all $1 \le i, j \le n$, and let's say that our light switches are labeled $x_1, \cdots, x_n, y_1, \cdots, y_n \in \{-1, 1\}$. Our goal is then to maximize the quantity

$$\sum_{i,j=1}^{n} a_{ij} x_i y_j,$$

since only the parity of how many times we flip each switch matters (not even the order).

Well, there are n^2 variables, so if we do our probabilistic method naively at random, we can guarantee a linear answer in n, since $\sqrt{n^2} = n$. But we can do better than that:

Theorem 2.25

Given fixed $a_{ij} \in \{-1, 1\}$, we can pick $x_1, \cdots, x_n, y_1, \cdots, y_n \in \{-1, 1\}$, such that

$$\sum_{i,j=1}^{n} a_{ij} x_i y_j \ge \left(\sqrt{\frac{2}{\pi}} + o(1) \right) n^{3/2}.$$

Proof. Choose $y_1, \cdots, y_n \in \{-1, 1\}$ randomly: this means that we pick a random way to flip our columns. Now, for each row, we can choose x_i such that the ith row sum is nonnegative (in other words, flip a row if the sum is negative). Each row sum is

$$R_i = \sum_{j=1}^{n} a_{ij} y_j,$$

and our final sum is just $R = \sum_{i=1}^{n} |R_i|$. Here we use linearity of expectation: the expected value of each R_i is the same, and each R_i is a sum of ± 1s. This gives a binomial distribution: we can use the Central Limit Theorem, since our quantity

$$\mathbb{E}\left(\frac{|R_1|}{\sqrt{n}} \right) \to \mathbb{E}|X| = \sqrt{\frac{2}{\pi}}.$$

(Alternatively, we can directly compute

$$\mathbb{E}[|R_1|] = n 2^{1-n} \binom{n-1}{\lfloor \frac{n-1}{2} \rfloor}$$

and use Stirling's formula.) Regardless, each row has expected value $\left(\sqrt{\frac{2}{\pi}} + o(1) \right) \sqrt{n}$, which is what we want. □

2.9 2-colorings of a hypergraph

Theorem 2.26

Let a k-uniform hypergraph have a vertex set V partitioned as

$$V = V_1 \cup \cdots \cup V_k,$$

where $|V_i| = n$ for all i. Suppose the edges of the complete k-uniform hypergraph on V are colored red and blue such that every edge that intersects all of $V_1, \cdots, V_k$ is colored blue. Then there exists a subset of the vertices $S \subset V$ such that

$$|\# \text{ blue edges} - \# \text{ red edges}| \geq c_k n^k$$

for some constant k.

For example, if $k = 2$, we're looking at a 2-coloring of a complete graph where all of the cross-edges between two halves are blue: our goal is to get a large difference in the number of red and blue edges. Similarly, if $k = 3$, we partition $3n$ vertices into three parts and draw triangles. All the triangles that intersect all three parts are blue, but everything else can be red or blue.

Proof. The idea here is to choose S by including each vertex in a given V_i with probability p_i. We'll leave $p_1, p_2, \cdots, p_k$ undetermined for now.

Let's do the proof for $k = 3$ for illustration, but this generalizes to any k. Let a_{xyz} be the difference in the number of blue and red edges in $V_x \times V_y \times V_z$. When we randomly pick our vertices, by linearity of expectation, the expected number of blue minus red edges is

$$n^3 p_1 p_2 p_3 + \sum_{\substack{x \leq y \leq z \\ \text{not all different}}} a_{xyz} p_x p_y p_z.$$

The first term here comes from the forced blue triangles between all V_is. Our goal is to show this absolute value of this expression is (at least) cubic, and then we'll be done by linearity of expectation.

We haven't chosen our p_is yet, and for each specific choice, we might end up with expected values that are pretty close to 0. So there is always a graph that beats a specific set of p_i, but we just want to find p_1, p_2, p_3 that work given a graph. This is now just an analysis problem:

Lemma 2.27

Let P_k denote the set of polynomials of the form $g(p_1, \cdots, p_k)$ with degree at least k and coefficients having absolute value at most 1, where the coefficient of $p_1 p_2 \cdots p_k$ is exactly 1. Then there exists a constant c_k such that for all polynomials in P_k, there exists $p_1, \cdots, p_k \in [0, 1]^k$ such that

$$g(p_1, p_2, \cdots, p_k) \geq c_k.$$

The proof of this is short: let $M(g)$ be the supremum

$$\sup_{p_1, \cdots, p_k \in [0,1]^k} |g(p_1, \cdots, p_k)|$$

By continuity and compactness, this is actually an achieved maximum, and it is always positive, since all polynomials are nonzero. Furthermore, this map $M : P_k \to \mathbb{R}$ is continuous on a compact domain, so it must achieve its minimum, which is nonzero.

This doesn't give a concrete value of c_k, but it tells us that one exists! And now we're done with the linearity of expectation argument, since all $a_{ijk} < n^3$. □

The main take-away here is that we decide probabilities for our random process in the last step, since no probabilities will work for every configuration.

2.10 High-dimensional sphere packings

Problem 2.28

What is the densest possible packing of unit balls in $\mathbb{R}^n$?

This has been solved for $n = 1$ (trivial), $n = 2$ (a rigorous proof wasn't found until the middle of the 20th century), and $n = 3$ (Kepler's conjecture; proved with computer assistance in the 1990s, and a formal computer proof was recently completed).

Recently, there was a breakthrough that found the answer for $n = 8$ and $n = 24$ as well; those answers come from the E_8 and Leech lattices respectively. However, the problem is open in all other dimensions.

The definition of "density" can be thought of pretty intuitively:

Definition 2.29

Let Δ_n be the maximum fraction of space occupied by non-overlapping unit balls in a large box in $\mathbb{R}^n$ as the volume of the box goes to infinity.

We wish to understand bounds on Δ_n. What are examples of good sphere-packings with high density?

Example 2.30

Consider a packing where we pack greedily: we keep throwing balls in wherever there is space. Alternatively, take any **maximal** packing: basically, find one where we can't fit any additional balls in $\mathbb{R}^n$ anymore without overlap.

What can we say about the density of such a maximal sphere packing? Well, double the radii of every ball, and suppose there is a spot not covered. Then we could just put a unit ball centered at that spot which doesn't intersect any of our initial balls, contradicting maximality of our packing. Thus, we must be able to cover all of $\mathbb{R}^n$ with doubled radii, and thus

$$2^n \Delta_n \geq 1, \text{ so } \Delta_n \geq 2^{-n}.$$

For comparison, what's the packing for $\mathbb{Z}^n$? We can put a ball with radius $\frac{1}{2}$ at every lattice point, and the density is just the volume of a ball of radius $\frac{1}{2}$. This is a pretty standard formula: it's

$$V = \frac{2^{-n}\pi^{n/2}}{(n/2)!} < n^{-cn},$$

so the integer lattice does very poorly compared to the "random" lattice. Are there better ways to construct lattices in higher dimensions? Here's the best bound we know at the moment:

Theorem 2.31 (Kabatiansky–Levenshtein, 1978)

The sphere-packing density in $\mathbb{R}^n$ is at most $2^{-(0.599\cdots+o(1))n}$.

Where does the probabilistic method come into our picture? Although we can't prove the above fact, we want to at least get a better bound than 2^{-n}.

Definition 2.32

A **lattice** is the $\mathbb{Z}$-span of a basis in $\mathbb{R}^n$: given $v_1, v_2, \cdots, v_n$, we can write a matrix with basis vectors as columns. A lattice is **unimodular** if the covolume (volume of the fundamental domain) is 1, which means the matrix has determinant ± 1.

Let's consider matrices A such that $\det A = 1$, so $A \in SL_n(\mathbb{R})$. On the other hand, given a lattice, there's different ways to represent it with a basis: we could always pick $(v_1 + v_2, v_2, \cdots, v_n)$ instead of $(v_1, v_2, \cdots, v_n)$. Any such transformation is matrix multiplication of $B \in SL_n(\mathbb{Z})$.

So the whole point is that lattices are matrices in $SL_n(\mathbb{R})/SL_n(\mathbb{Z})$ through row reduction. Our question: is there a way to pick a random lattice here?

Fact 2.33

The space has a finite Haar measure, so there exists a (normalized) probability Haar measure on $SL_n(\mathbb{R})/SL_n(\mathbb{Z})$, which allows us to choose a random point in the space. That random point will be our random lattice.

Theorem 2.34 (Siegel mean value theorem)

If L is a random unimodular lattice in $\mathbb{R}^n$ (chosen as above according to the Haar probability measure), and if S is any measurable subset of $\mathbb{R}^n$, then

$$\mathbb{E}\left(|L \cap (S \setminus \{0\})|\right) = \text{vol}(S).$$

The idea is that the average point density is 1, so the number of nonzero lattice points is the volume. We exclude 0 because it's always in the lattice.

Proof sketch. Observe that the function $S \to \mathbb{E}\left(|L \cap (S \setminus \{0\})|\right)$ is additive, so it is a measure. Because of how we chose our lattice, it is $SL_n(\mathbb{R})$-invariant, so the measure is also $SL_n(\mathbb{R})$ invariant. Therefore, the only measures that work are constant multiples of the Lebesgue measure.

Now imagine we take a very large ball, much larger than the size of our lattice: then the expected value is the volume minus some boundary errors. So $|S \cap L| \sim \text{vol } S$ and the normalizing constant must be 1. □

How do we use this to find dense lattices?

Proposition 2.35

There exist lattices with sphere packing density greater than 2^{-n}.

Proof. Let S be a ball of volume 1 centered at the origin, and pick a random lattice. By the Siegel mean value theorem, the expected number of nonzero lattice points of L that are in S is 1 (think of this as $1 - \varepsilon$). We can show, then, that there must exist L such that L has no nonzero lattice points in S, since there is a positive probability that there is more than 1 lattice point.

So now put $\frac{1}{2}S$ around every point of L; this gives us a packing with density 2^{-n}. But notice that the nonzero lattice points come in pairs $\{x, -x\}$! In other words, we can take S to be a ball of volume 2. Then we can guarantee the expected number of nonzero lattice points is 2, and we can't have exactly 1 lattice point, so we have the same conclusion as before. This yields a sphere packing with density 2^{1-n}, and this improvement is due to Minkowski. □

Can we do better? There's a lot of connections to the geometry of numbers here. There was a long sequence of improvements made, all of the form $\Delta n \ge cn2^{-n}$, over a few decades. c went from $\frac{1}{2}$ to about 2, but then Venkatesh realized that we can gain factors of k if we have additional symmetry in our lattices: number theory gives such lattices with k-fold symmetry!

For example, consider the lattice corresponding to a cyclotomic field: that is, look at the lattice spanned by a kth root of unity ω. This has a k-fold action, which is multiplication by ω. The end result is that a "random lattice" can be extended to a random unimodular lattice in dimensions $n = 2\phi(k)$, with k-fold symmetry, also satisfying the Siegel mean value theorem conditions. So now k-fold symmetry gives density

$$\Delta_n \ge k \cdot 2^{-n},$$

and this turns out to maximize the gain when $k = p_1 p_2 \cdots p_n$, where p_i is the ith prime. Number theoretic calculations give the following result:

Theorem 2.36 (Venkatesh, 2012)

There exists a lattice packing of unit balls of density

$$\Delta_n \ge cn \log\log n \cdot 2^{-n}$$

for infinitely many values of n and some $c > 0$.

These values of n are very sparse, but this is the state-of-the-art bound. Venkatesh also used a different method to show that (for all sufficiently large n)

$$\Delta_n \ge 60000n \cdot 2^{-n}.$$

It's an open problem whether or not we can get sphere packings of exponentially better density than this, though!

18.218 Probabilistic Method in Combinatorics
Spring 2019