The Ham Sandwich Theorem and Friends, Day 1

Teacher: Yuval

Mathcamp 2015

1 Cutting Ham Sandwiches

Suppose we have a ham sandwich:

Figure 1: A ham sandwich

This is a very simple ham sandwich, consisting of one slab of ham and two slices of bread. Our experience at delis and at the UPS dining hall suggests that we can cut such a sandwich evenly in half using a single straight knife cut, meaning that each piece gets half the volume of the ham and half the volume of each slice of bread. We can try to formalize this intuition as follows:

Theorem 1.1. *Let B_1, B_2, and H be three bounded sets in $\mathbb{R}^3$ (bread, bread, and ham, respectively). Then there is a linear polynomial*

$$P(x,y,z) = a_1x + a_2y + a_3z + a_0$$

such that the plane $\{P = 0\}$ bisects each set, meaning that

$$\mathrm{vol}(B_1 \cap \{P > 0\}) = \mathrm{vol}(B_1 \cap \{P < 0\}) = \frac{1}{2}\mathrm{vol}(B_1)$$
$$\mathrm{vol}(B_2 \cap \{P > 0\}) = \mathrm{vol}(B_2 \cap \{P < 0\}) = \frac{1}{2}\mathrm{vol}(B_2)$$
$$\mathrm{vol}(H \cap \{P > 0\}) = \mathrm{vol}(H \cap \{P < 0\}) = \frac{1}{2}\mathrm{vol}(H)$$

Note! What do we mean by the volume of a set? It's actually not that clear, since our sets could be really funky-shaped. For the purposes of this class, all our sets will be nice enough that we can intuitively understand their volume (but if you want more info, see Steve and Alfonso's class next week).

A natural guess for a generalization of this theorem is that we can do this in other dimensions apart from 3. Indeed, the more general theorem works in all dimensions:

Theorem 1.2 (The Ham Sandwich Theorem)**.** *Let $F_1, \ldots, F_n$ be n bounded sets in $\mathbb{R}^n$ (F stands for food). Then there is a linear function*

$$P(x_1, \ldots, x_n) = a_1x_1 + \cdots + a_nx_n + a_0$$

such that the hyperplane $\{P = 0\}$ bisects each F_i, meaning that for all $1 \le i \le n$,

$$\mathrm{vol}(F_i \cap \{P > 0\}) = \mathrm{vol}(F_i \cap \{P < 0\}) = \frac{1}{2}\mathrm{vol}(F_i)$$

How could we go about proving a theorem like this? A good technique when you don't know how to prove something is to try some special cases. For instance, if $n = 1$, then this theorem just tells us that for any bounded set $F \subset \mathbb{R}$, we can find a point that bisects it. This we can prove by the Intermediate Value Theorem, as follows. We define a function $f : \mathbb{R} \to \mathbb{R}$ by letting

$$f(t) = \mathrm{vol}(\{F < t\}) - \mathrm{vol}(\{F > t\})$$

(Note that vol here denotes the 1-dimensional volume, i.e. length). Then f is certainly continuous, and since F is bounded, we see that f takes on a negative value for small enough t and a positive value for large enough t. So by the Intermediate Value Theorem, there is some point t_0 for which $f(t_0) = 0$, which precisely tells us that the point t_0 bisects F.

In order to prove the general Ham Sandwich Theorem, we will need a higher-dimensional generalization of the Intermediate Value Theorem, called the Borsuk-Ulam Theorem. First, we need a definition.

Definition 1.1. The n-sphere S^n is the set of points

$$\{(x_0, \ldots, x_n) \in \mathbb{R}^{n+1} : x_0^2 + \cdots + x_n^2 = 1\}$$

We can now state the Borsuk-Ulam Theorem:

Theorem 1.3 (Borsuk-Ulam). *Let $f : S^n \to \mathbb{R}^n$ be a continuous map on the n-dimensional sphere. Then there exists some $x \in S^n$ for which $f(x) = f(-x)$.*

One corollary of this is that there are two antipodal points on Earth where both the temperature and pressure are exactly equal. For the map

$$\text{point on Earth} \mapsto (\text{its temperature}, \text{its pressure})$$

is a continuous map $S^2 \to \mathbb{R}^2$, so Borsuk-Ulam guarantees that it achieves the same value on at least one pair of antipodal points.

Also note that the Borsuk-Ulam Theorem is indeed a generalization of the Intermediate Value Theorem, since the $n = 1$ case follows directly from the Intermediate Value Theorem. For given a continuous map $f : S^1 \to \mathbb{R}$, we can define a map $g : [0, \pi] \to \mathbb{R}$ by defining

$$g(\theta) = f(\theta) - f(\pi + \theta)$$

and we get that g is continuous, since f was. If $g(0) = 0$, then we are done, since we get that $f(0) = f(\pi)$. If not, then $g(0)$ is either positive or negative. If it's positive, then $g(\pi) = -g(0)$ is negative, and the Intermediate Value Theorem says that $g(\theta) = 0$ for some θ. The exact same argument works if $g(0)$ is negative, so we can indeed prove the $n = 1$ case of Borsuk-Ulam via the Intermediate Value Theorem. Proving the general case (for any n) is much harder, but there's an outline of the proof in the homework.

Now that we have the Borsuk-Ulam Theorem, we can prove the Ham Sandwich Theorem.

Proof of the Ham Sandwich Theorem. Recall that we want to find a map

$$P(x_1, \ldots, x_n) = a_1 x_1 + \cdots + a_n x_n + a_0$$

such that the hyperplane $\{P = 0\}$ bisects each set F_i. We can generate such a map from a point in S^n by identifying the tuple $(b_0, \ldots, b_n)$ with the map $P_{\boldsymbol{b}}$ defined by

$$P_{\boldsymbol{b}}(x_1, \ldots, x_n) = b_1 x_1 + \cdots + b_n x_n + b_0$$

This allows us to define a map $f : S^n \to \mathbb{R}^n$ as follows. The ith coordinate of f is the map

$$f_i(b_0, \ldots, b_n) = \mathrm{vol}(F_i \cap \{P_{\boldsymbol{b}} < 0\}) - \mathrm{vol}(F_i \cap \{P_{\boldsymbol{b}} > 0\})$$

In other words, the ith coordinate of f measures how good the hyperplane $\{P_{\boldsymbol{b}} = 0\}$ is at bisecting the set F_i; in particular, this hyperplane bisects F_i if and only if $f_i(b_0, \dots, b_n) = 0$. Note that each f_i is continuous, since if we only vary the point $(b_0, \dots, b_n)$ by a little amount, we will only wiggle the plane $\{P_{\boldsymbol{b}} = 0\}$ by a little bit, which means that we will only change how well we cut F_i by a little bit. Since every coordinate map in f is continuous, we get that f is continuous. Finally, observe that

$$\begin{aligned} f_i(-b_0, \dots, -b_n) &= \mathrm{vol}(F_i \cap \{P_{-\boldsymbol{b}} < 0\}) - \mathrm{vol}(F_i \cap \{P_{-\boldsymbol{b}} > 0\}) \\ &= \mathrm{vol}(F_i \cap \{P_{\boldsymbol{b}} > 0\}) - \mathrm{vol}(F_i \cap \{P_{\boldsymbol{b}} < 0\}) \\ &= -f_i(b_0, \dots, b_n) \end{aligned}$$

Functions with this property are called *odd*. Since f_i is odd for all i, we get that f is odd as well.

We can apply the Borsuk-Ulam Theorem to f to get that there exists a point $(a_0, \dots, a_n) \in S^n$ such that $f(a_0, \dots, a_n) = f(-a_0, \dots, -a_n)$. However, since f is odd, we also know that $f(-a_0, \dots, -a_n) = -f(a_0, \dots, a_n)$. So $f(a_0, \dots, a_n)$ is its own negative, meaning that $f(a_0, \dots, a_n) = (0, \dots, 0)$. Let

$$P(x_1, \dots, x_n) = a_1 x_1 + \cdots + a_n x_n + a_0$$

Now recall that the ith component of f was 0 if and only if the hyperplane $\{P = 0\}$ bisected the set F_i, so we get that $\{P = 0\}$ bisects each F_i, which is what we wanted. $\square$

Hooray! We can now cut ham sandwiches! There's just one problem—most of the time, we want to eat sandwiches that are more interesting than two slices of bread and one slice of ham. For instance, suppose we add cheese to our sandwich; can we now bisect all four things (both breads, the ham, and the cheese) with one straight knife cut? On the homework, you'll see that the answer is no. However, we can do things like this if we abandon our fixation on straight knives.

2 Cutting More Interesting Sandwiches with More Interesting Knives

The basic idea is that we were restricting ourselves to cutting sandwiches with hyperplanes, which are simply the zero-sets of linear polynomials—but the linearity condition is kind of arbitrary. Given any multivariate polynomial $Q(x_1, \dots, x_n)$, we say that it bisects a bounded set $F \subset \mathbb{R}^n$ if

$$\mathrm{vol}(F \cap \{Q < 0\}) = \mathrm{vol}(F \cap \{Q > 0\}) = \frac{1}{2}\mathrm{vol}(F)$$

One question that we could ask is how many sets we could simultaneously bisect by allowing us to cut them with the zero-set of any polynomial, but it turns out that this is an uninteresting question; if we allow any polynomial, then it turns out that we can simultaneously bisect arbitrarily many sets. However, the question gets more interesting if we limit the degree of the polynomials that we can use:

Theorem 2.1 (The Polynomial Ham Sandwich Theorem). *Let $d \geq 1$, and let $m = \binom{n+d}{n} - 1$. Suppose we have m bounded sets $F_1, \dots, F_m$ in $\mathbb{R}^n$. Then there is an n-variate polynomial Q of degree at most d such that Q bisects all F_i, meaning that for all $1 \leq i \leq m$,*

$$\mathrm{vol}(F_i \cap \{Q < 0\}) = \mathrm{vol}(F_i \cap \{Q > 0\}) = \frac{1}{2}\mathrm{vol}(F)$$

This lets us cut really great sandwiches! For instance, if we set $n = 3$ and $d = 2$, we get that $m = 9$. So if we allow ourselves quadratic knives in $\mathbb{R}^3$ (e.g. paraboloids and hyperboloids), then we can cut sandwiches with up to 9 ingredients (e.g. bread, ham, cheese, tomatoes, lettuce, onion, mayonnaise, mustard, and a second slice of bread). And if we allow cubic knives, we can cut sandwiches with $\binom{6}{3} - 1 = 19$ ingredients! And so on. It's also worth noting that the $d = 1$ case of this theorem just reduces to the standard Ham Sandwich Theorem.

The Ham Sandwich Theorem and Friends, Day 2

Teacher: Yuval

Mathcamp 2015

3 The Polynomial Ham Sandwich Theorem, Again

Recall what we stated, but didn't prove, yesterday:

Theorem 3.1 (The Polynomial Ham Sandwich Theorem)**.** *Let $d \geq 1$, and let $m = \binom{n+d}{n} - 1$. Suppose we have m bounded sets $F_1, \ldots, F_m$ in $\mathbb{R}^n$. Then there is an n-variate polynomial Q of degree at most d such that Q bisects all F_i, meaning that for all $1 \leq i \leq m$,*

$$\mathrm{vol}(F_i \cap \{Q < 0\}) = \mathrm{vol}(F_i \cap \{Q > 0\}) = \frac{1}{2}\mathrm{vol}(F)$$

Proof of the Polynomial Ham Sandwich Theorem. Note that we have an acutal interpretation for the bizarre number $m = \binom{n+d}{n} - 1$; it is the number of non-constant monomials in n variables of degree at most d. This allows us to repeat a version of our original Borsuk-Ulam argument, as follows. We can make polynomials of degree $\leq d$ from points of S^m by identifying the tuple $(b_0, \ldots, b_m) \in S^m$ with the polynomial $Q_{\boldsymbol{b}}$ whose coefficients are precisely the values $b_0, \ldots, b_m$; we can do this since we picked m so that we have exactly as many coefficients as b's. We then define a map $f : S^m \to \mathbb{R}^m$ by defining its ith coordinate to be

$$f_i(b_0, \ldots, b_m) = \mathrm{vol}(F_i \cap \{Q_{\boldsymbol{b}} < 0\}) - \mathrm{vol}(F_i \cap \{Q_{\boldsymbol{b}} > 0\})$$

Now, the rest of the proof carries through exactly as before, except that we need to make sure that f is continuous in order to be able to apply the Borsuk-Ulam Theorem. This is indeed true: if we vary $(b_0, \ldots, b_m)$ by just a little bit, then the polynomial $Q_{\boldsymbol{b}}$ will only change a little bit, which means that the location of its zero-set $\{Q_{\boldsymbol{b}} = 0\}$ will only change a little bit, so the value of each f_i will only change a little bit. Making this argument fully rigorous is actually a bit tricky, and requires somewhat advanced machinery.

So f is continuous, and it's odd for the same reason as it was in the ordinary Ham Sandwich version. So by the Borsuk-Ulam Theorem, there is some $(a_0, \ldots, a_m) \in S^m$ for which $f(a_0, \ldots, a_m) = (0, \ldots, 0)$. This implies that the hypersurface $\{Q_{\boldsymbol{a}} = 0\}$ bisects each F_i. $\square$

Note! Morally, both the ordinary Ham Sandwich Theorem and its Polynomial version say that we can simultaneously bisect as many sets as we have degrees of freedom to work with. It is in fact possible to make a more precise statement of this intuition, but it gets a bit more complicated.

4 Cutting Spices

For much of the rest of the class, we will actually be interested in discrete things, rather than continuous objects like sandwiches. More concretely, we will have a bunch of spices, and we will want to simultaneously bisect each of them. Since we're dealing with discrete objects, we no longer care about the volume of each half (it'll be 0), but we do care about the number of points in each half.

Definition 4.1. Let $T \subset \mathbb{R}^n$ be a finite set of points, and let Q be a polynomial in n variables. We say that $\{Q = 0\}$ bisects T if

$$|T \cap \{Q < 0\}| \leq \frac{|T|}{2} \quad \text{and} \quad |T \cap \{Q > 0\}| \leq \frac{|T|}{2}$$

Note that we allow—and in general will need—some of the points of T to lie on the bisecting hypersurface $\{Q = 0\}$. With this concept, we can state the discrete analogue of the Ham Sandwich Theorem.

Theorem 4.1 (The Spice Sandwich Theorem)**.** *Let $T_1, \dots, T_n \subset \mathbb{R}^n$ be finite sets of points. Then there exists a linear polynomial P for which the hyperplane $\{P = 0\}$ simultaneously bisects T_i for each $1 \le i \le n$.*

Proof. We will again apply the Borsuk-Ulam Theorem in order to prove the Spice Sandwich Theorem. Recall that the way we did this in the Ham Sandwich case is that for every point $(b_0, \dots, b_n) \in S^n$, we constructed a linear polynomial $P_{\boldsymbol{b}} : \mathbb{R}^n \to \mathbb{R}$ by

$$P_{\boldsymbol{b}}(x_1, \dots, x_n) = b_1 x_1 + \cdots + b_n x_n + b_0$$

After that, we constructed a map $f : S^n \to \mathbb{R}^n$ whose ith coordinate is

$$\begin{aligned} f_i(b_0, \dots, b_n) &= \mathrm{vol}(F_i \cap \{P_{\boldsymbol{b}} < 0\}) - \mathrm{vol}(F_i \cap \{P_{\boldsymbol{b}} > 0\}) \\ &= \text{how good } \{P_{\boldsymbol{b}} = 0\} \text{ is at bisecting } F_i \end{aligned}$$

We would like to do exactly the same thing in the discrete case. The naive thing to try is to define

$$f_i(b_0, \dots, b_n) = |T_i \cap \{P_{\boldsymbol{b}} < 0\}| - |T_i \cap \{P_{\boldsymbol{b}} > 0\}|$$

Why doesn't this work? The key problem is that the Borsuk-Ulam Theorem only applies to *continuous* functions, and this is not continuous: as we move $(b_0, \dots, b_n)$ a tiny bit in such a way that $\{P_{\boldsymbol{b}} = 0\}$ crosses a point of S_i, the value of f_i will suddenly jump by ± 2. However, if we do things a bit more cleverly, we can make such an argument work. The basic idea is that we will construct a function that behaves like the f defined above, except that it is actually continuous. We need to do some preliminiary work first.

For a fixed set T and for a tuple $(c_1, \dots, c_n) \in \mathbb{R}^n$ and some $\alpha \in \mathbb{R}$, we denote by $P_{\boldsymbol{c},\alpha}$ the linear polynomial

$$P_{\boldsymbol{c},\alpha}(x_1, \dots, x_n) = c_1 x_1 + \cdots + c_n x_n + \alpha$$

We now define three functions $\mathbb{R}^n \to \mathbb{R}$:

$$\begin{aligned} m_*(c_1, \dots, c_n) &= \min\{\alpha \in \mathbb{R} : \{P_{\boldsymbol{c},\alpha} = 0\} \text{ bisects } T\} \\ m^*(c_1, \dots, c_n) &= \max\{\alpha \in \mathbb{R} : \{P_{\boldsymbol{c},\alpha} = 0\} \text{ bisects } T\} \\ m(c_1, \dots, c_n) &= \frac{m_*(c_1, \dots, c_n) + m^*(c_1, \dots, c_n)}{2} \end{aligned}$$

Intuitiviely, the idea is to slide the plane orthogonal to $(c_1, \dots, c_n)$ across $\mathbb{R}^n$. The moment we pick up at least half the points, that's the value of m_*; the last moment we have at least half on the other side is the value of m^*. Finally, m is just the average of these two, and is called the *midpoint function.* Note that by construction, if $m(c_1, \dots, c_n) = \alpha$, then the plane $\{P_{\boldsymbol{c},\alpha} = 0\}$ bisects T.

One important property to observe is that m is odd, simply since replacing $(c_1, \dots, c_n)$ by $(-c_1, \dots, -c_n)$ will flip the roles of the min and the max and would also turn α to $-\alpha$. We also claim that these three functions are continuous, which is good since the problem with our original attempt was that the function wasn't continuous. To see that m_* is continuous, consider what happens if we vary $(c_1, \dots, c_n)$ a little bit. This will shift each hyperplane $\{P_{\boldsymbol{c},\alpha} = 0\}$ by a little bit; why will this only affect the minimizing α by a little bit? If we don't alter the point where we start bisecting T, then it obviously won't alter it by a lot, since that point is the only one that matters. If, on the other hand, we change the crossing point, then α again can't change by much, since this necessarily implies that these two points must be both close to $\{P_{\boldsymbol{c},\alpha} = 0\}$, so again the minimizing α will only change a little bit. A similar argument shows that m^* is continuous, and the average of two continuous functions is continuous, so the claim is proved.

Now, let $T_1, \dots, T_n \subset \mathbb{R}^n$ be finite sets, and let m_i be the midpoint function of the set T_i, as defined above. We define a function $f : S^n \to \mathbb{R}^n$ whose ith coordinate is

$$f_i(b_0, \dots, b_n) = m_i(b_1, \dots, b_n) - b_0$$

Since we proved that each m_i is continuous, we get that f is continuous, and f is odd since

$$f_i(-b_0, \dots, -b_n) = m_i(-b_1, \dots, -b_n) - (-b_0) = b_0 - m_i(b_1, \dots, b_n) = -f_i(b_0, \dots, b_n)$$

where the middle equality comes from the fact that m is odd, as observed above. Therefore, we can apply the Borsuk-Ulam Theorem to f to get that f vanishes on some $(a_0, \dots, a_n) \in S^n$. This means that for each i,

$$m_i(a_1, \dots, a_n) = a_0$$

As we observed above, this tells us that the plane $\{P_{\boldsymbol{a}} = 0\}$ bisects each T_i, where

$$P_{\boldsymbol{a}}(x_1, \dots, x_n) = a_1 x_1 + \cdots + a_n x_n + a_0$$

Thus, we have found a simultaneously bisecting hyperplane for each of our spice-sets. $\square$

The Ham Sandwich Theorem and Friends, Day 3

Teacher: Yuval

Mathcamp 2015

5 Cutting More Interesting Spices with More Interesting Knives

Having figured out how to expand the Ham Sandwich Theorem to the realm of discrete sets, it's easy to guess what the next theorem will be:

Theorem 5.1 (The Polynomial Spice Sandwich Theorem). *Let $d \geq 1$ be an integer, and let $m = \binom{n+d}{d} - 1$. Let $T_1, \ldots, T_m \subset \mathbb{R}^n$ be finite sets. Then there exists an n-variate polynomial Q of degree $\leq d$ for which the hypersurface $\{Q = 0\}$ bisects each T_i.*

Note that as with the Ham Sandwich Theorem, this polynomial version specializes to our original theorem when $d = 1$.

Proof. We could prove this in a way analogous to how we proved the ordinary Spice Sandwich Theorem (again by picking a clever proxy for the non-continuous function that we want), but luckily we don't have to go through that messy argument again. Instead, we'll conclude the Polynomial Spice Sandwich Theorem from the ordinary Spice Sandwich Theorem using a powerful trick called the *Veronese mapping.*

Recall that m is the number of nonconstant monomials of degree $\leq d$ with coefficient 1 on n variables. So we label the basis of $\mathbb{R}^m$ by these monomials. For instance, if $n = 2$ and $d = 2$, then $m = 5$, and its basis is $\{x, y, x^2, xy, y^2\}$. We then define the Veronese mapping Φ as the function $\mathbb{R}^n \to \mathbb{R}^m$ that sends

$$(x_1, \ldots, x_n) \mapsto (p(x_1, \ldots, x_n))_{p \text{ runs over all monomials}}$$

For instance, in the $n = 2, d = 2$ example above, the Veronese mapping would be

$$\Phi(x, y) = (x, y, x^2, xy, y^2)$$

and thus e.g.

$$\Phi(2,3) = (2,3,4,6,9)$$

Why is the Veronese mapping great? It lets us convert a (hard) non-linear problem into an (easier) linear one. In our case, let $T_1, \ldots, T_m \subset \mathbb{R}^n$ be our m sets. Consider the sets $\Phi(T_1), \ldots, \Phi(T_m) \in \mathbb{R}^m$. Since these are spice-sets in $\mathbb{R}^m$, we can apply the ordinary Spice Sandwich Theorem to conclude that there is some hyperplane $\{P = 0\}$ that bisects each $\Phi(T_i)$. Then P is a linear polynomial on m variables. But recall that the coordinates in $\mathbb{R}^m$ are indexed by monomials in n variables, so we can actually view P as a polynomial on n variables; to avoid confusion, we will call this n-variate polynomial Q. Next, note that $\deg Q \leq d$; this is simply because all of the monomials that are the coordinates in $\mathbb{R}^m$ have degree $\leq d$.

Finally, we claim that $\{Q = 0\}$ bisects each $T_i \subset \mathbb{R}^n$. To see this, let $s \in T_i$ be a spice point, and let $\Phi(s)$ be its image in $\mathbb{R}^m$ under the Veronese mapping. Note that by the way we defined Q, we get that $Q(s) = P(\Phi(s))$. Therefore, if s was on the $\{P < 0\}$ side of T_i in $\mathbb{R}^m$, then it will be on the $\{Q < 0\}$ side of T_i in $\mathbb{R}^n$, and the analogous results are true for $\{P = 0\}$ and $\{P > 0\}$. Thus, since $\{P = 0\}$ bisected each $\Phi(T_i)$, we get that $\{Q = 0\}$ bisects each T_i. □

6 Dividing a Condiment between More People

We've now gotten really *really* good at cutting things. However, there remains one natural direction to extend our cutting acumen: cutting into more than two parts. It is possible to state results on this topic in the continuous (sandwich) case, but we will actually only care about the discrete (spice) case. This is primarily because our main goal in these last two days is to understand the Guth-Katz partitioning polynomial, a really new tool (c. 2010) that turns out to be stupidly useful for proving hard theorems, and this tool only makes sense for discrete sets.

The main object that we will be interested in is an r-partitioning polynomial. Though it can be defined in all dimensions, we will be primarily interested in the 2-dimensional case, so we will only define it in $\mathbb{R}^2$.

Definition 6.1. Let $C \subset \mathbb{R}^2$ be a set of n points in the plane (C for condiment), and let $1 < r \le n$ be a parameter. A polynomial $Q(x, y)$ is called an *r-partitioning polynomial* if every connected component of $\mathbb{R}^2 \setminus \{Q = 0\}$ contains at most n/r points of C.

This captures the idea of us dividing the condiment C among r people. Note that we might have more than r connected components and that some of the points of C might lie on the curve $\{Q = 0\}$, but it's still a pretty good notion of dividing among r people. At a high level, this is the theorem that we will prove:

"Theorem" 6.1. *Every set of points $C \subset \mathbb{R}^2$ has an r-partitioning polynomial Q. The degree of Q depends only on r (**not** on n), and this degree grows relatively slowly relative to r.*

In order to state a real version of this theorem, we need an important piece of notation.

Definition 6.2. Let $g, h : \mathbb{R} \to \mathbb{R}$ be two functions. We write $g(x) \le O(h(x))$, and say that g is big-o of h, if there exists some fixed constant M and some $x_0 \in \mathbb{R}$ for which $|g(x)| \le M|h(x)|$ for all $x \ge x_0$. Intuitively, what this says is that as long as we're only interested in the asymptotic growth of the two functions, then g grows at most as quickly as h, disregarding constants.

With this notation, we can state the actual theorem that we will prove:

Theorem 6.1 (Guth-Katz)**.** *Let $C \subset \mathbb{R}^2$ be a set of n points, and let $1 < r \le n$ be a parameter. Then there exists an r-partitioning polynomial Q of the set C with* $\deg Q \le O(\sqrt{r})$.

Proof. We will inductively construct a sequence $\mathscr{C}_0, \mathscr{C}_1, \mathscr{C}_2, \ldots$, where each $\mathscr{C}_j$ is a collection of disjoint subsets of C (i.e. each element of $\mathscr{C}_j$ is a subset of the points of C). We will also ensure that $|\mathscr{C}_j| \le 2^j$. The base case is $\mathscr{C}_0 = \{C\}$. Once we have a set $\mathscr{C}_j$ with $|\mathscr{C}_j| \le 2^j$, then we can apply the Polynomial Spice Sandwich Theorem with $m = 2^j$ to find a polynomial Q_j that bisects each of the elements of $\mathscr{C}_j$ (each of which is a spice-set). What is $\deg Q_j$? By the Polynomial Spice Sandwich Theorem, we can have $\deg Q_j \le d$, where $\binom{2+d}{2} - 1 \ge 2^j$. We can assume that d is the minimal integer that achieves this, meaning that

$$\begin{aligned} 2^j &\ge \binom{2+(d-1)}{2} \\ &= \frac{d^2+d}{2} \\ &> \frac{d^2}{2} \qquad\qquad \text{(for } d \ge 1\text{)} \end{aligned}$$

Therefore, $\deg Q_j \le \sqrt{2 \cdot 2^j}$. Now, we define

$$\mathscr{C}_{j+1} = \bigcup_{D \in \mathscr{C}_j} \{D \cap \{Q_j > 0\}, D \cap \{Q_j < 0\}\}$$

In other words, $\mathscr{C}_{j+1}$ consists of all the halves that we get by bisecting each of the sets in $\mathscr{C}_j$. Since $|\mathscr{C}_j| \le 2^j$ and since we're halving each set in $\mathscr{C}_j$, we get that $|\mathscr{C}_{j+1}| \le 2^{j+1}$. Also note that all of the sets in $\mathscr{C}_{j+1}$ are disjoint, since we assumed inductively that that was true for $\mathscr{C}_j$, and we have halved each set in $\mathscr{C}_j$.

Since $\mathscr{C}_j$ consists of repeated bisections of C, and since $|C| = n$, we get that each element of $\mathscr{C}_j$ has size at most $n/2^j$. Set $t = \lceil \log_2 r \rceil$, so that each element of $\mathscr{C}_t$ has size at most n/r. Define

$$Q = Q_1 \cdot Q_2 \cdots Q_t$$

Then we claim that Q is the r-partitioning polynomial that we want. First, we claim that each connected component of $\mathbb{R}^2 \setminus \{Q = 0\}$ contains at most n/r points of C. To see this, we claim that every component of $\mathbb{R}^2 \setminus \{Q = 0\}$ contains points from at most one set in $\mathscr{C}_t$, which suffices since each such set has size $\leq n/r$. To see this, let c_1 and c_2 be in distinct sets of $\mathscr{C}_t$. Then at some point, we must have bisected our set in such a way that we separate c_1 and c_2. But that means that for at least one $1 \leq j \leq t$, the curve $\{Q_j = 0\}$ separates c_1 and c_2. Because Q_j is a factor of Q, the same will be true of Q. So any path connecting c_1 and c_2 must pass through a point where $Q = 0$, meaning that c_1 and c_2 are in different connected components of $\mathbb{R}^2 \setminus \{Q = 0\}$. So each connected component contains points from at most one set of $\mathscr{C}_t$, and each such set has size $\leq n/r$, so Q is an r-partitioning polynomial.

Now, we have to show that $\deg Q \leq O(\sqrt{r})$. To see this, we observe that

$$\begin{aligned}
\deg Q &= \deg(Q_1) + \cdots + \deg(Q_t) \\
&\leq \sqrt{2} \sum_{j=1}^{t} 2^{j/2} \\
&= \sqrt{2}\left(\sqrt{2}\frac{2^{t/2} - 1}{\sqrt{2} - 1}\right) \\
&< \frac{2}{\sqrt{2} - 1} 2^{t/2} \\
&\leq \frac{2}{\sqrt{2} - 1} 2^{(\log_2 r + 1)/2} \\
&= \frac{2\sqrt{2}}{\sqrt{2} - 1} \sqrt{r} = O(\sqrt{r})
\end{aligned}$$

Thus, Q is an r-partitioning polynomial for C with degree $\deg Q \leq O(\sqrt{r})$, as desired. $\square$

The Ham Sandwich Theorem and Friends, Day 4

Teacher: Yuval

Mathcamp 2015

7 Sorting a Condiment into Lines

As an application of the Guth-Katz partitioning polynomial, we will prove the Szemerédi-Trotter Theorem, arguably the most important theorem in the field of combinatorial geometry. The setup is simple: we have a finite set of points $C \subset \mathbb{R}^2$, and a finite collection L of lines in the plane, with $|C| = n$ and $|L| = m$. We are interested in the quantity

$$I(C, L) = |\{(c, \ell) \in C \times L : c \text{ lies on } \ell\}|$$

$I(C, L)$ is called the *number of incidences* of C with L. Note that we count the number of pairs of a point and a line such that the point lies on the line; thus, if $C = \{c\}$ and every line of L passes through c, then $I(C, L) = m$.

We are interested in approximating the size of $I(C, L)$. Since we can always attain $I(C, L) = 0$ by making sure that none of our lines hits any of our points, there's no real point to trying to find lower bounds for $I(C, L)$. We can find the following upper bound:

Proposition 7.1. $I(C, L) \le m + n^2$.

Proof. We divide the set of lines L into two subsets: L' consists of all lines that are incident to at most one point of C, and L'' contains all the others. By definition, we have that

$$I(C, L') \le |L'| \le m$$

since each line in L' contributes at most one incidence. On the other hand, we claim that each $c \in C$ can have at most $n - 1$ incidences with L''; this is because any line in L'' that is incident with c must also be incident with at least one other point of C. Since all these lines are distinct, all these extra points are distinct, and there are at most $n - 1$ of them. So

$$I(C, L'') \le \sum_{c \in C} (n - 1) \le n^2$$

Adding these up gives our desired bound. $\square$

This is actually a pretty terrible bound—on the homework, you will prove that

$$I(C, L) \le O(mn^{1/2} + n) \qquad \text{and} \qquad I(C, L) \le O(m^{1/2}n + m)$$

which is in general a much stronger bound. For instance, if $m = n$, then the bound in the proposition is $O(n^2)$, where as this bound is $O(n^{3/2})$. If some forms of geometry, it turns out that this $O(n^{3/2})$ is the best we can do (see the homework for details), but the Euclidean plane has some very nice properties that allow us to do better:

Theorem 7.1 (Szemerédi-Trotter)**.** *Let C and L be as above, with $|C| = n$ and $|L| = m$. Then*

$$I(C, L) \le O(m^{2/3}n^{2/3} + m + n)$$

In particular, if $m = n$, then

$$I(C, L) \le O(n^{4/3})$$

We will prove this by applying the Guth-Katz partitioning polynomial. But before doing so, it's important to observe that this upper bound is actually tight, in the sense that we can construct an example that achieves this number of incidences.

Example 7.1. Fix some integer K. Let C be the set of points

$$C = \{(x, y) \in \mathbb{R}^2 : x, y \in \mathbb{Z}, 1 \le x \le K, 1 \le y \le 2K^2\}$$

Let L consist of all lines of the form $y = ax + b$, where $a, b \in \mathbb{Z}$ and $1 \le a \le K, 1 \le b \le K^2$. Then $n = |C| = 2K^3$ and $m = |L| = K^3$. Finally, observe that each $\ell \in L$ has exactly K incidences with C, since for any $1 \le x \le K$, we have that $y = ax + b \le 2K^2$, meaning that $(x, y) \in C$. So

$$I(C, L) = \sum_{\ell \in L} K = K^4$$

and $O(m^{2/3}n^{2/3}) = O((K^3)^{2/3}(K^3)^{2/3}) = O(K^4)$. So this example has (asymptotically) as many incidences as is possible, according to the Szemerédi-Trotter Theorem.

7.1 Proving Szemerédi-Trotter

Before proving the Szemerédi-Trotter Theorem, we need two basic results about polynomials in the plane.

Lemma 7.1. *If ℓ is a line in the plane and $Q(x, y)$ is a polynomial with* $\deg Q = d$*, then either $\ell \subseteq \{Q = 0\}$ or $|\ell \cap \{Q = 0\}| \le d$.*

This is a special case of a much stronger theorem called Bézout's Theorem, but we only need this version.

Proof. We can write ℓ as a parametrized curve

$$\ell = \{(a_1 t + b_1, a_2 t + b_2) : t \in \mathbb{R}\}$$

In that case, the points of $\ell \cap \{Q = 0\}$ are roots of the polynomial

$$q(t) = Q(a_1 t + b_1, a_2 t + b_2)$$

Then $q(t)$ is a polynomial of one variable, and its degree is d. So either it's identically zero (in which case $\ell \subseteq \{Q = 0\}$), or it has $\le d$ roots, meaning that $|\ell \cap \{Q = 0\}| \le d$. □

Lemma 7.2. *If* $\deg Q = d$*, then $\{Q = 0\}$ contains at most d distinct lines.*

Proof. Suppose we have k distinct lines $\ell_1, \dots, \ell_k \subseteq \{Q = 0\}$. We need to prove that $k \le d$. Since $\{Q = 0\}$ is not all of $\mathbb{R}^2$, we can pick some point $c \notin \{Q = 0\}$. Pick some line ℓ passing through c such that ℓ is not parallel to any ℓ_i and such that ℓ doesn't go through $\ell_i \cap \ell_j$; we can do this since there are infinitely many possible slopes for lines passing through c and only finitely many slopes that these restrictions make us avoid. Then ℓ intersects each ℓ_i exactly once, and all of these intersections are distinct. Then

$$k = \left| \bigcup_{i=1}^{k} \ell \cap \ell_i \right| \le |\ell \cap \{Q = 0\}| \le d$$

where the last inequality comes from the previous lemma. □

Note that we could alternately prove both of these theorems by observing that $\{Q = 0\}$ contains a line ℓ if and only if the linear polynomial defining ℓ divides Q.

With these two lemmas, we are ready to prove the Szemerédi-Trotter Theorem. We will, of course, apply the Guth-Katz partitioning polynomial theorem. The basic idea in any proof that applies this theorem is the following: we can cut our condiment-set into parts via a polynomial, and we end up with two types of points. The points that lie in the complement of our partitioning curve behave like "generic" points, in the sense that they are roughly evenly distributed between the connected components. All the points that lie on the partitioning curve are highly "non-generic", but they are also highly structured, since they all lie on a polynomial curve of low degree. So whenever we want to prove something by using Guth-Katz, we separate into these two cases, and often have to use very different techniques to attack each case. This is also the reason we were never really bothered by the fact that some of our theorems allowed points to land on the bisector, since we can just apply "non-generic" algebraic techniques to these points.

Proof of Szemerédi-Trotter. We will assume that $m = n$, since it is slightly simpler to write down and has the exact same ideas as the proof of the general version.

Let $r = n^{2/3}$ and let $Q(x, y)$ be an r-partitioning polynomial for C. By the Guth-Katz theorem, we have that $d := \deg Q = O(\sqrt{r}) = O(n^{1/3})$. Let $A_1, \dots, A_s$ be the connected components of $\mathbb{R}^2 \setminus \{Q = 0\}$, and let $C_i = C \cap A_i$ for all $1 \le i \le s$. Finally, let $C_0 = C \cap \{Q = 0\}$ denote our "non-generic" points. Since Q is r-partitioning, we have that $|C_i| \le n/r = n^{1/3}$ for $1 \le i \le s$. Also let $L_0 \subseteq L$ consist of all the lines in L that lie in $\{Q = 0\}$; by Lemma 7.2, $|L_0| \le d$.

Recall that we wish to estimate $I(C, L)$. We break this up as

$$I(C, L) = I(C_0, L_0) + I(C_0, L \setminus L_0) + \sum_{i=1}^{s} I(C_i, L)$$

We will estimate each of these terms separately. First,

$$I(C_0, L_0) \le |C_0| \cdot |L_0| \le |C_0| d \le nd = O(n^{4/3})$$

For the next term, observe that each line of $L \setminus L_0$ is not contained in $\{Q = 0\}$, so by Lemma 7.1, it intersects $\{Q = 0\}$ in at most d points. Since $P_0 \subset \{Q = 0\}$, we get that each such line is incident with P_0 in at most d points, so

$$I(C_0, L \setminus L_0) \le |L \setminus L_0| \cdot d \le nd = O(n^{4/3})$$

Finally, we need to bound $\sum_{i=1}^{s} I(C_i, L)$. For this, let $L_i \subseteq L$ consist of those lines in L that are incident with at least one point of C_i. In general, these L_i's will not be disjoint; however, we claim that each line $\ell \in L$ will be in at most $d + 1$ of the sets L_i. For by Lemma 7.1, ℓ can intersect $\{Q = 0\}$ at most d times, so it can meet at most $d + 1$ connected components of $\mathbb{R}^2 \setminus \{Q = 0\}$.

By definition of L_i, we have that $I(C_i, L) = I(C_i, L_i)$. By Proposition 7.1,

$$I(C_i, L_i) \le |L_i| + |C_i|^2$$

and thus,

$$\sum_{i=1}^{s} I(C_i, L_i) \le \sum_{i=1}^{s} |L_i| + \sum_{i=1}^{s} |C_i|^2$$

Since each line lies in at most $d + 1$ distinct L_i's, we get that

$$\sum_{i=1}^{s} |L_i| \le (d + 1)n = O(n^{4/3})$$

Finally,

$$\sum_{i=1}^{s} |C_i|^2 \le \left(\max_i |C_i|\right) \cdot \sum_{i=1}^{s} |C_i| \le \frac{n}{r} \cdot n = n^{4/3}$$

Since each of our partial bounds is $O(n^{4/3})$, we get that

$$I(C, L) = O(n^{4/3})$$

□

8 I Am Out of Food Metaphors

Let $C \subset \mathbb{R}^2$ be an n-point set in the plane. How many distinct distances are there between pairs of points in C? More precisely, define

$$\Delta(C) = |\{\text{dist}(c_1, c_2) : c_1, c_2 \in C\}|$$

Note that the maximal number of possible distances is $\binom{n}{2}$, and we can actually achieve that many distances by placing points generically in the plane. So a far more interesting problem is to find a lower bound for $\Delta(C)$, i.e. a number of distinct distances which *every* set of n points must have. In 1946, Erdős conjectured the following:

Conjecture: $\Delta(C) \geq Mn/\sqrt{\log n}$ for some constant M.

If this conjecture were proved, then the bound would be tight, since it turns out that a $\sqrt{n} \times \sqrt{n}$ grid achieves this number of distinct distances. This conjecture has become known as the Erdős distinct distances problem, and Guth and Katz originally studied partitioning polynomials in order to make progress on this theorem. In fact, what they proved is

Theorem 8.1 (Guth-Katz)**.** *For some constant M,*

$$\Delta(C) \geq M \frac{n}{\log n}$$

This is slightly smaller than Erdős's conjectured value, but is far better than any previously known result.

Guth and Katz managed to prove this using the polynomial partitioning theorem. They converted the Erdős distinct distances problem (which is a question about points in the plane) into a question about lines in $\mathbb{R}^3$. More specifically, they managed to prove that if Theorem 8.1 were false, then we could produce a set L of lines in $\mathbb{R}^3$ that had many incidences, i.e. many pairs $\ell_1, \ell_2 \in L$ with $\ell_1 \cap \ell_2 \neq \emptyset$. Finally, they managed to use the $\mathbb{R}^3$ version of the partitioning polynomial theorem to prove a version of Szemerédi-Trotter in $\mathbb{R}^3$; in particular, this theorem proves that such a collection L of lines cannot exist, so we must have that $\Delta(C) \geq Mn/\log n$.