POLYNOMIAL CELL DECOMPOSITIONS

1. Polynomial cell decompositions

A union of d planes is a special case of an algebraic surface of degree d. The main idea in this chapter is to cut space into pieces with a degree d algebraic surface. Allowing an arbitrary degree d surface instead of just d planes greatly increases our flexibility. (In $\mathbb{R}^3$, when we pick d planes, we have $3d$ parameters to play with, but when we pick a degree d surface we have $\sim (1/6)d^3$ parameters to play with!) With all this extra flexibility, we can do a much better job of decomposing space into evenly matched cells. On the other hand, if Z is a degree d surface, then a line either lies in Z or intersects Z in $\leq d$ points. Therefore, each line intersects $\leq d+1$ components of the complement of Z – exactly the same bound as if Z was a union of d planes.

Theorem 1.1. *If S is any finite subset of $\mathbb{R}^n$ and d is any degree, then there is a non-zero degree d polynomial P so that each component of $\mathbb{R}^n \setminus Z(P)$ contains $\leq C(n)|S|d^{-n}$ points of S.*

We will prove this theorem today. The proof is a cousin of finding a degree d polynomial that vanishes at $\sim d^n$ prescribed points, but it uses topology instead of linear algebra.

2. Ham sandwich theorems

We will build our polynomial cell decomposition using a tool from topology, the ham sandwich theorem. In this section, we develop the tools that we will use.

Theorem 2.1. *(Ham sandwich theorem) If $U_1, ..., U_n$ are finite volume open sets in $\mathbb{R}^n$, then there is a hyperplane that bisects each set U_i.*

This theorem was first proven by Banach in the late 30's (in the case $n = 3$). Then Stone and Tukey generalized the argument to higher dimensions, and they gave a much more general theorem (see below). We can get a heuristic sense of the situation by counting parameters. The set of hyperplanes in $\mathbb{R}^n$ is given by n parameters. Heuristically, we might expect that the subset of hyperplanes that bisect U_1 is given by $n-1$ parameters; that the subset of hyperplanes that bisect U_1 and U_2 is given by $n-2$ parameters etc. Another special case happens when each U_i is a round ball. In that case, the solution is a plane that goes through the center of each ball. If the centers are in general position, there will be exactly one solution.

The planes are exactly the zero sets of degree 1 polynomials (polynomials of the form $a_1 x_1 + ... + a_n x_n + b$). We can generalize this setup by allowing other functions, such as higher degree polynomials. Suppose that V is a vector space of functions from $\mathbb{R}^n$ to $\mathbb{R}$. Multiplication by a scalar doesn't change the zero set of a function f, so might say heuristically that the family of zero sets is given by $dimV - 1$ parameters. For example, if V is the polynomials of degree ≤ 1, then $dimV = n + 1$, and the dimension of the set of hyperplanes is n. Since we have $dimV - 1$ parameters to play with, we might hope to bisect $dimV - 1$ sets $U_i \subset \mathbb{R}^n$. This heuristic turns out to be correct under very mild conditions on the space V.

To state our theorem, we make a little basic notation. For any function $f : \mathbb{R}^n \to \mathbb{R}$, we let $Z(f) := \{x \in \mathbb{R}^n | f(x) = 0\}$. We say that f bisects a finite volume open set U if

$$Vol_n\{x \in U | f(x) > 0\} = Vol_n\{x \in U | f(x) < 0\} = (1/2) Vol_n U.$$

Theorem 2.2. *(General ham sandwich theorem, Stone and Tukey, 1942) Let V be a vector space of continuous functions on $\mathbb{R}^n$. Let $U_1, ..., U_N \subset \mathbb{R}^n$ be finite volume open sets with $N < dimV$. For any function $f \in V \setminus \{0\}$, suppose that $Z(f)$ has Lebesgue measure 0. Then there exists a function $f \in V \setminus \{0\}$ which bisects each set U_i.*

The ham sandwich theorem is one corollary, given by taking V to be the degree 1 polynomials. If we consider the space of polynomials with degree $\le d$, we get the following corollary.

Corollary 2.3. *(Polynomial ham sandwich theorem)*

Proof. We let $V(d)$ be the space of polynomials of degree $\le d$. We saw in the very beginning of the course that $dimV(d) = \binom{d+n}{n}$. It's also easy to check that for a non-zero polynomial P, $Z(P)$ has measure 0. We leave this as an exercise. $\square$

The polynomial ham sandwich theorem is analogous to the more basic polynomial existence lemma which we have been using throughout the course. We rewrite the lemma here to make the analogy clear.

Lemma 2.4. *(Polynomial existence lemma) If $p_1, ..., p_N \in \mathbb{R}^n$ are points and $N < \binom{d+n}{n}$, then there is a non-zero polynomial of degree $\le d$ that vanishes at each x_i.*

The polynomial existence lemma is analogous to the polynomial ham sandwich theorem. The first is based on linear algebra, and the second is based on topology. The polynomial existence lemma was a basic step in all of our arguments. Using the polynomial ham sandwich theorem instead gives a new direction to the polynomial method.

3. On the proof of the ham sandwich theorem

The heuristic argument above using parameter counting is definitely not a proof. The proof of the ham sandwich theorem is based on the Borsuk-Ulam theorem.

Theorem 3.1. *(Borsuk-Ulam) Suppose that $\phi : S^N \to \mathbb{R}^N$ is a continuous map that obeys the antipodal condition $\phi(-x) = -\phi(x)$ for all $x \in S^N$. Then the image of ϕ contains 0.*

For a proof of the Borsuk-Ulam theorem, the reader can look at Matousek's book *Using the Borsuk-Ulam theorem* or in the book *Differential Topology* by Guillemin and Pollack, Chapter 2.6. The book *Using the Borsuk-Ulam theorem* discusses some surprising applications of Borsuk-Ulam to combinatorics.

Proof of the general ham sandwich theorem. For each i from 1 to N, we define $\phi_i : V \setminus \{0\} \to \mathbb{R}$ by

$$\phi_i(F) := Vol\left(\{x \in U_i | F(x) > 0\}\right) - Vol\left(\{x \in U_i | F(x) < 0\}\right).$$

So $\phi_i(F) = 0$ if and only if f bisects U_i. Also, ϕ_i is antipodal, $\phi_i(-F) = -\phi_i(F)$.

We will check below that ϕ_i is a continuous function from $V \setminus \{0\}$ to $\mathbb{R}$. We assemble the ϕ_i into one function $\phi : V \setminus \{0\} \to \mathbb{R}^N$.

We know that $dim V > N$, and without loss of generality we can assume that $dim V = N + 1$. Now we choose an isomorphism of V with $\mathbb{R}^{N+1}$, and we think of S^N as a subset of V. The map $\phi : S^N \to \mathbb{R}^N$ is antipodal and continuous. By the Borsuk-Ulam theorem, there is a function $f \in S^N \subset V \setminus \{0\}$ so that $\phi(f) = 0$. This function f bisects each U_i. □

It only remains to check the technical point that ϕ_i is continuous. This follows from the next lemma. It's basically an exercise in measure theory.

Continuity Lemma. *Let V be a finite-dimensional vector space of continuous functions on $\mathbb{R}^n$. Suppose that for each $f \in V \setminus \{0\}$, the set $Z(f)$ has measure 0.*

If U is a finite volume open set, then the measure of the set $\{x \in U | f(x) > 0\}$ depends continuously on $f \in V \setminus \{0\}$.

Proof. Suppose that f is a function in $V \setminus \{0\}$ and $f_n \in V \setminus \{0\}$ with $f_n \to F$ in V. A priori, f_n converges to f in the topology of V. But then it follows that $f_n \to f$ pointwise. Pick any $\epsilon > 0$. We can find a subset $E \subset U$ so that $f_n \to f$ uniformly pointwise on $U \setminus E$, and $m(E) < \epsilon$.

The set $\{x \in U | f(x) = 0\}$ has measure zero. Therefore, we can choose δ so that the set $\{x \in U$ such that $|f(x)| < \delta\}$ has measure less than ϵ.

Next we choose n large enough so that $|f_n(x) - f(x)| < \delta$ on $U - E$. Then the measures of $\{x \in U | f_n(x) > 0\}$ and $\{x \in U | f(x) > 0\}$ differ by at most 2ϵ. But ϵ was arbitrary. □

4. Ham sandwich for finite sets

We now adapt the ham sandwich theorem to finite sets of points. Instead of open sets U_i, we will have finite sets S_i. We say that a polynomial P bisects a finite set S if at most half the points in S are in $\{P > 0\}$ and at most half the points in S are in $\{P < 0\}$. Note that P may vanish on some or all of the points of S. We will give an example below to illustrate why we want this definition.

Corollary 4.1. *Let $S_1, \ldots, S_N$ be finite sets of points in $\mathbb{R}^n$ with $N < \binom{n+d}{n}$. Then there is a non-zero polynomial of degree $\leq d$ that bisects each set S_i.*

Let us give an example now. Suppose that we take two sets S_1 and S_2 in the plane, both lying on the x-axis, with $S_1 \subset [0,1] \times \{0\}$ and $S_2 \subset [2,3] \times \{0\}$. Since $2 < \binom{2+1}{2} = 3$, we should be able to choose a degree 1 polynomial to bisect both S_1 and S_2. The only option is to choose the x_1-axis : any transverse line will fail to bisect one of the two sets. Because of this situation, we have to allow p to "bisect" a finite set S in the case that p vanishes on S.

The proof of the theorem is to replace the finite sets by finite unions of δ-balls, apply the polynomial ham sandwich theorem, and then take $\delta \to 0$. We include the details here, but this is again just an analysis exercise.

Proof. For each $\delta > 0$, define $U_{i,\delta}$ to be the union of δ-balls centered at the points of S_i. By the polynomial ham sandwich theorem we can find a non-zero polynomial P_δ of degree $\leq d$ that bisects each set $U_{i,\delta}$. In fact, the proof of the ham sandwich theorem tells us that $P_\delta \in S^N \subset V(d) \setminus \{0\}$.

Now we can find a sequence $\delta_m \to 0$ so that p_{δ_m} converges to a polynomial $PinS^N \subset V(d) \setminus \{0\}$. Since the coefficients of P_{δ_m} converge to the coefficients of P, it's easy to check that P_{δ_m} converges to P uniformly on compact sets.

We claim that P bisects each set S_i. We prove the claim by contradiction. Suppose instead that $P > 0$ on more than half of the points of S_i. (The case $P < 0$ is similar.) Let $S_i^+ \subset S_i$ denote the set of points of S_i where $P > 0$. By choosing ϵ sufficiently small, we can assume that $P > \epsilon$ on the ϵ-ball around each point of S_i^+. Also, we can choose ϵ small enough that the ϵ-balls around the points of S_i are disjoint. Since P_{δ_m} converges to p uniformly on compact sets, we can find m large enough that $P_{\delta_m} > 0$ on the ϵ-ball around each point of S_i^+. By making m large, we can also arrange that $\delta_m < \epsilon$. Therefore, $P_{\delta_m} > 0$ on more than half of U_{i,δ_m}. This contradiction proves that P bisects S_i. $\square$

5. Cell decompositions

Theorem 5.1. *If S is any finite subset of $\mathbb{R}^n$ and d is any degree, then there is a non-zero degree d polynomial P so that each component of $\mathbb{R}^n \setminus Z(P)$ contains $\leq C(n)|S|d^{-n}$ points of S.*

Proof. Find a polynomial P_0 of degree 1 that bisects S. Some points of S lie in $Z(P_0)$. The rest lie in S_+ and S_-, which each have $\le |S|/2$ points. The sets S_+ and S_- are in different components of the complement of $Z(P_0)$. Next we find a low degree polynomial P_1 that bisects S_+ and S_-. Neglecting the points in $Z(P_1)$ we have four subsets of S left each with $\le |S|/4$ points. These four subsets lie in different components of the complement of $Z(P_0 P_1)$. We continue in this way to define polynomials P_2, P_3, etc. The polynomial P_j bisects 2^j sets. By the polynomial ham sandwich theorem, we can find P_j with degree $\le C(n) 2^{j/n}$. Each component of the complement of $Z(P_0 \cdot ... \cdot P_j)$ has $\le |S| 2^{-j}$ points.

We repeat J times, and we let $P = P_0 \cdot ... \cdot P_J$. Each component of the complement of $Z(P)$ has $\le |S| 2^{-J}$ points of S. We need to choose d so that $deg(P) \le d$, which means that $C(n) \sum_{j=0}^{J} 2^{j/n} \le d$. The sum is a geometric sum, and the last term is comparable to the whole. Therefore, we can arrange that $deg P \le d$ and also $2^{J/n} \gtrsim d$. Therefore, $2^J \gtrsim d^n$, and each component of the complement of $Z(P)$ has $\lesssim |S| d^{-n}$ points of S. $\square$

We should also give a caveat. The theorem does NOT guarantee that the points of S lie in the complement of $Z(P)$. In fact it is possible that $S \subset Z(P)$. There are two extreme cases. If all the points of S lie in the complement of $Z(P)$, then we get optimal equidistribution, and we have a good tool to do a divide-and-conquer argument. If all the points of S lie in $Z(P)$, then we see that $deg(S) \le d$, and we get a good degree bound on S. Generally, S will have some points in $Z(P)$ and some points in the complement. One part of S has a low degree and the other part of S is spread out well among the cells.