Topics in Geometric Combinatorics

Francis Edward Su

August 13, 2004

Chapter 1

Introduction

1.1 What is geometric combinatorics?

Geometric combinatorics refers to a growing body of mathematics concerned with counting properties of geometric objects described by a finite set of building blocks. Primary examples include polytopes (which are bounded polyhedra and the convex hulls of finite sets of points) and complexes built up from them. Other examples include arrangements and intersections of convex sets and other geometric objects. As we'll see there are interesting connections to linear algebra, discrete mathematics, analysis, and topology, and there are many exciting applications to economics, game theory, and biology.

There are many topics that could be discussed in a course on geometric combinatorics; but in these lectures I have chosen what I consider to be my favorite topics for inclusion in such a course. Mostly the lectures reflect what I like in the subject. Some of the topics are of relatively recent development and reflect either my current interests (e.g., combinatorial fixed point theorems) or material I assimilated in the "Discrete and Computational Geometry" program at MSRI in Fall 2003.

Along the way, I will provide some exercises to allow the reader to play with the concepts, and some pointers to problems in the field.

Chapter 2

Combinatorial convexity and Helly's theorem

Convexity is a geometric notion with some interesting combinatorial consequences.

2.1 Application: Centerpoints

We all know what a *median* of a data set of real numbers is— it is a number for which half the data lies above it, and half the data lies below. In other words, the median gives a notion of a "center" of a set of points on the real line.

But what if the data set are points in the plane or points in $\mathbb{R}^d$ for $d > 1$? Can one define a kind of "median"? For instance, for a data set in the plane, is there always a point in the plane for which any line through that point cuts the data set in half? Investigate by constructing some examples.

If not, is there a weaker notion? For instance, is there always a point in the plane so that any line through that point has at least 1/3 of the points on each side?

2.2 Convex sets

Definition 2.1. A subset C of $\mathbb{R}^d$ is said to be *convex* if for any two points in C, the line segment between them is also contained in C.

Draw examples of convex and non-convex sets.

Definition 2.2. The *convex hull* of a set A in $\mathbb{R}^d$, denoted by $\operatorname{conv}(A)$, is the intersection of all convex sets that contain A.

Since the arbitrary intersection of convex sets is convex (see HW), the convex hull of set is always convex.

Definition 2.3. A *convex combination* of points $a_1, ..., a_n$ in A is a linear combination $\sum_{i=1}^n \lambda_i a_i$ in which the coefficients λ_i are non-negative and sum to 1. The set of all convex combinations of points in A is called the *convex span* of A.

Proposition 2.4. *For any set A in $\mathbb{R}^d$, $conv(A)$ is identical to the convex span of A.*

Actually, one never needs more than $d+1$ points of A to express any point in $\text{conv}(A)$ as a convex combination.

Theorem 2.5 (Carathéodory's theorem). *For A in $\mathbb{R}^d$, each point of $conv(A)$ is a convex combination of at most $d+1$ points of A.*

Carathéodory's theorem can be viewed as a combinatorial consequence of a geometric notion. It can be proved as a consequence of another theorem that relates geometry and combinatorics.

Theorem 2.6 (Radon's Lemma). *Let A be a set of size $d+2$ (or greater) in $\mathbb{R}^d$. Then A can be partitioned in two sets R and B (red and blue) such that $conv(R) \cap conv(B) \neq \emptyset$.*

Definition 2.7. A set of points $x_1, ..., x_m$ in $\mathbb{R}^d$ is said to be *affinely dependent* if $\sum_i \lambda_i x_i = 0$ and $\sum_i \lambda_i = 0$ for some $\lambda_1, ..., \lambda_m$ not all zero. Thus in this case one of the points can be expressed as an affine combination of the other points. (Otherwise, if the λ_i must all be zero, the points are said to be *affinely independent.*)

Corollary 2.8. *Any set of $d+2$ (or more) points in $\mathbb{R}^d$ must be affinely dependent.*

Theorem 2.9 (Helly's Theorem). *Suppose $A_1, ..., A_m$ are convex sets in $\mathbb{R}^d$, such that every subset of size $d+1$ has non-empty intersection. Then the intersection $cap_{i=1}^m A_i$ is non-empty.*

Is there a version of Helly's theorem for an infinite number of convex sets? Well, if the sets are not closed, the conclusion may not hold (find an example). And if the sets are not bounded, you could run into problems as well. However,

Theorem 2.10. *If $\mathcal{A} = \{A_\alpha\}$ is an infinite collection of compact convex sets in $\mathbb{R}^d$ such that any $d+1$ of them have a common point, then all the sets in $\mathcal{A}$ have a common point.*

Helly's theorem has many interesting consequences. Each of the following can be proved using Helly's Theorem.

Theorem 2.11 (Kirchberger's Theorem). *Let S is a set of sheep and G is a set of goats in $\mathbb{R}^d$, such that there are at least $d+2$ animals in total. Suppose that for every set C of $d+2$ animals, the sheep and the goats can be separated by a hyperplane. Then S and G can be separated by a hyperplane.*

Theorem 2.12 (Jung's theorem). *Every set of diameter 1 in* $\mathbb{R}^d$ *lies in a closed ball of radius* $\sqrt{d/2(d+1)}$.

Theorem 2.13 (Krasnosselsky's Theorem). *Let* K *be an infinite compact set in* $\mathbb{R}^d$*. Suppose that for every* $d+1$ *points in* K*, there's a point of* K *from which all these points are visible in* K*. Then there's a point of* K *from which all of* K *is visible.*

A set K with this property is said to be *star-shaped*. A special case of the above problem is the gallery problem in the homework exercises.

2.3 Return to Application

Let A be a set of n points in $\mathbb{R}^d$. A point c is called a *centerpoint* of A if each closed half-space containing c contains at least $\frac{n}{d+1}$ points of A. Note that c does not have to be a point of A.

Theorem 2.14. *If* A *is a finite point set in* $\mathbb{R}^d$*, then* A *has at least one centerpoint.*

2.4 Exercises

1. Let $\{A_\alpha\}$ be a family of convex sets in $\mathbb{R}^d$. Show that the intersection $A = \cap_\alpha A_\alpha$ is convex. (We use the letter α here as an index to indicate that the indexing set might possibly be infinite or uncountable.)

2. In this problem, we investigate the relationship between the notions of *linear span* and *affine span* of set of points.

   (a) Let w_1, w_2 be two linearly independent points in $\mathbb{R}^3$. Explain why the linear span of w_1 and w_2 is a plane P that passes through the three points: w_1, w_2, and the origin.

   (b) Let x_0, x_1, x_2 be three points in $\mathbb{R}^3$, and let $w_1 = x_1 - x_0$, $w_2 = x_2 - x_0$. Explain why the plane P' that passes through x_0, x_1, x_2 is just the plane P (defined as above) translated by x_0.

   (c) Thus any point in P' has the form $x_0 + \lambda_1 w_1 + \lambda_2 w_2$ for real numbers λ_1, λ_2. Rewrite this expression as a linear combination of x_0, x_1, x_2 and verify that the coefficients of the x_i sum to 1, and conclude that the affine span of the three points x_i is exactly the plane P' through those points.

3. Given a finite set S in $\mathbb{R}^d$, let R_S be the set of all Radon points of S (over all possible Radon partitions).

   (a) Is R_S necessarily convex? Prove or provide a counterexample.

(b) If $|S| = d+2$, does R_S necessarily consist of a single point?

(c) If $|S| = d+2$, and if no $d+1$ points are affinely dependent, prove that R_S is a single point, i.e., the Radon point is unique. [Hint: think about the rank of some suitable matrix.]

4. Use Radon's lemma to prove Carathéodory's theorem.

5. Let $K \subset \mathbb{R}^d$ be a convex set and and let $C_1, ..., C_n \subset \mathbb{R}^d$, $n \geq d+1$, be convex sets such that the intersection of every $d+1$ of them contains a translated copy of K. Prove that the intersection of all the sets C_i also contains a translated copy of K.

6. Show that if you have a finite collection of parallel line segments in $\mathbb{R}^2$ such that every 3 of them are cut by a common transversal, then there is a transversal that cuts all the line segments.

7. An art gallery is in a room that has the shape of a simple closed polygon (finitely many sides, but not necessarily convex), and there is one painting hung on each wall. Suppose that for any 3 paintings, there is a point in the gallery where those 3 paintings are visible. Show that there is some point of the gallery from which all paintings are visible.

If you enjoy a challenge, try to prove the more general case of Krasnosselsky's theorem for extra credit, and come see me if you do it.)

Chapter 3

Polytopes

As we shall see, a *polytope* is a convex object that can be defined as a convex hull of a finite set of points, or as the intersection of a finite set of half-spaces. Such a geometric object has a natural combinatorial structure.

3.1 Applications: in other areas

The book of Cromwell(1997) contains a nice discussion of the uses of polytopes in art (perspective in painting, ornaments), architecture (pyramids), computer-aided geometric design, nature (crystals, compressed cells, polyhedral molecules), cartography (grids), philosophy (Kepler's model of the planet distances).

3.2 Polytopes and Faces

Definition 3.1. A $\mathcal{V}$-*polytope* is the convex hull of finitely many points in $\mathbb{R}^d$. An $\mathcal{H}$-*polyhedron* is the intersection of finitely many half-spaces in $\mathbb{R}^d$. An $\mathcal{H}$-*polytope* is a bounded $\mathcal{H}$-polyhedron.

A basic but non-trivial fact (that we shall prove soon) is that any $\mathcal{V}$-polytope is an $\mathcal{H}$-polytope, and vice versa, so we may speak of a *polytope* and be assured that there is a description of either kind.

Convex sets have faces.

Definition 3.2. Let K in $\mathbb{R}^d$ be a convex set. Any set H in $\mathbb{R}^d$ of the form $H_{\mathbf{a}}(b) = \{\mathbf{x}|\mathbf{a}\cdot\mathbf{x} = b\}$ for some $\mathbf{a}$ and b is called an (affine) *hyperplane*. Let us denote the associated *half-spaces* by $H^+ = \{\mathbf{x}|\mathbf{a}\cdot\mathbf{x} \geq b\}$, and $H^- = \{\mathbf{x}|\mathbf{a}\cdot\mathbf{x} \leq b\}$.

Definition 3.3. A subset K of F is said to be a *face* if there exists a hyperplane H such that (a) $F = H \cap K$ and (b) K is contained in H^-. Such an H is called a *supporting hyperplane* and a face may have many supporting hyperplanes. The *dimension* of F is the dimension of its affine span.

Note a convex set K always has as faces itself and the empty set $\emptyset$. For instance, the empty set is supported by the degenerate hyperplane $H_{\mathbf{0}}(-1)$ and K is supported by $H_{\mathbf{0}}(0)$.

Polytopes have finitely many faces. Moreover, in a polytope P, if F is a face, then any face of F is a face of P. This is not true for arbitrary convex sets (construct an example!).

Let $\mathcal{F}(P)$ be the poset of faces of a polytope P, ordered by inclusion. (Recall that a *poset* is a partially ordered set.) We call $\mathcal{F}(P)$ the *face lattice* of P.

Definition 3.4. Two polytopes P,Q are said to be *combinatorial equivalent* if they have isomorphic face lattices: $\mathcal{F}(P) \cong \mathcal{F}(Q)$.

3.3 Examples

We now mention several examples of polytopes. In each case, thinking combinatorially, it is fun to ask: what does the face lattice look like? Which subsets of vertices form faces? It is also instructive to compare the $\mathcal{V}$-description and the $\mathcal{H}$-description of a given polytope.

In dimension 2, the 2-dimensional polytopes are just the (convex) polygons, and there is exactly one combinatorial type of 2-polytope for each n, namely, the n-gon.

The *standard d-simplex* is

$$\mathrm{conv}(\{\mathbf{0}, \mathbf{e}_1, \mathbf{e}_2, ..., \mathbf{e}_d\}),$$

and a *simplex* is any combinatorially equivalent polytope.

The *standard d-cube* is

$$\mathrm{conv}(\{+1, -1\}^d),$$

and a *cube* is any combinatorially equivalent polytope.

The *standard d-crosspolytope* is

$$\mathrm{conv}(\{+\mathbf{e}_1, -\mathbf{e}_1, +\mathbf{e}_2, -\mathbf{e}_2, ..., +\mathbf{e}_d, -\mathbf{e}_d\}),$$

and a *cube* is any combinatorially equivalent polytope. The 3-crosspolytope is called an *octahedron.*

The *moment curve* in $\mathbb{R}^d$ is the curve $\mathbf{x}(t) = (t, t^2, ..., t^d)$, and the *cyclic polytope* is defined by

$$C_d(t_1, ..., t_d) := \mathrm{conv}(\{\mathbf{x}(t_1), ..., \mathbf{x}(t_d)\}.$$

The *standard cyclic polytope* $C_d(n)$ is the n-vertex cyclic polytope $C_d(1, 2, ..., d)$. It is a (surprising) fact that the combinatorial equivalence class of a cyclic polytope only depends on d and n and not on the argument t_i. (We shall prove this later, too.)

The *permutahedron* Π_{d-1} is a $(d-1)$-dimensional polytope defined in $\mathbb{R}^d$ as the convex hull of vectors that are permutations of the coordinates of $(1, 2, ..., d)$. From a picture of Π_2 and Π_3, do you notice any patterns? Which pairs of vertices will form edges? Which subsets of vertices form faces?

3.4 New Polytopes from Old

If P is a polytope $\mathrm{conv}(V)$ where V is a finite set of points, then the *pyramid* over P is $\mathrm{conv}(V \cup \{x\})$ where x is some point affinely independent of V.

The *bipyramid* over P is $\mathrm{conv}(V \cup \{x, y\})$ where x and y are affinely independent of P and such that $conv(\{x, y\}) \cap int(P) \neq \emptyset$.

The *Minkowski sum* $P + Q$ of two polytopes P and Q in $\mathbb{R}^d$ is the set $\{p + q : p \in P, q \in Q\}$. It is a polytope (why?) in $\mathbb{R}^d$.

The *product* $P \times Q$ of a polytopes P in $\mathbb{R}^m$ and Q in $\mathbb{R}^n$ is the set $\{(p, q) : p \in P, q \in Q\}$ in $\mathbb{R}^{m+n}$. It is also a polytope (why?).

The *prism* over P is $P \times I$ where I is the unit interval.

3.5 Schlegel diagrams

Schlegel diagrams provide useful ways of visualizing a d-polytope by projecting (the skeleton of) a polytope onto one of its faces. This is especially useful for 4-dimensional polytopes.

3.6 Exercises

1. Draw the face lattice of a pyramid over a square. (Number the vertices of the square by $1, 2, 3, 4$ and the top vertex by 5.)

2. Give a facet description of the standard d-crosspolytope as an $\mathcal{H}$-polytope (i.e., tell me all the half-spaces that define the polytope).

3. (a) Consider the permutahedron Π_3. How many *facets* does it have?
   (b) How many ordered partitions are there of the set $\{1, 2, 3, 4\}$ into two non-empty parts? [An *ordered partition* of a set is a partition of the set into parts, in which the order of the parts matters, but the order within the parts does not, e.g., (123/45) is same as (132/45) but different from (45/123).]
   (c) Can you see a correspondence between the two questions? (Refer to the figure of the permutahedron.)
   (d) Make a conjecture about the number of k-faces of this permutahedron and ordered partitions of the set $\{1, 2, 3, 4\}$.

4. For each of the examples of polytopes described in this chapter, see if you can determine which subsets of vertices span faces.

5. Can you construct a 3-polytope with seven edges? Why or why not?

Chapter 4

Polar Duality

4.1 The Polar Dual

Let $H_{\mathbf{a}}$ denote the hyperplane $H_{\mathbf{a}}(1)$: $\{\mathbf{x} : \mathbf{a} \cdot \mathbf{x} = 1\}$. Any hyperplane that does not pass through zero can be put uniquely in this form. Define the *half-spaces*

$$H_{\mathbf{a}}^{-} = \{\mathbf{x} : \mathbf{a} \cdot \mathbf{x} \leq 1\},$$

$$H_{\mathbf{a}}^{+} = \{\mathbf{x} : \mathbf{a} \cdot \mathbf{x} \geq 1\}.$$

There is an interesting geometric duality between non-zero points and hyperplanes in $\mathbb{R}^d$, namely to each point $\mathbf{a}$ we can associate the hyperplane $H_{\mathbf{a}}$, and vice versa. Moreover, by considering the definitions, it is easy to see that

$$\mathbf{y} \in H_{\mathbf{x}} \iff \mathbf{x} \in H_{\mathbf{y}}$$

but this is somewhat surprising geometrically. Also, verify that

$$\mathbf{y} \in H_{\mathbf{x}}^{-} \iff \mathbf{x} \in H_{\mathbf{y}}^{-}. \tag{4.1}$$

Definition 4.1. Given a set X in $\mathbb{R}^d$, the *polar dual* $X^{\triangle}$ is defined by

$$X^{\triangle} := \{\mathbf{y} : \mathbf{y} \cdot \mathbf{x} \leq 1, \forall \mathbf{x} \in X\}.$$

Notice that this may be rewritten:

$$X^{\triangle} = \cap_{\mathbf{x} \in X} H_{\mathbf{x}}^{-}, \tag{4.2}$$

the intersection of half-spaces And because of (4.1), we may also say:

$$X^{\triangle} = \cup_{X \subseteq H_{\mathbf{y}}^{-}} \{\mathbf{y}\}.$$

(Yes, this could be written $\cup\{\mathbf{y} : X \subseteq H_{\mathbf{y}}^{-}\}$ but we write it in the format above to emphasize the duality between it and the previous displayed equation.)

Do a few examples here. What is the polar dual of a non-zero point? What is the polar dual of the origin? What is the polar dual of four points of a square? A filled in square? A circle? A halfspace $H^-_{\mathbf{x}}$? A halfspace $H^+_{\mathbf{x}}$?

After doing a few examples of polar duals of convex hulls of a set of points V, one may notice in (4.2) that one need not intersect all the half-spaces— it is sufficient to just use half-spaces corresponding to the vertices of V.

Theorem 4.2. *If $X = conv(V)$, then:*

$$X^{\triangle} = \cap_{\mathbf{v}\in V} H^-_{\mathbf{v}}.$$

Some other properties are easy to check. For any X in $\mathbb{R}^d$, $X^{\triangle}$ is closed, convex, and contains $\mathbf{0}$. If $X \subseteq Y$, then $Y^{\triangle} \subseteq X^{\triangle}$.

Also, if $X^{\triangle\triangle}$ denotes the polar dual of the polar dual of X, then $X \subseteq X^{\triangle\triangle}$. $X^{\triangle\triangle}$ is sometimes called the *double-polar* of X, and you may check that it is just $\{\mathbf{z} : \mathbf{y}\cdot\mathbf{x} \le 1, \forall \mathbf{x} \in X \implies \mathbf{y}\cdot\mathbf{z} \le 1\}$.

The double-polar of X won't always be equal to X, since even if X doesn't contain $\mathbf{0}$, $X^{\triangle\triangle}$ will contain $\mathbf{0}$. But one can show the following theorem. Recall the *closure* of a set A is the smallest closed set containing A.

Theorem 4.3. *$X^{\triangle\triangle}$ is the closure of $conv(X \cup \{\mathbf{0}\})$.*

This can be shown by appealing to another property of convex sets:

Theorem 4.4 (Hyperplane Separation). *If C, D are disjoint convex sets in $\mathbb{R}^d$, there exists a hyperplane separating them, i.e., there exists an $\mathbf{a}$ such that $C \subseteq H^-_{\mathbf{a}}$, and $D \subseteq H^+_{\mathbf{a}}$. Moreover, if C, D are closed sets, then the hyperplane separates them* strictly*, i.e., $H_{\mathbf{a}}$ does not intersect either set.*

For our purposes, this is most often used when C is a point and D is a polytope. In that case, the proof is quite easy; just take the line achieving the minimum distance between the point and polytope, and construct a hyperplane orthogonal to that line, intersecting it at its midpoint.

The following fact is a useful observation about polar duals:

Lemma 4.5. *If $C \subseteq \mathbb{R}^d$ is convex, then $C^{\triangle}$ is bounded if and only if $\mathbf{0}$ is in the interior of C.*

(A point x is in the *interior* of C if there exists a closed ball of positive radius around x that is contained in C.)

We can now show the following main theorem about polytopes, which dates back to Minkowski and Weyl:

Theorem 4.6. *Any $\mathcal{V}$-polytope is an $\mathcal{H}$-polytope, and vice versa.*

Why is this such a useful theorem? Try showing that the intersection of polytopes is a polytope. This is easy to see in the facet description, but not vertex description. Now try showing that the projection of a polytope on a plane is a polytope. This is easy to see in the vertex description, but not the facet description.

4.2 Polar Duals of Faces

We have seen from (4.2) that the vertices of a polytope P give rise to corresponding half-spaces in the polar $P^\triangle$. These in fact define facets of $P^\triangle$. An even stronger statement can be made.

Theorem 4.7. *The j-faces of a polytope P are in 1-1 correspondence with the $(d-j-1)$-faces of its polar dual $P^\triangle$. This correspondence is inclusion-reversing and hence the posets $\mathcal{F}(P)$ and $\mathcal{F}(P^\triangle)$ are opposite (i.e., the same posets but with partial orders reversed).*

This can be seen by constructing the correspondence explicitly. If F is a face of P, then define

$$F^\diamond := \{\mathbf{c} : \mathbf{c}\cdot\mathbf{x} \leq 1\ \forall \mathbf{x} \in P, \mathbf{c}\cdot\mathbf{x} = 1\ \forall \mathbf{x} \in F\}.$$

This is clearly a subset of $P^\triangle$ (same definition but with more constraints) and one can check that it is indeed a face of $P^\triangle$, that $F^{\diamond\diamond} = F$ and that $F \subseteq G$ iff $G^\diamond \subseteq F^\diamond$.

This means that in a 3-polytope, vertices of P correspond to facets of $P^\triangle$, and edges of P correspond to edges of $P^\triangle$, and facets of P correspond to vertices of $P^\triangle$. (Also, the empty face of P corresponds to all of $P^\triangle$, and all of P corresponds to the empty face of $P^\triangle$.)

Check that the polar of the d-cube is the d-crosspolytope, and their face lattices are opposite. What is the polar dual of the dodecahedron? The polar dual of the simplex?

4.3 Simple and Simplicial Polytopes

Some special classes of polytopes are the simple and simplicial polytopes.

Definition 4.8. A polytope P is said to be *simple* if each vertex is in exactly d facets. A polytope P is said to be *simplicial* if every facet is a simplex (i.e., every facet contains d vertices).

These two classes of polytopes are polar duals of each other (as witnessed by the fact that their definitions are polar to each other).

As examples: the crosspolytope is simplicial, the cube is simple, the simplex is both simple and simplicial, and the pyramid over a square base is neither.

4.4 Graphs of Polytopes

The *graph* $G(P)$ of a polytope P is the graph formed by just the vertices and edges of P. What can we learn about P from its graph? Can every graph appear as the graph of some polytope P? Here are a few famous theorems about graphs of polytopes.

A graph is *simple* if it contains no loops or doubled edges. A graph is said to be d-*connected* if removing any $d-1$ vertices leaves the graph connected.

Theorem 4.9 (Balinski). *If P is a d-polytope, $G(P)$ is d-connected.*

Project Idea 1. Understand the proof of this theorem and some of its generalizations.

Theorem 4.10 (Steinitz). *A finite graph G is isomorphic to $G(P)$ for some 3-polytope P if and only if G is simple, planar, and 3-connected.*

A very surprising fact holds for simple polytopes:

Theorem 4.11 (Blind and Mani-Levitska, 1987). *If P is simple, then $G(P)$ completely determines $\mathcal{F}(P)$.*

Project Idea 2. Understand the nice short proof of this result by Kalai(1988)

Theorem 4.12. *For a polytope P that is simple or simplicial, there exists some combinatorially equivalent polytope Q with integer vertex coordinates (i.e., it is a* lattice polytope*).*

This is not true in general— for instance, there is a 4-polytope which cannot be a lattice polytope in any realization!

4.5 Exercises

1. Show that $\mathbf{y} \in H_{\mathbf{x}}^- \iff \mathbf{x} \in H_{\mathbf{y}}^-$.

2. Let X, Y be sets in $\mathbb{R}^d$. Show that if $X \subseteq Y$, then $Y^\triangle \subseteq X^\triangle$.

3. Show that for any X in $\mathbb{R}^d$, $X \subseteq X^{\triangle\triangle}$.

4. Let $C \subseteq \mathbb{R}^d$ be a convex set. Show that $C^\triangle$ is bounded if and only if $\mathbf{0}$ is in the interior of C. (A point x is in the *interior* of C if there exists a closed ball of positive radius around x that is contained in C.)

5. Using (4.2), describe the polar dual $C^\triangle$ of each of the following convex polytopes C as the intersection of a finite number of half-spaces. (Specify the equations for those half-spaces.)

   (a) Let C be the d-cube: the convex hull of all points of all d-tuples of the form $(\pm 1, \pm 1, ..., \pm 1)$.

   (b) Let C be the 24-*cell*: the convex hull of all points of the form $\pm\mathbf{e}_i \pm \mathbf{e}_j$ for $i, j = 1, 2, 3, 4$ and $i \neq j$. (This is a very special polytope in $\mathbb{R}^4$ with a lot of symmetry.)

6. Draw the face lattices of a 3-cube and a 3-crosspolytope (octahedron), and verify that there is a 1-1 correspondence of faces that takes j-faces to $(3-j-1)$-faces and reverses inclusion.

7. Show that the set $S = \{x : 0 \leq x_1 \leq x_2 \leq ... \leq x_d \leq 1\}$ is a simplex in two ways:

   (a) by writing it as the convex hull of d+1 points.

   (b) by writing it as the intersection of d+1 half-spaces.

8. Consider the 4-polytope P that is the triangle cross the triangle. Which graph is the graph of the polar dual of P? (Recall that the *graph* of a polytope is the graph of its vertices and edges.)

9. Let Q be a polytope containing the origin in $\mathbb{R}^d$, with a half-space description: $Q = \cap H_{\mathbf{q}}$ for some collection of halfspaces $H_{\mathbf{q}}$. We now wish to construct the bipyramid over Q in $\mathbb{R}^{d+1}$ by taking the cones from $(0, 0, ..., 1)$ and $(0, 0, ..., -1)$ to $Q \times \mathbf{0}$ (which is just a copy of Q in the hyperplane $x_{d+1} = 0$).

   For each half-space $H_{\mathbf{q}} = \{\mathbf{x} : \mathbf{q} \cdot \mathbf{x} \leq 1\}$ that defines Q, there are two corresponding half-spaces that define the bipyramid over Q in $\mathbb{R}^{d+1}$. What are they? (When you define a vector, be careful to state whether your vectors live in $\mathbb{R}^d$ or $\mathbb{R}^{d+1}$.)

10. Let P be any d-polytope. Show that the bipyramid of the polar dual of P is the polar dual of the prism over P.

11. (do not need to turn in, but let me know if you make progress) Suppose a convex n-gon P in the plane has integer vertex coordinates. Let $f(n)$ be the smallest integer coordinates required to do this. Try to find upper and lower bounds for $f(n)$.

Chapter 5

Combinatorics of Faces

5.1 Which subsets of vertices span faces?

In the following examples, which subsets of vertices span faces? The d-simplex. The d-crosspolytope.

For cyclic polytopes, the combinatorics of the faces is very well understood, and quite interesting.

Theorem 5.1 (Gale's evenness condition). *The cyclic polytope $C_d(n)$ is simplicial and a d-subset $S \subseteq [n]$ forms a facet if and only if for all pairs $i, j \notin S$, the number of integers between i and j is even.*

So the combinatorial type of $C_d(n)$ doesn't depend on the parameters!

Corollary 5.2. *The cyclic polytope $C_d(n)$ is $\lfloor d/2 \rfloor$-neighborly, i.e., every subset $S \subseteq [n]$ of size $d/2$ vertices forms a face!*

This is counterintuitive for ≥ 4 in $C_d(n)$, because it says that every pair of vertices has an edge between them, even if n is much bigger than d. Thus the graph of cyclic polytope $C_4(n)$ is the complete graph on n vertices!

It is known from McMullen's Upper Bound Theorem that cyclic polytopes have the maximum possible number of k-faces for fixed n, d.

5.2 Euler's Formula

If P is a polytope, let $f_k = f_k(P)$ denote the number of k-faces of P. Thus f_0 is the number of vertices, f_1 the number of edges, f_{d-1} the number of facets, and $f_d = 1$ and $f_{-1} = 1$ since the full polytope is a face of dimension d and we can regard the empty set as a face of dimension -1. We call the vector $(f_{-1}, f_0, ..., f_{d-1}, f_d)$ the f-vector of a d-polytope.

A very important question is to characterize the f-vectors of polytopes. Which vectors can be f-vectors for some polytope?

There is a linear relation:

Theorem 5.3 (Euler-Poincaré formula). *For any d-polytope P,*

$$f_0 - f_1 + ... + (-1)^{d-1} f_{d-1} = 1 - (-1)^d.$$

This can also be written:

$$-f_{-1} + f_0 - f_1 + ... + (-1)^d = 0.$$

In dimension 3, this says that the number of vertices V minus the number of edges E plus the number of facets F must be 2. By stereographic projection, this can be viewed as a statement about planar graphs in the plane. In this setting, the formula is true for any connected planar graph:

Theorem 5.4. *For any connected planar graph, $V - E + F = 2$.*

For 3-polytopes, we can establish some other relations between the face numbers, known to Euler. For instance,

$$2E \geq 3F$$

and

$$2E \geq 3V.$$

Also, $V \geq 4$ and $F \geq 4$.

Also, from the equations above, we find $V \geq \frac{1}{2}F + 2$ and $V \leq 2F - 4$. In the V, F-plane we find that polytopes must lie in the cone formed by these two lines. Can all lattice points in this cone be achieved? Investigate, and try to construct examples for each such lattice point.

For d-polytopes, there are many inequalities that the f_i must satisfy, and huge literature on f-vectors.

5.3 Products

There are some nice properties of f-vectors and products of polytopes. What is the relationship between the face numbers f_i for the 2-simplex (triangle) and the face numbers f_i for the product of 2-simplex with itself? Compare to the algebraic expression $(3 + 3x + x^2)^2 = 9 + 18x + 15x^2 + 6x^3 + x^4$. Yes, generating functions are at work here, because the faces of the product are just products of the faces in the factors.

It is also interesting to explore the relationship between various other operations on polytopes (e.g., prisms, pyramids, etc.) and see what happens to the face numbers.

5.4 Exercises

1. Consider the standard d-simplex $P = \text{conv}(\{\mathbf{0}, \mathbf{e}_1, ..., \mathbf{e}_d\})$. Let S be a subset of the vertices that contains $\mathbf{0}$. Show that S forms a face of P by exhibiting a hyperplane that passes through the points of S and contains P in one half-space.

2. Based on inequalities developed thus far for 3-polytopes, show that $V \leq 2F - 4$.

3. Exhibit a 3-polytope for each lattice point on the line $V = \frac{F}{2} + 2$.

4. Exhibit a 3-polytope for each lattice point on the line $V = 2F - 4$.

5. A 3-polytope is said to be *regular* if (i) all facets are identical regular polygons, and (ii) the same number of polygons meet at each vertex. Show that there are only 5 possible regular 3-polytopes.

6. Show another theorem of Euler: that every 3-polytope contains at least one triangular, square, or pentagonal facet.

   [Hint: Let F_i be the number of i sided faces. Then $F = \sum_{i=3}^{\infty} F_i$. Now write $2E$ in terms of the F_i.]

7. Let P be the 4-polytope obtained by taking the product of a triangle T^2 with a square Q^2. Draw *two* different Schlegel diagrams for P.

   Justify your diagrams by also computing the number of faces that P should have in each dimension.

Chapter 6

Phylogenetic Trees

Here's an important problem in biology. Assuming that all life is related by some grand "tree of life", how can one reconstruct that tree from data?

For instance, one might have DNA sequences for various species. It is a reasonable to assume that if two species are close, then their DNA sequences will be similar; if so, then one might estimate a "distance" between species by looking at the number of places in which their DNA sequences differ.

This is just one example of how to measure a "distance" between species. Another method might involve comparing common morphological characteristics, for instance. We won't be concerned here with how we might actually convert data into a set of pairwise distances (although mathematics is very important in the modeling of this aspect of the problem).

What we shall focus on is: given pairwise distances $d(i, j)$ between every two species i and j, how can one find a tree that best represents the data? (Hopefully, such a tree will accurately model the evolutionary relationships between all the species.) For instance, can one tell when a metric d does represent distances along some tree?

6.1 Terminology

A *tree* is a graph that has no cycles. Let T represent a *weighted tree*, i.e., a tree with edge weights, and let L represent its *leaves*, i.e., all the nodes with degree 1. Let $d_T(i, j)$ be the metric on the leaves that takes the unique path from i to j and sums the tree weights along that path. We call d_T a *tree metric*.

A *cherry* in a tree T is a pair of leaves i, j which has just one intermediate node (the *cherry node*) on the unique path between them. Call the edge that connects the cherry node to the rest of the tree the *cherry stem*.

A *equidistant* tree T is (i) rooted, (ii) has the same distance from root to any leaf (called the *height* of T), and (iii) has non-negative weights on interior edges. (Note that weights on leaf edges can be non-negative.)

6.2 Ultrametrics

A metric d is an *ultrametric* on L if for any i, j, k in L,

$$d_{ij} \leq \max\{d_{ik}, d_{jk}\},$$

or equivalently, the maximum of $\{d_{ij}, d_{jk}, d_{ik}\}$ is achieved at least twice.

Theorem 6.1. *A metric d is an ultrametric if and only if $d = d_T$ for some equidistant tree T.*

It is easy to see that if $d = d_T$ for some equidistant tree, then d must satisfy the ultrametric property, since for any three leaves i, j, k, the subtree spanned by those leaves will contain two nodes i, j that are closest together and one leaf k whose distances from i, j must be the same. Thus the maximum of the three distances d_{ij}, d_{jk}, d_{ik} is achieved at least twice.

The other direction is harder, and involves constructing a tree T for any given ultrametric d. This can be done by induction, and the basic idea is to take i, j where d_{ij} is minimal and replace those leaves by a single leaf z. Then define a new leaf set (with one fewer leaf) by letting L' be the same as L but with i, j removed and z added. Define a metric d' on L' which agrees with the metric d on all leaves in $L' \setminus \{z\}$ and for each k in $L' \setminus \{z\}$, define $d'(z, k) := d(i, k)$. Then d' will be an ultrametric, and by the inductive hypothesis it will have a tree T' that represents it. Now replace the leaf z by a *cherry* using leaves i and j, whose with branch lengths from the cherry node equal to $d_{ij}/2$. The cherry stem length is chosen so that the total length of the stem and $d_{ij}/2$ is just the length of the original leaf edge for z. This gives the desired tree T.

We can thus define the *space of (equidistant) n-leaf trees* to be the set of all ultrametrics, which can be considered a subset of $\mathbb{R}^{\binom{n}{2}}$. What does this space look like?

For trees of fixed height H, the space of trees on 3 leaves is a subset of $\mathbb{R}^3$ consisting of all triples (d_{23}, d_{13}, d_{12}) in which the maximum is achieved at least twice. One may check that this space is topologically a cone over 3 points. The cone point is the triple (H, H, H), and the three rays that emanate from it represent 3 combinatorial types of trees, parametrized by the length of the interior edge.

Similarly, for fixed height H, the space of trees on 4 leaves turns out to be a cone over the Petersen graph.

6.3 The Tree Metric Theorem

A metric d satisfies the *four-point condition* if for all i, j, k, l in L, the maximum of $\{d_{ij} + d_{kl}, d_{ik} + d_{jl}, d_{il} + d_{jk}\}$ is achieved at least twice.

Theorem 6.2 (Tree Metric Theorem). *A metric d arises from a metric on a weighted tree if and only d satisfies the four-point condition.*

In one direction, the result is clear; any tree metric must satisfy the four-point condition, since any subtree generated by 4 leaves topologically consists of two cherries with a common stem.

In the other direction, the idea is to pick an arbitrary leaf l in L, and construct a new distance $\delta^l(i,j) := d_{ij} - d_{il} - d_{jl}$. One may check that this is an ultrametric by using the four-point condition and subtracting $d_{il} + d_{jl} + d_{kl}$ from it. Then one may construct an rooted equidistant tree that represents δ^l. To change this into a tree that represents d, we need to add weight d_{il} to leaf edge for i, and add an edge from the root to a new leaf l of length $-H$. The resulting tree will realize d.

References for this subject and these theorems can be found in the new books by Semple-Steel, *Phylogenetics*, and Felsenstein, *Inferring Phylogenies*.

6.4 Connections to Polytopes

Given a metric d, one may consider the set

$$P_d = \{x \in \mathbb{R}^n : x_i + x_j \geq d_{ij}, \quad 1 \leq i \leq j \leq n\}.$$

This is a convex polyhedron. Let B_d be the set points in bounded faces of P_d.

Then if d arises from a tree metric, then B_d is the desired metric tree, under the supremum norm!

Why is this? For any point x in a tree, let x_i be the distance of that point from vertex i along the tree. Thus every point x is represented by a n-tuple in $\mathbb{R}^n$. Then clearly for these x_i, the inequalities of P_d must be satisfied, and equality must hold for at least $n-1$ of the defining equalities and these inequalities are all independent. The intersection of $n-1$ hyperplanes will therefore correspond to 1-face, or an edge. The remaining inequalities will keep this edge bounded. Hence x lives in a bounded edge of P_d.

One may also check that the actual distance between two points is given by the supremum norm on difference in their coordinates.

6.5 Exercises

1. Show that these conditions (for being an ultrametric) are equivalent: for any i, j, k in L, (i) $d_{ij} \leq \max\{d_{ik}, d_{jk}\}$, and (ii) the maximum of $\{d_{ij}, d_{jk}, d_{ik}\}$ is achieved at least twice.

2. (a) Pick any weighted tree, and write out the corresponding tree metric d for each pair of leaves.

(b) Using the metric d that you found in part (a), reconstruct a tree that represents d.

Chapter 7

Tropical Geometry

In this section we introduce some very recent ideas about a funky geometry called *tropical geometry* that has connections to many areas of mathematics. Some ideas are motivated by considerations from algebraic geometry. And, as we shall, see there are some interesting connections to phylogenetic trees.

The subject is so recent that almost all of the papers that have been written on the subject can be found as preprints on the mathematics `arXiv` by doing a search on the word *tropical*. Many have been written by Bernd Sturmfels and co-authors. We shall attempt to build some intuition for this subject here.

7.1 The Tropical Semi-ring

The real numbers $\mathbb{R}$ come equipped with two operations $+$ and $\times$, which turns it into a ring. The beginnings of *tropical* arithmetic start with equipping $\mathbb{R}$ with two other operations $\oplus$ and $\otimes$ that turn it into a semi-ring. We let $\oplus$ represent the operation of taking minimums, and $\otimes$ represent the usual addition in $\mathbb{R}$. One may check that $\otimes$ distributes over $\oplus$. We may put in ∞ to get a "0"-like element, something that behaves like an additive identity. Note that 0 behaves like a multiplicative identity.

It will be convenient to work in *tropical projective space* $\mathbb{T}P^{n-1}$, which are equivalence classes of vectors in $\mathbb{R}^n$ under the relation that two vectors are equivalent if they differ (under usual arithmetic) by a real scalar multiple of $(1, 1, ..., 1)$. Thus $(6, 0, 7) = (0, -6, 1)$ in $\mathbb{T}P^2$. As a canonical representative we will usually make the first coordinate 0.

7.2 Tropical Convexity

We can define a *tropical linear combination* in the usual fashion, but with tropical operations. Thus $a \otimes \mathbf{v} \oplus b \otimes \mathbf{w}$ is a tropical linear combination of vectors $\mathbf{v}, \mathbf{w}$. What do the set of all tropical linear combinations (the *tropical linear span*) of two points look like? Investigate.

We can also define what it means to be *tropically convex*: we say that a set S in $\mathbb{R}^n$ is tropically convex if for any two points $\mathbf{v}, \mathbf{w}$ in S and scalars a, b, the combination $a \otimes \mathbf{v} \oplus b \otimes \mathbf{w}$ is also in S. You may be wondering at this point why we don't require $a \oplus b$ to equal the multiplicative identity 0, by analogy with usual geometry. Actually, we could require $a \oplus b = 0$ but a moment's reflection will reveal that this does not really add any restriction in $\mathbb{T}P^{n-1}$.

We can thus define a *tropical convex hull* of a set S as the tropical linear span of S. Similarly, a *tropical polytope* is then the tropical convex hull of a finite set of points.

Then we may ask which of the theorems in the usual geometry on $\mathbb{R}^n$ have analogues in tropical geometry. For instance, Develin-Sturmfels have shown that the intersection of tropically convex sets are still tropically convex, and one may also prove a *tropical Caratheodory* theorem as well as a *tropical hyperplane separation* theorem.

7.3 Tropical Hyperplanes

What does it mean to define a *tropical hyperplane*? We consider a linear form

$$a_1 \otimes x_1 \oplus a_2 \otimes x_2 \oplus \ldots \oplus a_n \otimes x_n$$

and instead of setting it equal to zero to get a hyperplane (as we would in usual geometry), we will demand (since this linear form represents the computation of some minimum) that *the minimum of this expression is achieved at least twice.*

Do some examples to see what a tropical hyperplane in $\mathbb{T}P^2$ looks like.

7.4 Connections to Phylogenetic Trees

As it turns out, a *tropical line* in $\mathbb{T}P^{n-1}$ corresponds to a tree on n-leaves. The space of such lines (the *tropical Grassmannian*) is the space of trees on n leaves!

Moreover, one may also define a notion of *tropical determinant*, which is like the usual determinant, interpreted tropically (and without minus signs). The following theorems, due to Develin-Sturmfels, shows an interesting connection between metrics, tree metrics, and minors of the corresponding distance matrices.

Theorem 7.1. *A symmetric matrix D is a metric if and only if the principal 3×3 minors of $-D$ are tropically singular.*

Theorem 7.2. *A metric D is a tree metric if and only if the principal 4×4 minors of $-D$ are tropically singular.*

The proof of the first result follows by noting that the 3×3 determinant is a minumum of six expressions, and to check that it is singular means that the minimum is achieved twice. But those six expressions are just $0, -2d_{12}, -2d_{23}, -2d_{13}, -d_{12} - d_{13} - d_{23}, -d_{12} - d_{13} - d_{23}$, the last two of which are identically equal. One may check that if the minimum is achieved twice, it must be achieved by the last two expressions, and the fact that it is less than or equal to all the other expressions gives the required triangle inequalities.

The proof of the second result is similar, with a lot more algebra.

Chapter 8

Triangulations

One way to try to understand smooth objects combinatorially is by breaking it into "nice" pieces, such as simplices.

Points in a simplex can be described by *barycentric coordinates*. If $\sigma = \{x : x = \sum t_i a_i, \sum t_i = 1\}$, then the t_i are called *barycentric coordinates* of x with respect to the points $\{a_i\}$.

A *simplicial complex* K in $\mathbb{R}^d$ is a collection of simplices such that: (i) $\sigma \in K$ implies every face of σ is in K, (ii) the intersection of any 2 simplices is a face of each.

A *subcomplex* is a subcollection of K that's also a complex. Let $K^{(p)}$ denote the subcomplex of K consisting of all faces with dimension at most p, called the *p-skeleton* of K. Thus $K^{(0)}$ are the vertices of K, and $K^{(1)}$ is the graph of edges and vertices.

Simplicial complexes are very nice because all the information about intersections of simplices is combinatorial. We can define an *abstract simplicial complex* to be a collection of finite sets such that if some subset A is in the collection, so are all subsets of A.

Theorem 8.1. *Every abstract simplicial complex can be realized.*

Thus is it enough to specify the maximal faces.

Why study simplicial complexes? (i) They are good ways of representing geometric or topological objects as combinatorial objects, (ii) they pop up in interesting places, and (iii) they provide good ways of computing things piece-by-piece.

A *triangulation* is a simplicial complex which is *pure*, i.e., all the maximal faces are of the same dimension. Some interesting questions to ask: given a space X, how many triangulations are there? What is the size of the minimal triangulation? What is the structure of the space of triangulations?

8.1 Triangulations of polygons

For instance, consider a polygon, and try to answer the questions above. How many triangulations does an n-gon have?

Theorem 8.2. *The number of triangulations of an n-gon is*

$$\frac{1}{n-1}\binom{2n-4}{n-2}.$$

For a direct proof of the above result, see the homework exercises.

You may recognize this last expression as the $(n-2)$-th Catalan number. As such, it is not surprising that there are many equivalent ways of counting the triangulations of an n-gon.

Theorem 8.3. *The number of triangulations of a (convex) n-gon is the same as:*

(i) the number of rooted, planar binary trees with $n-1$ leaves,

(ii) the number of ways to parenthesize the product of $n-1$ factors,

(iii) sequences of $n-2$ plus signs and $n-2$ minus signs such that in every initial segment the number of plus signs is never smaller than the number of minus signs, and

(iv) the number of paths from $(0,0)$ to $(n-2,n-2)$ in the integer grid not going above the diagonal and always moving right or up.

See if you can construction the bijections required by the statement of the theorem.

8.2 Triangulations of cubes

How can one triangulate a d-cube? There is a standard triangulation with $d!$ simplices; can you find it? The idea is to associate to each permutation a simplex of the cube in a natural way, an verify that these simplices meet face-to-face, and do not overlap.

What is the minimal triangulation of a 3-cube? Of a d-cube in general? This is a difficult question, but exact answers for various kinds of triangulations are known in low dimensions.

8.3 Flips

There are relationships between triangulations. For instance, in a planar triangulation, two triangles that share an edge can be replaced by two other triangles that represent the same quadrilateral, but with the internal diagonal "flipped". This leads to a *flip graph* of triangulations, which is a graph in which each node is a triangulation, and two nodes are connected by an edge if there is a flip that connects the corresponding triangulations.

What can we say about the flip-graph of a polygon? Try drawing one for a pentagon or a hexagon.

There a higher-dimensional analogue of a "flip". Can you figure out how it should be defined?

8.4 Triangulations of Point Sets

Given a point set in the plane, is there a good way to triangulate its convex hull so that the point set forms the vertices of a triangulation, and the triangulation is "nice"? For instance, we might not want long skinny triangles if we can use shorter and fatter ones.

Given a point set V in the plane, and a point p in V, let $R(p) = \{x : d(x,p) \leq d(x,q)\ \forall q \in V\}$. Thus $R(p)$ is the set of all points in the plane that are closer to p than to any other point in V.

The sets $R(p)$ for each p in V are called *Voronoi cells* and the subdivision of the plane into such cells is called a *Voronoi diagram*. The *Delaunay triangulation* is obtained from the Voronoi cells by connecting p, q in V by an edge iff the cells $R(p)$ and $R(q)$ are adjacent (share an edge). In some sense this is the "dual" of the Voronoi diagram.

As long as no 3 points of V are collinear and no 4 points are co-circular, this process will triangulate $\mathrm{conv}(V)$. The Delaunay triangulation has some very nice properties—it makes nice, non-skinny triangules for instance: it maximizes the minimum angle occurring in any triangulation of V, and it minimizes the maximum circumradius of any triangulation of V. Such results are usually proved by "local flipping".

8.5 Application: Robot Motion Planning

Given a finite set of obstacles in the plane, what path should a robot take to avoid these obstacles?

A nice answer is given by using a Voronoi diagram. If V is the set of locations of obstacles, and if motion is possible for the robot through the obstacle course, it should be possible along edges of the Voronoi diagram.

Project Idea 3. There are many other applications of Voronoi diagrams and Delaunay triangulations, in biology (medial axis transforms), chemistry (Wigner-Seitz zones), crystallography (domains of action), and meterology (Thiessen polygons). Investigate and explain.

8.6 Voronoi diagrams and polyhedra

There's a nice connection between Voronoi diagrams and polyhedra. Let $u(\mathbf{x}) = x_1^2 + ... + x_d^2$. Let U be the unit paraboloid given by the graph of u: $\{(x_1, ..., x_n, u(\mathbf{x}))\}$. Let H_p be the hyperplane tangent to U at $u(p)$.

Theorem 8.4. *The Voronoi diagram of V is the vertical projection of the facets of the polyhedron $\cap_{p \in V} H_p^+$ onto the hyperplane $x_{d+1} = 0$ in $\mathbb{R}^{d+1}$.*

8.7 Exercises

1. Consider a square with 4 vertices $\{a, b, c, d\}$, and look at the collection F of faces: $\{abcd, ab, ac, ad, bd, a, b, c, d, \emptyset\}$. Why is F not an abstract simplicial complex? Which condition does it not satisfy?

2. In any triangulation of an n-gon, show that: (a) the number of triangles is exactly $n-2$, and (b) the number of internal edges is exactly $n-3$. [Hint: use Euler's relation, and the fact that regions are triangles.]

3. Show that the number of triangulations of an n-gon is
$$\frac{1}{n-1}\binom{2n-4}{n-2}$$
by letting T_n deonte the set of all triangulations of the n-gon, and let t_n be the size of T_n. Label the vertices of an n-gon counterclockwise by $1, ..., n$. Define a function f from T_n to T_{n-1} by contracting the $(1, n)$ edge. Show that f is surjective, and show that t_n is the sum over all trees in T_n of the degree of vertex 1 in that tree.

   Since this can be done at all vertices, whos that $(n-1)t_n = 2(2n-5)t_{n-1}$, from which the conclusion follows.

4. Use one of the 5 ways of counting T, the number of triangulations of an n-gon to find this crude upper bound for the total number of such triangulations:
$$T \leq 2^{2n-4}.$$
[Thus without having to do the work of showing an exact formula for T as we did in class, you can easily see a bound.]

5. Look at the handout for the flip-graph of a hexagon. (a) what edge is missing? Draw it in. (b) The structure of this diagram suggests a division of the page into 2-faces, edges, and vertices.

   Consider any 2-face on your diagram, and determine what do all the triangulations on that face have in common? So then, what property characterizes the 2-faces on your diagram? What property characterizes the 1-faces (edges) on your diagram? [This is way cool when you see the pattern.]

6. Show that any triangulation of a 3-dimensional cube must have at least five simplices. Hint: what can you say about the volume of any simplex with an exterior face? [Recall that for a tetrahedron, volume = (1/3)(base)(height)]

Chapter 9

Minkowski's Theorem

Consider a convex set in the plane, and ask: how many integral points are there? This is another place where geometry and combinatorics intersect (literally).

9.1 Application: Seeing out of a Forest

You are at the center of a circular forest, 9 feet wide. Curiously, the trees are aligned in a very regular pattern— in fact, this forest can be placed on a grid (measured in feet) with you at the center in such a way that the center of the trees are at all the grid points with integer coordinates in the circle (except where you are situated at $(0,0)$). Each tree has a circular trunk exactly 6 inches (1/2 foot) wide. Can you see out of the forest?

9.2 Lattices

Definition 9.1. The *integer lattice* is the set $\mathbb{Z}^d$, and any point in $\mathbb{Z}^d$ is called a *lattice point.*

Theorem 9.2 (Minkowski). *Let C in $\mathbb{R}^d$ be a set that is convex, bounded, and symmetric about $\mathbf{0}$, such that $vol(C) > 2^d$. Then C contains at least one non-zero lattice point.*

The boundedness condition is not really needed (see exercises).

There are many nice applications in number theory, and because of this, Minkowski's theorem is part of a body of results known as the *geometry of numbers.*

For instance, we can use Minkowski's theorem to show how well an irrational can be approximated by rationals:

Theorem 9.3. *If $\alpha \in (0,1)$, and $N \in \mathbb{Z}_+$, then there exists $m, n \in \mathbb{Z}$ such that $n \leq N$ and $|\alpha - \frac{m}{n}| < \frac{1}{nN}$.*

Minkowski's result can also be generalized to more general lattices. Let $\Lambda = \Lambda(z_1, ..., z_d)$ in $\mathbb{R}^d$ be the set of all integer linear combinations of a linearly independent set $\{z_1, ..., z_d\}$ of vectors in $\mathbb{R}^d$. Let $\det(\Lambda)$ denote the determinant of the matrix Z whose columns are the z_i.

Theorem 9.4 (Minkowski's theorem for general lattices). *If Λ is a lattice in $\mathbb{R}^d$, and C in $\mathbb{R}^d$ is convex and symmetric about $\mathbf{0}$ such that vol$(C) > 2^d \cdot \det(\Lambda)$, then there exists at least one point in $C \cap (\Lambda - \mathbf{0})$.*

This can also provide a geometric proof of:

Theorem 9.5. *Every prime of the form $4k + 1$ can be written as the sum of two squares.*

In fact, similar ideas will prove Lagrange's theorem, that every natural number is the sum of (at most) 4 squares! Can you find such a proof?

9.3 Exercises

1. Answer the question in the capsule. Can you see your way out of the forest described in the capsule? Justify your answer.

2. Show that the boundedness condition in Minkowski's theorem is not really needed. Try to do this first without looking at my hints below (since maybe you can find a simpler proof than mine). Then if you give up (or your proof is too complicated), follow these steps:

   (a) In the plane, show that any convex set with positive area contains a little triangle T of positive area.

   (b) In the plane, show that any triangle T of positive area contains a little square C of positive area.

   (c) In the plane, show that the convex hull of such a square C with any point at distance D from the square must have area at least $sD/2\sqrt{2}$ where s is the sidelength of the square.

   (d) Conclude that any unbounded convex set in the plane cannot have finite area, and use this to show that Minkowski's thoerem holds for unbounded sets in the plane.

   (e) Then generalize your argument (a little hand-waving OK here) for Minkowski's theorem in $\mathbb{R}^d$.

Chapter 10

What are Ehrhart polynomials?

Many kinds of combinatorial problems can be expressed as the problem of counting lattice points in certain polytopes. For instance, if you wish to enumerate the number of 3×3 magic squares which have integer entries between 0 and 100, this can be phrased as finding integer lattice points in $\mathbb{R}^9$ which satisfy the linear equalities/inequalities arising from conditions on the row/column/diagonal sums and the bounds on the entries. This is a polytope.

A *lattice polytope* is a polytope that is the convex hull of lattice points.

In the plane, there is a particularly nice formula that relates the number of lattice points inside a lattice polygon to the area of the polygon. The polygon does not even have to be convex.

Theorem 10.1 (Pick's Theorem). *Let Q be a lattice polygon in $\mathbb{R}^2$ (not necessarily convex). Let n_i be the number of lattice points interior to Q and let n_b be the number of lattice points on the boundary of Q. Then the area of Q is*

$$n_i + \frac{1}{2}n_b - 1.$$

Pick's theorem thus gives the area of Q just by counting lattice points! The proof depends on Euler's formula ($V - E + F = 2$ for a planar graph) and the following lemma, whose proof is in the exercises.

Lemma 10.2. *If a triangle T is a lattice triangle and has no other lattice points inside, then the area of T is $1/2$.*

One may naturally wonder if there is any extension of Pick's theorem to a 3-polytope P? Well, there isn't quite a Pick-like formula, because it is easy to construct *bad* examples: arbitrarily large simplices with *no* interior lattice points, and the same number of boundary lattice points. Thus the volume of a polytope P cannot be simply given in terms of counting lattice points.

However, there is another possible direction in which to extend Pick's theorem, and that is to consider dilations of P. After all, even in the bad examples, it is still true that when you dilate P, the number of lattice points inside will be asympotically proportional to the volume of P.

This leads to the notion of *Ehrhart polynomials*. If P is lattice, then mP for an integer m is also a lattice polytope. Let

$$E_P(m) = \#(mP \cap \mathbb{Z}^d)$$

which counts the number of lattice points inside mP as a function of the dilation factor m. What kind of function is this? A theorem of Ehrhart says that, in fact, $E_P(m)$ is polynomial of degree d. Do some simple planar examples to convince yourself this is true.

What does this polynomial tell us about the original polytope?

Theorem 10.3 (Ehrhart, 1967). *The function $E_P(m)$ is a polynomial in m of degree d, with constant term 1 and leading coefficient equal to the volume of P.*

The proof of Ehrhart's theorem is deep and we shall not include a proof here. But it has some interesting consequences, which we shall explain in lecture. A related result is:

Theorem 10.4 (Ehrhart reciprocity). *The number of* interior *lattice points in mP is given by $(-1)^d E_p(-m)$.*

There are many things that are still not known about Ehrhart polynomials. For instance, what do the other coefficients mean? There are interpretations for the coeffients of m^d, m^{d-1} and the constant term, but all the others are unknown.

Another interesting question is: where are the roots of $E_p(m)$ and what do they mean? Only recently has it been possible to compute Ehrhart polynomials efficiently, and many of interesting conjectures of DeLoera et. al. were noted by doing computations in `LattE` and observing patterns. See the references for papers related to these questions.

10.1 Exercises

1. Prove Lemma 10.2 in the following steps:

   (a) If T has vertices p_0, p_1, p_2, let $v_1 = p_1 - p_0$ and $v_2 = p_2 - p_0$. What must be true about the entries of v_1, v_2 if T is a lattice triangle?

   (b) Show that v_1, v_2 form the sides of a parallelogram that has no lattice points inside it.

   (c) Argue that this parallelogram tiles the plane and integer combinations of v_1, v_2 must generate all lattice points.

   (d) Let A be the matrix with columns v_1, v_2. Use part (c) to show that there is some integer matrix Q such that AQ is the identity. What can you conclude about the determinant of A and the area of T?

Chapter 11

Combinatorial Fixed Point Theorems

11.1 Sperner's Lemma

Consider any triangulation T of a d-simplex Δ^d. Suppose that each vertex v of T is labelled by a label $\ell(v)$ chosen from $L = \{1, ..., d+1\}$. For any simplex σ in T, let $\ell(\sigma)$ be the set of labels of the vertices of σ. Call σ a *full cell* if all its labels are distinct, i.e, $\ell(\sigma) = L$.

For any point v in the simplex, express v in barycentric coordinates $v = (v_1, ..., v_{d+1})$. Let $\mathrm{carr}(v) = \{k : v_k > 0\}$. This set is called the *carrier* of v. Call a labelling ℓ a *Sperner-labelling* if for each vertex v, $\ell(v) \in \mathrm{carr}(v)$.

Theorem 11.1 (Sperner's Lemma). *Any Sperner-labelled triangulation of a simplex must contain an odd number of full cells.*

One of the primary reasons that Sperner's lemma is useful is that it is equivalent to the Brouwer Fixed Point Theorem. Thus it may be viewed as a combinatorial analogue. It is also equivalent to the following set intersection theorem of Knaster-Kuratowski-Mazurkiewicz, known as the "KKM lemma".

Theorem 11.2 (KKM lemma). *Suppose that the d-simplex Δ^d is covered by $d+1$ closed sets $C_1, ..., C_{d+1}$, such that for each point v in S, $v \in C_i$ for some $i \in carr(v)$. Then the intersection of all the sets is nonempty: $\cap_{i=1}^{d+1} C_i \neq \emptyset$.*

Compare this theorem to the statement of Helly's theorem.

Sperner's lemma has many proofs. See if you can find one for the case $d = 2$. In fact, there is a constructive proof of Sperner's lemma which yields a constructive method for finding fixed points of continuous functions!

There are also many neat applications of Sperner's lemma and the KKM theorem, especially in economics. And there are applications to cake-cutting and other "fair division" problems. We give a couple of examples from (Su, 1999).

Sperner's lemma can be used to prove the classical cake-cutting theorem, a connection that was first noted by Forest Simmons.

And there are many variants, such as the Polytopal Sperner Lemma (DeLoera-Peterson-Su 2002)and the KKM-Gale theorem.

11.2 Tucker's lemma

Consider any symmetric triangulation T of the d-sphere S^d, i.e., if σ is a simplex in T, then $-\sigma$ is also in T. Suppose that each vertex v of T is labelled by a label $\ell(v)$ chosen from $L = \{\pm 1, ..., \pm d\}$. For any edge e in T, let $\ell(e)$ be the set of labels of the endpoints of e. Call e a *oppositely signed edge* if its labels sum to zero.

A *Tucker-labelling* of a sphere is a labelling such that $\ell(-v) = -\ell(v)$ for all vertices v in T.

Theorem 11.3 (Tucker's Lemma). *Any Tucker-labelled triangulation a sphere must contain an oppositely signed edge.*

Tucker's lemma is the combinatorial analogue of the Borsuk-Ulam theorem from topology. In analogy with the Sperner/Brouwer relation, there is also a set-convering analogue in this setting, called the Lusternik-Schnirelman-Borsuk theorem [LSB theorem]. Can you figure out what it should say about closed sets covering the sphere?

Tucker's lemma also has an interesting application to cake-cutting.

11.3 Kneser colorings

Color the n-subsets of an $(2n+k)$-set. How many colors are needed to ensure that no 2 disjoint n-sets have the same color? Kneser showed in 1955 that $k+2$ colors were sufficient, and conjectured that $k+2$ were also necessary.

This combinatorial statement was proved in 1978 by Lovász using algebraic topology(!). That same year, Bárány gave a simpler proof using the LSB theorem, and a Gale theorem concering distribution of points on the sphere. In 2002, J. Greene (an undergraduate) gave an even simpler proof that avoided the use of the Gale theorem.

11.4 A combinatorial fixed point theorem for trees

Is there a combinatorial fixed point theorem for trees? Yes, and it can be used to prove a fixed point and KKM type result for trees, too!

The following theorem is well known from standard arguments in topology.

Theorem 11.4. *Let T be a tree. Every continuous $f : T \to T$ has a fixed point.*

The following result is a recent result of Berger (2002).

Theorem 11.5 (Tree-KKM Lemma). . *Suppose T is an n-leaf tree covered by n closed sets $C_1, ..., C_n$ such that for all leaves i and j, (1) leaf i is contained in set C_i and (2) $[i, j]$ is contained in $C_i \cup C_j$. Then all the sets C_i share a common point.*

Here is a combinatorial version, developed by Niedermaier-Su (2004).

Theorem 11.6 (Combinatorial Tree-KKM Lemma). . *Let T be an n-leaf tree which is segmented (triangulated into segments) such that nodes of the tree are vertices in the segmentation. Suppose that each vertex v has a non-empty set of labels $L(v)$ that are a subset of $\{1, ..., n\}$, such that for all leaves i and j,*

(1) $i \in L(\text{leaf } i)$, and

(2) if $v \in [i, j]$ then $L(v)$ contains either i or j (maybe both).

Then there is an edge (v_1, v_2) in the segmentation such that $L(v_1) \cup L(v_2) = \{1, ..., n\}$.

There are some cool applications of this theorem too.

What about other graphs?

11.5 Exercises

1. Prove the equivalence of Sperner's lemma, the KKM lemma, and the Brouwer Fixed Point Theorem

2. Show that Tucker's lemma is equivalent to the Borsuk-Ulam theorem.

An asterisk* denotes that the publication is available online at the ArXiv at http://front.math.ucdavis.edu/, or (likely) on someone's webpage (do a web search in Google, for instance).

Combinatorial Convexity

Proof of colorful Caratheodory: [1].

Text on convexity: [2].

Matousek's excellent book at the graduate level has a section on this: [3].

Helly's theorem

A nice section on Helly's theorem and applications: [4].

A new paper on a fractional Helly: [5].

Polytopes

A nice introduction is this chapter: [6]* in the book [7] which is a great overall introduction.

Ziegler's excellent book at the graduate level: [8].

Grübaum's classic text: [9].

Coxeter's classic text on regular polytopes: [10].

Kalai's proof that simple polytopes determined by their graph: [11]*.

Interesting history, nice pictures: [12].

McMullen's upper bound theorem: [13].

Triangulations

My paper on triangulations of cubes: [14]*

Upper bounds for triangulations of a planar point set: [15]*.

The space of triangulations is disconnected: [16]*.

The polytope of pointed pseudo-triangulations: [17]*.

Chamber complexes of polytopes: [18].

Riemmann hypothesis for triangulable manifolds: [19].

Combinatorial Fixed Point Theorems and Sperner's Lemma

Cake-cutting and rent-division applications: [20]*.

A polytopal generalization: [21]*.

Another analogue: [22].

Computation and applications to economics: [23].

A different generalization of Sperner's lemma: [24].

Tucker's lemma

The first constructive proof: [25]

Another constructive proof: [26]*.

Cake-cutting applications: [27]*.

Borsuk-Ulam implies Brouwer: [28]*.

Other papers on combinatorial fixed point theorems available at my webpage: `http://www.math.hmc.edu/ su/papers.html`

Kneser colorings

Joshua Greene's proof of the Kneser conjecture: [29].

Lovasz's famous proof of the Kneser conjecture: [30].

A survey by Björner: [31]*.

A paper by Ziegler on generalized Kneser colorings: [32]*.

Trees

A combinatorial fixed point theorems for trees: my paper will be available at `http://www.math.hmc.edu/~su/papers.html` later this summer.

The tree metric theorem: [33].

A biologist's book on phylogenies: [34].

Geometry of the space of phylogenetic trees: [35]*.

Tropical Geometry

A nice introduction is a chapter in [36].

Also: [37]*.

Tropical convexity: [38]*.

Rank of a tropical matrix: [39]*.

Tropical halfspaces: [40]*.

See other papers on the ArXiv at `http://front.math.ucdavis.edu/`.

Lattice point counting and Ehrhart polynomials

Ehrhart's paper: [41].

Coefficients and roots of Ehrhart Polynomials: [42]*.

An application of counting lattice points: [43]*.

A new book by Beck and Robins: see the webpage of Matthias Beck later this summer: `http://math.sfsu.edu/beck/papers.html`.

Unsolved problems

Book with lots of unsolved problems: [44].

Bibliography

[1] Imre Bárány. A generalization of Carathéodory's theorem. *Discrete Math.*, 40(2-3):141–152, 1982.

[2] Steven R. Lay. *Convex sets and their applications*. Robert E. Krieger Publishing Co. Inc., Malabar, FL, 1992. Revised reprint of the 1982 original.

[3] JiříMatoušek. *Lectures on discrete geometry*, volume 212 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2002.

[4] Roger Webster. *Convexity*. Oxford Science Publications. The Clarendon Press Oxford University Press, New York, 1994.

[5] Imre Bárány and JiříMatoušek. A fractional Helly theorem for convex lattice sets. *Adv. Math.*, 174(2):227–235, 2003.

[6] Martin Henk, Jürgen Richter-Gebert, and Günter M. Ziegler. Basic properties of convex polytopes. In *Handbook of discrete and computational geometry*, CRC Press Ser. Discrete Math. Appl., pages 243–270. CRC, Boca Raton, FL, 1997.

[7] Jacob E. Goodman and Joseph O'Rourke, editors. *Handbook of discrete and computational geometry*. CRC Press Series on Discrete Mathematics and its Applications. CRC Press, Boca Raton, FL, 1997.

[8] Günter M. Ziegler. *Lectures on polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995.

[9] Branko Grünbaum. *Convex polytopes*, volume 221 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 2003. Prepared and with a preface by Volker Kaibel, Victor Klee and Günter M. Ziegler.

[10] H. S. M. Coxeter. *Regular polytopes*. Dover Publications Inc., New York, third edition, 1973.

[11] Gil Kalai. A simple way to tell a simple polytope from its graph. *J. Combin. Theory Ser. A*, 49:381–383, 1988.

[12] Peter R. Cromwell. *Polyhedra.* Cambridge University Press, Cambridge, 1997. "One of the most charming chapters of geometry".

[13] Peter McMullen. The maximum numbers of faces of a convex polytope. *Mathematika*, 17:179–184, 1970.

[14] Adam Bliss and Francis Edward Su. Lower bounds for simplicial covers and triangulations of cubes.

[15] Francisco Santos and Raimund Seidel. A better upper bound on the number of triangulations of a planar point set. *J. Combin. Theory Ser. A*, 102(1):186–193, 2003.

[16] Francisco Santos. A point set whose space of triangulations is disconnected. *J. Amer. Math. Soc.*, 13(3):611–637 (electronic), 2000.

[17] Günter Rote, Francisco Santos, and Ileana Streinu. Expansive motions and the polytope of pointed pseudo-triangulations. In *Discrete and computational geometry*, volume 25 of *Algorithms Combin.*, pages 699–736. Springer, Berlin, 2003.

[18] Tatiana V. Alekseyevskaya. Combinatorial bases in systems of simplices and chambers. In *Proceedings of the 6th Conference on Formal Power Series and Algebraic Combinatorics (New Brunswick, NJ, 1994)*, volume 157, pages 15–37, 1996.

[19] K. S. Sarkaria. A "Riemann hypothesis" for triangulable manifolds. *Proc. Amer. Math. Soc.*, 90(2):325–326, 1984.

[20] Francis Edward Su. Rental harmony: Sperner's lemma in fair division. *Amer. Math. Monthly*, 106(10):930–942, 1999.

[21] Jesus A. De Loera, Elisha Peterson, and Francis Edward Su. A polytopal generalization of Sperner's lemma. *J. Combin. Theory Ser. A*, 100(1):1–26, 2002.

[22] Robert M. Freund. Combinatorial analogs of Brouwer's fixed-point theorem on a bounded polyhedron. *J. Combin. Theory Ser. B*, 47(2):192–219, 1989.

[23] Michael J. Todd. *The computation of fixed points and applications.* Springer-Verlag, Berlin, 1976. Lecture Notes in Economics and Mathematical Systems, Vol. 124.

[24] Duane W. DeTemple and Jack M. Robertson. An analog of Sperner's lemma. *Amer. Math. Monthly*, 83(6):465–467, 1976.

[25] Robert M. Freund and Michael J. Todd. A constructive proof of Tucker's combinatorial lemma. *J. Combin. Theory Ser. A*, 30(3):321–325, 1981.

[26] Timothy Prescott and Francis Edward Su. A Constructive Proof of Ky Fan's Generalization of Tucker's Lemma.

[27] Forest W. Simmons and Francis Edward Su. Consensus-halving via theorems of Borsuk-Ulam and Tucker. *Math. Social Sci.*, 45(1):15–25, 2003.

[28] Francis Edward Su. Borsuk-Ulam implies Brouwer: a direct construction. *Amer. Math. Monthly*, 104(9):855–859, 1997.

[29] Joshua E. Greene. A new short proof of Kneser's conjecture. *Amer. Math. Monthly*, 109(10):918–920, 2002.

[30] L. Lovász. Kneser's conjecture, chromatic number, and homotopy. *J. Combin. Theory Ser. A*, 25(3):319–324, 1978.

[31] A. Björner. Topological methods. In *Handbook of combinatorics, Vol. 1, 2*, pages 1819–1872. Elsevier, Amsterdam, 1995.

[32] Günter M. Ziegler. Generalized Kneser coloring theorems with combinatorial proofs. *Invent. Math.*, 147(3):671–691, 2002.

[33] Charles Semple and Mike Steel. *Phylogenetics*. Oxford Univ. Press, 2003.

[34] Joseph Felsenstein. *Inferring Phylogenies*. Sinauer Associates, Inc., 2003.

[35] Louis J. Billera, Susan P. Holmes, and Karen Vogtmann. Geometry of the space of phylogenetic trees. *Adv. in Appl. Math.*, 27(4):733–767, 2001.

[36] Bernd Sturmfels. *Solving systems of polynomial equations*, volume 97 of *CBMS Regional Conference Series in Mathematics*. Published for the Conference Board of the Mathematical Sciences, Washington, DC, 2002.

[37] Jürgen Richter-Gebert, Bernd Sturmfels, and Thorsten Theobald. First steps in tropical geometry.

[38] Mike Develin and Bernd Sturmfels. Tropical Convexity. AIM 2003-16.

[39] M. Develin, F. Santos, and B. Sturmfels. On the rank of a tropical matrix. AIM 2003-25.

[40] Michael Joswig. Tropical Halfspaces.

[41] E. Ehrhart. Sur un problème de géométrie diophantienne linéaire. II. Systèmes diophantiens linéaires. *J. Reine Angew. Math.*, 227:25–49, 1967.

[42] M. Beck, J. A. De Loera, M. Develin, J. Pfeifle, , and R. P. Stanley. Coefficients and Roots of Ehrhart Polynomials. AIM 2004-1.

[43] Matthias Beck, Moshe Cohen, Jessica Cuomo, and Paul Gribelyuk. The number of “magic” squares, cubes, and hypercubes. *Amer. Math. Monthly*, 110(8):707–717, 2003.

[44] Hallard T. Croft, Kenneth J. Falconer, and Richard K. Guy. *Unsolved problems in geometry*. Problem Books in Mathematics. Springer-Verlag, New York, 1994. Corrected reprint of the 1991 original. Unsolved Problems in Intuitive Mathematics, II.