Discrete Geometry

Austin Mohr

April 26, 2012

Problem 1

Theorem 1 (Linear Programming Duality)**.** *Suppose* $\mathbf{x}, \mathbf{y}, \mathbf{b}, \mathbf{c} \in \mathbb{R}^n$ *and* $A \in \mathbb{R}^{n \times n}$*,* $A\mathbf{x} \leq \mathbf{b}$*,* $\mathbf{x} \geq \mathbf{0}$*,* $A^\mathsf{T}\mathbf{y} \geq \mathbf{c}$*, and* $\mathbf{y} \geq \mathbf{0}$*. If* $\mathbf{x}$ *maximizes* $\mathbf{c}^\mathsf{T}\mathbf{x}$ *and* $\mathbf{y}$ *minimizes* $\mathbf{b}^\mathsf{T}\mathbf{y}$ *subject to these constraints, then*

$$\mathbf{c}^\mathsf{T}\mathbf{x} = \mathbf{y}^\mathsf{T}\mathbf{b},$$

where all inequalities are interpreted componentwise, i.e., $\mathbf{w} \leq \mathbf{z}$ *means that* $w_i \leq z_i$ *for all* i*.*

Proof. We show first that $\mathbf{c}^\mathsf{T}\mathbf{x} \leq \mathbf{y}^\mathsf{T}\mathbf{b}$. We have $A\mathbf{x} \leq \mathbf{b}$. Since $\mathbf{y} \geq 0$, it follows that $\mathbf{y}^\mathsf{T}A\mathbf{x} \leq \mathbf{y}^\mathsf{T}\mathbf{b}$. We also have $\mathbf{c} \leq A^\mathsf{T}\mathbf{y}$, which is equivalent to $\mathbf{c}^\mathsf{T} \leq \mathbf{y}^\mathsf{T}A$. Since $\mathbf{x} \geq 0$, we also have $\mathbf{c}^\mathsf{T}\mathbf{x} \leq \mathbf{y}^\mathsf{T}A\mathbf{x}$. Taken together, we have $\mathbf{c}^\mathsf{T}\mathbf{x} \leq \mathbf{y}^\mathsf{T}\mathbf{b}$, as desired.

To see the $\mathbf{c}^\mathsf{T}\mathbf{x} \geq \mathbf{y}^\mathsf{T}\mathbf{b}$, let

$$A' = \begin{bmatrix} A \\ -I \\ -\mathbf{c}^\mathsf{T} \end{bmatrix}$$

and

$$\mathbf{b}' = \begin{bmatrix} \mathbf{b} \\ 0 \\ -\mathbf{y}^\mathsf{T}\mathbf{b} \end{bmatrix}.$$

By (a version of) the Farkas Lemma, we may consider two cases.

In the first case, there exists a vector $\mathbf{z}$ with $A'\mathbf{z} \leq \mathbf{b}$. This implies that $A\mathbf{z} \leq \mathbf{b}$, $\mathbf{z} \geq 0$, and $\mathbf{y}^\mathsf{T}\mathbf{b} \leq \mathbf{c}^\mathsf{T}\mathbf{z}$. However, $\mathbf{c}^\mathsf{T}\mathbf{z} \leq \mathbf{c}^\mathsf{T}\mathbf{x}$, by definition of $\mathbf{x}$, and so $\mathbf{y}^\mathsf{T}\mathbf{b} \leq \mathbf{c}^\mathsf{T}\mathbf{x}$.

In the second case, there exists a vector

$$\mathbf{z} = [\mathbf{u}, \mathbf{v}, \lambda]$$

with $\mathbf{z} \geq 0$, $\mathbf{z}A' = 0$, and $\mathbf{z}\mathbf{b}' < 0$. Since $\mathbf{z}A' = 0$, we have

$$\begin{aligned} \mathbf{u}A &= \mathbf{v} + \lambda\mathbf{c} \\ &\geq \lambda\mathbf{c}. \end{aligned}$$

Since $\mathbf{z}\mathbf{b}' < 0$, we have $\mathbf{u}\mathbf{b} < \lambda\mathbf{y}^\mathsf{T}\mathbf{b}$.

Now, if $\lambda = 0$, $\mathbf{u}A \geq 0$ and $\mathbf{u}\mathbf{b} < 0$. So,

$$\begin{aligned} (\mathbf{u} + \mathbf{y})A &\geq \mathbf{y}A \\ &\geq \mathbf{c} \end{aligned}$$

and

$$(\mathbf{u} + \mathbf{y})\mathbf{b} < \mathbf{y}\mathbf{b},$$

which is a contradiction with the minimality of $\mathbf{y}$.

If $\lambda \neq 0$, then we may scale to $\mathbf{z}' = \left[\frac{1}{\lambda}\mathbf{u}, \frac{1}{\lambda}\mathbf{v}, 1\right]$. Proceeding as before, we get $\mathbf{u}A \geq \mathbf{c}$ and $\mathbf{u}\mathbf{b} < \mathbf{y}^\mathsf{T}\mathbf{b}$, which is a contradiction with the minimality of $\mathbf{y}$.

Therefore, the second case does not occur, and so we conclude that $\mathbf{c}^\mathsf{T}\mathbf{x} = \mathbf{y}^\mathsf{T}\mathbf{b}$. □

Problem 2

Given two sets A, B, their *Minkowski sum* is defined by

$$A + B = \{a + b \,|\, a \in A \wedge b \in B\}.$$

For $k \geq 3$, a *convex k-gon* in the plane is defined to be any set of the form $\text{conv}(X)$, where $|X| = k$ and X is in convex position.

Proposition 2. *If A and B are a convex n-gon and a convex m-gon, respectively, and no edge of A is parallel to an edge of B, then $A + B$ is a convex $(n+m)$-gon. (The "edges" of a convex k-gon $conv(X)$ are the vectors $\{x_1 - x_2, x_2 - x_3, \ldots, x_{k-1} - x_k, x_k - x_1\}$, where the elements of X have been ordered cyclically as $\{x_1, x_2, \ldots, x_k\}$.)*

Proof. Let X denote the set of vertices of A and Y the set of vertices of B. We show first that $\text{conv}(X) + \text{conv}(Y) = \text{conv}(X + Y)$.

If $z \in \text{conv}(X) + \text{conv}(Y)$, then

$$z = \sum_i a_i x_i + \sum_j b_j y_j$$

where $x_i \in X$, $y_j \in Y$, and $a_i, b_j \in \mathbb{R}$ with $\sum_i a_i = \sum_j b_j = 1$. It follows that

$$\begin{aligned} z &= \sum_i a_i \left(\sum_j b_j \right) x_i + \sum_j b_j \left(\sum_i a_i \right) y_j \\ &= \sum_{i,j} a_i b_j x_i + \sum_{i,j} a_i b_j y_j \\ &= \sum_{i,j} a_i b_j (x_i + y_j), \end{aligned}$$

and so $z \in \text{conv}(X + Y)$.

If $z \in conv(X + Y)$, then we similarly have

$$\begin{aligned} z &= \sum_i a_i (x_i + y_i) \\ &= \sum_i a_i x_i + \sum_i a_i y_i, \end{aligned}$$

and so $z \in \text{conv}(X) + \text{conv}(Y)$.

It remains to show that $\text{conv}(X + Y)$ is an $(m+n)$-gon. To that end, choose some orientation (say, clockwise) of the edges of each polygon. Considering the edges as vectors, each corresponds to some angle formed with the origin. Order the edges so that their corresponding angles are in increasing order (the order is strict, since no two edges have the same slope). An edge with slope m will be extreme in the direction of the vector with slope $\frac{1}{m}$ (with the convention that $\frac{1}{0} = \infty$ and $\frac{1}{\infty} = 0$). Thus, each edge of A and B appears precisely once in the sum $A + B$, and so give rise to $m + n$ distinct edges. $\square$

Problem 3

For a metric space (X, d), define the "unit distance graph" G to have $V(G) = X$, with $\{x, y\} \in E(G)$ iff $d(x, y) = 1$.

Lemma 3. *Let $p = \frac{n_1}{d_1}$ and $q = \frac{n_2}{d_2}$ be rational numbers in reduced form. If (p, q) is unit distance from the origin, then d_1 and d_2 are both odd and exactly one of n_1 and n_2 is odd.*

Proof. Since p and q are unit distance from the origin,

$$p^2 + q^2 = 1.$$

Rearranging terms gives

$$(n_1 d_2)^2 + (n_2 d_1)^2 = (d_1 d_2)^2.$$

We proceed by considering cases.

Suppose exactly one of d_1 and d_2 is even. Without loss of generality, let d_1 be even. Since p is in reduced form, n_1 is odd. Thus, $(n_1 d_2)^2$ is odd, $(n_2 d_1)^2$ is even, and $(d_1 d_2)^2$ is even, which is impossible.

Suppose both of d_1 and d_2 are even (and so both n_1 and n_2 are odd). Let $d_1 = 2^j s$ and $d_2 = 2^k t$ with s and t odd (without loss of generality, $j \leq k$). Thus,

$$(n_1 2^k t)^2 + (n_2 2^j s)^2 = (2^j s 2^k t)^2.$$

Rearranging terms gives

$$2^{2(k-j)} n_1^2 t^2 + n_2^2 s^2 = 2^{2k} s^2 t^2.$$

Thus, it must be that $j = k$ (or else the lefthand side is odd and the righthand side even). We now have

$$n_1^2 t^2 + n_2^2 s^2 = 2^{2k} s^2 t^2.$$

Now, the righthand side is divisible by 4, whereas the lefthand side is not (since all of n_1, n_2, s, and t are odd). We conclude that d_1 and d_2 cannot both be even.

We have now that both d_1 and d_2 are odd, and so $(d_1 d_2)^2$ is odd. Since $(n_1 d_2)^2 + (n_2 d_1)^2 = (d_1 d_2)^2$, it must be that exactly one of n_1 and n_2 is odd (and so the other even), as desired. □

Proposition 4. *The unit distance graph of $\mathbb{Q}^2$ with ordinary Euclidean distance is bipartite.*

Proof. Define a coloring f on $\mathbb{Q}^2$ via

$$f\left(\frac{a}{b}, \frac{c}{d}\right) = \begin{cases} \text{even} & \text{if } a + c \text{ is even} \\ \text{odd} & \text{if } a + c \text{ is odd,} \end{cases}$$

where $\frac{a}{b}$ and $\frac{c}{d}$ are written in reduced form. We show that f is a proper 2-coloring of the unit distance graph of $\mathbb{Q}^2$.

Let (p_1, q_1) and (p_2, q_2) be elements of $\mathbb{Q}^2$ at unit distance. Hence, $(0, 0)$ and $(p_2 - p_1, q_2 - q_1)$ are also at unit distance. We see that $(0, 0)$ is colored even. To determine the color of $(p_2 - p_1, q_2 - q_1)$, let $p_i = \frac{a_i}{b_i}$ and $q_i = \frac{c_i}{d_i}$. Thus,

$$p_2 - p_1 = \frac{a_2 b_1 - a_1 b_2}{b_1 b_2}$$

and

$$q_2 - q_1 = \frac{c_2 d_1 - c_1 d_2}{d_1 d_2}.$$

Since $(p_2 - p_1, q_2 - q_1)$ is unit distance from the origin, the lemma gives us immediately that exactly one of $a_2 b_1 - a_1 b_2$ and $c_2 d_1 - c_1 d_2$ is odd and both of $b_1 b_2$ and $d_1 d_2$ are odd. Without loss of generality, let it be that $a_2 b_1 - a_1 b_2$ is the sole even number among these four. It follows that both b_1 and b_2 are odd, and so both a_1 and a_2 are even. Thus, (p_1, p_2) is colored even. Similarly, both d_1 and d_2 are odd, and so exactly one of $c_2 d_1$ and $c_1 d_2$ is odd. Thus, (q_1, q_2) is colored odd. Therefore, no elements of $\mathbb{Q}^2$ at unit distance receive the same color, as desired. □

Problem 4

Given a convex set $S \subset \mathbb{R}^d$, $x \in S$ is called an *extreme point* of S if it does not lie in any open line segment joining two points of S. Denote the set of extreme points of S by $\text{Ext}(S)$.

Proposition 5. *If S is convex and compact, then $S = conv(Ext(S))$.*

Proof. We show first that any point p on the boundary of S belongs to $\text{conv}(\text{Ext}(S))$. We proceed by induction on the dimension d of the ambient space $\mathbb{R}^d$.

For $d = 1$, the ambient space is the real line, and so S is a closed interval $[a, b]$. As p is a point on the boundary of S, $p = a$ or $p = b$. Both a and b are extreme points, however, as any open line segment around either point intersects S^c. Thus, p can be written as a (trivial) convex combination of points belonging to $\text{Ext}(S)$.

For $d > 1$, choose any face F containing p. We see that F is convex and compact (otherwise, S fails to be convex and compact) and has dimension $d - 1$. By the inductive hypothesis, $p \in \text{conv}(Ext(F))$. Now, any open line segment containing a point of F must lie entirely in F, as F is on the boundary of S. Thus, $\text{Ext}(F) \subset \text{Ext}(S)$, and so $p \in \text{conv}(\text{Ext}(S))$.

Let now p be an arbitrary point belonging to $S \subset \mathbb{R}^d$ for any value of d. If p lies on the boundary of S, then we are finished by the previous remarks. Now, assume p lies on the interior of S. There exists an open (d-dimensional) ball containing p, and so also an open line segment containing p. Choose any such line segment L and extend it in both directions as far as possible while maintaining the property that $L \subset S$. Note here that we use the convexity of S to ensure that $L \subset S$ and the boundedness of S to ensure that L is indeed a line segment (and not a line). Now, since S is also closed, the topological closure $\overline{L}$ of L is also contained in S. Since $\overline{L}$ was taken to be as large as possible in this direction, it follows that the endpoints a and b of $\overline{L}$ lie on the boundary of S. Thus, p can be expressed as a convex combination of a and b. Since a and b lie on the boundary of S, $a, b \in \text{conv}(\text{Ext}(S))$, and so $p \in \text{conv}(\text{Ext}(S))$, thus completing the proof. □

Problem 5

Given a convex set $S \subset \mathbb{R}^d$, a hyperplane H is said to *support* S if one of its two closed half-spaces contains S. A point $x \in S$ is called an *exposed point* of S if there exists a supporting hyperplane H so that $H \cap S = \{x\}$. Denote the set of exposed points of S by $\text{Exp}(S)$.

Proposition 6. *For S compact and convex, the (topological) closure of $Exp(S)$ contains $Ext(S)$.*

Proof. Let $x \in \text{Ext}(S)$ and suppose, for contradiction, $x \notin \overline{\text{Exp}(S)}$. That is, there exists an open ball about x having empty intersection with $\text{Exp}(S)$.

We claim that $x \notin \text{conv}(\text{Exp}(S))$. Suppose, for contradiction, x could be written as a minimal convex combination $\sum_{i=1}^k a_i x_i$ with each $x_i \in \text{Exp}(S)$. Note that since $x \notin \text{Exp}(S)$, $k \geq 2$. Hence, we may partition $[k]$ into sets I and J and define $u = \sum_{i \in I} a_i x_i$ and $v = \sum_{j \in J} a_j x_j$ (so that $x = u + v$). Define also $A = \sum_{i \in I} a_i$ and $B = \sum_{j \in J} a_j$ and let $u' = \sum_{i \in I} \frac{a_i}{A} x_i$ and $v' = \sum_{j \in J} \frac{a_j}{B} x_j$. Now, $u', v' \in \text{conv}(\text{Exp}(S))$. We also have $x = Au' + Bv'$, and so x can be written as a convex combination of just two elements of $\text{Exp}(S)$. Since x is not itself a member of $\text{Exp}(S)$ and the convex combination was chosen minimally, $\{tu' + (1-t)v' \mid t \in (0, 1)\}$ is an open line segment containing x, which is a contradiction with the fact that $x \in \text{Ext}(S)$.

Now, since $\{x\}$ and $\text{conv}(\text{Exp}(S))$ are disjoint, closed, convex sets, they can be strictly separated by a hyperplane H. Let H_x denote the closed halfplane containing x. We show next that H_x contains a point of $\text{Exp}(S)$, which is a contradiction with our choice of hyperplane. To see this, construct a sphere T of sufficient radius such that x is on the boundary of T and $T \cap H \cap S = \emptyset$. Now, increase the radius of T as large as possible while maintaining $T \cap S \neq \emptyset$. Call the resulting sphere T'. Any point $y \in H_x$ on the boundary of T' will have a supporting hyperplane intersecting S precisely at y. Hence, $y \in \text{Exp}(S) \cap H_x$.

All told, the assumption $x \notin \overline{\text{Exp}(S)}$ results in the impossible situation that there is a hyperplane strictly separating $y \in \text{Exp}(S)$ from $\text{conv}(\text{Exp}(S))$, and so we must conclude that $x \in \overline{\text{Exp}(S)}$, as desired. □

Problem 6

A map $\overline{\cdot} : \mathcal{P}(\mathbb{R}^d) \to \mathcal{P}(\mathbb{R}^d)$ is called a "closure operator" if (a) $\forall X(X \subseteq \overline{X})$, (b) $\forall X \forall Y((X \subseteq Y) \Rightarrow (\overline{X} \subseteq \overline{Y}))$, and (c) $\forall X(\overline{\overline{X}} = \overline{X})$. ("$\mathcal{P}$" denotes the power set operation.)

Proposition 7. *The convex hull operation conv$(\cdot)$ is a closure operator.*

Proof. Recall that the convex hull of a set X is defined to be the intersection of all convex sets containing X. Let $\mathcal{C}$ be the class of all convex sets containing X, so that $\operatorname{conv}(X) = \bigcap \mathcal{C}$.

By definition, $X \subset C$ for all $C \in \mathcal{C}$, and so $X \subset \operatorname{conv}(X)$.

Let now Y be another set of points with $X \subset Y$ and let $\mathcal{D}$ be the collection of all convex sets containing Y. It follows that $X \subset Y \subset D$ for all $D \in \mathcal{D}$. Thus, $\mathcal{D} \subset \mathcal{C}$, and so $\operatorname{conv}(X) \subset \operatorname{conv}(Y)$.

For the remaining property, observe that $X \subset \operatorname{conv}(X)$ implies $\operatorname{conv}(X) \subset \operatorname{conv}(\operatorname{conv}(X))$. At the same time, $\operatorname{conv}(X)$ is itself a convex set containing $\operatorname{conv}(X)$, and thus $\operatorname{conv}(\operatorname{conv}(X)) \subset \operatorname{conv}(X)$. □

Problem 1

Define the *separation* of a point set X in a metric space with distance function d to be

$$\inf_{\substack{x,y \in X \\ x \neq y}} d(x,y).$$

For $X \subset \mathbb{R}^n$, define

$$\rho(X) = \sup_{\mathbf{x} \in \mathbb{R}^n} \inf_{\mathbf{y} \in X} \|\mathbf{x} - \mathbf{y}\|.$$

Proposition 8. *The minimum possible value of $\rho(X)$ over all sets $X \subset \mathbb{R}^2$ with separation 1 is $\frac{\sqrt{3}}{3}$. The class of all sets achieving this bound is precisely the class of (the vertices of) tilings by equilateral triangles.*

Proof. (Sketch) "Clearly", a set X minimizing $\rho(X)$ is a lattice with a basis. (Either every cluster of points admits the same value of ρ or it does not. If so, then there is no need for variation among the clusters (every cluster looks like every other). If not, then this configuration is not optimal.)

Consider the parallelogram *abcd* defined by basis vectors $\overrightarrow{v_1}$ and $\overrightarrow{v_2}$. The point $x \in \mathbb{R}^n$ obtaining the minimum value of ρ among points inside this parallelogram must be equidistant from opposing corners, and so lies on the intersection of the diagonals of the parallelogram. As a result, we see that $\overrightarrow{v_1}$ and $\overrightarrow{v_2}$ are exactly unit length. It remains to determine the angle between them.

For angles between 0 and $\frac{\pi}{2}$, a smaller angle gives a shorter diagonal (say, the one between b and d), so we should take the smallest angle not violating the separation of X. Some basic trigonometry shows that this is achieved at precisely $\frac{\pi}{3}$, and so the resulting lattice is a tiling by equilateral triangles (all are isomorphic up to translation).

For angles between $\frac{\pi}{2}$ and π, a larger angle gives a shorter diagonal (the one between a and c), so we should take the largest angle not violating the separation of X. Some basic trigonometry shows that this is achieved at precisely $\frac{2\pi}{3}$, and so the resulting lattice is a tiling by equilateral triangles (all are isomorphic up to translation). □

Problem 2

The diameter $\operatorname{diam}(X)$ of a set X in a metric space with distance function d is the number $\sup_{x,y \in X} d(x,y)$.

Proposition 9. *Every set $X \subset \mathbb{R}^2$ of diameter at most 1 is contained in some ball of radius $\frac{\sqrt{3}}{3}$, and this is smallest possible.*

Proof. Suppose first that $|X| = 3$. We may assume, without loss of generality, that the points of X form an equilateral triangle with side length 1, as any other configuration is contained in this one. Some basic trigonometry shows that this triangle can be inscribed in a circle of radius $\frac{\sqrt{3}}{3}$. Hence, if this radius works in general, it must be smallest possible.

Consider now an arbitrary diameter 1 subset X of the real plane. Let $\mathcal{C}$ denote the collection of all closed disks of radius $\frac{\sqrt{3}}{3}$ whose center is a point of X. Observe that any three circles of $\mathcal{C}$ have nonempty intersection. To see this, let $C_1, C_2, C_3 \in \mathcal{C}$ and consider two cases. If some $C_i \cap C_j =$, then their centers are at distance greater than 1, which is a contradiction with the fact that $\text{diam} X = 1$. Otherwise, there is some point x in the interior of the triangle induced by the centers of the three disks such that $x \notin C_1 \cap C_2 \cap C_3$. Thus, x has distance strictly greater than $\frac{\sqrt{3}}{3}$ from the center of at least one of the C_i. Using the same basic trigonometry as before, we conclude that on of the sides of the triangle is greater than 1, which again contradicts the fact that $\text{diam} X = 1$.

Since any three disks in $\mathcal{C}$ have nonempty intersection, the infinite version of Helly's Theorem guarantees that the intersection of *all* disks in $\mathcal{C}$ is nonempty. Let y be a point belonging to this intersection and consider a closed disk C of radius $\frac{\sqrt{3}}{3}$ centered at y. Since y is contained in every C_i, we know that y is at most distance $\frac{\sqrt{3}}{3}$ from any point of X. Therefore, every point of X is contained in C, thus completing the proof. □

Problem 3

Proposition 10. *Let $A \subset \mathbb{R}^2$ be a compact convex set of area α. If z is chosen uniformly at random from $[0, 1)^2$, then*

$$\mathbb{E}[|\mathbb{Z}^2 \cap (A + z)|] = \alpha,$$

where $\mathbb{E}$ is the expectation operator.

Proof. Suppose first that A is a square of side length $\beta < 1$. Since $x \in [0, 1)^2$, there are at most four possible points of $\mathbb{Z}^2$ that can belong to $A + x$, and these points form a 1×1 square in $\mathbb{Z}^2$. Denote these points by p_{00} (lower left), p_{01} (upper left), p_{10} (lower right), and p_{11} (upper right). For each of the aforementioned points, let $p_{ij} = (x_{ij}, y_{ij})$. For ease of notation, let $L_v = \overline{p_{00}p_{01}}$ and $L_h = \overline{p_{00}p_{10}}$. Finally, let $z = (x, y)$. We consider the four possible cases for intersections of L_v and L_h with A.

<u>Case 1</u>: If neither L_v nor L_h intersect A, then A is on the interior of the square formed by the p's. It follows that only p_{11} can belong to $A + z$. Moreover, for every point q of A, there is a choice of z so that $q + z = p_{11}$. Thus, $\mathbb{E}[|\mathbb{Z}^2 \cap (A + z)|] = \alpha$.

<u>Case 2</u>: If L_v intersects A but L_h does not, then A is split vertically into two rectangles of width x_1 and x_2. In this case, only p_{01} and p_{11} can belong to $A + z$. Now, both p_{01} and p_{11} can take on any vertical position in A for an appropriate choice of z. Thus, we have

$$\begin{aligned}
\mathbb{E}[|\mathbb{Z}^2 \cap (A + z)|] &= \beta P(x \in [0, x_1]) + \beta P(x \in [1 - x_2, 1]) \\
&= \beta x_1 + \beta x_2 \\
&= \beta(x_1 + x_2) \\
&= \beta^2 \\
&= \alpha.
\end{aligned}$$

<u>Case 3</u>: If L_h intersects A but L_v does not, then A is split horizontally into two rectangles of height y_1 and y_2. In this case, only p_{10} and p_{11} can belong to $A + z$. Now, both p_{10} and p_{11} can take on any horizontal position in A for an appropriate choice of z. Thus, we have

$$\begin{aligned}
\mathbb{E}[|\mathbb{Z}^2 \cap (A + z)|] &= \beta P(y \in [0, y_1]) + \beta P(y \in [1 - y_2, 1]) \\
&= \beta y_1 + \beta y_2
\end{aligned}$$

$$
\begin{aligned}
&= \beta(y_1 + y_2) \\
&= \beta^2 \\
&= \alpha.
\end{aligned}
$$

Case 4: If both L_v and L_h intersect A, then A is split vertically into two rectangles of width x_1 and x_2 and horizontally into two rectangles of height y_1 and y_2. It follows that,

$$
\begin{aligned}
\mathbb{E}[|\mathbb{Z}^2 \cap (A+z)|] &= P(p_{00} \in A+z) + P(p_{01} \in A+z) + P(p_{10} \in A+z) + P(p_{11} \in A+z) \\
&= P(x \in [0, x_1] \wedge y \in [0, y_1]) + P(x \in [0, x_1] \wedge y \in [1-y_2, 1]) \\
&\quad + P(x \in [1-x_2, 1] \wedge y \in [0, y_1]) + P(x \in [1-x_2, 1] \wedge y \in [1-y_2, 1]) \\
&= x_1 y_1 + x_1 y_2 + x_2 y_1 + x_2 y_2 \\
&= x_1(y_1 + y_2) + x_2(y_1 + y_2) \\
&= x_1 \beta + x_2 \beta \\
&= (x_1 + x_2)\beta \\
&= \beta^2 \\
&= \alpha.
\end{aligned}
$$

Taking these four cases together, we see that the proposition holds when A is square of side length $\beta < 1$.

For general convex, compact sets A, we may approximate collections $\mathcal{S}_n$ of internally disjoint squares of side length less than 1 so that

- $\bigcup_{S \in \mathcal{S}_n} S \subset A$ for each n, and
- the area of A differs from the combined area of the squares in $\mathcal{S}_n$ by less than $\frac{1}{n}$.

It follows that

$$
\begin{aligned}
\mathbb{E}[|\mathbb{Z}^2 \cap (A+z)|] &= \lim_{n \to \infty} \mathbb{E}[|\mathbb{Z}^2 \cap (\mathcal{S}_n + z)|] \\
&= \lim_{n \to \infty} \sum_{S \in \mathcal{S}} \mathbb{E}[|\mathbb{Z}^2 \cap (S + z)|] \\
&= \lim_{n \to \infty} \alpha - \frac{1}{n} \\
&= \alpha.
\end{aligned}
$$

□

Problem 4

Proposition 11. *If every two distinct points of $X \subset \mathbb{R}^d$ are distance 1 apart, then $|X| \leq d+1$.*

Proof. We proceed by induction on d. For the case $d = 1$, there are clearly at most 2 points that are pairwise distance 1 apart. For general d, begin by choosing a set X containing d points that are pairwise distance 1 apart (by induction, this can always be done). Let H be the hyperplane containing X. We show that there are precisely two points in $\mathbb{R}^d$ that are distance 1 from all points of X. Denote such a point by p. Since H is $d-1$-dimensional, it cannot be that p lies in H, as this would violate the inductive hypothesis. We also know that the projection of p onto H must be equidisant from all points of X (else p is not equidistant from all points of X). Thus, there is a unique line orthogonal to H on which p must lie. Now, p together with any two points of X form an equilateral triangle. Hence, p must lie on L at distance $\frac{\sqrt{3}}{2}$ from H, giving two possible choices for p. We may add either choice to X and maintain the property of all points having pairwise distance 1, but not both (the two choices are at distance $\sqrt{3}$). Therefore, there are at most $d+1$ points in $\mathbb{R}^d$ that are all pairwise distance 1 apart, as desired. □

Problem 5

Proposition 12. *Let $S \subset \mathbb{R}^2$ and $n = |S| < \infty$. There are $O(n^{7/3})$ triples of points in S forming the vertices of a triangle having area 1.*

Proof. Fix any point $p \in S$. For any other point $q \in S$, form the segment $\overline{pq}$. Define also the lines L and L' that are parallel to $\overline{pq}$, each at distance $\frac{2}{|pq|}$ from $\overline{pq}$. If p and q are to be vertices of a triangle with area 1, then the third vertex must lie somewhere on L or L' (since the area of the triangle is $\frac{1}{2}|pq| \cdot \frac{2}{|pq|}$). Ranging across choices of q, there are $2(n-1)$ such lines. The maximum number of triangles of area 1 with p as a vertex can thus be realized as the maximum number of point-line incidences with $n-1$ points and $2(n-1)$ lines. (Actually, each triangles is counted twice, but this does not affect the order of the count.) By the Szemerédi-Trotter Theorem, there are $O(n^{\frac{4}{3}})$ such incidences. Ranging over all choices for the initial point p, we thus have $O(n^{\frac{7}{3}})$ triangles having area 1. (Similar to the previous count, every triangle will be counted three times as we range over choies of p, but this does not affect the order.) ☐

Problem 1

Proposition 13. *For $n \geq d$ and $d \geq 4$, the cyclic polytope with n vertices in $\mathbb{R}^d$ has the property that every two vertices are on an edge together.*

Proof. Let v_1 and v_2 be any two vertices of the d-dimensional cyclic polytope and let t_1 and t_2 be the corresponding moment curve parameters. Consider the hyperplane H defined by

$$a_0 + a_1x_1 + a_2x_2 + a_3x_3 + a_4x_4 = 0$$

where each a_i is the coefficient of t^i in the polynomial $p(t) = (t-t_1)^2(t-t_2)^2$. As t_1 is a root of p, we have

$$a_0 + a_1t_1 + a_2t_1^2 + a_3t_1^3 + a_4t_1^4 = 0,$$

which is precisely to say that v_1 lies on the hyperplane H. Similarly, v_2 lies on H.

Let now w be any other vertex of the cyclic polytope and let s be its moment curve parameter. We have

$$\begin{aligned} p(s) &= (s-t_1)^2(s-t_2)^2 \\ &> 0. \end{aligned}$$

That is,

$$a_0 + a_1s + a_2s^2 + a_3s^3 + a_4s^4 > 0,$$

and so w lies on the positive side of H. Therefore, H supports an edge between v_1 and v_2, as desired. ☐

Problem 2

Proposition 14. *If $\{\mathbf{v}_1, \ldots, \mathbf{v}_k\} \subseteq vert(P)$ is a set of vertices of P, then $F = \bigvee_{j=1}^{k}\{\mathbf{v}_j\}$ if and only if $k^{-1}\sum_{j=1}^{k}\mathbf{v}_k \in relint(F)$, where $F \in L(P)$.*

Proof. Let x denote the convex combination $k^{-1}\sum_{j=1}^{k}\mathbf{v}_k$.

($\Rightarrow$) Since F is convex and contains $\{\mathbf{v}_1, \ldots, \mathbf{v}_k\}$, we know immediately that $x \in F$. It remains to show that x is not on the boundary of F. To see this, consider any supporting hyperplane H intersecting P precisely in a subface F' of F. Since F is the *smallest* face of P containing $\{\mathbf{v}_1, \ldots, \mathbf{v}_k\}$, at least one of the v_i lies off of H. Moreover, all such vertices lie in a common closed halfspace induced by H, as H supports P. Thus, x does not lie on H, and so does not lie on F'. As this holds for any subface of F, we conclude that x lies on no subface of F, and so lies in relint(F).

($\Leftarrow$) Let H_F be a supporting hyperplane intersecting P precisely in F. All vertices in the set $\{\mathbf{v}_1, \ldots, \mathbf{v}_k\}$ lie in a common closed halfspace induce by H_F, since H_F supports P. Reasoning as before, the fact that $x \in \operatorname{relint}(F)$ implies that $\{\mathbf{v}_1, \ldots, \mathbf{v}_k\}$ is contained in F.

To conclude that $F = \bigvee_{j=1}^{k} \{\mathbf{v}_j\}$, it remains to show that no subface of F contains $\{\mathbf{v}_1, \ldots, \mathbf{v}_k\}$. Suppose, for the purpose of contradiction, that $\{\mathbf{v}_1, \ldots, \mathbf{v}_k\}$ is contained in a subface F' of F. Let $H_{F'}$ be a supporting hyperplane intersecting P precisely in F'. Thus, $\{\mathbf{v}_1, \ldots, \mathbf{v}_k\} \subset F'$ is contained in $H_{F'}$, and so x is contained in $H_{F'}$, which is contrary to the fact that $x \in \operatorname{relint}(F)$. Hence, no subface of F contains the set $\{\mathbf{v}_1, \ldots, \mathbf{v}_k\}$, as desired. □

Problem 3

Let $(f_0^{(d)}, \ldots, f_d^{(d)})$ denote the f-vector of the d-dimensional hypercube $[0,1]^d$. Define

$$q(x, y) = \sum_{d=0}^{\infty} \sum_{j=0}^{d} f_j^{(d)} x^j y^d.$$

Proposition 15. *For $q(x, y)$ as defined above, we have $q(x, y) = 1/(1 - 2y - xy)$.*

Proof. We first obtain a closed form for $f_j^{(d)}$. A j-dimensional face of a d-dimensional hypercube is expressed using j "free" coordinates and $d-j$ "fixed" coordinated. The free coordinates range over all possible values for that coordinate (namely, 0 and 1). There are $\binom{d}{j}$ ways to select the free coordinates. The fixed coordinates are set to either 0 or 1, which can be accomplished in 2^{d-j} ways. Thus, $f_j^{(d)} = \binom{d}{j} 2^{d-j}$.

We have now

$$\begin{aligned}
q(x, y) &= \sum_{d=0}^{\infty} \sum_{j=0}^{d} f_j^{(d)} x^j y^d \\
&= \sum_{d=0}^{\infty} \sum_{j=0}^{d} \binom{d}{j} 2^{d-j} x^j y^d \\
&= \sum_{d=0}^{\infty} (2 + x)^d y^d && \text{(by the binomial theorem)} \\
&= \sum_{d=0}^{\infty} ((2 + x) y)^d \\
&= \frac{1}{1 - (2 + x) y} && \text{(sum of a geometric series)} \\
&= \frac{1}{1 - 2y - xy}.
\end{aligned}$$

□

Problem 4

Let $\Lambda \subset \mathbb{R}^n$ be a lattice, $S \subset \Lambda$ with $|S| < \infty$, $P = \operatorname{conv}(S)$, and $d = \dim(P)$. Define

$$L_P(t) = |\Lambda \cap tP|,$$

where, for any set $U \subset \mathbb{R}^n$, $tU = \{t\mathbf{u} | \mathbf{u} \in U\}$. It is a theorem of Ehrhart that there exist $a_0, \ldots, a_d \in \mathbb{Q}$ so that $L_P(t) = \sum_{j=0}^{d} a_j t^j$ for all $t \in \mathbb{Z}$.

Proposition 16. *Using the notation above,*

1. $a_0 = 1$;
2. $L_{[0,1]^n}(t) = (t+1)^n$, *where* $\Lambda = \mathbb{Z}^n$.
3. $a_d = vol_d(P)/\det(\Lambda)$ *(vol_d is the d-dimensional Lebesgue measure);*

Proof. For part 1, observe that $L_P(0) = a_0$ by Ehrhart's formula. At the same time

$$\begin{aligned} L_P(0) &= |\Lambda \cap (0 \cdot P)| \\ &= |\Lambda \cap \{0\}| \\ &= |\{0\}| \\ &= 1. \end{aligned}$$

For part 2, we proceed by induction on n. For $n = 1$, we have

$$\begin{aligned} L_{[0,1]}(t) &= |\Lambda \cap (t \cdot [0,1])| \\ &= |\Lambda \cap [0,t]| \\ &= |\{0,1,\dots,t\}| && \text{(here, } \Lambda = \mathbb{Z}^n) \\ &= t+1. \end{aligned}$$

Suppose now the result holds for $n \leq k$. We have

$$\begin{aligned} L_{[0,1]^{k+1}}(t) &= |\Lambda \cap (t \cdot [0,1]^{k+1})| \\ &= |\Lambda \cap [0,t]^{k+1}| \\ &= |\Lambda \cap [0,t]^k| \cdot |\Lambda \cap [0,t]| \\ &= (t+1)^k \cdot (t+1) \\ &= (t+1)^{k+1}. \end{aligned}$$

For part 3, we begin as in the proof of Minkowski's theorem. That is, we define a linear bijection f such that $f(\Lambda) = \mathbb{Z}^d$. Since P is convex, we have $\text{vol}_d(P) = \det(\Lambda)\text{vol}_d(f(P))$. Letting P' denote $f(P)$, the problem of showing $a_d = \frac{\text{vol}_d(P)}{\det(\Lambda)}$ for any lattice Λ is equivalent to showing $a_d = \text{vol}(P')$ in the lattice $\mathbb{Z}^d$.

Observe now by Erhart's formula

$$\begin{aligned} a_d &= \lim_{t\to\infty} \frac{\sum_{j=0}^{d} a_j t^j}{t^d} \\ &= \lim_{t\to\infty} \frac{|\mathbb{Z}^d \cap tP'|}{t^d}. \end{aligned}$$

Thus, we wish to establish $\lim_{t\to\infty} \frac{|\mathbb{Z}^d \cap tP'|}{t^d} = \text{vol}(P')$.

For any $\delta > 0$, let consider the cube $[0,\delta)^d$, which has volume δ^d. We have

$$\begin{aligned} a_d &= \lim_{t\to\infty} \frac{|\mathbb{Z}^d \cap (t \cdot [0,\delta)^d)|}{t^d} \\ &= \lim_{t\to\infty} \frac{|\mathbb{Z}^d \cap [0,\delta t)^d|}{t^d} \\ &= \lim_{t\to\infty} \frac{\delta^d t^d}{t^d} \\ &= \delta^d, \end{aligned}$$

and so the claim holds for a single cube of arbitrary size.

Let $\mathcal{C}_\delta$ be a collection of cubes of common size δ that approximate P'. Since the cubes are disjoint, we have

$$\lim_{\delta \downarrow 0} \sum_{C \in \mathcal{C}_d} \mathrm{vol}_d(C) = \mathrm{vol}(P').$$

Given a particular set of approximating cubes $\mathcal{C}_\delta$, let P'_δ denote the set $\bigcup_{C \in \mathcal{C}_\delta} C$. Finally, we have

$$\begin{aligned}
a_d &= \lim_{t \to \infty} \frac{|\mathbb{Z}^d \cap tP'|}{t^d} \\
&= \lim_{t \to \infty} \frac{|\mathbb{Z}^d \cap t\,(\lim_{\delta \downarrow 0} P'_\delta)\,|}{t^d} \\
&= \lim_{\delta \downarrow 0} \lim_{t \to \infty} \frac{|\mathbb{Z}^d \cap tP'_\delta|}{t^d} \\
&= \lim_{\delta \downarrow 0} \lim_{t \to \infty} \sum_{C \in \mathcal{C}_\delta} \frac{|\mathbb{Z}^d \cap tC|}{t^d} \\
&= \lim_{\delta \downarrow 0} \sum_{C \in \mathcal{C}_\delta} \lim_{t \to \infty} \frac{|\mathbb{Z}^d \cap tC|}{t^d} \\
&= \lim_{\delta \downarrow 0} \sum_{C \in \mathcal{C}_\delta} \mathrm{vol}_d(C) \\
&= \lim_{\delta \downarrow 0} \mathrm{vol}_d(P'_\delta) \\
&= \mathrm{vol}_d(P').
\end{aligned}$$

□