CMSC 754: Lecture 6
Halfplane Intersection and Point-Line Duality

Reading: Chapter 4 in the 4M's, with some elements from Sections 8.2 and 11.4.

Halfplane Intersection: Today we begin studying another fundamental topic in geometric computing and convexity. Recall that any line in the plane splits the plane into two regions, one lying on either side of the line. Each such region is called a *halfplane*. We say that a halfplane is either *closed* or *open* depending, respectively, on whether or not it contains the line. Unless otherwise stated, we will assume that halfplanes are closed.

In the *halfplane intersection problem*, we are given a collection of n halfplanes $H = \{h_1, \ldots, h_n\}$, and the objective is to compute their intersection. It is easy to see that the intersection of halfspaces is a convex polygon (see Fig. 1(a)), but this polygon may be unbounded (see Fig. 1(b)) or even empty (see Fig. 1(c)).

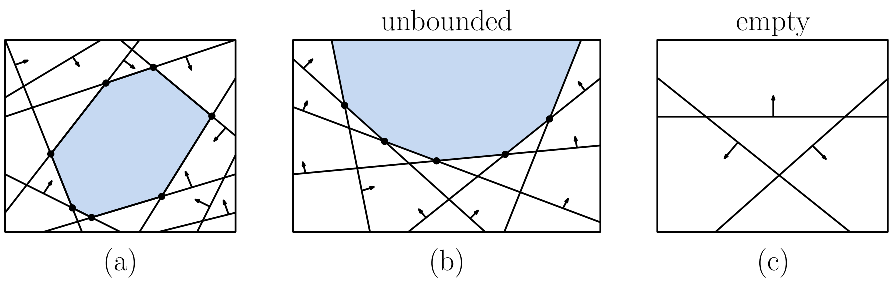

Fig. 1: Halfplane intersection.

Clearly, the number of sides of the resulting polygon is at most n, but may be smaller since some halfspaces may not contribute to the final shape.

Halfspace Intersection: In d-dimensional space the corresponding notion is a *halfspace*, which is the set of points lying to one side of a $(d-1)$-dimensional hyperplane. The intersection of halfspaces is a *convex polytope*. The resulting polytope will have at most n facets (at most one per halfspace), but (surprisingly) the overall complexity can be much higher.

A famous result, called *McMullen's Upper-Bound Theorem* states that a polytope with n facets in dimension d can have up to $O(n^{\lfloor d/2 \rfloor})$ vertices. (In dimensions 2 and 3, this is linear in the number of halfspaces, but even in dimension 4 the number of vertices can jump to $O(n^2)$.) Obtaining such a high number of vertices takes some care, but the bound is tight. There is a famous class of polytopes, called the *cyclic polytopes*, that achieve this bound. Symmetrically, the convex hull of n points in dimension d defines a convex polytope that can have $O(n^{\lfloor d/2 \rfloor})$ facets, and this bound is also tight.

Representing Lines and Hyperplanes: (Digression) While we will usually treat geometric objects rather abstractly, it may be useful to explore a bit regarding how lines, halfspaces, and their higher dimensional counterparts are represented. These topics would be covered in a more complete course on projective geometry or convexity.

Explicit Representation: If we think of a line as a linear function of the variable x, we can express any (nonvertical) line by the equation $y = ax + b$, where a is the slope and b is the y-intercept.

In dimension d, we can think of the dth coordinate as being special, and we will make the convention of referring to the d-th coordinate axis as pointing vertically upwards. We can express any "nonvertical" $(d-1)$-dimensional hyperplane by the set of points $(x_1, \ldots, x_d)$, where $x_d = \sum_{i=1}^{d-1} a_i x_i + b$, thus x_d is expressed "explicitly" as a linear function of the first $d-1$ coordinates.

The associated halfspaces arise replacing "=" with an inequality, e.g., the *upper halfplane* is the set (x, y) such that $y \geq ax + b$, and the *lower halfplane* is defined analogously.

Implicit Representation: The above representation has the shortcoming that it cannot represent vertical objects. A more general approach (which works for both hyperplanes and curved surfaces) is to express the object implicitly as the zero-set of some function of the coordinates. In the case of a line in the plane, we can represent the line as the set of points (x, y) that satisfy the linear function $f(x, y) = 0$, where $f(x, y) = ax + by + c$, for scalars a, b, and c. The corresponding halfplanes are just the sets of points such that $f(x, y) \geq 0$ and $f(x, y) \leq 0$.

This has the advantage that it can represent any line in the Euclidean plane, but the representation is not unique. For example, the line described by $5x - 3y = 2$ is the same as the line described by $10x - 6y = 4$, or any scalar multiple thereof. We could apply some normalization to overcome this, for example by requiring that $c = 1$ or $a^2 + b^2 = 1$.

Parametric Representation: The above representations describe $(d-1)$-dimensional hyperplanes in d-dimensional space. What if you want to represent a line, or more generally, a flat object some dimension $k < d - 1$? We can represent such an object as the affine span of a set of points. For example, to represent a line in 3-dimensional space, we can given two points p and q on the line, and then any point on this line can be expressed as an affine combination $(1 - \alpha)p + \alpha q$, for $\alpha \in \mathbb{R}$. This is called the *parametric representation*, since each point on the object is identified through the value of the parameter α. In general, we can represent any *k-dimensional affine subspace* (or *k-flat*) parametrically as the affine combination of $k + 1$ points, that is, $\sum_{i=1}^{k+1} \alpha_i p_i$, where $\sum_{i=1}^{k+1} \alpha_i = 1$. We can think of the function as being generated by k of the parameters, say α_1 through α_k, and α_{k+1} is determined by the constraint that the α values sum to 1.

Divide-and-Conquer Algorithm: Returning to the halfplane intersection problem, recall that we are given a set $H = \{h_1, \ldots, h_n\}$ of halplanes and wish to compute their intersection. Here is a simple divide-and-conquer algorithm.

(1) If $n = 1$, then just return this halfplane as the answer.

(2) Otherwise, partition H into subsets H_1 and H_2, each of size roughly $n/2$.

(3) Compute the intersections $K_1 = \bigcap_{h \in H_1} h$ and $K_2 = \bigcap_{h \in H_2} h$ recursively.

(4) If either either K_1 or K_2 is empty, return the empty set. Otherwise, compute the intersection of the convex polygons K_1 and K_2 (by the procedure described below).

If we let $I(n)$ denote the time needed to intersect two convex polygons, each with at most n vertices, we obtain the following recurrence for the overall running time:

$$T(n) = \begin{cases} 1 & \text{if } n = 1, \\ 2T(n/2) + I(n) & \text{if } n > 1, \end{cases}$$

We will show below that $I(n) \leq cn$, for some constant c. It follows by standard results (consult the *Master Theorem* in CLRS) that $T(n)$ is $O(n \log n)$.

Intersecting Two Convex Polygons: The only remaining task is the process of intersecting two convex polygons, K_1 and K_2 (see Fig. 2(a)). Note that these are somewhat special convex polygons because they may be empty or unbounded.

We can compute the intersection by a left-to-right plane sweep in $O(n)$ time (see Fig. 2(b)). We begin by breaking the boundaries of the convex polygons into their upper and lower chains. (This can be done in $O(n)$ time.) By convexity, the sweep line intersects the boundary of each convex polygon K_i in at most two points, one for the upper chain and one for the lower chain. Hence, the sweep-line status contains at most four points. This implies that updates to the sweep-line status can be performed in $O(1)$ time. Also, we need keep track of a constant number of events at any time, namely the right endpoints of the current segments in the sweep-line status, and the intersections between consecutive pairs of segments. Thus, each step of the plane-sweep process can be performed in $O(1)$ time.

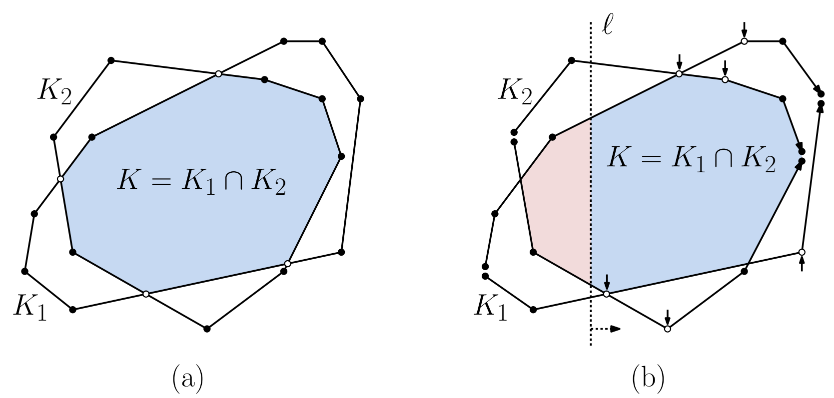

Fig. 2: Intersecting two convex polygons by plane sweep.

The total number of events is equal to the total number vertices, which is n, and the total number of intersection points. It is an easy exercise (which we leave to you) to prove that two convex polygons with a total of n sides can intersect at most $O(n)$ times. Thus, the overall running time is $O(n)$.

Lower Envelopes and Duality: Let's next consider a variant of the halfplane intersection problem. Given any set of nonvertical lines $L = \{\ell_1, \ell_2, \ldots, \ell_n\}$ in the plane. Each line defines two natural halfplanes, and upper and lower halfplane. The intersection of all the lower halfplanes is called the *lower envelope* of L and the *upper envelope* is defined analogously (see Fig. 3). Let's assume that each line ℓ_i is given explicitly as $y = a_i x - b_i$.

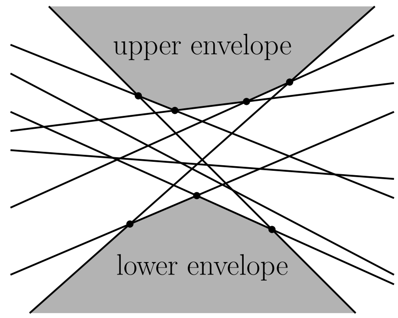

Fig. 3: Lower and upper envelopes.

The lower envelope problem is a restriction of the halfplane intersection problem, but it an interesting restriction. Notice that any halfplane intersection problem that does not involve any vertical lines can be rephrased as the intersection of two envelopes, a lower envelope defined by the lower halfplanes and an upper envelope defined by the upward halfplanes.

We will see that solving the lower envelope problem is very similar (in fact, essentially the same as) solving the upper convex hull problem. Indeed, they are so similar that exactly the same algorithm will solve both problems, without changing even a single character of code! All that changes is the way in which you interpret the inputs and the outputs.

Lines, Points, and Incidences: In order to motivate duality, let us discuss the representation of lines in the plane. Each line can be represented in a number of ways, but for now, let us assume the representation $y = ax - b$, for some scalar values a and b. (Why $-b$ rather than $+b$? The distinction is unimportant, but it will simplify some of the notation defined below.) We cannot represent vertical lines in this way, and for now we will just ignore them.

Therefore, in order to describe a line in the plane, you need only give its two coefficients (a, b). Thus, lines in the plane can be thought of as points in a new 2-dimensional space, in which the coordinate axes are labeled (a, b), rather than (x, y). For example, the line $\ell : y = 2x + 1$ corresponds to the point $(2, -1)$ in this space, which we denote by ℓ^*. Conversely, each point $p = (a, b)$ in this space of "lines" corresponds to a nonvertical line, $y = ax - b$ in the original plane, which we denote by p^*. We will call the original (x, y)-plane the *primal plane*, and the new (a, b)-plane the *dual plane*.

This insight would not be of much use unless we could say something about how geometric relationships in one space relate to the other. The connection between the two involves incidences between points and line.

Primal Relation	**Dual Relation**
Two (nonparallel) lines meet in a point	Two points join to form a line
A point may lie above/below/on a line	A line may pass above/below/through a point
Three points may be collinear	Three lines may pass through the same point

We'll show that these relationships are preserved by duality. For example, consider the two lines $\ell_1 : y = 2x + 1$ and the line $\ell_2 : y = -\frac{x}{2} + 6$ (see Fig. 4(a)). These two lines intersect at the point $p = (2, 5)$. The duals of these two lines are $\ell_1^* = (2, -1)$ and $\ell_2^* = \left(-\frac{1}{2}, -6\right)$. The line in the (a, b) dual plane passing through these two points is easily verified to be $b = 2a - 5$. Observe that this is exactly the dual of the point p (see Fig. 4(b)). (As an exercise, prove this for two general lines.)

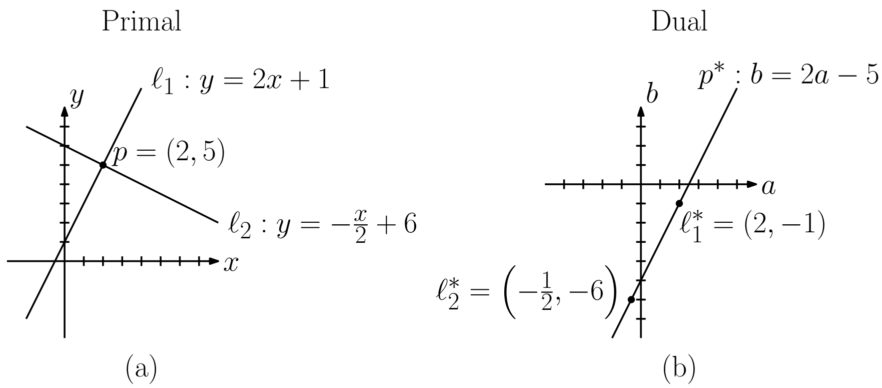

Fig. 4: The primal and dual planes.

Point-Line Duality: Let us explore this dual transformation more formally. Duality (or more specifically *point-line duality*) is a transformation that maps points in the plane to lines and lines to point. (More generally, it maps points in d-space to hyperplanes dimension d.) We denote this transformation using a asterisk $(*)$ as a superscript. Thus, given point p and line ℓ in the primal plane we define ℓ^* and p^* to be a point and line, respectively, in the dual plane defined as follows.[1]

$$\begin{array}{lll} \ell : y = \ell_a x - \ell_b & \Rightarrow & \ell^* = (\ell_a, \ell_b) \\ p = (p_x, p_y) & \Rightarrow & p^* : b = p_x a - p_y. \end{array}$$

It is convenient to define the dual transformation so that it is its own inverse (that is, it is an involution). In particular, it maps points in the dual plane to lines in the primal, and vice versa. For example, given a point $p = (p_a, p_b)$ in the dual plane, its dual is the line $y = p_a x - p_b$ in the primal plane, and is denoted by p^*. It follows that $p^{**} = p$ and $\ell^{**} = \ell$.

Properties of Point-Line Duality: Duality has a number of interesting properties, each of which is easy to verify by substituting the definition and a little algebra.

Self Inverse: $p^{**} = p$.

Order reversing: Point p is above/on/below line ℓ in the primal plane if and only if line p^* is below/on/above point ℓ^* in the dual plane, respectively (see Fig. 5).

Intersection preserving: Lines ℓ_1 and ℓ_2 intersect at point p if and only if the dual line p^* passes through points ℓ_1^* and ℓ_2^*.

[1] Duality can be generalized to higher dimensions as well. In $\mathbb{R}^d$, let us identify the y axis with the d-th coordinate vector, so that an arbitrary point can be written as $p = (x_1, \ldots, x_{d-1}, y)$ and a $(d-1)$-dimensional hyperplane can be written as $h : y = \sum_{i=1}^{d-1} a_i x_i - b$. The dual of this hyperplane is $h^* = (a_1, \ldots, a_{d-1}, b)$ and the dual of the point p is $p^* : b = \sum_{i=1}^{d-1} x_i a_i - y$. All the properties defined for point-line relationships generalize naturally to point-hyperplane relationships, where notions of above and below are based on the assumption that the y (or b) axis is "vertical."

Collinearity/Coincidence: Three points are collinear in the primal plane if and only if their dual lines intersect in a common point.

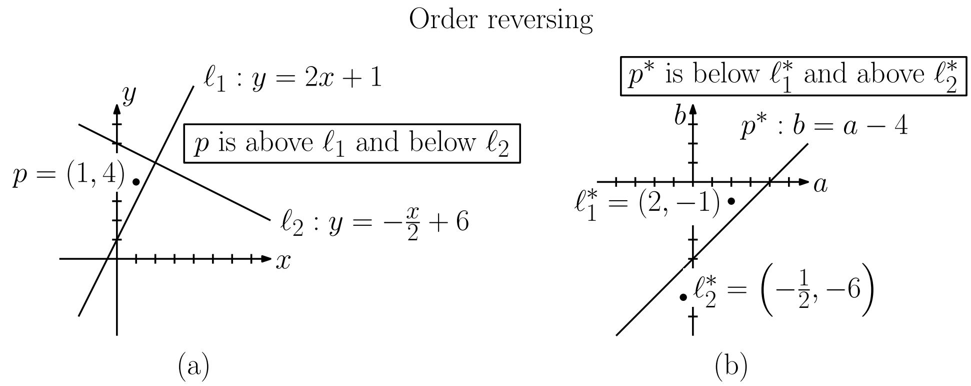

Fig. 5: The order-reversing property.

The self inverse property was already established (essentially by definition). To verify the order reversing property, consider any point p and any line ℓ.

$$p \text{ is on or above } \ell \iff p_y \geq \ell_a p_x - \ell_b \iff \ell_b \geq p_x \ell_a - p_y \iff p^* \text{ is on or below } \ell^*$$

(From this it should be apparent why we chose to negate the y-intercept when dualizing points to lines.) The other two properties (intersection preservation and collinearity/coincidence are direct consequences of the order reversing property.)

Convex Hulls and Envelopes: Let us return now to the question of the relationship between convex hulls and the lower/upper envelopes of a collection of lines in the plane. The following lemma demonstrates the, under the duality transformation, the convex hull problem is dually equivalent to the problem of computing lower and upper envelopes.

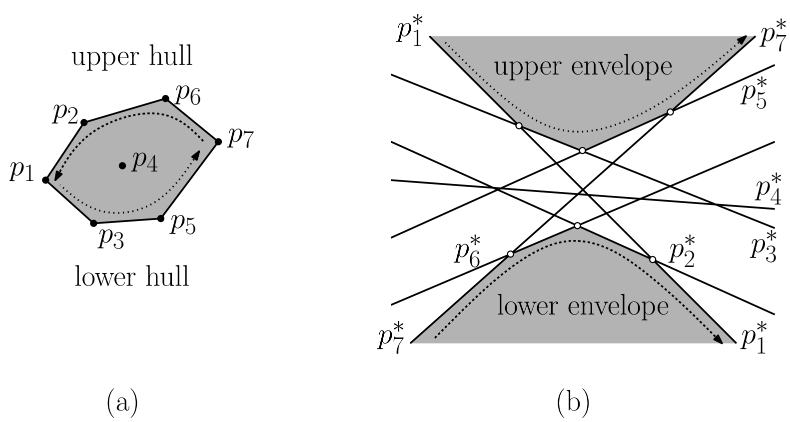

Fig. 6: Equivalence of hulls and envelopes.

Lemma: Let P be a set of points in the plane. The counterclockwise order of the points along the upper (lower) convex hull of P (see Fig. 6(a)), is equal to the left-to-right order of the sequence of lines on the lower (upper) envelope of the dual P^* (see Fig. 6(b)).

Proof: We will prove the result just for the upper hull and lower envelope, since the other case is symmetrical. For simplicity, let us assume that no three points are collinear.

Consider a pair of points p_i and p_j that are consecutive vertices on the upper convex hull. This is equivalent to saying that all the other points of P lie beneath the line ℓ_{ij} that passes through both of these points.

Consider the dual lines p_i^* and p_j^*. By the incidence preserving property, the dual point ℓ_{ij}^* is the intersection point of these two lines. (By general position, we may assume that the two points have different x-coordinates, and hence the lines have different slopes. Therefore, they are not parallel, and the intersection point exists.)

By the order reversing property, all the dual lines of P^* pass above point ℓ_{ij}^*. This is equivalent to saying the ℓ_{ij}^* lies on the lower envelope of P^*.

To see how the order of points along the hulls are represented along the lower envelope, observe that as we move counterclockwise along the upper hull (from right to left), the slopes of the edges increase monotonically. Since the slope of a line in the primal plane is the a-coordinate of the dual point, it follows that as we move counterclockwise along the upper hull, we visit the lower envelope from left to right.

One rather cryptic feature of this proof is that, although the upper and lower hulls appear to be connected, the upper and lower envelopes of a set of lines appears to consist of two disconnected sets. To make sense of this, we should interpret the primal and dual planes from the perspective of *projective geometry*, and think of the rightmost line of the lower envelope as "wrapping around" to the leftmost line of the upper envelope, and vice versa. The places where the two envelopes wraps around correspond to the vertical lines (having infinite slope) passing through the left and right endpoints of the hull. (As an exercise, can you see which is which?)

Primal/Dual Equivalencies: There are a number of computational problems that are defined in terms of affine properties of point and line sets. These can be expressed either in primal or in dual form. In many instances, it is easier to visualize the solution in the dual form. We will discuss many of these later in the semester. For each of the following, can you determine what the dual equivalent is?

- Given a set of points P, find the narrowest slab (that is, a pair of parallel lines) that contains P. Define the width of the slab to be the vertical distance between its bounding lines (see Fig. 7(a)).
- Given a convex polygon K, find the longest vertical line segment with one endpoint on K's upper hull and one on its lower hull (see Fig. 7(b)).
- Given a set of points P, find the triangle of smallest area determined by any three points of P (see Fig. 7(c)). (If three points are collinear, then they define a degenerate triangle of area 0.)

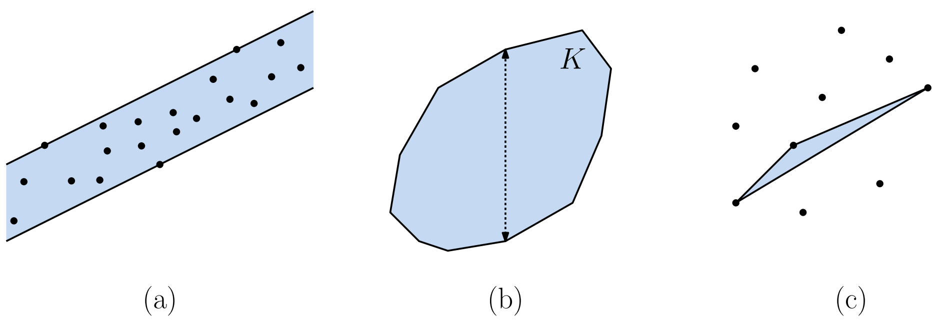

Fig. 7: sEquivalence of hulls and envelopes.