

CMSC 754: Lecture 2

Convex Hulls in the Plane

Reading: Some of the material of this lecture is covered in Chapter 1 in the 4M's (de Berg, Cheong, van Kreveld, Schwarzkopf). The divide-and-conquer algorithm is given in Joseph O'Rourke's, "Computational Geometry in C."

Convex Hulls: In this lecture we will consider a fundamental structure in computational geometry, called the *convex hull*. We will give a more formal definition later, but, given a set P of points in the plane, the convex hull of P , denoted $\text{conv}(P)$, can be defined intuitively by surrounding a collection of points with a rubber band and then letting the rubber band "snap" tightly around the points (see Fig. 1).




Fig. 1: A point set and its convex hull.

The (planar) *convex hull problem* is, given a discrete set of n points P in the plane, output a representation of P 's convex hull. The convex hull is a closed convex polygon, the simplest representation is a counterclockwise enumeration of the vertices of the convex hull. In higher dimensions, the convex hull will be a convex polytope. We will discuss the representation of polytopes in future lectures, but in 3-dimensional space, the representation would consist of a vertices, edges, and faces that constitute the boundary of the polytope.

There are a number of reasons that the convex hull of a point set is an important geometric structure. One is that it is one of the simplest shape approximations for a set of points. (Other examples include minimum area enclosing rectangles, circles, and ellipses.) It can also be used for approximating more complex shapes. For example, the convex hull of a polygon in the plane or polyhedron in 3-space is the convex hull of its vertices.

Also many algorithms compute the convex hull as an initial stage in their execution or to filter out irrelevant points. For example, the *diameter* of a point set is the maximum distance between any two points of the set. It can be shown that the pair of points determining the diameter are both vertices of the convex hull. Also observe that minimum enclosing convex shapes (such as the minimum area rectangle, circle, and ellipse) depend only on the points of the convex hull.

Convexity: Before getting to discussion of the algorithms, let's begin with a few standard definitions regarding convexity and convex sets. For any $d \geq 1$, let \mathbb{R}^d denote real d -dimensional space, that is, the set of d -dimensional vectors over the real numbers.

Affine and convex combinations: Given two points $p = (p_x, p_y)$ and $q = (q_x, q_y)$ in \mathbb{R}^d , we can generate any point on the line \overleftrightarrow{pq} as a linear combination of their coordinates,

where the coefficient sum to 1:

$$(1 - \alpha)p + \alpha q = ((1 - \alpha)p_x + \alpha q_x, (1 - \alpha)p_y + \alpha q_y).$$

This is called an *affine combination* of p and q (see Fig. 2(a)).




Fig. 2: Affine and convex combinations.

By adding the additional constraint that $0 \leq \alpha \leq 1$, the set of points generated lie on the *line segment* \overline{pq} (see Fig. 2(b)). This is called a *convex combination*. Notice that this can be viewed as taking a weighted average of p and q . As α approaches 1, the point lies closer to p and when α approaches zero, the point lies closer to q .

It is easy to extend both types of combinations to more than two points. For example, given k points $\{p_1, \dots, p_k\}$ an affine combination of these points is the linear combination

$$\sum_{i=1}^k \alpha_i p_i, \quad \text{such that } \alpha_1 + \dots + \alpha_k = 1.$$

When $0 \leq \alpha_i \leq 1$ for all i , the result is called a *convex combination*.

The set of all affine combinations of three (non-collinear) points generates a plane, and generally, the resulting set is called the *affine span* or *affine closure* of the points. The set of all convex combinations of a set of points is the *convex hull* of the point set.

Convexity: A set $K \subseteq \mathbb{R}^d$ is *convex* if given any points $p, q \in K$, the line segment \overline{pq} is entirely contained within K (see Fig. 3(a)). This is equivalent to saying that K is “closed” under convex combinations. Examples of convex sets in the plane include circular disks (the set of points contained within a circle), the set of points lying within any regular n -sided polygon, lines (infinite), line segments (finite), rays, and halfspaces (that is, the set of points lying to one side of a line).

Open/Closed: A set in \mathbb{R}^d is said to be *open* if it does not include its boundary. (The formal definition is a bit messy, so I hope this intuitive definition is sufficient.) A set that includes its boundary is said to be *closed*. (See Fig. 3(b).)

Boundedness: A convex set is *bounded* if it can be enclosed within a sphere of a fixed radius. Otherwise, it is *unbounded* (see Fig. 3(c)). For example, line segments, regular n -gons, and circular disks are all bounded. In contrast, lines, rays, halfspaces, and infinite cones are unbounded.

Convex body: A closed, bounded convex set is called a *convex body*.




Fig. 3: Basic definitions.

Support line/hyperplane: An important property of any convex set K in the plane is that at every point p on the boundary of K , there exists at least one line ℓ (or generally a $(d - 1)$ -dimensional hyperplane in higher dimensions) that passes through p such that K lies entirely in one of the closed halfplanes (halfspaces) defined by ℓ (see Fig. 3(d)). Such a line is called a *support line* for K . Observe that there may generally be multiple support lines passing through a given boundary point of K (e.g., when the point is a vertex of the convex hull).

Equivalent definitions: We can define the *convex hull* of a set of points P either in an additive manner as the closure of all convex combinations of the points or in a subtractive manner as the intersection of the set of all halfspaces that contain the point set.

When computing convex hulls, we will usually take P to be a finite set of points. In such a case, $\text{conv}(P)$ will be a convex polygon. Generally P could be an infinite set of points. For example, we could talk about the convex hull of a collection of circles. The boundary of such a shape would consist of a combination of circular arcs and straight line segments.

General Position: As in many of our algorithms, it will simplify the presentation to avoid lots of special cases by assuming that the points are in *general position*. This effectively means that degenerate configurations (e.g., two points sharing the same x or y coordinate, three points being collinear, etc.) do not arise in the input. More specifically, a point set fails to be in general position if it possesses some property (such as collinearity) that fails to hold if the point coordinates are perturbed infinitesimally. General position assumptions are almost never critical to the efficiency of an algorithm. They are merely a convenience to avoid the need of dealing with lots of special cases in designing our algorithms.

Graham's scan: We will begin with a presentation of a simple $O(n \log n)$ algorithm for the convex hull problem. It is a simple variation of a famous algorithm for convex hulls, called *Graham's scan*, which dates back to the early 1970's. The algorithm is loosely based on a common approach for building geometric structures called *incremental construction*. In such an algorithm object (points here) are added one at a time, and the structure (convex hull here) is updated with each new insertion.

An important issue with incremental algorithms is the order of insertion. If we were to add points in some arbitrary order, we would need some method of testing whether the newly added point is inside the existing hull. It will simplify things to add points in some appropriately sorted order, in our case, in increasing order of x -coordinate. This guarantees that each newly added point is outside the current hull. (Note that Graham's original algorithm sorted

points in a different way. It found the lowest point in the data set and then sorted points cyclically around this point. Sorting by x -coordinate seems to be a bit easier to implement, however.)

Since we are working from left to right, it would be convenient if the convex hull vertices were also ordered from left to right. As mentioned above, the convex hull is a convex polygon, which can be represented as a cyclic sequence of vertices. It will make matters a bit simpler for us to represent the boundary of the convex hull as two polygonal chains, one representing its upper part, called the *upper hull* and one representing the lower part, called the *lower hull* (see Fig. 4(a)).




Fig. 4: (a) Upper and lower hulls and (b) the left-hand turn property of points on the upper hull.

It suffices to show how to compute the upper hull, since the lower hull is symmetrical. (Just flip the picture upside down.) Once the two hulls have been computed, we can simply concatenate them with the reversal of the other to form the final hull.

Observe that a point $p \in P$ lies on the upper hull if and only if there is a support line passing through p such that all the points of P lie on or below this line. Our algorithm will be based on the following lemma, which characterizes the upper hull of P . This is a simple consequence of the convexity. The first part says that the line passing through each edge of the hull is a support line, and the second part says that as we walk from right to left along the upper hull, we make successive left-hand turns (see Fig. 4(b)).

Lemma 1: Let $\langle p_{i_1}, \dots, p_{i_m} \rangle$ denote the vertices of the upper hull of P , sorted from left to right. Then for $1 \leq j \leq m$, (1) all the points of P lie on or below the line $\overline{p_{i_j} p_{i_{j-1}}}$ joining consecutive vertices and (2) each consecutive triple $\langle p_{i_j} p_{i_{j-1}} p_{i_{j-2}} \rangle$ forms a left-hand turn.

Let $\langle p_1, \dots, p_n \rangle$ denote the sequence of points sorted by increasing order of x -coordinates. For i ranging from 1 to n , let $P_i = \langle p_1, \dots, p_i \rangle$. We will store the vertices of the upper hull of P_i on a stack S , where the top-to-bottom order of the stack corresponds to the right-to-left order of the vertices on the upper hull. Let $S[t]$ denote the stack's top. Observe that as we read the stack elements from top to bottom (that is, from right to left) consecutive triples of points of the upper hull form a (strict) left-hand turn (see Fig. 4(b)). As we push new points on the stack, we will enforce this property by popping points off of the stack that violate it.

Turning and orientations: Before proceeding with the presentation of the algorithm, we should first make a short digression to discuss the question of how to determine whether three points

form a “left-hand turn.” This can be done by a powerful primitive operation, called an *orientation test*, which is fundamental to many algorithms in computational geometry.

Given an ordered triple of points $\langle p, q, r \rangle$ in the plane, we say that they have *positive orientation* if they define a counterclockwise oriented triangle (see Fig. 5(a)), *negative orientation* if they define a clockwise oriented triangle (see Fig. 5(b)), and *zero orientation* if they are collinear, which includes as well the case where two or more of the points are identical (see Fig. 5(c)). Note that orientation depends on the order in which the points are given.




Fig. 5: Orientations of the ordered triple (p, q, r) .

Orientation is formally defined as the sign of the determinant of the points given in homogeneous coordinates, that is, by prepending a 1 to each coordinate. For example, in the plane, we define

$$\text{Orient}(p, q, r) = \det \begin{pmatrix} 1 & p_x & p_y \\ 1 & q_x & q_y \\ 1 & r_x & r_y \end{pmatrix}.$$

Observe that in the 1-dimensional case, $\text{Orient}(p, q)$ is just $q - p$. Hence it is positive if $p < q$, zero if $p = q$, and negative if $p > q$. Thus orientation generalizes the familiar 1-dimensional binary relations $<, =, >$.

Also, observe that the sign of the orientation of an ordered triple is unchanged if the points are translated, rotated, or scaled (by a positive scale factor). A reflection transformation (e.g., $f(x, y) = (-x, y)$) reverses the sign of the orientation. In general, applying any affine transformation to the point alters the *sign* of the orientation according to the *sign* of the determinant of the matrix used in the transformation. (By the way, the notion of orientation can be generalized to $d + 1$ points in d -dimensional space, and is related to the notion of *chirality* in Chemistry and Physics. For example, in 3-space the orientation is positive if the point sequence defines a right-handed screw.)

Given a sequence of three points p, q, r , we say that the sequence $\langle p, q, r \rangle$ makes a (strict) *left-hand turn* if $\text{Orient}(p, q, r) > 0$.

Graham’s algorithm continued: Returning to the algorithm, let us consider the insertion of the i th point, p_i (see Fig. 6(a)). First observe that p_i is on the upper hull of P_i (since it is the rightmost point seen so far). Let p_j be its predecessor on the upper hull of P_i . We know from Lemma 1 that all the points of P_i lie on or below the line $\overline{p_i p_j}$. Let p_{j-1} be the point immediately preceding p_j on the upper hull. We also know from this lemma that $\langle p_i p_j p_{j-1} \rangle$ forms a left-hand turn. Clearly then, if any triple $\langle p_i, S[t], S[t - 1] \rangle$ does *not* form a left-hand

turn (that is, $\text{Orient}(p_i, S[t], S[t - 1]) \leq 0$), we may infer that $S[t]$ is *not* on the upper hull, and hence it is safe to delete it by popping it off the stack. We repeat this until we find a left-turning triple (see Fig. 6(b)) or hitting the bottom of the stack. Once this happens, we push p_i on top of the stack, making it the rightmost vertex on the upper hull (see Fig. 6(c)). The algorithm is presented in the code block below.




Fig. 6: Graham's scan.

Graham's Scan

- (1) Sort the points according to increasing order of their x -coordinates, denoted $\langle p_1, p_2, \dots, p_n \rangle$.
 - (2) push p_1 and then p_2 onto S .
 - (3) for $i \leftarrow 3, \dots, n$ do:
 - (a) while ($|S| \geq 2$ and $\text{Orient}(p_i, S[t], S[t - 1]) \leq 0$) pop S .
 - (b) push p_i onto S .
-

Correctness: The correctness of the algorithm was essentially established by Lemma 1 and the above explanation. (Where we showed that it is safe to pop all right-turning triples off the stack, and safe to push p_i .) The only remaining issue is whether we might stop too early. In particular, might we encounter a left-turning triple *before* reaching p_j ? We claim that this cannot happen. Suppose to the contrary that before reaching p_j , we encounter triple $\langle p_i, S[t], S[t - 1] \rangle$ that forms a left-hand turn, but $S[t] \neq p_j$ (see Fig. 7). We know that $S[t]$ lies to the right of p_j . By Lemma 1, all the points of P_{i-1} (including p_j) lie on or below the line $\overline{S[t]S[t - 1]}$. But if p_j lies below this line, it follows that the triple $\langle p_i, S[t], p_j \rangle$ forms a left-hand turn, and this implies that $S[t]$ lies above the line $\overline{p_j p_i}$. This contradicts Lemma 1, because by our hypothesis, $\overline{p_j p_i}$ is an edge of the upper hull of P_i , and no point of P_i can lie above an edge of the upper hull.




Fig. 7: Correctness of Graham's scan.

How much detail? A question that often arises at this point of the semester is, “how much detail

is needed in giving a geometrical proof of correctness?” You might find the above proof to be a bit too vague. There is a bit of art between the extremes of producing proofs that are not convincing from those that contain excessive details. Whenever feasible, you should reduce base-level assertions to configurations involving just a constant number of points (e.g., the points involved in an orientation test). It may be helpful to add additional constructions (e.g., support lines) to help illustrate points. Don’t be fooled by your drawings. Finally, note that more detail is not always better. Your proof is intended to be read by a human, not a compiler or automated proof verifier. You should rely your (intelligent) reader to fill in low-level geometric reasoning.

Running-time analysis: We will show that Graham’s algorithm runs in $O(n \log n)$ time. Clearly, it takes this much time for the initial sorting of the points. After this, we will show that $O(n)$ time suffices for the rest of the computation.

Let d_i denote the number of points that are popped (deleted) on processing p_i . Because each orientation test takes $O(1)$ time, the amount of time spent processing p_i is $O(d_i + 1)$. (The extra +1 is for the last point tested, which is not deleted.) Thus, the total running time is proportional to

$$\sum_{i=1}^n (d_i + 1) = n + \sum_{i=1}^n d_i.$$

To bound $\sum_i d_i$, observe that each of the n points is pushed onto the stack once. Once a point is deleted it can never be deleted again. Since each of n points can be deleted at most once, $\sum_i d_i \leq n$. Thus after sorting, the total running time is $O(n)$. Since this is true for the lower hull as well, the total time is $O(2n) = O(n)$.

Convex Hull by Divide-and-Conquer: As with sorting, there are many different approaches to solving the convex hull problem for a planar point set P . Next, we will consider another $O(n \log n)$ algorithm, which is based on divide-and-conquer. It can be viewed as a generalization of the well-known MergeSort sorting algorithm (see any standard algorithms text). Here is an outline of the algorithm. As with Graham’s scan, we will focus just on computing the upper hull, and the lower hull will be computed symmetrically.

The algorithm begins by sorting the points by their x -coordinate, in $O(n \log n)$ time. It splits the point set in half at its median x -coordinate, computes the upper hulls of the left and right sets recursively, and then merges the two upper hulls into a single upper hull. This latter process involves computing a line, called the *upper tangent*, that is a line of support for both hulls. The remainder of the algorithm is shown in the code section below.

Divide-and-Conquer (Upper) Convex Hull

- (1) If $|P| \leq 3$, then compute the upper hull by brute force in $O(1)$ time and return.
 - (2) Otherwise, partition the point set P into two sets P' and P'' of roughly equal sizes by a vertical line.
 - (3) Recursively compute upper convex hulls of P' and P'' , denoted H' and H'' , respectively (see Fig. 8(a)).
 - (4) Compute the upper tangent $\ell = \overline{p'p''}$ (see Fig. 8(b)).
 - (5) Merge the two hulls into a single upper hull by discarding all the vertices of H' to the right of p' and the vertices of H'' to the left of p'' (see Fig. 8(c)).
-




Fig. 8: Divide and conquer (upper) convex hull algorithm.

Computing the upper tangent: The only nontrivial step is that of computing the common tangent line between the two upper hulls. Our algorithm will exploit the fact that the two hulls are separated by a vertical line. The algorithm operates by a simple “walking procedure.” We initialize p' to be the rightmost point of H' and p'' to be the leftmost point of H'' (see Fig. 9(a)). We will walk p' backwards along H' and walk p'' forwards along H'' until we hit the vertices that define the tangent line. As in Graham’s scan, it is possible to determine just how far to walk simply by applying orientation tests. In particular, let q' be the point immediately preceding p' on H' , and let q'' be the point immediately following p'' on H'' . Observe that if $\text{Orient}(p', p'', q'') \geq 0$, then we can advance p'' to the next point along H'' (see Fig. 9(a)). Symmetrically, if $\text{Orient}(p'', p', q') \leq 0$, then we can advance p' to its predecessor along H' (see Fig. 9(b)). When neither of these conditions applies, that is, $\text{Orient}(p', p'', q'') < 0$ and $\text{Orient}(p'', p', q') > 0$, we have arrived at the desired points of mutual tangency (see Fig. 9(c)).

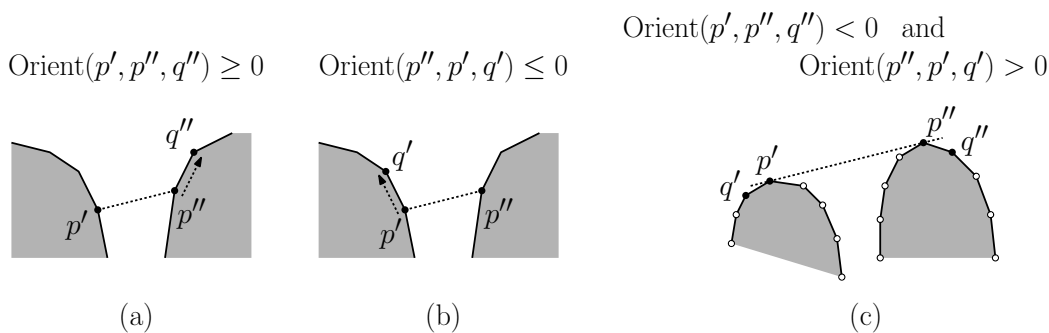


Fig. 9: Computing the upper tangent.

There is one rather messy detail in implementing this algorithm. This arises if either q' or q'' does not exist because we have arrived at the leftmost vertex of H' or the rightmost vertex of H'' . We can avoid having to check for these conditions by creating two *sentinel points*. We create a new leftmost vertex for H' that lies infinitely below its original leftmost vertex, and we create a new rightmost vertex for H'' that lies infinitely below its original rightmost vertex. The tangency computation will never arrive at these points, and so we do not need to add a special test for the case when q' and q'' do not exist. The algorithm is presented in the following code block.

A formal proof of correctness of this procedure is similar to that of Graham’s scan (but

UpperTangent(H', H'') :

- (1) Let p' be the rightmost point of H' , and let q' be its predecessor.
 - (2) Let p'' be the leftmost point of H'' , and let q'' be its successor.
 - (3) Repeat the following until $\text{Orient}(p', p'', q'') < 0$ and $\text{Orient}(p'', p', q') > 0$:
 - (a) while $(\text{Orient}(p', p'', q'') \geq 0)$ advance p'' and q'' to their successors on H'' .
 - (b) while $(\text{Orient}(p'', p', q') \leq 0)$ advance p' and q' to their predecessors on H' .
 - (4) return (p', p'') .
-

observe that there are now two tangency conditions to be satisfied, not just one). We will leave it as an exercise. Observe that the running time is $O(n)$, because with each step we spend $O(1)$ time and eliminate a point either from H' or from H'' as a candidate for the tangency points, and there are at most n points that can be so eliminated.

Running-time analysis: The asymptotic running time of the algorithm can be expressed by a recurrence. Given an input of size n , consider the time needed to perform all the parts of the procedure, ignoring the recursive calls. This includes the time to partition the point set, compute the upper tangent line, and return the final result. Clearly, each of these can be performed in $O(n)$ time, assuming any standard list representation of the hull vertices. Thus, ignoring constant factors, we can describe the running time by the following recurrence:

$$T(n) = \begin{cases} 1 & \text{if } n \leq 3 \\ n + 2T(n/2) & \text{otherwise.} \end{cases}$$

This is the same recurrence that arises in Mergesort. It is easy to show that it solves to $T(n) \in O(n \log n)$ (see any standard algorithms text).