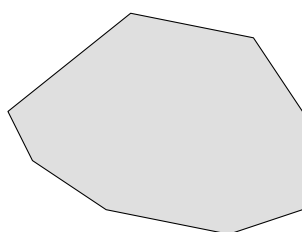



Chapter 5

Convexity and Convex Hulls

There exists an incredible variety of point sets and polygons. Among them, some have certain properties that make them “nicer” than others in some respect. For instance, look at the two polygons shown below.



(a) A convex polygon.



(b) A non-convex polygon.

Figure 5.1: *Examples of polygons: Which do you like better?*

As it is hard to argue about aesthetics, let us take a more algorithmic stance. When designing algorithms, the polygon shown on the left appears much easier to deal with than the visually and geometrically more complex polygon shown on the right. One particular property that makes the left polygon nice is that one can walk between any two vertices along a straight line without ever leaving the polygon. In fact, this statement holds true not only for vertices but for any two points within the polygon. A polygon or, more generally, a set with this property is called *convex*.

Definition 5.1. A set $P \subseteq \mathbb{R}^d$ is **convex** if $\overline{pq} \subseteq P$, for every pair $p, q \in P$.

An alternative, equivalent way to phrase convexity would be to demand that for every line $\ell \subset \mathbb{R}^d$ the intersection $\ell \cap P$ be connected. The polygon shown in Figure 5.1b is not convex because there are some pairs of points for which the connecting line segment is not completely contained within the polygon. An immediate consequence of the definition is the following

Observation 5.2. For any family $(P_i)_{i \in I}$ of convex sets, the intersection $\bigcap_{i \in I} P_i$ is convex.

Indeed there are many problems that are comparatively easy to solve for convex sets but very hard in general. We will encounter some particular instances of this phenomenon later in the course. However, not all polygons are convex and a discrete set of points is never convex, unless it consists of at most one point only. In such a case it is useful to make a given set P convex, that is, approximate P with or, rather, encompass P within a convex set $H \supseteq P$. Ideally, H differs from P as little as possible, that is, we want H to be a smallest convex set enclosing P .

At this point let us step back for a second and ask ourselves whether this wish makes sense at all: Does such a set H (always) exist? Fortunately, we are on the safe side because the whole space \mathbb{R}^d is certainly convex. It is less obvious, but we will see below that H is actually unique. Therefore it is legitimate to refer to H as **the** smallest convex set enclosing P or—shortly—the *convex hull* of P .

5.1 Convexity

In this section we will derive an algebraic characterization of convexity. Such a characterization allows to investigate convexity using the machinery from linear algebra.

Consider $P \subset \mathbb{R}^d$. From linear algebra courses you should know that the *linear hull*

$$\text{lin}(P) := \left\{ \sum_{i=1}^n \lambda_i p_i \mid n \in \mathbb{N} \wedge \forall i \in \{1, \dots, n\}: p_i \in P, \lambda_i \in \mathbb{R} \right\}$$

is the set of all *linear combinations* of P (smallest linear subspace containing P). For instance, if $P = \{p\} \subset \mathbb{R}^2 \setminus \{0\}$ then $\text{lin}(P)$ is the line through p and the origin.

A finite set P is *linearly independent* if no point of P is a linear combination of the others. In formulas, this means that $\sum_{i=1}^n \lambda_i p_i = 0$ implies that all λ_i are 0. Indeed, a nonzero λ_i would mean that p_i can be written as a linear combination of the other points: simply solve $\sum_{i=1}^n \lambda_i p_i = 0$ for p_i . Vice versa, if p_i is a linear combination of the other points, this gives us a solution to $\sum_{i=1}^n \lambda_i p_i = 0$ with $\lambda_i = 1$.

Similarly, the *affine hull*

$$\text{aff}(P) := \left\{ \sum_{i=1}^n \lambda_i p_i \mid n \in \mathbb{N} \wedge \forall i \in \{1, \dots, n\}: p_i \in P, \lambda_i \in \mathbb{R} \wedge \sum_{i=1}^n \lambda_i = 1 \right\}$$

is the set of all *affine combinations* of P (smallest affine subspace containing P). For instance, if $P = \{p, q\} \subset \mathbb{R}^2$ and $p \neq q$ then $\text{aff}(P)$ is the line through p and q .

A finite set P is *affinely independent* if no point of P is an affine combination of the others. In formulas, this means that $\sum_{i=1}^n \lambda_i p_i = 0$ and $\sum_{i=1}^n \lambda_i = 0$ implies that all λ_i are 0. As in the linear case, we can argue that a nonzero λ_i can be used to express p_i as an affine combination of the other points, and vice versa.

This immediately implies that affine and linear independence are related as follows.

Proposition 5.3. *Let $P \subseteq \mathbb{R}^d$ be a finite point set, and let $P' \subseteq \mathbb{R}^{d+1}$ be the point set obtained from P by appending to each point a new coordinate 1. For example, if $P = \{(2, 3), (0, 4)\} \subseteq \mathbb{R}^2$, then $P' = \{(2, 3, 1), (0, 4, 1)\} \subseteq \mathbb{R}^3$.*

Then the following two statements are equivalent.

(i) P is affinely independent.

(ii) P' is linearly independent.

It turns out that convexity can be described in a very similar way algebraically, which leads to the notion of *convex combinations*.

Proposition 5.4. *A set $P \subseteq \mathbb{R}^d$ is convex if and only if $\sum_{i=1}^n \lambda_i p_i \in P$, for all $n \in \mathbb{N}$, $p_1, \dots, p_n \in P$, and $\lambda_1, \dots, \lambda_n \geq 0$ with $\sum_{i=1}^n \lambda_i = 1$.*

Proof. “ \Leftarrow ”: obvious with $n = 2$.

“ \Rightarrow ”: Induction on n . For $n = 1$ the statement is trivial. For $n \geq 2$, let $p_i \in P$ and $\lambda_i \geq 0$, for $1 \leq i \leq n$, and assume $\sum_{i=1}^n \lambda_i = 1$. We may suppose that $\lambda_i > 0$, for all i . (Simply omit those points whose coefficient is zero.) We need to show that $\sum_{i=1}^n \lambda_i p_i \in P$.

Define $\lambda = \sum_{i=1}^{n-1} \lambda_i$ and for $1 \leq i \leq n-1$ set $\mu_i = \lambda_i/\lambda$. Observe that $\mu_i \geq 0$ and $\sum_{i=1}^{n-1} \mu_i = 1$. By the inductive hypothesis, $q := \sum_{i=1}^{n-1} \mu_i p_i \in P$, and thus by convexity of P also $\lambda q + (1-\lambda)p_n \in P$. We conclude by noting that $\lambda q + (1-\lambda)p_n = \lambda \sum_{i=1}^{n-1} \mu_i p_i + \lambda_n p_n = \sum_{i=1}^n \lambda_i p_i$. \square

Definition 5.5. *The convex hull $\text{conv}(P)$ of a set $P \subseteq \mathbb{R}^d$ is the intersection of all convex supersets of P .*

At first glance this definition is a bit scary: There may be a whole lot of supersets for any given P and it is not clear that taking the intersection of all of them yields something sensible to work with. However, by Observation 5.2 we know that the resulting set is convex, at least. The missing bit is provided by the following proposition, which characterizes the convex hull in terms of exactly those convex combinations that appeared in Proposition 5.4 already.

Proposition 5.6. *For any $P \subseteq \mathbb{R}^d$ we have*

$$\text{conv}(P) = \left\{ \sum_{i=1}^n \lambda_i p_i \mid n \in \mathbb{N} \wedge \sum_{i=1}^n \lambda_i = 1 \wedge \forall i \in \{1, \dots, n\} : \lambda_i \geq 0 \wedge p_i \in P \right\}.$$

The elements of the set on the right hand side are referred to as *convex combinations* of P .

Proof. “ \supseteq ”: Consider a convex set $C \supseteq P$. By Proposition 5.4 (only-if direction) the right hand side is contained in C . As C was arbitrary, the claim follows.

“ \subseteq ”: Denote the set on the right hand side by R . Clearly $R \supseteq P$. We show that R forms a convex set. Let $p = \sum_{i=1}^n \lambda_i p_i$ and $q = \sum_{i=1}^n \mu_i p_i$ be two convex combinations. (We may suppose that both p and q are expressed over the same p_i by possibly adding some terms with a coefficient of zero.)

Then for $\lambda \in [0, 1]$ we have $\lambda p + (1 - \lambda)q = \sum_{i=1}^n (\lambda \lambda_i + (1 - \lambda)\mu_i) p_i \in R$, as $\underbrace{\lambda \lambda_i}_{\geq 0} + \underbrace{(1 - \lambda)}_{\geq 0} \underbrace{\mu_i}_{\geq 0} \geq 0$, for all $1 \leq i \leq n$, and $\sum_{i=1}^n (\lambda \lambda_i + (1 - \lambda)\mu_i) = \lambda + (1 - \lambda) = 1$. \square

In linear algebra the notion of a basis in a vector space plays a fundamental role. In a similar way we want to describe convex sets using as few entities as possible, which leads to the notion of extremal points, as defined below.

Definition 5.7. *The convex hull of a finite point set $P \subset \mathbb{R}^d$ forms a convex polytope. Each $p \in P$ for which $p \notin \text{conv}(P \setminus \{p\})$ is called a vertex of $\text{conv}(P)$. A vertex of $\text{conv}(P)$ is also called an extremal point of P . A convex polytope in \mathbb{R}^2 is called a convex polygon.*

Essentially, the following proposition shows that the term vertex above is well defined.

Proposition 5.8. *A convex polytope in \mathbb{R}^d is the convex hull of its vertices.*

Proof. Let $P = \{p_1, \dots, p_n\}$ for $n \in \mathbb{N}$, such that without loss of generality p_1, \dots, p_k are the vertices of $\mathcal{P} := \text{conv}(P)$. We prove by induction on n that $\text{conv}(p_1, \dots, p_n) \subseteq \text{conv}(p_1, \dots, p_k)$. For $n = k$ the statement is trivial.

For $n > k$, the point p_n is not a vertex of \mathcal{P} and hence p_n can be expressed as a convex combination $p_n = \sum_{i=1}^{n-1} \lambda_i p_i$. Thus for any $x \in \mathcal{P}$ we can write $x = \sum_{i=1}^n \mu_i p_i = \sum_{i=1}^{n-1} \mu_i p_i + \mu_n \sum_{i=1}^{n-1} \lambda_i p_i = \sum_{i=1}^{n-1} (\mu_i + \mu_n \lambda_i) p_i$. As $\sum_{i=1}^{n-1} (\mu_i + \mu_n \lambda_i) = 1$, we conclude inductively that $x \in \text{conv}(p_1, \dots, p_{n-1}) \subseteq \text{conv}(p_1, \dots, p_k)$. \square

5.2 Classic Theorems for Convex Sets

Next we will discuss a few fundamental theorems about convex sets in \mathbb{R}^d . The proofs typically use the algebraic characterization of convexity and then employ some techniques from linear algebra.

Theorem 5.9 (Radon [8]). *Any set $P \subset \mathbb{R}^d$ of $d + 2$ points can be partitioned into two disjoint subsets P_1 and P_2 such that $\text{conv}(P_1) \cap \text{conv}(P_2) \neq \emptyset$.*

Proof. Let $P = \{p_1, \dots, p_{d+2}\}$. No more than $d + 1$ points can be affinely independent in \mathbb{R}^d . Hence suppose without loss of generality that p_{d+2} can be expressed as an affine combination of p_1, \dots, p_{d+1} , that is, there exist $\lambda_1, \dots, \lambda_{d+1} \in \mathbb{R}$ with $\sum_{i=1}^{d+1} \lambda_i = 1$ and $\sum_{i=1}^{d+1} \lambda_i p_i = p_{d+2}$. Let P_1 be the set of all points p_i for which λ_i is positive and let $P_2 = P \setminus P_1$. Then setting $\lambda_{d+2} = -1$ we can write $\sum_{p_i \in P_1} \lambda_i p_i = \sum_{p_i \in P_2} -\lambda_i p_i$,

where all coefficients on both sides are non-negative. Since $\sum_{i=1}^{i=d+2} \lambda_i = 0$ we have $s := \sum_{p_i \in P_1} \lambda_i = \sum_{p_i \in P_2} -\lambda_i$. Renormalizing by $\mu_i = \lambda_i/s$ and $\nu_i = -\lambda_i/s$ yields convex combinations $\sum_{p_i \in P_1} \mu_i p_i = \sum_{p_i \in P_2} \nu_i p_i$ that describe a common point of $\text{conv}(P_1)$ and $\text{conv}(P_2)$. \square

Theorem 5.10 (Helly). *Consider a collection $\mathcal{C} = \{C_1, \dots, C_n\}$ of $n \geq d + 1$ convex subsets of \mathbb{R}^d , such that any $d + 1$ pairwise distinct sets from \mathcal{C} have non-empty intersection. Then also the intersection $\bigcap_{i=1}^n C_i$ of all sets from \mathcal{C} is non-empty.*

Proof. Induction on n . The base case $n = d + 1$ holds by assumption. Hence suppose that $n \geq d + 2$. Consider the sets $D_i = \bigcap_{j \neq i} C_j$, for $i \in \{1, \dots, n\}$. As D_i is an intersection of $n - 1$ sets from \mathcal{C} , by the inductive hypothesis we know that $D_i \neq \emptyset$. Therefore we can find some point $p_i \in D_i$, for each $i \in \{1, \dots, n\}$. Now by Theorem 5.9 the set $P = \{p_1, \dots, p_n\}$ can be partitioned into two disjoint subsets P_1 and P_2 such that $\text{conv}(P_1) \cap \text{conv}(P_2) \neq \emptyset$. We claim that any point $p \in \text{conv}(P_1) \cap \text{conv}(P_2)$ also lies in $\bigcap_{i=1}^n C_i$, which completes the proof.

Consider some C_i , for $i \in \{1, \dots, n\}$. By construction $D_j \subseteq C_i$, for $j \neq i$. Thus p_i is the only point from P that may not be in C_i . As p_i is part of only one of P_1 or P_2 , say, of P_1 , we have $P_2 \subseteq C_i$. The convexity of C_i implies $\text{conv}(P_2) \subseteq C_i$ and, therefore, $p \in C_i$. \square

There is a nice application of Helly’s theorem showing the existence of so-called centerpoints of finite point sets. Basically, a centerpoint is one possible generalization of the median of one-dimensional sets to higher dimensions.

Definition 5.11. *Let $P \subset \mathbb{R}^d$ be a set of n points. A point $p \in \mathbb{R}^d$, not necessarily in P , is a centerpoint of P if every open halfspace that contains more than $\frac{dn}{d+1}$ points of P also contains p .*

Stated differently, every closed halfspace that contains a centerpoint also contains at least $\frac{n}{d+1}$ points of P (which is clearly equivalent to containing at least $\lceil \frac{n}{d+1} \rceil$ points). We have the following result.

Theorem 5.12. *For every set $P \subset \mathbb{R}^d$ of n points there exists a centerpoint.*

Proof. We may assume that P contains at least $d + 1$ affinely independent points (otherwise, a centerpoint can be found in a lower-dimensional affine sub-space).

Let \mathcal{A} be the family of subsets of P that are defined by the intersection of P with an open halfspace, and that contain more than $\frac{dn}{d+1}$ points. Note that since P is finite, also the number of sets in $\mathcal{A} = \{A_1, \dots, A_m\}$ is finite. Let $C_i := \text{conv}(A_i)$. If there exists a point c that is in the intersection $\bigcap_{i=1}^m C_i$, then c is contained in any open halfspace that contains more than $\frac{dn}{d+1}$ points of P and thus is a centerpoint of P . We show the existence of c by showing that any $d + 1$ elements of $\{C_1, \dots, C_m\}$ have a common point, and then apply Theorem 5.10.

For any $d + 1$ sets in \mathcal{A} , suppose that any point of P occurs in at most d of these subsets. Then in total we would have at most dn occurrences of points from P in these

subsets. However, by the choice of \mathcal{A} , each set contains more than $\frac{dn}{d+1}$ points, so the total number of occurrences is more than $(d+1)\frac{dn}{d+1} = dn$. Hence, any $d+1$ sets in \mathcal{A} have a common point, and thus $\bigcap_{i=1}^m C_i$ also contains a point c . \square

Exercise 5.13. Show that the number of points in Definition 5.11 is best possible, that is, for every n there is a set of n points in \mathbb{R}^d such that for any $p \in \mathbb{R}^d$ there is an open halfspace containing $\lfloor \frac{dn}{d+1} \rfloor$ points but not p .

Theorem 5.14 (Carathéodory [3]). For any $P \subset \mathbb{R}^d$ and $q \in \text{conv}(P)$ there exist $k \leq d+1$ points $p_1, \dots, p_k \in P$ such that $q \in \text{conv}(p_1, \dots, p_k)$.

Exercise 5.15. Prove Theorem 5.14.

Theorem 5.16 (Separation Theorem). Any two compact convex sets $C, D \subset \mathbb{R}^d$ with $C \cap D = \emptyset$ can be separated strictly by a hyperplane, that is, there exists a hyperplane h such that C and D lie in the opposite open halfspaces bounded by h .

Proof. Consider the distance function $\delta : C \times D \rightarrow \mathbb{R}$ with $(c, d) \mapsto \|c - d\|$. Since $C \times D$ is compact and δ is continuous and strictly bounded from below by 0, the function δ attains its minimum at some point $(c_0, d_0) \in C \times D$ with $\delta(c_0, d_0) > 0$. Let h be the hyperplane perpendicular to the line segment $\overline{c_0 d_0}$ that passes through the midpoint of c_0 and d_0 ; see Figure 5.2. We claim that h strictly separates C and D .




Figure 5.2: The hyperplane h strictly separates the compact convex sets C and D .

To see this, suppose there was a point, say, $c' \in C \cap h$. Then by convexity of C the whole line segment $\overline{c_0 c'}$ lies in C and some point along this segment is closer to d_0 than is c_0 , in contradiction to the choice of c_0 . The figure shown to the right depicts the situation in \mathbb{R}^2 . If, say, C has points on both sides of h , then by convexity of C it has also a point on h , but we just saw that there is no such point. Therefore, C and D must lie in different open halfspaces bounded by h . \square

The statement above is wrong for arbitrary (not necessarily compact) convex sets. Only if the separation is not required to be strict (the hyperplane may intersect the sets), such a separation always exists. However, the proof is a bit more involved (cf. Matoušek's book [7], but also check the errata on his webpage).

Exercise 5.17. *Show that the Separation Theorem does not hold in general if not both of the sets are convex.*

Exercise 5.18. *Prove or disprove:*

- a) *The convex hull of a compact subset of \mathbb{R}^d is compact.*
- b) *The convex hull of a closed subset of \mathbb{R}^d is closed.*

Altogether we obtain various equivalent definitions for the convex hull, summarized in the following theorem.

Theorem 5.19. *For a compact set $P \subset \mathbb{R}^d$ we can characterize $\text{conv}(P)$ equivalently as one of*

1. *the smallest (w. r. t. set inclusion) convex subset of \mathbb{R}^d that contains P ;*
2. *the set of all convex combinations of points from P ;*
3. *the set of all convex combinations formed by $d + 1$ or fewer points from P ;*
4. *the intersection of all convex supersets of P ;*
5. *the intersection of all closed halfspaces containing P .*

Exercise 5.20. *Prove Theorem 5.19.*

5.3 Planar Convex Hull

Although we know by now what is the convex hull of a point set, it is not yet clear how to construct it algorithmically. As a first step, we have to find a suitable representation for convex hulls. In this section we focus on the problem in \mathbb{R}^2 , where the convex hull of a finite point set forms a convex polygon. A convex polygon is easy to represent, for instance, as a sequence of its vertices in counterclockwise orientation. In higher dimensions finding a suitable representation for convex polytopes is a much more delicate task.

Problem 5.21 (Convex hull).

Input: $P = \{p_1, \dots, p_n\} \subset \mathbb{R}^2$, for some $n \in \mathbb{N}$.

Output: A sequence (q_1, \dots, q_h) of the vertices of $\text{conv}(P)$, ordered counterclockwise.

Another possible algorithmic formulation of the problem is to ignore the structure of the convex hull and just consider it as a point set.

Problem 5.22 (Extremal points).

Input: $P = \{p_1, \dots, p_n\} \subset \mathbb{R}^2$, for some $n \in \mathbb{N}$.

Output: The set $Q \subseteq P$ of the vertices of $\text{conv}(P)$.




Figure 5.3: *Convex Hull of a set of points in \mathbb{R}^2 .*

Degeneracies. A couple of further clarifications regarding the above problem definitions are in order.


First of all, for efficiency reasons an input is usually specified as a sequence of points. Do we insist that this sequence forms a set or are duplications of points allowed?

What if three points are collinear? Are all of them considered extremal? According to our definition from above, they are not and that is what we will stick to. But note that there may be cases where one wants to include all such points, nevertheless.

By the Separation Theorem, every extremal point p can be separated from the convex hull of the remaining points by a halfplane. If we take such a halfplane and translate its defining line such that it passes through p , then all points from P other than p should lie in the resulting open halfplane. In \mathbb{R}^2 it turns out convenient to work with the following “directed” reformulation.


Proposition 5.23. *A point $p \in P = \{p_1, \dots, p_n\} \subset \mathbb{R}^2$ is extremal for $P \iff$ there is a directed line g through p such that $P \setminus \{p\}$ is (strictly) to the left of g .*

The *interior angle* at a vertex v of a polygon P is the angle between the two edges of P incident to v whose corresponding angular domain lies in P° . If this angle is smaller than π , the vertex is called *convex*; if the angle is larger than π , the vertex is called *reflex*. For instance, the vertex c in the polygon depicted to the right is a convex vertex, whereas the vertex labeled r is a reflex vertex.



Exercise 5.24.

A set $S \subset \mathbb{R}^2$ is star-shaped if there exists a point $c \in S$, such that for every point $p \in S$ the line segment \overline{cp} is contained in S . A simple polygon with exactly three convex vertices is called a *pseudotriangle* (see the example shown on the right).



In the following we consider subsets of \mathbb{R}^2 . Prove or disprove:

- a) Every convex vertex of a simple polygon lies on its convex hull.
- b) Every star-shaped set is convex.
- c) Every convex set is star-shaped.
- d) The intersection of two convex sets is convex.
- e) The union of two convex sets is convex.
- f) The intersection of two star-shaped sets is star-shaped.
- g) The intersection of a convex set with a star-shaped set is star-shaped.
- h) Every triangle is a pseudotriangle.
- i) Every pseudotriangle is star-shaped.

5.4 Trivial algorithms

One can compute the extremal points using Carathéodory's Theorem as follows: Test for every point $p \in P$ whether there are $q, r, s \in P \setminus \{p\}$ such that p is inside the triangle with vertices q, r , and s . Runtime $O(n^4)$.

Another option, inspired by the Separation Theorem: test for every pair $(p, q) \in P^2$ whether all points from $P \setminus \{p, q\}$ are to the left of the directed line through p and q (or on the line segment \overline{pq}). Runtime $O(n^3)$.

Exercise 5.25. Let $P = (p_0, \dots, p_{n-1})$ be a sequence of n points in \mathbb{R}^2 . Someone claims that you can check by means of the following algorithm whether or not P describes the boundary of a convex polygon in counterclockwise order:

```
bool is_convex( $p_0, \dots, p_{n-1}$ ) {
    for  $i = 0, \dots, n - 1$ :
        if ( $p_i, p_{(i+1) \bmod n}, p_{(i+2) \bmod n}$ ) form a rightturn:
            return false;
    return true;
}
```

Disprove the claim and describe a correct algorithm to solve the problem.

Exercise 5.26. Let $P \subset \mathbb{R}^2$ be a convex polygon, given as an array $p[0] \dots p[n-1]$ of its n vertices in counterclockwise order.


- a) Describe an $O(\log(n))$ time algorithm to determine whether a point q lies inside, outside or on the boundary of P .
- b) Describe an $O(\log(n))$ time algorithm to find a (right) tangent to P from a query point q located outside P . That is, find a vertex $p[i]$, such that P is contained in the closed halfplane to the left of the oriented line $qp[i]$.

5.5 Jarvis' Wrap

We are now ready to describe a first simple algorithm to construct the convex hull. It is inspired by Proposition 5.23 and works as follows:

Find a point p_1 that is a vertex of $\text{conv}(P)$ (e.g., the one with smallest x -coordinate). “Wrap” P starting from p_1 , i.e., always find the next vertex of $\text{conv}(P)$ as the one that is rightmost with respect to the direction given by the previous two vertices.

Besides comparing x -coordinates, the only geometric primitive needed is an *orientation* test: Denote by $\text{rightturn}(p, q, r)$, for three points $p, q, r \in \mathbb{R}^2$, the predicate that is true if and only if r is (strictly) to the right of the oriented line pq .



Code for Jarvis' Wrap.

$p[0..N)$ contains a sequence of N points.
 p_{start} point with smallest x -coordinate.
 q_{next} some *other* point in $p[0..N)$.

```
int h = 0;
Point_2 q_now = p_start;
do {
    q[h] = q_now;
    h = h + 1;

    for (int i = 0; i < N; i = i + 1)
        if (rightturn_2(q_now, q_next, p[i]))
            q_next = p[i];

    q_now = q_next;
    q_next = p_start;
} while (q_now != p_start);
```

$q[0, h)$ describes a convex polygon bounding the convex hull of $p[0..N)$.

Analysis. For every output point the above algorithm spends n rightturn tests, which is $\Rightarrow O(nh)$ in total.

Theorem 5.27. [6] *Jarvis' Wrap computes the convex hull of n points in \mathbb{R}^2 using $O(nh)$ rightturn tests, where h is the number of hull vertices.*

In the worst case we have $h = n$, that is, $O(n^2)$ rightturn tests. Jarvis' Wrap has a remarkable property that is called *output sensitivity*: the runtime depends not only on the size of the input but also on the size of the output. For a huge point set it constructs the convex hull in optimal linear time, if the convex hull consists of a constant number of vertices only. Unfortunately the worst case performance of Jarvis' Wrap is suboptimal, as we will see soon.

Degeneracies. The algorithm may have to cope with various degeneracies.

- Several points have smallest x -coordinate \Rightarrow lexicographic order:

$$(p_x, p_y) < (q_x, q_y) \iff p_x < q_x \vee p_x = q_x \wedge p_y < q_y .$$

- Three or more points collinear \Rightarrow choose the point that is farthest among those that are rightmost.

Predicates. Besides the lexicographic comparison mentioned above, the Jarvis' Wrap (and most other 2D convex hull algorithms for that matter) need one more geometric predicate: the rightturn or—more generally—orientation test. The computation amounts to evaluating a polynomial of degree two, see the exercise below. We therefore say that the orientation test has *algebraic degree* two. In contrast, the lexicographic comparison has degree one only. The algebraic degree not only has a direct impact on the efficiency of a geometric algorithm (lower degree \leftrightarrow less multiplications), but also an indirect one because high degree predicates may create large intermediate results, which may lead to overflows and are much more costly to compute with exactly.

Exercise 5.28. *Prove that for three points $(p_x, p_y), (q_x, q_y), (r_x, r_y) \in \mathbb{R}^2$, the sign of the determinant*

$$\begin{vmatrix} 1 & p_x & p_y \\ 1 & q_x & q_y \\ 1 & r_x & r_y \end{vmatrix}$$

determines if r lies to the right, to the left or on the directed line through p and q .

Exercise 5.29. *The InCircle predicate is: Given three points $p, q, r \in \mathbb{R}^2$ that define a circle C and a fourth point s , is s located inside C or not? The goal of this exercise is to derive an algebraic formulation of the incircle predicate in form of a determinant, similar to the formulation of the orientation test given above in*

Exercise 5.28. To this end we employ the so-called *parabolic lifting map*, which will also play a prominent role in the next chapter of the course.


The parabolic lifting map $\ell : \mathbb{R}^2 \rightarrow \mathbb{R}^3$ is defined for a point $p = (x, y) \in \mathbb{R}^2$ by $\ell(p) = (x, y, x^2 + y^2)$. For a circle $C \subseteq \mathbb{R}^2$ of positive radius, show that the “lifted circle” $\ell(C) = \{\ell(p) \mid p \in C\}$ is contained in a unique plane $h_C \subseteq \mathbb{R}^3$. Moreover, show that a point $p \in \mathbb{R}^2$ is strictly inside (outside, respectively) of C if and only if the lifted point $\ell(p)$ is strictly below (above, respectively) h_C .

Use these insights to formulate the *InCircle* predicate for given points (p_x, p_y) , (q_x, q_y) , (r_x, r_y) , $(s_x, s_y) \in \mathbb{R}^2$ as a determinant.

5.6 Graham Scan (Successive Local Repair)

There exist many algorithms that exhibit a better worst-case runtime than Jarvis’ Wrap. Here we discuss only one of them: a particularly elegant and easy-to-implement variant of the so-called *Graham Scan* [5]. This algorithm is referred to as *Successive Local Repair* because it starts with some polygon enclosing all points and then step-by-step repairs the deficiencies of this polygon, by removing nonconvex vertices. It goes as follows:

Sort the points lexicographically to obtain a sequence p_0, \dots, p_{n-1} and build a corresponding circular sequence $p_0, \dots, p_{n-1}, \dots, p_0$ that walks around the point set in anticlockwise direction.



$p_0 p_1 p_2 p_3 p_4 p_5 p_6 p_7 p_8 p_7 p_6 p_5 p_4 p_3 p_2 p_1 p_0$

As long as there is a (consecutive) triple (p, q, r) such that r is to the right of or on the directed line \overrightarrow{pq} , remove q from the sequence.

Code for Graham Scan.

$p[0..N)$ lexicographically sorted sequence of pairwise distinct points, $N \geq 2$.

```

q[0] = p[0];
int h = 0;
// Lower convex hull (left to right):
for (int i = 1; i < N; i = i + 1) {
    while (h > 0 && !leftturn_2(q[h-1], q[h], p[i]))
        h = h - 1;

```

```

    h = h + 1;
    q[h] = p[i];
}

// Upper convex hull (right to left):
for (int i = N-2; i >= 0; i = i - 1) {
    while (!leftturn_2(q[h-1], q[h], p[i]))
        h = h - 1;
    h = h + 1;
    q[h] = p[i];
}

```

$q[0, h)$ describes a convex polygon bounding the convex hull of $p[0..N)$.

Correctness. We argue for the lower convex hull only. The argument for the upper hull is symmetric. A point p is on the lower convex hull of P if there is a rightward directed line g through p such that $P \setminus \{p\}$ is (strictly) to the left of g . A directed line is *rightward* if it forms an absolute angle of at most π with the positive x -axis. (Compare this statement with the one in Proposition 5.23.)

First, we claim that every point that the algorithm discards does not appear on the lower convex hull. A point q_h is discarded only if there exist points q_{h-1} and p_i with $q_{h-1} < q_h < p_i$ (lexicographically) so that $q_{h-1}q_hp_i$ does not form a leftturn. Thus, for every rightward directed line g through q_h at least one of q_{h-1} or p_i lies on or to the right of g . It follows that q_h is not on the lower convex hull, as claimed.

At the end of the (lower hull part of the) algorithm, in the sequence q_0, \dots, q_{h-1} every consecutive triple $q_i q_{i+1} q_{i+2}$, for $0 \leq i \leq h-3$, forms a leftturn with $q_i < q_{i+1} < q_{i+2}$. Thus, for every such triple there exists a rightward directed line g through q_{i+1} such that $P \setminus \{p\}$ is (strictly) to the left of g (for instance, take g to be perpendicular to the angular bisector of $\angle q_{i+2} q_{i+1} q_i$). It follows that every inner point of the sequence q_0, \dots, q_{h-1} is on the lower convex hull. The extreme points q_0 and q_{h-1} are the lexicographically smallest and largest, respectively, point of P , both of which are easily seen to be on the lower convex hull as well. Therefore, q_0, \dots, q_{h-1} form the lower convex hull of P , which proves the correctness of the algorithm.

Analysis.

Theorem 5.30. *The convex hull of a set $P \subset \mathbb{R}^2$ of n points can be computed using $O(n \log n)$ geometric operations.*

Proof. 1. Sorting and removal of duplicate points: $O(n \log n)$.

2. At the beginning we have a sequence of $2n - 1$ points; at the end the sequence consists of h points. Observe that for every positive orientation test, one point is discarded from the sequence for good. Therefore, we have exactly $2n - h - 1$ such

shortcuts/positive orientation tests. In addition there are at most $2n - 2$ negative tests (#iterations of the outer for loops). Altogether we have at most $4n - h - 3$ orientation tests.

In total the algorithm uses $O(n \log n)$ geometric operations. Note that the number of orientation tests is linear only, but $O(n \log n)$ lexicographic comparisons are needed. \square

5.7 Lower Bound

It is not hard to see that the runtime of Graham Scan is asymptotically optimal in the worst-case.

Theorem 5.31. $\Omega(n \log n)$ geometric operations are needed to construct the convex hull of n points in \mathbb{R}^2 (in the algebraic computation tree model).

Proof. Reduction from sorting (for which it is known that $\Omega(n \log n)$ comparisons are needed in the algebraic computation tree model). Given n real numbers x_1, \dots, x_n , construct a set $P = \{p_i \mid 1 \leq i \leq n\}$ of n points in \mathbb{R}^2 by setting $p_i = (x_i, x_i^2)$. This construction can be regarded as embedding the numbers into \mathbb{R}^2 along the x -axis and then projecting the resulting points vertically onto the unit parabola. The order in which the points appear along the lower convex hull of P corresponds to the sorted order of the x_i . Therefore, if we could construct the convex hull in $o(n \log n)$ time, we could also sort in $o(n \log n)$ time. \square

Clearly this reduction does not work for the Extremal Points problem. But using a reduction from Element Uniqueness (see Section 1.1) instead, one can show that $\Omega(n \log n)$ is also a lower bound for the number of operations needed to compute the set of extremal points only. This was first shown by Avis [1] for linear computation trees, then by Yao [9] for quadratic computation trees, and finally by Ben-Or [2] for general algebraic computation trees.

5.8 Chan's Algorithm

Given matching upper and lower bounds we may be tempted to consider the algorithmic complexity of the planar convex hull problem settled. However, this is not really the case: Recall that the lower bound is a worst case bound. For instance, the Jarvis' Wrap runs in $O(nh)$ time and thus beats the $\Omega(n \log n)$ bound in case that $h = o(\log n)$. The question remains whether one can achieve both output dependence and optimal worst case performance at the same time. Indeed, Chan [4] presented an algorithm to achieve this runtime by cleverly combining the "best of" Jarvis' Wrap and Graham Scan. Let us look at this algorithm in detail. The algorithm consists of two steps that are executed one after another.

Divide. *Input:* a set $P \subset \mathbb{R}^2$ of n points and a number $H \in \{1, \dots, n\}$.


1. Divide P into $k = \lceil n/H \rceil$ sets P_1, \dots, P_k with $|P_i| \leq H$.
2. Construct $\text{conv}(P_i)$ for all i , $1 \leq i \leq k$.

Analysis. Step 1 takes $O(n)$ time. Step 2 can be handled using Graham Scan in $O(H \log H)$ time for any single P_i , that is, $O(n \log H)$ time in total.

Conquer. *Output:* the vertices of $\text{conv}(P)$ in counterclockwise order, if $\text{conv}(P)$ has less than H vertices; otherwise, the message that $\text{conv}(P)$ has at least H vertices.

1. Find the lexicographically smallest point $p_{<}$ in P .
2. Starting from $p_{<}$ find the first H points of $\text{conv}(P)$ oriented counterclockwise (simultaneous Jarvis' Wrap on the sequences $\text{conv}(P_i)$).

Determine in every wrap step the point q_i of tangency from the current point of $\text{conv}(P)$ to $\text{conv}(P_i)$, for all $1 \leq i \leq k$. We have seen in Exercise 5.26 how to compute q_i in $O(\log |\text{conv}(P_i)|) = O(\log H)$ time. Among the k candidates q_1, \dots, q_k we find the next vertex of $\text{conv}(P)$ in $O(k)$ time.



Analysis. Step 1 takes $O(n)$ time. Step 2 consists of at most H wrap steps. Each wrap step needs $O(k \log H + k) = O(k \log H)$ time, which amounts to $O(Hk \log H) = O(n \log H)$ time for Step 2 in total.

Remark. Using a more clever search strategy instead of many tangency searches one can handle the conquer phase in $O(n)$ time, see Exercise 5.32 below. However, this is irrelevant as far as the asymptotic runtime is concerned, given that already the divide step takes $O(n \log H)$ time.

Exercise 5.32. Consider k convex polygons P_1, \dots, P_k , for some constant $k \in \mathbb{N}$, where each polygon is given as a list of its vertices in counterclockwise orientation. Show how to construct the convex hull of $P_1 \cup \dots \cup P_k$ in $O(n)$ time, where $n = \sum_{i=1}^k n_i$ and n_i is the number of vertices of P_i , for $1 \leq i \leq k$.

Searching for h . While the runtime bound for $H = h$ is exactly what we were heading for, it looks like in order to actually run the algorithm we would have to know h , which—in general—we do not. Fortunately we can circumvent this problem rather easily, by applying what is called a *doubly exponential search*. It works as follows.

Call the algorithm from above iteratively with parameter $H = \min\{2^{2^t}, n\}$, for $t = 0, \dots$, until the conquer step finds all extremal points of P (i.e., the wrap returns to its starting point).

Analysis: Let 2^{2^s} be the last parameter for which the algorithm is called. Since the previous call with $H = 2^{2^{s-1}}$ did not find all extremal points, we know that $2^{2^{s-1}} < h$, that is, $2^{s-1} < \log h$, where h is the number of extremal points of P . The total runtime is therefore at most

$$\sum_{i=0}^s cn \log 2^{2^i} = cn \sum_{i=0}^s 2^i = cn(2^{s+1} - 1) < 4cn \log h = O(n \log h),$$

for some constant $c \in \mathbb{R}$. In summary, we obtain the following theorem.

Theorem 5.33. *The convex hull of a set $P \subset \mathbb{R}^2$ of n points can be computed using $O(n \log h)$ geometric operations, where h is the number of convex hull vertices.*

Questions

17. *How is convexity defined? What is the convex hull of a set in \mathbb{R}^d ? Give at least three possible definitions and show that they are equivalent.*
18. *What is a centerpoint of a finite point set in \mathbb{R}^d ? State and prove the centerpoint theorem and the two classic theorems used in its proof (Theorem 5.12 along with Theorem 5.10 and Theorem 5.9).*
19. *What does it mean to compute the convex hull of a set of points in \mathbb{R}^2 ? Discuss input and expected output and possible degeneracies.*
20. *How can the convex hull of a set of n points in \mathbb{R}^2 be computed efficiently? Describe and analyze (incl. proofs) Jarvis' Wrap, Successive Local Repair, and Chan's Algorithm.*
21. *Is there a linear time algorithm to compute the convex hull of n points in \mathbb{R}^2 ? Prove the lower bound and define/explain the model in which it holds.*
22. *Which geometric primitive operations are used to compute the convex hull of n points in \mathbb{R}^2 ? Explain the two predicates and how to compute them.*

References

- [1] David Avis, [Comments on a lower bound for convex hull determination](#). *Inform. Process. Lett.*, 11/3, (1980), 126.
- [2] Michael Ben-Or, [Lower bounds for algebraic computation trees](#). In *Proc. 15th Annu. ACM Sympos. Theory Comput.*, pp. 80–86, 1983.
- [3] Constantin Carathéodory, [Über den Variabilitätsbereich der Fourierschen Konstanten von positiven harmonischen Funktionen](#). *Rendiconto del Circolo Matematico di Palermo*, 32, (1911), 193–217.

- [4] Timothy M. Chan, [Optimal output-sensitive convex hull algorithms in two and three dimensions](#). *Discrete Comput. Geom.*, 16/4, (1996), 361–368.
- [5] Ronald L. Graham, [An efficient algorithm for determining the convex hull of a finite planar set](#). *Inform. Process. Lett.*, 1/4, (1972), 132–133.
- [6] Ray A. Jarvis, [On the identification of the convex hull of a finite set of points in the plane](#). *Inform. Process. Lett.*, 2/1, (1973), 18–21.
- [7] Jiří Matoušek, [Lectures on discrete geometry](#), Springer-Verlag, New York, NY, 2002.
- [8] Johann Radon, [Mengen konvexer Körper, die einen gemeinsamen Punkt enthalten](#). *Math. Annalen*, 83/1–2, (1921), 113–115.
- [9] Andrew C. Yao, [A lower bound to finding convex hulls](#). *J. ACM*, 28/4, (1981), 780–787.