Chapter 5

Convexity and Convex Hulls

There is an incredible variety of point sets and polygons, but some of them are "nicer" than others in some respect. Look at the two polygons below, for instance:



(a) A convex polygon. **(b)** A non-convex polygon.

Figure 5.1: *Examples of polygons: Which one do you prefer?*

The polygon shown on the left is visually and geometrically much simpler than the one on the right. But let us take a more algorithmic stance, as aesthetics is hard to argue about. When designing algorithms, the left polygon turns out to be much easier to deal with. A particular exploitable property is that one can walk straight between any two points in it without ever leaving it. A polygon, or more generally a point set, with this property is called *convex*.

Definition 5.1. *A point set* $P \subseteq \mathbb{R}^d$ *is* **convex** *if* $\overline{pq} \subseteq P$ *for every pair* $p, q \in P$. *Equivalently, the intersection of* P *with any line is a connected segment.*

The polygon in Figure 5.1b is not convex because the line segment between some pair of points does not completely lie 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.*

Many problems are comparatively easy to solve for convex sets but very hard in general, and we will encounter some instances of this phenomenon in the course. However, many polygons are not convex, and a discrete point set is never convex (unless it contains one or no point). In such cases it is useful to approximate or encompass a given set P by a convex set $H \supseteq P$. Ideally, H should differ from P as little as possible, so we want it to be the smallest convex set encompassing P:

Definition 5.3. *The* **convex hull** $\mathrm{conv}(P)$ *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 can be infinitely many convex supersets, whose intersection might not yield something sensible to work with. But at least, (i) the intersection is well-defined, as the whole space $\mathbb{R}^d$ is certainly a convex superset which takes part in the intersection; (ii) the resulting intersection is convex due to Observation 5.2; and so (iii) the convex hull is the inclusion-wise smallest convex set containing P.

To see what it really *looks like*, we appeal to an algebraic characterization to be introduced in the next section.

5.1 Algebraic Characterizations

In this section we develop algebraic characterizations of convexity. They are indispensable tools in studying convex sets in general dimension d.

Consider $P \subseteq \mathbb{R}^d$. In linear algebra course you have learnt the notion of *linear hull* $\mathrm{lin}(P)$, which is the smallest linear subspace of $\mathbb{R}^d$ that contains P. For instance, the linear hull of $\{(1,2)\} \subset \mathbb{R}^2$ is the line through $(0,0)$ and $(1,2)$; the linear hull of $\{(1,2),(3,4)\}$ is the whole space $\mathbb{R}^2$. One can show that $\mathrm{lin}(P)$ is exactly the set of all *linear combinations* of P:

$$\mathrm{lin}(P) = \left\{ \sum_{i=1}^{n} \lambda_i p_i \;\middle|\; n \in \mathbb{N}; \quad p_i \in P,\ \lambda_i \in \mathbb{R} \text{ for } 1 \leqslant i \leqslant n \right\}.$$

A finite set $P = \{p_1, \ldots, p_N\}$ is *linearly independent* if no point in P is a linear combination of the others. Equivalently, the equation $\sum_{i=1}^{N} \lambda_i p_i = 0$ has only the trivial solution $\lambda_1 = \cdots = \lambda_N = 0$. Indeed, if some $\lambda_j \neq 0$ then p_j is a linear combination of the other points with coefficients $\{-\lambda_i/\lambda_j\}_{i \neq j}$. Vice versa, if p_j is a linear combination of the others, this gives us a non-trivial solution to the equation with $\lambda_j = -1$.

In analogue, the *affine hull* of P is the smallest affine subspace[1] of $\mathbb{R}^d$ that contains P. For instance, the affine hull of $\{(1,2),(3,4)\} \subset \mathbb{R}^2$ is the line through $(1,2)$ and $(3,4)$.

[1] An affine space is simply a linear space "shifted" by an offset. That is, adding a constant vector to all vectors in a linear space yields an affine space; conversely, subtracting a fixed vector in the affine space from all vectors sends us back to a linear space. In view of this correspondence, all concepts related to linear space can be translated directly to affine space.

One can show that $\operatorname{aff}(P)$ is exactly the set of all *affine combinations* of P.

$$\operatorname{aff}(P) = \left\{ \sum_{i=1}^{n} \lambda_i p_i \;\middle|\; n \in \mathbb{N};\quad p_i \in P,\ \lambda_i \in \mathbb{R} \text{ for } 1 \leqslant i \leqslant n;\quad \sum_{i=1}^{n} \lambda_i = 1 \right\}.$$

A finite set $P = \{p_1, \ldots, p_N\}$ is *affinely independent* if no point of P is an affine combination of the others. Equivalently, the equation system $\sum_{i=1}^{N} \lambda_i p_i = 0$, $\sum_{i=1}^{N} \lambda_i = 0$ has only the trivial solution $\lambda_1 = \cdots = \lambda_N = 0$. This equivalence can be argued as we did for linear independence. The following proposition is then immediate.

Proposition 5.4. *Let $P \subseteq \mathbb{R}^d$ be a finite point set, and obtain a point set $P' \subseteq \mathbb{R}^{d+1}$ by appending a new coordinate 1 to each point in P. For example, from $P = \{(2,3),(0,4)\} \subseteq \mathbb{R}^2$ we obtain $P' = \{(2,3,1),(0,4,1)\} \subseteq \mathbb{R}^3$. Then P is affinely independent if and only if P' is linearly independent.*

Corollary 5.5. *Any set of $d+2$ points in $\mathbb{R}^d$ is affinely dependent, as any set of $d+2$ points in $\mathbb{R}^{d+1}$ is linearly dependent.*

It turns out that convex hulls can be described algebraically in a very similar way.

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

$$\operatorname{conv}(P) = \left\{ \sum_{i=1}^{n} \lambda_i p_i \;\middle|\; n \in \mathbb{N};\quad p_i \in P,\ \lambda_i \geqslant 0 \text{ for } 1 \leqslant i \leqslant n;\quad \sum_{i=1}^{n} \lambda_i = 1 \right\}.$$

the set of all **convex combinations** *of P.*

To prove it, we need a powerful characterization of convexity.

Proposition 5.7. *A set $P \subseteq \mathbb{R}^d$ is convex if and only if it is closed under convex combination (i.e. any convex combination of P lands in P).*

Proof. "$\Leftarrow$": Convexity only requires closure under convex combination of $n = 2$ points, a special case of $n \in \mathbb{N}$.

"$\Rightarrow$": By induction on n, the number of points taking part in the convex combination. For $n = 1$ the statement is trivial. For $n \geqslant 2$, consider an arbitrary convex combination $p := \sum_{i=1}^{n} \lambda_i p_i$ where $p_i \in P$ and $\lambda_i > 0$ for $1 \leqslant i \leqslant n$, and $\sum_{i=1}^{n} \lambda_i = 1$. Here we assumed $\lambda_i > 0$ because otherwise we can just omit those points whose coefficients are zero. We need to show that $p \in P$.

Let us write

$$p = \left(\sum_{i=1}^{n-1} \lambda_i p_i \right) + \lambda_n p_n = \lambda \left(\sum_{i=1}^{n-1} \frac{\lambda_i}{\lambda} p_i \right) + (1-\lambda) p_n$$

where $\lambda := \sum_{i=1}^{n-1} \lambda_i = 1 - \lambda_n \in [0,1]$. Note that $q := \sum_{i=1}^{n-1} \frac{\lambda_i}{\lambda} p_i$ is a convex combination of $n-1$ points of P, so $q \in P$ by the inductive hypothesis. Consequently $p = \lambda q + (1-\lambda) p_n \in P$ by convexity of P. □

Proof of Proposition 5.6. Denote the set on the right hand side by R.

- $\operatorname{conv}(P) \supseteq R$: Consider an arbitrary convex superset $C \supseteq P$. By Proposition 5.7 ("$\Rightarrow$" direction), any convex combination of C (and in particular of P) is contained in C. Hence $C \supseteq R$, and it follows that $\operatorname{conv}(P) \supseteq R$.
- $\operatorname{conv}(P) \subseteq R$: Clearly R is a superset of P. We will show that R is convex, so it participates in the intersection that defines $\operatorname{conv}(P)$.

  To this end, take any two points $p, q \in R$. We may express $p =: \sum_{i=1}^{n} \lambda_i p_i$ and $q =: \sum_{i=1}^{n} \mu_i p_i$ as convex combinations of a common collection of points $p_1, \dots, p_n \in P$. This is always possible because we may take the union of their individual collections and set irrelevant coefficients to zero.

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

  Therefore $\overline{pq} \in R$, meaning that R is convex, indeed. □

In a linear space, the notion of a basis plays a fundamental role. It is a minimal description of the linear space of interest. Similarly, we want to describe convex sets using as few entities as possible, which leads to the notion of extreme points.

Definition 5.8. *The convex hull* $\operatorname{conv}(P)$ *of a finite point set* $P \subset \mathbb{R}^d$ *is called a* **convex polytope** *(or a* **convex polygon** *when* $d = 2$*). Every* $p \in P$ *such that* $p \notin \operatorname{conv}(P \setminus \{p\})$ *is called an* **extreme point** *of* P*.*

Exercise 5.9. *Show that a "convex polygon" defined above is really a "simple polygon that is convex".*

Proposition 5.10. *Any convex polytope* $\operatorname{conv}(P)$ *is the convex hull of the extreme points of* P*.*

Proof. Let $P = \{p_1, \dots, p_n\}$. Assume without loss of generality that its extreme points are $p_1, \dots, p_k$. We will prove by induction on $i = n, \dots, k$ that $\operatorname{conv}(P) = \operatorname{conv}\{p_1, \dots, p_i\}$.

For $i = n$ the statement is trivial. For $k \leqslant i < n$, we have $\operatorname{conv}(P) = \operatorname{conv}\{p_1, \dots, p_{i+1}\}$ by induction hypothesis. Since the point p_{i+1} is not extreme, it can be expressed as a convex combination $p_{i+1} = \sum_{j=1}^{i} \lambda_j p_j$. Thus any $x \in \operatorname{conv}(P)$ can be expressed as

$$x = \sum_{j=1}^{i+1} \mu_j p_j = \sum_{j=1}^{i} \mu_j p_j + \mu_{i+1} p_{i+1} = \sum_{j=1}^{i} (\mu_j + \mu_{i+1}\lambda_j) p_j.$$

Note that the coefficients are non-negative and sum up to 1, thus $x \in \operatorname{conv}\{p_1, \dots, p_i\}$. So we conclude $\operatorname{conv}(P) \subseteq \operatorname{conv}\{p_1, \dots, p_i\}$; the reverse inclusion is trivial. □

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 employ the algebraic characterization of convexity and some techniques from linear algebra.

Theorem 5.11 (Radon [8]). *Any set $P \subset \mathbb{R}^d$ of $d+2$ points can be partitioned into two disjoint subsets P^+ and P^- such that $\operatorname{conv}(P^+) \cap \operatorname{conv}(P^-) \neq \emptyset$.*

Proof. Let $P = \{p_1, \ldots, p_{d+2}\}$, which by Corollary 5.5 is affinely dependent. Hence $\sum_{i=1}^{d+2} \lambda_i p_i = 0$ and $\sum_{i=1}^{d+2} \lambda_i = 0$ for some $\lambda_1, \ldots, \lambda_{d+2} \in \mathbb{R}$ that are not all zero. In particular, there exist strictly positive and strictly negative coefficients.

Let P^+ be the set of all points p_i for which $\lambda_i > 0$, and denote $P^- := P \setminus P^+$. Then $P^+, P^- \neq \emptyset$ and $\sum_{p_i \in P^+} \lambda_i p_i = \sum_{p_i \in P^-} (-\lambda_i) p_i$. Observe that $\sum_{p_i \in P^+} \lambda_i = \sum_{p_i \in P^-} -\lambda_i =: s > 0$. So with renormalized coefficients

$$\mu_i := \begin{cases} \lambda_i / s & p_i \in P^+ \\ -\lambda_i / s & p_i \in P^- \end{cases} \geqslant 0$$

we have $\sum_{p_i \in P^+} \mu_i p_i = \sum_{p_i \in P^-} \mu_i p_i$, which describes a common point of $\operatorname{conv}(P^+)$ and $\operatorname{conv}(P^-)$. □

Theorem 5.12 (Carathéodory [3]). *For any $P \subset \mathbb{R}^d$ and $q \in \operatorname{conv}(P)$ there exist $k \leqslant d+1$ points $p_1, \ldots, p_k \in P$ such that $q \in \operatorname{conv}(p_1, \ldots, p_k)$.*

Exercise 5.13. *Prove Theorem 5.12.*

Exercise 5.14. *Given a finite point set $P \subset \mathbb{R}^d$, a point $q \in \operatorname{conv}(P)$, and a point $x \in \mathbb{R}^d$. Show that there exists a subset $P' \subseteq P$ of at most d points such that $q \in \operatorname{conv}(P' \cup \{x\})$.*

Theorem 5.15 (Helly). *Consider a collection $\mathcal{C} = \{C_1, \ldots, C_n\}$ of $n \geqslant d+1$ convex subsets of $\mathbb{R}^d$, such that any $d+1$ sets from $\mathcal{C}$ have non-empty intersection. Then $\bigcap_{i=1}^n C_i \neq \emptyset$, i.e. all sets from $\mathcal{C}$ have non-empty intersection.*

Proof. Induction on n. The base case $n = d+1$ holds by assumption. Hence suppose that $n \geqslant d+2$. Define sets $D_i = \bigcap_{j \neq i} C_j$, for $i \in \{1, \ldots, n\}$. As D_i is an intersection of $n-1$ sets from $\mathcal{C}$, by the inductive hypothesis we know that $D_i \neq \emptyset$. Hence we may take an arbitrary point $p_i \in D_i$, for each $i \in \{1, \ldots, n\}$. By Theorem 5.11 the set $\{p_1, \ldots, p_n\}$ can be partitioned into two disjoint subsets P^+ and P^- such that there exists a point $p \in \operatorname{conv}(P^+) \cap \operatorname{conv}(P^-)$. We claim that $p \in \bigcap_{i=1}^n C_i$, which completes the proof.

Fix any $i \in \{1, \ldots, n\}$ and consider C_i. By construction $p_{i'} \in D_{i'} \subseteq C_i$ for all $i' \neq i$. Suppose, say, $p_i \in P^-$, then $P^+ \subseteq \{p_{i'}\}_{i' \neq i} \subseteq C_i$. By convexity of C_i we see $\operatorname{conv}(P^+) \subseteq C_i$ and thus $p \in C_i$. The other case that $p_i \in P^+$ is symmetric. □

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 way to generalize the notion of a median to higher dimensions.

Definition 5.16. *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 containing more than $\frac{dn}{d+1}$ points of P also contains p.*

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

Theorem 5.17. *Every set $P \subset \mathbb{R}^d$ of n points has a centerpoint.*

Proof. We may assume that P contains at least $d+1$ points; otherwise, we may embed P in a lower-than-d-dimensional affine subspace and reduce d.

Define a family of subsets of P by

$$\mathcal{A} := \left\{ P \cap H \;\middle|\; H \text{ an open halfspace}, |P \cap H| > \frac{dn}{d+1} \right\}.$$

Since $|P| = n$, the number of subsets in $\mathcal{A} =: \{A_1, \ldots, A_m\}$ is also finite. For each $1 \leqslant i \leqslant m$, we denote $C_i := \operatorname{conv}(A_i)$ which, due to convexity, is contained in the open halfspaces that define A_i.

Suppose there is a point $c \in \bigcap_{i=1}^m C_i$, then c is also contained in every open halfspace $H : |P \cap H| > \frac{dn}{d+1}$ and thus is a centerpoint. So it suffices to show the existence of c. To this end, we will prove that any $d+1$ sets in $\mathcal{A}$ have a common point; so do any $d+1$ sets among $C_1, \ldots, C_m$. The claim then follows via Theorem 5.15.

For any $d+1$ sets in $\mathcal{A}$, each set by definition contains more than $\frac{dn}{d+1}$ points of P, so the total number of point occurrences is more than $(d+1)\frac{dn}{d+1} = dn$. Therefore, there exists a point $p \in P$ that occurs more than d times, that is, in all $d+1$ sets. This completes the proof. □

Exercise 5.18. *Show that the number of points in Definition 5.16 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 5.19 (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 $\delta : C \times D \to \mathbb{R}$ with $(c, d) \mapsto \|c - d\|$. Since $C \times D$ is compact and δ is continuous, the function δ attains its minimum at some point $(c_0, d_0) \in C \times D$. Note that $\delta(c_0, d_0) > 0$ because $C \cap D = \emptyset$. Let h be the hyperplane perpendicular to the line segment $\overline{c_0 d_0}$ and passing through its midpoint; see Figure 5.2. We claim that h strictly separates C and D.

To see this, suppose first that that there was a point $c' \in C \cap h$, say. Then by convexity of C we have $\overline{c_0 c'} \subseteq C$. But some point along this segment is closer to d_0 than is c_0, in contradiction to the choice of c_0. Suppose, then, that C has points on both

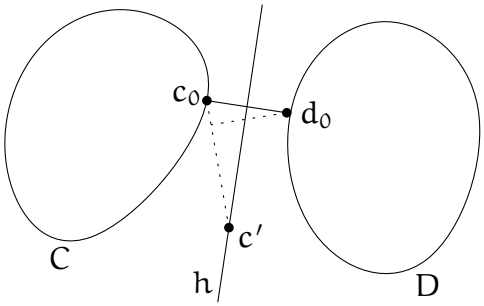

Figure 5.2: *The hyperplane* h *strictly separates the compact convex sets* C *and* D.

sides of h. Then by convexity of C it has also a point on h, but we just saw that it is impossible. The argument for D is symmetric. Therefore, C and D must lie in opposite open halfspaces bounded by h. □

The statement above is wrong for arbitrary (not necessarily compact) convex sets. Only if we allow non-strict separation (i.e. the hyperplane may intersect both sets), can we guarantee such a separation. However, the proof is a bit more involved (cf. Matoušek's book [7], but also check the errata on his webpage).

Exercise 5.20. *Show that the Separation Theorem does not hold in general if not both of the sets are convex.*

Exercise 5.21. *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 5.22. *For a compact set* $P \subset \mathbb{R}^d$ *we can characterize* $\operatorname{conv}(P)$ *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 5.23. *Prove Theorem 5.22.*

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 5.24 (Convex hull).

Input: $P = \{p_0, \ldots, p_{n-1}\} \subset \mathbb{R}^2$, for some $n \in \mathbb{N}$.

Output: A sequence $(q_0, \ldots, q_{h-1})$ of the vertices of $\mathrm{conv}(P)$, ordered counterclockwise.

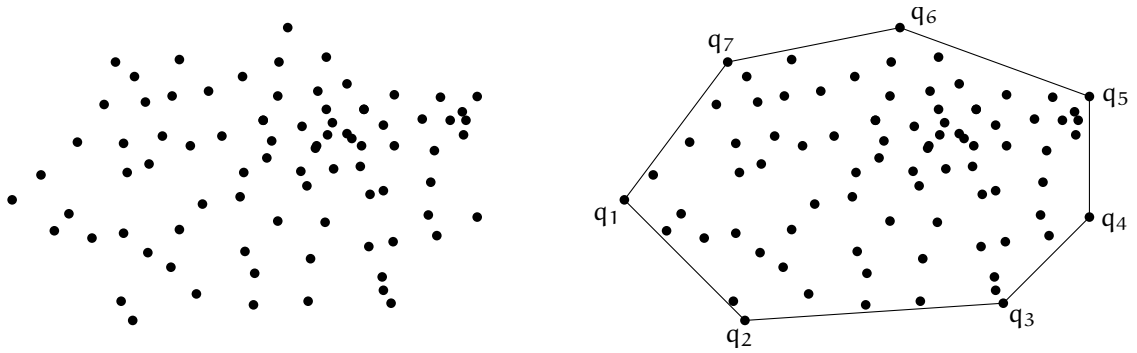

(a) Input. **(b)** Output.

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

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 5.25 (Extreme points).

Input: $P = \{p_0, \ldots, p_{n-1}\} \subset \mathbb{R}^2$, for some $n \in \mathbb{N}$.

Output: The set of vertices of $\mathrm{conv}(P)$.

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 extreme? They are not, according to our definition from above; and that is what we will stick to. But there may be situations where one wants to include these points nevertheless.

By the Separation Theorem, every extreme point p can be separated from the convex hull of the remaining points by a line. If we translate the line so that it passes through p, then every point in P other than p shall strictly lie in one side of it. In $\mathbb{R}^2$ it turns out convenient to work with the following "directed" reformulation.

Proposition 5.26. *Let $P \subset \mathbb{R}^2$ be a finite point set. A point $p \in P$ is extreme for P $\iff$ there is a directed line ℓ through p such that $P \setminus \{p\}$ is (strictly) to the left of ℓ.*

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*. In the polygon depicted on the right, the vertex c is convex whereas the vertex r is reflex.

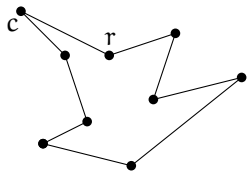

Exercise 5.27.

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).*

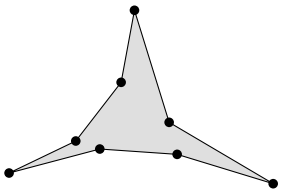

In the following we consider subsets of $\mathbb{R}^2$. Prove or disprove:

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 extreme 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 qrs. 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 $\overrightarrow{pq}$ (or on the line segment $\overline{pq}$). Runtime $O(n^3)$.

Exercise 5.28. *Let* $P = (p_0, \ldots, p_{n-1})$ *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 IsConvex(p_0, ..., p_{n-1}) {
    for i = 0, ..., n − 1:
        if (p_i, p_{(i+1) mod n}, p_{(i+2) mod n}) form a rightturn:
            return false;
    return true;
}
```

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

Exercise 5.29. *Let* $P \subset \mathbb{R}^2$ *be a convex polygon, given as an array* $p[0 \ldots n-1]$ *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 5.26 and works as follows:

> Find a vertex q_0 of $\text{conv}(P)$ (e.g., the point in P with smallest x-coordinate).
> "Wrap" P starting from q_0, i.e., always find the next vertex q_i of $\text{conv}(P)$ as the rightmost point with respect to the directed line $\overrightarrow{q_{i-2}q_{i-1}}$.

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

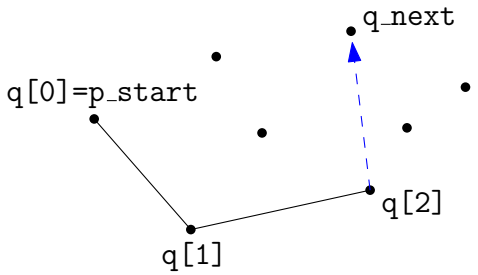

Code for Jarvis' Wrap.

`p[0..n-1]` contains a sequence of n points.
`p_start` is the point with smallest x-coordinate.
`q_next` is some *other* point in `p[0..n-1]`.

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

  for (int i = 0; i < n; i = i + 1)
    if (rightturn(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-1]` describes a convex polygon bounding the convex hull of `p[0..n-1]`.

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

Theorem 5.30. [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 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 whose convex hull consists of a constant number of vertices only, the algorithm constructs the convex hull in optimal linear time. But the worst case performance of Jarvis' Wrap is suboptimal, as we will see soon.

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

- Several points have smallest x-coordinate $\Rightarrow$ sort by lexicographical order:

$$(x_p, y_p) < (x_q, y_q) \iff (x_p < x_q) \vee (x_p = x_q \wedge y_p < y_q)\,.$$

- Three or more points collinear, so potentially multiple points are rightmost $\Rightarrow$ choose the farthest one.

Predicates. As mentioned above, the Jarvis' Wrap (and most other 2D convex hull algorithms) need the `rightturn` predicate, or more generally, orientation tests. The `rightturn` computation amounts to evaluating a polynomial of degree two, see the exercise below. We therefore say that it has *algebraic degree* two. In contrast, the lexicographic comparison has degree one only. Higher algebraic degree not only means more time-consuming multiplications, but also creates large intermediate results which may lead to overflows as well as challenges for storage and exact computation.

Exercise 5.31. *Prove that for three points* $(x_p, y_p), (x_q, y_q), (x_r, y_r) \in \mathbb{R}^2$*, the sign of the determinant*

$$\begin{vmatrix} 1 & x_p & y_p \\ 1 & x_q & y_q \\ 1 & x_r & y_r \end{vmatrix}$$

determines if r *lies to the right, to the left or on the directed line* $\overrightarrow{pq}$*.*

Exercise 5.32. *The* `InCircle` *predicate: Given four points* $p, q, r, s \in \mathbb{R}^2$*, is* s *located inside the circle defined by* p, q, r*? The goal of this exercise is to derive an algebraic formulation of this predicate as a determinant, similar to the* `rightturn` *predicate in Exercise 5.31. To this end we employ the so-called parabolic lifting map, which will also play a prominent role in later chapters.*

The parabolic lifting map $\ell : \mathbb{R}^2 \to \mathbb{R}^3$ *lifts a point* $p = (x, y) \in \mathbb{R}^2$ *to* $\ell(p) = (x, y, x^2 + y^2) \in \mathbb{R}^3$*. 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* (x_p, y_p)*,* (x_q, y_q)*,* (x_r, y_r)*,* $(x_s, y_s) \in \mathbb{R}^2$ *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 (possibly non-convex) polygon enclosing all the points and then step-by-step repairs the deficiencies by removing reflex vertices. It goes as follows:

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

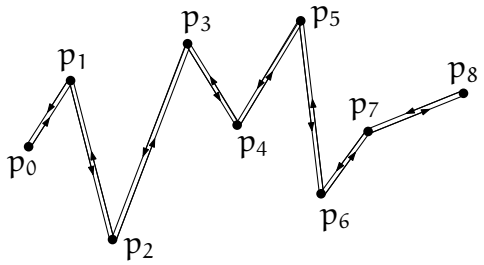

$p_0, p_1, \ldots, p_7, p_8, p_7, \ldots, p_1, p_0$

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-1]` is a lexicographically sorted sequence of $n \geqslant 2$ distinct points.

```
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(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(q[h-1], q[h], p[i]))
    h = h - 1;
  h = h + 1;
  q[h] = p[i];
}
```

`q[0..h-1]` describes a convex polygon bounding the convex hull of `p[0..n-1]`.

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 5.26.)

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.

Upon finishing the construction of lower hull, in the sequence $q_0, \ldots, q_{h-1}$ every consecutive triple $q_iq_{i+1}q_{i+2}$, for $0 \leqslant i \leqslant h-3$, 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 $q_0, \ldots, q_{h-1}$ is on the lower convex hull. The extreme points q_0 and q_{h-1} are the lexicographically smallest and largest point of P, respectively, both of which are easily seen to be on the lower convex hull as well. Therefore, $q_0, \ldots, q_{h-1}$ form the lower convex hull of P, which proves the correctness of the algorithm.

Analysis.

Theorem 5.33. *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 "false" `leftturn`, one point is discarded from the sequence for ever. Therefore, we have exactly $2n-h-1$ such tests. In addition there are at most $2n-2$ "true" `leftturn`, as bounded by the number of iterations of the outer `for` loop. Altogether we have at most $4n-h-3$ tests.

In total the algorithm uses $O(n \log n)$ geometric operations. Note that the number of `leftturn` tests is linear only, whereas we need worst-case $\Theta(n \log n)$ lexicographic comparisons which dominates the runtime. □

5.7 Lower Bound

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

Theorem 5.34. *$\Omega(n \log n)$ 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 the sorting problem, for which $\Omega(n \log n)$ comparisons are needed in the algebraic computation tree model. Given n real numbers $x_1, \ldots, x_n$, we construct a point set $P = \{(x_i, x_i^2) \mid 1 \leqslant i \leqslant n\} \subseteq \mathbb{R}^2$. This construction can be regarded as embedding the numbers into $\mathbb{R}^2$ along the x-axis and then lifting them vertically onto the unit parabola. The counterclockwise order in which the points appear along the lower convex hull of P corresponds to the sorted order of the x_i's. Therefore, if we could construct the convex hull in $o(n \log n)$ time, then we could also sort in $o(n \log n)$ time, a contradiction. □

Clearly this reduction does not work for the Extreme Points problem. But using a reduction from Element Uniqueness (see Section 1.1) instead, one can show that $\Omega(n \log n)$ operations is also needed for computing merely the set of extreme points. 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, there are fine-grained structures to be discovered: 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 whenever $h = o(\log n)$. The question remains whether one can achieve both output sensitivity and optimal worst case performance at the same time. Indeed, Chan [4] presented an algorithm to achieve this by cleverly combining the best of Jarvis' Wrap and Graham Scan. Let us look at this algorithm in detail. It first guesses an upper bound H for the number of vertices h. Then it proceeds in two phases that are executed one after another.

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

1. Divide P into $k = \lceil n/H \rceil$ sets $P_1, \ldots, P_k$ with $|P_i| \leqslant H$.
2. Construct $\text{conv}(P_i)$ using Graham Scan for $i \in \{1, \ldots, k\}$.

Analysis. Step 1 takes $O(n)$ time. Step 2 can be handled in $O(H \log H)$ time for each P_i, hence $O(kH \log H) = O(n \log H)$ time in total.

Conquer. *Output:* the first H vertices of $\text{conv}(P)$ in counterclockwise order.

1. Find the lexicographically smallest point $p_<$ in P.
2. Starting from $p_<$, find the first H vertices of $\text{conv}(P)$ in counterclockwise order by Jarvis' Wrap on the convex polygons $\text{conv}(P_1), \ldots, \text{conv}(P_k)$. Specifically, in each wrap step, determine for every i the right tangent t_i to $\text{conv}(P_i)$ from the current vertex (see Exercise 5.29 for definition, and the figure below for illustration). Select our next vertex among the k candidates $t_1, \ldots, t_k$ such that it is rightmost with respect to the direction of the last two vertices.

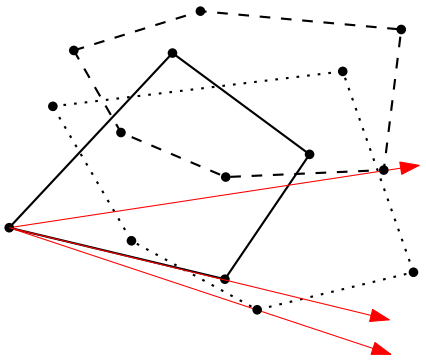

Analysis. Step 1 takes $O(n)$ time. Step 2 consists of at most H wrap steps. Each wrap step needs $O(k \log H)$ time for finding the right tangents (see Exercise 5.29) and $O(k)$ time for selecting the rightmost candidate. That amounts to $O(Hk \log H) = O(n \log H)$ time in total.

Remark. Using a more clever strategy instead of many tangency searches one can handle the conquer phase in $O(n)$ time, see Exercise 5.35 below. However, this is irrelevant as far as the asymptotic runtime is concerned, since already the divide phase takes $O(n \log H)$ time.

Exercise 5.35. *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 order. 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 $1 \leqslant i \leqslant k$.*

Searching for h. While the runtime bound for $H \approx h$ is exactly what we were heading for, we still need a means to estimate h closely, whose exact value is unknown in general. Fortunately we can address this problem rather easily, by applying what is called a *doubly exponential search*. It works as follows.

> Try the algorithm from above iteratively with parameter $H = \min\{2^{2^t}, n\}$, for $t = 0, 1 \dots$ until the conquer phase finds all vertices of $\operatorname{conv}(P)$ (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 trial with $H = 2^{2^{s-1}}$ did not find all vertices, we know that $2^{2^{s-1}} < h$, namely $2^{s-1} < \log h$, where h is the actual number of vertices of $\operatorname{conv}(P)$. The total runtime is therefore at most

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

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

Theorem 5.36. *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 are the three classic results about convex sets in* $\mathbb{R}^d$*?* State the three theorems (Theorems 5.11, 5.12 and 5.15) and prove one of them.
19. **(This topic was not covered in HS26. Thus, the question will not be asked in the exam.)** *What is a centerpoint of a finite point set in* $\mathbb{R}^d$*?* State and prove the centerpoint theorem (Theorem 5.17).
20. *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.
21. *How can the convex hull of a set of* n *points in* $\mathbb{R}^2$ *be computed efficiently?* Describe and analyze (including proofs) Jarvis' Wrap, Graham Scan, and Chan's Algorithm.
22. *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.
23. *Which geometric predicates 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, 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.