

Chapter 5

Delaunay Triangulations

In Chapter 3 we have discussed triangulations of simple polygons. A triangulation nicely partitions a polygon into triangles, which allows, for instance, to easily compute the area or a guarding of the polygon. Another typical application scenario is to use a triangulation T for interpolation: Suppose a function f is defined on the vertices of the polygon P , and we want to extend it “reasonably” and continuously to P° . Then for a point $p \in P^\circ$ find a triangle t of T that contains p . As p can be written as a convex combination $\sum_{i=1}^3 \lambda_i v_i$ of the vertices v_1, v_2, v_3 of t , we just use the same coefficients to obtain an interpolation $f(p) := \sum_{i=1}^3 \lambda_i f(v_i)$ of the function values.

If triangulations are a useful tool when working with polygons, they might also turn out useful to deal with other geometric objects, for instance, point sets. But what could be a triangulation of a point set? Polygons have a clearly defined interior, which naturally lends itself to be covered by smaller polygons such as triangles. A point set does not have an interior, except . . . Here the notion of convex hull comes handy, because it allows us to treat a point set as a convex polygon. Actually, not really a convex polygon, because points in the interior of the convex hull should not be ignored completely. But one way to think of a point set is as a convex polygon—its convex hull—possibly with some holes—which are points—in its interior. A triangulation should then partition the convex hull while respecting the points in the interior, as shown in the example in Figure 5.1b.




Figure 5.1: *Examples of (non-)triangulations.*

In contrast, the example depicted in Figure 5.1c nicely subdivides the convex hull

but should not be regarded a triangulation: Two points in the interior are not respected but simply swallowed by a large triangle.

This interpretation directly leads to the following adaption of Definition 3.7.

Definition 5.1 *A triangulation of a finite point set $P \subset \mathbb{R}^2$ is a collection \mathcal{T} of triangles, such that*

$$(1) \operatorname{conv}(P) = \bigcup_{T \in \mathcal{T}} T;$$

$$(2) P = \bigcup_{T \in \mathcal{T}} V(T); \text{ and}$$

(3) *for every distinct pair $T, U \in \mathcal{T}$, the intersection $T \cap U$ is either a common vertex, or a common edge, or empty.*

Just as for polygons, triangulations are universally available for point sets, meaning that (almost) every point set admits at least one.

Proposition 5.2 *Every set $P \subseteq \mathbb{R}^2$ of $n \geq 3$ points has a triangulation, unless all points in P are collinear.*

Proof. In order to construct a triangulation for P , consider the lexicographically sorted sequence p_1, \dots, p_n of points in P . Let m be minimal such that p_1, \dots, p_m are not collinear. We triangulate p_1, \dots, p_m by connecting p_m to all of p_1, \dots, p_{m-1} (which *are* on a common line), see Figure 5.2a.



Figure 5.2: *Constructing the scan triangulation of P .*

Then we add p_{m+1}, \dots, p_n . When adding p_i , for $i > m$, we connect p_i with all vertices of $C_{i-1} := \operatorname{conv}(\{p_1, \dots, p_{i-1}\})$ that it “sees”, that is, every vertex v of C_{i-1} for which $\overline{p_i v} \cap C_{i-1} = \{v\}$. In particular, among these vertices are the two points of tangency from p_i to C_{i-1} , which shows that we always add triangles (Figure 5.2b) whose union after each step covers C_i . \square

The triangulation that is constructed in Proposition 5.2 is called a *scan triangulation*. Such a triangulation (Figure 5.3a (left) shows a larger example) is usually “ugly”, though, since it tends to have many long and skinny triangles. This is not just an aesthetic deficit. Having long and skinny triangles means that the vertices of a triangle tend to be spread out far from each other. You can probably imagine that such a behavior is undesirable,

for instance, in the context of interpolation. In contrast, the *Delaunay triangulation* of the same point set (Figure 5.3b) looks much nicer, and we will discuss in the next section how to get this triangulation.




Figure 5.3: Two triangulations of the same set of 50 points.

Exercise 5.3 Describe an $O(n \log n)$ time algorithm to construct a scan triangulation for a set of n points in \mathbb{R}^2 .

On another note, if you look closely into the SLR-algorithm to compute planar convex hull that was discussed in Chapter 4, then you will realize that we also could have used this algorithm in the proof of Proposition 5.2. Whenever a point is discarded during SLR, a triangle is added to the polygon that eventually becomes the convex hull.

In view of the preceding chapter, we may regard a triangulation as a plane graph: the vertices are the points in P and there is an edge between two points $p \neq q$, if and only if there is a triangle with vertices p and q . Therefore we can use Euler's formula to determine the number of edges in a triangulation.

Lemma 5.4 Any triangulation of a set $P \subset \mathbb{R}^2$ of n points has exactly $3n - h - 3$ edges, where h is the number of points from P on $\partial \text{conv}(P)$.

Proof. Consider a triangulation T of P and denote by E the set of edges and by F the set of faces of T . We count the number of edge-face incidences in two ways. Denote $\mathcal{J} = \{(e, f) \in E \times F : e \subset \partial f\}$.

On the one hand, every edge is incident to exactly two faces and therefore $|\mathcal{J}| = 2|E|$. On the other hand, every bounded face of T is a triangle and the unbounded face has h edges on its boundary. Therefore, $|\mathcal{J}| = 3(|F| - 1) + h$.

Together we obtain $3|F| = 2|E| - h + 3$. Using Euler's formula ($3n - 3|E| + 3|F| = 6$) we conclude that $3n - |E| - h + 3 = 6$ and so $|E| = 3n - h - 3$. \square

In graph theory, the term “triangulation” is sometimes used as a synonym for “maximal planar”. But geometric triangulations are different, they are maximal planar in the sense that no straight-line edge can be added without sacrificing planarity.

Corollary 5.5 *A triangulation of a set $P \subset \mathbb{R}^2$ of n points is maximal planar, if and only if $\text{conv}(P)$ is a triangle.*

Proof. Combine Corollary 2.5 and Lemma 5.4. \square

Exercise 5.6 *Find for every $n \geq 3$ a simple polygon P with n vertices such that P has exactly one triangulation. P should be in general position, meaning that no three vertices are collinear.*

Exercise 5.7 *Show that every set of $n \geq 5$ points in general position (no three points are collinear) has at least two different triangulations.*

Hint: Show first that every set of five points in general position contains a convex 4-hole, that is, a subset of four points that span a convex quadrilateral that does not contain the fifth point.

5.1 The Empty Circle Property

We will now move on to study the ominous and supposedly nice Delaunay triangulations mentioned above. They are defined in terms of an empty circumcircle property for triangles. The *circumcircle* of a triangle is the unique circle passing through the three vertices of the triangle, see Figure 5.4.




Figure 5.4: *Circumcircle of a triangle.*

Definition 5.8 *A triangulation of a finite point set $P \subset \mathbb{R}^2$ is called a Delaunay triangulation, if the circumcircle of every triangle is empty, that is, there is no point from P in its interior.*

Consider the example depicted in Figure 5.5. It shows a Delaunay triangulation of a set of six points: The circumcircles of all five triangles are empty (we also say that the triangles satisfy the empty circle property). The dashed circle is not empty, but that is fine, since it is not a circumcircle of any triangle.





Figure 5.5: All triangles satisfy the empty circle property.

It is instructive to look at the case of four points in convex position. Obviously, there are two possible triangulations, but in general, only one of them will be Delaunay, see Figure 5.6a and 5.6b. If the four points are on a common circle, though, this circle is empty; at the same time it is the circumcircle of *all* possible triangles; therefore, both triangulations of the point set are Delaunay, see Figure 5.6c.



(a) Delaunay triangulation. (b) Non-Delaunay triangulation. (c) Two Delaunay triangulations.

Figure 5.6: Triangulations of four points in convex position.

Proposition 5.9 *Given a set $P \subset \mathbb{R}^2$ of four points that are in convex position but not cocircular. Then P has exactly one Delaunay triangulation.*

Proof. Consider a convex polygon $P = pqrs$. There are two triangulation of P : a triangulation \mathcal{T}_1 using the edge pr and a triangulation \mathcal{T}_2 using the edge qs .

Consider the family \mathcal{C}_1 of circles through pr , which contains the circumcircles $C_1 = pqr$ and $C'_1 = rsp$ of the triangles in \mathcal{T}_1 . By assumption s is not on C_1 . If s is outside of C_1 , then q is outside of C'_1 : Consider the process of continuously moving from C_1 to C'_1 in \mathcal{C}_1 (Figure 5.7a); the point q is “left behind” immediately when going beyond C_1 and only the final circle C'_1 “grabs” the point s .




Figure 5.7: Circumcircles and containment for triangulations of four points.

Similarly, consider the family \mathcal{C}_2 of circles through pq , which contains the circumcircles $C_1 = pqr$ and $C_2 = spq$, the latter belonging to a triangle in \mathcal{T}_2 . As s is outside of C_1 , it follows that r is inside C_2 : Consider the process of continuously moving from C_1 to C_2 in \mathcal{C}_2 (Figure 5.7b); the point r is on C_1 and remains within the circle all the way up to C_2 . This shows that \mathcal{T}_1 is Delaunay, whereas \mathcal{T}_2 is not.

The case that s is located inside C_1 is symmetric: just cyclically shift the roles of $pqrs$ to $qrsp$. \square

5.2 The Lawson Flip algorithm

It is not clear yet that every point set actually has a Delaunay triangulation (given that not all points are on a common line). In this and the next two sections, we will prove that this is the case. The proof is algorithmic. Here is the *Lawson flip algorithm* for a set P of n points.

1. Compute some triangulation of P (for example, the scan triangulation).
2. While there exists a subtriangulation of four points in convex position that is not Delaunay (like in Figure 5.6b), replace this subtriangulation by the other triangulation of the four points (Figure 5.6a).

We call the replacement operation in the second step a (*Lawson*) *flip*.

Theorem 5.10 *Let $P \subseteq \mathbb{R}^2$ be a set of n points, equipped with some triangulation \mathcal{T} . The Lawson flip algorithm terminates after at most $\binom{n}{2} = O(n^2)$ flips, and the resulting triangulation \mathcal{D} is a Delaunay triangulation of P .*

We will prove Theorem 5.10 in two steps: First we show that the program described above always terminates and, therefore, is an algorithm, indeed (Section 5.3). Then we show that the algorithm does what it claims to do, namely the result is a Delaunay triangulation (Section 5.4).

5.3 Termination of the Lawson Flip Algorithm: The Lifting Map

In order to prove Theorem 5.10, we invoke the (parabolic) *lifting map*. This is the following: given a point $p = (x, y) \in \mathbb{R}^2$, its *lifting* $\ell(p)$ is the point

$$\ell(p) = (x, y, x^2 + y^2) \in \mathbb{R}^3.$$

Geometrically, ℓ “lifts” the point vertically up until it lies on the *unit paraboloid*

$$\{(x, y, z) \mid z = x^2 + y^2\} \subseteq \mathbb{R}^3,$$

see Figure 5.8a.




Figure 5.8: *The lifting map: circles map to planes.*

Recall the following important property of the lifting map that we proved in Exercise 4.25. It is illustrated in Figure 5.8b.

Lemma 5.11 *Let $C \subseteq \mathbb{R}^2$ be a circle of positive radius. The “lifted circle” $\ell(C) = \{\ell(p) \mid p \in C\}$ is contained in a unique plane $h_C \subseteq \mathbb{R}^3$. Moreover, 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 .*

Using the lifting map, we can now prove Theorem 5.10. Let us fix the point set P for this and the next section. First, we need to argue that the algorithm indeed terminates (if you think about it a little, this is not obvious). So let us interpret a flip operation in the lifted picture. The flip involves four points in convex position in \mathbb{R}^2 , and their lifted images form a tetrahedron in \mathbb{R}^3 (think about why this tetrahedron cannot be “flat”).

The tetrahedron is made up of four triangles; when you look at it from the top, you see two of the triangles, and when you look from the bottom, you see the other two. In fact, what you see from the top and the bottom are the lifted images of the two possible triangulations of the four-point set in \mathbb{R}^2 that is involved in the flip.

Here is the crucial fact that follows from Lemma 5.11: The two top triangles come from the non-Delaunay triangulation before the flip, see Figure 5.9a. The reason is that both top triangles have the respective fourth point below them, meaning that in \mathbb{R}^2 , the circumcircles of these triangles contain the respective fourth point—the empty circle property is violated. In contrast, the bottom two triangles come from the Delaunay triangulation of the four points: they both have the respective fourth point above them, meaning that in \mathbb{R}^2 , the circumcircles of the triangles do not contain the respective fourth point, see Figure 5.9b.




Figure 5.9: *Lawson flip: the height of the surface of lifted triangles decreases.*

In the lifted picture, a Lawson flip can therefore be interpreted as an operation that replaces the top two triangles of a tetrahedron by the bottom two ones. If we consider the lifted image of the current triangulation, we therefore have a surface in \mathbb{R}^3 whose pointwise height can only decrease through Lawson flips. In particular, once an edge has been flipped, this edge will be strictly above the resulting surface and can therefore never be flipped a second time. Since n points can span at most $\binom{n}{2}$ edges, the bound on the number of flips follows.

5.4 Correctness of the Lawson Flip Algorithm

It remains to show that the triangulation of P that we get upon termination of the Lawson flip algorithm is indeed a Delaunay triangulation. Here is a first observation telling us that the triangulation is “locally Delaunay”.

Observation 5.12 *Let Δ, Δ' be two adjacent triangles in the triangulation \mathcal{D} that results from the Lawson flip algorithm. Then the circumcircle of Δ does not have any vertex of Δ' in its interior, and vice versa.*

If the two triangles together form a convex quadrilateral, this follows from the fact that the Lawson flip algorithm did not flip the common edge of Δ and Δ' . If the four vertices are not in convex position, this is basic geometry: given a triangle Δ , its circumcircle C can only contain points of $C \setminus \Delta$ that form a convex quadrilateral with the vertices of Δ .

Now we show that the triangulation is also “globally Delaunay”.

Proposition 5.13 *The triangulation \mathcal{D} that results from the Lawson flip algorithm is a Delaunay triangulation.*

Proof. Suppose for contradiction that there is some triangle $\Delta \in \mathcal{D}$ and some point $p \in P$ strictly inside the circumcircle C of Δ . Among all such pairs (Δ, p) , we choose one for which we the distance of p to Δ is minimal. Note that this distance is positive since \mathcal{D} is a triangulation of P . The situation is as depicted in Figure 5.10a.




Figure 5.10: *Correctness of the Lawson flip algorithm.*

Now consider the edge e of Δ that is facing p . There must be another triangle Δ' in \mathcal{D} that is incident to the edge e . By the local Delaunay property of \mathcal{D} , the third vertex q of Δ' is on or outside of C , see Figure 5.10b. But then the circumcircle C' of Δ' contains the whole portion of C on p 's side of e , hence it also contains p ; moreover, p is closer to

Δ' than to Δ (Figure 5.10c). But this is a contradiction to our choice of Δ and p . Hence there was no (Δ, p) , and \mathcal{D} is a Delaunay triangulation. \square

Exercise 5.14 *The Euclidean minimum spanning tree (EMST) of a finite point set $P \subset \mathbb{R}^2$ is a spanning tree for which the sum of the edge lengths is minimum (among all spanning trees of P). Show:*

- a) *Every EMST of P is a plane graph.*
- b) *Every EMST of P contains a closest pair, i.e., an edge between two points $p, q \in P$ that have minimum distance to each other among all point pairs in $\binom{P}{2}$.*
- c) *Every Delaunay Triangulation of P contains an EMST of P .*

5.5 The Delaunay Graph

Despite the fact that a point set may have more than one Delaunay triangulation, there are certain edges that are present in every Delaunay triangulation, for instance, the edges of the convex hull.

Definition 5.15 *The Delaunay graph of $P \subseteq \mathbb{R}^2$ consists of all line segments \overline{pq} , for $p, q \in P$, that are contained in every Delaunay triangulation of P .*

The following characterizes the edges of the Delaunay graph.

Lemma 5.16 *The segment \overline{pq} , for $p, q \in P$, is in the Delaunay graph of P if and only if there exists a circle through p and q that has p and q on its boundary and all other points of P are strictly outside.*

Proof. “ \Rightarrow ”: Let pq be an edge in the Delaunay graph of P , and let \mathcal{D} be a Delaunay triangulation of P . Then there exists a triangle $\Delta = pqr$ in \mathcal{D} , whose circumcircle C does not contain any point from P in its interior.

If there is a point s on ∂C such that \overline{rs} intersects \overline{pq} , then let $\Delta' = pqt$ denote the other ($\neq \Delta$) triangle in \mathcal{D} that is incident to pq (Figure 5.11a). Flipping the edge pq to rt yields another Delaunay triangulation of P that does not contain the edge pq , in contradiction to pq being an edge in the Delaunay graph of P . Therefore, there is no such point s .

Otherwise we can slightly change the circle C by moving away from r while keeping p and q on the circle. As P is a finite point set, we can do such a modification without catching another point from P with the circle. In this way we obtain a circle C' through p and q such that all other points from P are strictly outside C' (Figure 5.12b).

“ \Leftarrow ”: Let \mathcal{D} be a Delaunay triangulation of P . If \overline{pq} is not an edge of \mathcal{D} , there must be another edge of \mathcal{D} that crosses \overline{pq} (otherwise, we could add \overline{pq} to \mathcal{D} and still have