Voronoi Diagram

If the notion of Delaunay triangulation still sounds a bit arbitrary and uncomfortable, that's because we had deferred the truly historical motivation — until now. This section visits a rather natural geometric object named "Voronoi diagram" and argues its reduction to Delaunay triangulation.

Suppose we have, as usual, a finite point set $S \subseteq \mathbb{R}^2$ with $n := |S|$. Now given a query point $q \in \mathbb{R}^2$, how can we quickly find ~~the~~ its closest point in S?

The problem is boring if we have only one query — the best we could do is then inspecting all points $x \in S$, Computing $\|q-x\|$, and taking the minimum one. However, when the number of queries m is large, say $m \gg n$, the problem becomes quite interesting. The aforementioned approach takes $O(mn)$ time, but it seems like a lot of computations are redundant. After all, there are only n possible responses to any inquiry, so in principle we could "summarise" which point in $\mathbb{R}^2$ corresponds to which response <u>beforehand</u>, and look up for the response when an actual inquiry comes. This idea is exactly what Voronoi diagram does.

def Voronoi region.

For a finite $S \subseteq \mathbb{R}^2$ and $x \in S$, the Voronoi region of x with respect to S is defined as

$$\mathrm{Voronoi}_S(x) := \{ y \in \mathbb{R}^2 : \|y-x\| \le \|y-x'\| \ \forall x' \in S \}.$$

In other words, $\mathrm{Voronoi}_S(x)$ summarises all points in $\mathbb{R}^2$ that has a response "x" in our problem.

Clearly by definition $x \in \mathrm{Voronoi}_S(x)$, so the latter is non-empty. Two simple observations:

- Every point $y \in \mathbb{R}^2$ belongs to some $\mathrm{Voronoi}_S(x)$, $x \in S$.
- For any $x \neq x' \in S$, the interior of their Voronoi regions are disjoint; that is

$$\mathrm{Voronoi}^\circ_S(x) \cap \mathrm{Voronoi}^\circ_S(x') = \emptyset.$$

in particular, $x' \notin \mathrm{Voronoi}_S(x)$

Therefore,

$$\{ \mathrm{Voronoi}_S(x) : x \in S \}$$

forms a partition of $\mathbb{R}^2$.

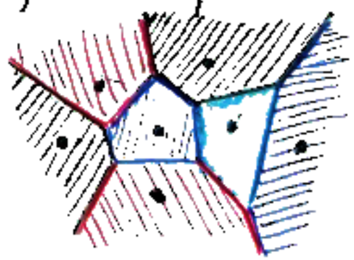

You might have noticed that the Voronoi regions are convex polyhedra (the term "polyhedron" generalises "polygon" by allowing unbounded edges). This observation is easy to formalise. For any two points $x, x' \in \mathbb{R}^2$, their bisector

$$B(x,x') := \{ y \in \mathbb{R}^2 : \|y-x\| = \|y-x'\| \}$$

contains all points that are equally-distant from x and x'. As an exercise, use basic linear algebra to prove $B(x,x')$ is an orthogonal line to xx' that passes through $\frac{1}{2}(x+x')$.

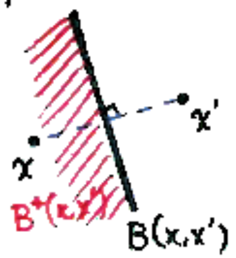

We use $B^+(x,x')$ to denote the halfplane bounded by $B(x,x')$ that contains point x.

Lemma 27.

$$\mathrm{Voronoi}_S(x) = \bigcap_{\substack{x' \in S \\ x' \neq x}} B^+(x,x').$$

Consequently, it is a convex polyhedron of $O(n)$ complexity (hence indeed a "region").

The proof is just by comparing definition, so we leave it as an exercise.

Now that we understand how an individual Voronoi region looks like, it's time to examine how different regions interact. For convenience, we collect the vertices of the Voronoi regions into set V_S, and the edges into set E_S.

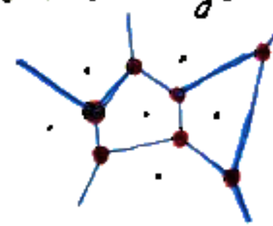

S, V_S and E_S

Lemma 28.

(1) $\forall e \in E_S$, $\mathring{e} \cap V_S = \emptyset$

(2) $\mathrm{Voronoi}_S(x)$ is incident to $\mathrm{Voronoi}_S(x')$
$\Leftrightarrow$ $\mathrm{Voronoi}_S(x) \cap \mathrm{Voronoi}_S(x')$ is exactly one common edge $e \in E_S$.

(3) $\forall v \in V_S$ is incident to ≥ 3 Voronoi regions, say $\mathrm{Voronoi}_S(x_1), \cdots, \mathrm{Voronoi}_S(x_k)$. Moreover, there is a circle centered at v that goes through $x_1, \cdots, x_k$, and whose interior $\cap S = \emptyset$.

Proof.

(3)

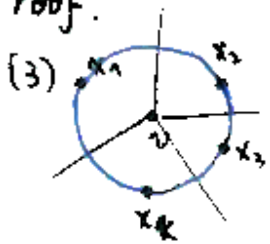

v is on the boundary of k Voronoi regions, all of which are convex by Lemma 27. So $k \geq 3$ for otherwise either a region is concave, or two regions are "flat" and v is not a vertex.

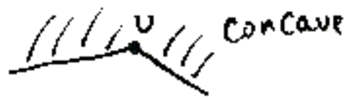

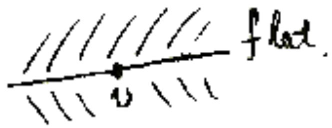

Now look at the points $x_1, \dots, x_k$. Since $v \in \mathrm{Voronoi}_S(x_i)$ for all $i \in [k]$, we see by definition that

$$\|v - x_1\| = \|v - x_2\| = \cdots = \|v - x_k\| =: r$$

so all x_i's lie on the circle of radius r centered at v.

Finally, note that no other point $x' \in S$ could lie in/on the circle, for otherwise $\|v - x'\| \leq r$, contradicting definitions of $x_1, \dots, x_k$.

(1) Suppose to the contrary that $v \in \mathring{e}$ for some $e \in E_S$, $v \in V_S$.

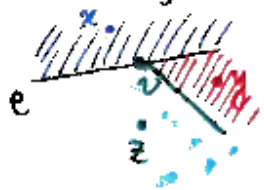

Then by (1) we know v is incident to at least 3 Voronoi regions, say $\mathrm{Voronoi}_S(x)$, $\mathrm{Voronoi}_S(y)$ and $\mathrm{Voronoi}_S(z)$. Without loss of generality we assume e is an edge of $\mathrm{Voronoi}_S(x)$.

From (3) we know that x, y, z are on a circle centered at v:

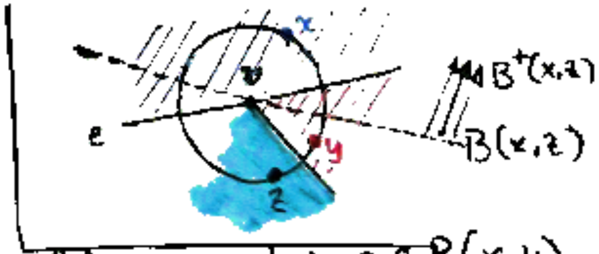

Also note that $e \subseteq B(x,y)$, so $e \perp xy$. The issue is, no matter where we put z, the bisector $B(x,z)$ is always a "tilted" line different from e, which means $\mathrm{Voronoi}(x)$ should have been bounded by $B(x,z)$ instead of e. This is a contradiction.

(2) The ($\Leftarrow$) part is trivial. The ($\Rightarrow$) part is true since incident regions could not share "half an edge" according to (1).

[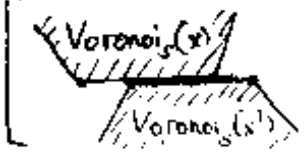

is excluded by (1).] ∎

Remark. If S is in general position (i.e. no four points are cocircular), then the "≥ 3" in statement (3) can be replaced by "$=3$". This would imply that v is the centre of the circumcircle of $x_1, x_2, x_3 \in S$ and that the circle is empty; in other words, the triangle $x_1x_2x_3$ satisfies the Delaunay property.

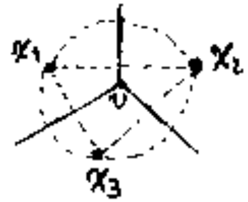

Now it should come at no surprise that the following connection exists:

Theorem 29.
Let $S \subseteq \mathbb{R}^2$ be a finite point set in general position. For each Voronoi vertex $v \in V_S$, associate a triangle $T_v := x_1x_2x_3$ where Voronoi$_S(x_1)$, Voronoi$_S(x_2)$, Voronoi$_S(x_3)$ are the incident Voronoi regions of v. Then $\mathcal{T} := \{T_v : v \in V_S\}$ is a Delaunay triangulation for S, and vice versa.

Proof. Let $\hat{\mathcal{T}}$ be a Delaunay triangulation of S.
($\Leftarrow$) We will show that $\forall T \in \hat{\mathcal{T}}$ also satisfies $T \in \mathcal{T}$.

Assume $T = x_1x_2x_3 \in \hat{\mathcal{T}}$. Then by Delaunay property C_T is empty. Let u be the centre of C_T. Since u must be in at least one Voronoi region, let us ask: Which Voronoi region(s) does u live? After some thoughts, one easily sees that $u \in$ Voronoi$_S(x_i)$, $i=1,2,3$ as C_T is empty. So $u \in V_S$ and is incident to three regions Voronoi$_S(x_1)$, Voronoi$_S(x_2)$ and Voronoi(x_3). Hence $T = T_u \in \mathcal{T}$.

(⇒) We will show that $\forall T_v \in \mathcal{T}$ also satisfies $T_v \in \hat{\mathcal{T}}$.

Assume $T_v = x_1x_2x_3$, then by definition v is incident to Voronoi regions $\text{Voronoi}_S(x_1)$, $\text{Voronoi}_S(x_2)$ and $\text{Voronoi}_S(x_3)$. By Lemma 28(3) and the remark, v is the centre of C_{T_v}, and the interior of C_{T_v} is empty. Moreover, no other points are on the circle. So we could "move C_{T_v}" a little to obtain an empty circle that goes through only x_1x_2. Then by Lemma 26, the segment $\overline{x_1x_2}$ is contained in all Delaunay triangulations of S — in particular $\hat{\mathcal{T}}$. Similarly, $\overline{x_2x_3} \in \hat{\mathcal{T}}$ and $\overline{x_1x_3} \in \hat{\mathcal{T}}$. Putting them together: $T_v = x_1x_2x_3 \in \hat{\mathcal{T}}$. ∎

Remark. The proof shows how natural it is to define the Delaunay property in the way we introduced it. Indeed, historically, Voronoi diagrams motivates the study of the research on Delaunay triangulations.

~~Besides, the proof shows the one-one correspondence between Voronoi diagram and Delaunay triangulation: a vertex $v \in V_S$ ↔ a face triangle $T \in DT$~~

To clarify the relation between Voronoi and Delaunay, we make a table:

Voronoi diagram		Delaunay triangulation
Voronoi region $\text{Voronoi}_S(x)$	$\overset{\text{obvious}}{\Longleftrightarrow}$	Vertex (i.e. point) x
Voronoi vertex $v \in V_S$	$\overset{\text{Thm 29}}{\Longleftrightarrow}$	Triangle
Voronoi edge $e \in E_S$	$\perp$	Edge
⇕ Lemma 28(2)		⇕
Incidence of two Voronoi regions $\text{Voronoi}_S(x)$ and $\text{Voronoi}_S(x')$	$\overset{\text{Thm 29}}{\Longleftrightarrow}$	Incidence of two vertices

There are of course more duality to exploit; we leave it as a fun exercise. Our main point is: Knowing one object would allow easy construction of the other.

- **Problem.** Describe an algorithm that takes $S \subseteq \mathbb{R}^2$ (in general position) and its Delaunay triangulation as inputs, and outputs the Voronoi diagram of S.

The duality readily implies that

$$|V_S| = \#\text{ of triangles in Delaunay triangulation of } S \le 2n-5 = O(n)$$

$$|E_S| = \#\text{ of edges in Delaunay triangulation of } S \le 3n-6 = O(n)$$

so the entire structure of Voronoi diagram can be stored in linear space.

Now that we have precomputed the Voronoi diagram, how does it help during the inquiries? The last missing piece is a data structure called "Kirkpatrick's hierarchy" built on top of the diagram. One could understand it as an indexing tree structure so that we don't need to go through all $x \in S$ when an inquiry comes.

The idea of Kirkpatrick's hierarchy is outlined below.

For simplicity we assume that the query points only lie inside a huge "wrapping triangle" so that we don't

have to deal with infinite regions.

- First, triangulate all regions arbitrarily:

$\mathcal{T}_0$

- Then, remove some (Voronoi) vertices to make the mesh coarser. But we also retriangulate the resulting mesh so as to keep things simple.

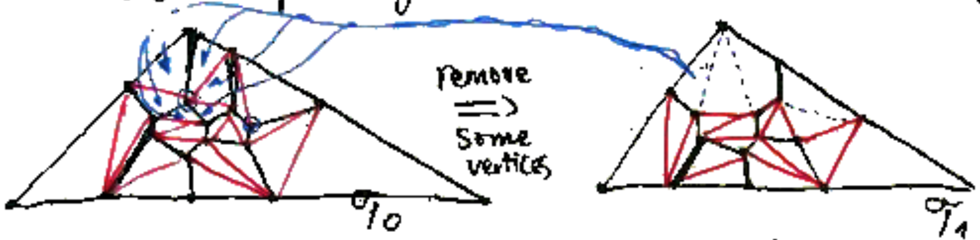

- Add pointers from the resulting mesh to the original one to indicate the "origin" of a triangle. The illustration above gives an example.

- Repeat the procedure until there are only three (external) vertices left:

- When a query q comes, we first look at the most coarse picture $\mathcal{T}_h$. Then, following the pointers, we ~~follow~~ trace back to $\mathcal{T}_{h-1}$ to see which triangle in $\mathcal{T}_{h-1}$ does q lies in. We repeat this search recursively until we arrive $\mathcal{T}_0$, where the complete picture is present.

To make the algorithm efficient, two factors have to be taken into account:

(1) The "height" h of the indexing structure should be small

(2) When we trace back from $\mathcal{T}_i$ to $\mathcal{T}_{i-1}$, the number of pointers we ~~inspect~~ traversed should be small.

Theorem 30 (Kirkpatrick)
We could build such an indexing structure in $O(n)$ time such that

(1) $h = O(\log n)$;

(2) For any $T \in \mathcal{T}_i$, the number of "trace-back pointers" to $\mathcal{T}_{i-1}$ ≤ 8.

Proof. The key insight of Kirkpatrick is:

"In a planar graph, the average degree is less than 6. Therefore, a good portion of the vertices actually have degree $\leq \Delta$ for some constant Δ. Let H be the subgraph induced by these vertices, then $\chi(H) \leq \Delta+1$ (and we could find a $(\Delta+1)$ colouring greedily). Hence $\alpha(G) \geq \frac{|V(H)|}{\chi(H)} \geq \frac{|V(H)|}{\Delta+1}$ is large.

Removing such an independent set would kill a $\left(\frac{1}{\Delta+1}\right)$ proportion of vertices, so $h \leq \log_{\frac{\Delta+1}{\Delta}} n$ as desired. Since we remove an independent set each time, the "holes" in the resulting mesh are independent, so the number of "track back pointers" from $T \in \mathcal{T}_i$ to $\mathcal{T}_{i-1}$ is bounded by Δ."

Now we simply carry out his plan by some calculations.

Claim: $U := \{ v \in V(G) : \deg_G(v) \leq 8 \}$ has at least $\frac{1}{2}|V(G)|$ vertices for maximal planar graph G.

Reason:

$$6|V(G)| \geq \sum_{v \in V(G)} \deg(v) = \sum_{v \in U} \deg(v) + \sum_{v \notin U} \deg(v)$$
$$\geq 3 \cdot |U| + 9 \cdot (|V(G)| - |U|)$$
$$= 9 \cdot |V(G)| - 6|U|$$

because max-planar graph is 3-connected

Moving terms give the claim.

Therefore, we could find greedily an independent set in $G[U]$ of size $\geq \frac{|U|}{9} \geq \frac{|V(G)|}{18}$.

Each step we remove such an independent set, so after at most $h \leq \log_{\frac{17}{18}} n$ steps only the wrapping triangle remains. Besides, by definition of the trace-back pointers, the number of pointers from $T \in \mathcal{T}_i$ to $\mathcal{T}_{i-1}$ is at most the degree of the removed vertex in step $(i-1)$, which is at most 8. This finishes the proof. $\blacksquare$