CMSC 754: Lecture 11
Delaunay Triangulations: General Properties

Reading: Chapter 9 in the 4M's.

Delaunay Triangulations: We discussed Voronoi diagrams. In this lecture, we consider a closely related structure, called the *Delaunay triangulation* (DT). The Voronoi diagram of a set of sites in the plane is a planar subdivision, in fact, a cell complex. The *dual* of such a subdivision is a cell complex that is defined as follows. For each face of the Voronoi diagram, we create a vertex (corresponding to the site). For each edge of the Voronoi diagram between two sites p_i and p_j, we create an edge in the dual connecting the corresponding vertices. Each Voronoi vertex corresponds to a face of the dual complex.

Recall that, under the assumption of general position (no four sites are collinear), the vertices of the Voronoi diagram all have degree three. Thus, the faces of the resulting dual complex (excluding the exterior face) are triangles, and the resulting dual graph is a *triangulation* of the sites. This is called the *Delaunay triangulation* (see Fig. 1(a)).

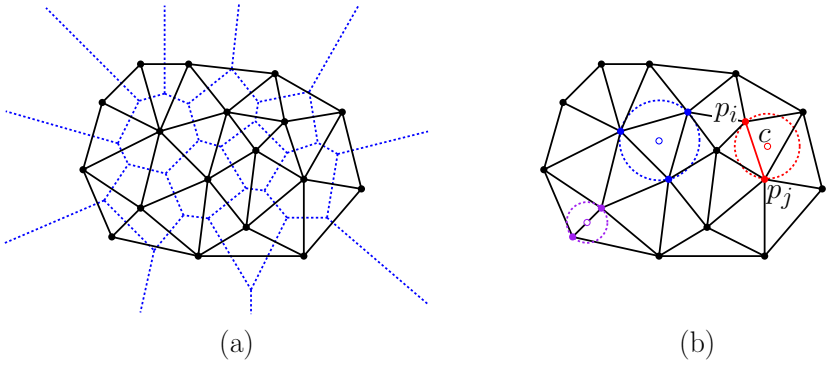


Fig. 1: (a) The Voronoi diagram of a set of sites (broken lines) and the corresponding Delaunay triangulation (solid lines), and (b) circle-related properties.

Delaunay triangulations have a number of interesting properties that are immediate consequences of the structure of the Voronoi diagram:

Convex hull: The boundary of the exterior face of the Delaunay triangulation is the boundary of the convex hull of the point set.

Circumcircle property: The circumcircle of any triangle in the Delaunay triangulation is "empty," that is, the interior of the associated circular disk contains no sites of P (see the blue circle in Fig. 1(b)).

Proof: The center of this circle is the corresponding dual Voronoi vertex, and by definition of the Voronoi diagram, the three sites defining this vertex are its nearest neighbors.

Empty circle property: Two sites p_i and p_j are connected by an edge in the Delaunay triangulation if and only if there is an empty circle passing through p_i and p_j (see the red circle in Fig. 1(b)).

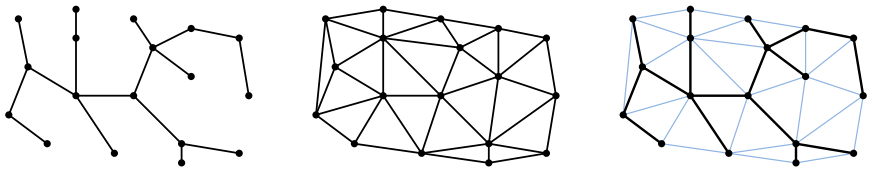
Proof: If two sites p_i and p_j are neighbors in the Delaunay triangulation, then their cells are neighbors in the Voronoi diagram, and so for any point on the Voronoi edge between these sites, a circle centered at this point passing through p_i and p_j cannot contain any other point (since they must be closest). Conversely, if there is an empty circle passing through p_i and p_j, then the center c of this circle is a point on the edge of the Voronoi diagram between p_i and p_j, because c is equidistant from each of these sites and there is no closer site (see Fig. 1(b)). Thus the Voronoi cells of two sites are adjacent in the Voronoi diagram, implying that this edge is in the Delaunay triangulation.

Closest pair property: The closest pair of sites in P are neighbors in the Delaunay triangulation (see the green circle in Fig. 1(b)).

Proof: Suppose that p_i and p_j are the closest sites. The circle with p_i and p_j as its diameter cannot contain any other site, since otherwise that site would be closer to one of these two points, violating the hypothesis that these points are the closest pair. Therefore, the center of this circle is on the Voronoi edge between these points, and so it is an empty circle.

Combinatorial Complexity: Given a point set P with n sites where there are h sites on the convex hull, it is not hard to prove by Euler's formula that the Delaunay triangulation has $2n - 2 - h$ triangles and $3n - 3 - h$ edges. Generally, in $\mathbb{R}^d$, the number of simplices (the d-dimensional generalization of a triangle) can range from $O(n)$ up to $O(n^{\lceil d/2 \rceil})$. For example, in $\mathbb{R}^3$ the Delaunay triangulation of n sites may have as many as $O(n^2)$ tetrahedra. (If you want a challenging exercise, try to create such a point set.)

Euclidean Minimum Spanning Tree: The Delaunay triangulation has several interesting properties that are not obviously related to the Voronoi diagram structure. One of these is its relation to the minimum spanning tree. Given a set of n points in the plane, we can think of the points as defining a *Euclidean graph* whose edges are all $\binom{n}{2}$ (undirected) pairs of distinct points, and edge (p_i, p_j) has weight equal to the Euclidean distance from p_i to p_j. Given a graph, the minimum spanning tree (MST) is a set of $n-1$ edges connecting these points (into a free tree) that minimizes the total edge weight. The MST of the Euclidean graph is called the *Euclidean minimum spanning tree* (EMST) (see Fig. 2(c)).



(a) (b) (c)

Fig. 2: (a) A point set and its EMST, (b) the Delaunay triangulation, and (c) the overlay of the two.

We could compute the EMST by brute force: construct the Euclidean graph and then run Kruskal's algorithm to find its MST. This would lead to a total running time of $O(n^2 \log n)$. However, a much faster method for planar point sets uses Delaunay triangulations. First compute the Delaunay triangulation of the point set. We will see later that it can be done

in $O(n \log n)$ time. Then compute the MST of the Delaunay triangulation by, say, Kruskal's algorithm and return the result. This leads to a total running time of $O(n \log n)$. The following theorem explains why this works.

Theorem: The minimum spanning tree of a set P of point sites (in any dimension) is a subgraph of the Delaunay triangulation (see Fig. 2(c)).

Proof: Let T be the EMST for P, let $w(T)$ denote the total weight of T. Let a and b be any two sites such that ab is an edge of T. Suppose to the contrary that ab is not an edge in the Delaunay triangulation. This implies that there is no empty circle passing through a and b, and in particular, the circle whose diameter is the segment $\overline{ab}$ contains another site, call it c (see Fig. 3.)

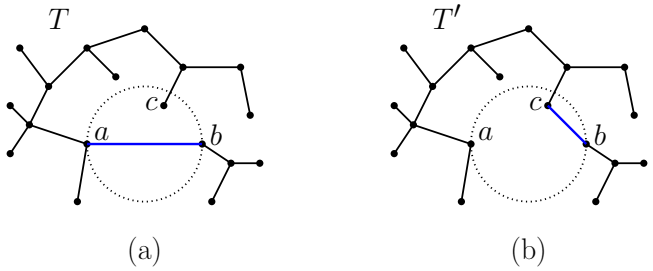


Fig. 3: The Delaunay triangulation and EMST.

The removal of $\overline{ab}$ from the EMST splits the tree into two subtrees. Assume without loss of generality that c lies in the same subtree as a. Now, remove the edge $\overline{ab}$ from the EMST and add the edge $\overline{bc}$ in its place. The result will be a spanning tree T' whose weight is

$$w(T') \;=\; w(T) + \|bc\| - \|ab\|.$$

Since ab is the diameter of the circle, any other segment lying within the circle is shorter. Thus, $\|bc\| < \|ab\|$. Therefore, we have $w(T') < w(T)$, and this contradicts the hypothesis that T is the EMST, completing the proof.

Minimum Weight Triangulation (Not!): Given the above result about EMSTs, it is tempting to conjecture that among all triangulations, the Delaunay triangulation minimizes the total (Euclidean) edge length. Unfortunately, this is not true and, indeed, there is a simple four-point counterexample. (This assertion was erroneously claimed in a highly cited paper on Delaunay triangulations, and you may still hear it quoted from time to time.)

The triangulation that minimizes total Euclidean edge weight is called the *minimum weight triangulation* (MWT). The computational complexity of MWT was open for many years, and many special cases were shown to be solvable in polynomial time. Finally, in 2008, it was proved by Mulzer and Rote[1] that this problem is NP-hard. The hardness proof is quite complex, and computer assistance was needed to verify the correctness of some of the constructions used in the proof.

[1] W. Mulzer and G. Rote, "Minimum-weight triangulation is NP-hard", *J. ACM*, 55 (2), 2008.

Spanner Properties: A natural observation about Delaunay triangulations is that their edges would seem to form a reasonable transportation road network between the points. From a few examples, it is natural to conjecture that the shortest path between two points in a planar Delaunay triangulation is not significantly longer than the straight-line distance between them.

This is closely related to the theory of *geometric spanners*, that is, geometric graphs whose shortest paths are not significantly longer than the straight-line distance. Consider any point set P and a straight-line graph G whose vertices are the points of P. For any two points $p, q \in P$, let $\delta_G(p, q)$ denote the length of the shortest path from p to q in G, where each edge has weight equal to its Euclidean length. Given any parameter $t \geq 1$, we say that G is a *t-spanner* if for any two points $p, q \in P$, the shortest path length between p and q in G is at most a factor t longer than the Euclidean distance between these points, that is

$$\delta_G(p, q) \leq t\|pq\|$$

Observe that when $t = 1$, the graph G must be the complete graph, consisting of $\binom{n}{2} = O(n^2)$ edges. The question is whether $O(1)$-spanners exist with $O(n)$ edges.

It can be proved that the edges of the Delaunay triangulation form a spanner (see Fig. 4). We will not prove the following result, which is due to Ge Xia.[2]

Theorem: Given a set of points P in the plane, the Delaunay triangulation of P is a t-spanner for $t = 1.998$.

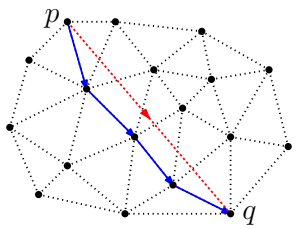


Fig. 4: Spanner property of the Delaunay Triangulation.

It had been conjectured for many years that the Delaunay triangulation is a $(\pi/2)$-spanner ($\pi/2 \approx 1.5708$). This was disproved in 2009, and the lower bound now stands at roughly 1.5932. Closing the gap between the upper and lower bound is an important open problem.

Maximizing Angles and Edge Flipping: Another interesting property of Delaunay triangulations is that among all triangulations, the Delaunay triangulation maximizes the minimum angle. This property is important because it implies that Delaunay triangulations tend to avoid skinny triangles. This is useful for many applications where triangles are used for the purposes of interpolation.

In fact, a stronger statement holds. Among all triangulations that maximize the smallest angle, the Delaunay triangulation maximizes the second smallest angle. Among all triangulations that maximize both the two smallest angles, the Delaunay triangulation maximizes

[2] G. Xia, "The stretch factor of the Delaunay triangulation is less than 1.998," *SIAM J. Comput.*, 42, 2013.

the third smallest angle. That is, the Delaunay triangulation maximizes the angle sequence lexicographically.

To make this more formal, consider any triangulation of a given set P of n sites. Let t be the number of triangles. Then the number of angles is $m = 3t$. We can associate this triangulation with a sorted *angle sequence*, by which we mean a non-decreasing sequence of angles $(\alpha_1 \leq \alpha_2 \leq \cdots \leq \alpha_m)$ appearing in all the triangles of the triangulation. (Note that the length of the sequence will be the same for all triangulations of the same point set, since the number of triangles depends only on the number of sites n and the number of points on the convex hull h.)

Theorem: Among all triangulations of a given planar point set, the Delaunay triangulation has the lexicographically largest angle sequence.

Proof: (Optional) Before the proof, let us recall a few basic facts about angles from basic geometry. First, recall that for the circumcircle of three points, each angle of the resulting triangle is exactly half the angle of the minor arc subtended by the opposite two points along the circumcircle. It also follows that if a point is inside this circle, it subtends a larger angle, and a point outside subtends a smaller angle. Thus, in Fig. 5(a) below, we have $\theta_1 > \theta_2 > \theta_3$.

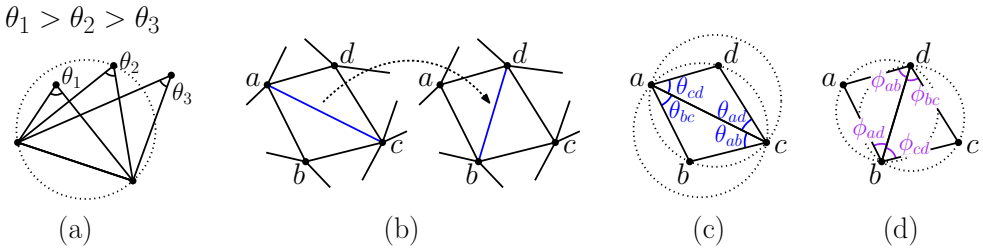


Fig. 5: Angles and edge flips.

We will not give a formal proof of the theorem. (A full proof appears in the text.) The main idea is to show that for any triangulation that fails to satisfy the empty circle property, it is possible to perform a local operation, called an *edge flip*, which increases the lexicographical sequence of angles. An edge flip is an important fundamental operation on triangulations in the plane. Given two adjacent triangles $\triangle abc$ and $\triangle cda$, such that their union forms a convex quadrilateral $abcd$, the edge flip operation replaces the diagonal ac with bd (see Fig. 5(b)). This is only possible when the quadrilateral is convex.

Suppose that the initial triangle pair violates the empty circle condition, in that point d lies inside the circumcircle of $\triangle abc$. (Note that this implies that b lies inside the circumcircle of $\triangle cda$.) If we flip the edge, the two circumcircles of the resulting triangles, $\triangle abd$ and $\triangle bcd$, are now empty (relative to these four points), and the observation above about circles and angles shows that the minimum angle increases at the same time. In particular, in Fig. 5(c) and (d), we have

$$\phi_{ab} > \theta_{ab}, \quad \phi_{bc} > \theta_{bc}, \quad \phi_{cd} > \theta_{cd}, \quad \text{and} \quad \phi_{da} > \theta_{da}.$$

Two other angles also need to be compared (can you spot them?). It is not hard to show that, after swapping, these other two angles cannot be smaller than the minimum of θ_{ab}, θ_{bc}, θ_{cd}, and θ_{da}. (Can you see why?)

Since there are only finitely many triangulations, this process must eventually terminate with the lexicographically maximum triangulation, which must satisfy the empty circle condition and hence is the Delaunay triangulation.

Polytopes and Spatial Subdivisions: At a glance, Delaunay triangulations and convex hulls seem quite different. One is based on metric properties (distances) and the other on affine properties (collinearity, coplanarity). On the other hand, if you look at the surface of the convex hull of a set of points in 3-dimensional space, the boundary structure looks much like a triangulation. (If the points are in general position, no four points are coplanar, so each face of the convex hull will be bounded by three vertices.)

Similarly, consider a boundary structure of a polytope defined by the intersection of a collection of halfplanes in 3-dimensional space. Assuming general position (no four planes intersecting at a common point), each vertex will be incident to exactly three faces, and hence to exactly three edges. Therefore, the boundary structure of this polytope will look very much like a Voronoi diagram.

We will show that there is a remarkably close relationship between these structures. In particular, we will show that:

- The Delaunay triangulation of a set of points in the plane is topologically equivalent to the boundary complex of the (lower) convex hull of an appropriate set of points in 3-space. Thus, we can reduce the problem of computing Delaunay triangulations in dimension d to computing convex hulls in dimension $d + 1$.
- The Voronoi diagram of a set of points in the plane is topologically equivalent to the boundary complex of the intersection of a set of halfspaces in 3-space. In general, we can reduce computing Voronoi diagrams in dimension d to computing the upper envelope of a set of hyperplanes in dimension $d + 1$.

We will demonstrate these results in 2-dimensional space, but the generalizations to higher dimensions are straightforward.

Delaunay Triangulations and Convex Hulls: Let us begin by considering the *paraboloid* Ψ defined by the equation $z = x^2 + y^2$. Observe that the vertical cross sections (constant x or constant y) are parabolas, and whose horizontal cross sections (constant z) are circles. For each point in $\mathbb{R}^2$, $p = (p_x, p_y)$, the *vertical projection* (also called the *lifted image*) of this point onto this Ψ is $p^\uparrow = (p_x, p_y, p_x^2 + p_y^2)$ in $\mathbb{R}^3$.

Given a set of points P in the plane, let $P^\uparrow$ denote the projection of every point in P onto Ψ. Consider the *lower convex hull* of $P^\uparrow$. This is the portion of the convex hull of $P^\uparrow$ which is visible to a viewer standing at $z = -\infty$. We claim that if we take the lower convex hull of $P^\uparrow$, and project it back onto the plane, then we get the Delaunay triangulation of P (see Fig. 6). In particular, let $p, q, r \in P$, and let $p^\uparrow$, $q^\uparrow$, $r^\uparrow$ denote the projections of these points onto Ψ. Then $\triangle p^\uparrow q^\uparrow r^\uparrow$ defines a *face* of the lower convex hull of $P^\uparrow$ if and only if $\triangle pqr$ is a triangle of the Delaunay triangulation of P.

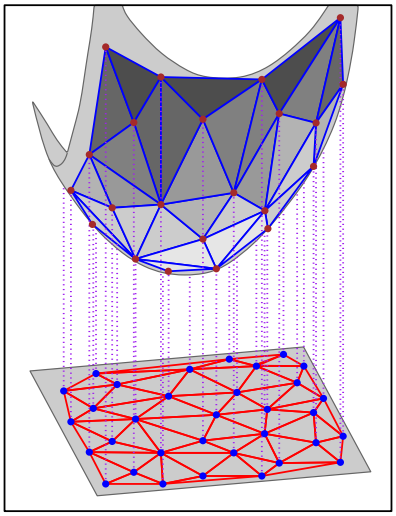


Fig. 6: The Delaunay triangulation and convex hull.

The question is, why does this work? To see why, we need to establish the connection between the triangles of the Delaunay triangulation and the faces of the convex hull of transformed points. In particular, recall that

Delaunay condition: Three points $p, q, r \in P$ form a Delaunay triangle if and only no other point of P lies *within the circumcircle* of the triangle defined by these points.

Convex hull condition: Three points $p^\uparrow, q^\uparrow, r^\uparrow \in P^\uparrow$ form a face of the convex hull of $P^\uparrow$ if and only if no other point of P lies *below the plane* passing through $p^\uparrow$, $q^\uparrow$, and $r^\uparrow$.

Clearly, the connection we need to establish is between the emptiness of circumcircles in the plane and the emptiness of lower halfspaces in 3-space. To do this, we will prove the following.

Lemma: Consider four distinct points p, q, r, and s in the plane, and let $p^\uparrow$, $q^\uparrow$, $r^\uparrow$, and $s^\uparrow$ denote their respective vertical projections onto Ψ, $z = x^2 + y^2$. The point s lies within the circumcircle of $\triangle pqr$ if and only if $s^\uparrow$ lies beneath the plane passing through $p^\uparrow$, $q^\uparrow$, and $r^\uparrow$.

To prove the lemma, first consider an arbitrary (nonvertical) plane in 3-space, which we assume is tangent to Ψ above some point (a, b) in the plane. To determine the equation of this tangent plane, we take derivatives of the equation $z = x^2 + y^2$ with respect to x and y giving

$$\frac{\partial z}{\partial x} = 2x, \qquad \frac{\partial z}{\partial y} = 2y.$$

At the point $(a, b, a^2 + b^2)$ these evaluate to $2a$ and $2b$. It follows that the plane passing through these point has the form

$$z = 2ax + 2by + \gamma, \quad \text{for some real } \gamma.$$

To solve for γ we know that the plane passes through $(a, b, a^2 + b^2)$ so we solve giving

$$a^2 + b^2 = 2a \cdot a + 2b \cdot b + \gamma,$$

Implying that $\gamma = -(a^2 + b^2)$. Thus the plane equation is

$$z = 2ax + 2by - (a^2 + b^2). \tag{1}$$

If we shift the plane upwards by some positive amount r^2 we obtain the plane

$$z = 2ax + 2by - (a^2 + b^2) + r^2.$$

How does this plane intersect Ψ? Since Ψ is defined by $z = x^2 + y^2$ we can eliminate z, yielding

$$x^2 + y^2 = 2ax + 2by - (a^2 + b^2) + r^2,$$

which after some simple rearrangements is equal to

$$(x - a)^2 + (y - b)^2 = r^2.$$

Hey! This is just a circle centered at the point (a, b). Thus, we have shown that the intersection of a plane with Ψ produces a space curve (which turns out to be an ellipse), which when projected back onto the (x, y)-coordinate plane is a circle centered at (a, b) whose radius equals the square root of the vertical distance by which the plane has been translated.

Thus, we conclude that the intersection of an arbitrary lower halfspace with Ψ, when projected onto the (x, y)-plane is the interior of a circle. Going back to the lemma, when we project the points p, q, r onto Ψ, the projected points $p^\uparrow$, $q^\uparrow$ and $r^\uparrow$ define a plane. Since $p^\uparrow$, $q^\uparrow$, and $r^\uparrow$, lie at the intersection of the plane and Ψ, the original points p, q, r lie on the projected circle. Thus this circle is the (unique) circumcircle passing through these p, q, and r. Thus, the point s lies within this circumcircle, if and only if its projection $s^\uparrow$ onto Ψ lies within the lower halfspace of the plane passing through p, q, r (see Fig. 7).

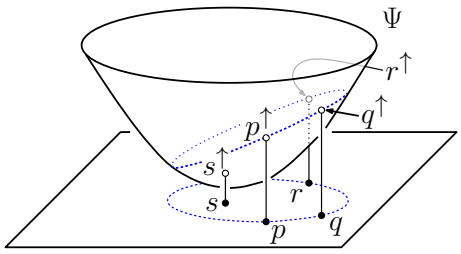


Fig. 7: Planes and circles.

Now we can prove the main result.

Theorem: Given a set of points P in the plane (assuming no four are cocircular), and given three points $p, q, r \in P$, the triangle $\triangle pqr$ is a triangle of the Delaunay triangulation of P if and only if triangle $\triangle p^\uparrow q^\uparrow r^\uparrow$ is a face of the lower convex hull of the lifted set $P^\uparrow$.

From the definition of Delaunay triangulations we know that $\triangle pqr$ is in the Delaunay triangulation if and only if there is no point $s \in P$ that lies within the circumcircle of pqr. From the previous lemma this is equivalent to saying that there is no point $s^\uparrow$ that lies in the lower convex hull of $P^\uparrow$, which is equivalent to saying that $\triangle p^\uparrow q^\uparrow r^\uparrow$ is a face of the lower convex hull. This completes the proof.

Voronoi Diagrams and Upper Envelopes: (Optional) Next, let us consider the relationship between Voronoi diagrams and envelopes. We know that Voronoi diagrams and Delaunay triangulations are dual geometric structures. We have also seen (informally) that there is a dual relationship between points and lines in the plane, and in general, between points and planes in 3-space. From this latter connection we argued that the problems of computing convex hulls of point sets and computing the intersection of halfspaces are somehow "dual" to one another. It turns out that these two notions of duality, are (not surprisingly) interrelated. In particular, in the same way that the Delaunay triangulation of points in the plane can be transformed to computing a convex hull in 3-space, the Voronoi diagram of points in the plane can be transformed into computing the upper envelope of a set of planes in 3-space.

Here is how we do this. For each point $p = (a, b)$ in the plane, recall from Eq. (1) that the tangent plane to Ψ passing through the lifted point $p^\uparrow$ is

$$z = 2ax + 2by - (a^2 + b^2).$$

Define $h(p)$ to be this plane. Consider an arbitrary point $q = (q_x, q_y)$ in the plane. Its vertical projection onto Ψ is (q_x, q_y, q_z), where $q_z = q_x^2 + q_y^2$). Because Ψ is convex, $h(p)$ passes below Ψ (except at its contact point $p^\uparrow$). The vertical distance from $q^\uparrow$ to the plane $h(p)$ is

$$\begin{aligned} q_z - (2aq_x + 2bq_y - (a^2 + b^2)) &= (q_x^2 + q_y^2) - (2aq_x + 2bq_y - (a^2 + b^2)) \\ &= (q_x^2 - 2aq_x + a^2) + (q_y^2 - 2bq_y + b^2) \;=\; \|qp\|^2. \end{aligned}$$

In summary, the vertical distance between $q^\uparrow$ and $h(p)$ is just the squared distance[3] from q to p (see Fig. 8(a)).

Now, consider a point set $P = \{p_1, \ldots, p_n\}$ and an arbitrary point q in the plane. From the above observation, we have the following lemma.

Lemma: Given a set of points P in the plane, let $H(P) = \{h(p) : p \in P\}$. For any point q in the plane, a vertical ray directed downwards from $q^\uparrow$ intersects the planes of $H(P)$ in the same order as the distances of the points of P from q (see Fig. 8(b)).

Consider the upper envelope $U(P)$ of $H(P)$. This is an unbounded convex polytope (whose vertical projection covers the entire x, y-plane). If we label every point of this polytope with the associated point of p whose plane $h(p)$ defines this face, it follows from the above lemma that p is the closest point of P to every point in the vertical projection of this face onto the plane. As a consequence, we have the following equivalence between the Voronoi diagram of P and $U(P)$ (see Fig. 9).

[3] If you have studied machine learning, you may have learned about the concept of a Bregman divergence. We have essentially proved that the squared Euclidean distance is the Bregman divergence induced by the function $F(x, y) = x^2 + y^2$.

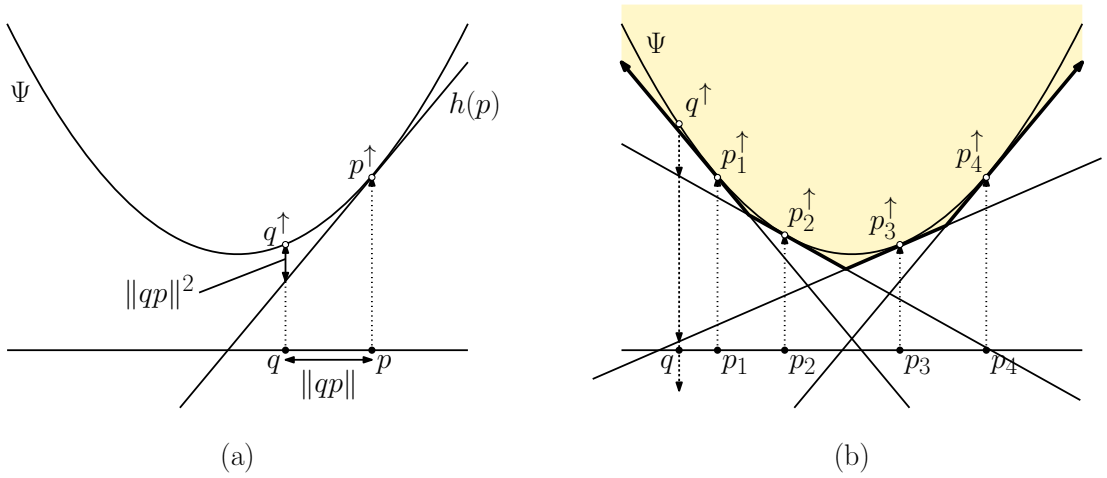


Fig. 8: The Voronoi diagram and the upper hull of tangent planes.

Theorem: Given a set of points P in the plane, let $U(P)$ denote the upper envelope of the tangent hyperplanes passing through each point $p^{\uparrow}$ for $p \in P$. Then the Voronoi diagram of P is equal to the vertical projection onto the (x, y)-plane of the boundary complex of $U(P)$ (see Fig. 9).

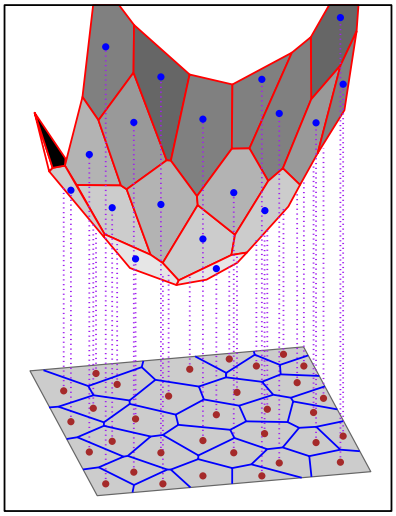


Fig. 9: The Voronoi diagram and an upper envelope.