

Basics for Reconstruction from Data_{taken from [1]}

Simply stated, the problem we study here is: how to approximate a shape from the coordinates of a given set of points from the shape. The set of points is called a point sample, or simply a *sample* of the shape. The specific shape that we will deal with are curves in two dimensions and surfaces in three dimensions. The problem is motivated by the availability of modern scanning devices that can generate a point sample from the surface of a geometric object. For example, a range scanner can provide the depth values of the sampled points on a surface from which the three dimensional coordinates can be extracted. Advanced hand held laser scanners can scan a machine or a body part to provide a dense sample of the surfaces. A number of applications in computer aided design, medical imaging, geographic data processing and drug designs, to name a few, can take advantage of the scanning technology to produce samples and then compute a digital model of a geometric shape with reconstruction algorithms. Figure 16 shows such an example for a sample on a surface which is approximated by a triangulated surface interpolating the input points.

The reconstruction algorithms described here produce a piecewise linear approximation of the sampled curves and surfaces. By approximation we mean that the output captures the topology and geometry of the sampled shape. This requires some concepts from topology which we already covered.

Clearly a curve or a surface cannot be approximated from a sample unless it is dense enough to capture the features of the shape. The notions of features and dense sampling are formalized in Section 18.

All reconstruction algorithms described here use the data structures called *Voronoi diagrams* and their duals called *Delaunay triangulations*. The key properties of these data structures are described in Section 19.

18 Feature size and sampling

We will mainly concentrate on smooth curves in \mathbb{R}^2 and smooth surfaces in \mathbb{R}^3 as the sampled spaces. The notation Σ will be used to denote this generic sampled space throughout. It is sufficient to assume that Σ is a 1-manifold in \mathbb{R}^2 and a 2-manifold in \mathbb{R}^3 for the definitions and results described in this chapter.

Obviously it is not possible to extract any meaningful information about Σ if it is not sufficiently sampled. This means features of Σ should be represented with sufficiently many sample points. Figure 17 shows a curve in the plane which is reconstructed from a sufficiently dense sample. But, this brings up the question of defining features. We aim for a measure that can tell us how complicated Σ is around each point $x \in \Sigma$. A geometric structure called the *medial axis* turns out to be useful to define such a measure.

For a set $P \subseteq \mathbb{R}^k$ and a point $x \in \mathbb{R}^k$, let $d(x, P)$ denote the Euclidean distance of x from P ; that is,

$$d(x, P) = \inf_{p \in P} \{\|p - x\|\}.$$




Figure 16: (a) A sample of MANNEQUIN, (b) a reconstruction, (c) rendered MANNEQUIN model.

We will also consider distances called

Definition 43. The *Hausdorff distance* between two sets $X, Y \subseteq \mathbb{R}^k$ is given by

$$\max\{\sup_{x \in X} d(x, Y), \sup_{y \in Y} d(y, X)\}.$$

Roughly speaking, the Hausdorff distance tells how much one set needs to be moved to be identical with the other set.

18.1 Medial axis

The medial axis of a curve or a surface Σ is meant to capture the middle of the shape bounded by Σ . There are slightly different definitions of the medial axis in the literature. We adopt one of them and mention the differences with the others.

Assume that Σ is embedded in \mathbb{R}^k . A ball $B \subset \mathbb{R}^k$ is empty if the interior of B is empty of points from Σ . A ball B is maximal if every empty ball that contains B equals B . The *skeleton* Sk_Σ of Σ is the set of centers of all maximal balls. Let M_Σ^o be the set of points in \mathbb{R}^k whose distance




Figure 17: (a) A curve in the plane, (b) a sample of it, (c) the reconstructed curve.

to Σ is realized by at least two points in Σ . The closure of M_Σ^o is M_Σ , that is, $M_\Sigma = Cl M_\Sigma^o$. The following inclusions hold:

$$M_\Sigma^o \subseteq Sk_\Sigma \subseteq M_\Sigma.$$

There are examples where the inclusions are strict. For example, consider the curve in Figure 18(a). The two end points u and v are not in M_Σ^o though they are in Sk_Σ . These are the centers of the curvature balls that meet the curve only at a single point. Consider the curve in Figure 18(b):

$$y = \begin{cases} 0 & \text{if } -1 \leq x \leq 0 \\ x^3 \sin \frac{1}{x} & \text{if } 0 < x \leq 1. \end{cases}$$

The two endpoints $(-1, 0)$ and $(1, \sin 1)$ can be connected with a smooth curve so that the resulting curve Σ is closed, that is, without any boundary point, see Figure 18(b). The set M_Σ^o has infinitely many branches, namely one for each oscillation of the $y = x^3 \sin \frac{1}{x}$ curve. The closure of M_Σ^o has a vertical segment at $x = 0$ which is not part of Sk_Σ and thus Sk_Σ is a strict subset of M_Σ . However, this example is a bit pathological since it is known that a large class of curves and




Figure 18: (a) The two endpoints on the middle segment are not in M_Σ^o , but are in Sk_Σ and M_Σ , (b) right half of the bottom curve is $y = x^3 \sin \frac{1}{x}$. Sk_Σ does not include the segment in M_Σ at $x = 0$.

surfaces have $Sk_\Sigma = M_\Sigma$. All curves and surfaces that are at least C^2 -smooth³ have $Sk_\Sigma = M_\Sigma$. The example we considered in Figure 18(b) is a C^1 -smooth curve which is tangent continuous but not curvature continuous.

In our case we will consider only the class of curves and surfaces where $Sk_\Sigma = M_\Sigma$ and thus define the *medial axis* of Σ as M_Σ . For simplicity we write M in place of M_Σ .

Definition 44. The medial axis M of a curve (surface) $\Sigma \subset \mathbb{R}^k$ is the closure of the set of points in \mathbb{R}^k that have at least two closest points in Σ .

Each point of M is the center of a ball that meets Σ only tangentially. We call each ball $B_{x,r}$, $x \in M$, a *medial ball* where $r = d(x, \Sigma)$. If a medial ball $B_{x,r}$ is tangent to Σ at $p \in \Sigma$, we say $B_{x,r}$ is a medial ball at p .

Figure 19(a) shows a subset of the medial axis of a curve. Notice that the medial axis may have a branching point such as v and boundary points such as u and w . Also, the medial axis need not be connected. For example, the part of the medial axis in the region bounded by the curve may be disjoint from the rest, see Figure 19(a). In fact, if Σ is C^2 -smooth, the two parts of the medial axis are indeed disjoint. The subset of the medial axis residing in the unbounded component of $\mathbb{R}^2 \setminus \Sigma$ is called the *outer* medial axis. The rest is called the *inner* medial axis.

It follows from the definition that if one grows a ball around a point on the medial axis, it will meet Σ for the first time tangentially in one or more points, see Figure 19(b). Conversely, for a point $x \in \Sigma$ one can start growing a ball keeping it tangent to Σ at x until it hits another point $y \in \Sigma$ or becomes maximally empty. At this moment the ball is medial and the segments joining the center m to x and y are normal to Σ at x and y respectively, see Figure 19.

If we move along the medial axis and consider the medial balls as we move, the radius of the medial balls increases or decreases accordingly to maintain the tangency with Σ . At the boundaries it coincides with the radius of the *curvature ball* where all tangent points merge into a single one. See Figure 19(b).

It will be useful for our proofs later to know the following property of balls intersecting the sampled space Σ . The proof of the lemma assumes that Σ is either a smooth curve or a smooth surface whose definitions are given in later chapters. Also, the proof uses some concepts from differential topology (critical point theory). The readers may skip the proof at this point if they are not familiar with these concepts.

³see the definition of C^i -smoothness for curves in the next chapter

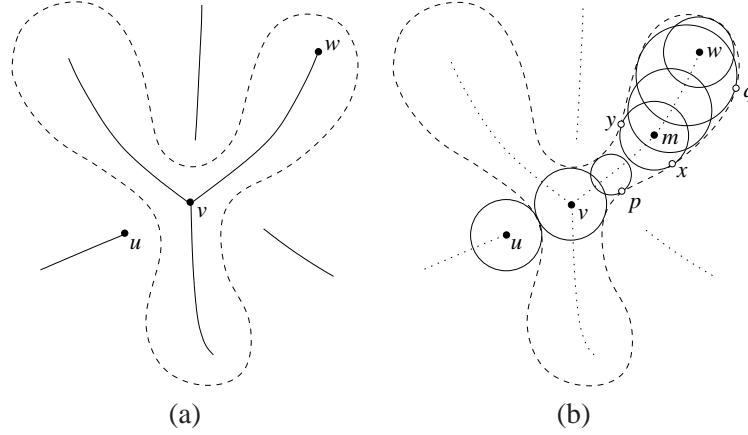


Figure 19: (a) A subset of the medial axis of the curve in Figure 17, (b) medial ball centered at v touches the curve in three points, whereas the ones with centers u and w touch it in only one point and coincide with the curvature ball.

We say that a topological space is a k -ball or a k -sphere if it is homeomorphic to \mathbb{B}^k or \mathbb{S}^k respectively.

Lemma 11 (Feature Ball.). *If a d -ball $B = B_{c,r}$ intersects a k -manifold $\Sigma \subset \mathbb{R}^d$ at more than one point where either (i) $B \cap \Sigma$ is not a k -ball or (ii) $\text{Bd}(B \cap \Sigma)$ is not a $(k-1)$ -sphere, then a medial axis point is in B .*

PROOF. First we show that if B intersects Σ at more than one point and B is tangent to Σ at some point, B contains a medial axis point. Let x be the point of this tangency. Shrink B further keeping it tangent to Σ at x . This means the center of B moves towards x along a normal direction at x . We stop when B meets Σ only tangentially. Observe that, since $B \cap \Sigma \neq x$ to start with, this happens eventually when B is maximally empty. At this moment B becomes a medial ball and its center is a medial axis point which must lie in the original ball B , refer to Figure 20.

Now consider when condition (ii) holds. Define a function $h: B \cap \Sigma \rightarrow \mathbb{R}$ where $h(x)$ is the distance of x from the center c of B . The function h is a scalar function defined over a smooth manifold. At the critical points of h where its gradient vanishes the ball B becomes tangent to Σ when shrunk appropriately.

Let m be a point in Σ so that $h(m)$ is a global minimum. If there is more than one such global minimum, the ball B meets Σ only tangentially at more than one point when radially shrunk to a radius of $h(m)$. Then, B becomes a medial ball which implies that the original B contains a medial axis point, namely its center. So, assume that there is only global minimum m of h .

We claim that the function h has a critical point p in $\text{Int}(B \cap \Sigma)$ other than m where B becomes tangent to Σ . If not, as we shrink B centrally the level set $\text{Bd}(B \cap \Sigma)$ does not change topology until it reaches the minimum m when it vanishes. This follows from the Morse theory of smooth functions over smooth manifolds. Since m is a minimum, there is a small enough $\delta > 0$ so that $B_{c,h(m)+\delta} \cap \Sigma$ is a k -ball. The boundary of this k -ball given by $(\text{Bd } B_{c,h(m)+\delta}) \cap \Sigma$ should be a $(k-1)$ -sphere. This contradicts the fact that $\text{Bd}(B \cap \Sigma)$ is not a $(k-1)$ -sphere and remains that way till the end. Therefore, there is a critical point, say $y \neq m$ of h in $\text{Int}(B \cap \Sigma)$. At this point y , the ball




Figure 20: (a) The ball B intersecting the upper right lobe is shrunk till it becomes tangent to another point other than x . The new ball B' intersects the medial axis. (b) The ball B intersecting the lower lobe is shrunk radially to the ball B' that is tangent to the curve at y and also intersects the curve in other points. B' can further be shrunk till it meets the curve only tangentially.

$B_{c, \|y-c\|}$ becomes tangent to Σ , see also Figure 20. Now we can apply our previous argument to claim that B contains a medial axis point.

Next, consider when condition (i) holds. If condition (ii) also holds, we have the previous argument. So, assume that $\text{Bd}(B \cap \Sigma)$ is a $(k-1)$ -sphere and $B \cap \Sigma$ is not a k -ball. Again, we claim that the function h as defined earlier has a critical point other than m . If not, consider the subset of Σ swept by B while shrinking it till it meets Σ only at m . This subset is homeomorphic to a space which is formed by taking the product of \mathbb{S}^{k-1} with the closed unit interval I in \mathbb{R} and then collapsing one of its boundary to a single point, i.e. the quotient space $(\mathbb{S}^{k-1} \times I)/(\mathbb{S}^{k-1} \times \{0\})$. This space is a k -ball which contradicts the fact that $B \cap \Sigma$ is not a k -ball to begin with. Therefore, as B is continually shrunk, it becomes tangent to Σ at a point $y \neq m$. Apply the previous argument to claim that B has a medial axis point.

□

□

Figure 21 illustrates the different cases of Feature Ball Lemma in \mathbb{R}^2 .

18.2 Local feature size

The medial axis M with the distance to Σ at each point $m \in M$ captures the shape of Σ . In fact, Σ is the boundary of the union of all medial balls centering points of the inner (or outer) medial axis. So, as a first attempt to capture local feature size one may define the following two functions on Σ .

$\rho_i, \rho_o : \Sigma \rightarrow \mathbb{R}$ where $\rho_i(x), \rho_o(x)$ are the radii of the inner and outer medial balls respectively both of which are tangent to Σ at x .




Figure 21: (a) $B \cap \Sigma$ is not a 1-ball, (b) $B \cap \Sigma$ is a 1-ball, but $\text{Bd } B \cap \Sigma$ is not a 0-sphere, (c) $\text{Bd } B \cap \Sigma$ is a 0-sphere, but $B \cap \Sigma$ is not a 1-ball.

The functions ρ_i and ρ_o are continuous for a large class of curves and surfaces. However, we need a stronger form of continuity on the local feature size function to carry out the proofs. This property, called the *Lipschitz property*, stipulates that the difference in the function values at two points is bounded by a constant times the distance between the points. Keeping this in mind we define the following.

Definition 45. The *local feature size* $f(x)$ at a point $x \in \Sigma$ is the distance of $x \in \Sigma$ to the medial axis M , that is, $f(x) = d(x, M)$.

Figure 22 illustrates how the local feature size can vary over a shape. As one can observe, the local feature sizes at the leg and tail are much smaller than the local feature sizes at the middle in accordance with our intuitive notion of features. For example, $f(b)$ is much smaller than $f(a)$. Local feature size can be determined either by the inner or outer medial axis. For example, $f(c)$ is determined by the outer medial axis whereas $f(d)$ is determined by the inner one.

It follows from the definitions that $f(x) \leq \min\{\rho_i(x), \rho_o(x)\}$. In Figure 22, $f(d)$ is much smaller than the radius of the drawn medial ball at d . Lipschitz property of the local feature size function f follows easily from the definition.

Lemma 12 (Lipschitz Continuity.). $f(x) \leq f(y) + \|x - y\|$ for any two points x and y in Σ .

PROOF. Let m be a point on the medial axis so that $f(y) = \|y - m\|$. By triangular inequality,

$$\begin{aligned} \|x - m\| &\leq \|y - m\| + \|x - y\|, \text{ and} \\ f(x) &\leq \|x - m\| \leq f(y) + \|x - y\|. \end{aligned}$$


□
□

Figure 22: Local feature sizes $f(a)$, $f(b)$, $f(c)$, and $f(d)$ are the lengths of the corresponding dotted line segments.

18.3 Sampling

A *sample* P of Σ is a set of points from Σ . Once we have quantized the feature size, we would require the sample respect the features, i.e., we require more sample points where the local feature size is small compared to the regions where it is not.

Definition 46. A sample P of Σ is an ε -sample if each point $x \in \Sigma$ has a sample point $p \in P$ so that $\|x - p\| \leq \varepsilon f(x)$.

The value of ε has to be smaller than 1 to have a dense sample. In fact, practical experiments suggest that $\varepsilon < 0.4$ constitutes a dense sample for reconstructing Σ from P . An ε -sample is also an ε' -sample for any $\varepsilon' > \varepsilon$. The definition of ε -sample allows a sample to be arbitrarily dense anywhere on Σ . It only puts a lower bound on the density. Figure 23 illustrates a sample of a circle which is a 0.2-sample. By definition, it is also a 0.3-sample of the same.

A useful application of the Lipschitz Continuity Lemma 12 is that the distance between two points expressed in terms of the local feature size of one can be expressed in terms of that of the other.

Lemma 13 (Feature Translation.). *For any two points x, y in Σ with $\|x - y\| \leq \varepsilon f(x)$ and $\varepsilon < 1$ we have*

$$(i) \quad f(x) \leq \frac{1}{1-\varepsilon} f(y) \text{ and}$$




Figure 23: Local feature size at any point on the circle is equal to the radius r . Each point on the circle has a sample point within $0.2r$ distance.

$$(ii) \|x - y\| \leq \frac{\varepsilon}{1-\varepsilon} f(y).$$

PROOF. We have

$$\begin{aligned} f(x) &\leq f(y) + \|x - y\| \\ \text{or, } f(x) &\leq f(y) + \varepsilon f(x). \end{aligned}$$

For $\varepsilon < 1$ the above inequality gives

$$f(x) \leq \frac{1}{1-\varepsilon} f(y) \text{ proving (i).}$$

Plug the above inequality in $\|x - y\| \leq \varepsilon f(x)$ to obtain (ii). □

□

Uniform sampling. The definition of ε -sample allows non-uniform sampling over Σ . A *globally uniform* sampling is more restrictive. It means that the sample is equally dense everywhere. Local feature size does not play a role in such sampling. There could be various definitions of globally uniform samples. We will say a sample $P \subset \Sigma$ is *globally δ -uniform* if any point $x \in \Sigma$ has a point in P within $\delta > 0$ distance. In between globally uniform and non-uniform samplings, there is another one called the *locally uniform sampling*. This sampling respects feature sizes and is uniform only locally. We say $P \subset \Sigma$ is *locally (ε, δ) -uniform* for $\delta > 1 > \varepsilon > 0$ if each point $x \in \Sigma$ has a point in P within $\varepsilon f(x)$ distance and no point $p \in P$ has another point $q \in P$ within $\frac{\varepsilon}{\delta} f(p)$ distance. This definition does not allow two points to be arbitrarily close which may become a severe restriction for sampling in practice. So, there is an alternate definition of local uniformity. A sample P is *locally (ε, κ) -uniform* for some $\varepsilon > 0$ and $\kappa \geq 1$ if each point $x \in \Sigma$ has at least one and no more than κ points within $\varepsilon f(x)$ distance.

$\tilde{O}(\varepsilon)$ notation Our analysis for different algorithms obviously involve the sampling parameter ε . To ease these analyses, sometimes we resort to \tilde{O} notation which provides the asymptotic dependences on ε . A value is $\tilde{O}(\varepsilon)$ if there exist two constants $\varepsilon_0 > 0$ and $c > 0$ so that the value is less than $c\varepsilon$ for any positive $\varepsilon \leq \varepsilon_0$. Notice that \tilde{O} notation is slightly different from the well known big- O notation since the latter would require ε greater than or equal to ε_0 .

19 Voronoi diagram and Delaunay triangulation

Voronoi diagrams and Delaunay triangulations are important geometric data structures that are built on the notion of ‘nearness’. Many differential properties of curves and surfaces are defined on local neighborhoods. Voronoi diagrams and their duals, Delaunay triangulations, provide a tool to approximate these neighborhoods in the discrete domain. They are defined for a point set in any Euclidean space. We define them in two dimensions and mention the extensions to three dimensions since the curve and surface reconstruction algorithms as dealt here are concerned with these two Euclidean spaces. Before the definitions we state a non-degeneracy condition for the point set P defining the Voronoi and Delaunay diagrams. This non-degeneracy condition not only makes the definitions less complicated but also makes the algorithms avoid special cases.

Definition 47. A point set $P \subset \mathbb{R}^k$ is *non-degenerate* if (i) the affine hull of any ℓ points from P with $1 \leq \ell \leq k$ is homeomorphic to $\mathbb{R}^{\ell-1}$ and (ii) no $k + 2$ points are co-spherical.

19.1 Two dimensions

Let P be a set of non-degenerate points in the plane \mathbb{R}^2 .

Voronoi diagrams. The Voronoi cell V_p for each point $p \in P$ is given as

$$V_p = \{x \in \mathbb{R}^2 \mid d(x, P) = \|x - p\|\}.$$

In words, V_p is the set of all points in the plane that have no other point in P closer to it than p . For any two points p, q the set of points closer to p than q are demarked by the perpendicular bisector of the segment pq . This means the Voronoi cell V_p is the intersection of the closed halfplanes determined by the perpendicular bisectors between p and each other point $q \in P$. An implication of this observation is that each Voronoi cell is a convex polygon since the intersection of convex sets remains convex.

Voronoi cells have *Voronoi faces* of different dimensions. A Voronoi face of dimension k is the intersection of $3 - k$ Voronoi cells. This means a k -dimensional Voronoi face for $k \leq 2$ is the set of all points that are equidistant from $3 - k$ points in P . A zero dimensional Voronoi face, called *Voronoi vertex* is equidistant from three points in P , whereas an one dimensional Voronoi face, called *Voronoi edge* contains points that are equidistant from two points in P . A Voronoi cell is a two dimensional Voronoi face.

Definition 48. The Voronoi diagram $\text{Vor } P$ of P is the cell complex formed by Voronoi faces.

Figure 24(a) shows a Voronoi diagram of a point set in the plane where u and v are two Voronoi vertices and uv is a Voronoi edge.

Some of the Voronoi cells may be unbounded with unbounded edges. It is a straightforward consequence of the definition that a Voronoi cell V_p is unbounded if and only if p is on the boundary of the convex hull of P . In Figure 24(a) V_p and V_q are unbounded and p and q are on the convex hull boundary.




Figure 24: (a) The Voronoi diagram and (b) the Delaunay triangulation of a point set in the plane.

Delaunay triangulations. There is a *dual* structure to the Voronoi diagram $\text{Vor } P$, called the *Delaunay triangulation*.

Definition 49. The Delaunay triangulation of P is a simplicial complex

$$\text{Del } P = \{\sigma = \text{conv}\{T\} \mid \bigcap_{p \in T \subseteq P} V_p \neq \emptyset\}.$$

In words, $k + 1$ points in P form a Delaunay k -simplex in $\text{Del } P$ if their Voronoi cells have nonempty intersection. We know that $k + 1$ Voronoi cells meet in a $(2 - k)$ -dimensional Voronoi face. So, each k -simplex in $\text{Del } P$ is dual to a $(2 - k)$ -dimensional Voronoi face. Thus, each Delaunay triangle pqr in $\text{Del } P$ is dual to a Voronoi vertex where V_p , V_q , and V_r meet, each Delaunay edge pq is dual to a Voronoi edge shared by Voronoi cells V_p and V_q , and each vertex p is dual to its corresponding Voronoi cell V_p . In Figure 24(b), the Delaunay triangle pqr is dual to the Voronoi vertex v and the Delaunay edge pr is dual to the Voronoi edge uv . In general, when μ is a dual Voronoi face of a Delaunay simplex σ we say $\mu = \text{dual } \sigma$ and conversely $\sigma = \text{dual } \mu$.

A *circumscribing ball* of a simplex σ is a ball whose boundary contains the vertices of the simplex. The smallest circumscribing ball of σ is called its *diametric ball*. A triangle in the plane has only one circumscribing ball, namely the diametric one. However, an edge has infinitely many circumscribing balls among which the diametric one is unique, namely the one with the center on the edge.

A dual Voronoi vertex of a Delaunay triangle is equidistant from its three vertices. This means that the center of the circumscribing ball of a Delaunay triangle is the dual Voronoi vertex. It implies that no point from P can lie in the interior of the circumscribing ball of a Delaunay triangle. These balls are called *Delaunay*. A ball is *empty* if its interior does not contain any point from P . Clearly, the Delaunay balls are *empty*. The converse also holds.

Property 1 (Triangle emptiness.). *A triangle is in the Delaunay triangulation if and only if its circumscribing ball is empty.*

The Triangle Emptiness Property of Delaunay triangles also implies a similar emptiness for Delaunay edges. Clearly, each Delaunay edge has an empty circumscribing ball passing through

its endpoints. It turns out that the converse is also true, that is, any edge pq with an empty circumscribing ball must also be in the Delaunay triangulation. To see this, grow the empty ball of pq always keeping p, q on its boundary. If it never meets any other point from P , the edge pq is on the boundary of $\text{conv}\{P\}$ and is in the Delaunay triangulation since V_p and V_q has to share an edge extending to infinity. Otherwise, when it meets a third point, say r from P , we have an empty circumscribing ball passing through p, q , and r . By the Triangle Emptiness Property pqr must be in the Delaunay triangulation and hence the edge pq .

Property 2 (Edge emptiness.). *An edge is in the Delaunay triangulation if and only if the edge has an empty circumscribing ball.*

The Delaunay triangulation form a planar graph since no two Delaunay edges intersect in their interiors. It follows from the property of planar graphs that the number of Delaunay edges is at most $3n - 6$ for a set of n points. The number of Delaunay triangles is at most $2n - 4$. This means that the dual Voronoi diagram also has at most $3n - 6$ Voronoi edges and $2n - 4$ Voronoi vertices. The Voronoi diagram and the Delaunay triangulation of a set of n points in the plane can be computed in $O(n \log n)$ time and $O(n)$ space.

Restricted Voronoi diagrams. When the input point set P is a sample of a curve or a surface Σ , the Voronoi diagram $\text{Vor } P$ imposes a structure on Σ . It turns out that this diagram plays an important role in reconstructing Σ from P . Formally, a restricted Voronoi cell $V_p|_\Sigma$ is defined as the intersection of the Voronoi cell V_p in $\text{Vor } P$ with Σ , i.e.,

$$V_p|_\Sigma = V_p \cap \Sigma \text{ where } p \in P.$$




Figure 25: (a) Restricted Voronoi diagram for a point set on a curve, (b) the corresponding restricted Delaunay triangulation.

Similar to the Voronoi faces, we can define *restricted Voronoi faces* as the intersection of the restricted Voronoi cells. They can also be viewed as the intersection of Voronoi faces with Σ . In Figure 25(a) the white circles represent restricted Voronoi faces of dimension zero. The curve segments between them are restricted Voronoi faces of dimension one which are restricted Voronoi cells in this case. Notice that the restricted Voronoi cell $V_p|_\Sigma$ in the figure consists of two curve segments whereas $V_r|_\Sigma$ consists of a single curve segment.

Definition 50. The restricted Voronoi diagram $\text{Vor}|_\Sigma P$ of P with respect to Σ is the collection of all restricted Voronoi faces.

Restricted Delaunay triangulations. As with Voronoi diagrams we can define a simplicial complex dual to a restricted Voronoi diagram $\text{Vor } P|_\Sigma$.

Definition 51. The restricted Delaunay triangulation of P with respect to Σ is a subcomplex $\text{Del } P|_\Sigma$ of the Delaunay complex $\text{Del } P$ where a Delaunay simplex $\sigma \in \text{Del } P|_\Sigma$ if and only if the dual Voronoi face of σ intersects Σ .

Figure 25(b) shows the restricted Delaunay triangulation for the restricted Voronoi diagram in (a). The vertex p is connected to q and r in the restricted Delaunay triangulation since $V_p|_\Sigma$ meets both $V_q|_\Sigma$ and $V_r|_\Sigma$. However, the triangle pqr is not in the triangulation since $V_p|_\Sigma$, $V_q|_\Sigma$ and $V_r|_\Sigma$ do not meet at a point.

19.2 Three dimensions

We chose the plane to explain the concepts of the Voronoi diagrams and the Delaunay triangulations in the previous subsection. However, these concepts extend to arbitrary dimensions. We will mention these extensions for three dimensions which will be important for later expositions.

Voronoi cells of a point set P in \mathbb{R}^3 are three dimensional convex polytopes some of which are unbounded. There are four types of Voronoi faces; Voronoi vertices, Voronoi edges, Voronoi facets, and Voronoi cells in increasing order of dimension starting with zero and ending with three. Four Voronoi cells meet at a Voronoi vertex which is equidistant from four points in P . Three Voronoi cells meet at a Voronoi edge, and two Voronoi cells meet at a Voronoi facet.

The Delaunay triangulation of P contains four types of simplices dual to each of the four types of Voronoi faces. The vertices are dual to the Voronoi cells, the Delaunay edges are dual to the Voronoi facets, the Delaunay triangles are dual to the Voronoi edges, and the Delaunay tetrahedra are dual to the Voronoi vertices. The circumscribing ball of each tetrahedron is empty. Conversely, any tetrahedron with empty circumscribing ball is in the Delaunay triangulation. Further, each Delaunay triangle and edge has an empty circumscribing ball. Conversely, an edge or a triangle belongs to the Delaunay triangulation if there exists an empty ball circumscribing it.

The number of edges, triangles, and tetrahedra in the Delaunay triangulation of a set of n points in three dimensions can be $O(n^2)$ in the worst case. By duality the Voronoi diagram can also have $O(n^2)$ Voronoi faces. Both of the diagrams can be computed in $O(n^2)$ time and space.

We can define the restricted Voronoi diagram and its dual restricted Delaunay triangulation for a point sample on a surface in \mathbb{R}^3 in the same way as we did for a curve in \mathbb{R}^2 . Figure 26 shows the restricted Voronoi diagram and its dual restricted Delaunay triangulation for a set of points on a surface. The triangle pqr is in the restricted Delaunay triangulation since $V_p|_\Sigma$, $V_q|_\Sigma$, and $V_r|_\Sigma$ meet at a common point v .

Exercises

1. Construct a triangulation of \mathbb{S}^2 and verify that $v - e + f = 2$ where v is the number of vertices, e is the number of edges, and f is the number of triangles. Prove that the number $v - e + f$ (Euler characteristic) is always 2 for any triangulation of \mathbb{S}^2 .
2. Let p be a vertex in $\text{Del } P$ in three dimensions. Show that a point $x \in V_p$ if and only if $\|p - x\| \leq \|q - x\|$ for each vertex q where pq is a Delaunay edge.




Figure 26: (a) The restricted Voronoi diagram and (b) the restricted Delaunay triangulation for a sample on a surface.

3. Show that for any Delaunay simplex σ and its dual Voronoi face $\mu = \text{dual } \sigma$, the affine hulls $\text{aff } \mu$ and $\text{aff } \sigma$ intersect orthogonally.
4. An edge e in a triangulation $T(P)$ of a point set $P \subset \mathbb{R}^2$ is called *locally Delaunay* if e is a convex hull edge or the circumscribing ball of one triangle incident to e does not contain the other triangle incident to e completely inside. Show that $T(P) = \text{Del } P$ if and only if each edge of $T(P)$ is locally Delaunay.
5. Given a point set $P \subset \mathbb{R}^2$, an edge connecting two points p, q in P is called a nearest neighbor edge if no point in P is closer to q than p is. Show that pq is a Delaunay edge.
6. Given a point set $P \subset \mathbb{R}^2$, an edge connecting two points in P is called *Gabriel* if its diametric ball is empty. The Gabriel graph for P is the graph induced by all Gabriel edges. Give an $O(n \log n)$ algorithm to compute the Gabriel graph for P where P has n points.
7. Let pq be a Delaunay edge in $\text{Del } P$ for a point set $P \subset \mathbb{R}^3$. Show that if pq does not intersect its dual Voronoi facet $g = \text{dual } pq$, the line of pq does not intersect g either.
8. For $\alpha > 0$, a function $f: \Sigma \rightarrow \mathbb{R}$ is called α -Lipschitz if $f(x) \leq f(y) + \alpha\|x - y\|$ for any two points x, y in Σ . Given an arbitrary function $f: \Sigma \rightarrow \mathbb{R}$, one can make it α -Lipschitz by considering the functions

$$f_m(x) = \min_{p \in \Sigma} \{f(p) + \alpha\|x - p\|\},$$

$$f_M(x) = \max_{p \in \Sigma} \{f(p) - \alpha\|x - p\|\}.$$

Show that both f_m and f_M are α -Lipschitz.

9. Consider the functions ρ_i and ρ_o as in Section 18.2. Show that these functions may be continuous but not 1-Lipschitz.

References

- [1] T. K. Dey. Curve and Surface Reconstruction: Algorithms with Mathematical Analysis. Cambridge U. Press, 2007.