Basics of Surface Reconstruction _{taken from} [3]

23 Surface Samples

In this chapter we introduce some of the properties of surfaces and their samples in three dimensions. The results developed in this chapter are used in later chapters to design algorithms for surface reconstruction and prove their guarantees. Before we talk about these results, let us explain what we mean by smooth surfaces.

Consider a map $\pi\colon U \to V$ where U and V are the open sets in $\mathbb{R}^2$ and $\mathbb{R}^3$ respectively. The map π has three components, namely $\pi(x) = (\pi_1(x), \pi_2(x), \pi_3(x))$ where $x = (x_1, x_2)$ is a point in $\mathbb{R}^2$. The three by two matrix of first order partial derivatives $(\frac{\partial \pi_i(x)}{\partial x_j})_{i,j}$ is called the *Jacobian* of π at x. We say π is *regular* if its Jacobian at each point of U has rank 2. The map π is C^i-continuous if the ith order $(i > 0)$ partial derivatives of π are continuous.

For $i > 0$, a subset $\Sigma \subset \mathbb{R}^3$ is a C^i-*smooth* surface if each point $x \in \Sigma$ satisfies the following condition. There is a neighborhood $W \subset \mathbb{R}^3$ of x and a map $\pi\colon U \to W \cap \Sigma$ of an open set $U \subset \mathbb{R}^2$ onto $W \cap \Sigma$ so that

(i) π is C^i-continuous,

(ii) π is a homeomorphism, and

(iii) π is regular.

The first condition says that π is continuously differentiable at least up to ith order. The second condition imposes one-to-one property which eliminates self intersections of Σ. The third condition together with the first actually enforce the smoothness. It makes sure that the tangent plane at each point in Σ is well defined. All of these three conditions together imply that the functions like π defined in the neighborhood of each point of Σ overlap smoothly. There are two extremes of smoothness. If the partial derivatives of π of all orders are continuous, we say Σ is C^∞-smooth. On the other hand if Σ is not C^1-smooth but is at least a 2-manifold, we say it is C^0-smooth or *nonsmooth*.

In this chapter and the chapters to follow, we assume that Σ is a C^2-smooth surface. Notice that, by the definition of smoothness (condition (ii)) Σ is a 2-manifold without boundary. We also assume that Σ is compact since we are interested in approximating Σ with a finite simplicial complex. We need one more assumption. Just like the curves, for a finite point set to be an ε-sample for some $\varepsilon > 0$, we need that $f(x) > 0$ for any point x in Σ. It is known that C^2-smooth surfaces necessarily have positive feature size everywhere. The example in Chapter **??** for curves can be extended to surfaces to claim that a C^1-smooth surface may not have positive local feature sizes everywhere.

As a C^2-smooth surface Σ has a tangent plane τ_x and a normal $\mathbf{n}_x$ defined at each point $x \in \Sigma$. We assume that the normals are oriented outward. More precisely, $\mathbf{n}_x$ points locally to the unbounded component of $\mathbb{R}^3 \setminus \Sigma$. If Σ is not connected, $\mathbf{n}_x$ points locally to the unbounded component of $\mathbb{R}^3 \setminus \Sigma'$ where x is in Σ', a connected component of Σ.

An important fact used in surface reconstruction is that, disregarding the orientation, the direction of the surface normals can be approximated from the sample. An illustration in $\mathbb{R}^2$ is helpful here. See Figure 30 in Chapter **??** which shows the Voronoi diagram of a dense sample on a smooth curve. This Voronoi diagram has a specific structure. Each Voronoi cell is elongated along the normal direction at the sample points. Fortunately, the same holds in three dimensions. The three dimensional Voronoi cells are long, thin, and the direction of the elongation matches with the normal direction at the sample points when the sample is dense, see Figure 36.
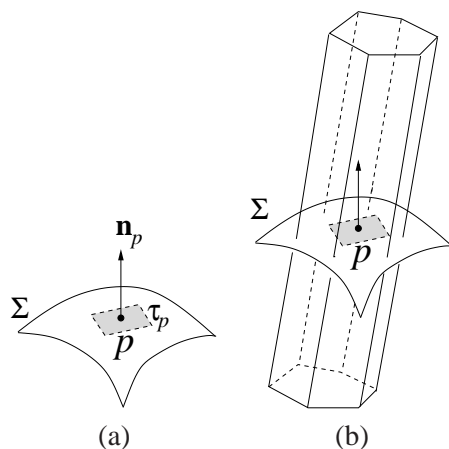

(a)   (b)

Figure 36: (a) Tangent plane and the normal at a point on a smooth surface, (b) a long thin Voronoi cell elongated along the normal direction.

24   Normals

Let $P \subset \mathbb{R}^3$ be an ε-sample of Σ. If P is all we know about Σ, it is impossible to know the line of direction of $\mathbf{n}_p$ exactly at a point $p \in P$. However, it is conceivable that as P gets denser, we should have more accurate idea about the direction of $\mathbf{n}_p$ by looking at the adjacent points. This is what is done using the Voronoi cells in Vor P.

For further developments we will often need to talk about how one vector approximates another one in terms of the angles between them. We denote the angle between two vectors $\mathbf{u}$ and $\mathbf{v}$ as $\angle(\mathbf{u}, \mathbf{v})$. For vector approximations that disregard the orientation, we use a slightly different notation. This approximation measures the acute angle between the line containing the vectors. We use $\angle_a(\mathbf{u}, \mathbf{v})$ to denote this acute angle between two vectors $\mathbf{u}$ and $\mathbf{v}$. Since any such angle is acute, we have the triangular inequality $\angle_a(\mathbf{u}, \mathbf{v}) \leq \angle_a(\mathbf{u}, \mathbf{w}) + \angle_a(\mathbf{v}, \mathbf{w})$ for any three vectors $\mathbf{u}, \mathbf{v}$ and $\mathbf{w}$.

24.1   Approximation of normals

It turns out that the structure of the Voronoi cells contains information about normals. Indeed, if the sample is sufficiently dense, the Voronoi cells become long and thin along the direction of the normals at the sample points. One reason for this structural property is that a Voronoi cell V_p

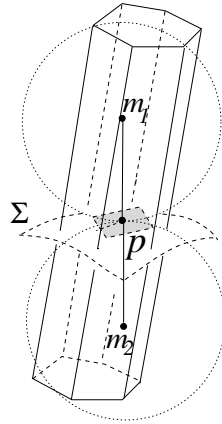must contain the medial axis points that are the centers of the medial balls tangent to Σ at p, see Figure 37.

Figure 37: Medial axis points m_1 and m_2 are in the Voronoi cell V_p.

Lemma 25 (Medial.)**.** *Let m_1 and m_2 be the centers of the two medial balls tangent to Σ at p. The Voronoi cell V_p contains m_1 and m_2.*

P   . Denote the medial ball with center m_1 as B. The ball B meets the surface Σ only tangentially at points, one of which is p. Thus, B is empty of any point from Σ and P in particular. Therefore, the center m_1 has p as the nearest point in P. By definition of Voronoi cells, m_1 is in V_p. A similar argument applies to the other medial axis point m_2.   □

□

We have already mentioned that the Voronoi cells are long and thin and they are elongated along the direction of the normals. The next lemma formalizes this statement by asserting that as we go further from p within V_p, the direction to p becomes closer to the normal direction.

Lemma 26 (Normal.)**.** *For $\mu > 0$ let $v \notin \Sigma$ be a point in V_p with $\|v - p\| > \mu f(p)$. For $\varepsilon < 1$, $\angle_a(\overrightarrow{vp}, \mathbf{n}_p) \leq arcsin\frac{\varepsilon}{\mu(1-\varepsilon)} + \arcsin\frac{\varepsilon}{1-\varepsilon}$.*

P   . Let m_1 and m_2 be the two centers of the medial balls tangent to Σ at p where m_1 is on the same side of Σ as v is. Both m_1 and m_2 are in V_p by the Medial Lemma 25. The line joining m_1 and p is normal to Σ at p by the definition of medial balls. Similarly, the line joining m_2 and p is also normal to Σ at p. Therefore, $m_1, m_2,$ and p are co-linear. See Figure 38. Consider the triangle pvm_2. We are interested in the angle $\angle m_1 pv$ which is equal to $\angle_a(\overrightarrow{pv}, \mathbf{n}_p)$. From the triangle pvm_2 we have

$$\angle m_1 pv = \angle pvm_2 + \angle vm_2 p.$$

To measure the two angles on the righthand side, drop the perpendicular px from p onto the segment vm_2. The line segment vm_2 intersects Σ, say at y, since m_1 and m_2 and hence v and m_2 lie on opposite sides of Σ. Furthermore, y must lie inside V_p since any point on the segment joining
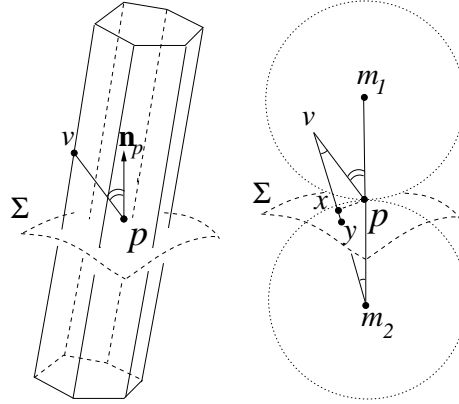
Figure 38: Illustration for the Normal Lemma 26.

two points v and m_2 in a convex set V_p must lie within the same convex set. This means y has p as the nearest sample point and thus

$$\|x - p\| \leq \|y - p\| \leq \varepsilon f(y) \text{ by the } \varepsilon\text{-sampling condition.}$$

Using the Feature Translation Lemma 13 we get

$$\|x - p\| \leq \frac{\varepsilon}{1 - \varepsilon} f(p)$$

when $\varepsilon < 1$. We have

$$\angle pvm_2 \quad = \quad \arcsin \frac{\|x - p\|}{\|v - p\|} \leq \arcsin \frac{\varepsilon}{\mu(1 - \varepsilon)} \text{ as } \|v - p\| \geq \mu f(p).$$

Similarly,

$$\angle vm_2p \quad = \quad \arcsin \frac{\|x - p\|}{\|m_2 - p\|} \leq \arcsin \frac{\varepsilon}{1 - \varepsilon} \text{ as } \|m_2 - p\| \geq f(p).$$

The assertion of the lemma follows immediately.   □

                                                                                        □

24.2   Normal variation

The directions of the normals at nearby points on Σ cannot vary too abruptly. In other words, the surface looks flat locally. This fact is used later in many proofs.

Lemma 27 (Normal Variation.)**.** *If* $x, y \in \Sigma$ *are any two points with* $\|x - y\| \leq \rho f(x)$ *for* $\rho < \frac{1}{3}$, $\angle(\mathbf{n}_x, \mathbf{n}_y) \leq \frac{\rho}{1 - 3\rho}$.

P   . See the book [1] on Delaunay mesh generation for a proof.   □

                                                                                        □

24.3   Edge and triangle normals

In Section 20, we saw that edges joining nearby points on a curve are almost parallel to the tangents at the endpoints of the edge. Similar results also hold for triangles connecting points on surfaces. But, the size is measured by circumradius. In fact, a triangle connecting three nearby points on a surface but with a large circumradius may lie almost perpendicular to the surface. However, if its circumradius is small compared to the local feature sizes at its vertices, it has to lie almost parallel to the surface. For an edge, half of its length is the same as its circumradius. Therefore, a small edge lies almost parallel to the surface. In essence if an edge or a triangle has a small circumradius, it must lie flat to the surface. We quantify these claims in the next two lemmas.

Lemma 28 (Edge Normal.). *For an edge pq with* $\|p-q\| \leq 2f(p)$*, the angle* $\angle_a(\overrightarrow{pq}, \mathbf{n}_p)$ *is at least* $\frac{\pi}{2} - \arcsin \frac{\|p-q\|}{2f(p)}$*.*

P   .   Consider the two medial balls sandwiching the surface Σ at p. The point q cannot lie
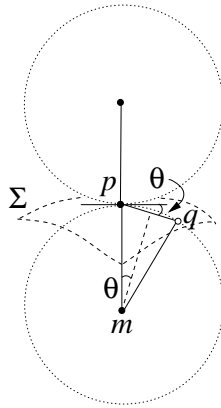

Figure 39: Illustration for the Edge Normal Lemma 28.

inside any of these two balls as they are empty of points from Σ. So, the smallest angle pq makes with $\mathbf{n}_p$ cannot be smaller than the angle pq makes with $\mathbf{n}_p$ when q is on the boundary of any of these two balls. In this case let θ be the angle between pq and the tangent plane at p. Clearly, (see Figure 39)

$$
\begin{aligned}
\sin \theta &= \frac{\|p-q\|}{2\|m-p\|} \\
&\leq \frac{\|p-q\|}{2f(p)}.
\end{aligned}
$$

Therefore,

$$
\begin{aligned}
\angle_a(\overrightarrow{pq}, \mathbf{n}_p) &= \frac{\pi}{2} - \theta \\
&\geq \frac{\pi}{2} - \arcsin \frac{\|p-q\|}{2f(p)}.
\end{aligned}
$$

□

□

It follows immediately from the Edge Normal Lemma 28 that small edges make a large angle with the surface normals at the vertices. For example, if pq has a length less than $\rho f(p)$ for $\rho < 2$, the angle $\angle_a(\vec{pq}, \mathbf{n}_p)$ is more than $\frac{\pi}{2} - \arcsin \frac{\rho}{2}$.

Next consider a triangle $t = pqr$ where p is the vertex subtending a maximal angle in pqr. Let R_{pqr} denote the circumradius of pqr.

Lemma 29 (Triangle Normal.). *If $R_{pqr} \leq \frac{f(p)}{\sqrt{2}}$,*

$$\angle_a(\mathbf{n}_{pqr}, \mathbf{n}_p) \leq \arcsin \frac{R_{pqr}}{f(p)} + \arcsin \left(\frac{2}{\sqrt{3}} \sin \left(2 \arcsin \frac{R_{pqr}}{f(p)} \right) \right)$$

where $\mathbf{n}_{pqr}$ is the normal of pqr.

P   . Consider the medial balls $B = B_{m,\ell}$ and $B' = B_{m',\ell'}$ that are tangent to Σ at p. Let D be the diametric ball of t (smallest circumscribing ball); refer to Figure 40. The radius of D is R_{pqr}. Let C and C' be the circles in which the boundary of D intersects the boundaries of B and B' respectively. The line normal to Σ at p passes through m, the center of B. Let α be the larger of the two angles this normal line makes with the normals to the planes containing C and C'. Since the radii of C and C' are at most R_{pqr} we have

$$\alpha \leq \arcsin \frac{R_{pqr}}{\|p - m\|} \leq \arcsin \frac{R_{pqr}}{f(p)}.$$

It follows from the definition of α that the planes containing C and C' make a wedge, say W, with an acute dihedral angle no more than 2α.
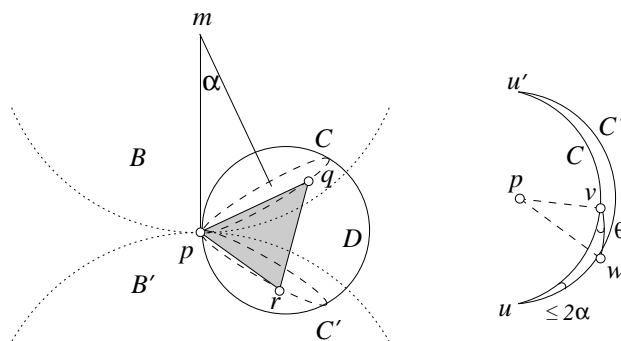

Figure 40: Illustration for the Triangle Normal Lemma 29. The two great arcs on the right picture are the intersections of the unit sphere with the planes containing C and C'.

The other two vertices q, r of t cannot lie inside B or B'. This implies that t lies completely in the wedge W. Let π_t, π, and π' denote the planes containing t, C, and C' respectively. Consider a unit sphere centered at p. This sphere intersects the line $\pi \cap \pi'$ at two points, say u and u'. Within W let the lines $\pi_t \cap \pi$ and $\pi_t \cap \pi'$ intersect the unit sphere at v and w respectively. See the

picture on the right in Figure 40. Without loss of generality, assume that the angle $\angle uvw \leq \angle uwv$. Consider the spherical triangle uvw. We are interested in the spherical angle $\theta = \angle uvw$ which is also the acute dihedral angle between the planes containing t and C. We have the following facts. The arc length of wv, denoted $|wv|$, is at least $\pi/3$ since p subtends the largest angle in t and t is in the wedge W. The spherical angle $\angle vuw$ is less than or equal to 2α. By standard sine laws in spherical geometry, we have

$$\sin \theta = \sin |uw| \frac{\sin \angle vuw}{\sin |wv|} \leq \sin |uw| \frac{\sin 2\alpha}{\sin |wv|}.$$

If $\pi/3 \leq |wv| \leq 2\pi/3$, we have

$$\sin |wv| \geq \sqrt{3}/2$$

and hence

$$\theta \leq \arcsin\left(\frac{2}{\sqrt{3}} \sin 2\alpha\right).$$

For the range $2\pi/3 < |wv| < \pi$, we use the fact that $|uw| + |wv| \leq \pi$. The arc length $|wv|$ cannot be longer than both $|wu'|$ and $|vu'|$ since $\angle vu'w \leq 2\alpha < \pi/2$ for $R_{pqr} \leq \frac{f(p)}{\sqrt{2}}$. If $|wv| \leq |wu'|$, we have

$$|uw| + |wv| \leq |uu'| = \pi.$$

Otherwise, $|wv| \leq |vu'|$. Then, we use the fact that $|uw| \leq |uv|$ as $\angle uvw \leq \angle uwv$. So, again

$$|uw| + |wv| \leq |uu'| = \pi.$$

Therefore, when $|wv| > \frac{2\pi}{3}$, we get

$$\frac{\sin |uw|}{\sin |wv|} < 1.$$

Thus, $\theta \leq \arcsin\left(\frac{2}{\sqrt{3}} \sin 2\alpha\right)$.

The normals to t and Σ at p make an acute angle at most $\alpha + \theta$ proving the lemma. $\qquad\square$

$\square$

25 Topology

The sample P as a set of discrete points does not have the topology of Σ. A connection between the topology of Σ and P can be established through the restricted Voronoi and Delaunay diagrams. In particular, one can show that the underlying space of the restricted Delaunay triangulation Del $P|_\Sigma$ is homeomorphic to Σ if the sample P is sufficiently dense. Although we will not be able to compute Del $P|_\Sigma$, the fact that it is homeomorphic to Σ will be useful in the surface reconstruction later.

25.1   Topological ball property

The underlying space of Del $P|_\Sigma$ becomes homeomorphic to Σ when the Voronoi diagram Vor P intersects Σ nicely. This condition is formalized by the topological ball property which says that the restricted Voronoi cells in each dimension is a ball.

Definition 52. Let F denote any Voronoi face of dimension k, $0 \leq k \leq 3$, in Vor P which intersects Σ and $F|_\Sigma = F \cap \Sigma$ be the corresponding restricted Voronoi face. The face F satisfies the topological ball property if $F|_\Sigma$ is a (i) $(k-1)$-ball and (ii) Int $F \cap \Sigma = $ Int $F|_\Sigma$. The pair (P, Σ) satisfies the topological ball property if all Voronoi faces $F \in$ Vor P satisfy the topological ball property.

Condition (i) means that Σ intersects a Voronoi cell in a single topological disk, a Voronoi facet in a single curve segment, a Voronoi edge in a single point, and does not intersect any Voronoi vertex; see Figure 41. Condition (ii) avoids any tangential intersection between a Voronoi face and Σ.

The following theorem is an important result relating the topology of a surface to a point sample.

Theorem 30. *The underlying space of* Del $P|_\Sigma$ *is homeomorphic to* Σ *if the pair* (P, Σ) *satisfies the topological ball property.*
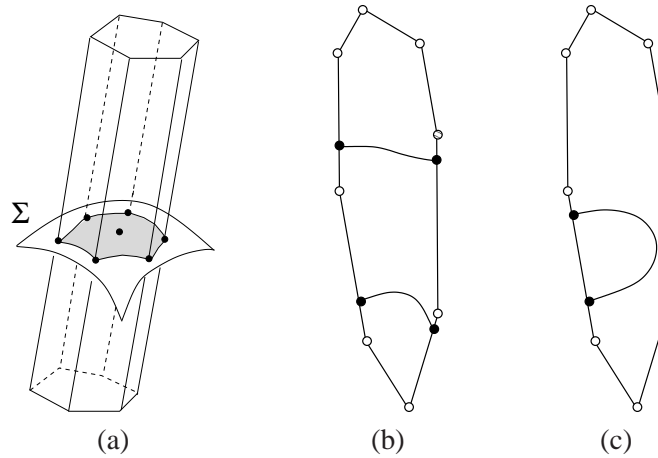

Figure 41: (a) A surface Σ intersects a Voronoi cell and its faces with the topological ball property, (b) a surface does not intersect a Voronoi facet in a 1-ball, (c) a surface does not intersect a Voronoi edge in a 0-ball.

Our aim is to show that, when P is a dense sample, the topology of Σ can be captured from P. Specifically, we prove that the pair (P, Σ) satisfies the topological ball property when ε is sufficiently small. The proof frequently uses the next two lemmas to reach a contradiction. The first one says that the points in a restricted Voronoi cell, that is, the points of Σ in a Voronoi cell, cannot be far apart. The second one says that any line almost normal to the surface cannot intersect it twice within a small distance.

Lemma 31 (Short Distance.). *Let x and y be any two points in a restricted Voronoi cell $V_p|_\Sigma$. For $\varepsilon < 1$, we have*

   (i) $\|x - p\| \le \frac{\varepsilon}{1-\varepsilon} f(p)$ and

   (ii) $\|x - y\| \le \frac{2\varepsilon}{1-\varepsilon} f(x)$.

P   . Since x has p as the nearest sample point, $\|x - p\| \le \varepsilon f(x)$ for $\varepsilon < 1$. Apply the Feature Translation Lemma 13 to claim (i). For (ii), observe that

$$\begin{aligned} \|x - y\| &\le \|x - p\| + \|y - p\| \\ &\le \varepsilon(f(x) + f(y)) \end{aligned}$$

By the Lipschitz Continuity Lemma 12

$$\begin{aligned} f(y) &\le f(x) + \|x - y\| \\ &\le f(x) + \varepsilon(f(x) + f(y)), \text{ or} \\ (1 - \varepsilon)f(y) &\le (1 + \varepsilon)f(x). \end{aligned}$$

Therefore, for $\varepsilon < 1$,

$$\|x - y\| \le \varepsilon\left(1 + \frac{1 + \varepsilon}{1 - \varepsilon}\right)f(x) \le \frac{2\varepsilon}{1 - \varepsilon}f(x).$$

□

□

A restricted Delaunay edge pq is dual to a Voronoi facet that intersects Σ. Any such intersection point, say x, is within $\frac{\varepsilon}{1-\varepsilon} f(p)$ distance from p by the Short Distance Lemma 31. The length of pq cannot be more than twice the distance between x and p. Hence $\|p - q\| \le \frac{2\varepsilon}{1-\varepsilon} f(p)$. We can extend this argument to the restricted Delaunay triangles too. A restricted Delaunay triangle t is dual to a Voronoi edge e that intersects Σ. The intersection point, say x, belongs to the Voronoi cells adjacent to e. Let V_p be any such cell. The point x is the center of a circumscribing ball of the triangle dual to e. By the Short Distance Lemma 31, x is within $\frac{\varepsilon}{1-\varepsilon} f(p)$ distance from p. The ball $B_{x,\|x-p\|}$ circumscribes t. The circumradius of t is no more than $\|x - p\|$ as the circumradius of a triangle cannot be more than any of its circumscribing ball (see Figure 42). Thus, the following corollary is immediate from the Short Distance Lemma 31.

Corollary 32. *For $\varepsilon < 1$, we have*

   (i) the length of a restricted Delaunay edge e is at most $\frac{2\varepsilon}{1-\varepsilon} f(p)$ where p is any vertex of e and

   (ii) the circumradius of any restricted Delaunay triangle t is at most $\frac{\varepsilon}{1-\varepsilon} f(p)$ where p is a vertex of t.

Lemma 33 (Long Distance.). *Suppose a line intersects Σ in two points x and y and makes an angle no more than ξ with $\mathbf{n}_x$. One has $\|x - y\| \ge 2f(x) \cos \xi$.*

P   . Consider the two medial balls at x as in Figure 43. The line meets the boundaries of these two balls at x and at points that must be at least $2r \cos \xi$ distance away from x where r is the radius of the smaller of the two balls. Since $r \ge f(x)$, the result follows as y cannot lie inside any of the two medial balls.
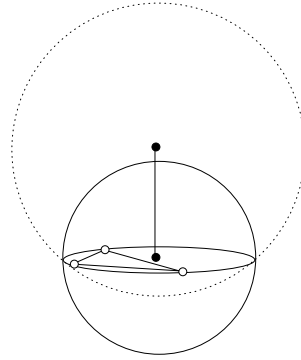
□

□

Figure 42: The circumradius of a triangle which is also the radius of its diamteric ball (shown with solid circle) is no more than the radius of a circumscribing ball (shown with dotted circle).
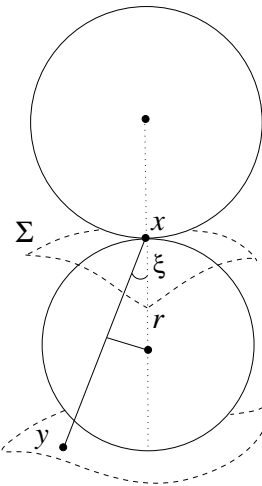

Figure 43: Illustration for the Long Distance Lemma 33.

25.2   Voronoi faces

Next we consider in turn the Voronoi edges, Voronoi facets, and Voronoi cells and show that they indeed satisfy the topological ball property if ε satisfies Condition A as stated below. For $\varepsilon < \frac{1}{3}$, let

$$\alpha(\varepsilon) = \frac{\varepsilon}{1 - 3\varepsilon}$$

$$\beta(\varepsilon) = \arcsin \frac{\varepsilon}{1 - \varepsilon} + \arcsin \left(\frac{2}{\sqrt{3}} \sin \left(2 \arcsin \frac{\varepsilon}{1 - \varepsilon} \right) \right).$$

Condition A. $\varepsilon < \frac{1}{3}$ and $\cos (\alpha(\varepsilon) + \beta(\varepsilon)) > \frac{\varepsilon}{1 - \varepsilon}.$

Figure 44 shows that, in the range $0 < \varepsilon \leq \frac{1}{3}$, Condition A holds for ε a little less than 0.2. So,
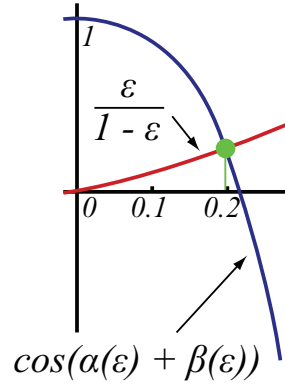
Figure 44: The graphs of the two functions on the left and right hand sides of the inequality in Condition A.

for example, $\varepsilon \le 0.18$ is a safe choice. Since Condition A stipulates $\varepsilon < \frac{1}{3}$ lemmas such as Normal Variation, Long Distance, Short Distance and Corollary 32 can be applied under Condition A.

Lemma 34 (Voronoi Edge.)**.** *A Voronoi edge intersects* Σ *transversally in a single point if Condition* A *holds.*

P   .   Suppose for the sake of contradiction there is a Voronoi edge e in a Voronoi cell V_p intersecting Σ at two points x and y, or at a single point tangentially, see Figure 45. The dual Delaunay triangle, say pqr, is a restricted Delaunay triangle. By Corollary 32, its circumradius is no more than $\frac{\varepsilon}{1-\varepsilon} f(p)$. By the Triangle Normal Lemma 29, $\angle_a(\mathbf{n}_{pqr}, \mathbf{n}_p) \le \beta(\varepsilon)$ if

$$\frac{1}{\sqrt{2}} > \frac{\varepsilon}{1-\varepsilon}$$

a restriction satisfied by Condition A.

The Normal Variation Lemma 27 puts an upper bound of $\alpha(\varepsilon)$ on the angle between the normals at p and x as $\|x - p\| \le \varepsilon f(x)$. Let ξ denote the angle between $\mathbf{n}_x$ and the Voronoi edge e. We have

$$\begin{aligned} \xi = \angle_a(\mathbf{n}_x, \mathbf{n}_{pqr}) &\le & \angle_a(\mathbf{n}_x, \mathbf{n}_p) + \angle_a(\mathbf{n}_p, \mathbf{n}_{pqr}) \\ &\le & \alpha(\varepsilon) + \beta(\varepsilon). \end{aligned} \tag{2}$$

If e intersects Σ tangentially at x, we have $\xi = \frac{\pi}{2}$ requiring $\alpha(\varepsilon) + \beta(\varepsilon) \ge \frac{\pi}{2}$. Condition A requires $\varepsilon < 0.2$ which gives $\alpha(\varepsilon) + \beta(\varepsilon) < \frac{\pi}{2}$. Therefore, when Condition A is satisfied, e cannot intersect Σ tangentially. So, assume that e intersects Σ at two points x and y.

By the Short Distance Lemma 31, $\|x - y\| \le \frac{2\varepsilon}{1-\varepsilon} f(x)$ and by the Long Distance Lemma 33, $\|x - y\| \ge 2f(x) \cos \xi$. A contradiction is reached when $2 \cos \xi > \frac{2\varepsilon}{1-\varepsilon}$, or

$$\cos(\alpha(\varepsilon) + \beta(\varepsilon)) > \frac{\varepsilon}{1 - \varepsilon}. \tag{3}$$

Condition A satisfies Inequality 3 giving the required contradiction. $\square$
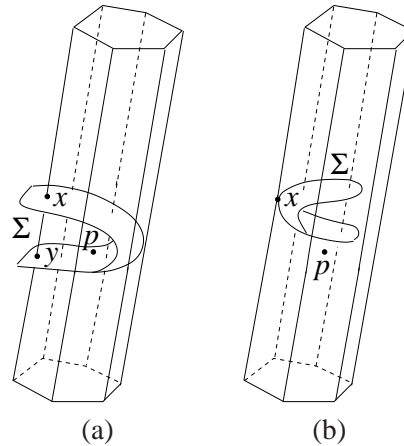
$\square$

Figure 45: Illustration for the Voronoi Edge Lemma 34. A Voronoi edge intersecting the surface (a) at two points, (b) tangentially in a single point.

Lemma 35 (Voronoi Facet.). *A Voronoi facet F intersects Σ transversally in a 1-ball if Condition A is satisfied.*

P   . The intersection of F with Σ may contradict the assertion of the lemma if (i) Σ touches F tangentially at a point, (ii) Σ intersects F in a 1-sphere, that is, a cycle, or (iii) Σ intersects F in more than one component.

The dual Delaunay edge, say pq, of F is in the restricted Delaunay triangulation. Let $\mathbf{n}_F$ denote the normal to F. Its direction is the same as that of pq up to orientation. We have $\|p-q\| \leq \frac{2\varepsilon}{1-\varepsilon} f(p)$ by Corollary 32. Therefore, the Edge Normal Lemma 28 gives

$$\angle_a(\mathbf{n}_p, \mathbf{n}_F) \geq \frac{\pi}{2} - \arcsin \frac{\varepsilon}{1-\varepsilon}$$

as long as $\varepsilon < 1$.

If Σ meets F tangentially at a point x, we have $\angle_a(\mathbf{n}_x, \mathbf{n}_F) = 0$ and by the Normal Variation Lemma 27 $\angle \mathbf{n}_p, \mathbf{n}_x \leq \frac{\varepsilon}{1-3\varepsilon}$ when $\varepsilon < \frac{1}{3}$. This means, for $\varepsilon < \frac{1}{3}$, we have

$$\frac{\pi}{2} - \arcsin \frac{\varepsilon}{1-\varepsilon} \leq \angle_a(\mathbf{n}_p, \mathbf{n}_F) \leq \frac{\varepsilon}{1-3\varepsilon} = \alpha(\varepsilon).$$

The above inequality contradicts the upper bound for ε given by Condition A.

If Σ meets F in a cycle, let x be any point on it and L be the line on F intersecting the cycle at x orthogonally, see Figure 46(a). This line must meet the cycle in another point, say y. The angle between L and $\mathbf{n}_x$ satisfies $\angle_a(L, \mathbf{n}_x) \leq \angle_a(L', \mathbf{n}_x)$ for any other line L' on F passing through x. Choose L' that minimizes the angle with $\mathbf{n}_p$. The line L' being on the Voronoi facet F makes exactly $\frac{\pi}{2}$ angle with the dual restricted Delaunay edge, say pq. We know by the Edge Normal Lemma 28

$$\angle_a(\overrightarrow{pq}, \mathbf{n}_p) \geq \frac{\pi}{2} - \arcsin \frac{\varepsilon}{1-\varepsilon}.$$

Therefore, for $\varepsilon < 1$,

$$\angle_a(L', \mathbf{n}_p) = \frac{\pi}{2} - \angle_a(\overrightarrow{pq}, \mathbf{n}_p) \leq \arcsin \frac{\varepsilon}{1-\varepsilon}.$$
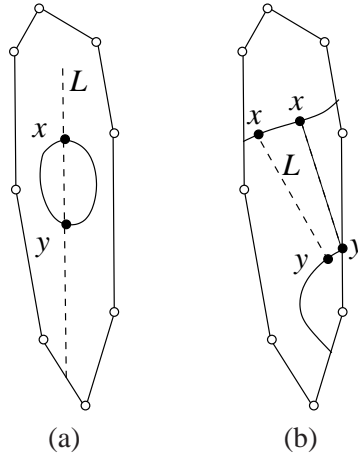
Figure 46: A Voronoi facet intersecting Σ (a) in a cycle, (b) in two segments.

These facts with the Normal Variation Lemma 27 give

$$\angle_a(L', \mathbf{n}_x) \leq \angle_a(L', \mathbf{n}_p) + \angle(\mathbf{n}_p, \mathbf{n}_x) \leq \arcsin \frac{\varepsilon}{1 - \varepsilon} + \alpha(\varepsilon) \tag{4}$$

for $\varepsilon < \frac{1}{3}$.

The right hand side of Inequality 4 is less than the upper bound for ξ in the proof of the Voronoi Edge Lemma 34. Thus, we reach a contradiction between distances implied by the Short Distance Lemma 31 and the Long Distance Lemma 33 when Condition A holds.

In the case Σ meets F in two or more components as in Figure 46(b), consider any point x in one of the components. Let y be the closest point to x on any other component, say C. If the line L joining x and y meets C orthogonally at y we have the situation as in the previous case with only x and y interchanged. In the other case, y lies on the boundary of C on a Voronoi edge. The angle between L and $\mathbf{n}_y$ is less than the angle between the Voronoi edge and $\mathbf{n}_y$ which is no more than $\alpha(\varepsilon) + \beta(\varepsilon)$ as proved in the Voronoi Edge Lemma 34 (Inequality 2). We reach a contradiction again between two distances using the same argument. □

□

Lemma 36 (Voronoi Cell.). *A Voronoi cell V_p intersects Σ in a 2-ball if Condition A holds.*

P   .   We have $W = V_p \cap \Sigma$ contained in a ball B of radius $\frac{\varepsilon}{1-\varepsilon} f(p)$ by the Short Distance Lemma 31. If W is a manifold without boundary, B contains a medial axis point m by the Feature Ball Lemma 11. Then the radius of B is at least

$$\frac{\|m - p\|}{2} \geq \frac{f(p)}{2}.$$

We reach a contradiction if $\varepsilon < \frac{1}{3}$ which is satisfied by Condition A. So, assume that W is a manifold with boundary. It may not be a 2-ball only if it is non-orientable, has a handle, or has more than one boundary cycle. If W were non-orientable, so would be Σ, which is impossible. In

case W has a handle, $B \cap \Sigma$ is not a 2-ball. By the Feature Ball Lemma 11, it contains a medial axis point reaching a contradiction again for $\varepsilon < \frac{1}{3}$ which is satisfied by Condition A.

The only possibility left is that W has more than one boundary cycles. Let L be the line of the normal at p. Consider a plane that contains L and intersects at least two boundary cycles. Such a plane exists since otherwise L must intersect W at a point other than p and we reach a contradiction between two distance lemmas. The plane intersects V_p in a convex polygon and W in at least two curves. We can argue as in the proof of the Voronoi Facet Lemma 35 to reach a contradiction between two distance lemmas.   □
                                                                                                                                                 □

Condition A holds for $\varepsilon \leq 0.18$. Therefore, the Voronoi Edge Lemma, Facet Lemma, and Cell Lemma hold for $\varepsilon \leq 0.18$. Then, Theorem 30 leads to the following result.

Theorem 37 (Topological Ball.). *Let P be an ε-sample of a smooth surface Σ. For $\varepsilon \leq 0.18$, (P, Σ) satisfies the topological ball property and hence the underlying space of $\mathrm{Del}\, P|_\Sigma$ is homeomorphic to Σ.*

26 Notes and exercises

The remarkable connection between ε-samples of a smooth surface and the Voronoi diagram of the sample points was first discovered by Amenta and Bern [1]. The Normal Lemma 26 and the Normal Variation Lemma 27 are two key observations made in this paper. The topological ball property that ensures the homeomorphism between the restricted Delaunay triangulation and the surface was discovered by Edelsbrunner and Shah [4]. Amenta and Bern showed that the Voronoi diagram of a sufficiently dense sample satisfies the topological ball property though the proof was not as rigorous as presented here. The proof presented here is adapted from Cheng, Dey, Edelsbrunner, and Sullivan [2].

Exercises

1. Let the restricted Voronoi cell $V_p|_\Sigma$ be adjacent to the restricted Voronoi cell $V_q|_\Sigma$ in the restricted Voronoi diagram $\mathrm{Vor}|_\Sigma P$. Show that the distance between any two points x and y from the union of $V_p|_\Sigma$ and $V_q|_\Sigma$ is $\tilde{O}(\varepsilon)f(x)$ when ε is sufficiently small.

2. A version of the Edge Normal Lemma 28 can be derived from the Triangle Normal Lemma 29, albeit with a slightly worse angle bound. Derive this angle bound and carry out the proof of the topological ball property with this bound. Find out an upper bound on ε for the proof.

3. The Topological Ball Property is a sufficient but not a necessary condition for the homeomorphism between a sampled surface and a restricted Delaunay triangulation of it. Establish this fact by an example.

4. Show an example where

    (i) all Voronoi edges satisfy the topological ball property, but the Voronoi cell does not,

(ii) all Voronoi facets satisfy the topological ball property, but the Voronoi cell does not.

5. Show that for any $n > 0$, there exists a C^2-smooth surface for which a sample with n points has the Voronoi diagram where no Voronoi edge intersects the surface.

6^h. Let F be a Voronoi facet in the Voronoi diagram Vor P where P is an ε-sample of a C^2-smooth surface Σ. Let Σ intersect F in a single interval and the intersection points with the Voronoi edges lie within $\varepsilon f(p)$ away from p where $F \subset V_p$. Show that all points of $F \cap \Sigma$ lie within $\varepsilon f(p)$ distance when ε is sufficiently small.

7. Let F and Σ be as described in exercise 6 but $F \cap \Sigma$ contains two or more topological intervals. Show that there exists a Voronoi edge $e \in F$ so that $e \cap \Sigma$ is at least $\lambda f(p)$ away from p where $\lambda > 0$ is an appropriate constant.

8^o. Let the pair (P, Σ) satisfy the topological ball property. We know that the underlying space of Del $P|_\Sigma$ and Σ are homeomorphic. Prove or disprove that they are isotopic.

References

[1] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discr. Comput. Geom.*, 22:481–504, 1999.

[2] H.-L. Cheng, T. K. Dey, H. Edelsbrunner and J. Sullivan. Dynamic skin triangulation. *Discrete Comput. Geom.*, 25:525–568, 2001.

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

[4] H. Edelsbrunner and N. Shah. Triangulating topological spaces. *Internat. J. Comput. Geom. Appl.*, 7:365–378, 1997.

Surface Reconstruction taken from [3]

27   Surface Reconstruction

In the previous chapter we learned that the restricted Delaunay triangulation is a good approximation of a densely sampled surface Σ from both topological and geometric view point. Unfortunately, we cannot compute this triangulation as the restricted Voronoi diagram $\text{Vor } P|_\Sigma$ cannot be computed without knowing Σ. As a remedy we approximate the restricted Voronoi diagram and compute a set of triangles that is a superset of all restricted Delaunay triangles. This set is pruned to extract a manifold surface which is output as an approximation to the sampled surface Σ.

28   Algorithm

First we observe that each restricted Voronoi cell $V_p|_\Sigma$ is almost flat if the sample is sufficiently dense. This follows from the Normal Variation Lemma 27 as the points in $V_p|_\Sigma$ cannot be far apart if ε is small. In particular, $V_p|_\Sigma$ lies within a thin neighborhood of the tangent plane τ_p at p. So we need two approximations, (i) an approximation to τ_p or equivalently to $\mathbf{n}_p$ and (ii) an approximation to $V_p|_\Sigma$ based on the approximation to $\mathbf{n}_p$. The following definitions of *poles* and *cocones* are used for these two approximations.

28.1   Poles and Cocones

Definition 53 (Poles.). The farthest Voronoi vertex, denoted p^+, in V_p is called the *positive pole* of p. The *negative pole* of p is the farthest point $p^- \in V_p$ from p so that the two vectors from p to p^+ and p^- make an angle more than $\frac{\pi}{2}$. We call $\mathbf{v}_p = p^+ - p$, the *pole vector* for p. If V_p is unbounded, p^+ is taken at infinity and the direction of $\mathbf{v}_p$ is taken as the average of all directions given by the unbounded Voronoi edges.

The following lemma is a direct consequence of the Normal Lemma 26. It says that the pole vectors approximate the true normals at the sample points.

Lemma 38 (Pole.). *For $\varepsilon < 1$, the angle between the normal $\mathbf{n}_p$ at p and the pole vector $\mathbf{v}_p$ satisfies the inequality*

$$\angle_a(\mathbf{n}_p, \mathbf{v}_p) \leq 2 \arcsin \frac{\varepsilon}{1 - \varepsilon}.$$

P   .   First, consider the case where V_p is bounded. Since the Voronoi cell V_p contains the centers of the medial balls at p, we have $\|p^+ - p\| \geq f(p)$. Thus, plugging $\mu = 1$ in the Normal Lemma 26 we obtain the result immediately.

Next, consider the case where V_p is unbounded. In this case $\mathbf{v}_p$ is computed as the average of the directions of the infinite Voronoi edges. The angle $\angle_a(\mathbf{v}_p, \mathbf{n}_p)$ in this case cannot be more than the worst angle made by an infinite Voronoi edge with $\mathbf{n}_p$. An infinite Voronoi edge e makes the same angle with $\mathbf{n}_p$ as the vector $\overrightarrow{pp_\infty}$ does, where the infinite endpoint of e is taken at p_∞.

Again we have $\|p - p_\infty\| \geq f(p)$ and the Normal Lemma 26 can be applied with $\mu = 1$ to give the result. □

□

The Pole Lemma 38 says that the pole vector approximates the normal $\mathbf{n}_p$. Thus, the plane $\tilde{\tau}_p$ passing through p with the pole vector as normal approximates the tangent plane τ_p. The following definition of *cocone* accommodates a thin neighborhood around $\tilde{\tau}_p$ to account for the small uncertainty in the estimation of $\mathbf{n}_p$.

Definition 54 (Cocone.)**.** The set $C_p = \{y \in V_p \; : \; \angle_a(\overrightarrow{py}, \mathbf{v_p}) \geq \frac{3\pi}{8}\}$ is called the cocone of p. In words, C_p is the complement of a double cone that is clipped within V_p. This double cone has p as the apex, the pole vector $\mathbf{v}_p$ as the axis, and an opening angle of $\frac{3\pi}{8}$ with the axis. See Figure 47 for an example of a cocone.
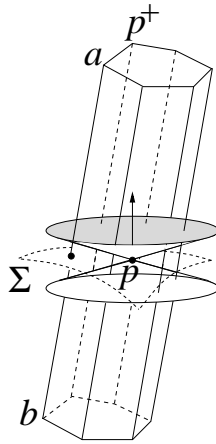

Figure 47: The positive pole p^+ helps estimating the normal. The double cone forming the cocone has the apex at p and axis pp^+. The Voronoi edge ab intersects the cocone. Its dual Delaunay triangle is a cocone triangle.

As an approximation to $V_p|_\Sigma$, cocones meet all Voronoi edges that are intersected by Σ. So, if we compute all triangles dual to the Voronoi edges intersected by cocones, we obtain all restricted Delaunay triangles and possibly a few others. We call this set of triangles *cocone triangles*. We will see later that all cocone triangles lie very close to Σ. A cleaning step is necessary to weed out some triangles from the set of cocone triangles so that a 2-manifold is computed as output. This is accomplished by a *manifold extraction* step.

```
C      (P)
 1   compute Vor P;
 2   T = ∅;
 3   for each Voronoi edge e ∈ Vor P do
 4     if C     T       (e)
 5         T := T ∪ dual e;
 6   endfor
```

7   $E :=$E   M (T);
8   output E.

Let us now look into the details of the two steps C   T   and E   M   .
In order to check if a Voronoi edge $e = (a, b)$ intersects C_p we consider the three vectors $\mathbf{v}_p$, $\mathbf{a} = \overrightarrow{pa}$, $\mathbf{b} = \overrightarrow{pb}$, and three conditions I, II, and III:

$$I. \quad \left| \frac{\mathbf{v}_p^T \mathbf{a}}{\|\mathbf{v}_p\| \|\mathbf{a}\|} \right| \leq \cos \frac{3\pi}{8} \text{ or } \left| \frac{\mathbf{v}_p^T \mathbf{b}}{\|\mathbf{v}_p\| \|\mathbf{b}\|} \right| \leq \cos \frac{3\pi}{8},$$

$$II. \quad \frac{\mathbf{v}_p^T \mathbf{a}}{\|\mathbf{v}_p\| \|\mathbf{a}\|} < 0 \text{ and } \frac{-\mathbf{v}_p^T \mathbf{b}}{\|\mathbf{v}_p\| \|\mathbf{b}\|} < 0,$$

$$III. \quad \frac{\mathbf{v}_p^T \mathbf{a}}{\|\mathbf{v}_p\| \|\mathbf{a}\|} > 0 \text{ and } \frac{-\mathbf{v}_p^T \mathbf{b}}{\|\mathbf{v}_p\| \|\mathbf{b}\|} > 0.$$

Condition I checks if any of the vertices a and b of the Voronoi edge e lies inside C_p. Conditions II and III check if both a and b lie outside C_p, but the edge e crosses it. The triangle $t = $ dual e is marked as a cocone triangle only if e intersects cocones of *all* three vertices of t.

C   T (e)
1   $t :=$ dual e;
2   flag :=   ;
3   for each vertex p of t do
4   if none of Conditions I, II, and III holds
5   flag:=   ;
6   endfor
7   return flag.

The set T of cocone triangles enjoys some interesting geometric properties which we exploit in the manifold extraction step as well as in the proofs of geometric and topological guarantees of C   . Of course, the sample has to be sufficiently dense for these properties to hold. In the rest of the chapter we assume that $\varepsilon \leq 0.05$ which satisfies Condition A stated in Chapter 23, enabling us to apply the results therein.

28.2   Cocone triangles

First we show that each triangle in T has a small empty ball circumscribing it, i.e., the radius of this ball is small compared to the local feature sizes at their vertices. Notice that the diametric ball of a triangle may not be empty. Hence, the smallest *empty* ball circumscribing a triangle may not be its diametric ball. Nevertheless, a small empty circumscribing ball also means that the circumradius of the triangle is small. This fact together with the Triangle Normal Lemma 29 implies that all cocone triangles lie almost flat to the surface.

Lemma 39 (Small Triangle.). *Let t be any cocone triangle and r denote the radius of the smallest empty ball circumscribing t. For each vertex p of t and $\varepsilon \leq 0.05$, one has*

(i) $r \leq \frac{1.18\varepsilon}{1-\varepsilon} f(p)$ *and*

(ii) *circumradius of t is at most* $\frac{1.18\varepsilon}{1-\varepsilon} f(p)$.

P   . Let z be any point in V_p so that

$$\angle_a(\mathbf{n}_p, \overrightarrow{pz}) \geq \frac{3\pi}{8} - 2\arcsin\frac{\varepsilon}{1-\varepsilon}. \tag{5}$$

First we claim that for any such point z, we have $\|z - p\| \leq \frac{1.18\,\varepsilon}{1-\varepsilon} f(p)$ if $\varepsilon \leq 0.05$.

If $\angle_a(\mathbf{n}_p, \overrightarrow{pz}) > \theta = \arcsin\frac{\varepsilon}{\mu(1-\varepsilon)} + \arcsin\frac{\varepsilon}{1-\varepsilon}$, then $\|z - p\| \leq \mu f(p)$ according to the Normal Lemma 26. With $\mu = \frac{1.18\varepsilon}{1-\varepsilon}$ and $\varepsilon \leq 0.05$ we have

$$\theta \;=\; \arcsin\frac{1}{1.18} + \arcsin\frac{\varepsilon}{1-\varepsilon} < \frac{3\pi}{8} - 2\arcsin\frac{\varepsilon}{1-\varepsilon}. \tag{6}$$

Thus, from Inequalities 5 and 6 we have

$$\angle_a(\mathbf{n}_p, \overrightarrow{pz}) \;\geq\; \frac{3\pi}{8} - 2\arcsin\frac{\varepsilon}{1-\varepsilon} > \theta. \tag{7}$$

Therefore, any point $z \in V_p$ satisfying Inequality 5 also satisfies

$$\|z - p\| \leq \frac{1.18\varepsilon}{1-\varepsilon} f(p).$$

Now let t be any cocone triangle with p being any of its vertices and $e = \text{dual } t$ being its dual Voronoi edge. For t to be a cocone triangle, it is necessary that there is a point $y \in e$ so that $\angle_a(\mathbf{v}_p, \overrightarrow{py}) \geq \frac{3\pi}{8}$. Taking into account the angle $\angle_a(\mathbf{v}_p, \mathbf{n}_p)$, this necessary condition implies

$$\angle_a(\mathbf{n}_p, \overrightarrow{py}) \geq \frac{3\pi}{8} - 2\arcsin\frac{\varepsilon}{1-\varepsilon}$$

which satisfies Inequality 5. Hence, we have

$$\|y - p\| \leq \frac{1.18\,\varepsilon}{1-\varepsilon} f(p) \text{ for } \varepsilon \leq 0.05.$$

The ball $B_{y,\|y-p\|}$ is empty and circumscribes t proving (i). The claim in (ii) follows immediately from (i) as the circumradius of t cannot be larger than the radius of any ball circumscribing it. □

□

The next lemma proves that all cocone triangles lie almost parallel to the surface. The angle bounds are expressed in terms of $\alpha(\varepsilon)$ and $\beta(\varepsilon)$ that are defined in Chapter 23.

Lemma 40 (Cocone Triangle Normal.)**.** *Let t be any cocone triangle and $\mathbf{n}_t$ be its normal. For any vertex p of t one has $\angle_a(\mathbf{n}_p, \mathbf{n}_t) \leq \alpha(\frac{2.36\varepsilon}{1-\varepsilon}) + \beta(1.18\varepsilon)$ when $\varepsilon \leq 0.05$.*

P   . Let q be a vertex of t with a maximal angle of t. The circumradius of t is at most $\frac{1.18\varepsilon}{1-\varepsilon} f(q)$ by the Small Triangle Lemma 39. Then, by the Triangle Normal Lemma 29,

$$
\begin{aligned}
\angle_a(\mathbf{n}_q, \mathbf{n}_t) \quad &\leq \quad \arcsin \frac{1.18\varepsilon}{1-\varepsilon} + \arcsin\left(\frac{2}{\sqrt{3}} \sin\left(2\arcsin \frac{1.18\varepsilon}{1-\varepsilon} \right) \right) \\
&\leq \quad \arcsin \frac{1.18\varepsilon}{1-1.18\varepsilon} + \arcsin\left(\frac{2}{\sqrt{3}} \sin\left(2\arcsin \frac{1.18\varepsilon}{1-1.18\varepsilon} \right) \right) \\
&= \quad \beta(1.18\varepsilon) \ \text{for } \varepsilon \leq 0.05.
\end{aligned}
$$

The distance between p and q is no more than the diameter of the circle circumscribing t, i.e., $\|p-q\| \leq \frac{2.36\varepsilon}{1-\varepsilon} f(p)$ (Small Triangle Lemma 39). By the Normal Variation Lemma 27, $\angle(\mathbf{n}_p, \mathbf{n}_q) \leq \alpha(\frac{2.36\varepsilon}{1-\varepsilon})$. The desired bound for $\angle_a(\mathbf{n}_p, \mathbf{n}_t)$ follows since it is no more than the sum $\angle(\mathbf{n}_p, \mathbf{n}_q) +$ $\angle_a(\mathbf{n}_q, \mathbf{n}_t)$. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

$\hfill\square$

28.3   Pruning

Prior to the extraction of a 2-manifold from the set of cocone triangles, some of them are pruned. An edge e is *sharp* if any two consecutive cocone triangles around it form an angle more than $\frac{3\pi}{2}$; see Figure 48. Edges with a single triangle incident to them are also sharp by default. We will show later that the cocone triangles include all restricted Delaunay triangles when a sample is sufficiently dense. The set of restricted Delaunay triangles cannot be incident to sharp edges. This implies that we can prune triangles incident to sharp edges and still retain the set of restricted Delaunay triangles. In fact, we can carry out this pruning in a cascaded manner. By deleting one triangle incident to a sharp edge, we may create other sharp edges. Since no restricted Delaunay triangle is pruned, none of their edges become sharp. Therefore, it is safe to delete the new sharp edges with all of their incident triangles.
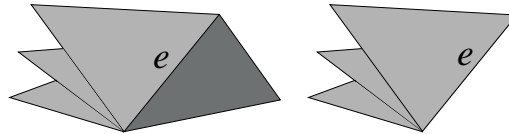

Figure 48: The edge e is not sharp in the left picture; it is sharp in the right picture.

This pruning step weeds out all triangles incident to sharp edges, but the remaining triangles still may not form a surface. They may form layers of thin pockets creating a non-manifold. A manifold surface is extracted from this possibly layered set by *walking* outside the space covered by them, see Figure 49. The manifold extraction step depends on the fact that cocone triangles contain all restricted Delaunay triangles none of whose edges is sharp. We prove this fact below.

Theorem 41 (Restricted Delaunay.). *For $\varepsilon \leq 0.05$, the following conditions hold:*

  (i) cocone triangles contain all restricted Delaunay triangles and

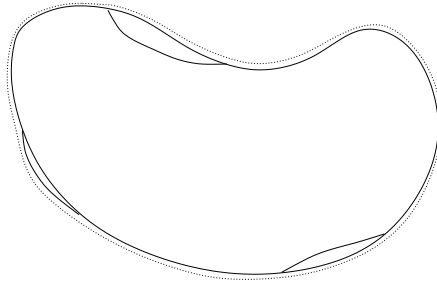 *(ii) no restricted Delaunay triangle has a sharp edge.*

Figure 49: Thin pockets left after pruning, a manifold is obtained by walking on the outside indicated by the dotted curve.

P . Consider (i). Let y be any point in any restricted Voronoi cell $V_p|_\Sigma$. We claim that $\angle_a(\mathbf{n}_p, \overrightarrow{py})$ is larger than $\frac{\pi}{2} - \arcsin \frac{\varepsilon}{2(1-\varepsilon)}$. We have $\|y - p\| \le \varepsilon f(y)$ since $y \in V_p|_\Sigma$ and P is an ε-sample of Σ. By the Feature Translation Lemma 13, $\|y - p\| \le \frac{\varepsilon}{1-\varepsilon} f(p)$. We can therefore apply the proof of the Edge Normal Lemma 28 to establish that

$$\angle_a(\mathbf{n}_p, \overrightarrow{py}) \ge \frac{\pi}{2} - \arcsin \frac{\varepsilon}{2(1-\varepsilon)}.$$

Let t be any restricted Delaunay triangle and $e = \text{dual } t$ be the dual Voronoi edge. Consider the point $y = e \cap \Sigma$. We have $y \in V_p|_\Sigma$ for each of the three points $p \in P$ determining e. For each such p, the angle $\angle_a(\mathbf{n}_p, \overrightarrow{py})$ is larger than $\pi/2 - \arcsin \frac{\varepsilon}{2(1-\varepsilon)}$. Therefore,

$$
\begin{aligned}
\angle_a(\overrightarrow{py}, \mathbf{v}_p) &\ge \angle_a(\overrightarrow{py}, \mathbf{n}_p) - \angle_a(\mathbf{n}_p, \mathbf{v}_p) \\
&\ge \frac{\pi}{2} - \arcsin \frac{\varepsilon}{2(1-\varepsilon)} - \angle_a(\mathbf{n}_p, \mathbf{v}_p).
\end{aligned}
\tag{8}
$$

By the Pole Lemma 38 we have

$$
\begin{aligned}
\angle_a(\mathbf{n}_p, \mathbf{v}_p) + \arcsin \frac{\varepsilon}{2(1-\varepsilon)} &\le 2 \arcsin \frac{\varepsilon}{1-\varepsilon} + \arcsin \frac{\varepsilon}{2(1-\varepsilon)} \\
&< \frac{\pi}{8} \text{ for } \varepsilon \le 0.05.
\end{aligned}
$$

So, by Inequality 8, $\angle_a(\overrightarrow{py}, \mathbf{v}_p) > \frac{3\pi}{8}$. Therefore, the point y is in the cocone C_p by definition. Hence t is a cocone triangle.

Consider (ii). Let t_1 and t_2 be adjacent triangles in the restricted Delaunay triangulation with e as their shared edge and let $p \in e$ be any of their shared vertices. Since t_1 and t_2 belong to the restricted Delaunay triangulation, they have circumscribing empty balls B_1 and B_2, respectively, centered at points, say v_1 and v_2 of Σ.

The boundaries of B_1 and B_2 intersect in a circle C contained in a plane H, with $e \subset H$. The plane H separates t_1 and t_2, since the third vertex of each triangle lies on the boundary of its circumscribing ball, and $B_1 \subseteq B_2$ on one side of H, while $B_2 \subseteq B_1$ on the other. See Figure 50. The line through v_1, v_2 is perpendicular to H. Both v_1 and v_2 belong to the Voronoi facet dual to e. This means v_1 and v_2 belong to a restricted Voronoi cell and the distance $\|v_1 - v_2\| \le \frac{2\varepsilon}{(1-\varepsilon)} f(v_1)$ by the Short Distance Lemma 31. So the segment $v_1 v_2$ forms an angle of at least $\pi/2 - \arcsin \frac{\varepsilon}{1-\varepsilon}$ with

$\mathbf{n}_{v_1}$ (Edge Normal Lemma 28). This normal differs, in turn, from $\mathbf{n}_p$ by an angle of at most $\frac{\varepsilon}{1-3\varepsilon}$ (Normal Variation Lemma 27). So, the angle between H and $\mathbf{n}_p$ is at most $\frac{\varepsilon}{1-3\varepsilon} + \arcsin \frac{\varepsilon}{1-\varepsilon}$. For small ε, they are nearly parallel. In particular, if $\varepsilon \leq 0.05$, H makes at most $7°$ with $\mathbf{n}_p$. Similarly, plugging $\varepsilon \leq 0.05$ in the angle upper bound of the Cocone Triangle Normal Lemma 40, one gets that the normals of both t_1 and t_2 differ from the surface normal at p by at most $24°$.
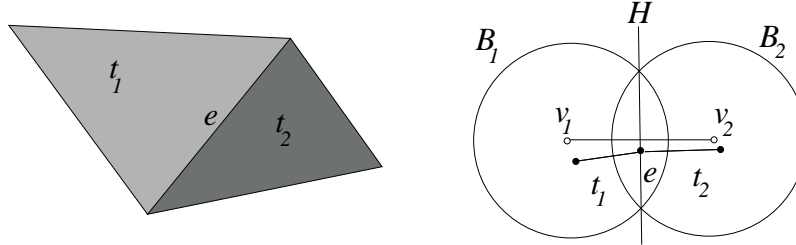

Figure 50: Illustration for the Restricted Delaunay Theorem 41.

Thus we have t_1 on one side of H, t_2 on the other and the smaller angle between H and either triangle is at least $59°$. Hence the smaller angle between t_1 and t_2 is at least $118°$ and e is not sharp. □

□

28.4 Manifold extraction

A simplicial complex with an underlying space of a 2-manifold is extracted out of the pruned set of cocone triangles. Let $\Sigma' \subseteq \Sigma$ be any connected component of the sampled surface. Since cocone triangles are small (Small Triangle Lemma 39), they cannot join points from different components of Σ. Let T' be the pruned set of cocone triangles with vertices in Σ'. Consider the medial axis of Σ'. The triangles of T' lie much closer to Σ' than to its medial axis. Furthermore, T' includes the restricted Delaunay triangulation $\text{Del }P|_{\Sigma'}$ (Restricted Delaunay Theorem 41). Therefore, if $|T'|$ denotes the underlying space of T', the space $\mathbb{R}^3 \setminus |T'|$ has precisely two disjoint open sets O_{in} and O_{out} containing the inner and outer medial axis of Σ' respectively. The manifold extraction step computes the boundary of the closure of O_{out}, which we simply refer to as the boundary of O_{out}.

Let E' be the boundary of O_{out}. We claim that E' is a 2-manifold. Let p be any vertex of E'. Orient the normal $\mathbf{n}_p$ so that it points toward O_{out}. Consider a sufficiently small ball B centering p. Call the point where the ray of $\mathbf{n}_p$ intersects the boundary of B the *north pole*. Obviously the north pole is in O_{out}. Let T_p denote the set of triangles in T' which are visible from the north pole within B. The triangles of T_p are in the boundary of O_{out}. Since there is no sharp edge in T', the set of triangles T_p makes a topological disk. We argue that T_p is the only set of triangles in the boundary of O_{out} which are incident to p.

Let $q \neq p$ be a vertex of a triangle $t \in T_p$. The triangle t is also in T_q. If not, the line of the normal $\mathbf{n}_p$, when moved parallelly through the edge pq toward q, must hit an edge in T' that is sharp. The assumption to this claim is that the normals $\mathbf{n}_p$ and $\mathbf{n}_q$ are almost parallel and hence the visibility directions at p and q are almost parallel. Since T' does not have any sharp edge, t is in T_q. This means that all topological disks at the vertices of E' are compatible and they form

a 2-manifold. This 2-manifold separates O_{out} from T' implying that E' cannot have any other triangles from T' other than the ones in the topological disks described above.

We compute E' from T' as a collection of triangles by a depth first walk in the Delaunay triangulation Del P. Recall that T' is disjoint from any other triangles on a component of Σ different from Σ'. The walk starts with a seed triangle in T'. The routine S   computes this seed triangle for each component T' of the pruned set by another depth first walk in the Delaunay triangulation. At any generic step, S   comes to a triangle t via a tetrahedron σ and performs the following steps. First, it checks if t is a cocone triangle. If so, it checks if it belongs to a component T' for which a seed has not yet been picked. If so, the pair (σ, t), also called the *seed pair*, is put into the seed set. Then, it marks all triangles of T' so that any subsequent check can identify that a seed for T' has been picked. The walk continues through the triangles and their adjacent tetrahedra in a depth first manner till a seed pair for each component such as T' of T is found. In a seed pair (σ, t) for a component T', the tetrahedron σ and the triangle t should be in O_{out} and on its boundary E' respectively. To ensure it S   starts the walk from any convex hull triangle in Del P and continue till it hits a cocone triangle. The initiation of the walk from a convex hull triangle ensures that the first triangle encountered in a component is on the outside of that component or equivalently on the boundary of O_{out} defined for that component. Assuming the function S   , a high level description of E   M   is given below.

```
E     M     (T)
 1   T := pruned T;
 2   SD := S    (T);
 3   for each tuple (σ,t) ∈ SD do
 4       E' := S    T         (σ,t);
 5       E := E ∪ E';
 6   endfor
 7   return the simplicial complex of E.
```

The main task in E   M   is done by S   T   which takes a seed pair (σ, t) as input. First, we initialize the surface E' with the seed triangle t which is definitely in E' (line 1). Next, we initialize a stack *Pending* with the triple (σ, t, e) where e is an edge of t (lines 3 and 4). As long as the stack *Pending* is not empty, we pop its top element (σ, t, e). If the edge e is not already processed we call the function S   N   to compute a tetrahedron-triangle pair (σ', t') (line 9). The tetrahedron σ' is adjacent to t' and intersects O_{out} where t' is in E' and is adjacent to t via e. The triangle t' is inserted in E'. Then two new triples (σ', t', e') are pushed on the stack *pending* for each edge $e' \neq e$ of t' (lines 11 to 13). Finally we return E' (line 16).

```
S     T          (σ,t)
 1   E' := {t};
 2   Pending := ∅;
 3   pick any edge e of t;
 4   push (σ, t, e) on Pending;
 5   while Pending ≠ ∅ do
 6       pop (σ, t, e) from Pending;
```

```
7      if e is not marked processed
8         mark e processed;
9         (σ′,t′) := S      N         (σ,t,e);
10        E′ := E′ ∪ {t′};
11        for each edge e′ ≠ e of t′ do
12           push (σ′,t′,e) on Pending;
13        endfor
14     endif
15  endwhile
16  return E′.
```

The question is how to implement the function S   N   . It has to output a tuple $(\sigma′,t′)$ where $t′$ is the neighbor of t on the surface given by $E′$ and $\sigma′$ is an adjacent tetrahedron intersecting O_{out}. One can compute the surface neighbor $t′$ of t using some numerical computations involving some dot product computations of vectors. However, these computations often run into trouble due to numerical errors with finite precision arithmetics. In particular, triangles of certain types of flat tetrahedra called *slivers* tend to contribute to these numerical errors and slivers are not uncommon in the Delaunay triangulation of a sample from a surface.

A robust and faster implementation of the function S   N   avoids numerical computations by exploiting the combinatorial structure of the Delaunay triangulation. Every triangle in the Delaunay triangulation has two incident tetrahedra if we account for the infinite ones incident to the convex hull triangles. S   N   is called with a triple (σ,t,e). It circles over the tetrahedra and triangles incident to the edge e starting from t and going towards the other triangle of σ incident to e. This circular walk stops when another cocone triangle $t′$ is reached. If $t′$ is reached via the tetrahedron $\sigma′$, we output the pair $(\sigma′,t′)$. Assuming inductively that σ intersects O_{out}, the tetrahedron $\sigma′$ also intersects O_{out}. For example, in Figure 51, S   N   - is passed on the triple (σ_1,t,e) and then it circles through the tetrahedra $\sigma_1,\sigma_2,\sigma_3$ and their triangles till it reaches $t′$. At this point it returns $(\sigma_3,t′)$ where both σ_1 and σ_3 lie outside, i.e., in O_{out}. S   T   with this implementation of S   N   is robust since no numerical decisions are involved, see Figure 51. Combinatorial computations instead of numerical ones make S   T   fast provided the Delaunay triangulation is given in a form which allows to answer queries for neighboring tetrahedra quickly.
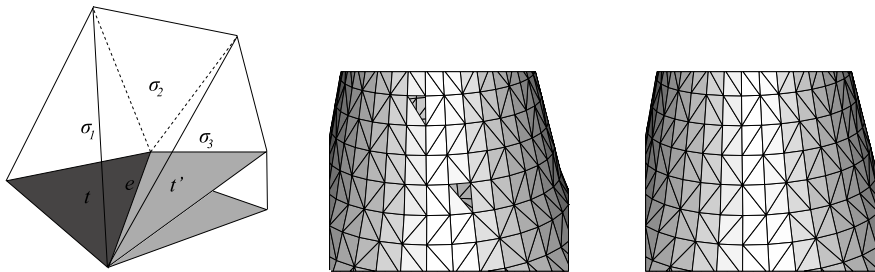

Figure 51: A stable computation of S   N   (left), a zoom on a reconstruction after an unstable computation with numerical errors (middle) and a stable computation without any numerical error (right).

29 Geometric guarantees

In this section we establish more properties of the cocone triangles which are eventually used to prove the geometric and topological guarantees of the output of C   . We introduce a map v that takes each point $x \in \mathbb{R}^3$ to its closest point in Σ. Notice that v is well defined everywhere in $\mathbb{R}^3$ except at the medial axis M of Σ. Mathematically, $v : \mathbb{R}^3 \setminus M \to \Sigma$ where $v(x) \in \Sigma$ is closest to x. Observe that the line containing x and $v(x)$ is normal to Σ at x. The map v will be used at many places in this chapter and the chapters to follow. Let

$$\tilde{x} = v(x) \text{ for any point } x \in \mathbb{R}^3 \setminus M \text{ and}$$
$$\tilde{U} = \{\tilde{x} : x \in U\} \text{ for any set } U \subset \mathbb{R}^3 \setminus M.$$

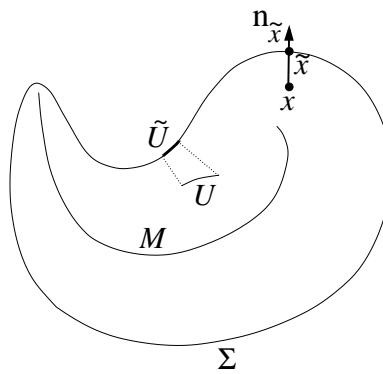See Figure 52 for an illustration.

Figure 52: Illustration for the map v.

First, we show that all points of the cocone triangles lie close to the surface. This, in turn, allows us to extend the Cocone Triangle Normal Lemma 40 to the interior points of the cocone triangles. The restriction of v to the underlying space $|T|$ of the set of cocone triangles T is a well defined function; refer to Figure 53. For if some point x had more than one closest point on the surface when $\varepsilon \leq 0.05$, x would be a point of the medial axis giving $\|p - x\| \geq f(p)$ for any vertex p of a triangle in T; but by the Small Triangle Lemma 39 every point $q \in |T|$ is within $\frac{1.18\,\varepsilon}{1-\varepsilon} f(p)$ distance of a triangle vertex $p \in \Sigma$ for $\varepsilon \leq 0.05$.
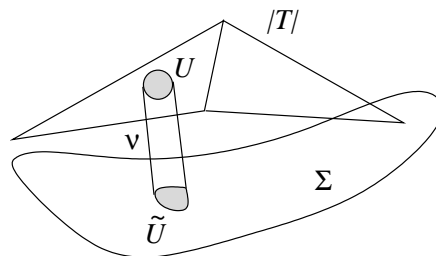

Figure 53: The map v restricted to $|T|$.

In the next two lemmas and also later we use the notation $\tilde{O}(\varepsilon)$ defined in Section 18.3.

Lemma 42. *Let q be any point in a cocone triangle $t \in T$. The distance between q and the point $\tilde{q}$ is $\tilde{O}(\varepsilon)f(\tilde{q})$ and is at most $0.08f(\tilde{q})$ for $\varepsilon \leq 0.05$.*

P   . By the Small Triangle Lemma 39 the circumradius of t is at most $\mu f(p)$ where $\mu = \frac{1.18\,\varepsilon}{1-\varepsilon} \leq$ .07 and p is any of its vertices. Let p be a vertex of t subtending a maximal angle of t. Since there is a sample point, namely a vertex of t, within $\mu f(p)$ distance from q, we have $\|q - \tilde{q}\| \leq \mu f(p)$. We are interested in expressing this bound in terms of $f(\tilde{q})$, so we need an upper bound on $\|p - \tilde{q}\|$.

The triangle vertex p has to lie outside the medial balls at $\tilde{q}$, while, since $\tilde{q}$ is the nearest sur-face point to q, q must lie on the segment between $\tilde{q}$ and the center of one of these medial balls. For any fixed $\|p - q\|$, these facts imply that $\|p - \tilde{q}\|$ is maximized when the angle $\angle pq\tilde{q}$ is a right angle. Thus, $\|p - \tilde{q}\| \leq \sqrt{5}\mu f(p) \leq 0.14f(p)$ for $\varepsilon \leq 0.05$. This implies that $f(p) = \tilde{O}(\varepsilon)f(\tilde{q})$ and in particular $f(p) \leq 1.17f(\tilde{q})$ by Lipschitz property of f. We have $\|q - \tilde{q}\| \leq \mu f(p) = \tilde{O}(\varepsilon)f(\tilde{q})$ and $\|q - \tilde{q}\| \leq 0.08f(\tilde{q})$ in particular.   □

                                                                                                                                                □

With a little more work, we can also show that the triangle normal agrees with the surface normal at $\tilde{q}$.

Lemma 43. *Let q be a point on triangle $t \in T$. The angle $\angle(\mathbf{n}_{\tilde{q}}, \mathbf{n}_p)$ is at most $14°$ where p is a vertex of t with a maximal angle. Also, the angle $\angle_a(\mathbf{n}_{\tilde{q}}, \mathbf{n}_t)$ is $\tilde{O}(\varepsilon)$ and is at most $38°$ for $\varepsilon \leq 0.05$.*

P   . We have already seen in the proof of Lemma 42 that $\|p - \tilde{q}\| = \tilde{O}(\varepsilon)f(p)$. In particu-lar, $\|p - \tilde{q}\| \leq 0.14f(p)$ when $\varepsilon \leq 0.05$. Applying the Normal Variation Lemma 27, and taking $\rho = \tilde{O}(\varepsilon)$ ($\rho = 0.14$ in particular), shows that the angle between $\mathbf{n}_{\tilde{q}}$ and $\mathbf{n}_p$ is $\tilde{O}(\varepsilon)$ and is less than $14°$. The angle between $\mathbf{n}_t$ and $\mathbf{n}_p$ is $\tilde{O}(\varepsilon)$ and is less than $24°$ for $\varepsilon \leq 0.05$ by the Cocone Triangle Normal Lemma 40. Thus, the triangle normal and $\mathbf{n}_{\tilde{q}}$ make $\tilde{O}(\varepsilon)$ angle which is at most $38°$ for $\varepsilon \leq 0.05$.   □

                                                                                                                                                □

Lemma 39, Lemma 42, and Lemma 43 imply that the output surface $|E|$ of C   is close to Σ both point-wise and normal-wise. The following theorem states this precisely.

Theorem 44. *The surface $|E|$ output by C   satisfies the following geometric properties for $\varepsilon \leq 0.05$.*

(i) *Each point $p \in |E|$ is within $\tilde{O}(\varepsilon)f(x)$ distance of a point $x \in \Sigma$. Conversely, each point $x \in \Sigma$ is within $\tilde{O}(\varepsilon)f(x)$ distance of a point in $|E|$.*

(ii) *Each point p in a triangle $t \in E$ satisfies $\angle_a(\mathbf{n}_{\tilde{p}}, \mathbf{n}_t) = \tilde{O}(\varepsilon)$.*

29.1   Additional properties

We argued in Section 28.4 that the underlying space of the simplicial complex output by C   is a 2-manifold. Let E be this simplicial complex output by C   . A pair of triangles $t_1, t_2 \in E$ are *adjacent* if they share at least one common vertex p. Since the normals to all triangles shar-ing p differ from the surface normal at p by at most $24°$ (apply the Cocone Triangle Normal

Lemma 40), and that normal in turn differs from the pole vector at p by less than $7°$ (apply the Pole Lemma 38), we can orient the triangles sharing p, arbitrarily but consistently. We call the normal facing the positive pole the *inside* normal and the normal facing away from it the *outside* normal. Let θ be the angle between the two inside normals of t_1, t_2. We define the angle at which the two triangles meet at p to be $\pi - \theta$.

P   I: Every two adjacent triangles in E meet at their common vertex at an angle greater than $\pi/2$.

Requiring this property excludes manifolds which contain sharp folds and, for instance, flat tunnels. Since the cocone triangles are all nearly perpendicular to the surface normals at their vertices (Cocone Triangle Normal Lemma 40) and the manifold extraction step eliminates triangles adjacent to sharp edges, E has this property.

P   II: Every point in P is a vertex of E.

The Restricted Delaunay Theorem 41 ensures that the set T of cocone triangles contains all restricted Delaunay triangles even after the pruning. Therefore at this point T contains a triangle adjacent to every sample point in P. Lemma 45 below says that each sample point is exposed to the outside for the component of T to which it belongs. This ensures that at least one triangle is selected for each sample point by the manifold extraction step. This implies that E has the second property as well.

Lemma 45 (Exposed.). *Let p be a sample point and let m be the center of a medial ball B tangent to Σ at p. No cocone triangle intersects the interior of the segment pm for $\varepsilon \leq 0.05$.*

P   . In order to intersect the segment pm, a cocone triangle t would have to intersect B and so would the smallest empty ball circumscribing t. Call it D. Let H be the plane of the circle where the boundaries of B and D intersect. See Figure 54. We argue that H separates the interior of pm and t.
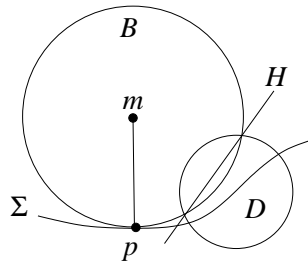

Figure 54: Illustration for the Exposed Lemma.

On one side of H, B is contained in D and on the other, D is contained in B. Since the vertices of t lie on Σ and hence not in the interior of B, t has to lie in the open halfspace, call it H^+, in which D is outside B. Since D is empty, p cannot lie in the interior of D; but since p lies on the boundary of B, it therefore cannot lie in H^+. We claim that $m \notin H^+$ either.

Since $m \in B$, if it lay in H^+, m would be contained in D. Since m is a point of the medial axis, the radius of D would be at least $\frac{f(p')}{2}$ for any vertex p' of t. For $\varepsilon \leq 0.05$, this contradicts the Small Triangle Lemma 39. Therefore p, m, and hence the segment pm cannot lie in H^+ and

H separates *t* and *pm*.   □

30   Topological guarantee

Recall that a function $h : \mathbb{X} \to \mathbb{Y}$ defines a homeomorphism between two compact Euclidean subspaces $\mathbb{X}$ and $\mathbb{Y}$ if *h* is continuous, one-to-one, and onto. In this section, we will show a homeomorphism between Σ and any piecewise-linear 2-manifold made up of cocone triangles from *T*. The piecewise-linear manifold *E* selected by the manifold extraction step is such a space thus completing the proof of homeomorphism.

30.1   The map ν

We define the homeomorphism explicitly, using the function $\nu : \mathbb{R}^3 \setminus M \to \Sigma$, as defined earlier. We will consider the restriction ν' of ν to the underlying space $|E|$ of *E*, i.e., $\nu' : |E| \to \Sigma$. Our approach will be first to show that ν' is well-behaved on the sample points themselves and then show that this property extends in the interior of each triangle in *E*.

Lemma 46. *For $\varepsilon \leq 0.05$, $\nu' : |E| \to \Sigma$ is a well defined continuous function.*

P   . By the Small Triangle Lemma 39, every point $q \in |E|$ is within $\frac{1.18\varepsilon}{1-\varepsilon} f(p)$ of a triangle vertex $p \in \Sigma$ when $\varepsilon \leq 0.05$. Therefore, $|E| \subset \mathbb{R}^3 \setminus M$ for $\varepsilon \leq 0.05$. It follows that ν' is well defined. It is continuous since it is a restriction of a continuous function.   □

                                                                                         □

Let *q* be any point such that $\tilde{q}$ is a sample point *p*. By the Exposed Lemma 45, *q* lies on the segment *pm* where *m* is the center of a medial ball touching Σ at *p*. We have the following.

Corollary 47. *For $\varepsilon \leq 0.05$, the function ν' is one-to-one from $|E|$ to every sample point p.*

In what follows, we will show that ν' is indeed one-to-one on all of $|E|$. The proof proceeds in three short steps. We show that ν' induces a homeomorphism on each triangle, then on each pair of adjacent triangles and finally on $|E|$ as a whole.

Lemma 48. *Let U be a region contained within one triangle $t \in E$ or in adjacent triangles of E. For $\varepsilon \leq 0.05$, the function ν' defines a homeomorphism between U and $\tilde{U} \subset \Sigma$.*

P   . We know that ν' is well defined and continuous on *U*, so it only remains to show that it is one-to-one. First, we prove that if *U* is in one triangle *t*, ν' is one-to-one. For a point $q \in t$, the vector $\mathbf{n}_q$ from $\tilde{q}$ to *q* is perpendicular to the surface at $\tilde{q}$; since Σ is smooth, the direction of $\mathbf{n}_q$ is unique and well defined. If there were some $y \in t$ with $\tilde{y} = \tilde{q}$, then *q*, $\tilde{q}$, and *y* would all be collinear and *t* itself would have to contain the line segment between *q* and *y*, see Figure 55. This implies that the normal $\mathbf{n}_q$ is parallel to the plane of *t*. In other words, $\mathbf{n}_q$ is orthogonal to the normal of *t*, contradicting the Cocone Triangle Normal Lemma 40 which says that the normal of *t* is nearly parallel to $\mathbf{n}_q$.

Now, we consider the case in which *U* is contained in more than one triangle. Let *q* and *y* be two points in *U* such that $\tilde{q} = \tilde{y} = x$ and let *v* be a common vertex of the triangles that contain *U*.
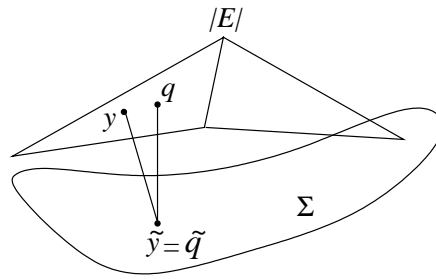
Figure 55: v' maps y and q to the same point which is impossible.

Since v' is one-to-one in one triangle, q and y must lie in the two distinct triangles t_q and t_y. The line l through x with direction $\mathbf{n}_x$ pierces the patch U at least twice; if y and q are not adjacent intersections along l, redefine q so that this is true ($\tilde{q} = x$ for any intersection q of l with U). Now consider the orientation of the patch U according to the direction to the positive pole at v. Either l passes from inside to outside and back to inside when crossing y and q, or from outside to inside and back to outside.

The acute angles between the triangle normals of t_q, t_y and $\mathbf{n}_x$ are less than 38° (Lemma 43), that is, the triangles are stabbed nearly perpendicularly by $\mathbf{n}_x$. But since the orientation of U is opposite at the two intersections, the angle between the two *oriented* triangle normals is greater than 104°, meaning that t_q and t_y must meet at v at an acute angle. This would contradict P I, which is that t_q and t_y meet at v at an obtuse angle. Hence there are no two points y, q in U with $\tilde{q} = \tilde{y}$. □

□

30.2 Homeomorphism proof

We finish the proof for homeomorphism guarantee using a theorem from topology.

Theorem 49 (Homeomorphism.). *The map v' defines a homeomorphism from the surface $|E|$ computed by* C *to the surface Σ for $\varepsilon \leq 0.05$.*

P . Let $\Sigma' \subset \Sigma$ be $v'(|E|)$. We first show that $(|E|, v')$ is a *covering space* of Σ'. Informally, $(|E|, v')$ is a covering space for Σ' if v' maps $|E|$ smoothly onto Σ', with no folds or other singularities. Showing that $(|E|, v')$ is a covering space is weaker than showing that v' defines a homeomorphism, since, for instance, it does not preclude several connected components of $|E|$ mapping onto the same component of Σ', or more interesting behavior, such as a torus wrapping twice around another torus to form a *double covering*.

For a set $X \subseteq \Sigma'$, let $\tau(X)$ denote the set in $|E|$ so that $v'(\tau(X)) = X$. Formally, the $(|E|, v')$ is a covering space of Σ' if, for every $x \in \Sigma'$, there is a path-connected *elementary neighborhood* V_x around x such that each path-connected component of $\tau(V_x)$ is mapped homeomorphically onto V_x by v'.

To construct such an elementary neighborhood, note that the set of points $\tau(x)$ corresponding to a point $x \in \Sigma'$ is non-zero and finite, since v' is one-to-one on each triangle of E and there are
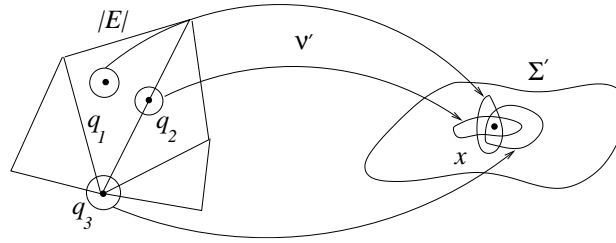
Figure 56: Proof of the Homeomorphism Theorem 49; $\tau(x) = \{q_1, q_2, q_3\}$.

only a finite number of triangles. For each point $q \in \tau(x)$, we choose an open neighborhood U_q of q, homeomorphic to a disk and small enough so that U_q is contained only in triangles that contain q. See Figure 56.

We claim that v' maps each U_q homeomorphically onto $\tilde{U}_q$. This is because it is continuous, it is onto $\tilde{U}_q$ by definition, and, since any two points x and y in U_q are in adjacent triangles, it is one-to-one by Lemma 48.

Let $U'(x) = \bigcap_{q \in \tau(x)} v'(U_q)$, the intersection of the maps of each of the U_q. $U'(x)$ is the intersection of a finite number of open neighborhoods, each containing x, so we can find an open disk V_x around x. V_x is path connected and each component of $\tau(V_x)$ is a subset of some U_q and hence is mapped homeomorphically onto V_x by v'. Thus $(|E|, v')$ is a covering space for Σ'.

We now show that v' defines a homeomorphism between $|E|$ and Σ'. Since $v' : |E| \to \Sigma'$ is onto by definition, we need only that v' is one-to-one. Consider one connected component G of Σ'. A theorem of algebraic topology says that when $(|E|, v')$ is a covering space of Σ', the sets $\tau(x)$ for all $x \in G$ have the same cardinality. We now use Corollary 47, that v' is one-to-one at every sample point. Since each connected component of Σ contains some sample points, it must be the case that v' is everywhere one-to-one and $|E|$ and Σ' are homeomorphic.

Finally, we show that $\Sigma' = \Sigma$. Since $|E|$ is a 2-manifold without boundary and is compact, Σ' must be as well. So Σ' cannot include part of a connected component of Σ, and hence Σ' must consist of a subset of the connected components of Σ. Since every connected component of Σ contains a sample p (actually many sample points) and $v'(p) = p$, all components of Σ belong to Σ'. Therefore, $\Sigma' = \Sigma$ and $|E|$ and Σ are homeomorphic. $\square$

                                                                                                                    $\square$

It can also be shown that $|E|$ and Σ are isotopic (Exercise 7). We will show a technique to prove isotopy in Section **??**.

31   Notes and exercises

The first algorithm for surface reconstruction with proved guarantees was devised by Amenta and Bern [1]. They generalized the C   algorithm for curve reconstruction to the surface reconstruction problem. The idea of poles and approximating the normals with the pole vector was a significant breakthrough. The crust triangles (Exercise 2) enjoy some nice properties that help the reconstruction. The C   algorithm as described here is a successor of C   . Devised by Amenta, Choi, Dey, and Leekha [2], this algorithm simplified the C   algorithm and its proof of

correctness. C   eliminated one of the two Voronoi diagram computations of C   and also a normal filtering step. The homeomorphism between the reconstructed surface and the original sampled surface was first established in [2].

Exercises

1. We know that Voronoi vertices for a dense sample from a curve in the plane lie near the medial axis. The same is not true for surfaces in three dimensions. Show an example where a Voronoi vertex for an arbitrarily dense sample lies arbitrarily close to the surface.

2^h. Let P be a sample from a C^2-smooth surface Σ and V be the set of poles in Vor P. Consider the following generalization of the C   . A triangle in the Del $(P \cup V)$ is a crust triangle if all of its vertices are in P. Show the following when P is an ε-sample for a sufficiently small ε.

   (i) All restricted Delaunay triangles in Del$|_\Sigma (P \cup V)$ are crust triangles.

   (ii) All crust triangles have circumradius $\tilde{O}(\varepsilon)f(p)$ where p is a vertex of the triangle.

3. Let t be a triangle in Del P where $B = B_{v,r}$ and $B' = B'_{v',r'}$ are two Delaunay balls circumscribing t. Let x be any point on the circle where the boundaries of B and B' intersect. Show that, if $\angle vxv' > \frac{\pi}{2}$, the triangle normal of t makes an angle of $\tilde{O}(\varepsilon)$ with the normals to Σ at its vertices when P is an ε-sample of Σ for a sufficiently small ε.

4. Recall that P is a locally (ε, δ)-uniform sample of a smooth surface Σ if P is an ε-sample of Σ and each sample point $p \in P$ is at least $\frac{\varepsilon}{\delta}f(p)$ distance away from all other points in P where $\delta > 1$ is a constant. Show that each triangle in the surface output by C   for such a sample has a bounded aspect ratio (circumradius to edge length ratio). Also prove that each vertex has no more than a constant number (determined by ε and δ) of triangles on the surface.

5^h. Let t be a cocone triangle. We showed that any point $x \in t$ is $\tilde{O}(\varepsilon)f(\tilde{x})$ away from its closest point $\tilde{x}$ in Σ. Prove that the bound can be improved to $\tilde{O}(\varepsilon^2)f(\tilde{x})$.

6. We defined a Delaunay triangle t as a cocone triangle if dual t intersects cocones of *all* of its three vertices. Relax the condition by defining t as a cocone triangle if dual t intersects the cocone of *any* of its vertices. Carry out the proofs of different properties of cocone triangles with this modified definition.

7. We showed that the surface $|E|$ computed by C   is homeomorphic to Σ when ε is sufficiently small. Prove that $|E|$ is indeed isotopic to Σ.

References

[1] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discr. Comput. Geom.*, 22:481–504, 1999.

[2] N. Amenta, S. Choi, T. K. Dey and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. *Internat. J. Comput. Geom. Appl.*, 12:125–141, 2002.

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