Restricted Constrained Delaunay Triangulations

Marc Khoury
University of California, Berkeley, USA

Jonathan Richard Shewchuk
University of California, Berkeley, USA

Abstract

We introduce the restricted constrained Delaunay triangulation (restricted CDT), a generalization of both the restricted Delaunay triangulation and the constrained Delaunay triangulation. The restricted CDT is a triangulation of a surface whose edges include a set of user-specified constraining segments. We define the restricted CDT to be the dual of a restricted Voronoi diagram defined on a surface that we have extended by topological surgery. We prove several properties of restricted CDTs, including sampling conditions under which the restricted CDT contains every constraining segment and is homeomorphic to the underlying surface.

2012 ACM Subject Classification Theory of computation → Randomness, geometry and discrete structures

Keywords and phrases restricted Delaunay triangulation, constrained Delaunay triangulation, surface meshing, surface reconstruction, topological surgery, portals

Related Version A complete version is available at `https://people.eecs.berkeley.edu/~jrs/papers/rcdt.pdf`.

Funding Supported in part by the National Science Foundation under Awards CCF-1423560 and CCF-1909204 and in part by a National Science Foundation Graduate Research Fellowship.

Acknowledgements This work was initiated at the Workshop on Geometric Algorithms in the Field hosted by the Lorentz Center in Leiden, the Netherlands during June 2014. We thank the organizers—Sándor Fekete, Maarten Löffler, Bettina Speckmann, and Jo Wood—and the Lorentz Center for providing accommodations. We especially thank Bruno Lévy for posing the problem this paper answers, Marc van Kreveld for helpful discussions, and the referees for improving the paper.

1 Introduction

The constrained Delaunay triangulation (CDT) in the plane [22, 29, 14] is a popular geometric construction that shares some of the advantages and mathematical properties of the Delaunay triangulation, but also permits users to constrain specified edges to be part of the triangulation. CDTs are used in applications such as computer graphics, geographical information systems, and guaranteed-quality mesh generation algorithms [13]. Our goal here is to offer a mathematically rigorous way to define a Delaunay-like triangulation on a curved surface embedded in three-dimensional space, with the same ability to constrain edges.

Another variant of the Delaunay triangulation, called the *restricted Delaunay triangulation* (RDT), has become a well-established way of generating triangulations on curved surfaces [17]. RDTs have equipped theorists to rigorously prove the correctness of algorithms for surface reconstruction [15] and surface mesh generation [13]. In this paper we introduce *restricted constrained Delaunay triangulations* (restricted CDTs), which combine ideas from CDTs and RDTs to enable the enforcement of constraining edges in RDTs.

Think of the restricted CDT as a function that takes in three inputs: a compact, smooth surface $\Sigma \subset \mathbb{R}^3$ without boundary; a finite set $V \subset \Sigma$ of points (called *sites* or *vertices*); and a finite set S of line segments whose endpoints are in V. If certain conditions on the density of V and the lengths of the segments are met then, as illustrated in Figure 1, the output

is a simplicial complex $\mathcal{T}$ such that the set of vertices of $\mathcal{T}$ is V, the set of edges of $\mathcal{T}$ is a superset of S, and $\mathcal{T}$ is a triangulation of Σ. The last phrase means that the *underlying space* of $\mathcal{T}$, written $|\mathcal{T}| = \bigcup_{\tau \in \mathcal{T}} \tau$, is homeomorphic to Σ.

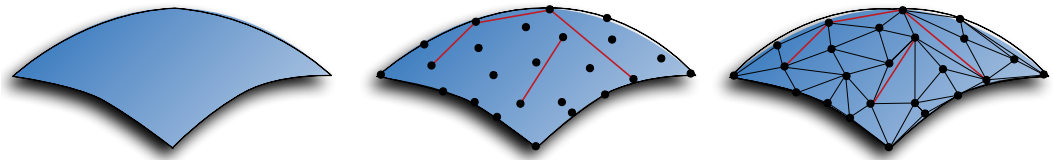

Figure 1 Given a set of points sampled from a surface Σ and a set of segments, red, we wish to compute a triangulation of Σ that contains all of the red segments.

Although Delaunay triangulations in the plane can be constrained to include arbitrary edges, the same is not true of three-dimensional Delaunay triangulations; consider the fact that not all nonconvex polyhedra can be tetrahedralized [28]. Nor is it always possible to constrain arbitrary edges to be part of a surface triangulation. Our challenge is to establish conditions on the input that guarantee that a suitable triangulation exists.

We follow the example of the RDT, which is defined by dualizing a *restricted Voronoi diagram*. Given inputs Σ and V (but no segments), the *restricted Voronoi cell* of a site $v \in V$, denoted $\mathrm{Vor}|_\Sigma v$, is the set of all points on Σ for which v is the closest site in V (possibly tied for closest), as measured by the Euclidean distance in $\mathbb{R}^3$. Equivalently, $\mathrm{Vor}|_\Sigma v = \mathrm{Vor}\, v \cap \Sigma$, where $\mathrm{Vor}\, v$ is v's standard Voronoi cell in $\mathbb{R}^3$. The name "restricted Voronoi cell" arises because $\mathrm{Vor}|_\Sigma v$ is the restriction of $\mathrm{Vor}\, v$ to the surface Σ.

A *restricted Voronoi face* is any nonempty set of points found by taking the intersection of one or more restricted Voronoi cells. The *restricted Voronoi diagram* $\mathrm{Vor}|_\Sigma V$ is the cell complex containing all the restricted Voronoi cells and faces.

The *restricted Delaunay triangulation* $\mathrm{Del}|_\Sigma V$ is the simplicial complex dual to $\mathrm{Vor}|_\Sigma V$. If the restricted Voronoi cells of two sites $v, w \in V$ have a nonempty intersection (typically a path on Σ), then vw is a *restricted Delaunay edge* in $\mathrm{Del}|_\Sigma V$. If the restricted Voronoi cells of three sites $u, v, w \in V$ have a nonempty intersection (typically a single point on Σ, called a *restricted Voronoi vertex*), then $\triangle uvw$ is a *restricted Delaunay triangle* in $\mathrm{Del}|_\Sigma V$. Every site in V is a vertex in $\mathrm{Del}|_\Sigma V$. Note that $\mathrm{Del}|_\Sigma V$ may not be a valid simplicial complex unless V is a sufficiently dense sample of Σ, perhaps with suitable perturbations of Σ and V. See Section 3 for a more nuanced discussion.

To modify RDTs so that we can constrain edges, we borrow from Seidel [29] the idea of an *extended Voronoi diagram*, which is the natural dual of the CDT in the plane. Seidel performs a topological surgery on the plane in which each segment in S becomes a slit cut in the plane; upon these slits he glues topological extensions called "secondary sheets" on which additional portions of the extended Voronoi diagram are drawn. Likewise, we perform surgery by cutting slits in the surface Σ and grafting independent new surfaces called *extrusions* onto Σ at these slits. We think of these slits as *portals*: an ant crawling on the surface across a constraining edge finds itself transported by the portal to a secondary space where the extended surface continues along an infinite extrusion.

A key contribution of this paper is our definition of the restricted constrained Delaunay triangulation, as the dual of the Voronoi diagram restricted to this surgically extended surface. Another contribution is to prove several properties of restricted CDTs, including conditions under which the restricted CDT contains every constraining edge, conditions under which the restricted CDT is homeomorphic to the underlying surface Σ, and a characterization of

which vertices must be considered to compute the triangles near a segment.

Shewchuk [30] demonstrates that for Delaunay mesh generators that create high-quality meshes of domains in the plane with constraining segments, the use of a CDT (rather than a pure Delaunay triangulation) reduces the number of triangles and vertices—on some domains, by as much as 25%. He also proves that there is a theoretical advantage: Delaunay meshing with a CDT offers a guarantee of a "size optimal" mesh with no angle less than 26.56°, whereas an *unconstrained* Delaunay triangulation offers a weaker guarantee, a size optimal mesh with no angle less than 20.7°. It is very likely that surface meshing algorithms based on restricted CDTs can offer the same advantages, compared to what pure RDTs can achieve.

An alternative approach sometimes suggested is to define a Voronoi diagram based on an intrinsic (geodesic) distance metric, then obtain a triangulation by duality. This idea is mathematically elegant, but computing a geodesic Voronoi diagram entails numerical approximation algorithms [21, 23, 24], which add coding complexity and running time. RDTs are popular in surface mesh generation because they are easier to compute. We emphasize that although our construction of restricted CDTs may seem complicated, it is in the service of producing simple algorithms. (In particular, Theorems 1 and 3 simplify computing the triangles near a segment.) See Section 6 for some speculation on prospective algorithms.

2 Portals and topological surgery

Informally, a *portal* P is a doorway between two topological spaces, with P shared by both. Our main topological construction starts with disjoint topological spaces Y and Z, then glues them together into a single space by specifying an equivalence relationship between a subset of points $P \subset Y$ and a subset $P' \subset Z$. For clarity, we explain Seidel's construction of portals in the plane [29] first, then our construction of portals and an extended surface in $\mathbb{R}^3$.

2.1 Portals and extended Voronoi diagrams in the plane

Let $X = \mathbb{R}^2$ and let S be a finite set of line segments in the plane; the segments may intersect each other only at their endpoints. Consider a segment $s = pq \in S$ (meaning s has endpoints p and q). The relative interior of s, denoted $\operatorname{relint} s$, consists of all points on s except p and q. Let the *slitted plane* $X_s = X - \operatorname{relint} s$ be the plane with the relative interior of s removed. The affine hull of s has two "sides." Our goal is to augment X_s by gluing it to two additional topological spaces, one for each side of s, along the slit created by removing $\operatorname{relint} s$. The three spaces are glued together along two portals, each of which is topologically a copy of s. Thus an ant crawling on the extended space that crosses s from one side finds itself in a *secondary branch*; and an ant that crosses s from the other side finds itself in a different secondary branch. After repeating this augmentation for every segment in S, we can draw on the extended space an *extended Voronoi diagram* whose dual is the CDT.

Topologically, X_s has a hole such that X_s is *almost* an open set, except that X_s has two boundary points, p and q. We want to glue two additional spaces to X_s—one for each side of s—so we augment X_s with additional points that serve as two portals to those additional spaces. We define a closed topological space $\overline{X}_s$ by augmenting X_s with two connected curves ζ^+ and ζ^-, called *portals*, that together serve as the boundary of the hole. Each of ζ^+ and ζ^- has p and q as its endpoints, but the two curves share no other points. In essence, the portals are copies of s with shared endpoints. Formally, $\overline{X}_s$ is the completion of the incomplete metric space X_s with respect to the shortest-path metric in X_s.

The points in X_s inherit Cartesian coordinates from the plane, and the points on the portals ζ^+ and ζ^- inherit Cartesian coordinates from the segment s. Two points in $\overline{X}_s$—one

on ζ^+ and one on ζ^-—can have the same (x, y)-coordinate values yet be topologically distinct.

Let $\mathbb{R}^2_-$ and $\mathbb{R}^2_+$ be two copies of $\mathbb{R}^2$. We treat $\overline{X}_s$, $\mathbb{R}^2_-$, and $\mathbb{R}^2_+$ as three distinct topological spaces that all inherit the Cartesian coordinate system—so two points in two different spaces can have the same coordinate values yet be topologically distinct.

Informally, we glue $\mathbb{R}^2_+$ to $\overline{X}_s$ along ζ^+ and glue $\mathbb{R}^2_-$ to $\overline{X}_s$ along ζ^-. Formally, we write $x \equiv y$ if x and y have the same coordinate values, even though they may lie in different spaces. Let p and q be the endpoints of s. Define an equivalence relation $\sim$ as

$$x \sim y \iff \begin{cases} x = y & x, y \in \overline{X}_s \text{ or } x, y \in \mathbb{R}^2_+ \text{ or } x, y \in \mathbb{R}^2_-, \\ x \equiv y & x \in \mathbb{R}^2_+ \text{ and } y \in \zeta^+, \\ x \equiv y & x \in \mathbb{R}^2_- \text{ and } y \in \zeta^-, \\ x \equiv p \equiv y \text{ or } x \equiv q \equiv y & x \in \mathbb{R}^2_+ \text{ and } y \in \mathbb{R}^2_-. \end{cases} \tag{1}$$

With $\sim$ we construct the quotient space $\widetilde{X} = (\overline{X}_s \sqcup \mathbb{R}^2_+ \sqcup \mathbb{R}^2_-)/\sim$. We refer to $\overline{X}_s$ as the *principal branch* and refer to $\mathbb{R}^2_+$ and $\mathbb{R}^2_-$ as *secondary branches*. Figures 2 and 3 illustrate this construction. Note that in the quotient space, the endpoints p and q of the segment s are present in, and shared by, all three of the original spaces.

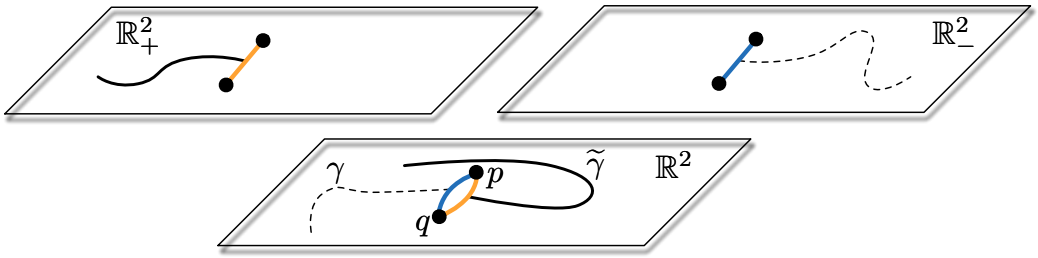

Figure 2 The completion of the slitted plane has a topological hole bounded by two portals, marked in blue and orange. (Geometrically, the two portals are straight line segments that occupy exactly the same coordinates.) The equivalence relation $\sim$ identifies the blue path in the principal branch with the blue path in $\mathbb{R}^2_-$; likewise the two orange paths become one. A path in the principal branch (bottom) that enters a portal continues in the appropriate secondary branch.

The construction works for any finite number $m = |S|$ of non-crossing segments. Let $X_S = X - \bigcup_{s \in S} \operatorname{relint} s$. Let $\overline{X}_S$ be the completion of X_S with respect to the shortest-path metric in X_S, which adds two portals for each segment. Then we construct a quotient space $\widetilde{X}$ composed of $\overline{X}_S$ and $2m$ copies of $\mathbb{R}^2$ glued along the $2m$ portals bounding the m holes in $\overline{X}_S$.

For the sake of defining the Voronoi diagram of a finite set of sites in $\widetilde{X}$, Seidel [29] defines a distance function on $\widetilde{X}$ which is essentially the Euclidean distance, except that the distance between two points is infinite if they are not visible from each other. (Note that this distance function is not a metric.) A path $\gamma \subset \widetilde{X}$ may pass through portals and visit secondary branches, but because of the slits we have cut in X_S, γ cannot cross the relative interior of a segment without being transported by a portal. We call a path *straight* if its Cartesian embedding is a straight line segment. Two points $p, q \in \widetilde{X}$ are *visible* from each other if there is a straight path $\gamma \subset \widetilde{X}$ with endpoints p and q. The distance $\widehat{d}(p, q)$ from p to q is the Euclidean distance if p and q are visible from each other; otherwise, $\widehat{d}(p, q) = \infty$.

The extended Voronoi diagram assigns each point in $\widetilde{X}$ to (the Voronoi cells of) one or more sites in V. Those sites must be visible from the point; no site can claim a point it cannot see. If a point on a secondary branch is claimed by a site other than the branch's portal's

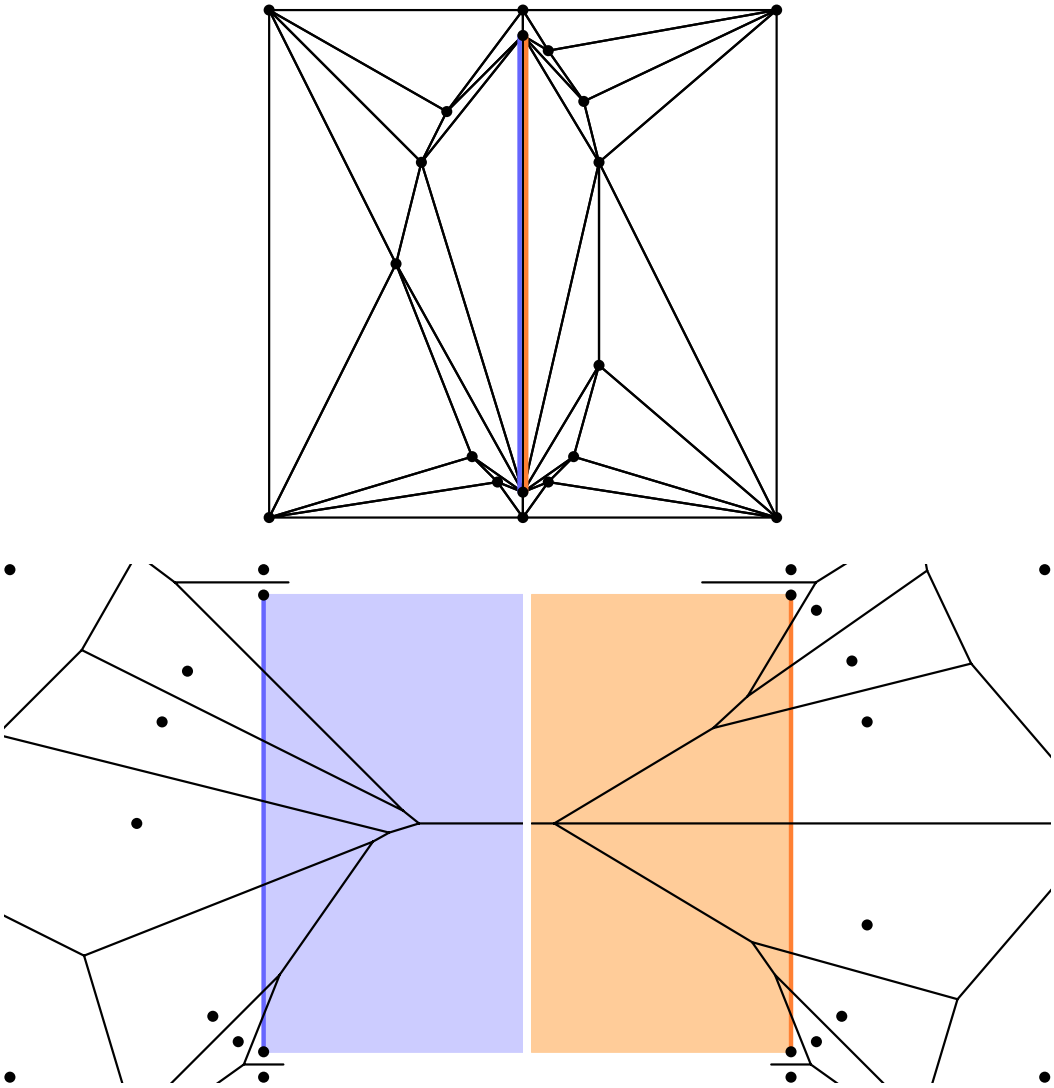

Figure 3 A one-segment CDT (top) and its dual extended Voronoi diagram (bottom). The blue and orange regions show the portions of the Voronoi diagram on the secondary branches.

endpoints, the site is visible from the point through the portal, as Figure 3 illustrates. Seidel gives an algorithm for constructing the extended Voronoi diagram, and by duality the CDT.

2.2 Portals on surfaces embedded in $\mathbb{R}^3$

A similar construction works for a compact, smooth surface without boundary $\Sigma \subset \mathbb{R}^3$. However, whereas in the plane we construct one new topological space, here we will require two. We surgically augment the surface Σ by cutting slits along *portal curves*, one for each segment, and gluing two *extrusions* onto each portal curve, yielding an *extended surface* $\widetilde{\Sigma}$. The purpose of this extended surface is to serve as a canvas upon which we can draw an extended restricted Voronoi diagram, which we dualize to define a restricted CDT.

To define a Voronoi diagram we need a distance function, and $\widetilde{\Sigma}$ alone does not provide one that is easily computed. While an intrinsic (geodesic) distance might be ideal in principle,

for the sake of speed and a simple implementation, we use the Euclidean distance in $\mathbb{R}^3$ as RDTs do; but we must modify the Euclidean distance so that the restricted Voronoi diagram respects the input segments. Hence most of our work will be to construct a surgically modified three-dimensional space $\widetilde{X}$ in which we embed $\widetilde{\Sigma} \subset \widetilde{X}$. Like Seidel's extended space in Section 2.1, $\widetilde{X}$ obstructs (and supports) visibility in a manner that is suitable for defining a restricted Voronoi diagram on $\widetilde{\Sigma}$ and makes it easy to compute restricted CDTs.

To define $\widetilde{X}$, we specify portals in $\mathbb{R}^3$ where points will be removed, analogous to cutting slits in the plane. Each portal is a two-dimensional ruled surface with boundary (not generally flat), approximately perpendicular to Σ. The intersection of a portal with Σ is a portal curve. Each portal has two "sides," and on each side we glue an additional copy of $\mathbb{R}^3$ to form $\widetilde{X}$. In each copy of $\mathbb{R}^3$ we embed an extrusion to form $\widetilde{\Sigma}$. The extended Voronoi diagram assigns each point x on $\widetilde{\Sigma}$ to one or more sites in V that are visible from x along straight paths in $\widetilde{X}$.

To define portal geometry, we need several definitions. The *medial axis* M of Σ is the closure of the set of all points in $\mathbb{R}^3$ for which the closest point on Σ is not unique. Intuitively, the medial axis of Σ is meant to capture the "middle" of the region bounded by Σ. A *medial ball* is a ball whose center lies on M and whose boundary intersects Σ (tangentially), but the interior of the ball does not. For any point $x \in \Sigma$, there are one or two medial balls that have x on their boundaries, called *medial balls at* x. If there are two, there is one on each side of Σ. If there is only one, it is enclosed by Σ.

For $x \in \Sigma$, the *normal line* $\mathcal{L}_x$ at x is the line orthogonal to Σ at x with $x \in \mathcal{L}_x$. The *normal segment* ℓ_x at x is a line segment or ray whose endpoints lie on M, satisfying $x \in \ell_x \subset \mathcal{L}_x$. If there are two medial balls at x, the endpoints of ℓ_x are the centers of those two medial balls. If there is only one medial ball at x, then ℓ_x is a ray originating at the medial ball's center.

The *local feature size* function is $\mathrm{lfs} : \Sigma \to \mathbb{R}$, $x \mapsto d(x, M)$ where $d(x, M)$ denotes the Euclidean distance from x to M. We require that Σ is smooth in the sense that $\inf_{x \in \Sigma} \mathrm{lfs}(x) > 0$. A finite point set $V \subset \Sigma$ is an *ϵ-sample of* Σ if for every point $x \in \Sigma$, $d(x, V) \leq \epsilon\, \mathrm{lfs}(x)$. That is, the ball with center x and radius $\epsilon\, \mathrm{lfs}(x)$ contains at least one point in V. See Figure 4.

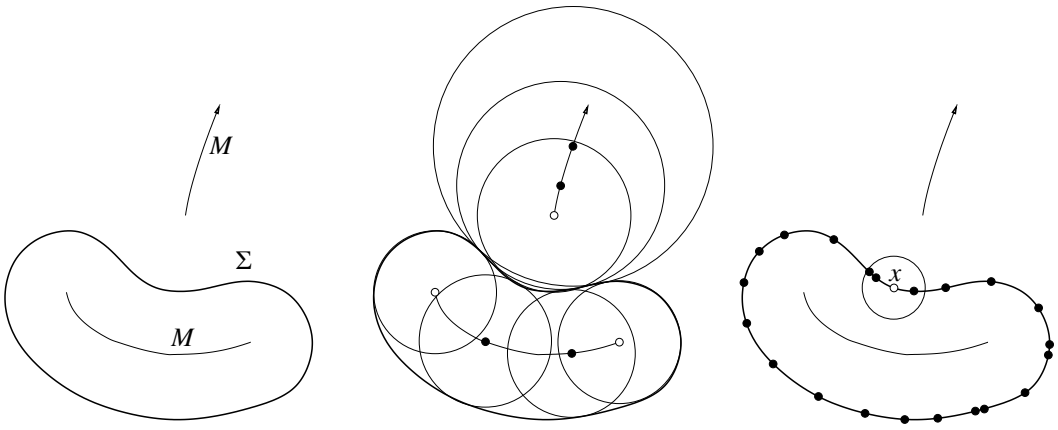

Figure 4 Left: A 1-manifold Σ and its medial axis M (as medial axes in three dimensions are hard to draw or understand). This medial axis is unbounded; one of its components extends infinitely far away. Center: Some of the medial balls that define M. Right: A 0.5-sample of Σ (filled circles). The ball with center x and radius $0.5\,\mathrm{lfs}(x)$ contains a site.

Let S be a finite set of line segments whose endpoints are in V, called the *segments*, which constrain the restricted CDT. Consider a segment $s = pq \in S$ (its endpoints are $p, q \in V$). Let B_s be the *diametric ball* of s—the smallest closed ball such that $s \subset B_s$, so that s is a

diameter of B_s. Suppose that $d(p,q) \leq \rho\,\mathrm{lfs}(p)$ for some $\rho \in (0,1)$; that is, s is short relative to the local feature size. Then $B_s \cap \Sigma$ is a topological disk [13, Lemma 12.6].

Suppose that we know (or can approximate) the unit vector n_p normal to Σ at any site p. We choose a *cutting plane* $h_s \supset s$ that is locally orthogonal to the surface Σ at p or q (or perhaps somewhere between p and q). We use h_s to specify a *portal curve* $\zeta_s = h_s \cap B_s \cap \Sigma$, which is a single connected curve from p to q on Σ. There is not a canonical choice of cutting plane (and thus portal curve) for s, and the user might be presented with a range of choices, but for our presentation here, we choose $h_s = \mathrm{span}\{n_p, \vec{pq}\}$. We require that the portal curves do not cross each other. More precisely, the relative interior of a portal curve may not intersect another portal curve nor a site in V.

Our requirement that each portal curve must lie on a plane has both a theoretical motivation and a practical one. The fact that every constraining segment is an edge in the restricted CDT (Theorem 2) depends on the fact that each portal curve lies in a plane and its extrusions are orthogonal to that plane. The requirement simplifies algorithms for computing a restricted CDT, because the Voronoi cells on an extrusion are solely influenced by sites on the other side of the cutting plane—plus the segment endpoints p and q. (See Theorem 1.)

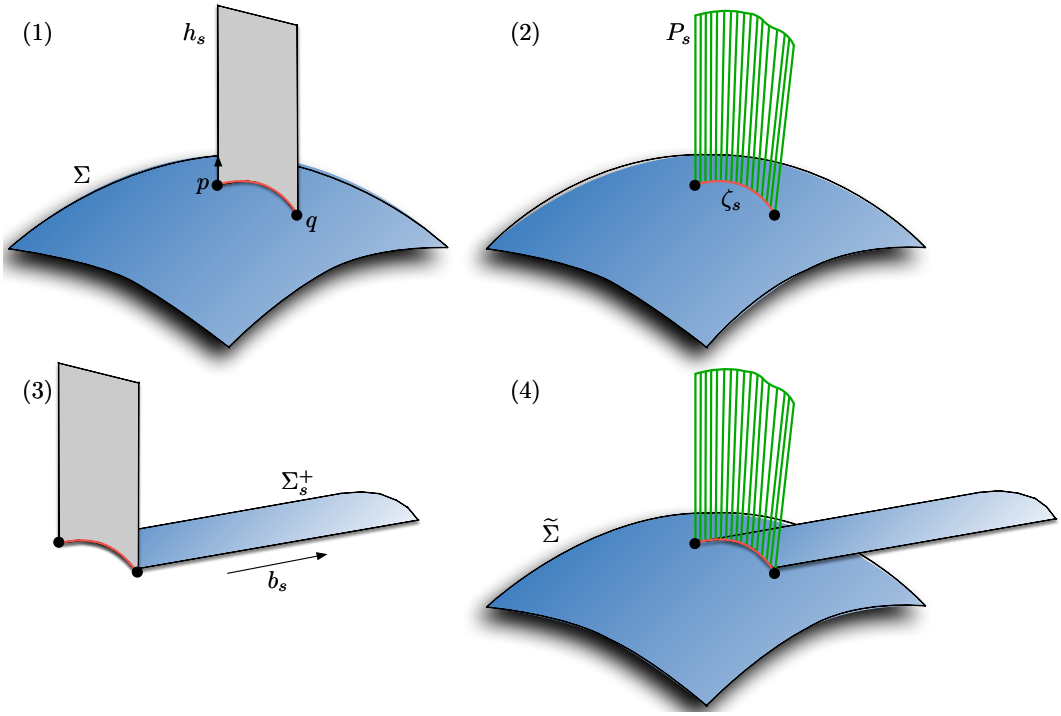

Figure 5 (1) The plane h_s intersects Σ in a curve; the portal curve ζ_s (red) is the portion of this curve in the diametric ball B_s of the segment $s = pq$. (2) Our portal P_s, shown in green, is the union of the normal segments (locally orthogonal to Σ) of the points on the portal curve ζ_s. The normal segments terminate on the medial axis M. (3) We extrude the portal curve ζ_s into $\mathbb{R}^3_+$ in a direction b_s orthogonal to h_s, thus defining Σ_s^+. (4) We glue the extrusion Σ_s^+ to $\overline{\Sigma}_S$ (the surface Σ with slits cut into it) along ζ_s^+ at the entrance to the portal P_s^+.

Figure 5 illustrates our portal construction. For each segment s, the portal $P_s = \bigcup_{x \in \zeta_s} \ell_x$ is the union of the normal segments of the points on the portal curve ζ_s. Hence a portal is a ruled surface, topologically two-dimensional but not lying in a plane. Each portal reaches to

the medial axis, thereby obstructing visibility so that sites on one "side" of a segment do not influence the restricted Delaunay triangles on the other "side."

If two segments share an endpoint p, then their portals share the boundary segment ℓ_p. The other location where portals' boundaries may intersect each other is at the medial axis. However, no portal intersects the relative interior of another portal.

We construct the extended space $\widetilde{X}$ as we did in Section 2.1, with P_s replacing s and $\mathbb{R}^3$ replacing $\mathbb{R}^2$. Let $X = \mathbb{R}^3$. Let $X_S = X \setminus \bigcup_{s\in S} \operatorname{relint} P_s$, which is $\mathbb{R}^3$ with the relative interior of each portal removed. Let $\overline{X}_S$ be the completion of the incomplete metric space X_S endowed with the shortest path metric. The effect of completing X_S is to augment each "slit" P_s with two portals P_s^+ and P_s^-, one for each side of P_s. These two portals are distinct copies of P_s, but P_s^+ and P_s^- share a common boundary $\partial P_s = P_s^+ \cap P_s^- = P_s^+ \cap X_S = P_s^- \cap X_S$.

For each segment $s \in S$, let $\mathbb{R}^3_{s+}$ and $\mathbb{R}^3_{s-}$ be two topologically distinct copies of $\mathbb{R}^3$, called *secondary branches*. The points in each secondary branch and the points in the *principal branch* $\overline{X}_S$ all inherit Cartesian coordinates, but points with the same coordinates in different branches are topologically distinct. Define an equivalence relation $\sim$ analogous to (1) that identifies (glues) the points of the portal $P_s^+ \subset \overline{X}_S$ with the points in $\mathbb{R}^3_{s+}$ having the same coordinates, and identifies the points of $P_s^- \subset \overline{X}_S$ with the corresponding points in $\mathbb{R}^3_{s-}$. Thus we glue $2m$ copies of $\mathbb{R}^3$ along the $2m$ portals bounding the m holes in $\overline{X}_S$. The *extended space* is the quotient space $\widetilde{X} = (\overline{X}_S \sqcup \bigsqcup_{s\in S} \mathbb{R}^3_{s+} \sqcup \bigsqcup_{s\in S} \mathbb{R}^3_{s-})/\!\sim$.

Similarly, we surgically modify Σ to construct an extended surface $\widetilde{\Sigma} \subset \widetilde{X}$, as shown in the bottom two illustrations in Figure 5. Let $\Sigma_S = \Sigma \setminus \bigcup_{s\in S} \operatorname{relint} \zeta_s$ be the surface with the portal curve interiors removed, and let the *principal surface* $\overline{\Sigma}_S \subset \overline{X}_S$ be its completion. For each $s \in S$, $\overline{\Sigma}_S$ includes two portal curves $\zeta_s^+ \subset P_s^+$ and $\zeta_s^- \subset P_s^-$, one for each side of the cutting plane h_s. We extrude ζ_s^+ into $\mathbb{R}^3_{s+}$ and ζ_s^- into $\mathbb{R}^3_{s-}$, each in one of the two directions normal to the cutting plane h_s. Let b_s be a unit vector normal to h_s, directed to pass through P_s^+ from the principal branch to $\mathbb{R}^3_{s+}$. For each point $x \in \zeta_s$ we define a ray $x_s^+ = \{x + \omega b_s \in \mathbb{R}^3_{s+} : \omega \in [0,\infty)\}$ and a ray $x_s^- = \{x - \omega b_s \in \mathbb{R}^3_{s-} : \omega \in [0,\infty)\}$ (specifying points by their coordinates). We then define two *extrusions*, the ruled surfaces $\Sigma_s^+ = \{x_s^+ : x \in \zeta_s\} \subset \mathbb{R}^3_{s+}$ and $\Sigma_s^- = \{x_s^- : x \in \zeta_s\} \subset \mathbb{R}^3_{s-}$. The *extended surface* is $\widetilde{\Sigma} = (\overline{\Sigma}_S \sqcup \bigsqcup_{s\in S} \Sigma_s^+ \sqcup \bigsqcup_{s\in S} \Sigma_s^-)/\!\sim$.

3 Restricted constrained Delaunay triangulations

To define the restricted constrained Delaunay triangulation, we first define the *extended restricted Voronoi diagram* (or just *extended Voronoi diagram* for short) on the extended surface $\widetilde{\Sigma}$. As in Section 2.1, we define a distance function $\widehat{d}(p,q)$ that is the Euclidean distance in $\mathbb{R}^3$ if p and q are visible to each other along a straight path in $\widetilde{X}$, or ∞ if they cannot see each other. For any $v \in V$, the *extended restricted Voronoi cell* of v is

$$\operatorname{Vor}|_{\widetilde{\Sigma}} v = \{x \in \widetilde{\Sigma} : \widehat{d}(x,v) \le \widehat{d}(x,u),\ \forall u \in V\}.$$

An *extended restricted Voronoi face* is any nonempty intersection of one or more extended restricted Voronoi cells. The *extended restricted Voronoi diagram* $\operatorname{Vor}|_{\widetilde{\Sigma}} V$ is the cell complex containing all the extended restricted Voronoi cells and faces.

We define the *restricted constrained Delaunay subdivision* $\operatorname{Del}|_{\widetilde{\Sigma}} V$ to be the polyhedral complex dual to the extended Voronoi diagram in this sense: for each extended Voronoi face $f \in \operatorname{Vor}|_{\widetilde{\Sigma}} V$, let $W \subseteq V$ be the set of sites whose restricted Voronoi cells include f and let f^* be the convex hull of W. We say that f^* is the face *dual* to f. Then $\operatorname{Del}|_{\widetilde{\Sigma}} V = \{f^* : f \in \operatorname{Vor}|_{\widetilde{\Sigma}} V\}$.

A one-point face in $\operatorname{Vor}|_{\widetilde{\Sigma}} V$ is called an *extended restricted Voronoi vertex*, and its dual is a polygonal or polyhedral face in $\operatorname{Del}|_{\widetilde{\Sigma}} V$, usually a triangle. If an intersection of two extended

restricted Voronoi cells includes a path on Σ, we call it an *extended restricted Voronoi edge*, and its dual is a (straight) *restricted constrained Delaunay edge* in $\mathrm{Del}|_{\widetilde{\Sigma}} V$. Figure 6 illustrates an extended Voronoi vertex on an extrusion and its dual restricted Delaunay triangle, as well as three extended Voronoi edges and their dual restricted constrained Delaunay edges.

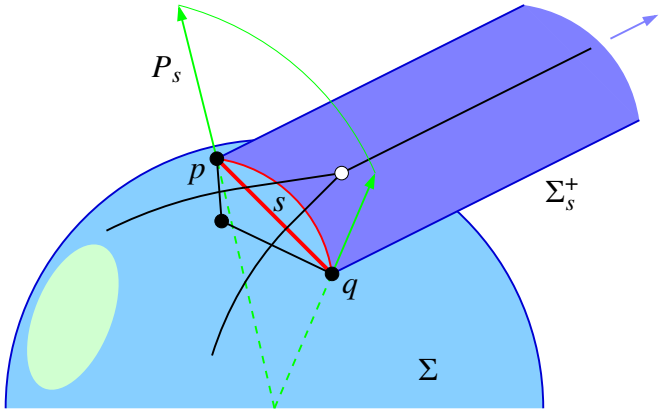

Figure 6 An extended Voronoi vertex on an extrusion and its dual restricted Delaunay triangle.

If $\mathrm{Del}|_{\widetilde{\Sigma}} V$ contains a polyhedron, we can perturb $\widetilde{\Sigma}$ infinitesimally so that $\widetilde{\Sigma}$ does not pass through the polyhedron's circumcenter; thus with suitable perturbations, $\mathrm{Del}|_{\widetilde{\Sigma}} V$ contains no polyhedra. Relatedly, just as a standard Delaunay triangulation in the plane can be ambiguous if four vertices lie on a common circle, if V has four or more cocircular vertices then $\mathrm{Del}|_{\widetilde{\Sigma}} V$ might contain polygons with four or more sides. If desired, a triangulation can be obtained by subdividing each polygonal face into triangles or by an infinitesimal perturbation of V. We recommend the former in practice, but for the sake of our proofs, we will exploit the latter. For simplicity, we will assume in this paper that no point on $\widetilde{\Sigma}$ is equidistant from four visible vertices in V; then every polygonal face is a triangle and we can call $\mathrm{Del}|_{\widetilde{\Sigma}} V$ a *restricted constrained Delaunay triangulation* (restricted CDT).

Whereas a restricted Delaunay triangulation (RDT) is a subcomplex of a three-dimensional Delaunay triangulation, we know no natural three-dimensional complex that has the restricted CDT as a subcomplex. One could define a Voronoi diagram over $\widetilde{X}$ and dualize it, but there is no reason to suppose the dual will be a valid polyhedral complex: the Voronoi cells that are supposed to be kept apart by portals are likely to meet near the medial axis. (The dual complex would also be difficult to compute.) The rest of this paper seeks sampling conditions that tame the extended Voronoi diagram (over $\widetilde{\Sigma}$, not $\widetilde{X}$) and its dual restricted CDT.

Now we present several useful properties of extended Voronoi diagrams and restricted CDTs, supposing that no segment is too long. The following theorem shows that the sites whose extended Voronoi cells lie in part on an extrusion Σ_s^+ must lie on the side of the cutting plane h_s strictly opposite Σ_s^+ (excepting the endpoints of s, which lie on h_s). Thus the restricted Voronoi vertices on Σ_s^+ dualize to restricted Delaunay triangles that are also on the side of h_s opposite Σ_s^+. This theorem simplifies computing the restricted CDT, because an algorithm only needs to look at sites in one halfspace when computing the portion of $\mathrm{Vor}|_{\widetilde{\Sigma}} V$ that lies on Σ_s^+. Unfortunately, the proof is five pages long; see the full-length article [20].

▶ **Theorem 1** (Cutting Plane Theorem). *Let $s \in S$ be a segment with endpoints $p, q \in V$ such that $d(p,q) \le \rho\,\mathrm{lfs}(p)$ for $\rho \le 0.47$. Consider a point $x \in \Sigma_s^+$ and a site $v \in V \setminus \{p,q\}$ such that $x \in \mathrm{Vor}|_{\widetilde{\Sigma}} v$. Then v is strictly on the side of h_s opposite Σ_s^+. (The symmetric claim holds for any $x \in \Sigma_s^-$.)*

The next theorem shows that the restricted CDT $\mathrm{Del}|_{\widetilde{\Sigma}} V$ contains every edge in S.

▶ **Theorem 2** (Constraint Theorem). *Let $s \in S$ be a segment with endpoints $p, q \in V$ such that $d(p,q) \le \rho\,\mathrm{lfs}(p)$ for $\rho \le 0.47$. Then $\mathrm{Vor}|_{\widetilde{\Sigma}} p \cap \mathrm{Vor}|_{\widetilde{\Sigma}} q \neq \emptyset$. Hence pq is an edge in $\mathrm{Del}|_{\widetilde{\Sigma}} V$.*

Moreover, the rays p_s^+ and p_s^- lie in the interior of $\mathrm{Vor}|_{\widetilde{\Sigma}} p$ ("interior" with respect to the space $\widetilde{\Sigma}$), and neither ray intersects any other extended restricted Voronoi cell. Likewise, q_s^+ and q_s^- lie in the interior of $\mathrm{Vor}|_{\widetilde{\Sigma}} q$, and neither ray intersects another cell.

Proof. We will show that $\mathrm{Vor}|_{\widetilde{\Sigma}} p$ meets $\mathrm{Vor}|_{\widetilde{\Sigma}} q$ on the extrusion Σ_s^+, as Figures 3 and 6 show. (The same is true on Σ_s^-.) Let Π be the plane orthogonally bisecting s. Consider the point $z = \Pi \cap \zeta_s$ on the portal curve and the ray $z_s^+ = \Pi \cap \Sigma_s^+$, whose origin is z. Let x be a point on z_s^+. Note that z is the point closest to x on the portal plane h_s, and xz is perpendicular to zp. We will show that for all $x \in z_s^+$ sufficiently far from z, $x \in \mathrm{Vor}|_{\widetilde{\Sigma}} p \cap \mathrm{Vor}|_{\widetilde{\Sigma}} q$.

Theorem 1 states that for every site $v \in V \setminus \{p, q\}$ whose extended Voronoi cell $\mathrm{Vor}|_{\widetilde{\Sigma}} v$ intersects Σ_s^+, v is strictly on the side of h_s opposite Σ_s^+. Therefore, there exists some $\delta > 0$ such that $d(x,v) \ge d(x,z) + \delta$ for *every* such site v. Consider any point $x \in z_s^+$ such that $d(x,z) \ge d(z,p)^2/(2\delta)$. By Pythagoras' Theorem, for every site v whose cell intersects Σ_s^+,

$$d(x,p)^2 = d(x,z)^2 + d(z,p)^2 \le d(x,z)^2 + 2\delta\, d(x,z) < (d(x,z) + \delta)^2 \le d(x,v)^2.$$

Hence $d(x,q) = d(x,p) < \widehat{d}(x,v)$ for every site $v \in V \setminus \{p,q\}$. As x is visible from p and q, $x \in \mathrm{Vor}|_{\widetilde{\Sigma}} p$ and $x \in \mathrm{Vor}|_{\widetilde{\Sigma}} q$. Hence $\mathrm{Vor}|_{\widetilde{\Sigma}} p \cap \mathrm{Vor}|_{\widetilde{\Sigma}} q \neq \emptyset$.

To prove the final claim, consider a point $x \in p_s^+$. Observe that p is the point nearest x on h_s and $d(x,p) < d(x,q)$. Repeating the reasoning above, there exists some $\delta > 0$ such that $d(x,v) \ge d(x,p) + \delta$ for every site $v \in V \setminus \{p\}$ such that $\mathrm{Vor}|_{\widetilde{\Sigma}} v$ intersects Σ_s^+. Therefore, there is an open neighborhood $N \subset \widetilde{\Sigma}$ of x such that $N \subset \mathrm{Vor}|_{\widetilde{\Sigma}} p$ and N intersects no other cell. The same reasoning applies to points on p_s^-, q_s^+, and q_s^-. Hence p_s^+ and p_s^- lie in the interior of $\mathrm{Vor}|_{\widetilde{\Sigma}} p$ and do not intersect another extended Voronoi cell. ◀

The shape of our extrusions Σ_s^+ and Σ_s^- is motivated in part by Theorem 2, which justifies the word "constrained" in "restricted constrained Delaunay triangulation."

The following theorem shows that the sites whose extended Voronoi cells lie in part on an extrusion Σ_s^+ must lie in a ball (of modest radius) centered on the midpoint of the segment s. This helps us to efficiently compute the restricted CDT, because the portion of $\mathrm{Vor}|_{\widetilde{\Sigma}} V$ that lies on Σ_s^+ or Σ_s^- depends only on sites near s.

▶ **Theorem 3** (Possession Theorem). *Let $s \in S$ be a segment with endpoints $p, q \in V$ such that $d(p,q) \le \rho\,\mathrm{lfs}(p)$ for $\rho \le 0.47$. Let c be the midpoint of s. Let $v \in V$ be a site whose extended Voronoi cell $\mathrm{Vor}|_{\widetilde{\Sigma}} v$ contains a point $x \in \Sigma_s^+$ or $x \in \Sigma_s^-$. Then v lies in the ball $B(c, \lambda\,\mathrm{lfs}(p))$ with center c and radius $\lambda\,\mathrm{lfs}(p)$, where*

$$\lambda = \sqrt{1-2\rho}\left(1 - \sqrt{1 - \frac{\rho^2}{4\,(1-2\rho)}}\right) + \sqrt{(2-4\rho)\left(1 - \sqrt{1 - \frac{\rho^2}{4\,(1-2\rho)}}\right)}.$$

For the limiting case $\rho = 0.47$, $\lambda \doteq 0.4694$; $B(c, \lambda\,\mathrm{lfs}(p))$ has almost twice the radius of s.

4 Extended Voronoi cell boundaries

There are only two phenomena that can determine the boundary of an extended Voronoi cell. (1) Portions of a cell's boundary may be determined by hyperplanes, each hyperplane

being equidistant from two sites. For example, a point on the boundaries of two cells $\mathrm{Vor}|_{\widetilde{\Sigma}} v$ and $\mathrm{Vor}|_{\widetilde{\Sigma}} w$ might lie on the hyperplane that orthogonally bisects the line segment vw. (2) Portions of a cell's boundary may be determined by a shadow cast by a portal. For example, consider a point $x \in \mathrm{Vor}|_{\widetilde{\Sigma}} v$ such that the line segment xv intersects the boundary of a portal P_s. Portal boundaries do not block visibility; hence the set of points on $\widetilde{\Sigma}$ visible from v is closed. But an infinitesimal perturbation of x might cause x to be no longer visible from v. If x is in the principal branch, this may happen because the perturbed xv intersects the relative interior of P_s; if x is in a secondary branch, it may happen because the perturbed xv no longer intersects P_s. We say that a portal *casts a shadow at* x if $x \in \mathrm{Vor}|_{\widetilde{\Sigma}} v$ lies on the boundary of the points on $\widetilde{\Sigma}$ visible from v. A Voronoi cell $\mathrm{Vor}|_{\widetilde{\Sigma}} w$ is not necessarily closed, because its boundary might contain a shadow point $x \in \mathrm{Vor}|_{\widetilde{\Sigma}} v$ such that $\widehat{d}(v,x) < \widehat{d}(w,x)$.

The following theorem states sampling conditions under which the second phenomenon cannot happen, so the boundaries of all the extended Voronoi cells are determined solely by bisecting hyperplanes, all the extended Voronoi cells are closed point sets, and every point on $\widetilde{\Sigma}$ is in an extended Voronoi cell. For a site $v \in V$, v's *principal Voronoi cell* $\mathrm{Vor}|_{\overline{\Sigma}_S} v = \overline{\Sigma}_S \cap \mathrm{Vor}|_{\widetilde{\Sigma}} v$ is the subset of v's extended Voronoi cell in the principal branch, including the portal curves but excluding the remainder of the extrusions.

▶ **Theorem 4** (Shadow Theorem)**.** *Let S be a set of segments (with endpoints in V) such that for every segment $s = pq \in S$, $d(p,q) \le 0.47\,\mathrm{lfs}(p)$. Suppose that for every site $v \in V$ and every point x in the principal Voronoi cell $\mathrm{Vor}|_{\overline{\Sigma}_S} v$, $d(v,x) \le \max\{\mathrm{lfs}(v), \mathrm{lfs}(x)\}$. (This last condition holds if V is a constrained ϵ-sample, as defined in Section 5, for some $\epsilon \le 1$.)*

Then for every site $v \in V$ and every point x in the extended Voronoi cell $\mathrm{Vor}|_{\widetilde{\Sigma}} v$, the relative interior of the line segment xv does not intersect the boundary of a portal.

▶ **Corollary 5.** *Under the conditions of Theorem 4, every extended Voronoi cell is a closed point set (closed with respect to the topological space $\widetilde{\Sigma}$ or $\widetilde{X}$).*

▶ **Corollary 6.** *Under the conditions of Theorem 4, for every site $v \in V$ and every point x on the boundary of $\mathrm{Vor}|_{\widetilde{\Sigma}} v$, there is a site $w \in V \setminus \{v\}$ such that $x \in \mathrm{Vor}|_{\widetilde{\Sigma}} w$ and $d(v,x) = d(w,x)$.*

▶ **Corollary 7.** *Under the conditions of Theorem 4, if every connected component of Σ contains at least one site in V, then every point on $\widetilde{\Sigma}$ is in at least one extended Voronoi cell.*

The proofs of the Shadow Theorem and its corollaries appear in the full-length article [20].

5 Topological guarantees

Here we introduce conditions under which a restricted CDT is homeomorphic to the surface Σ, with a view toward applications in guaranteed-quality surface mesh generation. The *nearest point map* ν (nu) maps a point $x \in \mathbb{R}^3 \setminus M$ to the point $\nu(x)$ nearest x on Σ. We show that the nearest point map (with its domain restricted to $|\mathrm{Del}|_{\widetilde{\Sigma}} V|$) is a homeomorphism from the underlying space of the restricted CDT $\mathrm{Del}|_{\widetilde{\Sigma}} V$ to the surface Σ.

Our proof has three conditions: *a segment length condition*, that each segment $s \in S$ with endpoints p and q satisfies $d(p,q) \le 0.3647\,\mathrm{lfs}(p)$; a *sampling condition* requiring the sites V to be sufficiently dense; and an *encroachment condition* that prevents vertices in V from being too close to a segment, to prevent the possibility of triangles with excessively large circumcircles. We build on a long line of theoretical work for proving that certain triangulations are topologically correct approximations of surfaces [1, 2, 5, 6, 7, 9, 13, 15, 17], developed to support provably good algorithms for surface reconstruction and mesh generation. Many RDT-based surface mesh generation algorithms enforce a sampling condition by inserting new

vertices on Σ [7, 8, 12, 13, 25], and some support constraining segments by inserting additional vertices that subdivide segments until the RDT naturally respects them [10, 11, 13, 16, 27]. Our three conditions can likewise be enforced by inserting new vertices, but restricted CDTs will often reduce the number of new vertices needed on the segments.

To understand the sampling condition, consider a surface $\Sigma \subset \mathbb{R}^3$ without boundary, a set of segments S with their endpoints on Σ, and a set of portal curves $Z = \{\zeta_s : s \in S\}$. Recall the principal surface $\overline{\Sigma}_S$, defined in Section 2 to be the completion of $\Sigma - \bigcup_{s \in S} \operatorname{relint} \zeta_s$. We say that a finite vertex set $V \subset \Sigma$ is a *constrained ϵ-sample of* (Σ, S, Z) if V contains every endpoint of every $s \in S$ and for every point $x \in \overline{\Sigma}_S$, there is a site $v \in V$ such that $\widehat{d}(x, v) \leq \epsilon \operatorname{lfs}(x)$. That is, the ball with center x and radius $\epsilon \operatorname{lfs}(x)$ contains at least one site visible from x. Here, visibility and $\widehat{d}$ are as defined in Section 2.2; they are what differentiates a constrained ϵ-sample from a standard ϵ-sample. (If S is empty, the two are the same.) Our homeomorphism proof requires that V be a constrained 0.3202-sample of (Σ, S, Z).

The encroachment condition applies only to restricted Delaunay triangles whose dual faces intersect an extrusion, as in Figure 6. (A triangle's dual face is usually a single point, called an extended Voronoi vertex, but our homeomorphism proof does not depend on it.) Let τ be such a triangle. The *circumradius* r of τ is the radius of the unique circle that passes through τ's three vertices. Let w be the vertex of τ at τ's largest plane angle. We require that $r \leq 0.3606 \operatorname{lfs}(w)$. The purpose of this restriction is to prevent the existence of "inverted" triangles in $\operatorname{Del}|_{\widetilde{\Sigma}} V$, which create *foldovers*, points where the nearest point map ν from $|\operatorname{Del}|_{\widetilde{\Sigma}} V|$ to Σ is not locally injective (hence ν is not a homeomorphism).

To put the encroachment condition into perspective, suppose that Σ is a sphere and consider a segment $s \in S$ having the maximum safe length of 0.3647 times the sphere's radius. A triangle τ whose dual vertex lies on Σ_s^+ can exceed the safe radius of 0.3606 times the sphere's radius only if τ has an angle greater than 149.62°. If s is shorter, this angle is larger: in the limit as the segment lengths approach zero (or the radius of Σ approaches infinity), the encroachment condition falls away and restricted CDTs on Σ behave like CDTs in the plane. By contrast, in standard approaches using an RDT that conforms to the segments, no triangle with edge s can have an angle opposite s greater than 90°.

The sampling and encroachment conditions both rule out triangles with circumradii that are excessively large relative to the local feature size. A large circumradius implies either that the triangle is large, or that it has a large plane angle (close to 180°). Imposing these conditions is consistent with a mesh generator's goal of producing only well-shaped triangles, so our conditions are not onerous. Nevertheless, there are other applications such as surface reconstruction where the encroachment condition is not a natural condition. The restricted CDT may nevertheless still be useful in that context; see the Conclusions for speculations.

Our main topological result is the next theorem. Unfortunately, the proof is over twenty pages long; see the full-length article [20]. To keep the proof from being even longer, we assume that no point on $\widetilde{\Sigma}$ is equidistant from four visible vertices in V, which implies that no point is in more than three cells. (This assumption is not actually necessary.)

▶ **Theorem 8** (Homeomorphism). *Let V be a constrained ϵ-sample of (Σ, S, Z) for some $\epsilon \leq 0.3202$. Suppose that for every segment $pq \in S$, $d(p, q) \leq 0.3647 \operatorname{lfs}(p)$. Suppose that no portal curve in Z has a relative interior that intersects another portal curve in Z or a site in V. Suppose that no point on $\widetilde{\Sigma}$ is equidistant from four visible vertices in V. Suppose that for every restricted Delaunay triangle τ whose dual extended Voronoi face intersects an extrusion, τ satisfies $r \leq 0.3606 \operatorname{lfs}(w)$, where r is τ's circumradius and w is the vertex of τ at τ's largest plane angle. Then the nearest point map $\nu : |\operatorname{Del}|_{\widetilde{\Sigma}} V| \to \Sigma$ is a homeomorphism.*

The proof is related to proofs by Amenta et al. [3] and Boissonnat and Ghosh [6] that also use the nearest point map. We sketch a few ideas. To make every extended Voronoi cell become a topological disk and to clarify the duality between the extended Voronoi diagram and the restricted CDT, it is convenient to define a compact 2-manifold without boundary $\mathring{\widetilde{\Sigma}}$, obtained from $\widetilde{\Sigma}$ by gluing each pair Σ_s^+ and Σ_s^- together along their boundaries as illustrated in Figure 7. For each segment $s = pq$, we topologically identify the ray p_s^+ with the ray p_s^-, and q_s^+ with q_s^-. (Theorem 2 shows that p_s^+ and p_s^- are subsets of p's extended Voronoi cell, so this gluing does not confuse which points are in which Voronoi cell.) We create a single point s^∞ "at infinity" (one such point per segment s) at the end of every ray x_s^+ and x_s^- for all $x \in \zeta_s$, thereby making $\mathring{\widetilde{\Sigma}}$ compact. Thus the hole created in $\overline{\Sigma}_S$ by cutting a slit at ζ_s is filled in $\mathring{\widetilde{\Sigma}}$ with a topological disk $\Sigma_s^+ \cup \Sigma_s^- \cup s^\infty$, as illustrated. Clearly, $\mathring{\widetilde{\Sigma}}$ is homeomorphic to Σ. (Note that unlike $\widetilde{\Sigma}$, $\mathring{\widetilde{\Sigma}}$ is not embedded in $\widetilde{X}$ and has neither coordinates nor distances.)

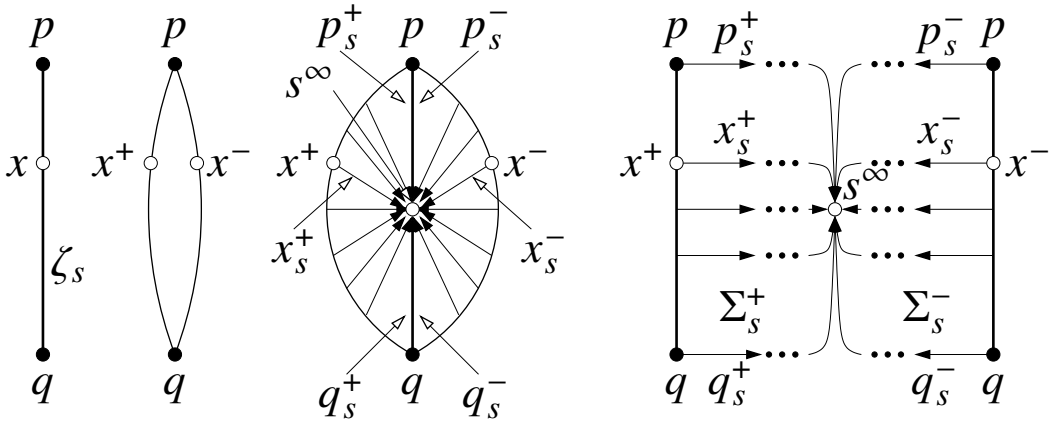

Figure 7 After we remove relint ζ_s from Σ, we glue two extrusions Σ_s^+ and Σ_s^- in the hole to create $\widetilde{\Sigma}$. Additional gluing can transform $\widetilde{\Sigma}$ into a compact 2-manifold without boundary $\mathring{\widetilde{\Sigma}}$, restoring the topology of Σ, by gluing the ray p_s^+ to p_s^-, gluing q_s^+ to q_s^-, and filling the hole with a point s^∞.

If V is a constrained 0.44-sample and S satisfies the segment length condition, we show that every extended Voronoi cell on $\mathring{\widetilde{\Sigma}}$ is homeomorphic to a closed disk. With the assumption that no point is in more than three cells, this implies that the adjacency graph of $\mathrm{Vor}|_{\widetilde{\Sigma}} V$, which is the graph of $\mathrm{Del}|_{\widetilde{\Sigma}} V$, can be drawn on $\mathring{\widetilde{\Sigma}}$ (and therefore on $\widetilde{\Sigma}$) with no crossings.

We call an extended Voronoi vertex a *principal vertex* if it lies in the principal branch (on $\overline{\Sigma}_S$), or a *secondary vertex* if it lies on an extrusion but not on a portal curve. We show that if V is a constrained 0.3202-sample, each principal vertex dualizes to a triangle whose circumradius is not large (relative to the local feature size). The encroachment condition implies that each secondary vertex dualizes to a triangle whose circumradius is not large.

The bounds on circumradii allow us to prove that the nearest point map restricted to any single restricted Delaunay triangle is a homeomorphism. Moreover, there is a sense in which the map preserves orientation: for any extended Voronoi vertex u whose dual extended Delaunay triangle is $\tau = \triangle pp'p''$, the sites p, p', and p'' are in counterclockwise order around $v(\tau)$ if and only if the cells $\mathrm{Vor}|_{\widetilde{\Sigma}} p$, $\mathrm{Vor}|_{\widetilde{\Sigma}} p'$, and $\mathrm{Vor}|_{\widetilde{\Sigma}} p''$ adjoin u in counterclockwise order around u (as seen from outside Σ). From this, we argue that along each of its edges, each restricted Delaunay triangle adjoins another restricted Delaunay triangle with a consistent orientation, and therefore the triangles must cover the whole surface Σ—that is, the nearest point map is a surjection from $|\mathrm{Del}|_{\widetilde{\Sigma}} V|$ to Σ. As the boundary of a Voronoi cell $\mathrm{Vor}|_{\widetilde{\Sigma}} v$ is a simple loop, the restricted Delaunay triangles adjoining v form a fan of triangles around v

whose union is a topological disk. From these facts we prove that the nearest point map is an injection from $|\mathrm{Del}|_{\widetilde{\Sigma}} V|$ to Σ (there are no foldovers) and therefore a homeomorphism.

6 Conclusions

The restricted constrained Delaunay triangulation is a rigorous generalization of the constrained Delaunay triangulation to surfaces. Under suitable conditions on the vertex density and the segment lengths, the restricted CDT is homeomorphic to Σ and contains every constraining segment. We believe that the restricted CDT will become a useful tool for enforcing specified boundaries in guaranteed-quality algorithms for surface meshing. But first and foremost, we think the existence of restricted CDTs is a beautiful mathematical fact.

Several algorithms suggest themselves for computing the restricted CDT. The classical gift-wrapping algorithm [13, Section 3.11] [29] can be adapted. Another approach, likely faster in practice, is to start with the RDT and then incrementally insert the segments one by one [31]. It is an open problem to design an algorithm that runs in $O(|V|^2)$ time or better.

Another open problem is to design a guaranteed-quality algorithm that uses the restricted CDT to mesh surfaces with prescribed boundaries. The algorithm must generate new vertices on Σ with the goal of enforcing the sampling and encroachment conditions, in addition to the customary goal of constructing high-quality triangles. As we have said, we believe that restricted CDTs will require fewer triangles and vertices than algorithms based on RDTs.

Although our encroachment condition is reasonable in a surface mesh generator, it is undesirable in some applications such as surface reconstruction. Unfortunately, without this condition, we cannot prove a homeomorphism because the nearest point map is not necessarily injective. Figure 8 illustrates the problem. Suppose that we place two segments s and s' close together and we place a vertex r very close to the midpoint of $\zeta_{s'}$ (violating the encroachment condition), as shown in Figure 8. Consider the triangle formed by r and s', and its dual 3D Voronoi edge e; e can be arbitrarily close to perpendicular to $P_{s'}$. Then e may enter both portals P_s and $P_{s'}$ and generate two extended Voronoi vertices (illustrated in red) where it intersects $\Sigma^+_{s'}$ (which is desirable) and Σ^+_s (which is not). This circumstance is possible because the segments are close together, the portals are tilted toward each other, and Σ^+_s is extruded infinitely far. Increasing the sampling density does not fix this problem.

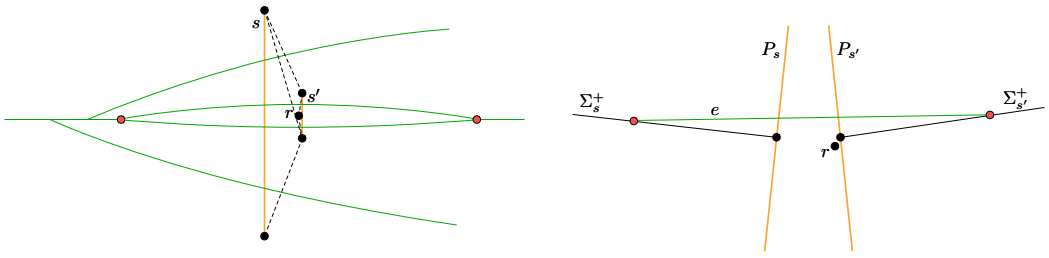

Figure 8 The left figure shows a top view of a Voronoi diagram drawn on $\widetilde{\Sigma}$; the right figure shows a side view. The segments s and s' are placed close together on Σ and their portals P_s and $P_{s'}$ are tilted toward each other in the side view. If r is arbitrarily close to $P_{s'}$, the Voronoi edge e dual to $\triangle rs'$ is tilted nearly tangent to the surface and can leave $P_{s'}$, enter P_s, and intersect Σ^+_s (perhaps far down the extrusion). The dual triangulation contains a dangling triangle $\triangle rs'$.

If we drop the encroachment condition (but retain the other two conditions), we conjecture that the restricted Delaunay triangles still form a watertight enclosure such that the nearest

point map is a surjection from the restricted CDT to Σ. However, it is not necessarily injective; there may be foldovers where sites brush up against segments. There may also be "dangling" triangles, connected to the remainder of the triangulation by only a single edge; an example is the triangle formed by r and s' in Figure 8. Such triangles are easily pruned.

References

1 Nina Amenta and Marshall Bern. Surface Reconstruction by Voronoi Filtering. *Discrete & Computational Geometry*, 22(4):481–504, June 1999.

2 Nina Amenta, Marshall W. Bern, and David Eppstein. The Crust and the β-Skeleton: Combinatorial Curve Reconstruction. *Graphical Models and Image Processing*, 60(2):125–135, March 1998.

3 Nina Amenta, Sunghee Choi, Tamal K. Dey, and Naveen Leekha. A Simple Algorithm for Homeomorphic Surface Reconstruction. *International Journal of Computational Geometry and Applications*, 12(1–2):125–141, 2002.

4 Dominique Attali, Herbert Edelsbrunner, and Yuriy Mileyko. Weak Witnesses for Delaunay Triangulations of Submanifolds. In *Proceedings of the 2007 ACM Symposium on Solid and Physical Modeling*, pages 143–150, Beijing, China, June 2007. ACM.

5 Jean-Daniel Boissonnat, Frédéric Chazal, and Mariette Yvinec. *Geometric and Topological Inference*, volume 57. Cambridge University Press, September 2018.

6 Jean-Daniel Boissonnat and Arijit Ghosh. Manifold Reconstruction Using Tangential Delaunay Complexes. *Discrete & Computational Geometry*, 51(1):221–267, January 2014.

7 Jean-Daniel Boissonnat and Steve Oudot. Provably Good Surface Sampling and Approximation. In *Symposium on Geometry Processing*, pages 9–18. Eurographics Association, June 2003.

8 Jean-Daniel Boissonnat and Steve Oudot. Provably Good Sampling and Meshing of Surfaces. *Graphical Models*, 67(5):405–451, September 2005.

9 Ho-Lun Cheng, Tamal K. Dey, Herbert Edelsbrunner, and John Sullivan. Dynamic Skin Triangulation. *Discrete & Computational Geometry*, 25(4):525–568, December 2001.

10 Siu-Wing Cheng, Tamal K. Dey, and Joshua A. Levine. A Practical Delaunay Meshing Algorithm for a Large Class of Domains. In *Proceedings of the 16th International Meshing Roundtable*, pages 477–494, Seattle, Washington, October 2007.

11 Siu-Wing Cheng, Tamal K. Dey, and Edgar A. Ramos. Delaunay Refinement for Piecewise Smooth Complexes. *Discrete & Computational Geometry*, 43(1):121–166, 2010.

12 Siu-Wing Cheng, Tamal K. Dey, Edgar A. Ramos, and Tathagata Ray. Sampling and Meshing a Surface with Guaranteed Topology and Geometry. *SIAM Journal on Computing*, 37(4):1199–1227, 2007.

13 Siu-Wing Cheng, Tamal Krishna Dey, and Jonathan Richard Shewchuk. *Delaunay Mesh Generation*. CRC Press, Boca Raton, Florida, December 2012.

14 L. Paul Chew. Constrained Delaunay Triangulations. *Algorithmica*, 4(1):97–108, 1989.

15 Tamal K. Dey. *Curve and Surface Reconstruction: Algorithms with Mathematical Analysis*. Cambridge University Press, New York, 2007.

16 Tamal K. Dey and Joshua A. Levine. Delaunay Meshing of Piecewise Smooth Complexes without Expensive Predicates. *Algorithms*, 2(4):1327–1349, 2009.

17 Herbert Edelsbrunner and Nimish R. Shah. Triangulating Topological Spaces. *International Journal of Computational Geometry and Applications*, 7(4):365–378, August 1997.

18 Ronald H. Hardin, Neil J. A. Sloane, and Warren D. Smith. Coverings by Points on a Sphere. Web page `http://neilsloane.com/coverings/`, February 1994.

19 Marc Khoury and Jonathan Richard Shewchuk. Approximation Bounds for Interpolation and Normals on Triangulated Surfaces and Manifolds. *CoRR*, abs/1911.03424, November 2019. URL: `http://arxiv.org/abs/1911.03424`.

20 Marc Khoury and Jonathan Richard Shewchuk. Restricted Constrained Delaunay Triangulations. The complete version of this paper, March 2021. URL: `https://people.eecs.berkeley.edu/~jrs/papers/rcdt.pdf`.

21 Ron Kimmel and James A. Sethian. Computing Geodesic Paths on Manifolds. *Proceedings of the National Academy of Sciences*, 95(15):8431–8435, July 1998.

22 Der-Tsai Lee and Arthur K. Lin. Generalized Delaunay Triangulations for Planar Graphs. *Discrete & Computational Geometry*, 1:201–217, 1986.

23 Takashi Maekawa. Computation of Shortest Paths on Free-Form Parametric Surfaces. *Journal of Mechanical Design*, 118(4):499–508, December 1996.

24 Manish Mandad, David Cohen-Steiner, Leif Kobbelt, Pierre Alliez, and Mathieu Desbrun. Variance-Minimizing Transport Plans for Inter-Surface Mapping. *ACM Transactions on Graphics*, 36(4), July 2017.

25 Steve Oudot, Laurent Rineau, and Mariette Yvinec. Meshing Volumes Bounded by Smooth Surfaces. In *Proceedings of the 14th International Meshing Roundtable*, pages 203–219, San Diego, California, September 2005. Springer.

26 Andrew Pressley. *Elementary Differential Geometry (Second Edition)*. Springer Science & Business Media, 2010.

27 Laurent Rineau and Mariette Yvinec. Meshing 3D Domains Bounded by Piecewise Smooth Surfaces. In *Proceedings of the 16th International Meshing Roundtable*, pages 443–460, Seattle, Washington, October 2007. Springer.

28 Erich Schönhardt. Über die Zerlegung von Dreieckspolyedern in Tetraeder. *Mathematische Annalen*, 98:309–312, March 1928.

29 Raimund Seidel. Constrained Delaunay Triangulations and Voronoi Diagrams with Obstacles. In H. S. Poingratz and W. Schinnerl, editors, *1978–1988 Ten Years IIG*, pages 178–191. Institute for Information Processing, Graz University of Technology, 1988.

30 Jonathan Richard Shewchuk. Delaunay Refinement Algorithms for Triangular Mesh Generation. *Computational Geometry: Theory and Applications*, 22(1–3):21–74, May 2002.

31 Jonathan Richard Shewchuk and Brielin C. Brown. Fast Segment Insertion and Incremental Construction of Constrained Delaunay Triangulations. *Computational Geometry: Theory and Applications*, 48(8):554–574, September 2015.

A Useful Known Results

We include here some geometric lemmas that are used in this paper or the forthcoming appendices.

▶ **Lemma 9** (Feature Ball Lemma [15, 13]). *If a geometric closed d-ball B intersects a k-manifold $\Sigma \subset \mathbb{R}^d$ without boundary at more than one point and either (i) $\Sigma \cap B$ is not a topological k-ball or (ii) $\Sigma \cap \partial B$ is not a topological $(k-1)$-sphere, then B contains a medial axis point.*

▶ **Lemma 10.** *Let s be a line segment with endpoints $p, q \in \Sigma$ where $\Sigma \subset \mathbb{R}^d$ is a k-manifold without boundary. Let B_s be the* diametric ball *of s, the smallest closed d-ball such that $B_s \supset s$ (for which s is a diameter of B_s). If $d(p,q) \le \rho\,\mathrm{lfs}(p)$ for some $\rho < 1$, then $B_s \cap \Sigma$ is a topological $(k-1)$-ball.*

Proof. Let B_s be the diametric ball of s. Suppose that $B_s \cap \Sigma$ is not a topological 2-ball. As $B_s \cap \Sigma$ contains more than one point (p and q), by the Feature Ball Lemma (Lemma 9), there exists a medial axis point $m \in B_s$. This implies that $\mathrm{lfs}(p) \le d(p,m) \le d(p,q) \le \rho\,\mathrm{lfs}(p)$, which is a contradiction. ◀

▶ **Lemma 11** (Triangle Normal Lemma [19]). *Let $\Sigma \subset \mathbb{R}^3$ be a smooth 2-manifold without boundary. Let τ be a triangle whose vertices lie on Σ. Let r be τ's circumradius. Let v be a vertex of τ and let ϕ be τ's plane angle at v. Then*

$$\angle(N_\tau, N_v\Sigma) = \angle(\mathrm{aff}\,\tau, T_v\Sigma) \le \arcsin\left(\frac{R}{\mathrm{lfs}(v)} \max\left\{\cot\frac{\phi}{2}, 1\right\}\right).$$

(Note that the argument $\cot\frac{\phi}{2}$ dominates if ϕ is acute and the argument 1 dominates if ϕ is obtuse.) In particular, if v is the vertex at τ's largest plane angle (so $\phi \ge 60°$) and $r < \mathrm{lfs}(v)/\sqrt{3} \doteq 0.577\,\mathrm{lfs}(v)$, then

$$\angle(N_\tau, N_v\Sigma) = \angle(\mathrm{aff}\,\tau, T_v\Sigma) \le \arcsin\frac{\sqrt{3}R}{\mathrm{lfs}(v)}.$$

▶ **Lemma 12** (Feature Translation Lemma [3, 15]). *Let $\Sigma \subset \mathbb{R}^3$ be a smooth surface and let $p, q \in \Sigma$ be points on Σ such that $d(p,q) \le \epsilon\,\mathrm{lfs}(p)$ for some $\epsilon < 1$. Then*

$$\mathrm{lfs}(p) \le \frac{1}{1-\epsilon}\mathrm{lfs}(q) \quad \text{and} \quad d(p,q) \le \frac{\epsilon}{1-\epsilon}\mathrm{lfs}(q).$$

Proof. By the definition of the local feature size, there is a medial axis point m such that $|qm| = \mathrm{lfs}(q)$. By the Triangle Inequality, $\mathrm{lfs}(p) \le |pm| \le |pq| + |qm| \le \epsilon\,\mathrm{lfs}(p) + \mathrm{lfs}(q)$. Rearranging terms gives $\mathrm{lfs}(p) \le \mathrm{lfs}(q)/(1-\epsilon)$. The second claim follows immediately. ◀

▶ **Lemma 13** (Normal Variation Lemma [19]). *Let $\Sigma \subset \mathbb{R}^3$ be a bounded, smooth 2-manifold without boundary. Consider two points $p, q \in \Sigma$ and let $\delta = d(p,q)/\mathrm{lfs}(p)$. Let n_p and n_q be outside-facing vectors normal to Σ at p and q, respectively. If $\delta < \sqrt{4\sqrt{5}-8} \doteq 0.9717$, then $\angle(n_p, n_q) \le \eta(\delta)$ where*

$$\eta(\delta) = \arccos\left(1 - \frac{\delta^2}{2\sqrt{1-\delta^2}}\right) \approx \delta + \frac{7}{24}\delta^3 + O(\delta^5).$$

B The Geometry of Portal Curves

In this section we derive some facts about the geometry of portal curves that we use to prove this paper's main results, including Theorems 1, 3, and 8.

B.1 Some Facts About Plane Geometry and Curvature

Let s be a line segment in the plane with endpoints p and q, let C be the circle whose diameter is s, let c be the center of C (the midpoint of s), and let v be a vector normal to s. The bisector of s is the line $l = \{c + \lambda v : \lambda \in \mathbb{R}\}$. Let $\mathcal{P}(s)$ denote the pencil of circles that pass through both p and q. All these circles have their centers on l.

Without loss of generality, we choose a coordinate system so that c (the midpoint of s/center of C) is the origin and s lies on the x-axis. C is the unique smallest circle in the pencil $\mathcal{P}(s)$, and its radius is $d(p,q)/2$. For every value $r > d(p,q)/2$, there are two circles $C_r, C_{-r} \in \mathcal{P}(s)$ with radius r; here we define C_r to have its center above s (i.e., with positive y-coordinate) and C_{-r} to have its center below s, as illustrated in Figure 9.

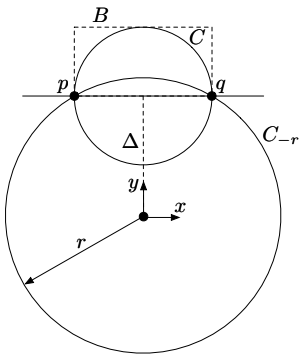

Figure 9 The circles C and C_{-r} are members of the pencil of circles $\mathcal{P}(pq)$ induced by the segment pq. We define a coordinate system whose origin is the center of C_{-r} and whose x-axis is parallel to pq.

▶ **Lemma 14.** *Let $p, q \in \mathbb{R}^2$ be two points and let C be the circle with diameter pq. Let $C_{-1/\kappa} \in \mathcal{P}(pq)$ be a circle with radius $1/\kappa$ (curvature κ). Then the distance from the north pole of $C_{-1/\kappa}$ to the line segment pq is*

$$\frac{1 - \sqrt{1 - \kappa^2\, d(p,q)^2/4}}{\kappa}.$$

By symmetry, the same bound holds for the south pole of $C_{1/\kappa}$.

Proof. Let $c = (0, \Delta)$ be the center of $C_{-1/\kappa}$. As $C_{-1/\kappa}$ passes through p and q, its radius is $1/\kappa = d(c,p) = d(c,q) = \sqrt{d(p,q)^2/4 + \Delta^2}$. Solving for Δ gives

$$\Delta = \pm\frac{\sqrt{1 - \kappa^2\, d(p,q)^2/4}}{\kappa}.$$

The negative solution is the relevant one. (The positive solution provides the symmetric result for the south pole of $C_{1/\kappa}$.) The distance between the north pole of $C_{-1/\kappa}$ and the line

segment pq is

$$\left|\Delta + \frac{1}{\kappa} - 0\right| = \left|\frac{1 - \sqrt{1 - \kappa^2\, d(p,q)^2/4}}{\kappa}\right| = \frac{1 - \sqrt{1 - \kappa^2\, d(p,q)^2/4}}{\kappa}.$$

◀

Before moving on we prove a supporting result.

▶ **Lemma 15.** *Let $x \in \mathbb{R}^2$. Define two balls $B_1 = B(c_1, R), B_2 = B(c_2, R)$ with centers $c_1, c_2 \in \mathbb{R}^2$ and radii R. Suppose that B_1, B_2 are tangent to x with c_1xc_2 collinear. Let $B = B(x, r)$ be a ball centered at x of radius r with $r < R$. Let $\gamma : [0,1] \to \mathbb{R}^2$ be a regular curve in $\mathbb{R}^2$ with $\gamma(0) = p$, $\dot\gamma(0) \perp c_1pc_2$, and with curvature everywhere at most $\kappa = \frac{1}{R}$. Then γ must exit B before entering the interior of $B_1 \cup B_2$.*

Proof. Imagine moving along γ starting at x in the direction $\dot\gamma(0)$ toward some point y. The rate at which γ can curve away from the line in the direction $\dot\gamma(0)$ is bounded by κ. Specifically, by Property II in [4], the angle $\angle(\dot\gamma(0), \dot\gamma(t)) \le \kappa\ell$, where ℓ is the distance along γ from $\gamma(0)$ to $\gamma(t)$. The maximum distance that can be traveled before exiting B is given by moving along arcs defined by the intersection of B with the boundaries $\partial B_1, \partial B_2$. These arc lengths are $2R \arcsin \frac{r}{2R}$, which gives

$$\angle(\dot\gamma(0), \dot\gamma(t)) \le \kappa\ell \le 2 \arcsin \frac{r}{2R} \le 2 \arcsin \frac{1}{2} < \frac{\pi}{2}.$$

Thus γ cannot remain in B indefinitely, since its tangent vector always has some component in the direction $\dot\gamma(0)$.

Suppose, for the sake of contradiction, that γ enters $B_1 \cup B_2$ at the point $\gamma(t)$ with tangent vector v. At this point $\angle(\dot\gamma(0), v) < \angle(\dot\gamma(0), \dot\gamma(t))$. This is impossible since the angle is proportional to the distance from x. ◀

Given a circle C, let $B(C)$ denote the closed disk whose boundary is C. The following lemma states that if the curvature of a portal curve ζ_s is bounded, then $\zeta_s \subset B(C_{1/\kappa}) \cap B(C_{-1/\kappa})$. Intuitively, ζ_s cannot stray far from the line segment $s = pq$. This result allows us to bound the distance between ζ_s and s on a portal P_s.

▶ **Lemma 16.** *Let $p, q \in \mathbb{R}^2$ be two points and let C be the circle with diameter pq. Let $\gamma : [0,1] \to \mathbb{R}^2$ be a regular curve in $\mathbb{R}^2$ with $\gamma(0) = p$ and $\gamma(1) = q$ and with curvature everywhere at most κ. Suppose that $\gamma \subset B(C)$ and that $|pq| < \frac{2}{\kappa}$. Then $\gamma \subset B(C_{1/\kappa}) \cap B(C_{-1/\kappa})$.*

Proof. First we show that $\gamma \subset B(C_{-1/\kappa})$; suppose for the sake of contradiction that $\gamma \not\subset B(C_{-1/\kappa})$. Let $\gamma(t')$ be a point on γ that maximizes the distance from $\gamma(t')$ to the center of $C_{-1/\kappa}$; then $\gamma(t') \notin B(C_{-1/\kappa})$. Let C' be a circle of radius $1/\kappa$ tangent to γ at $\gamma(t')$ with the center of C' lying on the line segment connecting $\gamma(t')$ to the center of $C_{-1/\kappa}$. As C' has the same radius as $C_{-1/\kappa}$ (which is greater than the radius of C) but passes through a point in $C \setminus C_{-1/\kappa}$, the circle C' must enclose either p or q (possibly both). Now apply Lemma 15, with C' as B_2 and a ball centered at $\gamma(t')$ with diameter $|pq|$ as B. Then γ cannot return to either p or q (which ever is contained in C') without first leaving B. Since the radius of B is equal to the radius of $B(C)$, this implies that γ must leave $B(C)$, a contradiction. It follows that $\gamma \subset B(C_{-1/\kappa})$. By a symmetrical argument, $\gamma \subset B(C_{1/\kappa})$. ◀

For any point in $B(C_{1/\kappa}) \cap B(C_{-1/\kappa})$, we can bound the distance between that point and the nearest endpoint of pq. The bound is at its worst on the boundary of $B(C_{1/\kappa}) \cap B(C_{-1/\kappa})$

so we compute the maximum distance to one of the endpoints for every point along the relevant arcs of $C_{1/\kappa}, C_{-1/\kappa}$.

For the next lemma, we find it convenient to define a coordinate system different from the one used in the previous two lemmas. Consider a coordinate system where the center of C_{-r} is the origin, the center of C lies on the positive y-axis, and the segment pq is parallel with the x-axis, as in Figure 9. In this coordinate system, $p = (-d(p,q)/2, \Delta)$ and $q = (d(p,q)/2, \Delta)$.

▶ **Lemma 17.** *Let $p, q \in \mathbb{R}^2$ be two points and let C denote the circle with diameter pq. Parameterize C as*

$$C(\theta) = \left(-\frac{d(p,q)}{2}\cos\theta, \Delta + \frac{d(p,q)}{2}\sin\theta\right).$$

For each value of $\theta \in [0,\pi]$, the line segment from the center of C to $C(\theta)$ intersects a unique point on C_{-r}. The distance from this point on C_{-r} to p or q (whichever is closest) can be expressed in terms of θ as

$$\min\{d(C_{-r}(\theta),p), d(C_{-r}(\theta),q)\} = \begin{cases} [r^2 + \frac{d(p,q)^2}{4} - \frac{d(p,q)}{2}\sqrt{r^2\sec^2\theta - \Delta^2} \\ \quad - \cos 2\theta(\Delta^2 + \frac{d(p,q)}{2}\sqrt{r^2\sec^2\theta - \Delta^2}) \\ \quad -\Delta \sin 2\theta(-\frac{d(p,q)}{2} + \sqrt{r^2\sec^2\theta - \Delta^2}]^{1/2} & \theta \in [0,\pi/2), \\ \sqrt{r^2 + \Delta^2 - 2r\Delta + \frac{d(p,q)^2}{4}} & \theta = \pi/2, \\ [r^2 + \frac{d(p,q)^2}{4} - \frac{d(p,q)}{2}\sqrt{r^2\sec^2\theta - \Delta^2} \\ \quad - \cos 2\theta(\Delta^2 + \frac{d(p,q)}{2}\sqrt{r^2\sec^2\theta - \Delta^2}) \\ \quad +\Delta \sin 2\theta(-\frac{d(p,q)}{2} + \sqrt{r^2\sec^2\theta - \Delta^2})]^{1/2} & \theta \in (\pi/2,\pi]. \end{cases} \tag{2}$$

A symmetric result holds for C_r.

Proof. C_{-r} can be parameterized as

$$C_{-r}(\varphi) = (-r\cos\varphi, r\sin\varphi).$$

Then the squared distance from a point on C_{-r} to p is

$$d(C_{-r}(\varphi), p)^2 = r^2 + \Delta^2 + \frac{d(p,q)^2}{4} - rd(p,q)\cos\varphi - 2r\Delta\sin\varphi.$$

Similarly, the squared distance from a point on C_{-r} to q is

$$d(C_{-r}(\varphi), q)^2 = r^2 + \Delta^2 + \frac{d(p,q)^2}{4} + rd(p,q)\cos\varphi - 2r\Delta\sin\varphi.$$

We wish to reparameterize these two distances in terms of θ; that is, for each $\theta \in [0,\pi]$ we need an expression for the corresponding value of φ. It follow from basic trigonometry that

$$\tan\theta = \frac{|r\sin\varphi - \Delta|}{|r\cos\varphi|}. \tag{3}$$

We'll consider the solutions to this equation in two steps; first for $\varphi \in [\arcsin(\Delta/r), \pi/2)$ and second for $\varphi \in (\pi/2, \pi - \arcsin(\Delta/r)]$. In both of these intervals, $r\sin\varphi - \Delta \geq 0$, so we can remove the absolute value from the numerator in Equation 3. In the case

where $\varphi \in [\arcsin(\Delta/r), \pi/2)$, $r\cos\varphi > 0$ so we can also remove the absolute values in the denominator. In this case,

$$\varphi = \arctan \frac{\Delta + \tan\theta\sqrt{r^2\sec^2\theta - \Delta^2}}{-\Delta\tan\theta + \sqrt{r^2\sec^2\theta - \Delta^2}}.$$

In the second case, $\varphi \in (\pi/2, \pi - \arcsin\frac{\Delta}{r}]$, $r\cos\varphi < 0$. Thus we replace the absolute value in the denominator with a negative sign, and

$$\varphi = \arctan \frac{\Delta - \tan\theta\sqrt{r^2\sec^2\theta - \Delta^2}}{\Delta\tan\theta + \sqrt{r^4\sec^2\theta - \Delta^2}}.$$

However, we're interested in the angle that the hypotenuse makes with the negative x-axis. Thus the correct value of φ is

$$\varphi = \pi - \arctan \frac{\Delta - \tan\theta\sqrt{r^2\sec^2\theta - \Delta^2}}{\Delta\tan\theta + \sqrt{r^4\sec^2\theta - \Delta^2}}.$$

Plugging our two solutions for φ into our two distance equations respectively, we have

$$\begin{aligned} r^2 + \frac{d(p,q)^2}{4} - \frac{d(p,q)}{2}\sqrt{r^2\sec^2\theta - \Delta^2} \\ - \cos 2\theta\left(\Delta^2 + \frac{d(p,q)}{2}\sqrt{r^2\sec^2\theta - \Delta^2}\right) \\ - \Delta\sin 2\theta\left(-\frac{d(p,q)}{2} + \sqrt{r^2\sec^2\theta - \Delta^2}\right) \end{aligned}$$

and

$$\begin{aligned} r^2 + \frac{d(p,q)^2}{4} - \frac{d(p,q)}{2}\sqrt{r^2\sec^2\theta - \Delta^2} \\ - \cos 2\theta\left(\Delta^2 + \frac{d(p,q)}{2}\sqrt{r^2\sec^2\theta - \Delta^2}\right) \\ + \Delta\sin 2\theta\left(-\frac{d(p,q)}{2} + \sqrt{r^2\sec^2\theta - \Delta^2}\right). \end{aligned}$$

The first equation holds in the range $[0, \pi/2)$, while the second holds in the range $(\pi/2, \pi]$. Taking square roots gives Equation 2.

The point corresponding to $\theta = \varphi = \pi/2$ is equidistant from both p and q. Its value can be computed by considering the left and right hand limits and confirming that they are identical. Indeed the limits are equal and have value

$$d(C_{-r}(\theta), p) = d(C_{-r}(\theta), q) = \sqrt{r^2 + \Delta^2 - 2r\Delta + \frac{d(p,q)^2}{4}}.$$

◀

B.2 Curvature Bounds

For a segment s and plane h_s, defined as in Section 2.2, we wish to bound the curvature κ of the curve ζ_s. Using the fact that ζ_s is defined as the intersection of h_s with Σ, we bound the curvature κ in terms of the maximum principal curvatures of Σ along ζ_s.

▶ **Lemma 18.** *Let $\Sigma \subset \mathbb{R}^3$ be a smooth surface and let $p, q \in \Sigma$ be points on Σ. Let $h = \text{span}\{n_p, \vec{pq}\}$ be a plane and define $b_h = \frac{n_p \times \vec{pq}}{\|n_p \times \vec{pq}\|}$, the unit vector normal to h. Define a curve $\gamma : [0,1] \to \Sigma$ as the intersection $h \cap \Sigma$ such that $\gamma(0) = p$ and $\gamma(1) = q$. Let κ be the curvature along γ. Then the geodesic curvature is $\kappa_g = \kappa\langle b_h, n\rangle$.*

Proof. Let T denote the tangent field along γ and N denote the principal normal along γ. Then the geodesic curvature is

$$\begin{aligned}\kappa_g &= \kappa\langle N, n \times T\rangle \\ &= \kappa\langle b_h \times T, n \times T\rangle \\ &= \kappa \begin{vmatrix} \langle b_h, n\rangle & \langle b_h, T\rangle \\ \langle T, n\rangle & 1 \end{vmatrix} \\ &= \kappa(\langle b_h, n\rangle - \langle b_h, T\rangle\langle T, n\rangle).\end{aligned}$$

The third identity uses the scalar quadruple product. Notice that $\langle b_h, T\rangle = 0$, since T always lies in h, and $\langle T, n\rangle = 0$, since T is also a tangent field along a curve in Σ. The result follows. ◀

▶ **Lemma 19.** *Let $\Sigma \subset \mathbb{R}^3$ be a smooth surface and let $p, q \in \Sigma$ be points on Σ. Let $h = \text{span}\{n_p, \vec{pq}\}$ be a plane and define $b_h = \frac{n_p \times \vec{pq}}{\|n_p \times \vec{pq}\|}$, the unit vector normal to h. Define a curve $\gamma : [0,1] \to \Sigma$ as the intersection $h \cap \Sigma$ such that $\gamma(0) = p$ and $\gamma(1) = q$. Suppose that, for every $t \in [0,1]$, the angle $0 \le \angle(n_p, n_{\gamma(t)}) \le \alpha < \pi/2$ in radians. Then the curvature κ of the curve γ is bounded by*

$$\frac{\max\{|\kappa_1|, |\kappa_2|\}}{\sqrt{1 - \sin^2\alpha}}$$

where κ_1, κ_2 are the principal curvatures of Σ.

Proof. The angle $\angle(n_p, n_{\gamma(t)}) \le \alpha$ for all t, which implies that the the angle $\angle(b_h, n_{\gamma(t)}) \in \left[\frac{\pi}{2} - \alpha, \frac{\pi}{2} + \alpha\right]$. It follows that the cosine of the angle is in the range $[-\sin\alpha, \sin\alpha]$.

Assume without loss of generality, reparameterizing γ if necessary, that γ is a unit-speed curve. Recall from elementary differential geometry [26, Page 166, Equation 7.7] that $\ddot{\gamma} = \kappa_n n + \kappa_g n \times \dot{\gamma}$, where n is the unit normal vector to the surface in which γ is embedded, $\dot{\gamma}$ is the unit tangent vector, and κ_n, κ_g are the normal and geodesic curvatures respectively. (Obviously these quantities are all parameterized by the specific point along γ). This fact follows from the fact that, since γ is unit-speed, $\ddot{\gamma}$ is orthogonal to $\dot{\gamma}$, and thus must lie in the plane spanned by n and $n \times \dot{\gamma}$. ($\dot{\gamma}(t) \in T_{\gamma(t)}\Sigma$ is also orthogonal to n.)

Recall further that the curvature $\kappa = \|\ddot{\gamma}\|$. From this, and the above equations, it is easy to prove that $\kappa_n = \ddot{\gamma} \cdot n = \kappa\cos\varphi$, where φ is the angle between the surface normal n and the principal normal m of γ, which, in this case since γ is a plane curve, lies in h [26, Page 166, Proposition 7.3.2].

In the special case where γ is defined by the intersection of a plane h with Σ, the angle $\varphi = \pi/2 - \theta$, where $\theta = \angle(T_{\gamma(t)}\Sigma, h)$ is the angle between the tangent space $T_{\gamma(t)}\Sigma$ and h. (Note that this is the exact same argument used to prove Meusnier's Theorem; we refer the reader to [26, Page 168, Proposition 7.3.4] for the proof.) Since $\angle(T_{\gamma(t)}\Sigma, h) = \angle(n_{\gamma(t)}, b_h)$, it follows that $\kappa_n = \kappa\cos(\pi/2 - \theta) = \kappa\sin\theta$.

Square both sides of $\kappa_n = \kappa\sin\theta$ and apply the bound on the cosine of $\angle(b_h, n_{\gamma(t)})$ to get

$$\kappa^2 = \frac{\kappa_n^2}{\sin^2\theta} = \frac{\kappa_n^2}{1 - \cos^2\theta} \le \frac{\kappa_n^2}{1 - \sin^2\alpha} \implies \kappa \le \frac{|\kappa_n|}{\sqrt{1 - \sin^2\alpha}}.$$

(We recognize that $\sin^2\alpha + \cos^2\alpha = 1$; we keep this form because it is more convenient later on.)

Now we simply apply the fact that κ_n is bounded by the minimum and maximum principal curvatures [26, Page 189, Corollary 8.2.5], and thus $|\kappa_n| \leq \max\{|\kappa_1|, |\kappa_2|\}$, to get the result. (For the reader not acquainted with this result, the statement that $\min\{\kappa_1, \kappa_2\} \leq \kappa_n \leq \max\{\kappa_1, \kappa_2\}$ is, at its core, simply a statement that the convex combination of the minimum and maximum eigenvalues of a symmetric matrix lies between the minimum and maximum eigenvalues.)

Finally a note to the reader on the use of the absolute value. The signs of κ_n and the principal curvatures κ_1, κ_2 are *not* intrinsic to the quantities. Only the magnitudes of these quantities are well-defined. The signs of these quantities are determined by the choice of orientation of the normal field of Σ. Flipping the orientation flips the sign of all of these quantities. The use of the absolute value reflects this fact. ◀

B.3 Bounds in Terms of Local Feature Size

A normal plane Π at a point $x \in \Sigma$ is a plane containing the normal vector n_x. It follows that Π must necessarily contain a unique unit tangent vector of $T_x\Sigma$. The plane Π cuts Σ in a plane curve. The curvature of this plane curve is defined as the reciprocal of the radius of the osculating circle at x. Notice that the radius of the osculating circle is at least $\mathrm{lfs}(x)$ since there is an empty ball of radius $\mathrm{lfs}(x)$ tangent to Σ at x. This holds for any normal plane at x. In particular it holds for the planes of principal curvature. Thus the maximum principal curvature at x is at most $\frac{1}{\mathrm{lfs}(x)}$. Combining this observation with Lemma 19 we derive an upper bound on the curvature κ of the plane curve γ.

▶ **Lemma 20.** *Let $\Sigma \subset \mathbb{R}^3$ be a smooth surface and let $p, q \in \Sigma$ be points on Σ such that $d(p,q) \leq \rho\,\mathrm{lfs}(p)$ for $\rho < 1/2$. Let $h = \mathrm{span}\{n_p, \vec{pq}\}$ be a plane and define $b_h = \frac{n_p \times \vec{pq}}{\|n_p \times \vec{pq}\|}$, the unit vector normal to h. Define a curve $\zeta : [0,1] \to \Sigma$ as the intersection $h \cap \Sigma$ such that $\zeta(0) = p$ and $\zeta(1) = q$ and ζ is included in the ball $B = B((p+q)/2, d(p,q)/2)$. Then the curvature κ of the curve ζ is at most*

$$\frac{1}{\mathrm{lfs}(p)\sqrt{1-2\rho}}.$$

Proof. For any point r on Σ, the medial balls at r impose a lower bound on the radius of any osculating circle generated by the intersection of Σ with a plane normal to Σ at r. In particular, the radii of the osculating circles that define the principal curvatures are also at least the local feature size. Applying the bound in Lemma 19, we have

$$\begin{aligned}\kappa &\leq \frac{\max\{\kappa_1, \kappa_2\}}{\sqrt{1-\sin^2\alpha}} \\ &\leq \frac{1}{\min_{r\in\zeta}\mathrm{lfs}(r)}\frac{1}{\sqrt{1-\sin^2\alpha}}.\end{aligned}$$

Suppose that $r^* = \arg\min_{r\in\zeta}\mathrm{lfs}(r)$, minimizes the local feature size along ζ. Since ζ is entirely included in the ball B, $d(p,r^*) \leq \rho\,\mathrm{lfs}(p)$. Then, by the Feature Translation Lemma (Lemma 12), we have that

$$\begin{aligned}\mathrm{lfs}(p) &\leq \frac{1}{1-\rho}\mathrm{lfs}(r^*) \\ (1-\rho)\mathrm{lfs}(p) &\leq \mathrm{lfs}(r^*).\end{aligned}$$

Thus we make the bound weaker by writing it as

$$\kappa \leq \frac{1}{(1-\rho)\mathrm{lfs}(p)} \frac{1}{\sqrt{1-\sin^2\alpha}}.$$

Finally, by the Normal Variation Lemma (Lemma 13), $\angle(n_p, n_q) \leq \frac{\rho}{1-\rho}$. Furthermore, $\sin\frac{\rho}{1-\rho} \leq \frac{\rho}{1-\rho}$ for $\rho < 1$. (The values are essentially identical for $\rho < \frac{2}{5}$.) Our final bound is

$$\begin{aligned}\kappa &\leq \frac{1}{(1-\rho)\mathrm{lfs}(p)\sqrt{1-\left(\frac{\rho}{1-\rho}\right)^2}} \\ &= \frac{1}{(1-\rho)\mathrm{lfs}(p)\sqrt{\frac{1-2\rho}{(1-\rho)^2}}} \\ &= \frac{1}{\mathrm{lfs}(p)\sqrt{1-2\rho}}.\end{aligned}$$

◀

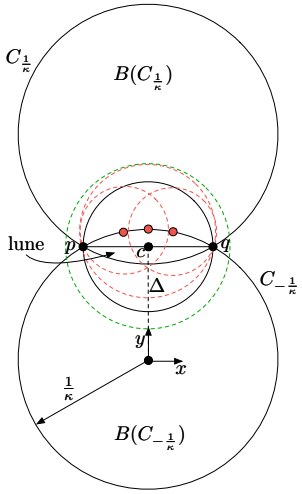

Figure 10 The circles $C_{1/\kappa}$ and $C_{-1/\kappa}$, and the disks $B(C_{1/\kappa})$ and $B(C_{-1/\kappa})$ bounded by these circles, all have radius $1/\kappa$. By Lemma 16, if the curvature of ζ_{pq} is at most κ, ζ_{pq} lies in the lune $B(C_{1/\kappa}) \cap B(C_{-1/\kappa})$. At each point z in the lune, we place a circle whose radius is $\min\{d(p,z), d(q,z)\}$. Three such circles are drawn in red. The union of all these circles is included in the green circle, centered at the midpoint of pq.

Recall Lemma 16, which states that γ is included in $B(C_{1/\kappa}) \cap B(C_{-1/\kappa})$, shown pictorially in Figure 10. We call this region the *lune*. One important feature of the lune is its *width*, the distance from the midpoint of pq to the north pole of $C_{-1/\kappa}$ (or equivalently to the south pole of $C_{1/\kappa}$). This distance was derived in terms of κ in Lemma 14. Using the bound in Lemma 20, we derive an upper bound on the width in terms of the local feature size.

▶ **Lemma 21.** *Let $\Sigma \subset \mathbb{R}^3$ be a smooth surface and let $p, q \in \Sigma$ be points on Σ such that $d(p,q) \leq \rho\,\mathrm{lfs}(p)$ for $\rho < 1/2$. Let $h = \mathrm{span}\{n_p, \vec{pq}\}$ be a plane and define $b_h = \frac{n_p \times \vec{pq}}{\|n_p \times \vec{pq}\|}$, the unit vector normal to h. Define a curve $\gamma : [0,1] \to \Sigma$ as the intersection $h \cap \Sigma$ such that*

$\gamma(0) = p$ and $\gamma(1) = q$ and γ is included in the ball $B = B((p+q)/2, d(p,q)/2)$. Then the width of the lune $B(C_{1/\kappa}) \cap B(C_{-1/\kappa})$ is at most

$$\sqrt{1-2\rho}\left(1 - \sqrt{1 - \frac{\rho^2}{4(1-2\rho)}}\right)\mathrm{lfs}(p).$$

Proof. The result follows by substituting the bounds for κ and $d(p,q)$ into the expression for the width of the lune given by Lemma 14. This expression is an upper bound because increasing the radius (equivalently, decreasing κ) decreases the width of the lune. ◀

The proof of Lemma 14 gives an expression for Δ. The following lemma gives a lower bound on Δ in terms of the local feature size.

▶ **Lemma 22.** *Let $\Sigma \subset \mathbb{R}^3$ be a smooth surface. Let $p, q \in \Sigma$ be points such that $d(p,q) \leq \rho\,\mathrm{lfs}(p)$ for some $\rho < 2\sqrt{5} - 4 \doteq 0.472136$. Let $h = \mathrm{span}\{n_p, \vec{pq}\}$ be a plane. Parameterize the curve $h \cap \Sigma$ as $\gamma : [0,1] \to \Sigma$ such that $\gamma(0) = p$, $\gamma(1) = q$, and γ is a bijection from $[0,1]$ to $h \cap \Sigma$. Let B be the diametric ball of the segment pq, and suppose that $\gamma \subset B$. Let κ be the maximum curvature of γ. Let $C_{1/\kappa}$ and $C_{-1/\kappa}$ be the circles on h defined in Lemma 16, both of which enclose γ. Let Δ be the distance from the center of $C_{1/\kappa}$ (or the center of $C_{-1/\kappa}$) to the center of B. Then*

$$\Delta \geq \sqrt{1 - 2\rho - \frac{\rho^2}{4}}\,\mathrm{lfs}(p).$$

Proof. Substituting the upper bound on κ derived in Lemma 20 into the expression for Δ derived in Lemma 14 gives

$$\begin{aligned}
\Delta &= \frac{\sqrt{1 - \kappa^2\, d(p,q)^2/4}}{\kappa} \\
&\geq \sqrt{1-2\rho}\sqrt{1 - \frac{1}{\mathrm{lfs}(p)^2(1-2\rho)}\frac{\rho^2\mathrm{lfs}(p)^2}{4}}\,\mathrm{lfs}(p) \\
&= \sqrt{1 - 2\rho - \frac{\rho^2}{4}}\,\mathrm{lfs}(p).
\end{aligned}$$

◀

Combining Lemmas 17, 20, and 22 gives an upper bound on the distance between any point on γ and the nearest endpoint of the segment pq.

▶ **Lemma 23.** *Let $\Sigma \subset \mathbb{R}^3$ be a smooth surface. Let $p, q \in \Sigma$ be points on Σ such that $d(p,q) \leq \rho\,\mathrm{lfs}(p)$ for some $\rho < 2\sqrt{5} - 4 \doteq 0.472136$. Let $h = \mathrm{span}\{n_p, \vec{pq}\}$ be a plane. Parameterize the curve $h \cap \Sigma$ as $\gamma : [0,1] \to \Sigma$ such that $\gamma(0) = p$, $\gamma(1) = q$, and γ is a bijection from $[0,1]$ to $h \cap \Sigma$. Let B be the diametric ball of the segment pq, and suppose that $\gamma \subset B$. Let κ be the maximum curvature of γ. Let $C_{1/\kappa}$ and $C_{-1/\kappa}$ be the circles on h defined in Lemma 16, both of which enclose γ. Then for every point $x \in \gamma$,*

$$\min\{d(x,p), d(x,q)\} \leq \beta\,\mathrm{lfs}(p), \qquad \text{where } \beta = \sqrt{2 - 4\rho - \sqrt{(1-2\rho)(4-8\rho-\rho^2)}}.$$

Proof. By Lemma 16, $\gamma \subset B(C_{1/\kappa}) \cap B(C_{-1/\kappa})$, so $x \in B(C_{1/\kappa}) \cap B(C_{-1/\kappa})$. The maximum distance between any point $y \in B(C_{1/\kappa}) \cap B(C_{-1/\kappa})$ and the nearest of p or q is achieved where the bisector of pq intersects the boundary of $B(C_{1/\kappa}) \cap B(C_{-1/\kappa})$. By Lemma 17

(with $\theta = \pi/2$), this distance is $\sqrt{r^2 + \Delta^2 - 2r\Delta + d(p,q)^2/4}$, where $r = 1/\kappa$ and Δ is the distance from the center of $C_{1/\kappa}$ (or the center of $C_{-1/\kappa}$) to the center of B. By Lemma 20, $\kappa \leq 1/\big(\sqrt{1-2\rho}\,\mathrm{lfs}(p)\big)$. By Lemma 22, $\Delta \geq \sqrt{1-2\rho-\rho^2/4}\,\mathrm{lfs}(p)$. The result follows by substitution. ◂

C Proof of the Cutting Plane Theorem

Recall the Cutting Plane Theorem (Theorem 1) from Section 3:

Let $s \in S$ be a segment with endpoints $p, q \in V$ such that $d(p,q) \leq \rho\,\mathrm{lfs}(p)$ for $\rho \leq 0.47$. Consider a point $x \in \Sigma_s^+$ and a site $v \in V \setminus \{p,q\}$ such that $x \in \mathrm{Vor}|_{\widetilde{\Sigma}} v$. Then v is strictly on the side of h_s opposite Σ_s^+. (The symmetric claim holds for any $x \in \Sigma_s^-$.)

Proof. Recall that s's cutting plane h_s induces a portal curve $\zeta_s \subset h_s$ and a unit vector b_s normal to h_s, with b_s determining the direction of the extrusions. For any point z, let $z_\perp \in \mathbb{R}$ denote the signed distance from z to h_s, being positive if z is on the same side of h_s as Σ_s^+ and negative if z is on the same side as Σ_s^-. As x lies on an extrusion Σ_s^+, we can write $x = \bar{x} + x_\perp b_s$ for some point $\bar{x} \in \zeta_s$ and some scalar $x_\perp \geq 0$. (To apply this proof to $x \in \Sigma_s^-$, simply reverse the direction of b_s.) Observe that $\bar{x}$ is both the point on h_s closest to x and the point on ζ_s closest to x. By assumption, no site lies on ζ_s except p and q, so v cannot lie on ζ_s nor anywhere else on the portal P_s (as $v \in \Sigma$).

Suppose for the sake of contradicting Theorem 1 that $v_\perp \geq 0$ (i.e., v is on the same side of h_s as Σ_s^+ or $v \in h_s$). Our proof has two parts: first we show that $x_\perp \leq 0.17\,\mathrm{lfs}(p)$; then we show by different means that $x_\perp \geq 0.48\,\mathrm{lfs}(p)$. The theorem holds by contradiction. Both parts work by identifying a medial ball that touches ζ_s (a different ball for each part) that constrains the location of v, thereby constraining the location of x.

As $x \in \mathrm{Vor}|_{\widetilde{\Sigma}} v$, there is a line segment $vx \subset \widetilde{X}$ that intersects the portal P_s at some point $c \in vx \cap P_s$ (because x lies on the extrusion Σ_s^+ and v is in the principal branch), illustrated in Figure 11. If there is more than one such intersection point, let c be the intersection point nearest v. Observe that $c_\perp \geq 0$ because $v_\perp \geq 0$ and $x_\perp \geq 0$. Recall that P_s is the union of the normal segments of the points on ζ_s. Let $\tilde{m} \in \zeta_s$ be the point whose normal segment $\ell_{\tilde{m}}$ contains c. Let m be the endpoint of $\ell_{\tilde{m}}$ such that $c \in \tilde{m}m$. Recall that m (like every normal segment endpoint) lies on Σ's medial axis M and is the center of an open medial ball B_m whose boundary is tangent to the surface Σ at $\tilde{m}$. As $B_m \cap \Sigma = \emptyset$, v cannot lie in B_m.

(Note that sometimes a medial "ball" is degenerate, effectively having infinite radius, thus being an open halfspace whose boundary is tangent to Σ at $\tilde{m}$. In this case, $\ell_{\tilde{m}}$ is a ray rather than a line segment, and although the point m is not defined, we can replace $\tilde{m}m$ with a ray originating at $\tilde{m}$ and lying on $\ell_{\tilde{m}}$. However, we note that a degenerate ball B_m cannot arise in this circumstance, because that would place the site v in the halfspace B_m. We won't prove that, because a degenerate B_m causes no difficulty for the rest of this proof.)

Let B_c be the open ball centered at c with $\tilde{m}$ on its boundary, and let $r_c = d(c,\tilde{m})$ be the radius of B_c, as illustrated in Figure 11. Observe that $B_c \subseteq B_m$, so B_c also does not intersect Σ nor contain v. Let θ be the angle at which the normal segment $\ell_{\tilde{m}}$ meets the cutting plane h_s, and observe that $c_\perp = r_c \sin\theta$. As $v \notin B_c$, $d(c,v) \geq r_c = c_\perp/\sin\theta$.

As $x \in \mathrm{Vor}|_{\widetilde{\Sigma}} v$, $\widehat{d}(x,v) \leq \min\{d(x,p), d(x,q)\}$. As p, q, and $\bar{x}$ all lie on h_s, we have by Pythagoras' Theorem that $d(x,p)^2 = x_\perp^2 + d(\bar{x},p)^2$ and $d(x,q)^2 = x_\perp^2 + d(\bar{x},q)^2$. As $\bar{x} \in \zeta_s$, we have by Lemma 23 that $\min\{d(\bar{x},p), d(\bar{x},q)\} \leq \beta\,\mathrm{lfs}(p)$ where $\beta = \sqrt{2 - 4\rho - \sqrt{(1-2\rho)\,(4-8\rho-\rho^2)}}$.

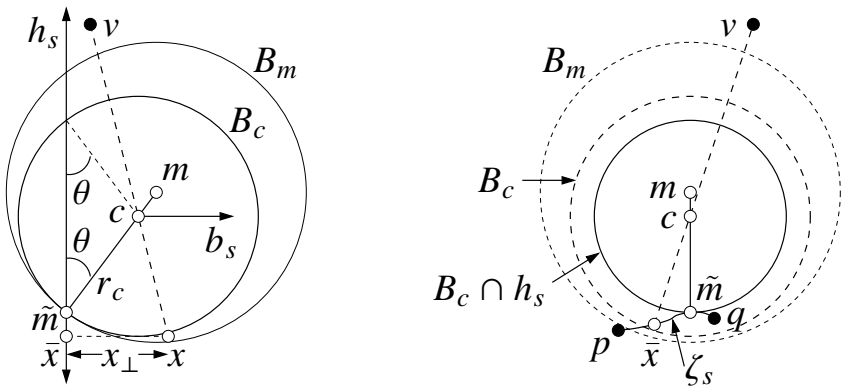

Figure 11 Left: side view of a cutting plane h_s with the extrusion Σ_s^+ extending to the right (with $\bar{x}x \subset \Sigma_s^+$). For a point $x \in \Sigma_s^+$ to lie in the Voronoi cell of a site v, the line segment vx must pass through the portal P_s at a point c, which lies on a line segment $m\tilde{m}$ connecting the center m of a medial ball B_m to a tangency point $\tilde{m} \in \Sigma$. (This figure is drawn so that c, m, and $\tilde{m}$ lie on the plane of the paper and h_s is orthogonal to the paper. However, v, x, and $\bar{x}$ do not generally lie on the plane of the paper.) As $v \notin B_m$ and $v_\perp \geq 0$ (that is, v is not left of h_s), the angle between the line $\overleftrightarrow{\bar{x}cv}$ and the plane h_s is at most θ. Therefore, $x_\perp$ cannot be large. Right: another view of the same configuration; here h_s is the plane of the paper.

Hence

$$\begin{aligned} d(x,v)^2 &\leq \min\{d(x,p)^2, d(x,q)^2\} \\ &= x_\perp^2 + \min\{d(\bar{x},p)^2, d(\bar{x},q)^2\} \\ &\leq x_\perp^2 + \beta^2 \operatorname{lfs}(p)^2. \end{aligned} \tag{4}$$

We consider two cases: one for $v_\perp \leq c_\perp \leq x_\perp$ and one for $v_\perp \geq c_\perp \geq x_\perp$.

In the case where $v_\perp \leq c_\perp \leq x_\perp$ as illustrated in Figure 11, the angle between the vector b_s and the vector $\vec{cv}$ is at least 90°. As $v_\perp \geq 0$ and $v \notin B_c$, the angle between b_s and $\vec{cv}$ cannot exceed $90° + \theta$. Thus the angle between the line $\overleftrightarrow{\bar{x}cv}$ and the plane h_s cannot exceed θ. So $x_\perp - c_\perp \leq d(x,c)\sin\theta$ and

$$\begin{aligned} \beta^2 \operatorname{lfs}(p)^2 &\geq d(x,v)^2 - x_\perp^2 \\ &= (d(x,c) + d(c,v))^2 - x_\perp^2 \\ &\geq \left(\frac{x_\perp - c_\perp}{\sin\theta} + \frac{c_\perp}{\sin\theta}\right)^2 - x_\perp^2 \\ &= \frac{x_\perp^2}{\sin^2\theta} - x_\perp^2 \\ &= x_\perp^2 \cot^2\theta. \end{aligned}$$

In the case where $v_\perp \geq c_\perp \geq x_\perp$,

$$\begin{aligned} \beta^2 \operatorname{lfs}(p)^2 &\geq d(x,v)^2 - x_\perp^2 \\ &\geq d(c,v)^2 - c_\perp^2 \\ &\geq \frac{c_\perp^2}{\sin^2\theta} - c_\perp^2 \\ &= c_\perp^2 \cot^2\theta \\ &\geq x_\perp^2 \cot^2\theta. \end{aligned}$$

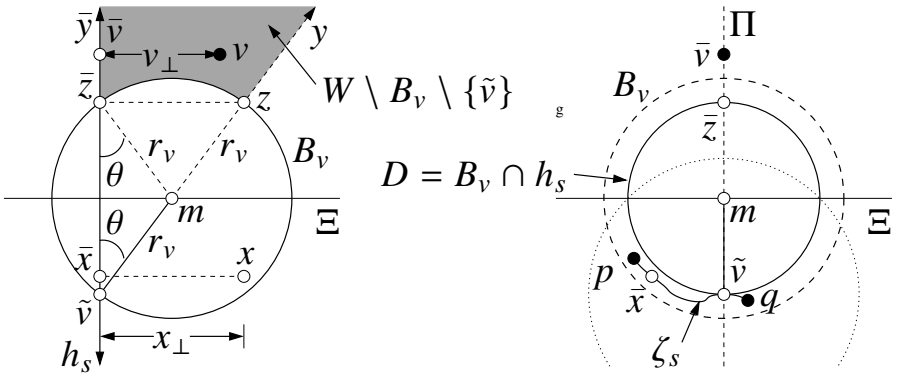

Figure 12 Left: side view of a cutting plane h_s. The paper is the plane Π orthogonal to the portal curve ζ_s at $\tilde{v}$, where $\tilde{v}$ is the point on ζ_s closest to the site v. Assuming that $v \in \Pi$, v must lie in the shaded region $W \setminus B_v \setminus y$. Note that x and $\bar{x}$ do not necessarily lie on the plane of the paper, but the other points do. For x to be in v's Voronoi cell rather than p's or q's, $x_\perp$ must be large. Right: another view of the same configuration; here the paper is the plane h_s, and we see side views of Π and Ξ.

In either case, $x_\perp \le \beta \tan\theta \,\mathrm{lfs}(p)$. As $d(p,\tilde{m}) \le d(p,q) \le \rho\,\mathrm{lfs}(p)$, the Normal Variation Lemma (Lemma 13) implies that $\theta \le \angle(n_p, n_{\tilde{m}}) \le \eta(\rho)$ where $\eta(\delta) = \arccos\left(1 - \delta^2/(2\sqrt{1-\delta^2})\right)$. Hence $x_\perp \le \beta \tan\eta(\rho)\,\mathrm{lfs}(p)$. (Note that β is a function of ρ.) From this inequality, it follows that $x_\perp \le 0.17\,\mathrm{lfs}(p)$ for all $\rho \in (0, 0.47]$, which completes the first part of our claim.

For the second part, we will show that $x_\perp \ge 0.48\,\mathrm{lfs}(p)$, yielding a contradiction.

Let $\tilde{v}$ be the point on ζ_s closest to v. As ζ_s is a smooth curve, there is a unique plane Π through $\tilde{v}$ that is locally orthogonal to ζ_s at $\tilde{v}$. As $\zeta_s \subset h_s$, Π is also orthogonal to the plane h_s. If $\tilde{v}$ is not an endpoint of ζ_s (p or q), then $v \in \Pi$, as $v\tilde{v}$ is orthogonal to ζ_s at $\tilde{v}$. However, if $\tilde{v}$ is an endpoint of ζ_s, then v might or might not lie on Π. If $v \notin \Pi$, then Π separates v from ζ_s.

Let $\ell_{\tilde{v}} \subset P_s$ be the normal segment through $\tilde{v}$, and let the line $L_{\tilde{v}}$ be the affine hull of $\ell_{\tilde{v}}$. Let $y \subset L_{\tilde{v}}$ be the ray that originates at $\tilde{v}$ and is half of $L_{\tilde{v}}$, on the positive side of h_s (the same side as Σ_s^+), as illustrated in Figure 12. As y is orthogonal to ζ_s at $\tilde{v}$, $y \subset \Pi$. Let the ray $\bar{y}$ be the orthogonal projection of y onto h_s. As the projection direction lies in Π, $\bar{y} \subset \Pi$. The rays y and $\bar{y}$ bound an infinite wedge $W \subset \Pi$ that has apex $\tilde{v}$, as illustrated.

As $x \in \mathrm{Vor}|_{\tilde{\Sigma}} v$, the site v is positioned so that the sightline vx enters the portal P_s from the correct side to access the extrusion Σ_s^+. Therefore, if $v \in \Pi$, then v lies in the wedge W because $v_\perp \ge 0$ (by assumption) but v lies on the same side of y as W Note that $v \ne \tilde{v}$ and $v \notin y$ because $v \notin P_s$. On the other hand, if $v \notin \Pi$, let w be the point where the line segment vx intersects Π. Observe that $w_\perp \ge 0$ because $v_\perp \ge 0$ and $x_\perp \ge 0$. Hence, w lies in the wedge W.

The ray y passes through an endpoint m of the normal segment $\ell_{\tilde{v}}$, and $m \in M$ is the center of an open medial ball B_v that is tangent to Σ at $\tilde{v}$. As $B_v \cap \Sigma = \emptyset$, v cannot lie in B_v. (Note that the medial ball B_v cannot degenerate to a halfspace because then B_v would include $W \setminus \{\tilde{v}\}$, contradicting $v \in W \setminus B_v \setminus y$.)

Let $r_v = d(m, \tilde{v})$ be the radius of B_v. Observe that $\mathrm{lfs}(\tilde{v}) \le r_v$. The rays y and $\bar{y}$ each intersect the boundary of B_v at two points: they both enter the ball at the ray origin $\tilde{v}$ and they each exit at another point. Let z be the point where y exits, and let $\bar{z}$ be the point where $\bar{y}$ exits, as illustrated in Figure 12. By circle geometry, $\angle z\bar{z}\tilde{v} = 90°$, so $\bar{z}$ is the point closest to z on h_s. Let θ be the wedge angle at which y meets $\bar{y}$. Then $d(\tilde{v}, z) = 2r_v$ and

$d(\tilde{v}, \bar{z}) = 2r_v \cos\theta$.

As $\bar{x}$ and $\tilde{v}$ both lie on ζ_s, $d(\bar{x}, \tilde{v}) \leq d(p, q) \leq \rho\,\mathrm{lfs}(p)$ and $d(p, \tilde{v}) \leq d(p, q) \leq \rho\,\mathrm{lfs}(p)$. The Feature Translation Lemma (Lemma 12) implies that $\mathrm{lfs}(p) \leq \mathrm{lfs}(\tilde{v})/(1-\rho)$ and the Normal Variation Lemma (Lemma 13) implies that $\theta \leq \angle(n_p, n_{\tilde{v}}) \leq \eta(\rho)$.

Consider the open disk $D = B_v \cap h_s$, illustrated at right in Figure 12. The line segment $\tilde{v}\bar{z}$ is a diameter of D with length $2r_v \cos\theta$. D and $\bar{x}$ both lie on the cutting plane h_s but $\bar{x} \notin D$, as D does not intersect Σ. Hence we have by circle geometry that $\angle\tilde{v}\bar{x}\bar{z} \leq 90°$ and by Pythagoras' Theorem that $d(\bar{x}, \tilde{v})^2 + d(\bar{x}, \bar{z})^2 \geq d(\tilde{v}, \bar{z})^2 = 4r_v^2 \cos^2\theta$.

Let Ξ be the plane that bisects $\tilde{v}\bar{z}$; note that $m \in \Xi$ and Ξ cuts both B_v and D in half. We claim that $\tilde{v}$, x, and $\bar{x}$ all lie on the same side of Ξ; that is, $d(\tilde{v}, x) < d(\bar{z}, x)$ and $d(\tilde{v}, \bar{x}) < d(\bar{z}, \bar{x})$. This claim holds because for all $\rho \in [0, 0.53]$, $\rho < \sqrt{2}(1-\rho)\cos\eta(\rho)$ and therefore $d(\tilde{v}, \bar{x}) \leq d(p, q) \leq \rho\,\mathrm{lfs}(p) < \sqrt{2}(1-\rho)\,\mathrm{lfs}(p)\cos\eta(\rho) \leq \sqrt{2}\,\mathrm{lfs}(\tilde{v})\cos\theta \leq \sqrt{2}r_v\cos\theta = d(\tilde{v}, \bar{z})/\sqrt{2}$. This means that $\bar{x}$ is in the open ball with center $\tilde{v}$ and radius equal to $\sqrt{2}$ times the radius of D (see the dotted sphere at right in Figure 12). The boundary of this open ball intersects the boundary of D at two points that lie on Ξ, as illustrated. As $\bar{x} \notin D$, $\bar{x}$ must lie on the same side of Ξ as $\tilde{v}$. Furthermore, x must lie on the same side of Ξ as $\bar{x}$, because $\bar{x}$ is defined by projecting x in a direction parallel to Ξ.

We consider two cases, depending on whether $v \in \Pi$. For the case where $v \in \Pi$, v lies in the region $W \setminus B_v \setminus \{\tilde{v}\}$, which is shaded in Figure 12. The boundary of $W \setminus B_v \setminus \{\tilde{v}\}$ consists of three curves: a circular arc connecting z to $\bar{z}$, a ray that originates at z and is a subset of the ray y, and a ray $\bar{y}_{\bar{z}}$ that originates at $\bar{z}$ and is a subset of the ray $\bar{y}$. Let $\bar{v}$ be the point on h_s closest to v. As the figure makes clear, $\bar{v} \in \bar{y}_{\bar{z}}$. We see that $d(\bar{x}, \bar{v}) \geq d(\bar{x}, \bar{z})$, because $\bar{z}$ is the point in $\bar{y}_{\bar{z}}$ that is closest to Ξ (which is orthogonal to $\bar{y}_{\bar{z}}$) and $\bar{x}$ is on the other side of Ξ.

From Inequality (4) we have

$$\begin{aligned}
x_\perp^2 &\geq d(x, v)^2 - \beta^2\,\mathrm{lfs}(p)^2 \\
&= d(\bar{x}, \bar{v})^2 + (x_\perp - v_\perp)^2 - \beta^2\,\mathrm{lfs}(p)^2 \\
&\geq d(\bar{x}, \bar{z})^2 - \beta^2\mathrm{lfs}(p)^2; \\
&\geq d(\tilde{v}, \bar{z})^2 - d(\bar{x}, \tilde{v})^2 - \beta^2\,\mathrm{lfs}(p)^2 \\
&\geq 4r_v^2\cos^2\theta - \rho^2\,\mathrm{lfs}(p)^2 - \beta^2\,\mathrm{lfs}(p)^2 \\
&\geq 4\,\mathrm{lfs}(\tilde{v})^2\cos^2\theta - (\rho^2 + \beta^2)\,\mathrm{lfs}(p)^2 \\
&\geq \left(4(1-\rho)^2\cos^2\eta(\rho) - \rho^2 - \beta^2\right)\mathrm{lfs}(p)^2.
\end{aligned}$$

From this inequality, it follows that $x_\perp \geq 0.74\,\mathrm{lfs}(p)$ for all $\rho \in (0, 0.47]$.

Now consider the case where $v \notin \Pi$. In this case, the point closest to v on the portal curve ζ_s is an endpoint of ζ_s; suppose without loss of generality that it is $\tilde{v} = q$. Recall that the plane Π is orthogonal to ζ_s at q. The site v and the curve ζ_s are on opposite sides of Π (otherwise, q could not be the point closest to v). Hence, v and x are on opposite sides of Π or $x \in \Pi$.

Let ℓ_q be q's normal segment, let B_q be the open medial ball tangent to Σ at q whose center m satisfies $m_\perp \geq 0$, and let $r_q = d(q, m)$ be the radius of B_q. (Note that as $\tilde{v} = q$, ℓ_q is the same normal segment we have already been calling $\ell_{\tilde{v}}$, and B_q is the same ball we have been calling B_v.) We define W, z, $\bar{z}$, y, and $\bar{y}$ as before; see Figure 13. As $B_q \cap \Sigma = \emptyset$, v cannot lie in B_q.

Let the line $\ell \supset \ell_q$ be the affine hull of ℓ_q. The line ℓ cuts Π into two halfplanes; let Π^+ be the open halfplane on the positive side of ℓ, where the extrusion goes, and let Π^- be the closed halfplane on the negative side, so that $\bar{y} \subset \Pi^-$ and $W \subset \Pi^-$.

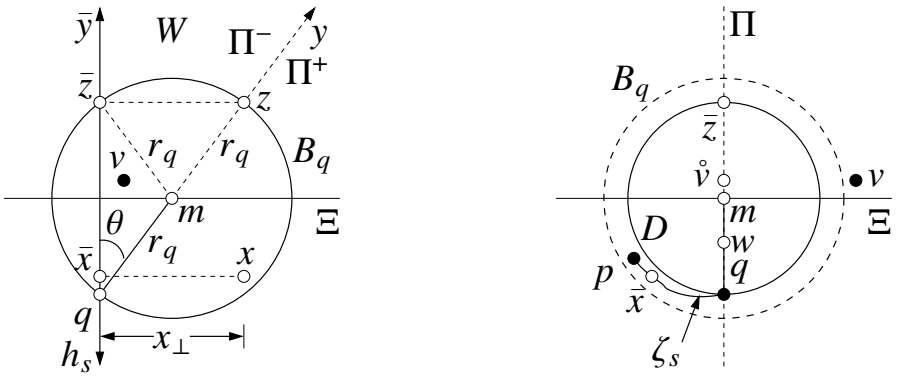

Figure 13 The circumstances of this figure are similar to those of Figure 12, but v does not lie on the plane Π (so it is not necessarily restricted to the region shaded in Figure 12).

Define a coordinate system with q at the origin, such that for any point $p \in \mathbb{R}^3$, $p_\perp$ is p's signed distance from h_s (as we have already defined), p_Π is p's signed distance from Π, with the sign defined so that $v_\Pi > 0$ and $x_\Pi \leq 0$, and $p_\parallel$ is the coordinate in the direction $\bar{y}$ (the vertical axis in Figure 13). As $x \in \mathrm{Vor}|_{\widetilde{\Sigma}} v$, $d(x,v) \leq d(x,q)$. Therefore, $\|v\|^2 - 2x \cdot v \leq \|q\|^2 - 2x \cdot q$. We have chosen q as the origin, so we can write this as $\|v\|^2/2 \leq x \cdot v = x_\perp v_\perp + x_\parallel v_\parallel + x_\Pi v_\Pi$. As $x_\Pi v_\Pi \leq 0$, we can shorten this to $\|v\|^2/2 \leq x_\perp v_\perp + x_\parallel v_\parallel$. As $v \neq q$, this implies that $0 < x_\perp v_\perp + x_\parallel v_\parallel$.

As $v \notin B_q$, $d(m,v) \geq r_q$. Therefore, $\|m\|^2 - 2m \cdot v + \|v\|^2 \geq r_q^2$. But $\|m\| = r_q$, so we can write $\|v\|^2/2 \geq m \cdot v = m_\perp v_\perp + m_\parallel v_\parallel$. Combining this with an inequality from the previous paragraph gives $m_\perp v_\perp + m_\parallel v_\parallel \leq x_\perp v_\perp + x_\parallel v_\parallel$, or equivalently, $0 \leq (x_\perp - m_\perp)v_\perp + (x_\parallel - m_\parallel)v_\parallel$.

Let $\mathring{x}$ be the point closest to x on Π, and let $\mathring{v}$ be the point closest to v on Π. (Hence $\mathring{x}$ sets $\mathring{x}_\Pi = 0$ while retaining the coordinates $\mathring{x}_\perp = x_\perp$ and $\mathring{x}_\parallel = x_\parallel$.) We claim that $\mathring{x} \in \Pi^+$. (That is, with respect to the figure, $\mathring{x}$ is to the right of y.) To see this, suppose for the sake of contradiction that $\mathring{x} \in \Pi^-$. This implies that $\mathring{x}$ lies on or left of the ray $\vec{qm}$; equivalently, $m_\perp x_\parallel \geq m_\parallel x_\perp$. We know that $x_\perp \geq 0$, $m_\parallel > 0$, and $m_\perp > 0$, so it follows that $x_\parallel \geq 0$. Hence we can transform the inequality at the end of the last paragraph to $0 \leq (x_\perp x_\parallel - m_\perp x_\parallel)v_\perp + (x_\parallel - m_\parallel)x_\parallel v_\parallel$, and then to $0 \leq (x_\parallel - m_\parallel)x_\perp v_\perp + (x_\parallel - m_\parallel)x_\parallel v_\parallel$. As $m \in \Xi$ and x lies on the same side of Ξ as q, we have $x_\parallel - m_\parallel < 0$, so it follows that $0 \geq x_\perp v_\perp + x_\parallel v_\parallel$. But this contradicts the fact that $0 < x_\perp v_\perp + x_\parallel v_\parallel$ (from two paragraphs ago).

By this contradiction, we establish our claim that $\mathring{x} \in \Pi^+$. But the point $w = vx \cap \Pi$ must lie in Π^-; if it did not, the sightline from x to v would not exit the secondary branch and enter the principal branch. Observe that w lies on the line segment $\mathring{v}\mathring{x}$, which implies that at least one of $\mathring{x}$ or $\mathring{v}$ is in Π^-. As $\mathring{x}$ is not, we must have $\mathring{v} \in \Pi^-$. Equivalently, $m_\perp v_\parallel \geq m_\parallel v_\perp$. As $v_\perp \geq 0$, $\mathring{v}$ lies in the wedge W and thus $v_\parallel \geq 0$. Moreover, it is not possible that $\mathring{v} = q$, because then we would have $d(x,q) < d(x,v)$ and x could not be in v's Voronoi cell. Therefore, $\mathring{v}$ lies in $W \setminus \{q\}$ and thus $v_\parallel > 0$.

As $m_\parallel > 0$, we can transform the inequality from three paragraphs ago to $0 \leq (x_\perp - m_\perp)m_\parallel v_\perp + (x_\parallel - m_\parallel)m_\parallel v_\parallel$, and then to $0 \leq (x_\perp - m_\perp)m_\perp v_\parallel + (x_\parallel - m_\parallel)m_\parallel v_\parallel$, and then to $0 \leq (x_\perp - m_\perp)m_\perp + (x_\parallel - m_\parallel)m_\parallel = x \cdot m - \|m\|^2$. (The last identity follows because $m_\Pi = 0$.)

In this coordinate system with origin q, $z = 2m$. We claim that $d(x,z) \leq d(x,q)$. To see this, observe that $d(x,z)^2 - d(x,q)^2 = \|z\|^2 - 2x \cdot z = 4\|m\|^2 - 4x \cdot m \leq 0$.

Therefore,

$$\begin{aligned}
0 &\geq d(x,z)^2 - d(x,q)^2 \\
&= d(\bar{x},\bar{z})^2 + (x_\perp - z_\perp)^2 - d(\bar{x},q)^2 - x_\perp^2 \\
&\geq d(\tilde{v},\bar{z})^2 - d(\bar{x},\tilde{v})^2 - 2x_\perp z_\perp + z_\perp^2 - \rho^2 \operatorname{lfs}(p)^2 \\
&\geq 4r_q^2 \cos^2\theta - \rho^2 \operatorname{lfs}(p)^2 - 4r_q x_\perp \sin\theta + 4r_q^2 \sin^2\theta - \rho^2 \operatorname{lfs}(p)^2 \\
&= 4r_q^2 - 2\rho^2 \operatorname{lfs}(p)^2 - 4r_q x_\perp \sin\theta \\
&\geq 4r_q^2 - 2\rho^2 \operatorname{lfs}(p)^2 - 4r_q x_\perp \sin\eta(\rho).
\end{aligned}$$

Therefore,

$$\begin{aligned}
x_\perp &\geq \frac{1}{\sin\eta(\rho)} \left(r_q - \frac{\rho^2}{2r_q} \operatorname{lfs}(p)^2 \right) \\
&\geq \frac{1}{\sin\eta(\rho)} \left(\operatorname{lfs}(\tilde{v}) - \frac{\rho^2}{2\operatorname{lfs}(\tilde{v})} \operatorname{lfs}(p)^2 \right) \\
&\geq \frac{1}{\sin\eta(\rho)} \left(1 - \rho - \frac{\rho^2}{2(1-\rho)} \right) \operatorname{lfs}(p).
\end{aligned}$$

From this inequality, it follows that $x_\perp \geq 0.48 \operatorname{lfs}(p)$ for all $\rho \in (0, 0.5]$. ◀

D Proof of the Possession Theorem

Recall the Possession Theorem (Theorem 3) from Section 3:

Let $s \in S$ be a segment with endpoints $p, q \in V$ such that $d(p,q) \leq \rho \operatorname{lfs}(p)$ for $\rho \leq 0.47$. Let c be the midpoint of s. Let $v \in V$ be a site whose extended Voronoi cell $\operatorname{Vor}|_{\widetilde{\Sigma}} v$ contains a point $x \in \Sigma_s^+$ or $x \in \Sigma_s^-$. Then v lies in the ball $B(c, \lambda \operatorname{lfs}(p))$ with center c and radius $\lambda \operatorname{lfs}(p)$, where

$$\lambda = \sqrt{1-2\rho} \left(1 - \sqrt{1 - \frac{\rho^2}{4(1-2\rho)}} \right) + \sqrt{(2-4\rho)\left(1 - \sqrt{1 - \frac{\rho^2}{4(1-2\rho)}} \right)}.$$

We will need a more general version of the Possession Theorem to prove the Extrusion Shadow Lemma (Lemma 30) in Appendix F. The Possession Theorem is a special case of the following lemma.

▶ **Lemma 24.** *Let $s \in S$ be a segment with endpoints $p, q \in V$ such that $d(p,q) \leq \rho \operatorname{lfs}(p)$ for $\rho \leq 0.47$. Let c be the midpoint of s. Let $x \in \Sigma_s^+$. Let $y \in \widetilde{X}$ be a point that is not on the same side of the cutting plane h_s as Σ_s^+ (though y may lie on h_s, and so may x) such that $d(x,y) \leq \min\{d(x,p), d(x,q)\}$. ($\Sigma_s^+$ can be replaced by Σ_s^-. Note that by Theorem 1, these conditions hold for any site $y \in V$ such that $x \in \operatorname{Vor}|_{\widetilde{\Sigma}} y$). Then y lies in the ball $B(c, \lambda \operatorname{lfs}(p))$ with center c and radius $\lambda \operatorname{lfs}(p)$, where*

$$\lambda = \sqrt{1-2\rho} \left(1 - \sqrt{1 - \frac{\rho^2}{4(1-2\rho)}} \right) + \sqrt{(2-4\rho)\left(1 - \sqrt{1 - \frac{\rho^2}{4(1-2\rho)}} \right)}.$$

Proof. Let $\bar{x}$ be the point nearest to x on h_s, and observe that $\bar{x} \in \zeta_s$. Consider a closed ball B_x centered at x with radius $d(x,y)$. As $d(x,y) \leq \min\{d(x,p), d(x,q)\}$, neither p nor q is in the interior of B_x. The intersection of B_x with h_s is a closed disk with center $\bar{x}$. As $p, q \in h_s$,

the radius of the disk $B_x \cap h_s$ is at most $\min\{d(\bar{x}, p), d(\bar{x}, q)\}$. As y and x are not on the same side of h_s, y lies in the ball centered at $\bar{x}$ with the same radius as $B_x \cap h_s$.

It follows that y must be included in the union of balls constructed by centering a ball at each point $w \in \zeta_s$ with radius $\min\{d(w, p), d(w, q)\}$. Lemma 16 states that ζ_s is a subset of the lune $B(C_{1/\kappa}) \cap B(C_{-1/\kappa})$, where κ is the curvature of ζ_s. (See Figure 10.) Let $\mathcal{B}$ be the union of all balls centered at every point in the lune. $\mathcal{B}$ is included in a ball centered at the midpoint c of pq with radius λ. To see this, we show that the point farthest from c in $\mathcal{B}$ is at most distance $\lambda\,\mathrm{lfs}(p)$ away. To find the farthest point, we focus on the balls centered at points on the boundary of the lune. By symmetry we need consider only one of the two arcs of the lune, and only the portion from p to the point directly above c in Figure 10.

We define a local coordinate system as shown in Figure 10, with the origin at the center of $C_{-1/\kappa}$, $c = (0, \Delta)$, and $p = (-d(p, q)/2, \Delta)$. In this coordinate system, we can parameterize $C_{-1/\kappa}$ as $C_{-1/\kappa}(\theta) = \sqrt{d(p, q)^2/4 + \Delta^2}(-\cos\theta, \sin\theta)$. The distance from c to the farthest point of each ball along the lune can be expressed as the sum of two distances, the distance from c to $C_{-1/\kappa}(\theta)$ plus the radius of the ball at $C_{-1/\kappa}(\theta)$.

$$d(c, C_{-1/\kappa}(\theta)) + d(C_{-1/\kappa}(\theta), p) = \sqrt{\frac{d(p,q)^2}{4} + 2\Delta^2 - 2\Delta\sqrt{\frac{d(p,q)^2}{4} + \Delta^2}\sin\theta} + \sqrt{\frac{d(p,q)^2}{2} + 2\Delta^2 - d(p,q)\sqrt{\frac{d(p,q)^2}{4} + \Delta^2}\cos\theta - 2\Delta\sqrt{\frac{d(p,q)^2}{4} + \Delta^2}\sin\theta}.$$

The first derivative of this sum with respect to θ is

$$-\frac{\Delta\sqrt{\frac{d(p,q)^2}{4} + \Delta^2}\cos\theta}{\sqrt{\frac{d(p,q)^2}{4} + 2\Delta^2 - 2\Delta\sqrt{\frac{d(p,q)^2}{4} + \Delta^2}\sin\theta}} + \frac{\frac{1}{2}d(p,q)\sqrt{\frac{d(p,q)^2}{4} + \Delta^2}\sin\theta - \Delta\sqrt{\frac{d(p,q)^2}{4} + \Delta^2}\cos\theta}{\sqrt{\frac{d(p,q)^2}{2} + 2\Delta^2 - d(p,q)\sqrt{\frac{d(p,q)^2}{4} + \Delta^2}\cos\theta - 2\Delta\sqrt{\frac{d(p,q)^2}{4} + \Delta^2}\sin\theta}}.$$

By symmetry, we are interested in the zeros of the derivative in the range $\theta \in [\pi/2 - \arctan(d(p, q)/(2\Delta)), \pi/2]$. The derivative has a zero at $\theta = \pi/2 - \arctan(d(p, q)/(2\Delta))$, where $C_{-1/\kappa}(\theta) = p$, and no others in the range $[\pi/2 - \arctan(d(p, q)/(2\Delta)), \pi/2]$. Furthermore, the function is at a minimum at $\theta = \pi/2 - \arctan(d(p, q)/(2\Delta))$, with value $d(p, q)/2$. The derivative is positive at every other point in the range $[\pi/2 - \arctan(d(p, q)/(2\Delta)), \pi/2]$. Thus the maximum is achieved at one of the limits of the range, when $\theta = \pi/2$.

All that remains is to bound the sum of the distances when $\theta = \pi/2$. At $\theta = \pi/2$, Lemma 21 states that

$$d(c, C_{-1/\kappa}(\pi/2)) \leq \sqrt{1 - 2\rho}\left(1 - \sqrt{1 - \frac{\rho^2}{4(1 - 2\rho)}}\right)\mathrm{lfs}(p).$$

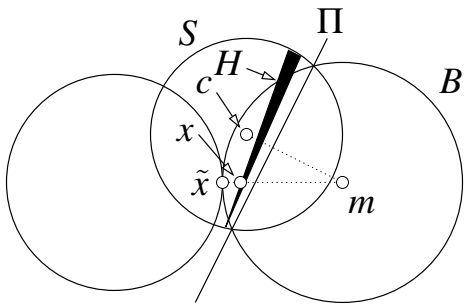

Figure 14 For every point $x \in H \setminus S$, $\tilde{x}$ is strictly inside S.

With the bound $d(c,p) = d(p,q)/2 \leq \rho\,\mathrm{lfs}(p)/2$, Pythagoras' Theorem implies that

$$\begin{aligned} d(C_{-1/\kappa}(\pi/2), p) &= \sqrt{d(c,p)^2 + d(c, C_{-1/\kappa}(\pi/2))^2} \\ &\leq \sqrt{\frac{\rho^2}{4} + (1-2\rho)\left(1 - \sqrt{1 - \frac{\rho^2}{4\,(1-2\rho)}}\right)^2}\,\mathrm{lfs}(p) \\ &= \sqrt{(2-4\rho)\left(1 - \sqrt{1 - \frac{\rho^2}{4\,(1-2\rho)}}\right)}\,\mathrm{lfs}(p). \end{aligned}$$

So $\mathcal{B}$ is a subset of a ball centered at c with radius $d(c, C_{-1/\kappa}(\pi/2)) + d(C_{-1/\kappa}(\pi/2), p) \leq \lambda\,\mathrm{lfs}(p)$. ◀

E The Nearest Point Map and Circumscribing Spheres

Recall that the nearest point map ν maps any point $x \in \mathbb{R}^3 \setminus M$ to the point $\tilde{x} = \nu(x)$ nearest x on Σ, where M is Σ's medial axis. The following lemma helps to constrain where $\tilde{x}$ can lie. We will use it later for three purposes: to prove conditions under which a mapped triangle $\nu(\tau)$ does not intersect any site other than τ's vertices (at the end of this section), to help prove conditions under which the boundaries of Voronoi cells are not determined by shadows cast by portals (in Section 4), and to help prove conditions under which ν defines a homeomorphism from τ to its image $\nu(\tau)$ on Σ (in Section G.3).

▶ **Lemma 25.** *Consider a sphere $S \subset \mathbb{R}^3$ and a set $P \subset \Sigma$ of points that lie on or inside S. Let H be the convex hull of P. Let r be the radius of S, and suppose that $r \leq \mathrm{lfs}(p)/2$ for some $p \in P$. Then for every point $x \in H$ that is strictly inside S, $\tilde{x} = \nu(x)$ is strictly inside S.*

Proof. Consider a point $x \in H \setminus S$, which is inside S. If $\tilde{x} = x$ the lemma follows immediately, so suppose that $\tilde{x} \neq x$ and thus $x \notin \Sigma$. There are two open medial balls tangent to Σ at $\tilde{x}$; let B be the one that contains x, as illustrated in Figure 14. Let m be the center of B; m lies on the medial axis of Σ. Observe that x lies on the line segment $m\tilde{x}$.

If the entire closure of B is strictly inside S, the lemma follows immediately; so assume it is not. The entirety of B cannot be outside S, as $x \in B$ and x is inside S. Nor is $S \subset B$ possible, as the points in P lie on Σ and thus are not in B. Hence the intersection of S with B's boundary is a circle or a point. If it is a circle, let Π be the affine hull of that circle, as illustrated; if it is a point, let Π be the plane tangent to S and B at that point. Let $\bar{\Pi}_S$ be the closed halfspace bounded by Π that includes $S \setminus B$, and let Π_S be the open version of the

same halfspace. The portion of B's boundary in Π_S is entirely enclosed by S. The portion of S in the open halfspace complementary to $\bar{\Pi}_S$ is entirely included in B. Every point in P lies on or inside S but not in B, hence $P \subset \bar{\Pi}_S$. Therefore, $H \subset \bar{\Pi}_S$ and $x \in \bar{\Pi}_S$.

Let c be the center of S. Observe that the plane Π is orthogonal to the line segment cm. By assumption, $r \leq \mathrm{lfs}(p)/2$, so $d(p,m) \geq \mathrm{lfs}(p) \geq 2r$. As p lies on or inside S and $d(p,m)$ is at least twice the radius r of S, it follows that m lies outside of S or on S. Hence $m \notin \bar{\Pi}_S$.

Given the facts that x lies on the line segment $m\tilde{x}$, $m \notin \bar{\Pi}_S$, $x \in \bar{\Pi}_S$, and $\tilde{x} \neq x$, it follows that $\tilde{x} \in \Pi_S$. As $\tilde{x}$ is also on B's boundary, $\tilde{x}$ is strictly inside S. ◀

A corollary of Lemma 25 is that if the vertices of a restricted Delaunay triangle τ are sufficiently close to τ's dual restricted Voronoi vertex, then the image of τ under the nearest point map ν does not intersect any site other than τ's vertices. We will use this fact in Appendix G.7 to help prove that under suitable sampling conditions, the nearest point map is a homeomorphism from (the underlying space of) a restricted CDT to the surface Σ.

▶ **Corollary 26.** *Let V be a finite set of sites on Σ. Let $p, p', p'' \in V$ be three sites that generate a restricted Voronoi vertex $u \in \mathrm{Vor}|_{\widetilde{\Sigma}} V$ and its dual restricted Delaunay triangle $\tau = \triangle pp'p''$. Suppose that $d(p,u) \leq \mathrm{lfs}(p)/2$. (Note that $d(p,u) = d(p',u) = d(p'',u)$.) Then $\nu(\tau)$ intersects no site in $V \setminus \{p, p', p''\}$.*

Proof. Suppose for the sake of contradiction that for some site $w \in V \setminus \{p, p', p''\}$, $w \in \nu(\tau)$. Let $x \in \tau$ be the point for which $\nu(x) = w$. As τ is a restricted Delaunay triangle dual to u, τ's vertices lie on a sphere S that has center u and encloses no site, particularly not w. The radius of S is $d(p,u)$. By Lemma 25 (with $P = \{p, p', p''\}$ and $H = \tau$), w is strictly inside S. The result follows by contradiction. ◀

F Proofs of the Shadow Theorem and its Corollaries

Recall the Shadow Theorem (Theorem 4) from Section 4:

Let $\Sigma \subset \mathbb{R}^3$ be a smooth surface without boundary and let $V \subset \Sigma$ be a finite set of points (sites). Let S be a set of segments (with endpoints in V) such that for every segment $s = pq \in S$, $d(p,q) \leq 0.47\,\mathrm{lfs}(p)$. Suppose that for every site $v \in V$ and every point x in the principal Voronoi cell $\mathrm{Vor}|_{\widetilde{\Sigma}_S} v$, $d(v,x) \leq \max\{\mathrm{lfs}(v), \mathrm{lfs}(x)\}$. (This last condition holds if V is a constrained ϵ-sample, as defined in Section 5, for some $\epsilon \leq 1$.)

Then for every site $v \in V$ and every point x in the extended Voronoi cell $\mathrm{Vor}|_{\widetilde{\Sigma}} v$, the relative interior of the line segment xv does not intersect the boundary of a portal.

Proof. Follows from Lemmas 28 (for $x \in \mathrm{Vor}|_{\widetilde{\Sigma}_S} v$) and 30 (for x on an extrusion) below. ◀

Before we prove the two parts of the Shadow Theorem—one part for points in the principal branch, and one for points on extrusions—we give the much shorter proofs of the three corollaries of the Shadow Theorem, introduced in Section 4. We first prove Corollary 5:

Under the conditions of Theorem 4, every extended Voronoi cell is a closed point set (closed with respect to the topological space $\widetilde{\Sigma}$ or $\widetilde{X}$).

Proof. Suppose for the sake of contradiction that there is a site $v \in V$ and a point $x \in \widetilde{\Sigma}$ that lies on the boundary of $\mathrm{Vor}|_{\widetilde{\Sigma}} v$, but $x \notin \mathrm{Vor}|_{\widetilde{\Sigma}} v$. Every point in $\mathrm{Vor}|_{\widetilde{\Sigma}} v$ is visible from v, and the set of points on $\widetilde{\Sigma}$ visible from v is closed (because portal boundaries do not block visibility), so x is visible from v.

The fact that x is visible from v but $x \notin \mathrm{Vor}|_{\widetilde{\Sigma}} v$ implies that there is a site $w \in V$ such that $x \in \mathrm{Vor}|_{\widetilde{\Sigma}} w$, $d(w,x) < d(v,x)$, and x is visible from w. By Theorem 4, the relative

interior of wx does not intersect the boundary of a portal. Therefore, there is an open neighborhood $N \subset \widetilde{\Sigma}$ of x such that every point in N is visible from w. Moreover, there is an open neighborhood $N' \subset \widetilde{\Sigma}$ of x such that for every point $y \in N'$, $d(w,y) < d(v,y)$. No point in $N \cap N'$ can lie in $\mathrm{Vor}|_{\widetilde{\Sigma}} v$, contradicting the supposition that x is on the boundary of $\mathrm{Vor}|_{\widetilde{\Sigma}} v$. By contradiction, every extended Voronoi cell is closed. ◀

We now prove Corollary 6:

Under the conditions of Theorem 4, for every site $v \in V$ and every point x on the boundary of $\mathrm{Vor}|_{\widetilde{\Sigma}} v$, there is a site $w \in V \setminus \{v\}$ such that $x \in \mathrm{Vor}|_{\widetilde{\Sigma}} w$ and $d(v,x) = d(w,x)$.

Proof. Consider a site $v \in V$ and a point $x \in \widetilde{\Sigma}$ on the boundary of $\mathrm{Vor}|_{\widetilde{\Sigma}} v$. By Corollary 5, $x \in \mathrm{Vor}|_{\widetilde{\Sigma}} v$. By Theorem 4, the relative interior of vx does not intersect the boundary of a portal. Therefore, there is an open neighborhood $N \subset \widetilde{\Sigma}$ of x such that every point in N is visible from v. As x is on the boundary of $\mathrm{Vor}|_{\widetilde{\Sigma}} v$, x is also on the boundary of $N \setminus \mathrm{Vor}|_{\widetilde{\Sigma}} v$. Every point in N belongs to some extended Voronoi cell (as every point in N can see at least one site, v), so x is on the boundary of some other extended Voronoi cell $\mathrm{Vor}|_{\widetilde{\Sigma}} w$. By Corollary 5, $x \in \mathrm{Vor}|_{\widetilde{\Sigma}} w$. As x is in both v's and w's cells, $d(v,x) = d(w,x)$. ◀

We now prove Corollary 7:

Under the conditions of Theorem 4, if every connected component of Σ contains at least one site in V, then every point on $\widetilde{\Sigma}$ is in at least one extended Voronoi cell.

Proof. Suppose for the sake of contradiction that some point on a connected component $\widetilde{\sigma}$ of $\widetilde{\Sigma}$ belongs to no extended Voronoi cell. As $\widetilde{\sigma}$ contains at least one site, there is a point $x \in \widetilde{\sigma}$ that is on both the boundary of an extended Voronoi cell $\mathrm{Vor}|_{\widetilde{\Sigma}} v$ and the boundary of the set of points on $\widetilde{\sigma}$ that belong to no cell. Therefore, $x \in \mathrm{Vor}|_{\Sigma} v$ by Corollary 5, but x is on the boundary of the set of points not visible from v. Hence the relative interior of xv intersects a portal boundary, contradicting Theorem 4. ◀

The proof of the Shadow Theorem divides into two parts: we prove it first for the case where x is in the principal branch (the Principal Shadow Lemma, Lemma 28), then for the case where x lies on an extrusion (the Extrusion Shadow Lemma, Lemma 30). Both results use the following brief lemma.

▶ **Lemma 27.** *Let $\Sigma \subset \mathbb{R}^3$ be a closed point set and let M be its medial axis. Consider three points $x \in \mathbb{R}^3$ and $y, z \in \mathbb{R}^3 \setminus M$. Suppose that $\tilde{y} \neq \tilde{z}$ and y lies on the line segment xz. Then $d(x,\tilde{y}) < d(x,\tilde{z})$.*

Proof. Let Π be the plane that bisects the line segment $\tilde{y}\tilde{z}$, illustrated in Figure 15. As $\tilde{y}$ is the unique point nearest y on Σ, $d(y,\tilde{y}) < d(y,\tilde{z})$, so y lies on the same side of Π as $\tilde{y}$. Symmetrically, z lies on the same side of Π as $\tilde{z}$. As $y \in xz$, x lies on the same side of Π as $\tilde{y}$. Therefore, $d(x,\tilde{y}) < d(x,\tilde{z})$. ◀

▶ **Lemma 28** (Principal Shadow Lemma)**.** *Let $\Sigma \subset \mathbb{R}^3$ be a smooth surface without boundary. Let $V \subset \Sigma$ be a finite set of sites. Let $\mathrm{Vor}|_{\widetilde{\Sigma}} V$ be the extended Voronoi diagram (for a suitable S and Z). Consider a site $v \in V$ and a point x in v's principal Voronoi cell $\mathrm{Vor}|_{\overline{\Sigma}_S} v$. If $d(x,v) \leq \mathrm{lfs}(x)$ or $d(x,v) \leq \mathrm{lfs}(v)$, then the relative interior of the line segment xv does not intersect the boundary of a portal.*

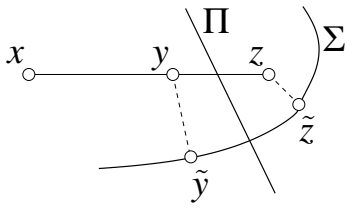

Figure 15 As y lies on xz, $d(x, \tilde{y}) < d(x, \tilde{z})$.

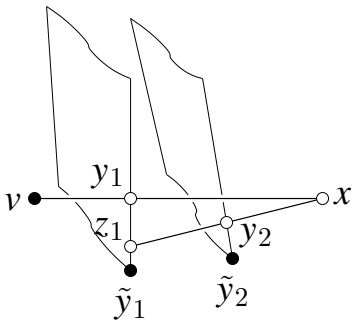

Figure 16 If vx intersects the boundary of a portal, there is another site $\tilde{y}_1$ that is closer to x than v is. If $\tilde{y}_1$ is not visible from x, there is another site $\tilde{y}_2$ that is closer to x than $\tilde{y}_1$ is.

Proof. Suppose for the sake of contradiction that the relative interior of xv intersects one or more portal boundaries. Let y_1 be the intersection point closest to x. Observe that x is visible from y_1 (because x is visible from v). By assumption, either $d(x, v) \leq \text{lfs}(x)$ or $d(x, v) \leq \text{lfs}(v)$, so either $d(x, y_1) < d(x, v) \leq \text{lfs}(x)$ or $d(y_1, v) < d(x, v) \leq \text{lfs}(v)$. In either case, it follows from the definition of lfs that y_1 does not lie on the medial axis. But y_1 lies on a portal boundary, so y_1 must lie on the normal segment of a site in V, and that site is located at the point $\tilde{y}_1 \neq v$. By Lemma 27, $d(x, \tilde{y}_1) < d(x, v)$.

Although $d(x, \tilde{y}_1) < d(x, v)$, x is in v's principal Voronoi cell rather than $\tilde{y}_1$'s cell, so x is not visible from the site at $\tilde{y}_1$, even though x is visible from y_1. Let z_1 be the point nearest $\tilde{y}_1$ on the line segment $y_1\tilde{y}_1$ that can see x, as illustrated in Figure 16. Observe that $\tilde{z}_1 = \tilde{y}_1$ (all three points y_1, $\tilde{y}_1$, and z_1 lie on the normal segment of $\tilde{y}_1$). The relative interior of the line segment xz_1 intersects one or more portal boundaries (because as you slide from y_1 to $\tilde{y}_1$, z_1 is the last point that can see x). Let y_2 be the intersection point closest to x, as illustrated. Then x is visible from y_2, and y_2 lies on the boundary of a portal.

Let S be the smallest sphere that passes through v and x (thus the line segment vx is a diameter of S). Let P be a set containing v, x, and every site strictly inside S. Let H be the convex hull of P. Observe that the diameter of S is at most $\max\{\text{lfs}(v), \text{lfs}(x)\}$ and therefore no medial axis point is strictly inside S. By Lemma 25, $\tilde{y}_1$ is strictly inside S, hence $\tilde{y}_1 \in P$ and $\tilde{y}_1 \in H \setminus S$, hence $z_1 \in H \setminus S$, hence $y_2 \in H \setminus S$. Therefore, y_2 does not lie on the medial axis.[1] But y_2 lies on a portal boundary, so y_2 must lie on the normal segment of a site in V, and that site is located at the point $\tilde{y}_2 \neq \tilde{z}_1$. By Lemma 27, $d(x, \tilde{y}_2) < d(x, \tilde{z}_1)$; hence $d(x, \tilde{y}_2) < d(x, \tilde{y}_1) < d(x, v)$.

[1] In the case where $d(x, v) \leq \text{lfs}(x)$, we can show this in a simpler way without the need for Lemma 25: $d(x, y_2) < d(x, z_1) \leq \max\{d(x, y_1), d(x, \tilde{y}_1)\} < d(x, v) \leq \text{lfs}(x)$, so y_2 does not lie on the medial axis. However, we seem to need Lemma 25 for the case where $d(x, v) \leq \text{lfs}(v)$.

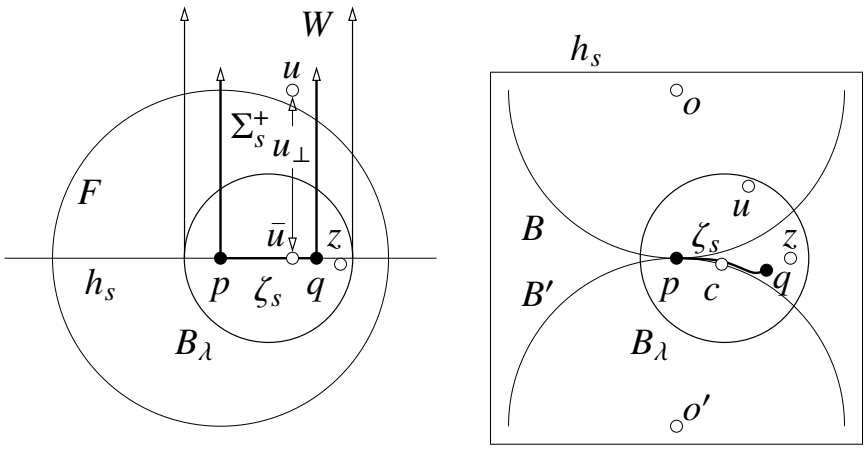

Figure 17 Left: a view from which the surface Σ and extrusion Σ_s^+ lie approximately in the page, and the cutting plane h_s is orthogonal to the page. The point u is in the wiener W but not in the ball F. Right: a view from which the cutting plane h_s lies in the page, and Σ_s^+ is extruded away from the reader. (The points c and u do not necessarily lie on Σ, but z does.)

As $x \in \mathrm{Vor}|_{\overline{\Sigma}_S} v$ (rather than $x \in \mathrm{Vor}|_{\overline{\Sigma}_S} \tilde{y}_2$) x is not visible from the site at $\tilde{y}_2$, even though x is visible from y_2. By inductively repeating the argument we can construct an infinite sequence of sites not visible from x such that each successive site is closer to x than the previous site. But V contains finitely many sites. The result follows by contradiction. ◀

Next, we prove another Shadow Lemma for any point x on an extrusion. The proof of the Extrusion Shadow Lemma is almost the same as the proof of the Principal Shadow Lemma, but there is one major complication: both proofs use the fact that no point in the sequence $y_1, y_2, \ldots$ lies on the medial axis. This fact is more difficult to prove when x lies on an extrusion (rather than on Σ).

Consider a segment $s \in S$ with endpoints $p, q \in V$ such that $d(p,q) \leq \rho\,\mathrm{lfs}(p)$ for some $\rho \leq 0.47$. Let F be the open ball with center p and radius $\mathrm{lfs}(p)$. By the definition of lfs, F does not intersect the medial axis of Σ. Let c be the midpoint of s. Let B_λ be the closed ball with center c and radius $\lambda\,\mathrm{lfs}(p)$, where λ is the function of ρ defined in Theorem 3. Observe that $B_\lambda \subset F$, as the distance from p to any point in B_λ is at most $(\rho/2 + \lambda)\,\mathrm{lfs}(p) < 0.705\,\mathrm{lfs}(p)$.

Let h_s be the cutting plane for s, let $\zeta_s \subset h_s \cap \Sigma$ be the portal curve for s, and let b_s be a unit vector normal to h_s. Let Σ_s^+ be the extrusion of s's portal curve in the direction b_s. Recall that by Theorem 3, every site whose extended Voronoi cell intersects Σ_s^+ must lie in B_λ.

Let W be the set of points $\{x + \omega b_s : x \in B_\lambda \text{ and } \omega \geq 0\}$; that is, every point that is in B_λ or in the direction b_s from a point in B_λ. The set W has the shape of a wiener that is infinite in one direction, as illustrated in Figure 17 (left). The premise of W is that it is a convex point set that encloses both the ball B_λ and the extrusion Σ_s^+. However, W is defined with respect to the space $\mathbb{R}^3$ whereas Σ_s^+ is embedded in a secondary branch of $\widetilde{X}$. If Σ_s^+ were in $\mathbb{R}^3$ (but the point coordinates were unchanged), we could write $\Sigma_s^+ \subset W$.

▶ **Lemma 29.** *Define s, p, q, F, B_λ, h_s, and W as above, with $d(p,q) \leq \rho\,\mathrm{lfs}(p)$ for some $\rho \leq 0.47$. Let h_s^- be the closed halfspace on the side of h_s opposite from Σ_s^+. Consider a point $z \in \Sigma \cap B_\lambda \cap h_s^-$ and let ℓ_z be its normal segment. Then $W \cap \ell_z \subset F$.*

Proof. Suppose for the sake of contradiction that there is a point $u \in W \cap \ell_z$ such that $u \notin F$, as illustrated in Figure 17 (left). Then $d(p,u) \geq \text{lfs}(p)$ and $u \notin B_\lambda$. Let B_u be the open ball with center u whose boundary passes through z. As u lies on z's normal segment, B_u does not intersect Σ.

Let B and B' be the two open balls of radius $\text{lfs}(p)$ tangent to Σ at p, as illustrated in Figure 17 (right); neither ball intersects Σ. As Σ is a surface without boundary that passes through p, Σ partitions $\mathbb{R}^3$ into two disjoint components (by the Jordan–Brouwer Separation Theorem), with B included in one and B' in the other. One component must include B_u too; choose the labels B and B' so that B' is in the same component as B_u. It follows that B_u is disjoint from B.

Let o and o' be the centers of B and B', respectively. As B_u and B are disjoint, $d(u,o) \geq d(u,z) + \text{lfs}(p)$. Let $u_\perp$ denote the distance from u to the cutting plane h_s. Whereas $z \in h_s^-$, $u \in W \setminus F$ must lie on the positive side of h_s (opposite from h_s^-). Therefore, $d(u,z) \geq u_\perp$ and $d(u,o) \geq u_\perp + \text{lfs}(p)$.

As the cutting plane h_s is parallel to p's normal vector, $o, o' \in h_s$. Let $\bar{u}$ be the orthogonal projection of u onto h_s. By Pythagoras' Theorem, $d(u,o)^2 = u_\perp^2 + d(\bar{u},o)^2$. Combining this with the last inequality gives $d(\bar{u},o)^2 = d(u,o)^2 - u_\perp^2 \geq 2u_\perp \text{lfs}(p) + \text{lfs}(p)^2$. Therefore, $u_\perp \leq (d(\bar{u},o)^2/\text{lfs}(p) - \text{lfs}(p))/2$.

The fact that $u \in W$ implies that $\bar{u} \in B_\lambda$, so we can write $d(\bar{u},o) \leq d(\bar{u},c) + d(c,o) \leq \lambda \text{lfs}(p) + d(c,o)$. To find an upper bound for $d(c,o)$, consider a coordinate system that places the site p at the origin, the cutting plane h_s in the x-y plane, the point o at the coordinate $(0, \text{lfs}(p), 0)$, and the point o' at the coordinate $(0, -\text{lfs}(p), 0)$, as illustrated. Then we write $q = (q_x, q_y, 0)$ and $c = (q_x/2, q_y/2, 0)$. With this notation, $d(q,o')^2 = q_x^2 + (q_y + \text{lfs}(p))^2 = \|q\|^2 + 2\text{lfs}(p)q_y + \text{lfs}(p)^2$ and $d(c,o)^2 = (q_x/2)^2 + (q_y/2 - \text{lfs}(p))^2 = \|q\|^2/4 - \text{lfs}(p)q_y + \text{lfs}(p)^2$. Adding half the first equation to the second (to eliminate q_y) gives $d(q,o')^2/2 + d(c,o)^2 = 3\|q\|^2/4 + 3\,\text{lfs}(p)^2/2$. Observe that $\|q\| = d(p,q) \leq \rho\,\text{lfs}(p)$. As $q \notin B'$ (because $q \in \Sigma$), $d(q,o') \geq \text{lfs}(p)$. Hence $d(c,o)^2 \leq (3\rho^2/4 + 1)\,\text{lfs}(p)^2$.

Therefore, $d(\bar{u},o) \leq (\lambda + \sqrt{3\rho^2/4 + 1})\,\text{lfs}(p) < 1.55\,\text{lfs}(p)$ and $u_\perp < 0.702\,\text{lfs}(p)$. By Pythagoras' Theorem, $d(p,u)^2 = d(p,\bar{u})^2 + u_\perp^2 \leq (d(p,c) + d(c,\bar{u}))^2 + u_\perp^2 < (\rho\,\text{lfs}(p)/2 + \lambda\,\text{lfs}(p))^2 + 0.702^2\,\text{lfs}(p)^2 < 0.99\,\text{lfs}(p)^2$. But this contradicts the fact that $u \notin F$. Hence there is no point $u \in W \cap \ell_z$ such that $u \notin F$. ◀

▶ **Lemma 30** (Extrusion Shadow Lemma). *Let $\Sigma \subset \mathbb{R}^3$ be a smooth surface without boundary. Let $V \subset \Sigma$ be a finite set of sites. Let $s \in S$ be a segment with endpoints $p, q \in V$, and suppose that $d(p,q) \leq 0.47\,\text{lfs}(p)$. Consider a site $v \in V$ and a point x in v's extended Voronoi cell $\text{Vor}|_{\widetilde{\Sigma}} v$ such that x lies on the extrusion Σ_s^+ (or Σ_s^-). Consider the line segment $xv \subset \widetilde{X}$. The relative interior of xv does not intersect the boundary of a portal (including the boundary of P_s).*

To clarify the interpretation of Lemma 30, note that xv lies partly in a secondary branch and partly in the principal branch. The portion of xv that is solely in the secondary branch cannot intersect any portal (regardless of matching point coordinates). The lemma focuses on the portion in the principal branch.

Proof. The proof is identical to that of Lemma 28, except that we employ an entirely different argument to establish that no point in the sequence $y_1, y_2, \ldots$ lies on the medial axis.

If v is a vertex of s then the result follows immediately, so assume that $v \in V \setminus \{p,q\}$. Define F, B_λ, h_s, ζ_s, and W as in the preamble before Lemma 29. Let h_s^- be the closed halfspace on the side of h_s opposite from Σ_s^+. By Theorem 1, $v \in h_s^-$. By Theorem 3, $v \in B_\lambda \subset W \cap F$.

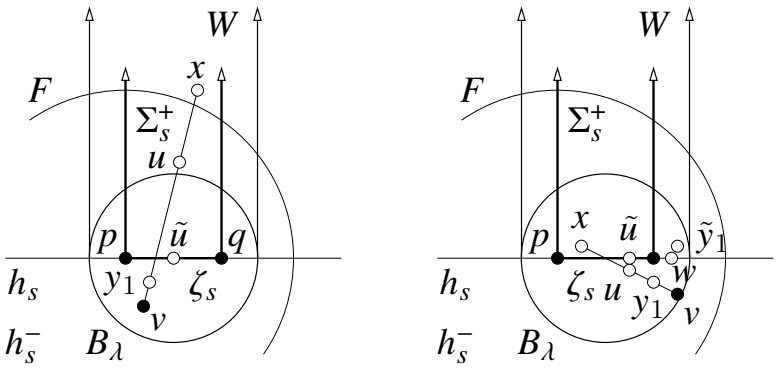

Figure 18 Two views from which the surface Σ and extrusion Σ_s^+ lie approximately in the page, and the cutting plane h_s is orthogonal to the page. Left: The line segment xv passes through the portal at u and intersects a portal boundary at y_1. We see that $u, y_1 \in W$ and $u, y_1 \in F$. Right: the circumstance $\tilde{y}_1 \notin h_s^-$, depicted here, cannot happen.

Although the wiener W is defined in the Euclidean space $\mathbb{R}^3$, we say that W *encloses* a point set A if for every point $a \in A$ there is a point in W with the same coordinates as a, regardless of whether a is in the principal branch, a secondary branch, or $\mathbb{R}^3$. Observe that W encloses Σ_s^+ and B_λ. As $x \in \Sigma_s^+$, $v \in B_\lambda$, and W is convex, it follows that W encloses xv.

Suppose for the sake of contradicting the lemma that the relative interior of xv intersects one or more portal boundaries. Let y_1 be the intersection point closest to x. As y_1 is in the principal branch, the line segment xy_1 passes through the portal P_s at one or more points; let u be one of them, as illustrated in Figure 18 (left). Then $y_1 \in uv$. Observe that W encloses y_1 and u, and u lies on the normal segment $\ell_{\tilde{u}}$ of a point $\tilde{u}$ on the portal curve ζ_s. By Lemma 29, $W \cap \ell_{\tilde{u}} \subset F$, so $u \in F$. Recall that F is a medial-free open ball and $v \in F$, so F encloses uv and y_1. This confirms that y_1 is not a medial axis point.

As y_1 is on a portal boundary but not on the medial axis, there is a site at $\tilde{y}_1$. We claim that $\tilde{y}_1 \in h_s^-$; suppose for the sake of contradiction that $\tilde{y}_1 \notin h_s^-$, as illustrated in Figure 18 (right). Observe that as uv does not intersect the medial axis, the nearest point map ν is continuous over uv and $\nu(uv)$ is a connected path on Σ. As $\tilde{y}_1$ is in the principal branch and lies on $\nu(uv)$, the path $\nu(uy_1)$ must somewhere (on the way from $\tilde{u} \in \zeta_s$ to $\tilde{y}_1$) exit h_s^- without entering the portal P_s before reaching $\tilde{y}_1$. Let $w \in \nu(uy_1) \cap h_s$ be a point where the path exits h_s^- without entering a secondary branch, as illustrated. Then $w \in \Sigma \cap h_s$ but w is not in the relative interior of the portal curve ζ_s. Let $\bar{x}$ be the orthogonal projection of x onto h_s, and recall that $\bar{x} \in \zeta_s$. As w is not in the relative interior of ζ_s, $d(\bar{x}, w) \geq \min\{d(\bar{x}, p), d(\bar{x}, q)\}$. As $x\bar{x}$ is orthogonal to h_s and $\bar{x}, w, p, q \in h_s$, we also have $d(x, w) \geq \min\{d(x, p), d(x, q)\}$. But by Lemma 27, $d(x, w) < d(x, v)$; and as $x \in \mathrm{Vor}|_{\tilde{\Sigma}} v$, $d(x, w) < d(x, v) \leq \min\{d(x, p), d(x, q)\}$. This is a contradiction; hence $\tilde{y}_1 \in h_s^-$ as claimed.

As $\tilde{y}_1 \in h_s^-$, by Theorem 3, $\tilde{y}_1 \in B_\lambda \subset W \cap F$. As $y_1 \in W \cap F$, $\tilde{y}_1 \in W \cap F$, and W and F are convex, it follows that the point $z_1 \in y_1\tilde{y}_1$ discussed in the proof of Lemma 28 is also in $W \cap F$. Recall the point $y_2 \in xz_1$ defined in the proof of Lemma 28. By repeating the argument of the previous two paragraphs with v replaced by z_1, we show that $y_2 \in W \cap F$, and hence y_2 is not a medial axis point. We repeat the argument inductively for y_3, y_4, etc. The rest of the proof proceeds as the proof of Lemma 28. ◀

G The Nearest Point Map Is a Homeomorphism

Here we prove that for a sufficiently dense sample V of a smooth surface $\Sigma \subset \mathbb{R}^3$ without boundary, with suitable conditions on segment length and encroachment, the nearest point map (restricted to the restricted CDT) is a homeomorphism from the underlying space of the restricted CDT $\mathrm{Del}|_{\widetilde{\Sigma}} V$ to Σ. (Recall that the nearest point map ν maps any point $x \in \mathbb{R}^3 \setminus M$ to the point $\nu(x)$ nearest x on Σ. We use the abbreviation $\tilde{x}$ to denote $\nu(x)$.) Specifically, we suppose that V is a constrained 0.3202-sample of (Σ, S, Z), that for every segment $pq \in S$, $d(p,q) \leq 0.47\,\mathrm{lfs}(p)$, and that the encroachment condition described in Section 5 holds. We note that in the unconstrained case (i.e., for restricted Delaunay triangulations), our sampling constant is substantially better than those in the classical proofs from the literature on provably good surface reconstruction [15]: we prove homeomorphism for a 0.3202-sample instead of merely for a 0.18-sample. This reduces the number of sample points required by a factor of about 3.16 (the square of 0.32/0.18).

Unlike the classical proofs that the restricted Delaunay triangulation is homeomorphic to the underlying manifold, our proof does not use the Topological Ball Theorem of Edelsbrunner and Shah [17]. The Topological Ball Theorem cannot be applied to the restricted *constrained* Delaunay triangulation because it depends on the barycentric subdivision of the Delaunay triangulation in $\mathbb{R}^3$, but we know no subdivision of space into a three-dimensional triangulation that conforms to a restricted CDT. That is why we use the nearest point map to prove a homeomorphism, as Boissonnat and Ghosh [6] do. (The homeomorphism that Edelsbrunner and Shah use is not the nearest point map; it is a different map based on the barycentric subdivision of the restricted Delaunay triangles.)

Amenta et al. [3] also use the nearest point map to prove that their Cocone Algorithm produces a triangulation homeomorphic to Σ; but their proof relies on the assumption that the RDT is homeomorphic to Σ, as proven by Amenta and Bern [1] with the Topological Ball Theorem. Hence our proof faces some hurdles that Amenta et al. could avoid.

Our main new idea is a direct proof that, under the right sampling conditions, each principal Voronoi cell is a "star-shaped" topological disk (Lemma 35 and Theorem 40). Another interesting part of this result is a proof that, under the right conditions, the nearest point map over any single restricted Delaunay triangle is an orientation-preserving homeomorphism from the triangle to its image on Σ (Theorem 46).

For the sake of brevity, we use the notation $|pq|$ to denote $d(p,q)$ throughout this section.

G.1 Relationships between Surface Points and Nearby Tangent Planes

In this section we establish a relationship between a point $v \in \Sigma$ and the tangent planes of points on Σ close to v; a relationship between v and the tangent planes of points lying on extrusions close to v (even if the points themselves are far from v); and a relationship between neighborhoods of those points and the tangent plane at v. These relationships prepare us to prove the main theorem of Section G.2: that under suitable sampling conditions, every extended Voronoi cell is a topological disk.

Recall that for a point $x \in \Sigma$, the tangent plane $T_x\Sigma \subset \mathbb{R}^3$ is the plane tangent to Σ at x, and the normal vector n_x is the outside-facing vector orthogonal to Σ and to $T_x\Sigma$ at x. In the same manner, we define tangent planes and normal vectors for extrusions; e.g., $T_y\Sigma_s^+$ is the plane tangent to the extrusion Σ_s^+ at a point $y \in \Sigma_s^+$. Each extrusion Σ_s^+ inherits an orientation from Σ through gluing, so we define an "outside-facing" normal vector n_y, even though Σ_s^+ itself does not bound a volume like Σ does. The extended surface $\widetilde{\Sigma}$ is not generally smooth where an extrusion meets the principal surface $\overline{\Sigma}_S$, so a point on a portal

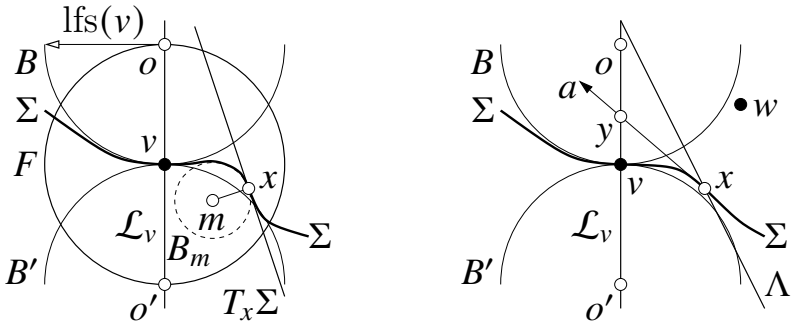

Figure 19 Left: for two sufficiently close points $v, x \in \Sigma$, suppose that the tangent plane $T_x\Sigma$ does not intersect oo', as shown. This leads to a contradiction; hence $T_x\Sigma$ must intersect oo'. The medial ball B_m is tangent to Σ at x. The center m of B_m cannot lie in the open ball F. In this figure, the points v, x, o, and o' lie on the plane of the page, but m and w generally do not; imagine m floating above the page. The dashed circle shows the page's cross section of B_m, but B_m is larger. The surface Σ cannot intersect the open balls B, B', and B_m, so B and B_m are disjoint. Right: the plane Λ bisects vw. The ray $a \in T_x\Sigma$ intersects the relative interior of oo' and is strictly on the same side of Λ as v.

curve can have two tangent planes and two normal vectors. (A site can have many, as it can adjoin many extrusions, but we will only care about a site's tangent plane and normal vector with respect to the original surface Σ.)

In this and subsequent sections, let φ be the continuous map that orthogonally projects $\mathbb{R}^3$ or $\widetilde{X}$ onto $T_v\Sigma$ (with the projection direction being parallel to n_v).

The following lemma is the key that will give us the power to prove that theorem for relatively coarse point samples. (To understand why, it is helpful to know that the line segment oo' discussed in the proof is in v's three-dimensional Voronoi cell Vor v; see also Lemma 34.)

▸ **Lemma 31.** *Consider two points $v, x \in \Sigma$ such that $|vx| < \xi\,\mathrm{lfs}(v)$, where $\xi = \sqrt{(\sqrt{5}-1)/2} \doteq 0.786151$. Let B and B' be the two open balls of radius $\mathrm{lfs}(v)$ tangent to Σ at v, and let o and o' be their centers, respectively. Let $T_x\Sigma$ be the plane tangent to Σ at x. Then the ball centers o and o' lie on strictly opposite sides of $T_x\Sigma$.*

Proof. Suppose for the sake of contradiction that o and o' do not lie on strictly opposite sides of $T_x\Sigma$. Then $T_x\Sigma \neq T_v\Sigma$ and $x \neq v$. Moreover, either $oo' \subset T_x\Sigma$ or $T_x\Sigma$ does not intersect the relative interior of oo'.

Let B_m be the open medial ball tangent to Σ at x that is on the same side of $T_x\Sigma$ as v (an arbitrary side if $v \in T_x\Sigma$), as illustrated in Figure 19. As $B_m \cap \Sigma = \emptyset$, if B_m is an open halfspace then $v \in T_x\Sigma$ and $T_x\Sigma = T_v\Sigma$ and the lemma follows; so assume B_m is bounded. Its center m lies on the medial axis of Σ. The line segment xm is perpendicular to $T_x\Sigma$.

Observe that v is in the relative interior of oo'. If $oo' \subset T_x\Sigma$, then $\angle vxm = \angle oxm = \angle o'xm = 90°$. Otherwise, $T_x\Sigma$ does not intersect the relative interior of oo', so $\angle vxm < 90°$, $\angle oxm \leq 90°$, and $\angle o'xm \leq 90°$ (the last two because v cannot be on the side of $T_x\Sigma$ opposite from o or o'). By Pythagoras' Theorem, $|ox|^2 + |mx|^2 \geq |om|^2$ and $|vx|^2 + |mx|^2 \geq |vm|^2$. As m lies on the medial axis, $|vm| \geq \mathrm{lfs}(v)$ (by the definition of lfs), so $|vx|^2 + |mx|^2 \geq \mathrm{lfs}(v)^2$.

The surface Σ intersects none of the open balls B, B', or B_m, but it passes between B and B' at v. As Σ has no boundary and divides space into two pieces, one containing B and one containing B', the ball B_m must lie in one of those two pieces. Choose the labels B and B' so

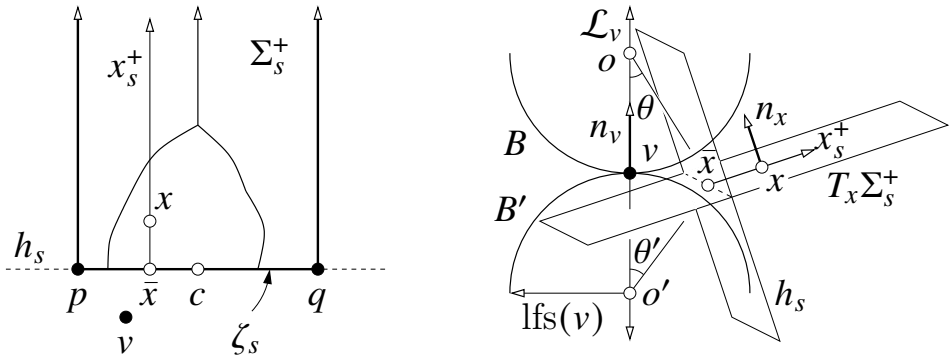

Figure 20 The point x lies on a ray x_s^+, which is orthogonal to the cutting plane h_s and is a subset of both the extrusion Σ_s^+ and the plane $T_x\Sigma_s$ tangent to Σ_s^+ at x. Left: view from above Σ_s^+; h_s is orthogonal to the page. Right: view with n_v vertical.

that B_m lies in the same piece as B', as illustrated; then B_m must be disjoint from B. The radii of B and B_m are $\mathrm{lfs}(v)$ and $|mx|$ respectively, so $|om| \geq \mathrm{lfs}(v) + |mx|$. Combining this with the inequality $|ox|^2 + |mx|^2 \geq |om|^2$ gives $|ox|^2 \geq \mathrm{lfs}(v)^2 + 2\,\mathrm{lfs}(v)\,|mx|$. Combining this with the inequality $|mx|^2 \geq \mathrm{lfs}(v)^2 - |vx|^2$ gives $|ox|^2 \geq \mathrm{lfs}(v)^2 + 2\,\mathrm{lfs}(v)\sqrt{\mathrm{lfs}(v)^2 - |vx|^2}$.

Create a coordinate system with v at the origin such that x_v is the coordinate of x in the direction parallel to the line $\mathcal{L}_v$ through o', v, and o (the vertical axis in Figure 19) and x_h is the distance from x to $\mathcal{L}_v$ (the horizontal axis in Figure 19). Then $|ox|^2 + |o'x|^2 = x_h^2 + (x_v - \mathrm{lfs}(v))^2 + x_h^2 + (x_v + \mathrm{lfs}(v))^2 = 2x_h^2 + 2x_v^2 + 2\,\mathrm{lfs}(v)^2 = 2\,|vx|^2 + 2\,\mathrm{lfs}(v)^2$. As $x \notin B'$, $|o'x|^2 \geq \mathrm{lfs}(v)^2$. Combining these with the inequality $|ox|^2 \geq \mathrm{lfs}(v)^2 + 2\,\mathrm{lfs}(v)\sqrt{\mathrm{lfs}(v)^2 - |vx|^2}$ gives $|vx|^2 = (|ox|^2 + |o'x|^2 - 2\,\mathrm{lfs}(v)^2)/2 \geq \mathrm{lfs}(v)\sqrt{\mathrm{lfs}(v)^2 - |vx|^2}$.

As $|vx| < \xi\,\mathrm{lfs}(v)$, we have $\xi^2 > |vx|^2/\mathrm{lfs}(v)^2 \geq \sqrt{1 - |vx|^2/\mathrm{lfs}(v)^2} > \sqrt{1-\xi^2}$, which is equivalent to $\xi^4 + \xi^2 - 1 > 0$, hence $\xi > \sqrt{(\sqrt{5}-1)/2}$. The result follows by contradiction. ◀

The next lemma proves a result similar to Lemma 31 for a point x lying on an extrusion.

▶ **Lemma 32.** *Consider a segment s with endpoints $p, q \in \Sigma$ such that $d(p,q) \leq \rho\,\mathrm{lfs}(p)$ for $\rho \leq 0.3647$. Let $v \in \Sigma$ be a site whose extended Voronoi cell $\mathrm{Vor}|_{\widetilde{\Sigma}} v$ includes a point on Σ_s^+ (or Σ_s^-). Let $B, B' \subset \mathbb{R}^3$ be the two open balls of radius $\mathrm{lfs}(v)$ tangent to Σ at v, and let $o, o' \in \mathbb{R}^3$ be their centers, respectively. Let x be any point on Σ_s^+ (respectively, Σ_s^-). Let $T_x\Sigma_s \subset \mathbb{R}^3$ be the plane tangent to Σ_s^+ (or Σ_s^-) at x. Let n_x be a vector normal to Σ_s^+ (or Σ_s^-) at x; thus n_x is normal to $T_x\Sigma_s$ as illustrated in Figure 20 (right). Then the ball centers o and o' lie on opposite sides of $T_x\Sigma_s$. Moreover, $\angle(n_v, n_x) < 51°$.*

Proof. Let h_s be the cutting plane for s and let $\zeta_s \subset h_s \cap \Sigma$ be the portal curve for s. Let $\bar{x} \in \zeta_s$ be the orthogonal projection of x onto h_s. The ray x_s^+ originating at $\bar{x}$ and passing through x is a subset of the extrusion Σ_s^+ and a subset of $T_x\Sigma_s$, as illustrated in Figure 20.

Let c be the midpoint of the segment s. By Theorem 3, $d(v,c) \leq \lambda\,\mathrm{lfs}(p)$, where λ is defined in Theorem 3 and satisfies $\lambda < 0.2184$ for $\rho \leq 0.3647$. Hence $d(v,p) \leq d(v,c) + d(c,p) \leq (\lambda + \rho/2)\,\mathrm{lfs}(p)$. Likewise, $d(v,\bar{x}) \leq (\lambda + \rho/2)\,\mathrm{lfs}(p)$. Observe that $d(p,\bar{x}) \leq d(p,q) \leq \rho\,\mathrm{lfs}(p)$. By the Feature Translation Lemma (Lemma 12), $\mathrm{lfs}(p) \leq \mathrm{lfs}(v)/(1 - \lambda - \rho/2)$ and $\mathrm{lfs}(p) \leq \mathrm{lfs}(\bar{x})/(1-\rho)$. Therefore, $d(v,\bar{x}) \leq (\lambda + \rho/2)\,\mathrm{lfs}(v)/(1 - \lambda - \rho/2)$ and $d(v,\bar{x}) \leq (\lambda+\rho/2)\,\mathrm{lfs}(\bar{x})/(1-\rho)$. By the Normal Variation Lemma (Lemma 13), $\angle(n_v, n_p) \leq \eta(\lambda+\rho/2)$ and $\angle(n_v, n_{\bar{x}}) \leq \eta((\lambda+\rho/2)/(1-\rho))$.

To establish an upper bound on $\angle(n_v, n_x)$, we divide the angle into portions transverse to h_s and orthogonal to h_s. The former portion is equal to the transverse portion of $\angle(n_v, n_{\bar{x}})$. The latter portion is equal to the orthogonal portion of $\angle(n_v, n_p)$. Therefore, $\tan^2 \angle(n_v, n_x) \leq \tan^2 \angle(n_v, n_{\bar{x}}) + \tan^2 \angle(n_v, n_p) \leq \tan^2 \eta((\lambda+\rho/2)/(1-\rho)) + \tan^2 \eta(\lambda+\rho/2)$. For $\rho \leq 0.3647$, this implies that $\angle(n_v, n_x) < 51°$.

The upper bound on $d(v, \bar{x})$ given above implies an upper bound on the angles $\theta = \angle vo\bar{x}$ and $\theta' = \angle vo'\bar{x}$, illustrated in Figure 20 (right). Specifically, as $v, \bar{x} \notin B$, we have $d(v, \bar{x}) \geq 2\,\mathrm{lfs}(v)\sin(\theta/2)$, so $\theta \leq 2\arcsin(d(v,\bar{x})/(2\,\mathrm{lfs}(v))) \leq 2\arcsin((\lambda+\rho/2)/(2-2\lambda-\rho))$. Symmetrically, as $v, \bar{x} \notin B'$, the same inequality holds with θ replaced by θ'. Recall that $T_x\Sigma_s$ passes through $\bar{x}$ with normal vector n_x, and that n_v is parallel to vo and vo'. Therefore, if $\angle(n_v, n_x) < 90° - 2\arcsin(d(v,\bar{x})/(2\,\mathrm{lfs}(v)))$, then o and o' lie on opposite sides of $T_x\Sigma_s$. This inequality must hold if $\eta((\lambda+\rho/2)/(1-\rho)) + \eta(\lambda+\rho/2) < 90° - 2\arcsin((\lambda+\rho/2)/(2-2\lambda-\rho))$, and the latter inequality holds for all $\rho \in [0, 0.3647]$, so o and o' lie on opposite sides of $T_x\Sigma_s$. ◀

▶ **Lemma 33.** *Consider a site $v \in V$ and a point $y \in \mathrm{Vor}|_{\widetilde{\Sigma}} v \setminus \{v\}$. where y does not lie on one of the rays (of the form v_s^+ or v_s^-) that bound an extrusion. If y is in the principal branch, suppose that $|vy| \leq 0.9101\,\mathrm{lfs}(v)$. If y lies on an extrusion Σ_s^+ or Σ_s^- of a segment $s = pq$, suppose that $|pq| \leq 0.3647\,\mathrm{lfs}(p)$. If y lies on a portal curve, suppose that both conditions hold. Let φ be the function that orthogonally projects points onto v's tangent plane $T_v\Sigma$.*

Then there exists an open neighborhood $N \subset \widetilde{\Sigma}$ of y such that $\varphi|_N$ is a homeomorphism from N to its image $\varphi(N)$ on $T_v\Sigma$.

Proof. Consider three cases: y lies in the principal branch but not on a portal curve; y lies on an extrusion but not on a portal curve; or y lies on a portal curve. (Under no circumstance can y be a site.)

In the first case, as $|vy| \leq 0.9101\,\mathrm{lfs}(v)$, by the Normal Variation Lemma (Lemma 13), $\angle(n_v, n_y) < 90°$ where n_v and n_y are outside-facing vectors normal to Σ at v and y, respectively. Hence n_y is not parallel to $T_v\Sigma$. It follows from the smoothness of Σ that if N is sufficiently small, $\varphi|_N$ is injective.

In the second case, let $s = pq$ be the segment such that $y \in \Sigma_s^+$ or $y \in \Sigma_s^-$. As $d(p,q) \leq 0.3647\,\mathrm{lfs}(p)$, by Lemma 32, $\angle(n_v, n_y) < 51°$ where n_y is an "outward"-facing vector normal to the extrusion at y. Again, it follows from the smoothness of the extrusion that if we choose N to be sufficiently small, $\varphi|_N$ is injective.

In the third case, as y lies on a portal curve, $\widetilde{\Sigma}$ is not generally smooth at y and there are two normal vectors at y; call them n_y for Σ and n_y^+ for Σ_s^+. Then $\angle(n_v, n_y) \leq 90°$ and $\angle(n_v, n_y^+) < 51°$. The smooth portal curve $\zeta_s \cap N$ cuts the neighborhood N into two pieces, $\overline{\Sigma}_S \cap N$ and $\Sigma_s^+ \cap N$ (or $\Sigma_s^- \cap N$). If N is sufficiently small, φ is injective over each piece due to smoothness. As the vectors n_y and n_y^+ are directed to the same side of $T_v\Sigma$, φ preserves the orientation of the manifold N at all points, so $\varphi(\overline{\Sigma}_S \cap N)$ and $\varphi(\Sigma_s^+ \cap N)$ adjoin opposite sides of the curve $\varphi(\zeta_s \cap N)$ in the plane $T_v\Sigma$; hence if N is sufficiently small, $\varphi(\overline{\Sigma}_S \cap N)$ does not overlap $\varphi(\Sigma_s^+ \cap N)$ and $\varphi|_N$ is injective.

As $\varphi|_N$ is injective and $\varphi|_N$ and its inverse are continuous, $\varphi|_N$ is a homeomorphism from N to $\varphi(N)$. ◀

G.2 Extended Voronoi Cells Are Topological Disks

This section investigates sampling conditions that guarantee that every extended Voronoi cell has the topology of a closed disk. Assuming that a suitable condition on segment lengths is

satisfied, the forthcoming Corollary 42 shows that a constrained 0.44-sample suffices, whereas Theorem 40 gives a sampling condition more suitable for mesh generation algorithms: every point in the principal Voronoi cell of a site v must be within a distance of $0.786\,\mathrm{lfs}(v)$ from v.

Consider a site $v \in V$ and its extended Voronoi cell $\mathrm{Vor}|_{\widetilde{\Sigma}} v$. Let φ be the map that orthogonally projects points onto v's tangent plane $T_v\Sigma$. Note that $\varphi(v) = v$.

We define a *radial path* to be a closed topological interval (i.e., a topological 1-ball) $\gamma \subset \mathrm{Vor}|_{\widetilde{\Sigma}} v$ such that

- one endpoint of γ is the site v,
- the other endpoint—call it z—lies on the boundary of $\mathrm{Vor}|_{\widetilde{\Sigma}} v$ (and may be a point at infinity),
- no point on $\gamma \setminus \{z\}$ lies on the boundary of $\mathrm{Vor}|_{\widetilde{\Sigma}} v$, and
- $\varphi|_\gamma$ is a homeomorphism from γ to a line segment (or a ray) on $T_v\Sigma$ with endpoints v and $\varphi(z)$.

A radial path is not necessarily smooth, because it can contain points where it makes a transition from the principal branch to an extrusion or vice versa; but a radial path is piecewise smooth. One radial path can contain an arbitrary number of these transitions—a fact that complicates proving that extended Voronoi cells are topological disks.

As a special case that deviates slightly from the fourth criterion above, we declare that each extrusion boundary ray v_s^+ is a radial path in $\mathrm{Vor}|_{\widetilde{\Sigma}} v$. Recall that topologically, $v_s^+ = v_s^-$, because we have glued the two rays together; so we have one topological ray with two embeddings, which suggests that φ has two values for each point on v_s^+. For our purposes, we simply declare that $v_s^+ = v_s^-$ is one radial path. Boundary extrusion rays are the only radial paths that have an endpoint at infinity.

Below, we show that under suitable sampling conditions, every point in $\mathrm{Vor}|_{\widetilde{\Sigma}} v$ lies on exactly one radial path, except v itself (Lemma 35). It follows that we can decompose $\mathrm{Vor}|_{\widetilde{\Sigma}} v$ into radial paths such that no two share a point besides v (Corollary 36). That is, if we remove v from each radial path, then we have a partition of $\mathrm{Vor}|_{\widetilde{\Sigma}} v \setminus \{v\}$ into paths. Therefore, $\varphi(\mathrm{Vor}|_{\widetilde{\Sigma}} v)$ is *star-shaped*: for every point $y \in \varphi(\mathrm{Vor}|_{\widetilde{\Sigma}} v)$, the line segment connecting v to y is a subset of $\varphi(\mathrm{Vor}|_{\widetilde{\Sigma}} v)$. Although $\mathrm{Vor}|_{\widetilde{\Sigma}} v$ is not itself star-shaped, its decomposition into radial paths is a curvy variant of the star-shaped property. As the lengths of the projected radial paths vary continuously, $\mathrm{Vor}|_{\widetilde{\Sigma}} v$ is homeomorphic to a closed disk (Theorem 40).

The following lemma implies that if you are standing on the boundary of $\mathrm{Vor}|_{\widetilde{\Sigma}} v$ and you try to walk toward v on a radial path, you immediately enter the interior of $\mathrm{Vor}|_{\widetilde{\Sigma}} v$ and stay there until you reach v.

▶ **Lemma 34.** *Consider two distinct sites $v, w \in V$ and a point $x \in \mathrm{Vor}|_{\widetilde{\Sigma}} v \cap \mathrm{Vor}|_{\widetilde{\Sigma}} w$. If x is in the principal branch, suppose that $|vx| < \xi\,\mathrm{lfs}(v)$ where $\xi = \sqrt{(\sqrt{5}-1)/2} \doteq 0.786151$; whereas if x lies on an extrusion Σ_s^+ or Σ_s^- of a segment $s = pq$, suppose that $d(p,q) \le 0.3647\,\mathrm{lfs}(p)$. Let Λ be the plane that bisects the line segment vw (and thus $x \in \Lambda$). Let $\mathcal{L}_v$ be the line normal to Σ at v. Let $a \subset T_x\widetilde{\Sigma}$ be the open ray that originates at x, is tangent to $\widetilde{\Sigma}$ at x, and passes through $\mathcal{L}_v$ (i.e., $\varphi(a)$ passes through v). Then v and a are strictly on the same side of Λ.*

(If x lies on a portal curve, x has two tangent spaces; then the result holds for $f \subset T_x\Sigma$ if $|vx| \le \xi\,\mathrm{lfs}(v)$, and it holds for $f \subset T_x\Sigma_s$ if $d(p,q) \le 0.3647\,\mathrm{lfs}(p)$.)

Proof. Let B and B' be the two open balls of radius $\mathrm{lfs}(v)$ tangent to Σ at v, and let o and o' be their centers, respectively. By Lemma 31 (if x is in the principal branch) or Lemma 32 (if x lies on an extrusion), $T_x\widetilde{\Sigma}$ intersects the relative interior of oo' at a lone point y. As $\mathcal{L}_v \supset oo'$, $T_x\widetilde{\Sigma}$ intersects $\mathcal{L}_v$ only at y. As $a \subset T_x\widetilde{\Sigma}$ and a intersects $\mathcal{L}_v$, a is the ray $\vec{xy}$.

Neither B nor B' intersects Σ, hence neither ball contains any site, hence $|vo| \leq |wo|$ and $|vo'| \leq |wo'|$. Therefore, each of o and o' lies either on Λ or on the same side of Λ as v, as illustrated in Figure 19. As $v \in oo'$ and $v \notin \Lambda$, Λ does not intersect the relative interior of oo'. By contrast, a does intersect the relative interior of oo' (at y). As $x \in \mathrm{Vor}|_{\widetilde{\Sigma}} v \cap \mathrm{Vor}|_{\widetilde{\Sigma}} w$, a's origin x lies on Λ. Therefore, the open ray a is strictly on the same side of Λ as the relative interior of oo', which contains v. ◀

The next lemma shows that under suitable sampling conditions, every point in $\mathrm{Vor}|_{\widetilde{\Sigma}} v \setminus \{v\}$ lies on one and only one radial path. Recall that for a site $w \in V$, w's *principal Voronoi cell* $\mathrm{Vor}|_{\overline{\Sigma}_S} w = \overline{\Sigma}_S \cap \mathrm{Vor}|_{\widetilde{\Sigma}} w$ excludes the portion of w's extended Voronoi cell on the extrusions (except points on the portal curves, which are included).

▶ **Lemma 35.** *Suppose that for every site $w \in V$ and every point y in the principal Voronoi cell $\mathrm{Vor}|_{\overline{\Sigma}_S} w$, $|wy| \leq \xi\, \mathrm{lfs}(w)$, where $\xi = \sqrt{(\sqrt{5}-1)/2} \doteq 0.786151$. Suppose that for every segment $pq \in S$, $|pq| \leq 0.3647\, \mathrm{lfs}(p)$. Consider a site $v \in V$ and a point $x \in \mathrm{Vor}|_{\widetilde{\Sigma}} v \setminus \{v\}$. Then there is a unique radial path $\gamma \subset \mathrm{Vor}|_{\widetilde{\Sigma}} v$ such that $x \in \gamma$. (Note that if x lies on an extrusion boundary ray v_s^+ or v_s^-, we consider $\gamma = v_s^+ = v_s^-$ to be a unique path topologically.)*

Moreover, let z be the endpoint of γ on the boundary of $\mathrm{Vor}|_{\widetilde{\Sigma}} v$; then no point in $\gamma \setminus \{z\}$ is in any other site's extended Voronoi cell.

Proof. Let $r \in T_v\Sigma$ be the closed ray with origin v that passes through $\varphi(x)$, with a topological "point at infinity" r^∞ at the end of r to make it compact. Let $C = \mathrm{Vor}|_{\widetilde{\Sigma}} v$ be a shorthand for v's extended Voronoi cell. As C may include parts of portal curves and extrusions, $\varphi|_C$ is not generally injective. We define the point set $\varphi|_C^{-1}(r) = \{z \in C : \varphi(z) \in r\}$. (One may think of $\varphi|_C^{-1}(r)$ as the intersection of C with a closed halfplane, namely, the halfplane whose orthogonal projection onto $T_v\Sigma$ is the ray r. Recall that each pair of extrusions Σ_s^+ and Σ_s^- has a point at infinity s^∞ to make $\widetilde{\Sigma}$ compact; we set the convention that $s^\infty \in \varphi|_C^{-1}(r)$ if and only if $\varphi(v_s^+) = r$ or $\varphi(v_s^-) = r$.) Let γ' be the connected component of $\varphi|_C^{-1}(r) \setminus \{v\}$ that contains x, and let $\gamma = \gamma' \cup \{v\}$. We will show that γ is a radial path. (More broadly, $\varphi|_C^{-1}(r)$ is a union of radial paths meeting at v.)

Every radial path containing x is a subset of $\varphi|_C^{-1}(r)$, and moreover is a subset of γ (because a radial path is connected and has v as an endpoint). The fact that γ is a radial path implies (by definition) that γ contains exactly one point on the boundary of C, which is γ's other endpoint besides v, so γ is the only radial path containing x.

If x lies on an extrusion boundary ray v_s^+ or v_s^-, then $\gamma = v_s^+ = v_s^-$ is a radial path, because by Theorem 2 every point on v_s^+ and v_s^- is in the interior of C except the point at infinity s^∞. (This is the only circumstance where γ is unbounded.)

Otherwise, as $\gamma' \subset C \setminus \{v\}$, by Lemma 33, for every point $y \in \gamma'$ there exists an open neighborhood $N \subset \widetilde{\Sigma}$ of y such that $\varphi|_N$ is a homeomorphism, so $\varphi|_{N \cap \gamma'}$ is a homeomorphism from $N \cap \gamma'$ to its image $\varphi(N \cap \gamma')$. Hence, $\varphi|_{\gamma'}$ is a local homeomorphism from the path γ' to its image $\varphi(\gamma')$. As γ' is connected and $\varphi(\gamma')$ is embedded in the ray r, $\varphi|_{\gamma'}$ is a (global) homeomorphism. (Intuitively, the map φ cannot cause the path to double back on itself, so $\varphi|_{\gamma'}$ is an injection.) Therefore, γ' is a topological interval or a single point.

Let q and z be the endpoints of γ' (in the topological space $\widetilde{\Sigma}$). As $\varphi|_{\gamma'}$ is a homeomorphism, $\varphi(q)$ and $\varphi(z)$ are the endpoints of $\varphi(\gamma')$ (in the topological space r). We choose the labels q and z so that $\varphi(q)$ is closer to v than $\varphi(z)$ is. As $v \notin \varphi(\gamma')$, $z \neq v$. By Corollary 5, C is closed (with respect to the space $\widetilde{\Sigma}$), hence $\varphi|_C^{-1}(r)$ is closed, hence $z \in \gamma'$ and if $q \neq v$ then $q \in \gamma'$. Hence $\gamma = \gamma \cup \{v\}$ is closed.

We will show that $q = v$ and that z is on the boundary of C, thereby establishing two of the criteria for γ to be a radial path. (Note that the fact that $q = v$ is what guarantees that $\varphi(C)$ is star-shaped.)

But first, we show the lemma's final claim: no point in $\gamma \setminus \{z\}$ is in the extended Voronoi cell of any site but v. Suppose for the sake of contradiction that some point $y \in \gamma \setminus \{z\}$ is in the cell $\mathrm{Vor}|_{\widetilde{\Sigma}} w$ of some site $w \in V \setminus \{v\}$. Then there is a plane Λ that orthogonally bisects vw such that $y \in \Lambda$. Clearly, $v \notin \Lambda$, so $y \neq v$. Recall that γ' might not be smooth at y, but γ' is piecewise smooth. Hence, let $\gamma_{yz} \subseteq \gamma$ be the closed subpath with endpoints y and z; then γ_{yz} has a unique tangent line at y. (As $y \neq z$, γ_{yz} has nonzero length.) Let f be the open ray that originates at y, is tangent to γ_{yz} (and thus tangent to $\widetilde{\Sigma}$) at y, and passes through $\mathcal{L}_v$. Note that f and γ_{yz} leave $y \in \Lambda$ in opposite directions. By Lemma 34, v and f are strictly on the same side of Λ. Hence, if you walk along γ_{yz} from y to z (in the direction opposite to the direction in which f points), you enter w's side of Λ at the instant you leave y.

By Theorem 4, neither vy nor wy intersects a portal boundary, so there is an open neighborhood $H \subset \gamma_{yz}$ of y such that every point in H is visible from both v and w. As $H \subset C$, no point in H is on the same side of Λ as w. But this contradicts the fact that H enters w's side of Λ where it leaves y. Hence, no point in $\gamma \setminus \{z\}$ is in any extended Voronoi cell besides $\mathrm{Vor}|_{\widetilde{\Sigma}} v$, as claimed.

It follows that no point on $\gamma \setminus \{z\}$ is on the boundary of C, establishing one of the criteria for γ to be a radial path.

Next, we show that every endpoint of γ' that is not v is on the boundary of C. Let y be an endpoint of γ' (either q or z) that is not v, and suppose for the sake of contradiction that y is in the interior of C. By Lemma 33, there exists an open neighborhood $N \subset \widetilde{\Sigma}$ of y such that $\varphi|_N$ is a homeomorphism from N to $\varphi(N)$; we can choose N such that $N \subset C \setminus \{v\}$. Then $\varphi(N) \subset T_v\Sigma$ is an open neighborhood of $\varphi(y)$ that does not contain v. Therefore, $\varphi(N) \cap r$ is an open neighborhood of $\varphi(y)$ in the space r. Observe that $\varphi(N) \cap r$ is not necessarily connected; let $I \subseteq \varphi(N) \cap r$ be an open interval with $\varphi(y)$ in its interior. As $I \subset \varphi(N)$ and $\varphi|_N^{-1}$ is a homeomorphism from $\varphi(N)$ to N, $\varphi|_N^{-1}(I)$ is a topological interval with y in its relative interior. As $\varphi|_N^{-1}(I) \subseteq \varphi|_C^{-1}(r) \setminus \{v\}$ and γ' is the connected component of $\varphi|_C^{-1}(r) \setminus \{v\}$ containing y, $\varphi|_N^{-1}(I) \subset \gamma'$. Hence y is in the relative interior of γ'. But this contradicts the fact that y is an endpoint of γ'. Hence y is on the boundary of C as claimed.

As $z \neq v$, it follows that z is on the boundary of C. We have seen that every point in $\gamma \setminus \{z\}$ is in the interior of C, so another consequence is that $q = v$ or $q = z$.

Next, we show that $q = v$; suppose for the sake of contradiction that $q = z$ (and thus $\gamma' = \{z\}$ has length zero). By Lemma 33, there exists an open neighborhood $N \subset \widetilde{\Sigma}$ of z such that $\varphi|_N$ is a homeomorphism. The intersection of $\varphi(N)$ with the line segment $v\varphi(z)$ is composed of one or more intervals; let $I \subseteq \varphi(N) \cap v\varphi(z) \setminus \{v\}$ be a closed interval with endpoint $\varphi(z)$ and nonzero length. Let $P = \varphi|_N^{-1}(I)$; as $\varphi|_N^{-1}$ is a homeomorphism from $\varphi(N)$ to N, P is a path (topological interval) with endpoint z and nonzero length. Let f be the open ray that originates at z, is tangent to P (and thus tangent to $\widetilde{\Sigma}$) at z, and passes through $\mathcal{L}_v$. Observe that f and P leave z in the *same* direction.

Recall that γ' is the connected component of $\varphi|_C^{-1}(r) \setminus \{v\}$ that contains z, and that $I \subset r$. As γ contains no point in P besides z, P must exit C at z. By Theorem 4, vz does not intersect a portal boundary, so P's exit cannot be explained by a portal shadow; it can only be explained by a bisecting plane. Specifically, there exists a site $w \in V \setminus \{v\}$ and a plane Λ that orthogonally bisects vw such that $z \in \Lambda$ and P enters w's side of Λ at z. Therefore, f lies entirely on w's side of Λ. But this contradicts the fact that, by Lemma 34, v and f are strictly on the same side of Λ. This contradiction establishes that $q = v$.

Hence, one endpoint of γ is v, one endpoint z is on the boundary of C, no point on $\gamma \setminus \{z\}$ lies on the boundary of C, and $\varphi(\gamma) \subset r$; therefore γ is a radial path. ◀

It follows that we can decompose $\mathrm{Vor}|_{\widetilde{\Sigma}} v$ into radial paths.

▶ **Corollary 36.** *Suppose the conditions specified in Lemma 35 hold. Let $v \in V$ be a site. Let Γ be the set of all radial paths for all points in $\mathrm{Vor}|_{\widetilde{\Sigma}} v \setminus \{v\}$. Then $\bigcup_{\gamma\in\Gamma} \gamma = \mathrm{Vor}|_{\widetilde{\Sigma}} v$ and no two paths in Γ share a common point besides v.*

Moreover, there is a one-to-one correspondence between paths in Γ and points on the boundary of $\mathrm{Vor}|_{\widetilde{\Sigma}} v$. Moreover, no other extended Voronoi cell intersects the interior of $\mathrm{Vor}|_{\widetilde{\Sigma}} v$.

Proof. By Lemma 35, for each point $x \in \mathrm{Vor}|_{\widetilde{\Sigma}} v \setminus \{v\}$, there is a unique radial path $\gamma \subset \mathrm{Vor}|_{\widetilde{\Sigma}} v$ such that $x \in \gamma$. By definition, $v \in \gamma$. Therefore, $\bigcup_{\gamma\in\Gamma} \gamma = \mathrm{Vor}|_{\widetilde{\Sigma}} v$. As each $x \in \mathrm{Vor}|_{\widetilde{\Sigma}} v \setminus \{v\}$ lies on only one radial path, no two paths in Γ share a common point besides v.

By definition, each radial path contains exactly one point on the boundary of $\mathrm{Vor}|_{\widetilde{\Sigma}} v$, and each point on the boundary is in exactly one radial path, so there is a one-to-one correspondence between them. By Lemma 35, for every path $\gamma \in \Gamma$, no point in $\gamma \setminus \{z_\gamma\}$ is in any other site's extended Voronoi cell, where z_γ is the endpoint of γ on the boundary of $\mathrm{Vor}|_{\widetilde{\Sigma}} v$. Therefore, no other extended Voronoi cell intersects the interior of $\mathrm{Vor}|_{\widetilde{\Sigma}} v$. ◀

Consider a site $v \in V$. Let $C = \mathrm{Vor}|_{\widetilde{\Sigma}} v$ be a shorthand for v's extended Voronoi cell. We will use the orthogonal projection φ (onto $T_v\Sigma$) to prove that under suitable sampling conditions, C is homeomorphic to a closed disk. However, as C may include parts of portal curves and extrusions, $\varphi|_C$ is not generally injective. Because we define $\widetilde{\Sigma}$ by gluing together rays on the boundaries of extrusions, $\varphi|_C$ is not generally continuous, either.

For simplicity, first consider the case where v does not adjoin a segment: then $\varphi|_C$ is continuous (as C intersects no extrusion boundary ray). If the conditions specified in Lemma 35 hold, we will see that $\varphi|_C$ is also injective. Under those conditions, by Corollary 36, we can decompose C into a set Γ of radial paths. Imagine walking along a radial path $\gamma \in \Gamma$ away from v; although you might eventually enter a portal, you walk initially on Σ (because v does not adjoin a segment). As Σ is smooth, every radial path $\gamma \in \Gamma$ is tangent to $T_v\Sigma$ at v. We can assign each radial path $\gamma \in \Gamma$ an angle $\theta \in [0°, 360°)$ according to the direction in which the line segment $\varphi(\gamma)$ leaves v on $T_v\Sigma$, and there is a one-to-one correspondence between angles and radial paths (because Σ is a smooth 2-manifold). Hence, for any two distinct radial paths $\gamma_1, \gamma_2 \in \Gamma$, the line segments $\varphi(\gamma_1)$ and $\varphi(\gamma_2)$ do not leave v in the same direction; hence $\varphi(\gamma_1) \cap \varphi(\gamma_2) = \{v\}$. Recall that by Lemma 35, $\varphi|_\gamma$ is injective for every $\gamma \in \Gamma$; it follows that $\varphi|_C$ is injective.

Now consider the general case, where v may adjoin segments in S. Again we can decompose C into a set Γ of radial paths (under the conditions of Lemma 35), but there may be distinct radial paths $\gamma_1, \gamma_2 \in \Gamma$ such that the line segments $\varphi(\gamma_1)$ and $\varphi(\gamma_2)$ leave v in the same direction; for example, if γ_1 lies on Σ whereas γ_2 lies on an extrusion. In that case, $\varphi|_C$ is not injective.

Our strategy for proving that C is homeomorphic to a disk is to slice the cell C like a pizza along radial paths, yielding several slices, each homeomorphic to a sector of a closed disk. Therefore, C (the whole pizza) is homeomorphic to a closed disk. We choose the cutting paths so that φ is continuous and injective over each pizza slice. The cutting paths include every extrusion boundary ray—in essence, we temporarily undo the gluing of rays. Then φ is continuous over each pizza slice.

Additional cuts may be needed to make φ be injective over each slice. As φ is injective over each radial path, it suffices that no slice includes two radial paths γ_1 and γ_2 such that $\varphi(\gamma_1)$ and $\varphi(\gamma_2)$ leave v in the same direction.

To show that cuts can always achieve this, we categorize radial paths based on whether they initially travel on Σ or on an extrusion where they leave v. Note that a radial path might pass through a portal, perhaps multiple times; its categorization depends solely on how it adjoins v. A radial path is a *portal path* if it is tangent to a portal curve at v and leaves v in the same direction as the portal curve; there are two portal paths for each portal curve adjoining v (one for each side of the portal). We will cut along every portal path. (Note that if the portal is parallel to n_v, then the portal paths follow the portal curve exactly; otherwise, they usually do not, and they might even cross the portal curve multiple times.) If a radial path is not a portal path, it is a *principal path* if it initially travels on Σ, or a *secondary path* if it initially travels on an extrusion. No two principal paths leave v in the same direction. Likewise, if two secondary paths γ_1 and γ_2 initially travel on the same extrusion, then $\varphi(\gamma_1)$ and $\varphi(\gamma_2)$ do not leave v in the same direction. However, the two portal paths that are tangent to a single portal curve both leave v in the same direction; thus we must ensure that no slice includes both of those portal paths.

Let $\Gamma_{\text{cut}} \subset \Gamma$ be a finite set of radial paths that contains every extrusion boundary ray in Γ, every portal path in Γ, and perhaps one or two additional radial paths chosen to ensure that every slice spans a sector strictly less than 360° (one is needed if v intersects only one segment; two are needed if v intersects none). These cutting paths ensure that Γ_{cut} cuts C into slices such that no slice includes two radial paths γ_1 and γ_2 for which $\varphi(\gamma_1)$ and $\varphi(\gamma_2)$ leave v in the same direction. This follows because no slice includes both principal and secondary paths, nor secondary paths from two different extrusions.

Let $C_{\text{cut}} = C \setminus \bigcup_{\gamma \in \Gamma_{\text{cut}}} \gamma$. Consider any connected component $\mathring{\Delta}$ of C_{cut}. Observe that $\mathring{\Delta}$ is equal to a union of some of the radial paths in Γ minus the site v. We call $\mathring{\Delta}$ a *half-open slice*, because $\mathring{\Delta}$ is closed along the portion of its boundary that lies on the boundary of C, but $\mathring{\Delta}$ is open where C was sliced along two radial paths in Γ_{cut}.

Let Δ denote the union of $\mathring{\Delta}$ with the two radial paths whose removal separates $\mathring{\Delta}$ from the rest of C. We call Δ a *closed slice*. (Lemma 38, below, implies that Δ is the closure of $\mathring{\Delta}$ with respect to the space C.) Observe that the union of all the closed slices is C.

We must address a minor technicality: recall that each extrusion boundary ray is embedded twice in $\widetilde{X}$. For a segment s adjoining v, the extrusion boundary rays v_s^+ and v_s^- are embedded in two different secondary branches, but topologically they are a single ray in $\widetilde{\Sigma}$, because we explicitly glued them together to construct $\widetilde{\Sigma}$. A consequence of this gluing is that φ is not continuous over C. Deviations from continuity occur only on extrusion boundary rays. When we cut C into slices, we want to undo this gluing so that φ is a continuous map over each closed slice. Hence a slice might have v_s^+ or v_s^- on its boundary, but not both. We unglue solely to clarify the proof that the slice is homeomorphic to a disk; the slice is a topological disk both before and after the gluing. Therefore, for each closed slice Δ, we define a slightly modified orthogonal projection φ_Δ that is almost the same as $\varphi|_\Delta$, but it projects only the "correct" embedding for the boundary of Δ: the embedding such that φ_Δ is continuous.

▶ **Lemma 37.** *Suppose the conditions specified in Lemma 35 hold. Consider a site $v \in V$ and its extended Voronoi cell $\mathrm{Vor}|_{\widetilde{\Sigma}} v$. Let $\mathring{\Delta}$ be a half-open slice of $\mathrm{Vor}|_{\widetilde{\Sigma}} v$ that does not intersect an extrusion boundary ray, though the corresponding closed slice Δ might intersect some. Suppose that Δ does not include two radial paths γ_1 and γ_2 such that $\varphi_\Delta(\gamma_1)$ and $\varphi_\Delta(\gamma_2)$ leave v in the same direction. Then φ_Δ is a homeomorphism from Δ to its image $\varphi_\Delta(\Delta)$ on $T_v\Sigma$.*

Proof. By Corollary 36, there is a decomposition Γ of $\mathrm{Vor}|_{\widetilde{\Sigma}} v$ into radial paths. The closed slice Δ is the union of a subset of those radial paths.

We first show that φ_Δ is an injection. Consider two points $x, y \in \Delta$ such that $\varphi_\Delta(x) = \varphi_\Delta(y)$; we will show that $x = y$. For any radial path $\gamma \subset \Delta$ that is not a subset of an extrusion

boundary ray, $\varphi_\Delta|_\gamma = \varphi|_\gamma$ is an injection by the definition of radial path (proven in Lemma 35). For any radial path $\gamma \subset \Delta$ that is a subset of an extrusion boundary ray, $\varphi_\Delta|_\gamma$ is clearly an injection (mapping a ray to a ray). Hence if x and y lie on the same radial path, then $x = y$. If x and y lie on two distinct radial paths $\gamma_1, \gamma_2 \subset \Delta$, then by assumption, the line segments $\varphi_\Delta(\gamma_1)$ and $\varphi_\Delta(\gamma_2)$ do not leave v in the same direction, so $\varphi_\Delta(\gamma_1) \cap \varphi_\Delta(\gamma_2) = \{v\}$ and $\varphi_\Delta(x) = \varphi_\Delta(y) = v$. But every radial path contains v and $\varphi_\Delta(v) = v$, so $x = y = v$. Hence φ_Δ is injective.

By the definition of φ_Δ and the fact that $\mathring{\Delta}$ does not intersect an extrusion boundary ray, φ_Δ is continuous. As Δ is compact and φ_Δ is injective and continuous, φ_Δ is a homeomorphism. ◀

Lemma 37 shows that φ_Δ is a homeomorphism from Δ to its image $I_\Delta = \varphi_\Delta(\Delta)$ on $T_v\Sigma$, but what is the shape of I_Δ? The next two lemmas show that I_Δ, and therefore Δ, is homeomorphic to a sector of a closed disk (equivalently, homeomorphic to a closed disk, but we will construct an explicit homeomorphism to a sector, which will be handy for gluing the pizza slices together into one pizza, thereby showing that $\mathrm{Vor}|_{\widetilde{\Sigma}} v$ is a topological disk).

Recall that the closed slice Δ is (by definition) the union of some of the radial paths in Γ. Let $\Lambda = \{\varphi_\Delta(\gamma) : \gamma \in \Gamma \text{ and } \gamma \subset \Delta\}$ be the set of the orthogonal projections of those radial paths onto $T_v\Sigma$. Then we can write $I_\Delta = \bigcup_{e\in\Lambda} e$, a decomposition of I_Δ into line segments with endpoint v, no two leaving v in the same direction. Note that if a radial path γ is an extrusion boundary ray, it has a point at infinity at its end for compactness, and its projection $\varphi_\Delta(\gamma)$ is also a ray with a point at infinity at its end for compactness. As the half-open slice $\mathring{\Delta}$ is connected (by definition), the interior of I_Δ is connected, so the line segments in Λ leave v at angles that form one continuous interval of angles.

For each edge $e \in \Lambda$, let $l(e)$ be the length of e (which is infinite if e is a ray). For every point $x \in I_\Delta \setminus \{v\}$, let e_x denote the unique line segment in Λ that contains x, and let $l(x) = l(e_x)$. Hence l is defined over the domain $I_\Delta \setminus \{v\}$ (but $l(v)$ is not defined). We will show that l is continuous, then we will use that fact to show that I_Δ is homeomorphic to a sector. The continuity of l follows directly from two facts: I_Δ is a closed point set and every point on the boundary of I_Δ is an endpoint of a line segment $e \in \Lambda$ or a point on one of the two boundary line segments.

▶ **Lemma 38.** *Suppose the conditions specified in Lemmas 35 and 37 hold. Then l is continuous over $I_\Delta \setminus \{v\}$.*

Proof. We show that l is continuous at any point $x \in I_\Delta \setminus \{v\}$ by showing that for every $\delta > 0$, there is a neighborhood $N \subset I_\Delta \setminus \{v\}$ of x such that for every point $y \in N$, $l(y) \in [l(x) - \delta, l(x) + \delta]$. By Corollary 36, there is a decomposition Γ of $\mathrm{Vor}|_{\widetilde{\Sigma}} v$ into radial paths, and by assumption, Δ is a union of some of those radial paths. Hence there is a decomposition Λ of I_Δ into line segments as described above. Recall that e_x denotes the unique line segment in Λ that contains x and $l(e_x)$ is its length. One endpoint of e_x is v; let z be the other endpoint. As $\varphi|_C$ is a homeomorphism by Lemma 37, both v and z are on the boundary of I_Δ. The distance between v and z is $l(x)$.

Let $z^- \in e_x$ be the point at a distance of $l(x) - \delta$ from v in the direction of z. (We assume $\delta < l(x)$.) Let $z^+ \in T_v\Sigma$ be the point at a distance of $l(x) + \delta$ from v in the direction of z. (Hence z is the midpoint between z^- and z^+.) As the slice Δ is closed and compact (with respect to the space $\widetilde{X}$), its projection I_Δ is also closed and compact (with respect to the space $T_v\Sigma$). As $z^+ \notin I_\Delta$, there is a neighborhood $N^+ \subset T_v\Sigma$ of z^+ that is disjoint from I_Δ. Let $C^+ \subset T_v\Sigma$ be a circle with center v and radius $l(x) + \delta$. Then $z^+ \in C^+$ and $C^+ \cap N^+$ is an arc or a union of arcs with z^+ in its relative interior. Let $a^+ \subseteq C^+ \cap N^+$ be the connected component (one arc) that contains z^+. Let p_1^+ and p_2^+ be the endpoints of a^+

and let $\theta^+ = \min\{\angle xvp_1^+, \angle xvp_2^+\}$, which is strictly positive. As a^+ is disjoint from I_Δ, every line segment $e \in \Lambda$ such that $\angle(e_x, e) < \theta^+$ has length less than $l(x) + \delta$. Therefore, for every point $y \in I_\Delta \setminus \{v\}$ such that $\angle xvy < \theta^+$, $l(y) < l(x) + \delta$.

We will use a similar argument about z^- to show that for some of these points y, $l(y) > l(x) - \delta$, but the details are more complicated.

Let $\gamma_x \in \Gamma$ be the radial path such that $e_x = \varphi_\Delta(\gamma_x)$. As z^- lies on e_x but is not an endpoint of e_x and $\varphi_\Delta|_{\gamma_x}$ is a homeomorphism (by the definition of radial path, and as established by Lemma 35), the point $z^* = \varphi_\Delta^{-1}(z^-)$ lies on γ_x but is not an endpoint of γ_x. Hence z^* lies in $\mathrm{Vor}|_{\widetilde{\Sigma}} v$ but not on its boundary (by the definition of radial path, and as established by Lemma 35) nor on v. Therefore, there is a neighborhood $N^* \subset \Delta$ of z^* that does not intersect v nor the boundary of $\mathrm{Vor}|_{\widetilde{\Sigma}} v$; that is, N^* does not intersect any endpoint of a radial path. Let $N^- = \varphi_\Delta(N^*)$. By Lemma 37, φ_Δ is a homeomorphism, so N^- is a neighborhood of z^- with respect to the space I_Δ, and N^- does not intersect any endpoint of a line segment in Λ. (Note that this observation is a key to this proof!)

Let $C^- \subset T_v\Sigma$ be a circle with center v and radius $l(x) - \delta$. Then $z^- \in C^-$ and $C^- \cap N^-$ is an arc or a union of arcs that contains z^-. Let $a^- \subseteq C^- \cap N^-$ be the connected component (one arc) that contains z^-. Let p_1^- and p_2^- be the endpoints of a^-. There are two cases. If e_x is not one of the two line segments on the boundary of I_Δ, then let $\theta^- = \min\{\angle xvp_1^-, \angle xvp_2^-\}$, which is strictly positive in that case. If e_x is one of the two line segments on the boundary of I_Δ, then one of the endpoints of a^- is z^-, but the other endpoint is not z^-, as the arc a^- cannot degenerate to a single point, because N^- does not intersect any endpoint of a line segment in Λ. In that case, suppose $p_1^- \neq z^-$ without loss of generality; let ϕ be the internal angle subtended by I_Δ at v, which is strictly less than 360° (as Λ does not contain two line segments that leave v in the same direction); and let $\theta^- = \min\{\angle xvp_1^-, 360° - \phi\}$.

As $a^- \subseteq I_\Delta$, every line segment $e \in \Lambda$ such that $\angle(e_x, e) < \theta^-$ intersects a^- and thus has length at least $l(x) - \delta$. Therefore, for every point $y \in I_\Delta \setminus \{v\}$ such that $\angle xvy < \theta^-$, $l(y) \geq l(x) - \delta$.

Let $N = I_\Delta \cap \{y \in T_v\Sigma \setminus \{v\} : \angle xvy < \min\{\theta^+, \theta^-\}\}$. As x is in the interior of the latter set, N is a neighborhood of x with respect to the space $I_\Delta \setminus \{v\}$. For every point $y \in N$, $l(y) \in [l(x) - \delta, l(x) + \delta]$. As we can identify such a neighborhood for every $\delta > 0$ and every $x \in I_\Delta \setminus \{v\}$, l is continuous. ◀

It is not hard to see that the continuity of l implies that Δ is homeomorphic to a sector of the unit disk, but for completeness, we construct an explicit homeomorphism. We wish to remap each segment in Λ to have unit length (while preserving its direction), thereby mapping I_Δ to a sector of the unit disk centered at v. A linear map might seem like an obvious choice, but we require a map that is continuous and injective even at a point x where $l(x) = \infty$.

To get around this technicality, we define $\chi(\omega, l) = \frac{1+1/l}{1+1/\omega} = \omega\frac{1+1/l}{\omega+1}$. Observe that χ is continuous over all values of $l \in (0, \infty]$ and $\omega \in [0, l]$ (use the first fraction for $\omega = \infty$, and the second for $\omega = 0$). The interval $\omega \in [0, l]$ is mapped bijectively to $\chi(\omega, l) \in [0, 1]$ (even if $l = \infty$). We define a point mapping $\mathring{\chi} : I_\Delta \to T_v\Sigma$ that maps each line segment in Λ to a line segment with unit length. For any point $x \in I_\Delta \setminus \{v\}$, let $\vec{e}_x$ be a unit vector directed from v along e_x. Let $\mathring{\chi}(x) = v + \chi(|vx|, l(x))\vec{e}_x$ for $x \neq v$ and let $\mathring{\chi}(v) = v$.

▶ **Lemma 39.** *Suppose the conditions specified in Lemmas 35 and 37 hold. Then $\mathring{\chi} \circ \varphi_\Delta$ is a homeomorphism from Δ to a closed sector of the unit disk on $T_v\Sigma$, and $\mathring{\chi} \circ \varphi_\Delta$ maps the two radial paths on the boundary of Δ to the two unit-length radii on the boundary of the sector.*

Proof. As $\chi(\omega, l)$ is continuous over $l \in (0, \infty], \omega \in [0, l]$ and l is continuous and positive over $I_\Delta \setminus \{v\}$, $\mathring{\chi}$ also is continuous over $I_\Delta \setminus \{v\}$. The map $\mathring{\chi}$ is continuous at v as well, because $\mathring{\chi}(x)$ converges to v in the limit as $x \in I_\Delta \setminus \{v\}$ approaches v.

To see that $\mathring{\chi}$ is injective, consider two points $x, y \in I_\Delta$ such that $\mathring{\chi}(x) = \mathring{\chi}(y)$. If $\mathring{\chi}(x) = \mathring{\chi}(y) = v$, then $x = y = v$. Otherwise, $\mathring{\chi}(x) = \mathring{\chi}(y) \neq v$, $x \neq v$, and $y \neq v$. If we express each point in polar coordinates relative to the origin v, we see that $\mathring{\chi}$ preserves a point's angle while mapping each point's radius so that each line segment in Λ is mapped to the unit interval; hence $x = y$ again. In either case, $\mathring{\chi}(x) = \mathring{\chi}(y)$ implies $x = y$, so $\mathring{\chi}$ is injective on I_Δ.

As I_Δ is compact and $\mathring{\chi}$ is injective and continuous, $\mathring{\chi}$ is a homeomorphism from I_Δ to its image $\mathring{\chi}(I_\Delta)$. As $\mathring{\chi}(I_\Delta)$ is a union of unit-length line segments subtending an interval of angles (the same range of angles that Λ subtends), $\mathring{\chi}(I_\Delta)$ is a sector of a unit disk.

As φ_Δ is a homeomorphism, $\mathring{\chi} \circ \varphi_\Delta$ is a homeomorphism from Δ to a sector of the unit disk. As this homeomorphism maps boundary points to boundary points and radial paths to unit-length line segments, it maps the two radial paths on the boundary of Δ to the two unit-length radii on the boundary of the sector. ◀

We are now ready to prove the main theorem of this section: that each extended Voronoi cell is a topological disk. (A summary of the proof: "Glue the sectors of a disk together along their straight edges, and you get a disk.")

▶ **Theorem 40.** *Suppose that for every site $w \in V$ and every point y in w's principal Voronoi cell $\mathrm{Vor}|_{\overline{\Sigma}_S} w$, $|wy| \leq \xi \operatorname{lfs}(w)$, where $\xi = \sqrt{(\sqrt{5}-1)/2} \doteq 0.786151$. Suppose that for every segment $pq \in S$, $|pq| \leq 0.3647 \operatorname{lfs}(p)$. Then for every site $v \in V$, its extended Voronoi cell $\mathrm{Vor}|_{\overline{\Sigma}} v$ is homeomorphic to a closed disk.*

Proof. Let $C = \mathrm{Vor}|_{\overline{\Sigma}} v$ be a shorthand for v's extended Voronoi cell. By Corollary 36, there is a decomposition Γ of C into radial paths. As we have argued above, there is a finite set of radial paths $\Gamma_{\mathrm{cut}} \subset \Gamma$ such that if $C_{\mathrm{cut}} = C \setminus \bigcup_{\gamma \in \Gamma_{\mathrm{cut}}} \gamma$, for every connected component $\mathring{\Delta}$ of C_{cut}, the corresponding closed slice Δ does not include two radial paths γ_1 and γ_2 such that $\varphi(\gamma_1)$ and $\varphi(\gamma_2)$ leave v in the same direction.

By Lemma 39, every such closed slice Δ is homeomorphic to a sector of a unit disk, and the two radial paths in Γ_{cut} that lie on the boundary of Δ are mapped by the homeomorphism to the two unit-length line segments on the boundary of the sector.

Suppose we glue together the sectors (one sector per slice) in a manner that corresponds to how the slices are glued together to form C. That is, each cutting path $\gamma \in \Gamma_{\mathrm{cut}}$ is included in exactly two closed slices, and the homeomorphisms of Lemma 39 map γ to a unit-length boundary edge for each of two sectors. We glue together those two boundary edges (i.e., topologically identify them with each other). Let D be the topological space produced by gluing all the sectors together. Then there is a one-to-one mapping from radial paths in C to radial line segments in D, and thus we construct a homeomorphism from C to D.

A space D created by gluing together sectors in this way, with every sector having the same base vertex v, is either a closed disk with v in its interior or several closed disks that have been glued together at a shared interior point v. (In other words, we could glue the sectors together to get one boundary loop or to get several boundary loops.) But by construction, $\widetilde{X}$ is a 2-manifold without boundary and $v \in C \subset \widetilde{X}$, so C cannot have the latter topology. Hence, C is homeomorphic to a closed disk. ◀

It is notable that the condition $|wy| < \xi \operatorname{lfs}(w)$ of Lemma 35 and Theorem 40 implies, by the Normal Variation Lemma (Lemma 13), that $\angle(n_w, n_y) < \eta(\xi) = 60°$. The sixty-degree

bound is exact. (We think it is just a coincidence that that number comes out so cleanly. The bound is probably not tight.)

The following corollary shows that if the principal Voronoi cells meet the condition of Theorem 40 (the condition of the forthcoming Corollary 42), every connected component of Σ has at least six sites on it.

▶ **Corollary 41.** *Let $V \subset \Sigma$ be a nonempty, finite set of points (sites) on Σ. Suppose that for every site $v \in V$ and every point x in v's principal Voronoi cell $\mathrm{Vor}|_{\overline{\Sigma}_S} v$, $|vx| < \xi \operatorname{lfs}(v)$, where $\xi = \sqrt{(\sqrt{5}-1)/2} \doteq 0.786151$.*

Then every connected component of Σ has at least six sites and at least six principal Voronoi cells on it.

Proof. Every point x in v's principal Voronoi cell $\mathrm{Vor}|_{\overline{\Sigma}_S} v$ lies on the same connected component of Σ as v does, because if they lay on different connected components, then the line segment vx would intersect the medial axis, contradicting the fact that $|vx| < \operatorname{lfs}(v)$. By the Normal Variation Lemma (Lemma 13), for every site $v \in V$ and every point $x \in \mathrm{Vor}|_{\overline{\Sigma}_S} v$, $\angle(n_v, n_x) < \eta(\xi) = 60°$.

Let σ be a connected component of Σ. For any unit vector u on the unit sphere, let y be a point on σ that is most extreme in the direction u. As σ is a smooth surface without boundary in $\mathbb{R}^3$, u is normal to σ at y and oriented to the outside of σ; that is, the unit vector $n_y = u$ is an outside-facing normal vector at y. It follows that the outside-facing unit normal vectors on σ constitute the entire sphere of directions. A principal Voronoi cell $\mathrm{Vor}|_{\overline{\Sigma}_S} v$ can contain only points on Σ whose outside-facing normals are less than 60° from n_v. At least six sites are required on a sphere so that every point on the sphere is less than 60° from one of the sites; five do not suffice [18]. (The six points where the coordinate axes intersect the unit sphere suffice.) Hence there are at least six sites and six principal Voronoi cells on σ.

The same reasoning applies to every connected component of Σ. ◀

Constrained 0.44-samples satisfy the sampling condition of Theorem 40.

▶ **Corollary 42.** *Let V be a constrained ϵ-sample of (Σ, S, Z) for $\epsilon < \frac{\xi}{\xi+1} \doteq 0.440137$. Suppose that for every segment $pq \in S$, $|pq| \le 0.3647 \operatorname{lfs}(p)$.*

Then every extended Voronoi cell in $\mathrm{Vor}|_{\overline{\Sigma}} V$ is homeomorphic to a closed disk, every connected component of Σ has at least six sites and at least six principal Voronoi cells on it, and no extended Voronoi cell intersects the interior of another extended Voronoi cell.

Proof. Consider any site $w \in V$ and any point $y \in \mathrm{Vor}|_{\overline{\Sigma}_S} w$. As V is a constrained ϵ-sample, $|wy| \le \epsilon \operatorname{lfs}(y)$. By the Feature Translation Lemma (Lemma 12), $|wy| \le (\epsilon/(1-\epsilon)) \operatorname{lfs}(w) < \xi \operatorname{lfs}(w)$.

The first claim follows from Theorem 40. The second claim follows from Corollary 41. The third claim follows from Corollary 36. ◀

G.3 The Nearest Point Map on a Triangle is a Homeomorphism

The forthcoming Theorem 46 establishes conditions under which the nearest point map, restricted to a restricted Delaunay triangle, is a homeomorphism; so there are no foldovers within a single triangle's map.

Assuming Σ is a 2-manifold, we define an *extended Voronoi vertex* to be a point in the intersection of three distinct extended Voronoi cells. However, without suitable sampling

conditions, such an intersection might include one or more line segments. Theorem 46 also establishes conditions that guarantee that an extended Voronoi vertex is isolated from any other points in the intersection, thereby justifying the name "vertex."

Recall that a *principal vertex* is an extended Voronoi vertex that lies on $\overline{\Sigma}_S$ (i.e., on Σ or a portal curve). A *secondary vertex* is an extended Voronoi vertex that is not principal; it lies on an extrusion but not on a portal curve.

Given an extended Voronoi vertex u and its dual restricted Delaunay triangle τ, Lemmas 44 and 45 below relate the normal vectors n_u and $n_{\tilde{x}}$ at any point $\tilde{x} \in \nu(\tau)$, showing that all these normals point to the same side of τ. The two lemmas are quite similar but differ in two ways: the sampling condition for Lemma 44 depends on $\operatorname{lfs}(p)$ at a vertex p of τ, whereas the sampling condition for Lemma 45 depends on $\operatorname{lfs}(u)$; and Lemma 44 applies to any extended Voronoi vertex u, principal or secondary, whereas Lemma 45 applies only if u is a principal vertex (otherwise $\operatorname{lfs}(u)$ is not defined). Theorem 46 applies if either of the two sampling conditions is met. Although a secondary vertex u does not lie on Σ, but instead lies on an extrusion, we still speak of an "outside-facing" normal vector n_u consistent with the outside-facing normal vectors on Σ, as we can extend Σ's orientation onto the extrusions.

In the unconstrained case, Theorem 46 implies that the nearest point map, restricted to a restricted Delaunay triangle, is a homeomorphism for a 0.3202-sample. Amenta et al. [3] proved the same for a 0.06-sample. We start with a technical lemma.

▶ **Lemma 43.** *Let $\tau \subset \mathbb{R}^3$ be a simplex. Let $x \in \tau$ be a point that does not lie on the medial axis of Σ. Let $\tilde{x}$ be the point on Σ nearest x. Let u be a point on Σ. There is a vertex p of τ such that $|p\tilde{x}| \le |pu|$, and such that $|p\tilde{x}| < |pu|$ if $\tilde{x} \neq u$.*

Proof. If $\tilde{x} = u$ then the result follows immediately, so assume that $\tilde{x} \neq u$. As x does not lie on the medial axis, $\tilde{x}$ is the unique point on Σ nearest x. As u also lies on Σ, $|x\tilde{x}| < |xu|$. Let Π be the plane that bisects the line segment $\tilde{x}u$, and observe that x lies on the same side of Π as $\tilde{x}$. As $x \in \tau$ and τ is a simplex, some vertex p of τ lies on the same side of Π as $\tilde{x}$, thus $|p\tilde{x}| < |pu|$. ◀

▶ **Lemma 44.** *Let u be an extended Voronoi vertex (principal or secondary) and let $\tau = \triangle pp'p''$ be the restricted Delaunay triangle dual to u, where p is the vertex of τ at τ's largest plane angle. Let μ be the positive root of $4\mu^4 = (1 - 4\mu^2)(1 - \sqrt{3}\mu)^2$, with approximate value $\mu \doteq 0.3606001$. Let $R = |pu| = |p'u| = |p''u|$ and suppose that $R < \mu \operatorname{lfs}(p)$. If u is a secondary vertex on an extrusion of a segment s, suppose also that the length of s is at most $\rho \operatorname{lfs}(a)$, where $\rho \le 0.47$ and a is an endpoint of s. Let x be any point on τ, and let $\tilde{x} = \nu(x)$ be the point on Σ nearest x. Let n_u be an outside-facing vector normal to $\widetilde{\Sigma}$ at u, let $n_{\tilde{x}}$ be an outside-facing vector normal to Σ at $\tilde{x}$, and let n_τ be a vector normal to τ.*

Then the angles $\angle(n_u, n_\tau)$ and $\angle(n_{\tilde{x}}, n_\tau)$ are either both less than $90°$ or both greater than $90°$ (depending on which way n_τ is directed). Equivalently, the dot products $n_u \cdot n_\tau$ and $n_{\tilde{x}} \cdot n_\tau$ are either both positive or both negative.

Proof. Let r be τ's circumradius. Let S be the sphere with center u and radius R, which passes through all three vertices of τ. As τ's circumcircle is a cross section of S, $r \le R$.

By the Triangle Normal Lemma (Lemma 11), $\sin \angle(n_p, n_\tau) \le \sqrt{3}r/\operatorname{lfs}(p) < \sqrt{3}\mu$. Suppose without loss of generality that n_τ is directed so that $\angle(n_p, n_\tau)$ is acute; then $\angle(n_p, n_\tau) < 38.652°$.

By Lemma 25 (with S as defined above, $P = \{p, p', p''\}$, and $H = \tau$), $\tilde{x}$ is inside or on S; hence $|u\tilde{x}| \le R$ and $|p\tilde{x}| \le |pu| + |u\tilde{x}| \le 2R < 2\mu \operatorname{lfs}(p)$. By the Normal Variation Lemma (Lemma 13), $\angle(n_p, n_{\tilde{x}}) < \eta(2\mu)$ where $\eta(\delta) = \arccos\left(1 - \frac{\delta^2}{2\sqrt{1-\delta^2}}\right)$. Therefore, $\angle(n_p, n_{\tilde{x}}) <$

$\arccos\left(1 - \frac{2\mu^2}{\sqrt{1-4\mu^2}}\right) = \arccos\left(1 - (1-\sqrt{3}\mu)\right) = \arccos(\sqrt{3}\mu)$ and $\angle(n_{\tilde{x}}, n_\tau) \le \angle(n_p, n_{\tilde{x}}) + \angle(n_p, n_\tau) < \arccos(\sqrt{3}\mu) + \arcsin(\sqrt{3}\mu) = 90°$.

If u is a principal vertex, then as $|pu| = R < \mu\,\mathrm{lfs}(p)$, by the Normal Variation Lemma, $\angle(n_p, n_u) < \eta(\mu) < 21.52°$. Therefore, $\angle(n_u, n_\tau) \le \angle(n_p, n_u) + \angle(n_p, n_\tau) < 21.52° + 38.652° = 60.172°$. So $\angle(n_u, n_\tau)$ and $\angle(n_{\tilde{x}}, n_\tau)$ are both less than 90°, and the lemma holds.

If u is a secondary vertex, let s be the segment on whose extrusion u lies, let h_s be the cutting plane for s, and let $\zeta_s \subset h_s \cap \Sigma$ be the portal curve for s. Let $\bar{u}$ be the point nearest u on h_s, and note that $\bar{u} \in \zeta_s$ and $\bar{u} \in \Sigma$. It follow from Theorem 1 that the plane h_s separates τ from u; therefore, $|p\bar{u}| < |pu|$.

As $|p\bar{u}| < |pu| = R < \mu\,\mathrm{lfs}(p)$, by the Normal Variation Lemma, $\angle(n_p, n_{\bar{u}}) < \eta(\mu) < 21.52°$. As the length of s is at most $\rho\,\mathrm{lfs}(a)$, we have $|a\bar{u}| \le \rho\,\mathrm{lfs}(a)$ and, by the Normal Variation Lemma, $\angle(n_a, n_{\bar{u}}) \le \eta(\rho) \le \eta(0.47) < 28.971°$. As the site u lies on an extrusion from ζ_s, the vector n_u normal to the extrusion at u is the projection of $n_{\bar{u}}$ onto h_s. As h_s is parallel to n_a, $\angle(n_u, n_{\bar{u}}) \le \angle(n_a, n_{\bar{u}}) < 28.971°$.

Therefore, $\angle(n_u, n_\tau) \le \angle(n_u, n_{\bar{u}}) + \angle(n_p, n_{\bar{u}}) + \angle(n_p, n_\tau) < 28.971° + 21.52° + 38.652° = 89.143°$. Hence, $\angle(n_u, n_\tau)$ and $\angle(n_{\tilde{x}}, n_\tau)$ are both less than 90° as claimed. ◂

▸ **Lemma 45.** *Let u be a principal vertex and let $\tau = \triangle pp'p''$ be the restricted Delaunay triangle dual to u. Let x be any point on τ, and let $\tilde{x} = \nu(x)$ be the point on Σ nearest x. Let n_u be an outside-facing vector normal to Σ at u, let $n_{\tilde{x}}$ be an outside-facing vector normal to Σ at $\tilde{x}$, and let n_τ be a vector normal to τ. Let $R = |pu| = |p'u| = |p''u|$ and suppose that $R \le 0.3202\,\mathrm{lfs}(u)$.*

Then the angles $\angle(n_u, n_\tau)$ and $\angle(n_{\tilde{x}}, n_\tau)$ are either both less than 90° or both greater than 90° (depending on which way n_τ is directed). Equivalently, the dot products $n_u \cdot n_\tau$ and $n_{\tilde{x}} \cdot n_\tau$ are either both positive or both negative.

Proof. Let r be τ's circumradius. Let S be the sphere with center u and radius R, which passes through all three vertices of τ. As τ's circumcircle is a cross section of S, $r \le R$.

Suppose without loss of generality that p is the vertex of τ nearest $\tilde{x}$. Let $q \in \{p, p', p''\}$ be the vertex at τ's largest plane angle. As $|pu| = |qu| = R \le 0.3202\,\mathrm{lfs}(u)$, by the Feature Translation Lemma (Lemma 12), $\mathrm{lfs}(u) \le \mathrm{lfs}(p)/(1 - 0.3202)$ and likewise $\mathrm{lfs}(u) \le \mathrm{lfs}(q)/(1 - 0.3202)$, so $r \le R \le \frac{0.3202}{0.6798}\,\mathrm{lfs}(p) < 0.47103\,\mathrm{lfs}(p)$ and likewise $r \le 0.47103\,\mathrm{lfs}(q)$. By Lemma 25 (with S as defined above, $P = \{p, p', p''\}$, and $H = \tau$), $\tilde{x}$ is inside or on S; hence $|\tilde{x}u| \le R \le 0.3202\,\mathrm{lfs}(u)$. By Lemma 43, there is a vertex $\dot{p}$ of τ such that $|\dot{p}\tilde{x}| \le |\dot{p}u|$; hence $|p\tilde{x}| \le |\dot{p}\tilde{x}| \le |\dot{p}u| = R \le 0.47103\,\mathrm{lfs}(p)$. By the Normal Variation Lemma (Lemma 13), $\angle(n_p, n_u) \le \eta(0.3202)$, $\angle(n_q, n_u) \le \eta(0.3202)$, $\angle(n_{\tilde{x}}, n_u) \le \eta(0.3202)$, and $\angle(n_p, n_{\tilde{x}}) \le \eta(0.47103)$, where $\eta(\delta) = \arccos\left(1 - \frac{\delta^2}{2\sqrt{1-\delta^2}}\right)$, $\eta(0.3202) < 18.94°$, and $\eta(0.47103) < 29.05°$.

If τ's plane angle at the vertex p is 56.653° or greater, then by the Triangle Normal Lemma (Lemma 11), $\sin\angle(n_p, n_\tau) \le r \cot 28.3265°/\mathrm{lfs}(p) < 0.47103 \cdot 1.8552 < 0.8739$. Suppose without loss of generality that n_τ is directed so that $\angle(n_p, n_\tau)$ is acute; then $\angle(n_p, n_\tau) < 60.92°$. Thus $\angle(n_u, n_\tau) \le \angle(n_p, n_u) + \angle(n_p, n_\tau) < 18.94° + 60.92° = 79.86°$ and $\angle(n_{\tilde{x}}, n_\tau) \le \angle(n_p, n_{\tilde{x}}) + \angle(n_p, n_\tau) < 29.05° + 60.92° = 89.97°$, so both angles are less than 90° as claimed.

Otherwise, τ's plane angle at p is less than 56.653°, so τ's plane angle at q (τ's largest plane angle) is greater than $(180° - 56.653°)/2 = 61.6735°$. By the Triangle Normal Lemma (Lemma 11), $\sin\angle(n_q, n_\tau) \le r \cot 30.83675°/\mathrm{lfs}(q) < 0.47103 \cdot 1.6751 < 0.7891$. Suppose without loss of generality that n_τ is directed so that $\angle(n_q, n_\tau)$ is acute; then $\angle(n_q, n_\tau) < 52.11°$. Thus $\angle(n_u, n_\tau) \le \angle(n_q, n_u) + \angle(n_q, n_\tau) < 18.94° + 52.11° = 71.05°$ and $\angle(n_{\tilde{x}}, n_\tau) \le$

$\angle(n_{\tilde{x}}, n_u) + \angle(n_q, n_u) + \angle(n_q, n_\tau) < 18.94° + 18.94° + 52.11° = 89.99°$, confirming that both angles are less than 90°. ◀

▶ **Theorem 46.** *Consider* $\mathrm{Vor}|_{\widetilde{\Sigma}} V$*, where* $\Sigma \subset \mathbb{R}^3$ *is a smooth 2-manifold without boundary,* $V \subset \Sigma$ *is a finite sample, and each segment* $s \in S$ *has length at most* $0.47\,\mathrm{lfs}(a)$ *for some endpoint* a *of* s*. Consider three distinct sites* $p, p', p'' \in V$ *and the triangle* $\tau = \triangle pp'p''$*, where* p *is the vertex of* τ *at* τ*'s largest plane angle. Let* $U = \mathrm{Vor}|_{\widetilde{\Sigma}} p \cap \mathrm{Vor}|_{\widetilde{\Sigma}} p' \cap \mathrm{Vor}|_{\widetilde{\Sigma}} p''$ *(which is the extended Voronoi face dual to* τ*) and suppose that* $U \neq \emptyset$ *(so* τ *is a restricted Delaunay triangle). Let* u *be a point in* U *and let* $R = |pu| = |p'u| = |p''u|$*. Suppose that at least one of the following holds: either* $R \leq 0.3606\,\mathrm{lfs}(p)$*, or* u *is a principal vertex and* $R \leq 0.3202\,\mathrm{lfs}(u)$*.*

Then the nearest point map ν *restricted to* τ*, denoted* $\nu|_\tau$*, is a homeomorphism from* τ *to its image* $\nu(\tau)$ *on* Σ*. Moreover, no normal segment of* Σ *that intersects* τ *is parallel to* τ*. Moreover,* u *is isolated from the other points in* U*.*

Proof. First we show that τ does not intersect the medial axis M of Σ, so the (restricted) nearest point map $\nu|_\tau$ is defined and continuous over τ. If the condition $R \leq 0.3202\,\mathrm{lfs}(u)$ holds, then the distance from u to any point in τ is at most $0.3202\,\mathrm{lfs}(u)$, whereas the distance from u to M is $\mathrm{lfs}(u)$, so τ is disjoint from M. If the condition $R \leq 0.3606\,\mathrm{lfs}(p)$ holds, then the distance from p to any point in τ is at most $2R \leq 0.7212\,\mathrm{lfs}(p)$, whereas the distance from p to M is $\mathrm{lfs}(p)$; again τ is disjoint from M. In either case, $\nu|_\tau$ is continuous over τ.

By Lemma 44 (if $R \leq 0.3606\,\mathrm{lfs}(p)$) or Lemma 45 (if $R \leq 0.3202\,\mathrm{lfs}(u)$), for every point $x \in \tau$, $\angle(n_{\tilde{x}}, n_\tau) \neq 90°$; therefore, the unique normal segment $\ell_{\tilde{x}}$ that passes through x is not parallel to τ. This proves our claim that no normal segment that intersects τ is parallel to τ. It follows that the nearest point map $\nu|_\tau$ is injective: if two distinct points $x, x' \in \tau$ could map to the same point $\tilde{x} \in \Sigma$, then $\tilde{x}$'s normal segment $\ell_{\tilde{x}}$ would intersect both x and x' and thus be parallel to τ, but that is not possible.

The nearest point map $\nu|_\tau$ is a continuous bijection between a compact set τ and its image $\nu(\tau)$ on a bounded manifold. Its inverse $\nu|_\tau^{-1}$ is also continuous over $\nu(\tau)$, as the normal lines are a continuous function of the points on Σ, and the intersection of a line with τ's affine hull Π is a continuous function over the domain of lines that are not parallel to Π. Hence $\nu|_\tau$ is a homeomorphism (proving our first claim).

To address our third claim, let ℓ_τ be the set containing every point z in the extended three-dimensional space $\widetilde{X}$ such that $d(z, p) = d(z, p') = d(z, p'')$, and observe that $U \subset \ell_\tau$. By Lemma 44 or Lemma 45, $\angle(n_u, n_\tau) \neq 90°$. As ℓ_τ is parallel to the normal vector n_τ, ℓ_τ does not intersect Σ tangentially at u, so u is isolated from the other points in U. ◀

G.4 Extended Voronoi Edges Are Topological Line Segments

Theorem 40 and Corollary 42 give conditions under which each principal Voronoi cell is a topological closed disk, and no cell intersects the interior of another cell. Theorem 46 gives conditions under which the intersection of any three distinct extended Voronoi cells is composed of isolated points, which we call "extended Voronoi vertices." What about an intersection of two distinct extended Voronoi cells? We call such an intersection an *extended Voronoi edge* if it contains a connected curve (and thus it is not merely a set of isolated points). The following lemma helps to justify this name.

▶ **Lemma 47.** *Consider an extended Voronoi diagram* $\mathrm{Vor}|_{\widetilde{\Sigma}} V$*. Suppose that every extended Voronoi cell is a topological closed disk and every intersection of three distinct extended Voronoi cells is a set of isolated points (i.e., no two distinct points in the intersection are*

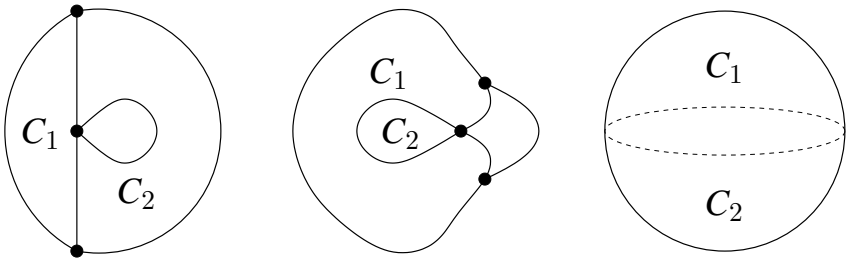

Figure 21 Left: Cells C_1 and C_2 share two adjoining edges; either C_1 or C_2 is not a topological disk. Center: Cells C_1 and C_2 share a circle with one vertex on it; either C_1 or C_2 is not a topological disk. Right: Cells C_1 and C_2 share a circle with no vertex on it; if they are topological disks, their union is a connected component of Σ.

path-connected). Suppose also that no extended Voronoi cell intersects the interior of another extended Voronoi cell. Then for every pair of distinct sites $p, q \in V$, $\mathrm{Vor}|_{\widetilde{\Sigma}} p \cap \mathrm{Vor}|_{\widetilde{\Sigma}} q$ is either a topological circle containing no extended Voronoi vertex or a union of disjoint topological closed 1-balls and isolated points, where each isolated point is an extended Voronoi vertex and each 1-ball contains exactly two extended Voronoi vertices which are its endpoints.

Moreover, if at least three sites in V lie on each connected component of Σ, then the possibility that $\mathrm{Vor}|_{\widetilde{\Sigma}} p \cap \mathrm{Vor}|_{\widetilde{\Sigma}} q$ is a topological circle is eliminated, every extended Voronoi cell has at least two extended Voronoi vertices on its boundary, and every connected component of Σ has at least two extended Voronoi vertices on it.

Proof. As $\widetilde{\Sigma}$ is a surface without boundary and $\widetilde{\Sigma}$ is also a union of extended Voronoi cells, which are topological closed disks, each point on the boundary of each extended Voronoi cell is shared with at least one other extended Voronoi cell. By assumption, no interior point of a cell is shared with another cell.

Consider a site p, its extended Voronoi cell $\mathrm{Vor}|_{\widetilde{\Sigma}} p$, and the cell's boundary C, which is a topological circle. If two or more extended Voronoi vertices lie on C, they subdivide C into two or more topological closed 1-balls (as extended Voronoi vertices are isolated points by assumption). Let I be one of these topological 1-balls, or let $I = C$ if C contains fewer than two extended Voronoi vertices. The subset of I obtained by removing its extended Voronoi vertices is path-connected, so the points in that subset are all shared with one and only one other site q; hence $I \subseteq \mathrm{Vor}|_{\widetilde{\Sigma}} p \cap \mathrm{Vor}|_{\widetilde{\Sigma}} q$. It follows that for every $p, q \in V$, $\mathrm{Vor}|_{\widetilde{\Sigma}} p \cap \mathrm{Vor}|_{\widetilde{\Sigma}} q$ is either a topological circle or a union of 1-balls and extended Voronoi vertices.

The intersection of two cells C_1 and C_2 cannot include two 1-balls that are not disjoint—that is, two 1-balls that share an extended Voronoi vertex—because the shared vertex lies on the boundary of a third cell, which implies that either C_1 or C_2 is not a topological closed disk, as Figure 21 (left) illustrates. For the same reason, if the intersection of two cells is a topological circle, no extended Voronoi vertex can lie on the circle, as Figure 21 (center) illustrates. This establishes the lemma's first claim.

If the intersection of two cells is a topological circle (with no extended Voronoi vertex), then as the two cells are topological disks, their union is a topological sphere covering an entire connected component of Σ, as Figure 21 (right) illustrates. Hence, if at least three sites in V lie on each connected component of Σ, then no two cells have a circle as their intersection. The lemma's second claim follows. ◂

G.5 The Nearest Point Map Preserves Orientation

Section G.3 and the forthcoming Lemmas 48, 49, and 51 use the idea of assigning an orientation to Σ, which manifests as both a normal vector direction and a rotary spin in the tangent plane. There are two directions in which a normal vector n_u can point; let us choose the normal vectors so they all point to the outside of Σ. Then we define a counterclockwise ordering of the extended Voronoi cells adjoining an extended Voronoi vertex u according to a *right-hand rule*: with the thumb of your right hand pointing in the direction of n_u, your fingers curl in a direction that defines the *counterclockwise* ordering of cells adjoining u.

We can extend this notion of orientation to all the points on the normal segments. Let M be the medial axis of Σ. Each point $y \in \mathbb{R}^3 \setminus M$ lies on the normal segment ℓ_x of a point $x = \nu(y) \in \Sigma$. The point y inherits both aspects of x's orientation: the orientation direction n_x, parallel to ℓ_x, and the counterclockwise ordering around ℓ_x, derived from n_x by the right-hand rule.

Let $\tau \subset \mathbb{R}^3 \setminus M$ be a triangle whose vertices lie on Σ. Theorem 46 guarantees (under the stated conditions) that the nearest point map ν is a homeomorphism from τ to $\nu(\tau)$. By Lemmas 45 and 44, the orientations of the points on τ are all mutually consistent: their orientation directions all point to the same side of τ. Let n_τ be a unit vector orthogonal to τ that points to the same side as well. The right-hand rule induces a counterclockwise ordering of τ's vertices around τ's boundary, which is also a counterclockwise ordering of $\nu(\tau)$'s vertices around $\nu(\tau)$'s boundary.

Consider an extended Voronoi vertex $u \in \mathrm{Vor}|_{\widetilde{\Sigma}} V$ generated by three sites $p, p', p'' \in V$ and u's dual restricted Delaunay triangle $\tau = \triangle pp'p''$. Let $\ell_\tau \subset \mathbb{R}^3$ be the line comprising the points equidistant from p, p', and p''; ℓ_τ passes through both τ's circumcenter and u. (However, ℓ_τ is *not* necessarily parallel to n_u.) Imagine the three-site Voronoi diagram $\mathrm{Vor}\,\{p, p', p''\}$: it subdivides $\mathbb{R}^3$ into three wedges W_p, $W_{p'}$, and $W_{p''}$ whose mutual intersection is ℓ_τ.

In the special case where there are no segments and no portals—that is, the case of a restricted (but unconstrained) Delaunay triangulation—the restricted Voronoi cells in $\mathrm{Vor}|_\Sigma V$ satisfy $\mathrm{Vor}|_\Sigma p \subset W_p$, $\mathrm{Vor}|_\Sigma p' \subset W_{p'}$, and $\mathrm{Vor}|_\Sigma p'' \subset W_{p''}$. In the presence of portals, the extended Voronoi cells do not necessarily satisfy these constraints, but they hold in a sufficiently small neighborhood of u. Specifically, it follows from Theorem 4 that there is an open neighborhood $N \subset \widetilde{\Sigma}$ of u such that every point in N is visible from p, p', and p''. Hence $N \cap \mathrm{Vor}|_{\widetilde{\Sigma}} p \subset W_p$, $N \cap \mathrm{Vor}|_{\widetilde{\Sigma}} p' \subset W_{p'}$, and $N \cap \mathrm{Vor}|_{\widetilde{\Sigma}} p'' \subset W_{p''}$. Therefore, the cyclical ordering of τ's vertices around ℓ_τ is consistent with the cyclical ordering of their extended Voronoi cells around ℓ_τ where they touch u.

You might expect that if τ's vertices p, p', and p'' occur in counterclockwise order around τ's perimeter, then $\mathrm{Vor}|_{\widetilde{\Sigma}} p$, $\mathrm{Vor}|_{\widetilde{\Sigma}} p'$, $\mathrm{Vor}|_{\widetilde{\Sigma}} p''$ occur in counterclockwise order around u. However, if the surface is twisted enough that the angle between n_u and n_τ exceeds 90°, this expectation is violated, and $\mathrm{Vor}|_{\widetilde{\Sigma}} p$, $\mathrm{Vor}|_{\widetilde{\Sigma}} p'$, $\mathrm{Vor}|_{\widetilde{\Sigma}} p''$ are in *clockwise* order around u. Intuitively, this causes a nasty "foldover" in the restricted Delaunay triangulation, which may prevent ν from being injective over the triangulation. We shall show that such foldovers can be prevented by use of a sufficiently dense sample.

▶ **Lemma 48.** *Let u be an extended Voronoi vertex and let $\tau = \triangle pp'p''$ be the restricted Delaunay triangle dual to u. Let $R = |pu| = |p'u| = |p''u|$ and suppose that at least one of the following holds: either $R \leq 0.3606\,\mathrm{lfs}(q)$ where q is the vertex of τ at τ's largest plane angle, or u is a principal vertex and $R \leq 0.3202\,\mathrm{lfs}(u)$. Then p, p', and p'' are in counterclockwise order around τ if and only if $\mathrm{Vor}|_{\widetilde{\Sigma}} p$, $\mathrm{Vor}|_{\widetilde{\Sigma}} p'$, and $\mathrm{Vor}|_{\widetilde{\Sigma}} p''$ adjoin u in counterclockwise order around u.*

Proof. By Lemma 44 (if $R \leq 0.3606\,\mathrm{lfs}(q)$) or Lemma 45 (if u is a principal vertex and $R \leq 0.3202\,\mathrm{lfs}(u)$), the outside-facing normal n_u at u and the outside-facing normals $n_{\tilde{x}}$ for every point $x \in \tau$ are all directed to the same side of τ. Therefore, we can assign a counterclockwise orientation to every point on τ that is consistent over τ and induces a counterclockwise ordering of τ's vertices.

Consider the wedges W_p, $W_{p'}$, and $W_{p''}$ and their line ℓ_τ of mutual intersection, defined above. Recall that ℓ_τ is parallel to n_τ and passes through u. Recall that $N \cap \mathrm{Vor}|_{\widetilde{\Sigma}} p \subset W_p$, $N \cap \mathrm{Vor}|_{\widetilde{\Sigma}} p' \subset W_{p'}$, and $N \cap \mathrm{Vor}|_{\widetilde{\Sigma}} p'' \subset W_{p''}$. The ordering of the sites p, p', and p'' around τ (counterclockwise or clockwise) is determined by the orientation of the normals $n_{\tilde{x}}$ relative to τ; in turn, this ordering determines the ordering of the wedges around ℓ_τ. The ordering of the cells $\mathrm{Vor}|_{\widetilde{\Sigma}} p$, $\mathrm{Vor}|_{\widetilde{\Sigma}} p'$, and $\mathrm{Vor}|_{\widetilde{\Sigma}} p''$ around u is determined by the ordering of the wedges around ℓ_τ and the orientation of the normal n_u relative to τ. As the normals n_u and $n_{\tilde{x}}$ point to the same side of τ, the result follows. ◀

G.6 The Nearest Point Map Is Surjective

Consider two restricted Delaunay triangles $\tau_1 = \triangle pp'w_1$ and $\tau_2 = \triangle p'pw_2$ that satisfy the requirements of Theorem 46; hence $\nu|_{\tau_1}$ is a homeomorphism and so is $\nu|_{\tau_2}$. The two triangles share an edge pp'. The next lemma shows that their images under ν do not overlap each other; in particular, their images fall on opposite "sides" of the image $\nu(pp')$.

▶ **Lemma 49.** *Let $e \in \mathrm{Vor}|_{\widetilde{\Sigma}} V$ be an extended Voronoi edge with vertices u_1 and u_2, whose dual restricted Delaunay triangles are $\tau_1 = \triangle pp'w_1$ and $\tau_2 = \triangle p'pw_2$, with $p, p', w_1, w_2 \in V$. Suppose that the extended Voronoi cells of p, p', and w_1 adjoin u_1 in counterclockwise order around u_1, and the cells of p', p, and w_2 adjoin u_2 in counterclockwise order around u_2. Let $R_1 = |pu_1| = |p'u_1| = |w_1u_1|$ and $R_2 = |pu_2| = |p'u_2| = |w_2u_2|$. Suppose that at least one of the following holds: either $R_1 \leq 0.3606\,\mathrm{lfs}(q_1)$ where q_1 is the vertex of τ_1 at τ_1's largest plane angle, or u_1 is a principal vertex and $R_1 \leq 0.3202\,\mathrm{lfs}(u_1)$. Moreover, suppose that at least one of the following holds: either $R_2 \leq 0.3606\,\mathrm{lfs}(q_2)$ where q_2 is the vertex of τ_2 at τ_2's largest plane angle, or u_2 is a principal vertex and $R_2 \leq 0.3202\,\mathrm{lfs}(u_2)$. Let $Q = \tau_1 \cup \tau_2$. Then $\nu|_Q$ is a homeomorphism from Q to its image $\nu(Q)$ on Σ.*

Proof. By Theorem 46, $\nu|_{\tau_1}$ is a continuous bijection with a continuous inverse that preserves the orientation of every projected point, and so is $\nu|_{\tau_2}$. Hence it remains only to show that $\nu|_Q$ is injective.

Suppose for the sake of contradiction that there are two distinct points $x \in \tau_1 \setminus pp'$ and $y \in \tau_2 \setminus pp'$ such that $\nu(x) = \nu(y)$, as illustrated in Figure 22. Let $\tilde{x} = \nu(x) = \nu(y)$. Let $\ell_{\tilde{x}}$ be the normal segment of $\tilde{x}$, which passes through $\tilde{x}$, x, and y. Let $T_{\tilde{x}}\Sigma$ be the plane tangent to Σ at $\tilde{x}$, which is perpendicular to $\ell_{\tilde{x}}$. For any point $p \in \mathbb{R}^3$, let $\bar{p}$ denote the orthogonal projection of p onto $T_{\tilde{x}}\Sigma$. (The projection direction is parallel to $\ell_{\tilde{x}}$.) By Theorem 46, neither τ_1 nor τ_2 is parallel to $\ell_{\tilde{x}}$, so the orthogonal projections of τ_1 and τ_2 onto $T_{\tilde{x}}\Sigma$ are triangles (rather than line segments). As $\tilde{x}$ lies in both orthogonal projections but not on the orthogonal projection of pp', it follows that if $\bar{p}$, $\bar{p}'$, and $\bar{w}_1$ occur in counterclockwise order on $T_{\tilde{x}}\Sigma$, then $\bar{p}'$, $\bar{p}$, and $\bar{w}_2$ occur in clockwise order on $T_{\tilde{x}}\Sigma$; and if the former sites occur in clockwise order, then the latter sites occur in counterclockwise order. Therefore, the orientation of p, p', and w_1 is opposite to the orientation of p', p, and w_2.

By Lemma 48, p, p', and w_1 occur in counterclockwise order around the boundary of τ_1; likewise, p', p, and w_2 occur in counterclockwise order around the boundary of τ_2. But this contradicts the conclusion of the previous paragraph.

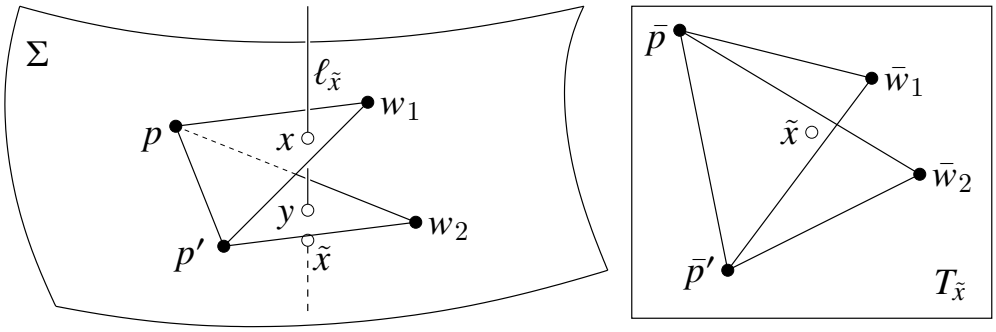

Figure 22 If $Q = \tau_1 \cup \tau_2$ is not injective, then one of the triangles has the wrong orientation.

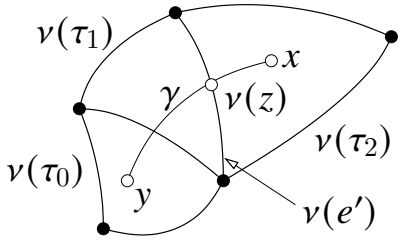

Figure 23 By walking a path $\gamma \subset \Sigma$ from y to x, we find a triangle on Σ that contains x.

Hence $\nu|_Q$ is an injection; hence $\nu|_Q$ is a bijection from Q to $\nu(Q)$. As $\nu|_{\tau_1}$ is continuous and has a continuous inverse over $\nu(\tau_1)$, and $\nu|_{\tau_2}$ is continuous and has a continuous inverse over $\nu(\tau_2)$, $\nu|_Q$ is continuous and has a continuous inverse over $\nu(Q)$. Therefore $\nu|_Q$ is a homeomorphism from Q to $\nu(Q)$. ◀

▶ **Lemma 50.** *Consider* $\mathrm{Vor}|_{\overline{\Sigma}} V$*, where* $\Sigma \subset \mathbb{R}^3$ *is a smooth 2-manifold without boundary and* $V \subset \Sigma$ *is a nonempty, finite sample. Suppose that for every site* $p \in V$ *and every point* $x \in \mathrm{Vor}|_{\overline{\Sigma}_S} p$*,* $|px| < \xi \,\mathrm{lfs}(p)$*, where* $\xi = \sqrt{(\sqrt{5}-1)/2} \doteq 0.786151$*. Moreover, suppose that for every extended Voronoi vertex* $u \in \mathrm{Vor}|_{\overline{\Sigma}_S} V$*, at least one of the following holds: either* $R \leq 0.3606\,\mathrm{lfs}(q)$ *or* u *is a principal vertex and* $R \leq 0.3202\,\mathrm{lfs}(u)$*, where* τ *is the restricted Delaunay triangle dual to* u*,* R *is the distance from* u *to each vertex of* τ*, and* q *is the vertex of* τ *at* τ*'s largest plane angle. Suppose that every extended Voronoi vertex in* $\mathrm{Vor}|_{\overline{\Sigma}} V$ *has degree three. Let* T *be the set of triangles in the restricted Delaunay triangulation* $\mathrm{Del}|_{\overline{\Sigma}} V$*. Then the nearest point map* $\nu : |T| \to \Sigma$ *is a surjection.*

Proof. Suppose for the sake of contradiction that some point $x \in \Sigma$ is not in $\nu(|T|)$. Let σ be the connected component of Σ that contains x. By Corollary 41, there are at least six sites on σ, so by Lemma 47, there is at least one extended Voronoi vertex on (the extended version of) σ; let $\tau_0 \in T$ be its dual restricted Delaunay triangle. By Theorem 46, for every triangle $\tau \in T$, $\nu|_\tau$ is a homeomorphism from τ to $\nu(\tau)$, so $\nu(\tau_0)$ is a topological disk that is closed with respect to Σ. Let y be a point in the interior of $\nu(\tau_0)$ that is not in V (not a site). As σ is a connected 2-manifold, there exists a directed path $\gamma \subset \sigma$ from y to x that does not intersect any site in V (except x, if x is a site), as illustrated in Figure 23.

As $\nu(\tau)$ is closed with respect to Σ for every $\tau \in T$, $\nu(|T|)$ is closed with respect to Σ. Let $z \in |T|$ be a point such that γ leaves $\nu(|T|)$ for the last time at $\nu(z)$, never to re-enter, as illustrated. By supposition, $z \neq x$ and z is not a site. Let $\tau_1 \in T$ be a triangle that contains z. As γ leaves $\nu(|T|)$ at $\nu(z)$, γ leaves $\nu(\tau_1)$ at $\nu(z)$ and z lies on the relative interior of

an edge e' of τ_1. There is an extended Voronoi edge e dual to e' (because every extended Voronoi vertex has degree three, and by Lemma 47, every extended Voronoi edge has distinct endpoints). One of e's vertices is dual to τ_1; let $\tau_2 \in T$ be the restricted Delaunay triangle dual to the other vertex. Let $Q = \tau_1 \cup \tau_2$. By Lemma 49, $\nu|_Q$ is a homeomorphism from Q to $\nu(Q)$, so $\nu(z)$ is in the interior of $\nu(Q)$. Therefore, where the path γ leaves $\nu(\tau_1)$ for the last time at $\nu(z)$, γ enters the interior of $\nu(\tau_2)$. This contradicts the claim that γ leaves $\nu(|T|)$ for the last time at $\nu(z)$. Therefore, for every point $x \in \Sigma$, $x \in \nu(|T|)$. ◀

G.7 The Nearest Point Map Is Injective

▶ **Lemma 51.** *Let V be a nonempty, finite sample of Σ. Let p be a site in V. Let W be the set of extended Voronoi vertices in $\mathrm{Vor}|_{\overline{\Sigma}} p$. Let $T_p \subseteq \mathrm{Del}|_{\overline{\Sigma}} V$ be the set of restricted Delaunay triangles that have vertex p—that is, the set of restricted Delaunay triangles dual to the vertices in W. Suppose that for every point $x \in \mathrm{Vor}|_{\overline{\Sigma}_S} p$, $|px| < \xi \,\mathrm{lfs}(p)$ where $\xi = \sqrt{(\sqrt{5}-1)/2} \doteq 0.786151$. Moreover, suppose that for each vertex $u \in W$, at least one of the following holds: $R \leq 0.3606\,\mathrm{lfs}(q)$, or u is a principal vertex and $R \leq 0.3202\,\mathrm{lfs}(u)$, where $\tau \in T_p$ is the restricted Delaunay triangle dual to u, R is the distance from u to each vertex of τ, and q is the vertex at τ's largest plane angle. Lastly, suppose that every vertex in W has degree three.*

Then the triangles in T_p intersect each other only at p and along their shared edges, and $|T_p|$ is a topological closed disk with p in its interior. Moreover, there exists an open neighborhood $N \subset |T_p|$ of p such that $\nu|_N$ is a homeomorphism from N to its image $\nu(N)$ on Σ.

Proof. By Theorem 40, $\mathrm{Vor}|_{\overline{\Sigma}_S} p$ is homeomorphic to a closed disk. By Lemma 35, the orthogonal projection of $\mathrm{Vor}|_{\overline{\Sigma}_S} p$ onto the tangent plane $T_p\Sigma$ is star-shaped.

Choose an arbitrary axis on $T_p\Sigma$ with origin p such that each point in $T_p\Sigma \setminus \{p\}$ can be assigned an angle counterclockwise from the axis in the range $[0°, 360°)$. The rotary direction deemed "counterclockwise" is consistent with p's orientation. For convenience, we use a rotary equivalence class of angles in which θ and $\theta + 360°$ denote the same angle; so, for instance, the range $[350°, 370°]$ denotes a 20° interval of angles, proceeding counterclockwise from 350° and stopping at 10°. We extend these assigned angles to $\mathbb{R}^3$: we assign any point in $\mathbb{R}^3$ the same angle as its orthogonal projection onto $T_p\Sigma$, unless the projected point is p (in which case its angle is undefined). The star-shaped property implies that no two principal vertices of $\mathrm{Vor}|_{\overline{\Sigma}_S} p$ have the same angle, and that the principal vertices of $\mathrm{Vor}|_{\overline{\Sigma}_S} p$ sorted by increasing angle match their counterclockwise ordering around the boundary of $\mathrm{Vor}|_{\overline{\Sigma}_S} p$.

Each Voronoi vertex u of $\mathrm{Vor}|_{\overline{\Sigma}} p$ has a dual triangle $\tau \in T_p$ such that, by Theorem 46, $\nu|_\tau$ is an orientation-preserving homeomorphism from τ to $\nu(\tau)$ and p's normal segment ℓ_p is not parallel to τ. Therefore, there is an angle $\psi > 0°$ such that $\angle(n_p, n_\tau) \leq 90° - \psi$ for every $\tau \in T_p$. For any $\tau \in T_p$, let $\bar{\tau}$ be the orthogonal projection of τ onto $T_p\Sigma$, and observe that $\bar{\tau}$ is also a triangle with straight edges. As $\angle(n_p, n_\tau) < 90°$, if the vertices p, p', and p'' of τ occur in counterclockwise order around the boundary of τ, then their projections $\bar{p}$, $\bar{p}'$, and $\bar{p}''$ on $T_p\Sigma$ occur in counterclockwise order around the boundary of $\bar{\tau}$. (Intuitively, projecting τ onto $T_p\Sigma$ does not "invert" the triangle.) If the angle assigned to p' and $\bar{p}'$ (both angles are the same) is ϕ, then the angle of p'' and $\bar{p}''$ is $\phi + \angle\bar{p}''\bar{p}\bar{p}'$, where $\angle\bar{p}''\bar{p}\bar{p}' \in (0°, 180°)$.

These facts hold not only for $T_p\Sigma$, but also if $T_p\Sigma$ is replaced by any plane whose normal vector n satisfies $\angle(n, n_p) < \psi$, because any such n also satisfies $\angle(n, n_\tau) < 90°$ for every $\tau \in T_p$.

Hence the counterclockwise ordering of the principal vertices around $\mathrm{Vor}|_{\widetilde{\Sigma}_S} p$ implies a counterclockwise ordering of the corresponding projected triangles around p. However, it does not imply that the triangles wind only once around p; we must eliminate the possibility that the triangles wind around p two or more times.

Observe that if an extended Voronoi vertex u is assigned an angle θ, the range of angles spanned by its dual restricted Delaunay triangle τ cannot include $\theta + 180°$, because if it did, the sphere with center u that passes through p could not enclose τ. Let $\theta_0, \theta_1, \ldots, \theta_{j-1}$ be the sorted angles of the j extended Voronoi vertices on the boundary of $\mathrm{Vor}|_{\widetilde{\Sigma}} p$. For each such angle θ_i, let ϕ_i and ϕ_{i+1} (where the subscript $i+1$ is interpreted modulo j) be the angles of the vertices (except p) of the corresponding dual restricted Delaunay triangle. For the triangles in T_p to wind around p two or more times, there must be an i such that $\theta_i + 180° \in [\phi_i, \phi_{i+1}]$. As this is impossible, the triangles wind around p only once.

Each projected triangle is confined to its own range of angles $[\phi_i, \phi_{i+1}]$, and the only points it shares with other triangles lie on the shared vertex p and on the shared edges at angles ϕ_i and ϕ_{i+1}. Therefore, the projection of $|T_p|$ onto $T_p\Sigma$ is an injection, and $|T_p|$ is a topological closed disk with p in its interior.

To show that ν restricted to some open neighborhood $N \subset |T_p|$ of p is a homeomorphism, we take advantage of the fact that for every plane Π whose normal vector n satisfies $\angle(n, n_p) < \psi$, the orthogonal projection of $|T_p|$ onto Π is an injection. Therefore, every line $\ell \subset \mathbb{R}^3$ such that $\angle(\ell, n_p) < \psi$ intersects $|T_p|$ in at most one point. We choose $N = |T_p| \cap B$ where B is an open ball centered at p, small enough that for every point $q \in \nu(N)$, $\angle(n_q, n_p) < \psi$. A sufficiently small ball satisfies this condition, because ν is continuous over $|T_p|$ and $\angle(n_q, n_p)$ is continuous for $q \in \Sigma$, so their composition is continuous. The condition guarantees that $\nu|_N$ is injective; hence $\nu|_N$ is a bijection from N to $\nu(N)$. As $\nu|_N$ is continuous and has a continuous inverse over $\nu(N)$, $\nu|_N$ is a homeomorphism from N to $\nu(N)$. ◀

▶ **Lemma 52.** *Consider* $\mathrm{Vor}|_{\widetilde{\Sigma}} V$*, where* $\Sigma \subset \mathbb{R}^3$ *is a smooth 2-manifold without boundary and* $V \subset \Sigma$ *is a nonempty, finite sample. Suppose that for every site* $p \in V$ *and every point* $x \in \mathrm{Vor}|_{\widetilde{\Sigma}_S} p$*,* $|px| < \xi\,\mathrm{lfs}(p)$*, where* $\xi = \sqrt{(\sqrt{5}-1)/2} \doteq 0.786151$*. Moreover, suppose that for every vertex* $u \in \mathrm{Vor}|_{\widetilde{\Sigma}} V$*, at least one of the following holds: either* $R \leq 0.3606\,\mathrm{lfs}(q)$*, or* u *is a principal vertex and* $R \leq 0.3202\,\mathrm{lfs}(u)$*, where* τ *is the restricted Delaunay triangle dual to* u*,* R *is the distance from* u *to each vertex of* τ*, and* q *is the vertex of* τ *at* τ*'s largest plane angle. Suppose that every extended Voronoi vertex of* $\mathrm{Vor}|_{\widetilde{\Sigma}} V$ *has degree three. Then the nearest point map* $\nu : |\mathrm{Del}|_{\widetilde{\Sigma}} V| \to \Sigma$ *is an injection.*

Proof. Suppose for the sake of contradiction that there are two distinct points $x, y \in |\mathrm{Del}|_{\widetilde{\Sigma}} V|$ such that $\nu(x) = \nu(y)$. By Corollary 41, there are at least six sites on each connected component of Σ, so there are at least two extended Voronoi vertices in each extended Voronoi cell; hence every restricted Delaunay vertex and every restricted Delaunay edge is a subset of some restricted Delaunay triangle. It follows that x lies on some restricted Delaunay triangle, and likewise for y.

Let T be the set of triangles in the restricted Delaunay triangulation $\mathrm{Del}|_{\widetilde{\Sigma}} V$. Let $\tau_x, \tau_y \in T$ be triangles containing x and y, respectively. Theorem 46 implies that no triangle in T contains both x and y; one implication is that $\tau_x \neq \tau_y$. By Corollary 26, no site intersects $\nu(\tau_x)$ except the three vertices of τ_x; hence y is not a site. Symmetrically, x is not a site.

Let p be a vertex of τ_x not shared by τ_y. Let $T_p \subset T$ be the set of restricted Delaunay triangles that have p for a vertex (including τ_x). Let $\gamma \subset \nu(\tau_x)$ be a directed path on Σ from $\nu(x) = \nu(y)$ to p such that $\gamma \setminus \{\nu(x), p\}$ lies in the relative interior of $\nu(\tau_x)$. By Lemma 51, there exists an open neighborhood $N \subset |T_p|$ of p such that the nearest point map $\nu|_N$ is a

homeomorphism from N to its image $\nu(N)$ on Σ. By Corollary 26, p does not intersect $\nu(\tau)$ for any restricted Delaunay triangle $\tau \in T \setminus T_p$, so we can assume without loss of generality that N is sufficiently small that $\nu(N)$ does not intersect $\nu(\tau)$ for any triangle $\tau \in T \setminus T_p$. Let $w \in \tau_x$ be a point such that $\nu(w) \in \gamma \cap \nu(N) \setminus \{\nu(x), p\}$. Observe that w is in the relative interior of τ_x, because $\gamma \setminus \{\nu(x), p\}$ is a subset of the relative interior of $\nu(\tau_x)$.

Let $T_{\neg x} = T \setminus \{\tau_x\}$, the set containing all the restricted Delaunay triangles except τ_x. We claim that there is a point $w' \in |T_{\neg x}|$ such that $\nu(w') = \nu(w)$. We establish the claim with essentially the same method used to prove Lemma 50. Suppose for the sake of contradicting this claim that $\nu(w) \notin \nu(|T_{\neg x}|)$. Observe that the path γ starts at $\nu(y)$, which is a subset of $\nu(|T_{\neg x}|)$ because $\tau_y \in T_{\neg x}$. Hence there exists a point $z \in |T_{\neg x}|$ such that γ leaves $\nu(|T_{\neg x}|)$ at $\nu(z)$ for the last time before reaching $\nu(w)$. Let $\tau_1 \in T_{\neg x}$ be a triangle that contains z. As γ leaves $\nu(|T_{\neg x}|)$ at $\nu(z)$, γ leaves $\nu(\tau_1)$ at $\nu(z)$, so z lies on the relative interior of an edge e' of τ_1. Let e be the extended Voronoi edge dual to e'. By Lemma 47, e has two distinct vertices. One of e's vertices is dual to τ_1; let $\tau_2 \in T$ be the restricted Delaunay triangle dual to the other vertex. As $\nu(z)$ lies on γ in the relative interior of $\nu(\tau_x)$ and $\nu(z)$ also lies on $\nu(e')$, e' is not an edge of τ_x. Therefore, $\tau_2 \neq \tau_x$ and $\tau_2 \in T_{\neg x}$. By Lemma 49, $\nu|_{\tau_1 \cup \tau_2}$ is a homeomorphism. Therefore, where the path γ leaves $\nu(\tau_1)$ for the last time at $\nu(z)$, γ enters the interior of $\nu(\tau_2)$. This contradicts the fact that γ leaves $\nu(|T_{\neg x}|)$ for the last time at $\nu(z)$. Hence, $\nu(w) \in \nu(|T_{\neg x}|)$; that is, there is a triangle $\tau_{w'} \in T_{\neg x}$ and a point $w' \in \tau_{w'}$ such that $\nu(w') = \nu(w)$.

As $\nu(w') \in \nu(N) \cap \nu(\tau_{w'})$ and $\nu(N)$ does not intersect $\nu(\tau)$ for any restricted Delaunay triangle $\tau \notin T_p$, $\tau_{w'} \in T_p$. But $\tau_{w'} \neq \tau_x$. By Lemma 51, the triangles in T_p intersect each other only at p and along their shared edges, so $\tau_{w'}$ does not intersect the relative interior of τ_x. As $w' \in \tau_{w'}$ and w is in the relative interior of τ_x, $w' \neq w$. Hence there exist two distinct points $w, w' \in N$ such that $\nu(w) = \nu(w')$, contradicting the fact that $\nu|_N$ is a homeomorphism. By this contradiction, we conclude that there are no two distinct points $x, y \in |\mathrm{Del}|_{\widetilde{\Sigma}} V|$ such that $\nu(x) = \nu(y)$, and therefore $\nu : |\mathrm{Del}|_{\widetilde{\Sigma}} V| \to \Sigma$ is an injection. ◀

▶ **Theorem 53.** *Consider* $\mathrm{Vor}|_{\widetilde{\Sigma}} V$*, where* $\Sigma \subset \mathbb{R}^3$ *is a smooth* 2*-manifold without boundary and* $V \subset \Sigma$ *is a nonempty, finite sample. Suppose that for every segment* $pq \in S$*,* $|pq| \leq 0.3647 \operatorname{lfs}(p)$*. Suppose that for every site* $p \in V$ *and every point* $x \in \mathrm{Vor}|_{\widetilde{\Sigma}_S} p$*,* $|px| < \xi \operatorname{lfs}(p)$*, where* $\xi = \sqrt{(\sqrt{5}-1)/2} \doteq 0.786151$*. Suppose that every extended Voronoi vertex in* $\mathrm{Vor}|_{\widetilde{\Sigma}} V$ *has degree three. Moreover, suppose that for every vertex* $u \in \mathrm{Vor}|_{\widetilde{\Sigma}} V$*, at least one of the following holds: either* $R \leq 0.3606 \operatorname{lfs}(q)$*, or* u *is a principal vertex and* $R \leq 0.3202 \operatorname{lfs}(u)$*, where* τ *is the restricted Delaunay triangle dual to* u*,* R *is the distance from* u *to each vertex of* τ*, and* q *is the vertex at* τ*'s largest plane angle. Then the nearest point map* $\nu : |\mathrm{Del}|_{\widetilde{\Sigma}} V| \to \Sigma$ *is a homeomorphism.*

Proof. By Lemma 50, $\nu : |\mathrm{Del}|_{\widetilde{\Sigma}} V| \to \Sigma$ is a surjection. By Lemma 52, it is an injection too. Hence it has an inverse defined over Σ. Both ν and its inverse are continuous, so it is a homeomorphism. ◀

This brings us to our main result, Theorem 8. Recall its statement:

Let V *be a constrained* ϵ*-sample of* (Σ, S, Z) *for some* $\epsilon \leq 0.3202$*. Suppose that for every segment* $pq \in S$*,* $|pq| \leq 0.3647 \operatorname{lfs}(p)$*. Suppose that every extended Voronoi vertex in* $\mathrm{Vor}|_{\widetilde{\Sigma}} V$ *has degree three. Suppose that for every restricted Delaunay triangle* τ *whose dual extended Voronoi face intersects an extrusion,* τ *satisfies* $r \leq 0.3606 \operatorname{lfs}(w)$*, where* r *is* τ*'s circumradius and* w *is the vertex of* τ *at* τ*'s largest plane angle. Then the nearest point map* $\nu : |\mathrm{Del}|_{\widetilde{\Sigma}} V| \to \Sigma$ *is a homeomorphism.*

Proof. For every point $x \in \overline{\Sigma}_S$, the nearest site $p \in V$ satisfies $|px| \leq 0.3202\,\mathrm{lfs}(x)$ by the definition of constrained ϵ-sample. By the Feature Translation Lemma (Lemma 12), $|px| \leq 0.3202/(1-0.3202)\,\mathrm{lfs}(p) < 0.48\,\mathrm{lfs}(p)$. Therefore, for every site $p \in V$ and every point $x \in \mathrm{Vor}|_{\overline{\Sigma}_S} p$, $|px| < \xi\,\mathrm{lfs}(p)$, satisfying one of the conditions of Theorem 53.

Let $u \in \mathrm{Vor}|_{\overline{\Sigma}} V$ be an extended Voronoi vertex and let $\tau \in \mathrm{Del}|_{\overline{\Sigma}} V$ be the restricted Delaunay triangle dual to u. If u lies on the principal surface $\overline{\Sigma}_S$, let s be the distance from u to any vertex of τ. As τ's vertices are the sites closest to u (that are visible from u), for every vertex q of τ, $s = |qu| < 0.3202\,\mathrm{lfs}(u)$, which satisfies a condition of Theorem 53. If u lies on an extrusion, let r be τ's circumradius and let w be the vertex of τ at τ's largest plane angle. By assumption, $r \leq 0.3606\,\mathrm{lfs}(w)$, which satisfies the remaining condition of Theorem 53.

By Theorem 53, $\nu : |\mathrm{Del}|_{\overline{\Sigma}} V| \to \Sigma$ is a homeomorphism. ◀