Separation of Critical Points and Flow-Complex-Based Surface Reconstruction

Tamal K. Dey*
Department of CSE
The Ohio State University
Columbus, OH 43210, USA

Joachim Giesen†
Theoretische Informatik
ETH Zürich
Switzerland

Edgar A. Ramos
Department of Computer Science
University of Illinois
Urbana, IL 61801, USA

Bardia Sadri
Department of Computer Science
University of Illinois
Urbana, IL 61801, USA

Abstract

The distance function to surfaces in three dimensions plays a key role in many geometric modeling applications such as medial axis approximations, surface reconstructions, offset computations, feature extractions among others. In many cases, the distance function induced by the surface can be approximated by the distance function induced by a discrete sample of the surface. The critical points of the distance functions are known to be closely related to the topology of the sets inducing them. However, no earlier theoretical result has found a link between topological properties of a geometric object and critical points of the distance to a discrete sample of it. We provide this link by showing that the critical points of the distance function induced by a discrete sample of a surface fall into two disjoint classes: those that lie very close to the surface and those that are near its medial axis. This closeness is precisely quantified and is shown to depend on the sampling density. It turns out that critical points near medial axis can be used to extract topological information about the sampled surface. With this result, we provide a new flow-complex-based surface reconstruction algorithm that, given a tight ε-sample of a surface, approximates the surface geometrically, both in distance and normals, and captures its topology. Furthermore, we show that the same algorithm can be used for curve reconstruction.

*Research partly supported by NSF grants DMS-0138456, DMS-0310642 and ARO grant DAAD19-02-1-0347.

†Research partially supported by the Swiss National Foundation under the project "Non-linear manifold learning".

1 Introduction

Given a compact surface Σ smoothly embedded in three dimensional Euclidean space $\mathbb{R}^3$, a distance function

$$h_\Sigma : \mathbb{R}^3 \to \mathbb{R}, \qquad x \mapsto \inf_{y \in \Sigma} \|x - y\|.$$

can be defined over $\mathbb{R}^3$ by assiging to each point in space its distance to Σ. h_Σ carries a great deal of information about Σ and its embedding. The surface Σ itself is trivially the zero level set $h_\Sigma^{-1}(0)$. Less trivial is the encoding of the embedding of Σ in space in the inner and outer components of its medial axis. For example, the homotopy type of the outer medial axis of an embedded torus can be used to determine if the the torus is knotted or not. It is well known that medial axis consists of exactly those points in space where h_Σ is not differentiable except for the points in Σ itself.

In many applications Σ is only known through a finite sample $P \subset \Sigma$ from which one desires to learn about Σ and its embedding. This is of course a reasonable goal only if P is a *dense* enough sample of Σ. There are various models for specifying the density of a sampling. A *uniform* ε-sample of a surface Σ is one that includes a sample point within ε distance from every point of Σ. In other words, the sample and surface must have Hausdorff distance of ε or less. For a uniform ε-sample to be capable of capturing the topology of the sampled surface, ε should be chosen proportional to the size of the smallest feature of the surface. As a result, a uniform ε-sample can be rather excessive. This has motivated the introduction of models that allow the sampling density to vary locally depending on the size of the features and therefore result much smaller samples.

The well-known ε-sampling theory of Amenta and Bern [1] is one of the most prominent models for relative sampling of surfaces. An ε-sample of a surface Σ in this model is one that contains a point within ε times the *local feature size* of every point $x \in \Sigma$. The local feature size of the point x is defined as the distance between x the medial axis $M(\Sigma)$ of the surface and appears to be a dependable measure for the required local density of sampling. The results of the present paper are based on this sampling model which we describe formally in section 2.

Sampling relative to local feature size has been used successfully to analyze algorithms that reconstruct Σ from P or approximate $M(\Sigma)$. Given an ε-sample of a smooth surface Σ, for a small enough value of ε, the algorithms of Amenta and Bern [1], Amenta, Choi, Dey and Leekha [3], Amenta, Choi and Kolluri [4], and Boissonnat and Cazals [5] reconstruct surfaces with the same topology as Σ that geometrically approximate its in terms of ε. Likewise, progress on the medial axis approximation problem was made in this model by the algorithm of Amenta, Choi and Kolluri [4] which succeds to capture the homotopy type of the medial axis $M(\Sigma)$ of the sampled surface Σ and the algorithm of Dey

and Zhao [10] that geometrically approximates $M(\Sigma)$ in terms of ε when ε approaches zero.

Considering the wealth of information encoded in h_Σ, when a sample P of Σ is at hand, it is natural to try to approximate h_Σ by the function

$$h_P : \mathbb{R}^3 \to \mathbb{R}, \qquad x \mapsto \min_{p \in P} \|x - p\|,$$

and try to extract the desired topological information from h_P instead of h_Σ. This idea has been used by Edelsbrunner [11], Chaine [6], and Giesen and John [12] to reconstruct Σ from P. Although all three of these algorithms work relatively well in practice, none provides a guarantee on the geometry and topology of the generated ouput.

The distance functions h_Σ and h_P are not smooth everywhere. Nevertheless, there is a well developed theory of critical points of such functions [13]. The critical points of h_Σ are all points in Σ (minima) plus a subset of its medial axis $M(\Sigma)$ (consisting of saddle points and maxima). All surface reconstruction algorithms based on h_P make use of its critical points, i.e., its local extrema and saddle points. These points are easily computable from the Delaunay triangulation of P [12]. A first contribution of our paper is to associate these critical points for an ε-sampling of Σ to either Σ itself or its medial axis M. We can show that for an ε-sampling of a surface Σ, for a small enough value of ε, all critical points of h_P either reside very close to Σ or rather very close to M. That is, we can label the critical points of h_P as either surface critical points if they are close to Σ or medial axis critical points if they are close to M. Interestingly, all types of critical points, including local maxima, can be close to Σ. This is particularly remarkable since the 2-skeleton of the Voronoi complex of P is exactly the medial axis of P, but its vertices, edges, or facets cannot be unambiguously assigned to either Σ or M even if ε becomes arbitrarily small. It is well know that Voronoi vertices can reside almost anywhere in $\mathbb{R}^3 \setminus P$.

This separation of the critical points of h_P leads to an algorithm for reconstruction of Σ from P by considering only medial axis critical points and regarding the surface critical points as an artifact of the discretization of Σ into P. This is the second main contribution of our paper. We can show that the reconstructed surface is homeomorphic to Σ and geometrically close to it both in distance and deviations of normals provided the input is a tight ε-sample of Σ. Similar results hold for curves and curve reconstruction.

The structure of the paper is as follows. Section 2 introduces the basic concepts including flow complex and critical points of the distance function induced by a point set. In section 3, we state and prove our separation of critical points lemma. Section 4 describes how it can be algorithmically determined for a critical point whether it is close to the surface or to the medial axis and uses this to build a surface reconstruction algorithm. Section 5 analyzes the quality of the produced reconstruction and establishes its geometric closeness and topological correctness. Section 6 studies the critical points of a smooth curve in $\mathbb{R}^3$ and

gives algorithms for classification of critical points and reconstruction of the curve analogous to those for the surface. An example of our experiments and concluding remarks are given in Section 7. Appendix A quotes necessary background information about the structure of stable manifolds of index 2 saddle points from [12].

2 Basic concepts

Throughout this paper Σ is a connected and compact smooth 2-manifold without boundary embedded in $\mathbb{R}^3$. Since it does not have a boundary, Σ separates $\mathbb{R}^3$ into a bounded region and an unbounded region. With a slight abuse of terminology we refer to the bounded region as the *interior* of Σ and to the unbounded region as its *exterior*. Since Σ is smooth, the normal to Σ at any point $x \in \Sigma$ is well defined. For $x \in \Sigma$, we denote by n_x^+ and n_x^-, the normal vectors at x pointing to the exterior and interior of Σ respectively. By n_x (with no "+" or "−" superscript) we denote either of n_x^+ or n_x^-, i.e. the direction of the line normal to Σ at x without a particular orientation. Also, throughout, $P \subset \Sigma$ is a discrete sample satisfying certain conditions to be specified shortly. To simplify our exposition we assume that P is in general position.

Any point set $S \subset \mathbb{R}^3$ induces a *distance function*

$$h_S : \mathbb{R}^3 \to \mathbb{R}, \qquad x \mapsto \inf_{p \in S} \|x - p\|,$$

where $\|\cdot\|$ denotes the Euclidean norm. It is easy to check that every distance function in the above sense is 1-Lipschitz, i.e. for all $x, y \in \mathbb{R}^3$, $|h_S(x) - h_S(y)| \leq \|x - y\|$. In this paper, we work with two major distance functions, one induced by Σ and the other by P. To simplify our notation, in the sequel, we use $s(\cdot)$ instead of $h_\Sigma(\cdot)$ and $h(\cdot)$ instead $h_P(\cdot)$.

2.1 Surface samples

The medial axis $M = M(\Sigma)$ of Σ is the set of all points in $\mathbb{R}^3$ that have at least 2 distinct closest points in Σ, i.e.

$$M = \left\{ x \in \mathbb{R}^3 : |\{y \in \Sigma : \|x - y\| = h_\Sigma(x)\}| \geq 2 \right\}.$$

For a point $c \in \mathbb{R}^3$ and real number r, the *ball* with center c and radius r, denoted $B_{c,r}$, is the set of all points $x \in \mathbb{R}^3$ at distance no more than r from c. We call a ball empty, if its interior does not contain any point from Σ. A *medial ball* is a maximal empty ball, i.e. an empty ball that is not contained in any other empty ball.

Medial feature size. For any point $x \in \mathbb{R}^3 \setminus \Sigma$ we denote by $\hat{x}$ the unique closest surface point to x, i.e.,

$$\hat{x} = \operatorname*{argmin}_{y \in \Sigma} \|x - y\|,$$

and by $\check{x} \in M$ we denote the center of the medial ball tangent to Σ at $\hat{x}$ and at the same side of Σ as x. The medial feature size is the function

$$\mu : \mathbb{R}^3 \setminus (\Sigma \cup M) \to \mathbb{R} \cup \{\infty\}, \qquad x \mapsto \|\hat{x} - \check{x}\|.$$

Note that for a point x in the unbounded component of $\mathbb{R}^3 \setminus \Sigma$, $\check{x}$ can be at infinity. This happens exactly when $\hat{x}$ lies on the boundary of the convex hull of Σ and the boundary of the medial ball tangent to Σ at $\hat{x}$ turns into the plane tangent to Σ at $\hat{x}$. In such a case we declare that the medial feature size of x is ∞.

Besides the medial feature size we will also use the function

$$m : \mathbb{R}^3 \setminus (\Sigma \cup M) \to \mathbb{R} \cup \{\infty\}, \qquad x \mapsto \|x - \check{x}\|,$$

which we call the *medial projection length*. Notice that for every $x \in \mathbb{R}^3 \setminus (\Sigma \cup M)$ we have the identity $\mu(x) = m(x) + s(x)$.

Feature size. The function

$$f : \Sigma \to \mathbb{R}, \qquad x \mapsto \inf_{y \in M} \|x - y\|,$$

which assigns to each point in Σ its distance to the medial axis M, is called the *local feature size*. Notice that for $x \in \mathbb{R}^3 \setminus (\Sigma \cup M)$ it always holds that $f(\hat{x}) \leq \mu(x)$. Notice also that $f(\cdot)$, being a distance function, is 1-Lipschitz, i.e. $|f(x) - f(y)| \leq \|x - y\|$ for all $x, y \in \Sigma$.

Sampling conditions. For a constant $\varepsilon > 0$, a finite sample $P \subset \Sigma$ is called an ε-sample if

$$\forall x \in \Sigma \ \exists p \in P \text{ such that } \|x - p\| \leq \varepsilon f(x).$$

An ε-sample P is called an (ε, δ)-sample or a *tight* ε-sample if it satisfies the additional condition

$$\forall p, q \in P \text{ it holds that } \|p - q\| \geq \delta f(p)$$

for some δ, with $0 < \delta < \varepsilon$.

Poles. For a sample point $p \in P$ we denote by V_p the closed Voronoi cell of p. If V_p is bounded, the *positive pole* of p, denoted p^+, is the Voronoi vertex of V_p farthest away from p. The *positive pole vector* ν_p^+ is the vector $p^+ - p$ if V_p is bounded or is taken as the unit vector in the direction which is the average of all unbounded Voronoi edges in V_p. In the latter case we informally refer to a

point at infinity in the direction ν_p^+ as the positive pole. The *negative pole* p^- of p is the farthest Voronoi vertex of V_p from p for which the smaller angle between the vectors ν_p^+ and $\nu_p^- = p^- - p$ is greater than than $\pi/2$ (or equivalently the inner product $\langle \nu_p^-, \nu_p^+ \rangle$ is negative). We call ν_p^- the *negative pole vector* at p.

Notation. The angle between two vectors u and v, denoted $\angle(u, v)$ is always smaller than π. For three points x, y, and z, we denote by $\angle xyz$ the angle between vectors $x - y$ and $z - y$, i.e., $\angle(x - y, z - y)$. The angle between two lines is the nonobtuse angle formed by them. The *acute angle* between vectors u and v is the smaller of the two angles made by the lines through u and v.

We will extensively use the following two lemmas due to Amenta and Bern [1] and the corollary following them.

Lemma 1 *Let x and y be points on Σ with $\|x - y\| \leq \xi f(x)$ for $\xi \leq 1/3$. Then*

$$\angle(n_x^+, n_y^+) = \angle(n_x^-, n_y^-) \leq \frac{\xi}{1 - 3\xi}.$$

Lemma 2 *Let p be a sample point in an ε-sample P. Let x be any point in V_p with $\|x - p\| \geq \xi f(p)$ for $\xi > 0$. Then*

$$\angle(x - p, n_p) \leq \arcsin \frac{\varepsilon}{1 - \varepsilon} + \arcsin \frac{\varepsilon}{\xi(1 - \varepsilon)}.$$

Corollary 1 *For any point p of an ε-sampling P of a surface Σ, the acute angle between the normal to surface at p, n_p, and either of ν_p^+ and ν_p^- is at most $2 \arcsin(\varepsilon/(1 - \varepsilon))$.*

2.2 Induced flows

Critical points. Our results involve the critical points, i.e., its local extrema and saddle points of the distance function h. For every point $x \in \mathbb{R}^3$ we define

$$A(x) = \{p \in P : \|x - p\| = h(x)\}$$

as the set of all sample points closest to x. There is a simple characterization for critical points of h [12], namely, a point c is a critical point if and only if $c \in \operatorname{conv} A(c)$. It turns out that these points are exactly the intersection points of Voronoi faces and their dual Delaunay simplices. A local maximum is a Voronoi vertex that is contained in its dual Delaunay tetrahedron. Every sample point is a minimum and there are no other minima. The remaining critical points are saddle points. The dimension of the Delaunay simplex that contains a critical point c is interpreted as the *index* of c. We call a non-critical point of h a *regular* point.

Flow. Although the gradient of h is not defined everywhere, there is a unique direction of steepest ascent at every regular point of h. The direction of steepest ascent at $x \in \mathbb{R}^3$ is given by the vector $x - d(x)$ where

$$d(x) = \operatorname*{argmin}_{y \in \operatorname{conv} A(x)} \|x - y\|.$$

We call the point $d(x)$ the *driver* of x. Assigning to the critical points of h the zero vector and to every other point in $\mathbb{R}^3$ the unique unit vector in the direction of steepest ascent results a vector field

$$v : \mathbb{R}^3 \to \mathbb{S}^2 \cup \{0\}$$

on $\mathbb{R}^3$. This vector field is not continuous but nevertheless can be integrated to give a *flow*, i.e., a map

$$\phi : [0, \infty) \times \mathbb{R}^3 \to \mathbb{R}^3,$$

such that at every point $(t, x) \in [0, \infty) \times \mathbb{R}^3$ the right derivative

$$\lim_{t \leftarrow t'} \frac{\phi(t', x) - \phi(t, x)}{t' - t}$$

exists and equals $v(\phi(t, x))$. The flow tells how a point would move if it always followed the direction of steepest ascent of the distance function h. The flow curve (integral line) that a point x follows is given by $\phi_x : \mathbb{R} \to \mathbb{R}^3, t \mapsto \phi(t, x)$ and is called the *orbit* of x. We denote by $\phi(x)$, the set $\{\phi_x(t) : t \in [0, +\infty)\}$.

Stable manifolds. Given a critical point c of h the set of all points whose orbit ends in c, i.e. the set of all points that flow into c, is called the *stable manifold* of c and is denoted $S(c)$. The collection of all stable manifolds forms a cell complex which is known as the *flow complex*. The dimension of each cell in the flow complex is the index of its associated critical point. The cells have a recursive structure, namely, the boundary of the stable manifold of a critical point is made up of stable manifolds of critical points of lower index. As in [12] we assume throughout this paper that Voronoi and their dual Delaunay faces intersect in their interiors if they intersect at all. Other intersections are unstable under small perturbation of the point set and can therefore be considered degenerate. Here we summarize the basic facts of the stable manifolds for the different indices of the critical points.

INDEX-0. The stable manifold of an index-0 critical point, i.e., a local minimum, is just the minimum itself.

INDEX-1. The stable manifold of an index-1 critical point, also called a *1-saddle*, i.e., the intersection point of a Delaunay edge with its dual Voronoi facet, is the Delaunay edge which in this case is a Gabriel edge.

INDEX-2. The stable manifolds of an index-2 critical point, also called a *2-saddle*, is a piecewise linear surface patch. See [12] and Appendix A for details on structure and computation of these patches.

INDEX-3. The stable manifolds of index-3 critical points, i.e. local maxima, are the bounded regions in the complex built by the stable manifolds of critical points of indices 0, 1 and 2.

3 Separation of critical points

In the following P is always an ε-sample (with ε to be specified) of a smooth closed surface Σ embedded in $\mathbb{R}^3$. Also, h is the distance function associated with P and ϕ is the flow induced by P following the vector field v.

Lemma 3 *Let x be a point in $\mathbb{R}^3 \setminus (\Sigma \cup M)$ with $\mu(x) = \infty$. Then, x is not a critical point of h and the angle between the vectors $\check{x} - x$ and $v(x)$ is strictly less than $\pi/2$.*

Proof. If $m(x) = \infty$ then $\check{x}$ is at infinity and the plane H tangent to Σ at $\hat{x}$ does not have any point from Σ on the same side as $\hat{x}$. Therefore, H separates x from Σ and in particular from $\operatorname{conv} A(x)$. Consequently, x cannot be a critical point of h and for every point y except $\hat{x}$, on the ray from $\hat{x}$ through x, the angle between the vectors $v(y)$ and $\hat{y} - y$ is strictly less than $\pi/2$. □

The next lemma states that for $x \in \mathbb{R}^3 \setminus (\Sigma \cup M)$ with $\mu(x), m(x) < \infty$ any critical point of h on the line segment $\hat{x}\check{x}$ must be very close to one of the two ends of this segment.

Lemma 4 *Let $\varepsilon < 1/3$ and let x be a point in $\mathbb{R}^3 \setminus (\Sigma \cup M)$ for which $\mu(x)$ is bounded. Moreover assume that $2\varepsilon\mu(x) < m(x)$ and $\varepsilon^2 f(\hat{x}) < s(x)$. Then $v(x) \neq 0$ and $\angle(\check{x} - x, v(x)) < \pi/2$.*

Proof. By the definition of local feature size $f(\hat{x}) \leq \mu(x)$. Thus the closest sample point in P to $\hat{x}$ lies inside a ball centered at $\hat{x}$ with radius at most $\varepsilon f(\hat{x})$. Hence the distance from x to its closest sample point is at most $s(x) + \varepsilon f(\hat{x})$. Consequently, the set $A(x) \subseteq P$, of sample points at minimum distance from x, is contained in the ball B centered at x with radius $s(x) + \varepsilon f(\hat{x})$. Let B' be the open ball centered at $\check{x}$ with radius $\mu(x)$. Since B' is empty of any sample points, every point of $A(x)$ is contained in $B \setminus B'$. The driver $d(x)$ of the flow induced by P at x is by definition in the convex hull of $A(x)$ and is therefore contained in the convex hull of $B \setminus B'$.

Consider the disk $D = H \cap B'$ where H is the plane containing x perpendicular to the line through $\hat{x}$ and $\check{x}$. Let r be the radius of this disk. If $s(x) + \varepsilon f(\hat{x}) < r$, then $x \notin \operatorname{conv}(B \setminus B')$. Consequently, x is not a critical point of h and the angle between the vectors $\check{x} - x$ and $v(x)$ is strictly less than $\pi/2$, see Figure 1. Thus for the statement of the lemma to hold, it suffices to show that $s(x) + \varepsilon f(\hat{x}) < r$. By the Pythagorean theorem

$$r^2 = \mu(x)^2 - m(x)^2.$$

Therefore, since

$$s(x) + \varepsilon f(\hat{x}) = \mu(x) - m(x) + \varepsilon f(\hat{x}),$$

$s(x) + \varepsilon f(\hat{x}) < r$ is equivalent to

$$(\mu(x) - m(x) + \varepsilon f(\hat{x}))^2 < \mu(x)^2 - m(x)^2,$$

which in turn is equivalent to,

$$m(x)^2 - \big(\mu(x) + \varepsilon f(\hat{x})\big)m(x) + \varepsilon\mu(x)f(\hat{x}) + \frac{\varepsilon^2}{2}f(\hat{x})^2 < 0.$$

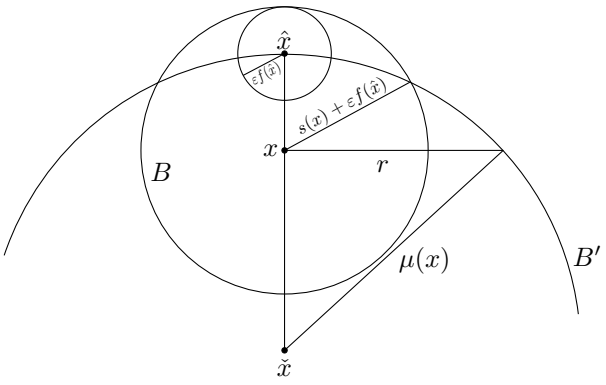

Figure 1: Flow can escape the center of the medial ball only close to surface or close to center.

This inequality holds between the two roots of the quadratic function of $m(x)$ on the left hand side of the inequality, i.e., for

$$\begin{aligned} m(x) &> \frac{1}{2}\Big(\mu(x) + \varepsilon f(\hat{x}) - \\ &\qquad \sqrt{\mu(x)^2 - \varepsilon^2 f(\hat{x})^2 - 2\varepsilon\mu(x)f(\hat{x})}\Big) \\ m(x) &< \frac{1}{2}\Big(\mu(x) + \varepsilon f(\hat{x}) + \\ &\qquad \sqrt{\mu(x)^2 - \varepsilon^2 f(\hat{x})^2 - 2\varepsilon\mu(x)f(\hat{x})}\Big). \end{aligned}$$

Using $f(\hat{x}) < \mu(x)$ the lower bound on $m(x)$ can be weakened as follows

$$\begin{aligned} m(x) &> \frac{1}{2}\Big(\mu(x) + \varepsilon\mu(x) - \\ &\qquad \sqrt{\mu(x)^2 - \varepsilon^2\mu(x)^2 - 2\varepsilon\mu(x)^2}\Big) \\ &= \frac{\mu(x)}{2}\left(1 + \varepsilon - \sqrt{1 - 2\varepsilon - \varepsilon^2}\right). \end{aligned}$$

Notice that our assumption that $\varepsilon < 1/3$ implies

$$0 < 1 - 2\varepsilon - \varepsilon^2 < 1,$$

i.e., we have a real lower bound on $m(x)$ and can further weaken this bound, taking into account that $\sqrt{x} > x$ for $x \in (0,1)$, to get

$$\begin{aligned} 1 + \varepsilon - \sqrt{1 - 2\varepsilon - \varepsilon^2} &< 1 + \varepsilon - (1 - 2\varepsilon - \varepsilon^2) \\ &= 3\varepsilon + \varepsilon^2 < 4\varepsilon, \end{aligned}$$

which implies the weaker lower bound $m(x) > 2\varepsilon\mu(x)$. Next we want to show that $s(x) > \varepsilon^2 f(\hat{x})$ implies the upper bound on $m(x)$, i.e., $s(x) < \varepsilon^2 f(\hat{x})$ is a weaker form of this upper bound. Assuming $s(x) > \varepsilon^2 f(\hat{x})$ we get

$$m(x) = \mu(x) - s(x) < \mu(x) - \varepsilon^2 f(\hat{x}).$$

Thus it is enough to show that

$$\begin{aligned} \mu(x) - \varepsilon^2 f(\hat{x}) < \frac{1}{2}\Big(&\mu(x) + \varepsilon f(\hat{x}) + \\ &\sqrt{\mu(x)^2 - \varepsilon^2 f(\hat{x})^2 - 2\varepsilon\mu(x) f(\hat{x})}\Big). \end{aligned}$$

The latter inequality is equivalent to

$$\begin{aligned} -\frac{3}{4}\mu(x)^2 + \left(\frac{3}{2}\varepsilon - \varepsilon^2\right)\mu(x) f(\hat{x}) + \\ \left(\frac{5}{4}\varepsilon^2 + \varepsilon^3 + \varepsilon^4\right) f(\hat{x})^2 < 0. \end{aligned}$$

Plugging in $f(\hat{x}) \leq \mu(x)$ we get the stronger inequality

$$-\frac{3}{4} + \frac{3}{2}\varepsilon + \frac{1}{4}\varepsilon^2 + \varepsilon^3 + \varepsilon^4 < 0,$$

which in turn gives by summarizing

$$\frac{1}{4}\varepsilon^2 + \varepsilon^3 + \varepsilon^4 \text{ into } \frac{9}{4}\varepsilon^2$$

the even stronger inequality

$$-1 + 2\varepsilon + 3\varepsilon^2 < 0$$

which is satisfied through our assumption that $\varepsilon < 1/3$. Thus $s(x) > \varepsilon^2 f(\hat{x})$ implies the upper bound on $m(x)$. □

η-tubular neighborhoods. For a constant $\eta \leq 1$, define the η-neighborhood of the medial axis, denoted M_η as

$$M_\eta = \{x \in \mathbb{R}^3 \setminus (\Sigma \cup M) \,:\, m(x) < \eta\mu(x)\} \cup M.$$

Similarly let

$$\Sigma_\eta = \{x \in \mathbb{R}^3 \setminus (\Sigma \cup M) \,:\, s(x) < \eta f(\hat{x})\} \cup \Sigma.$$

Corollary 2 *Every critical point of the distance function h either belongs to Σ_{ε^2} or $M_{2\varepsilon}$ provided $\varepsilon < 1/3$.*

Surface and medial axis critical points. Let $\varepsilon < 1/3$. We call a critical point of h a *surface critical point* if it is contained in Σ_{ε^2} and we call it a *medial axis critical point* if it is contained in $M_{2\varepsilon}$.

4 Algorithms

4.1 Separating the critical points

Our goal here is to devise an algorithm that can separate the surface critical points from the medial axis critical points. To prove the correctness of our algorithm we will frequently make use of the following definition.

Associating critical point to sample points. A critical point c is said to be *associated* with a sample point $p \in P$ if $\|p - c\| = h(c)$, i.e., if c is contained in the closed Voronoi cell of p.

At first we want to deal with surface critical points. The following two lemmas turn out to be useful to this end.

Lemma 5 *For any sample point $p \in P$ the ball B with diameter pp^+ does not contain any critical point associated with p in its interior. The analogous statement holds for the negative pole p^-.*

Proof. If p^+ lies at infinity then the ball B becomes a half-space with normal ν_p^+. The boundary of this half-space is a plane that supports the convex hull of P. Since any critical point of h must be contained in the convex hull of P it follows that the the interior of B, i.e., the open half space, cannot contain any critical point of h. Thus we can assume that p^+ is finite. Let c be a critical point associated with p. If c is a minimum, then $c = p$ and there is nothing to prove. Otherwise, $|A(c)| > 1$. All points in $A(c)$ lie on the boundary of the ball B' of radius $\|c-p\|$ centered at c. Let B'' be the open ball of radius $\|p-p^+\|$ centered at p^+. By construction there can be no points of P in B''. Thus all points of $A(c)$ must belong to $\partial B' \setminus B''$. On the other hand, for c to be a critical point, it must be in the convex hull of $A(c)$. This happens only if the angle $\angle pcp^+$ is smaller than $\pi/2$. The latter condition is in turn identical to c being outside the ball B, see Figure 2. The proof of the analogous statement for p^- follows along the same lines. □

Lemma 6 *Let c be a surface critical point associated with sample point $p \in P$. If $\varepsilon < 0.1$ then*

$$\|c - p\| \leq \frac{1.1\varepsilon}{1 - 1.2\varepsilon}\|p - p^-\| \leq \frac{1.1\varepsilon}{1 - 1.2\varepsilon}\|p - p^+\|.$$

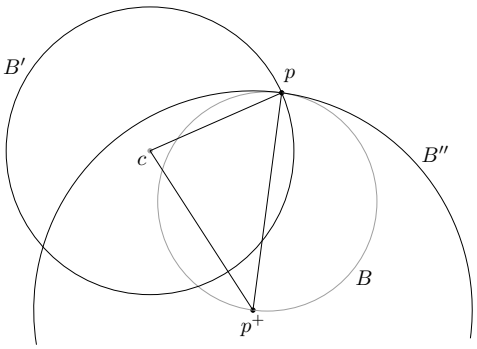

Figure 2: The ball with diameter pp^+ contains no critical point associated with p.

Proof. Let $q \in P$ be the closest sample point to $\hat{c}$. By the sampling condition it holds that $\|\hat{c} - q\| \leq \varepsilon f(\hat{c})$. Using the triangle inequality we obtain

$$\begin{aligned} \|c - p\| &\leq \|c - q\| \leq s(c) + \|\hat{c} - q\| \\ &\leq \varepsilon^2 f(\hat{c}) + \varepsilon f(\hat{c}) = \varepsilon(1+\varepsilon) f(\hat{c}). \end{aligned} \tag{1}$$

Since the local feature size f is 1-Lipschitz we can write

$$\begin{aligned} f(\hat{c}) &\leq f(p) + \|\hat{c} - p\| \\ &\leq f(p) + s(c) + \|c - p\| \\ &< f(p) + \varepsilon^2 f(\hat{c}) + \varepsilon(1+\varepsilon) f(\hat{c}) \\ &= f(p) + \varepsilon(1+2\varepsilon) f(\hat{c}) \\ &< f(p) + 1.2\varepsilon f(\hat{c}). \end{aligned}$$

Rearranging, we obtain

$$f(\hat{c}) < \frac{1}{1 - 1.2\varepsilon} f(p). \tag{2}$$

Combining (1) and (2) we finally get

$$\begin{aligned} \|c - p\| &\leq \varepsilon(1+\varepsilon) f(\hat{c}) < \frac{\varepsilon(1+\varepsilon)}{1 - 1.2\varepsilon} f(p) \\ &< \frac{1.1\varepsilon}{1 - 1.2\varepsilon} f(p) \leq \frac{1.1\varepsilon}{1 - 1.2\varepsilon} \|p - p^-\| \\ &\leq \frac{1.1\varepsilon}{1 - 1.2\varepsilon} \|p - p^+\|. \end{aligned}$$

□

From Lemmas 5 and 6 we derive the following corollary.

Corollary 3 *Let c be a surface critical point associated with sample point $p \in P$. If $\varepsilon < 0.1$ then then the acute angle between the vector $c - p$ and each of the vectors ν_p^- and ν_p^+ is at least 75.5 degrees.*

Proof. By Lemma 5 neither of the two balls with diameters pp^+ and pp^- respectively contain any critical points associated with p. Thus we get from using Lemma 6 and Thales theorem that

$$\cos \angle(c-p, \nu_p^+) \leq \frac{1.1\varepsilon}{1-1.2\varepsilon}.$$

The same bound holds for $\cos \angle(c-p, \nu_p^-)$. Thus for $\varepsilon < 0.1$, the resulting angles are between 82.5 and 180 degrees. By Corollary 1,

$$\angle(\nu_p^+, \nu_p^-) \geq \pi - 4 \arcsin \frac{\varepsilon}{1-\varepsilon}.$$

Since $\varepsilon < 0.1$ this angle is at least 173 degrees. Thus

$$\angle(c-p, \nu_p^+) \leq 360 - 173 - 82.5 = 104.5 \text{ degrees},$$

with the same bound holding for $\angle(c-p, \nu_p^-)$. From this, both $\angle(c-p, \nu_p^+)$ and $\angle(c-p, \nu_p^-)$ and their complements are at least $180 - 104.5 = 75.5$ degrees. □

Next we deal with medial axis critical points.

Lemma 7 *Let c be a medial axis critical point associated with $p \in P$. If $\varepsilon < 1/3$ then $\|p-c\| \geq (1-2\varepsilon) f(p)$.*

Proof. We have $s(c) \leq \|c-p\|$ since $s(c)$ is the distance from c to Σ and we have $m(c), \mu(c) < \infty$ by Lemma 3. From the definition of a medial axis critical point we have

$$m(c) \leq 2\varepsilon\mu(c) = 2\varepsilon(m(c) + s(c)).$$

Combining these inequalities we obtain

$$m(c) \leq \frac{2\varepsilon}{1-2\varepsilon} s(c) \leq \frac{2\varepsilon}{1-2\varepsilon} \|c-p\|,$$

which gives

$$\begin{aligned} f(p) &\leq \|\check{c}-p\| \leq m(c) + \|c-p\| \\ &\leq \left(\frac{2\varepsilon}{1-2\varepsilon} + 1\right) \|c-p\| = \frac{1}{1-2\varepsilon} \|c-p\| \end{aligned}$$

and thus $\|p-c\| \geq (1-2\varepsilon) f(p)$. □

Corollary 4 *Let c be a medial axis critical point associated with $p \in P$. If $\varepsilon < 0.1$ then the acute angle between the vector $c-p$ and each of the vectors ν_p^- and ν_p^+ is at most 28 degrees.*

Proof. Using Lemma 7 we can plug in $1-2\varepsilon$ for ξ in Lemma 2 to obtain that the acute angle between $c-p$ and n_p is at most

$$\arcsin \frac{\varepsilon}{1-\varepsilon} + \arcsin \frac{\varepsilon}{(1-2\varepsilon)(1-\varepsilon)}.$$

```
RECONSTRUCT(P)
 1  C := set of the critical points of h.
 2  (C_M, C_Σ) := SEPARATE(P, C).
 3  for each c ∈ C_M do
 4     U.make(c).
 5  for each c ∈ C_M do
 6     for all c' ∈ ∂S(c) ∩ C_M
 7        U.union(c, c').
 8  O := union of all stable manifolds in one component of U.
 9  T := ∂O.
10  return T
```

Figure 3: The surface-reconstruction algorithm.

By Corollary 1, the acute angle between n_p and each of ν_p^+ or ν_p^- is at most $2\arcsin(\varepsilon/(1-\varepsilon))$. Thus the acute between $c-p$ and each of ν_p^+ or ν_p^- is at most

$$3\arcsin\frac{\varepsilon}{1-\varepsilon}+\arcsin\frac{\varepsilon}{(1-2\varepsilon)(1-\varepsilon)},$$

which amounts to less than 28 degrees when $\varepsilon < 0.1$. □

Corollary 3 and Corollary 4 show that for a critical point c associated with a sample point p, the angle between $c-p$ and a pole of p falls into one of the two disjoint ranges of 0 to 45 degrees or 75.5 to 90 degrees, depending on whether c is a surface or a medial axis critical point, respectively. Thus the two types of critical points can be distinguished by simply measuring this angle.

4.2 Reconstruction

The algorithmic classification of the critical points of h as either surface or medial axis critical points suggests the algorithm of Figure 3 for reconstruction of Σ from P. By SEPARATE(P, C) we refer the to the algorithm described above that partitions the input set of critical points C into two sets C_M, of medial axis critical points, and C_Σ, of surface critical points, and returns the pair (C_M, C_Σ).

The algorithm RECONSTRUCT maintains a Union-Find data structure U on the set of medial axis critical points. In line 4, U.make(c) adds a singleton set $\{c\}$ to U and in line 7, U.union(c, c') combines the sets containing c and c' into a single set. The Union-Find data structure is used to find all connected components of stable manifolds $S(c)$ of medial axis critical points $c \in C_M$. In the end the boundary of one arbitrary component is returned. Notice that this boundary is made of stable manifolds of surface critical points.

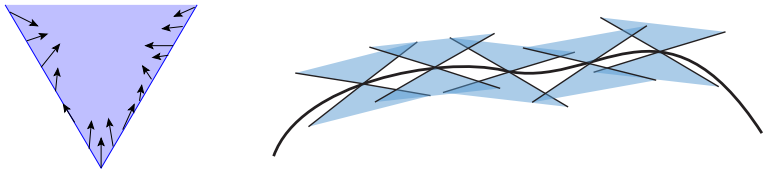

Figure 4: A sink cone (left) and enclosing of the surface with sink cone envelopes (right).

5 Reconstruction Properties

In this section, we give geometric and topological guarantees for the output of the algorithm RECONSTRUCT under (ε, δ)-sampling. We summarize the results in the following Theorem.

Theorem 1 *For any $0 < \rho < 1$ there exists ε_0 such that given an (ε, δ)-sample P from a smooth closed surface Σ, where $\varepsilon < \varepsilon_0$ and $\delta/\varepsilon = \rho$, the algorithm* RECONSTRUCT *outputs a sub-complex $\mathcal{T}$ of the flow complex of P with the following properties:*

(i) *$\mathcal{T}$ is contained in the tubular neighborhood $\Sigma_{3\varepsilon^2}$.*

(ii) *The normal to every triangle pqr in $\mathcal{T}$, with $p \in P$, forms an angle of $O(\varepsilon)$ with the normal to Σ at p, n_p.*

(iii) *$\mathcal{T}$ is homeomorphic (in fact, isotopic) to Σ.*

In particular these claims hold for $\rho \geq 1/3$ and $\varepsilon_0 \leq 0.01$.

5.1 Closeness

To analyze local geometry of the flow near the surface, we place at sample points $p \in P$ cones that open along inner and outer normal directions at p. We show that, under certain conditions, such cones are *sinks*, i.e., on their surfaces and sufficiently close to Σ, the flow is either tangential or points to the inside of the cones. By overlapping together these close-reaching cones (see Figure 4), we obtain inner and outer envelopes that enclose the surface and are in a sense "one-way" for the flow. This means that flow cannot escape from these envelopes leading to properties of the flow complex central to the analysis of the output of our algorithm.

Sink Cones. For a point p, and a direction vector n, let $C = \text{cone}(p, n, \theta, r)$ be the *cone-patch* consisting of points x for which $\|x - p\| \leq r$ and $\angle(n, x - p) = \theta$. We call θ, and r, the *angle* and the *reach* of C, respectively. A cone-patch is

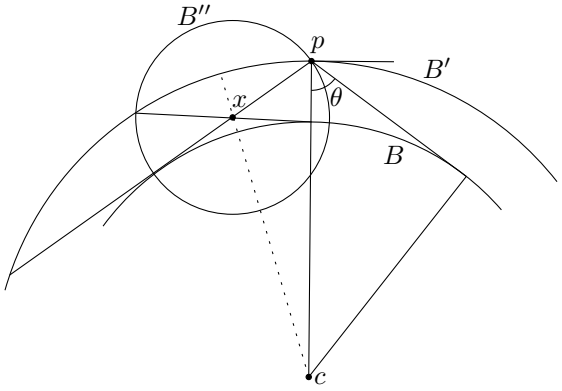

Figure 5: Sink cone for a sample point p.

essentially part of the surface of an infinite cone. With a slight abuse of notation we refer to a cone-patch also as a cone. The *boundary* of C consists of points $x \in C$ for which $\|x - p\| = r$. We say that C is a *sink* if at every point in the relative interior of C (thus excluding the boundary of C) the flow is either tangential or directed toward the interior of the convex hull of C.

Lemma 8 *For any point $p \in P$, the two cones $C_p^+ = \mathrm{cone}(p, n_p^+, \theta, r)$ and $C_p^- = \mathrm{cone}(p, n_p^-, \theta, r)$ are sinks for any $0 < \theta < \frac{\pi}{2}$ and $r = f(p)\cos\theta$. Furthermore, the interior of the convex hulls of C_p^+ and C_p^- do not contain any critical points.*

Proof. We only prove that C_p^+ is a sink. The proof for C_p^- is similar. Let c be the point on the ray in direction n_p^+ at distance $f(p)$ from p. Let B' be the ball with radius $f(p) = \|p - c\|$ centered at c. Note that B' cannot contain any points from Σ in its interior since it is contained in a medial ball. Let B be the ball centered at c with radius $f(p)\sin\theta$. Then C_p^+ is the cone tangent to B with apex p (see Figure 5). For every $x \in C_p^+$, the set $A(x)$ of closest sample points to x are inside the ball B'' of radius $\|x-p\|$ centered at x but outside B'. Let H be the plane tangent to C_p^+ at x. Since $r = f(p)\cos\theta$ we have that $\mathrm{conv}(B'' \setminus B')$ is entirely on the opposite side of H with respect to B and therefore, $v(x)$ points toward the interior of C_p^+. Thus C_p^+ is a sink and no point in the relative interior of C_p^+ is a critical point. Every point in the interior of the convex hull of C_p^+ is on a cone $\mathrm{cone}(p, n_p^+, \theta', r)$ for some $\theta' < \theta$ or in the relative interior of the line segment pc. As we have seen the points on the cones cannot be critical. But neither can be any point y in the relative interior of the line segment pc since for such a point $A(y) = \{p\}$ and $y \neq p$. □

Fixing cone angles. Notice that the above lemma puts at every sample point two sink cones with the same apex, angle, reach, and axis but in opposite direction. For the rest of the paper, we shall consider only such cones with a fixed cone angle that depends only on the density of sampling. Indeed, we fix

$\theta = \theta(\varepsilon) = \frac{\pi}{2} - 1.1\varepsilon$ and we respectively denote the outer and inner cones at a vertex p by $C_p^- = \text{cone}(p, n_p^-, \theta, f(p)\cos\theta)$ and $C_p^+ = \text{cone}(p, n_p^+, \theta, f(p)\cos\theta)$. We also denote by C_p the union of the two cones C_p^+ and C_p^-.

Lemma 9 *For any point $x \in V_p \cap \Sigma$, the ray shot in direction n_x^+ (n_x^-) hits C_p^+ (C_p^-) provided that $\varepsilon \leq 0.05$.*

Proof. Let β be the angle between n_x^+ and n_p^+. Since $x \in V_p \cap \Sigma$, we have $\|x - p\| \leq \varepsilon f(x)$ and therefore $\beta \leq \frac{\varepsilon}{1-3\varepsilon} < \theta$. Therefore, the ray shot from x in direction of n_x^+ hits $\text{cone}(p, n_p^+, \theta, \infty)$, the infinite extension of C_p^+, at some point x'. Let $\eta = \|x' - p\|/f(p)$. If $\eta \leq \cos\theta$, then $x' \in C_p^+$ and there is nothing to prove. Otherwise $\|x' - p\| > f(p)\cos\theta$. It can be easily observed that the closest point to p on the line through x' and x is at distance no less than

$$\|x' - p\|\sin(\theta - \beta) > \cos\theta \sin(\theta - \beta) f(p)$$

from p. On the other hand, since $x \in V_p \cap \Sigma$, $\|x - p\| \leq \frac{\varepsilon}{1-\varepsilon} f(p)$. This implies

$$\begin{aligned}
\frac{\varepsilon}{1-\varepsilon} f(p) &\geq \|x - p\| \\
&> \cos\theta \sin(\theta - \beta) f(p) \\
&= \sin(1.1\varepsilon)\cos(1.1\varepsilon + \beta) f(p) \\
&\geq \sin(1.1\varepsilon)\cos\left(1.1\varepsilon + \frac{\varepsilon}{1-3\varepsilon}\right) f(p)
\end{aligned}$$

which is a contradiction for $\varepsilon \leq 0.05$. □

Lemma 10 *Let x be a point on C_p. Then $\|x - \hat{x}\| \leq 3\varepsilon^2 f(\hat{x})$ when $\varepsilon \leq 0.05$.*

Proof. Without loss of generality assume that $x \in C_p^+$ (the proof for the case where $x \in C_p^-$ is similar). Shoot a ray from x, parallel to n_p^- until it hits C_p^- at a point x'. Since each of C_p^+ and C_p^- is completely contained in its corresponding medial ball tangent to Σ at p, the line segment xx' intersects Σ. Therefore

$$\|x - \hat{x}\| \leq \|x - x'\| = 2\|x - p\|\cos\theta \leq 2f(p)\cos^2\theta. \tag{3}$$

On the other hand, by the triangle inequality,

$$\begin{aligned}
\|\hat{x} - p\| &\leq \|x - \hat{x}\| + \|x - p\| \\
&\leq 2f(p)\cos^2\theta + f(p)\cos\theta \\
&\leq 3f(p)\cos\theta.
\end{aligned}$$

Since $f(\cdot)$ is 1-Lipschitz

$$\begin{aligned}
f(\hat{x}) &\geq f(p) - \|\hat{x} - p\| \\
&\geq f(p) - 3f(p)\cos\theta,
\end{aligned}$$

which gives us

$$f(p) \leq \frac{f(\hat{x})}{1 - 3\cos\theta}.$$

Plugging this into (3) we get

$$\|x - \hat{x}\| \leq \frac{2\cos^2\theta}{1 - 3\cos\theta} \cdot f(\hat{x}) \leq 3\varepsilon^2$$

for $\varepsilon \leq 0.05$ and $\theta = \pi/2 - 1.1\varepsilon$.

□

Cone envelopes. Let $C^+ = \bigcup_{p \in P} C_p^+$ and let Σ^+ be the set of points $x \in C^+$ for which the segment $x\hat{x}$ has no point of C^+ in its relative interior. We call Σ^+ the *outer cone envelope* of Σ. The *inner cone envelope* Σ^- is defined analogously.

Lemma 11 *The surface Σ is homeomorphic to both Σ^- and Σ^+. Both cone envelopes divide $\mathbb{R}^3$ into a bounded and an unbounded component. The bounded component of the inner cone envelope and the unbounded component of the outer cone envelope are closed under the flow ϕ.*

Proof. We only present the proof for the outer cone envelope. The proof for the inner cone envelope follows the same lines.

Let $\pi : \mathbb{R}^3 \setminus M \to \Sigma$ project into Σ, i.e. π maps every point to its closest point in Σ. It is well known that this map is continuous. Let π^+ be the restriction of π to Σ^+. Since $\Sigma^+ \cap M = \emptyset$ (because $\Sigma^+ \subset \Sigma_{3\varepsilon^2}$ by Lemma 10), π^+ is defined and continuous on Σ^+. By definition of Σ^+, π^+ is injective and by Lemma 9, π^+ is surjective. This along with the compactness of Σ^+ (inferred from the compactness of Σ and the continuous one to one mapping given by π^+) implies that π^+ is a homeomorphism.

Since we assumed that Σ is a manifold without boundary, so is Σ^+. Thus Σ^+ divides $\mathbb{R}^3$ into a bounded and an unbounded component. By Lemma 8 the bounded component has to be closed under the flow ϕ. □

Cone neighborhood. We call the closed volume sandwiched between Σ^- and Σ^+ the *cone neighborhood* of Σ and denote it by $\tilde{\Sigma}$.

Theorem 2 *The output of the algorithm* RECONSTRUCT *lies in $\tilde{\Sigma}$ and the latter itself is contained in $\Sigma_{3\varepsilon^2}$.*

Proof. By Lemmas 8 and 11, the stable manifold $S(c)$ of any surface critical point c has to be contained in $\tilde{\Sigma}$. Thus the output of RECONSTRUCT completely

lies in $\tilde{\Sigma}$. By Lemma 10, Σ^+ and Σ^- are contained in $\Sigma_{3\varepsilon^2}$. This implies that $\tilde{\Sigma}$ is also contained in $\Sigma_{3\varepsilon^2}$. □

The following Corollary will be needed later.

Corollary 5 *For every $x \in \Sigma_\xi$, i.e., for every x satisfying $\|x - \hat{x}\| \leq \xi f(\hat{x})$, and for every $p \in A(x)$, $\|x - p\| \leq \frac{\varepsilon+\xi}{1-\varepsilon-2\xi} f(p)$. In particular, when $\varepsilon \leq 0.05$, for every point $x \in \tilde{\Sigma}$, and every $p \in A(x)$, $\|x - p\| \leq 1.23\varepsilon f(p)$, and for every surface critical point c, and every $p \in A(c)$, $\|c - p\| \leq 1.12\varepsilon f(p)$.*

Proof. Let x be a point in Σ_ξ. From the definition of ε-sampling, $\|\hat{x} - q\| \leq \varepsilon f(\hat{x})$, where $q \in A(\hat{x})$, i.e. q is a closest sample point to $\hat{x}$. For any sample point $p \in A(x)$, by the triangle inequality

$$\|x - p\| \leq \|x - q\| \leq \|x - \hat{x}\| + \|\hat{x} - q\| \leq (\varepsilon + \xi) f(\hat{x}). \tag{4}$$

Thus we get

$$\|p - \hat{x}\| \leq \|x - p\| + \|x - \hat{x}\| \leq (\varepsilon + \xi) f(\hat{x}) + \xi f(\hat{x}) \leq (\varepsilon + 2\xi) f(\hat{x}),$$

and from this and because the local feature size is 1-Lipschitz,

$$f(p) \geq f(\hat{x}) - \|\hat{x} - p\| \geq f(\hat{x}) - (\varepsilon + 2\xi) f(\hat{x}). \tag{5}$$

Combining (4) and (5) we get

$$\|x - p\| \leq \frac{\varepsilon + \xi}{1 - \varepsilon - 2\xi} f(p).$$

Using $\varepsilon \leq 0.05$ along with Theorem 2 implies the bounds for the case where $x \in \tilde{\Sigma}$. In the case of x being a surface critical point we invoke Corollary 2, instead. □

5.2 Convergence of Normals

The output $\mathcal{T}$ produced by the algorithm RECONSTRUCT consists of stable manifolds of index-2 saddle points that lie in a small tubular neighborhood of the surface. We refer to these stable manifolds as *surface patches*. We want to show that under (ε, δ)-sampling, with a fixed $\rho = \delta/\varepsilon$, the normal of triangles in these surface patches is within $O(\varepsilon)$ from the normal to surface at a nearby point, for sufficiently small ε. We use the following two lemmas from [3].

Lemma 12 *For any two points $p, q \in \Sigma$, the angle between segment pq and either of n_p^+ and n_p^- is greater than $\frac{\pi}{2} - \arcsin \frac{\|p-q\|}{2f(p)}$.*

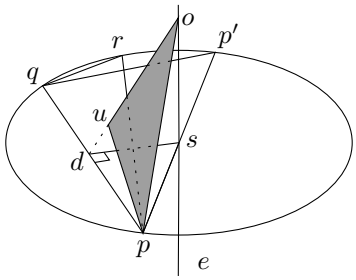

Figure 6: A generic patch triangle on the stable manifold of a surface 2-saddle.

Lemma 13 *For points $p, q, r \in \Sigma$, let p be a vertex of the triangle pqr with the largest angle and let r be its circumradius. If $r = \lambda f(p)$, then the acute angle between the normal to pqr and the normal to surface at p is at most $\beta(\lambda)$ where*

$$\beta(\lambda) = \arcsin(\lambda) + \arcsin\left(\frac{2}{\sqrt{3}}\sin(2\arcsin\lambda)\right) \leq 4\lambda,$$

for $\lambda \leq \frac{1}{4}$.

The stable manifold $S(c)$ of every 2-saddle is a piece-wise linear surface made of a finite number of triangles, which we call *patch triangles*. Each patch triangle t has exactly one vertex in P. Note that for every point x in a patch triangle t, the vertex of t that belongs to P is a closest sample point to x (refer to [12] for details on the structure of stable manifolds of critical points). If x is on the boundary of t, it can have more than one closest sample point as it belongs to more than one patch triangle. The following lemma shows that under tight sampling, each patch triangle must have a normal close to surface normal at its vertex in P.

Lemma 14 *For any $0 < \rho < 1$, there exists ε_0 such that if P is an (ε, δ)-sample of Σ with $\varepsilon \leq \varepsilon_0$ and $\delta = \rho\varepsilon$, then for any $x \in S(c)$, the stable manifold of a surface 2-saddle c, the acute angle between n_p, where p is a closest sample point to x, and n_t, the normal direction of the patch triangle $t \subset S(c)$ that contains x and has p as a vertex, is at most*

$$\arcsin\left(\frac{\sin(1.23\varepsilon)}{2\sin\left(\frac{1}{2}\arcsin\left(\rho/2.46\right)\right)}\right) = O(\varepsilon/\rho).$$

Proof. Let P be an ε-sample of Σ for $\varepsilon \leq \varepsilon_0$. By Corollary 5, $\|x-p\| \leq 1.23\varepsilon f(p)$ for every closest sample point p to x. Every point x on $S(c)$ is on a patch triangle $t = pou$ of $S(c)$ with the following structure (see Figure 6): t has exactly one vertex p in P. The edge uo of t opposite to p is on the Voronoi facet dual to a

Delaunay edge pq and ends on the dual Voronoi edge e of a Delaunay triangle pqr in which $\|r-p\| > \|q-p\|$. The mid-point d of pq is the driver of the points on uo. Furthermore, the line containing e does not intersect the triangle pqr except when o is the critical point c in which case the patch triangle t is coplanar with (an in fact contained in) the Delaunay triangle t_c containing c. We postpone the study of this special case for later. Let s be the circumcenter of pqr and let p' be a point on the circumcircle of pqr opposite to p with respect to s. Then $\angle p'qp = \pi/2$ and that $\|d-s\| = \frac{1}{2}\|q-p'\|$. Furthermore, $\|q-p'\| > \|q-r\|$. Therefore we get $\|d-s\| \geq \frac{1}{2}\delta f(r)$. On the other hand, since $r \in A(o)$, by Corollary 5, $\|s-p\| = \|s-r\| < \|r-o\| \leq 1.23\varepsilon f(r)$. Combining these we get for the angle $\alpha = \angle qps$:

$$\sin\alpha \geq \frac{\|d-s\|}{\|p-s\|} \geq \frac{\delta}{2.46\varepsilon} = \frac{\rho}{2.46}.$$

On the other hand, $\pi/2 > \angle qpo > \angle qps \geq \alpha$. Also, $o \in S(c)$ and $S(c)$ is contained in $\tilde{\Sigma}$ and therefore, po makes an angle of at least $\theta = \frac{\pi}{2} - 1.1\varepsilon$ with n_p. Moreover, $\|p-q\| \leq 2.46\varepsilon f(p)$ and therefore by Lemma 12, $q-p$ makes an angle of at least $\frac{\pi}{2} - 1.23\varepsilon < \theta$ with n_p. Thus, the three points p, q, and o, make a triangle $t' = pqo$ with an angle of at least α at vertex p satisfying $\arcsin(\rho/2.46) < \alpha < \frac{\pi}{2}$, and with both of the edges incident to p making an angle of at least $\pi/2 - 1.23\varepsilon$ with n_p. It can be shown through elementary calculations that under these conditions, $n_{t'}$, the normal to t', and n_p, make an angle of at most

$$\arcsin\left(\frac{\cos(\pi/2 - 1.23\varepsilon)}{2\sin(\alpha/2)}\right),$$

which matches the bound in the statement of the lemma. We now consider the special case when o coincides with c. This happens when the Voronoi edge e intersects its dual Delaunay triangle pqr at c and the patch triangle t in question becomes coplanar with the Delaunay triangle $t_c = pqr$. Notice that $A(c) = \{p, q, r\}$ and by Corollary 5, $\|p-c\| \leq 1.12\varepsilon f(p)$. Similar inequalities hold for q and r. As in the previous case, let p' be the point on the circumcircle of pqr opposite to p with respect to circumcenter c. We denote the angles $\angle qpp'$ and $\angle rpp'$ by β and γ respectively, and their sum by α. Since the angles $\angle pqp'$ and $\angle prp'$ are each 90 degrees, we have $\sin\beta = \frac{1}{2}\|q-p'\|/\|p-c\|$ and $\sin\gamma = \frac{1}{2}\|r-p'\|/\|p-c\|$. In order for c to be a critical point, all the angles of the triangle t_c must be acute. Since the sine function is concave for acute angles we

have

$$\begin{aligned}
\sin\frac{\alpha}{2} &\geq \frac{\sin\beta+\sin\gamma}{2} \\
&= \frac{\|r-p'\|+\|q-p'\|}{4\cdot\|p-c\|} \\
&\geq \frac{\|r-q\|}{4\cdot\|p-c\|} \\
&\geq \frac{\delta f(r)}{4\cdot 1.12\varepsilon f(r)} \\
&= \frac{\rho}{4.48}.
\end{aligned}$$

On the other hand $\|p-q\| \leq \|p-c\| + \|q-c\| \leq 2.24\varepsilon f(r)$ by Lemma 12. Thus pq makes an angle of at most $\pi/2 - 1.12\varepsilon$ with n_p. A similar argument establishes the same bound for the angle between pr and n_p. Similar to the previous case, We have shown for the triangle pqr that the angle α at p is at least $2\arcsin(\rho/4.48)$ and that the edges pq and pr make an angle of at least $\pi/2 - 1.12\varepsilon$ with n_p. This implies that the angle between the normal to the plane of this triangle and the normal to Σ at p is at most

$$\arcsin\left(\frac{\cos(\pi/2-1.12\varepsilon)}{2\rho/4.48}\right).$$

It can be verified that the this bound results a smaller angle than the one obtained above (matching the bound in the statement of the lemma) for the general case, whenever the two bounds are defined for any $0 < \varepsilon \leq 0.05$ and $0 < \rho < 1$. □

Corollary 6 *For $\varepsilon \leq 0.05$ and $\rho \geq 1/2$, we have for every point x on the stable manifold of a surface 2-saddle c that the acute angle between normal n_t to any patch triangle t of $S(c)$ that contains x, and the normal $n_{\hat{x}}$ to Σ at $\hat{x}$ is at most* 23 *degrees.*

Proof. Plugging $\varepsilon \leq 0.05$ and $\rho \geq 1/2$ in Lemma 14, gives an upper bound of 18 degrees for the angle between n_t and n_p, where n_p is the normal to Σ at a closest sample point p to x. Since by Theorem 2, $S(c)$ is contained in $\Sigma_{3\varepsilon^2}$, we have $\|x-\hat{x}\| \leq 3\varepsilon^2 f(\hat{x})$. Let q be a closest sample point to $\hat{x}$. By sampling condition, $\|\hat{x}-q\| \leq \varepsilon f(\hat{x})$ and therefore

$$\|\hat{x}-p\| \leq \|\hat{x}-x\| + \|x-p\|.$$

On the other hand $\|x-p\| \leq \|x-\hat{x}\| + \|\hat{x}-q\|$. Therefore we get

$$\|\hat{x}-p\| \leq 2\|x-\hat{x}\| + \|\hat{x}-q\| \leq (6\varepsilon^2+\varepsilon)f(\hat{x}) \leq 1.3\varepsilon f(\hat{x}),$$

for $\varepsilon \leq 0.05$. Therefore by Lemma 1 the angle between n_p^+ and $n_{\hat{x}}^+$ is at most $1.3\varepsilon/(1-3\cdot 1.3\varepsilon) \leq 5°$. □

The following proposition is directly based on the structure of the stable manifolds of 2-saddles [12] and the fact that flow lines never bend by more than 90 degrees.

Proposition 1 *If t_1 and t_2 are patch triangles of $S(c)$ for a surface 2-saddle c such that t_1 and t_2 have one edge in common, then the dihedral angle between t_1 and t_2 is no less than $\pi/2$.*

Lemma 15 *Let c be a surface 2-saddle. Suppose we orient the patch triangles in $S(c)$ arbitrarily but consistently so that for any patch triangle t, n_t^+ and n_t^- are respectively the outer and inner normal directions on t with respect to the applied orientation. Then, under the assumptions of Corollary 6 exactly one of the following cases holds.*

1. *$\angle(n_t^+, n_{\hat{x}}^+) \leq 23°$, for every patch triangle t of $S(c)$ and for every $x \in t$.*
2. *$\angle(n_t^+, n_{\hat{x}}^-) \leq 23°$, for every patch triangle t of $S(c)$ and for every $x \in t$.*

Proof. First notice that as was shown in the proof of Corollary 6, for any point $x \in t$, where t is a patch triangle of $S(c)$, $\angle(n_{\hat{x}}^+, n_p^+) \leq 5°$, where p is the vertex of t that is a sample point. Thus, if for the arbitrary orientation of t and for a point $x \in t$, $\angle(n_t^+, n_{\hat{x}}^+) = \alpha$, the same holds for every other point y in t, modulo changing α by 5 degree.

Let t_c be the Delaunay triangle that contains c. All patch triangles $t \subset t_c$ of $S(c)$, have the same n_t^+ which agrees with one of the two orientations of the direction normal to t_c. By Corollary 6, the normal direction of t_c makes an angle of at most 23° with either $n_{\hat{c}}^+$ or $n_{\hat{c}}^-$. Assume without loss of generality that the first case holds, i.e. $\angle(n_t^+, n_{\hat{c}}^+) \leq 23°$. We show now that this will imply that that for every patch triangle t of $S(c)$ and every $x \in t$, $\angle(n_t^+, n_{\hat{x}}^+) \leq 23°$. We prove this by extending the result for the triangles we already have this property for to their neighboring patch triangles. Thus, assume t and t' are two patch triangles with an edge e in common. Let z be a point on e. Since t and t' are oriented consistently, the dihedral angle between t and t' is $\pi - \angle(n_t^+, n_{t'}^+)$. By Proposition 1 this angle is at least $\pi/2$ and therefore $\angle(n_t^+, n_{t'}^+) \leq \pi/2$. Therefore using triangle inequality for angles we get $\angle(n_{\hat{z}}^+, n_{t'}^+) \leq \angle(n_t^+, n_{\hat{z}}^+) + \angle(n_t^+, n_{t'}^+) \leq 90 + 23 = 113°$. But by Corollary 6, $\angle(n_{\hat{z}}^+, n_{t'}^+)$ is either less than 23° or more than $180° - 23°$ and we have just shown that the latter case does not hold. □

5.3 Orientation of surface patches

The output the algorithm RECONSTRUCT is a collection of stable manifolds of surface 2-saddles (patches) that are attached to each other at Gabriel edges (stable manifolds of surface 1-saddles). In order to establish that this reconstruction

has the same topology as the original surface Σ, we shall provide a homeomorphism between the two surfaces. We have shown in the previous section that, roughly speaking, each patch is almost flat and lies almost orthogonal to the normal to Σ at the 2-saddle into which the patch points flow. In order to achieve the desired homeomorphism we need to show that neighboring patches do not fold over each other. Our way of showing this can be summarized as follows: we observe that the normal to a patch induced by a 2-saddle c at c itself is close to the normal to Σ at a near surface point to c. This gives a natural orientation of the patch. We prove afterward that the side of the patch that faces the union of stable manifolds $\mathcal{O}$ of the critical points in the component computed by the algorithm RECONSTRUCT is consistently determined by the given orientation.

We will need the following two technical lemmas.

Lemma 16 *Let C_1 and C_2 be two infinite cones with cone angle θ, with the same apex p and same axis, extended in opposite directions. Let x be a point not in the interior of either of the convex hulls of C_1 or C_2. Consider a line ℓ passing through x, making an angle of $\alpha < \theta$ with the common axis of C_1 and C_2, and hitting C_1 and C_2 in points x_1 and x_2, respectively. Then*

$$\|x_1 - x_2\| \leq 2 \cdot \|x - p\| \cdot \frac{\cos\theta}{\sin(\theta - \alpha)}.$$

Proof. Without loss of generality assume that p is the origin and that the common axis of C_1 and C_2 is the z-axis. By the assumptions of the lemma, x_1 and x_2 are in opposite sides of x on ℓ. Consider the vertical plane H containing x and the z axis. When x is fixed, if we consider an arbitrary line ℓ through x making an angle of α with the z-axis, it is easy to observe that $\|x - x_1\|$ is maximized when ℓ is contained in H, in which case by the law of sines $\|x - x_1\| = \|x - y_1\| \cdot \sin\theta / \sin(\theta - \alpha)$, where y_1 is the vertical projection of x to C_1. Thus in general

$$\|x - x_1\| \leq \|x - y_1\| \cdot \sin\theta / \sin(\theta - \alpha). \tag{6}$$

Similarly we get for the distance between x and x_2

$$\|x - x_2\| \leq \|x - y_2\| \cdot \sin\theta / \sin(\theta - \alpha), \tag{7}$$

where y_2 is the vertical projection of x to C_2. On the hand, when $\|x - p\|$ is fixed, $\|y_1 - y_2\|$ is maximized when x is in the plane $z = 0$, in which case $\|y_1 - y_2\| = 2 \cdot \|x - p\| \cdot \cot\theta$. So, in general

$$\|y_1 - y_2\| \leq 2 \cdot \|x - p\| \cdot \cot\theta. \tag{8}$$

Combining (6), (7), and (8) we get

$$\begin{aligned}
\|x_1 - x_2\| &= \|x - x_1\| + \|x - x_2\| \\
&\leq (\|x - y_1\| + \|x - y_2\|) \cdot \sin\theta / \sin(\theta - \alpha) \\
&= \|y_1 - y_2\| \cdot \sin\theta / \sin(\theta - \alpha) \\
&\leq 2 \cdot \|x - p\| \cdot \cos\theta / \sin(\theta - \alpha).
\end{aligned}$$

□

Lemma 17 *Assume $\varepsilon \leq 0.01$. Let x be a point in $\tilde{\Sigma}$ with $\{p, q, r\} \subset A(x)$. Then the acute angle between normal to the Delaunay triangle pqr and each of the normals n_p, n_q, and n_r is at most 8ε.*

Proof. Assume without loss of generality that p is the vertex of pqr with the largest face angle. Since $x \in \tilde{\Sigma} \subset \Sigma_{3\varepsilon^2}$, by Corollary 5, $\|x-p\| = \|x-q\| = \|x-r\| \leq 1.23\varepsilon f(p)$. On the other hand, $\|x-p\|$ is an upper bound for the circumradius of pqr. Thus using Lemma 13, the acute angle between n_p and normal to pqr is at most $\beta(1.23\varepsilon) \leq 5\varepsilon$. On the other hand, $\|p-q\| \leq \|p-x\| + \|x-q\| \leq 2 \cdot 1.23 f(p)$. Therefore by Lemma 1, $\angle(n_p^+, n_q^+) \leq \frac{2.46\varepsilon}{1-3\cdot 2.46\varepsilon} \leq 3\varepsilon$ for $\varepsilon \leq 0.01$. The same argument can be repeated with r instead q. □

We say that a triangle $t = pqr$ with $p, q, r \in P$ *lies flat to surface* or simply *is flat* if the normal of t is within 8ε from one of n_p, n_q, or n_r. For such a triangle, it is meaningful to distinguish between the side that faces the interior of Σ and the one that faces its exterior. We refer to these sides as *inner* and *outer* sides, respectively.

Let c be a surface 2-saddle. By definition, c is the intersection point of a Delaunay triangle t_c and its dual Voronoi edge e_c. Thus $|A(c)| = 3$ and by Lemma 17, the normal to t_c makes an angle of 8ε or less with the surface normal at any of the vertices of t_c. Thus, t_c lies flat to surface. Since t_c intersects e_c in a point of its relative interior (by our non-degeneracy assumption), we can distinguish between the two endpoints of e_c as its inner and outer vertices and refer to them as v_c^- and v_c^+, respectively. We denote the the segment cv_c^+ excluding c by e_c^+, and define e_c^- similarly. Notice that c is the driver for points on e_c and therefore the flow direction on $e_c \setminus \{c\}$ is toward its endpoints at each side of c. Therefore, every point of e_c between c and v_c^+ flows to the same maximum that v_c^+ flows into. A similar statement holds for the points between c and v_c^-. We define $U_c^+ = e_c^+ \cup \phi(v_c^+)$ and $U_c^- = e_c^- \cup \phi(v_c^-)$. In fact, U_c^+ and U_c^- together constitute the *unstable manifold* of c [12]. Thus, if U_c^+ intersects Σ^+ then the flow originated at any point of e_c^+, arbitrarily close to c must end up in an exterior medial axis maximum m implying that $S(c)$ is incident to $S(m)$ through the outer side of t_c. Similar statements can be made by replacing U_c^+ with U_c^- and Σ^+ with Σ^-.

Lemma 18 *For any $0 < \rho < 1$, there exists ε_0 small enough such that if P is an (ε, δ)-sampling of Σ for $\varepsilon \leq \varepsilon_0$ and $\delta = \rho\varepsilon$, then for any $x \in U_c^+ \cap \tilde{\Sigma}$, $\angle(v(x), n_p^+) \leq 8\varepsilon$, and for every point $x \in U_c^- \cap \tilde{\Sigma}$, $\angle(v(x), n_p^-) \leq 8\varepsilon$, where p is any point in $A(x)$. In particular, for $\rho = \delta/\varepsilon \geq 1/3$, $\varepsilon_0 \leq 0.01$ suffices.*

Proof. We only prove the lemma for points in $U_c^+ \cap \tilde{\Sigma}$. The proof for points in $U_c^- \cap \tilde{\Sigma}$ is analogous. For simplicity we enforce $\varepsilon_0 \leq 0.01$ although the statement of the lemma may hold for larger values of ε_0. Let P be an (ε, δ)-sampling of Σ with $\varepsilon \leq \varepsilon_0 \leq 0.01$ and $\delta = \rho\varepsilon$.

Since $x \in \tilde{\Sigma} \subset \Sigma_{3\varepsilon^2}$, by Corollary 5, $\|x-p\| \leq 1.23\varepsilon f(p)$ for every $p \in A(x)$.

Notice that it suffices to prove that $\angle(v(x), n_p^+) \leq 5\varepsilon$ for only one point $p \in A(x)$. This is because for any other point $q \in A(x)$, $\|p-q\| \leq \|p-x\| + \|q-x\| \leq 2 \cdot 1.23\varepsilon \max\{f(p), f(q)\}$ and therefore by Lemma 1, $\angle(n_p^+, n_q^+) \leq \frac{2.46\varepsilon}{1-3\cdot 2.46\varepsilon} \leq 3\varepsilon$ for $\varepsilon \leq 0.01$.

As above, let t_c and e_c be the Delaunay triangle and its dual Voronoi edge for which $\{c\} = t_c \cap e_c$. Notice that U_c^+ is a piece-wise linear curve. Let $u_0, u_1, \ldots, u_k$ be the vertices of this curve with $u_0 = c$, $u_1 = v_c^+$, and $u_k = m$, where m, the maximum at which U_c^+ ends. Notice of course that u_0 itself does not belong to U_c^+ as $v(u_0) = v(c) = 0$. We prove the lemma inductively starting from the segment u_0u_1 and going up to $u_{i-1}u_i$ for the smallest i for which $u_{i-1} \in \tilde{\Sigma}$ but $u_{i-1}u_i$ has a point outside $\tilde{\Sigma}$. For this last line segment $u_{i-1}u_i$, the argument we provide will be true only for the initial part $u_{i-1}u^*$, where u^* is the first point of U_c^+ (starting from u_{i-1}) not in $\tilde{\Sigma}$. Since no flow enters $\tilde{\Sigma}$, no point of U_c^+ past u^* will be in $\tilde{\Sigma}$. From the structure of flow complex [12], it is easy to see that every vertex in $\{u_1, \ldots, u_k\}$ is either a Voronoi vertex or lies on a Voronoi edge. Furthermore, the relative interior of every segments $u_{i-1}u_i$, $i = 1, \ldots, k$, falls entirely inside a Voronoi edge or facet.

For the base case of our induction we observe that the lemma holds for points $x \in u_0u_1$ (excluding u_0). To see this, notice that the direction of $v(x)$ for such points agrees with the vector $v_c^+ - c$. Let p be the vertex of t_c with the largest angle in t_c. Using Lemma 17, and taking into account that x is on the outer side of t, implies that the angle between $n^+(p)$ and $v_c^+ - c$ is at most 8ε.

In fact, by Lemma 17, for any point $x \in U_c^+ \cap \tilde{\Sigma}$ that flows on a Voronoi edge e, the Delaunay triangle t dual to e must lie flat to surface and thus we can distinguish between its side facing outward and the one facing inward. Informally, we will say that in such a case x is *above* t if x is on the side of t facing outward, or *below* t otherwise.

For the induction step, we assume that the statement of the lemma holds for points on a segment $u_{i-1}u_i$ of $U_c^+ \cap \tilde{\Sigma}$ and show that this entails the same for the point on the segment u_iu_{i+1}. Let f_1 be the Voronoi face of dimension d_1 that contains the relative interior of $u_{i-1}u_i$, and let f_2 be the Voronoi face of dimension d_2 containing u_i. Finally let f_3 be the Voronoi face that contains the relative interior of u_iu_{i+1} and let d_3 be its dimension. Notice that f_1 and f_3 are cofaces of f_2 and therefore d_1 and d_3 are both greater than d_2. We prove the induction step by going over all possible combinations of f_1, f_2, and f_3.

1. Edge-vertex-edge. First we study the case in which the flow on a Voronoi edge e, reaches a Voronoi vertex v and enters another Voronoi edge e'. We assume that the statement of the lemma holds for points on e and show that it remains true as the flow moves on to e'. To see this, let $t = pqr$ be the Delaunay triangle dual to e and let $t' = qrs$ be the one dual to e' (See Figure 7 (left)). The Voronoi vertex v must be dual to the Delaunay tetrahedron Δ with vertex set $\{p, q, r, s\}$. Since the flow through v continues on e', the driver of points in e' must lie in the interior of the triangle t' or in other words, the line through e' must intersect

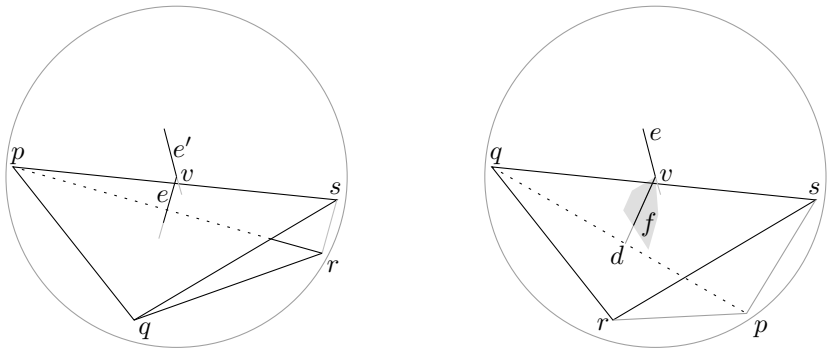

Figure 7: Proof of Lemma 18, cases 1 (left) and 2 (right).

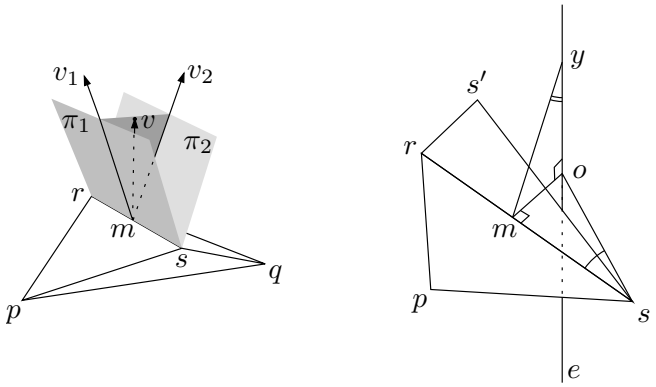

Figure 8: Proof of Lemma 18, cases 3 (left) and 5 (right).

t'. As discussed above both t and t' are flat and v is above t. It is easy to see that extending the statement of the Lemma to e', is identical to showing that v is also above e'. By Lemma 17, the outward normal direction for both t and t' are within 8ε from n_q^+. If v is not above t', it must be in Δ. But in that case v is a maximum and the flow does not leave it to enter e'.

2. Facet-vertex-edge. Next, we consider the case where the flow through a Voronoi facet f dual to Delaunay edge pq reaches a Voronoi vertex v dual to tetrahedron $\Delta = pqrs$ and continues on a Voronoi edge e dual to Delaunay triangle $t = qrs$ (See Figure 7 (right)). We assume that the lemma holds for the points x on $U_c^+ \cap f$ and show that this extends to the points on e. As in the previous case t is flat and we only need to show that v is above t. By induction hypothesis, $\angle(v - d, n_q^+) \leq 8\varepsilon$ where $d = \frac{1}{2}(p + q)$ is the driver of the points on f. It can be verified that if v is not above t, it must be that $v \in \Delta$ making v a local maximum, a contradiction.

3. Edge-vertex-facet. Consider now the case where the flow through a Voronoi edge e reaches a Voronoi vertex v and enters a Voronoi facet f incident to v. Let pqr be the Delaunay triangle dual to e and let rs be the Delaunay edge dual to f. Note that v is the circumcenter of the Delaunay tetrahedron $\Delta = pqrs$. We assume that the lemma holds for points on e which is identical to assuming that v is above t. Since the flow through v continues on f, the closest point of Δ to v is the midpoint $m = \frac{1}{2}(r+s)$ of the edge rs (see Figure 8, left). For this to happen, v must be in the wedge made by two half-planes π_1 and π_2 both having the line through rs as boundary and respectively being orthogonal to triangles $t_1 = prs$ and $t_2 = qrs$. Since $A(v) = \{p, q, r, s\}$, by Lemma 17, the normals to both t_1 and t_2 make an angle of at most 8ε with n_r. Since v is above t but not contained in Δ, it must be above both t_1 and t_2. If we base at m, two vectors v_1 and v_2, respectively normal to t_1 and t_2 in their outward directions, v_1 will lie in π_1 and v_2 in π_2. The segment vm is on the plane bisecting rs and so are v_1 and v_2. It follows from the triangle inequality for angles that $v - m = v(x)$, for $x \in f \cap U_c^+$, also makes and angle of at most 8ε with n_r^+. Notice that with exactly the same argument but with using s instead of r, we get the same bound with respect to n_s^+.

4. Facet-vertex-facet. The proof of this case is a simple combination of the proofs of cases 2 and 3.

5. Facet-edge-facet. We show now that under tight enough sampling, i.e. by choosing ρ large enough, if the flow through a Voronoi facet f arrives at a Voronoi edge e of f, it will continue on e and does not enter another facet f' incident to e, given that the statement of lemma holds for the points of $U_c^+ \cap f$. Suppose to the contrary that this is not the case, i.e. (see Figure 8, right) the flow crosses e and enters another facet f' incident to e. Let rs be the Delaunay edge dual to f. The driver of the flow on f is $m = \frac{1}{2}(r+s)$. Let y be the point where the flow reaches e. The dual Delaunay triangle t to e has r and s for vertices plus another vertex s'. For the flow to cross e and enter f', f' must be dual to the Delaunay edge ss'. Furthermore, the line of e must not intersect t. Let o be the circumcenter of t. By our assumption, the flow direction on f, which coincides with $y - m$ makes an angle of no more than 8ε with n_r^+. On the other hand, $y \in \tilde{\Sigma}$ and thus t lies flat to surface and since the largest angle in t is at r, the normal to t, i.e. direction of e, makes an acute angle of at most 5ε with n_r (See proof of Lemma 17). This in particular implies that y is above t. Therefore, $\angle myo$ is at most 13ε. In order for the line of e not to intersect t, it must hold that $\angle mss' < \angle mso$. The two triangles mso and myo both share the edge mo and both have a right angle on one of the end-points of this edge. We will show below that $\|m - y\| < \|m - s\|$. Since $\|m - s\| < \|o - y\|$, this will imply that $\angle mso < \angle myo$. Using exactly the same argument as in the proof of Lemma 14, we get $\angle mso \geq \arcsin(\rho/2.46)$, and therefore since $\angle myo \leq 13\varepsilon$, we must have

$$13\varepsilon > \arcsin\left(\frac{\rho}{2.46}\right).$$

This inequality is violated for $\varepsilon \leq \varepsilon_0$ for $\varepsilon_0 \leq \frac{1}{13}\arcsin(\rho/2.46)$ (in particular

for $\varepsilon_0 \leq 0.01$ when $\rho \geq 1/3$) giving us the desired contradiction.

Now we prove that $\|m-y\| \leq \|m-s\|$. Notice that s is a closest sample point to y and by our assumption $y \in \Sigma$. Therefore, y is between the cones C_s^+ and C_s^-. On the other hand by Corollary 5, $\|r-s\| = 2\|m-s\| \leq 2\|y-s\| \leq 2 \cdot 1.23\varepsilon f(s)$ and therefore using Lemma 12, ms makes an angle of at least $\frac{\pi}{2} - 1.23\varepsilon > \theta$ with the normal to Σ at s. This implies that m is also between C_s^+ and C_s^-. By our inductive hypothesis, my makes an angle of at most 8ε with n_s^+. Lemma 16 can now be used to get

$$\|m-y\| \leq 2 \cdot 1.23\varepsilon \cdot f(s) \frac{\sin(1.1\varepsilon)}{\cos(9.1\varepsilon)} < \frac{1}{2}\delta f(s) \leq \|m-s\|,$$

where the middle inequality holds for $\varepsilon \leq 0.01$ and $\rho \geq 1/3$. In fact for any constant $0 < \rho < 1$ the above inequality holds (and the desired contradiction is achieved) for any $\varepsilon \leq \varepsilon_0$ for small enough ε_0 since the left hand side has a quadratic dependence on ε.

Thus we have proved that whenever the flow on U_c^+ moves to a Voronoi facet f, it leaves f by either hitting a Voronoi edge e and continuing on e, or by hitting a vertex v. Thus we have covered all cases in the inductive step and this completes the proof of the lemma. □

In the following Lemma we show that if $S(c)$ is incident to the stable manifold $S(m)$ of an interior (exterior) medial axis maximum m, then the part of $S(c)$ that is contained in t_c is incident to $S(m)$ at the inner (outer) side of t_c.

Lemma 19 *For any surface 2-saddle c, U_c^+ does not intersect Σ^- and U_c^- does not intersect Σ^+.*

Proof. We prove the claim for U_c^+. The other claim is proved analogously. Suppose to the contrary that U_c^+ intersects Σ^- at x. Let v be the last turning point of U_c^+ before reaching x. Let q be a sample point for which $x \in C_q^-$ and let p be a closest sample point to x. Then $\|x-p\| \leq \|x-q\| \leq f(q)\cos\theta = f(q)\sin(1.1\varepsilon)$. Therefore, $\|p-q\| \leq 2f(q)\sin(1.1\varepsilon)$ and therefore by Lemma 1,

$$\angle(n_p^+, n_q^+) \leq \frac{2\sin(1.1\varepsilon)}{1 - 3 \cdot 2\sin(1.1\varepsilon)} = O(\varepsilon).$$

On the other hand, by Lemma 18, the vector $x-v$ makes an angle of $O(\varepsilon)$ with n_p^+. It is easy to observe that this contradicts the assumption that the flow hits C_q^-. □

The following lemma is a direct consequence of Lemma 18 and Lemma 15.

Lemma 20 *Let c_1 and c_2 be two surface critical points with $S(c_1)$ and $S(c_2)$ put by* Reconstruct *into $\mathcal{T}$, such that boundaries of $S(c_1)$ and $S(c_2)$ have a Gabriel edge e in common. Let t_1 and t_2 be the patch triangles incident to e in $S(c_1)$ and $S(c_2)$, respectively. Then, the dihedral angle between t_1 and t_2 is larger than $\pi/2$.*

Proof. Orient patch triangles of $S(c_1)$ by taking for every patch triangle t of $S(c_1)$, the normal to t pointing to the side of t incident to the interior of the reconstruction. Denote this normal by n_t^-. Lemma 18 implies that in this case for every point $x \in t$ where t is a patch triangle of $S(c_1)$, $\angle(n_{\hat{x}}^+, n_t^+) < 23°$. In particular, by letting $t = t_1$ and choosing x to be a point on e, we get $\angle(n_{\hat{x}}^+, n_{t_1}^+) < 23°$

If we do a similar orientation on $S(c_2)$, we get $\angle(n_{\hat{x}}^+, n_{t_2}^+) < 23°$. Thus, the dihedral angle between t_1 and t_2 is at least $180° - 46° = 134°$. □

5.4 Homeomorphism

Theorem 3 *Under assumptions of Theorem 1, the output $\mathcal{T}$ produced by the algorithm* RECONSTRUCT *is a 2-manifold without boundary homeomorphic to Σ.*

Proof. First we observe that the complex $\mathcal{T}$ produced by RECONSTRUCT is the boundary of the union of stable manifolds of either the inner or outer medial axis critical points. Let m_1 and m_2 be medial axis maxima such that $S(m_1)$ and $S(m_2)$ are neighboring 3-cells in the flow complex, i.e. they both have $S(c)$ contained in their boundaries, where c is a 2-saddle. If m_1 is an inner medial axis maximum and m_2 an outer one, then c must be a surface critical point as the common boundary of $S(m_1)$ and $S(m_2)$ must lie in $\tilde{\Sigma} \subset \Sigma_{3\varepsilon^2}$. On the other hand, if m_1 and m_2 are both inner (outer) medial axis maxima then $S(c)$ cannot be a surface critical point since otherwise both U_c^+ and U_c^- arrive at inner (outer) medial axis maxima and therefore both must have crossed Σ^- (Σ^+) and this violates Lemma 19. This in particular implies that the algorithm RECONSTRUCT in fact partitions the medial axis critical points into two subsets.

We consider in this proof the case where $\mathcal{T}$ is the boundary of the union of stable manifolds of the inner medial axis critical points (the outer case being analogous). We argue that $\mathcal{T}$ and Σ are homeomorphic. Consider the restriction $\zeta\colon \mathcal{T} \to \Sigma$ of the closest point map $x \mapsto \hat{x}$. We prove that ζ is a homeomorphism. Since both $\mathcal{T}$ and Σ are compact, it is sufficient to show that ζ is continuous, one-to-one and onto.

First, we argue that ζ is one-to-one. Orient the normal to each patch triangle t so that it makes an angle less than $\frac{\pi}{2}$ with the oriented normal n_p^+ at the vertex p of t which is a sample point. Because of Lemma 15 and Lemma 20, the triangles of $\mathcal{T}$ can be oriented consistently satisfying this condition. We denote this oriented normal for a patch triangle t by n_t. Notice that although Lemmas 15 and 20 are stated for the special case where $\rho \geq 1/3$ and $\varepsilon \leq 0.01$, they can effectively be reproduced for any smaller ρ provided that ε is chosen small enough accordingly.

By Lemma 14, for every point x in a patch triangle t the oriented triangle normal n_t makes an angle of $O(\varepsilon/\rho)$ with $n_{\hat{x}}^+$. In particular when $\varepsilon \leq 0.01$ and $\rho \geq 1/3$,

this angle is at least 23°. Suppose ζ is not one-to-one. Then, there are two points x and x' in $\mathcal{T}$ that are both mapped to the same point $\hat{x}$ by ζ. Consider the line ℓ normal to Σ at $\hat{x}$. This line passes through both x and x'. Assume without loss of generality that x and x' are consecutive intersection points of ℓ and $\mathcal{T}$. Then, at one of x and x' the line ℓ enters and at the other exits the interior bounded by $\mathcal{T}$. In other words, if we orient ℓ along $n_{\hat{x}}^+$, it makes an angle at least $\frac{\pi}{2}$ with one of the oriented normals of $\mathcal{T}$ at x or x', an impossibility.

Next, we argue that $\mathcal{T}$ is a manifold. For this we first observe that each edge in $\mathcal{T}$ is incident to at least two triangles. This of course holds by definition for the interior edges of each surface patch. If a Gabriel edge on the boundary of a surface patch is incident only to that patch, the patch must be incident to the stable manifold of the same inner (or outer) medial axis maximum on both sides. This contradicts Lemma 19. We show now that the triangles incident to each vertex v of $\mathcal{T}$ form a topological disk and hence $\mathcal{T}$ is a 2-manifold. If not, there are two triangles incident to v so that a normal line stabs both of them at points arbitrarily close to v since they lie almost parallel to Σ. This is in contradiction with ζ being one-to-one.

We are left to show that ζ is continuous and onto. The continuity of ζ follows from the fact that the original closest point function $x \mapsto \hat{x}$ is continuous everywhere except at the medial axis. To show that ζ is onto, consider $\zeta(\mathcal{T}) \subseteq \Sigma$. We claim that $\zeta(\mathcal{T}) = \Sigma$. Since $\mathcal{T}$ is a 2-manifold without boundary and ζ maps it homeomorphically to $\zeta(\mathcal{T})$, we have $\zeta(\mathcal{T})$ as a compact 2-manifold without boundary and $\zeta(\mathcal{T}) \subseteq \Sigma$. This is only possible if $\zeta(\mathcal{T}) = \Sigma$ as both $\zeta(\mathcal{T})$ and Σ are compact 2-manifolds without boundary. □

6 Curves in $\mathbb{R}^3$

The separation of critical points and the resulting surface reconstruction algorithm both have analogues when we consider the distance function to an ε-sample of a smooth curve Γ in $\mathbb{R}^3$, instead of that of a surface.

6.1 Separation of critical points

Let P be an ε-sample of a smooth closed curve $\Gamma \subset \mathbb{R}^3$. We analyze the critical points of the distance function h induced by P.

Lemma 4 still holds, i.e., all critical points of h are either near the curve (called the *curve critical points*), or near the medial axis (called the *medial axis critical points*). However, unlike surfaces, not all types of critical points can be near the curve.

Lemma 21 *If the boundary of a ball B intersects Γ in three or more points, then it contains a medial axis point.*

Proof. Shrink B centrally until its boundary becomes tangent to a point, say x of Γ. Then keeping x fixed on the boundary shrink it further by moving its center toward x. Stop when the interior of B becomes empty of Γ. At this moment B is tangent to Γ. If x is the only point of tangency, B is a curvature ball and its center is on the medial axis. If it is tangent to two or more points of Γ, its center is again a medial axis point. In both cases the medial axis point is in the original ball B. □

Lemma 22 *Let c be a critical point of h. If $c \in \Gamma_{\varepsilon^2}$, then c is either an index-0 or an index-1 critical point provided that $\varepsilon < 0.3$.*

Proof. If c is an index-2 or index-3 critical point, we have a ball B centered at c whose boundary intersects Γ in three or more points. Let r be the radius of B. By Lemma 21, B contains a medial axis point and hence $r > \frac{1}{2}f(p)$ for any point $p \in P \cap B$. Therefore, c is at least $f(p)$ distance away from its closest point in P. We claim that c is also at least $\varepsilon^2 f(\hat{c})$ distance away from $\hat{c}$ proving that $c \notin \Gamma_{\varepsilon^2}$.

To reach a contradiction assume that $s(c) = \|c - \hat{c}\|$ is no more than $\varepsilon^2 f(\hat{c})$. The closest sample point, say p, to $\hat{c}$ is within $\varepsilon f(\hat{c})$ distance from it. This point p is within $(\varepsilon + \varepsilon^2)f(\hat{c})$ distance from c. Applying the Lipschitz property of the feature size f as in the proof of Corollary 5, we get that $\|p - c\| \leq \frac{\varepsilon+\varepsilon^2}{1-\varepsilon-2\varepsilon^2} f(p)$. On the other hand we know $\|p - c\| > \frac{1}{2}f(p)$. Thus we reach a contradiction if $\varepsilon < 0.3$. □

6.2 Reconstruction

We will state some more results regarding the critical points of the distance from a curve. These results lead straightforwardly to a reconstruction algorithm. The edges that connect two consecutive points on Γ are called *correct* edges. All other edges are *incorrect*. It is known that all correct edges are Delaunay edges if $\varepsilon < 1/3$. Also, it is easy to show that they intersect their dual Voronoi facets, i.e., they contain index-1 critical points. It is further known that the length of any correct edge pq is at most $\frac{2\varepsilon}{1-\varepsilon} f(p)$ [9]. This means the index-1 critical point which is the midpoint of pq is at most $\frac{\varepsilon}{1-\varepsilon} f(p)$ distance away from its closest sample point which suggests that this critical point cannot lie in $M_{2\varepsilon}$ and hence resides in Γ_{ε^2}. On the other hand, as the next lemma shows, the incorrect edges containing index-1 critical points are longer.

Lemma 23 *Let pq be a Delaunay edge containing an index-1 critical point c.*

(i) If pq is correct then the distance of c from p is at at most $\frac{\varepsilon}{1-\varepsilon} f(p)$.

(ii) If pq is incorrect then the distance of c from p is at least $f(p)/2$.

Proof. If pq is correct, its length is at most $\frac{2\varepsilon}{1-\varepsilon} f(p)$ by Lemma 3.4 of [9]. So, its midpoint c has distance at most $\frac{\varepsilon}{1-\varepsilon} f(p)$ from p proving (i).

Now assume that pq is incorrect. If the boundary of the ball B centered at c with radius $\|c-p\|$ intersects Γ only in p and q then B is a medial ball and the distance of c from p is at least $f(p)$. Otherwise the boundary of B intersects Γ in more than two points and hence B contains a medial axis point by Lemma 21. Therefore, the diameter of B is at least $f(p)$ and the distance of c from p is at least $f(p)/2$ proving (ii). □

The curve critical points can be separated from the medial critical points using an algorithm similar to the one given for surfaces. For a sample point p we determine the nearest critical point. By Lemma 22 and Lemma 23 this critical point has index 1 and is the midpoint of a correct edge. A result of Amenta et al. [2] implies that a correct edge makes an angle of at most $\arcsin(\varepsilon/2)$ with the tangent t_p at p. Thus the vector from p to its nearest critical point makes an angle no larger than $\arcsin(\varepsilon/2)$ with t_p. On the other hand any index-1 critical point in $M_{2\varepsilon}$ is at least $f(p)/2$ distance away from p as a result of Lemma 23. This fact along with the following result due to Dey et al. [8] give the required critical point separation. For a point $p \in \Gamma$, the space spanned by the vectors normal to t_p is called its *normal space*.

Lemma 24 [8] *Let x be a point in V_p where $\|p-x\| \geq \xi f(p)$. Then there is a vector v_p in the normal space of p so that $\angle(v_p, x-p) \leq \arcsin \frac{\varepsilon}{\xi(1-\varepsilon)} + \arcsin \frac{\varepsilon}{1-\varepsilon}$.*

Combining the results of [2], Lemma 23 and Lemma 24 we get the following corollary.

Corollary 7 *Let c be any index-1 critical point on an edge pq. Let c' be the nearest index-1 critical point of p. Then, for $\varepsilon < 0.3$*

(i) $\angle(c'-p, c-p) < 3\pi/4$ *if pq is incorrect, and*

(ii) $\angle(c'-p, c-p) > 3\pi/4$ *if pq is correct.*

We get an immediate separation algorithm for index-1 critical points which also gives a curve reconstruction algorithm: For every point p determine the shortest Gabriel edge pq incident to p. Choose the other Gabriel edge pr incident to p satisfying $\angle(q-p, r-p) > 3\pi/4$. These are the two correct edges for p. Note the similarity between this algorithm and that of [9].

7 Experiments and conclusion

We provide the first theoretical results that link the critical points of the distance function to a tight ε-sampling of a curve or surface embedded in $\mathbb{R}^3$ to either

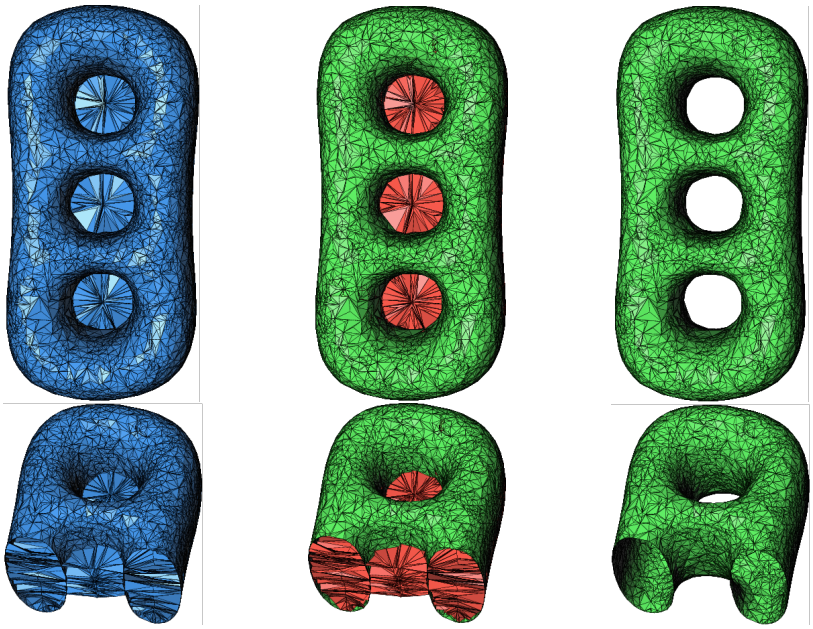

Figure 9: Form left to right: the 2-skeleton of the flow complex; the segmentation of the index-2 stable manifolds into surface (green) and medial axis critical points (red); the stable manifolds of surface critical points only.

the surface (curve) or its medial axis. This allows us to derive reconstruction algorithms for curves and surfaces embedded in $\mathbb{R}^3$ that come with topological and geometric reconstruction guarantees. Although our proofs for normal convergence and orientation of surface patches and, because of those, our proof of homeomorphism depend on the assumption that the given ε-sample is tight, we believe that these guarantees can be strengthened to the case of general ε-sampling.

The output of our algorithm is not a Delaunay sub-complex, a property sometimes desired in practice. However, this output can be easily modified to satisfy this requirement. A natural way is to replace the stable manifold of a surface 2-saddle with the union of Delaunay triangles corresponding to its flow complex triangles (See Appendix A). Notice that doing this, we replace each patch with a patch made of the Delaunay triangles which shares the same boundary of Gabriel edges with the original patch. Furthermore, these triangles are asymptotically as close to the surface as our reconstruction. For a related heuristic method for approximating stable manifolds of maxima with collections of Delaunay tetrahedra see [7].

We experimented with the separation of the critical points for surface samples using an implementation that computes the 2-skeleton of the flow complex, i.e.,

the union of the stable manifolds of index-2 saddle points. It turned out that the union of the stable manifolds of surface critical points (see Figure 9) already gives a good reconstruction. However, this reconstruction is not guaranteed to be (in fact, due to slivers, it is rather unlikely to be) a manifold.

References

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

[2] N. Amenta, M. Bern and D. Eppstein. The crust and the beta-skeleton: combinatorial curve reconstruction. *Graphical Models and Image Processing*, **60**, pp. 125–135, 1998.

[3] N. Amenta, S. Choi, T. K. Dey and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. *Internat. J. Comput. Geom. & Applications*, vol. 12, pp. 125–141, 2002.

[4] N. Amenta, S. Choi and R. Kolluri. The power crust, unions of balls, and the medial axis transform. *Computational Geometry: Theory and Applications*, **19**, pp. 127–153, 2001.

[5] J. D. Boissonnat and F. Cazals. Smooth Surface Reconstruction via Natural Neighbour Interpolation of Distance Functions. *Computational Geometry: Theory and Applications*, **22**, pp. 185–203, 2002.

[6] R. Chaine. A geometric convection approach of 3-D reconstruction. In *Proc. Eurographics Sympos. on Geometry Processing*, pp. 218–229, 2003.

[7] T. K. Dey, J. Giesen and S. Goswami. Shape Segmentation and Matching with Flow Discretization. In *Proc. 8th Workshop on Algorithms Data Strucutres*, pp. 25–36, 2003.

[8] T. K. Dey, J. Giesen, S. Goswami and W. Zhao. Shape dimension and approximation from samples. *Discr. Comput. Geom.*, **29**, pp. 419–434, 2003.

[9] T. K. Dey and P. Kumar. A simple provable algorithm for curve reconstruction. In *Proc. 10th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 893–894, 1999.

[10] T. K. Dey and W. Zhao. Approximating the Medial Axis from the Voronoi Diagram with a Convergence Guarantee. *Algorithmica*, **38**, pp. 179–200, 2004.

[11] H. Edelsbrunner. Surface reconstruction by wrapping finite point sets in space. *Discr. Comput. Geom.*, **32**, pp. 231–244, 2004.

[12] J. Giesen and M. John. The Flow Complex: A Data Structure for Geometric Modeling. In *Proc. 14th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 285–294, 2003.

[13] K. Grove. Critical Point Theory for Distance Functions. In *Proceedings of Symposia in Pure Mathematics* **54**(3), pp. 357–385, 1993.

Appendix A. Stable manifolds of index-2 critical points

An index-2 critical point, i.e., a saddle point s, is the intersection point of a Delaunay triangle t with its dual Voronoi edge e. Under a mild non-degeneracy condition, the stable manifold of s is a surface patch that can be constructed explicitly, see [12]. The degeneracy condition is that the inflow of s does not contain a Voronoi vertex which can be always achieved by an arbitrarily small perturbation of the sample points. We start by constructing a polygon P whose interior points all flow into s. This polygon contains s and is contained in t. To simplify our exposition assume that there are three Voronoi facets incident to every Voronoi edge. We are going to construct a polyline for each of the three Voronoi facets incident to e. These three polylines together make up the boundary of the polygon P. The drivers of the Voronoi facets incident to e are points on their dual Delaunay edges. These Delaunay edges are all in the boundary of t. Note that it is possible that such a driver is a index-1 critical point. First, consider a driver d which is not an index-1 critical point. The line segment that connects d with s is contained in t and intersects the boundary of the corresponding Voronoi facet in two points, namely in s and in a second point s'. We get a polyline from the two segments that connect s' to the two Delaunay vertices incident to the Delaunay edge that contains d. Second, if the driver of the Voronoi facet is a saddle of index-1 we take its dual Delaunay edge as the polyline. That is, we get three polylines all contained in t, one for each Voronoi facet incident to e. Let P be the polygon whose boundary is given by these polylines. P is contained in t and all its interior points flow into s. It can be triangulated by connecting s with the points s' and the Delaunay vertices incident to t. Figure 10 shows two examples of two such polygons P.

Let s' be a point as constructed above for a Voronoi facet that is not driven by an index-1 critical point. By construction s' is contained in a Voronoi edge e'. Furthermore, by our assumption it has to be an interior point of e'. We can assume again that e' is incident to three Voronoi facets. For one of these Voronoi facets we have already computed a polyline. For the remaining two we do it exactly the same way we did it above for P. Thus we have again three polylines, one for each Voronoi facet incident to e'. Two of these polylines always intersect in a common Delaunay vertex. That is, the three polylines together form a polyline which is homeomorphic to $\mathbb{S}^1$. The latter polyline need not be

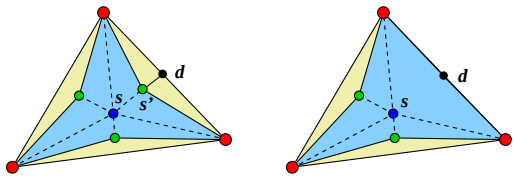

Figure 10: Two examples of polygons that are contained in a Delaunay triangle that is intersected by its dual Voronoi edge in s. The interior points of these polygons flow into s. The polygon in the figure on the right has one index-1 critical point on its boundary.

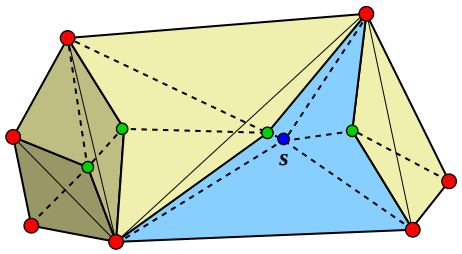

Figure 11: In this example the stable manifold of s is made up from five surface patches. Note that the surface patches need not be planar.

contained in a hyperplane but it can be triangulated by connecting the point s' with newly computed points s' and to the Delaunay vertices incident to the Delaunay facet dual to e'. This gives us a new triangulated surface patch whose interior points all flow into s.

We continue with the above construction until there are no more points s' left for which we have not already constructed a surface patch. The surface of points that flow into the index-2 saddle s is made up from all the patches. By construction the boundary of this surface consists of Gabriel edges, i.e. Delaunay edges. Figure 11 shows an example of the stable manifold of some index-2 critical point.