Surface Reconstruction by Wrapping Finite Sets in Space

Herbert Edelsbrunner*

Abstract

Given a finite point set in $\mathbb{R}^3$, the surface reconstruction problem asks for a surface that passes through many but not necessarily all points. We describe an unambiguous definition of such a surface in geometric and topological terms, and sketch a fast algorithm for constructing it. Our solution overcomes past limitations to special point distributions and heuristic design decisions.

Keywords. Computational geometry, computer graphics, geometric modeling; Delaunay complexes, Morse functions, vector fi elds, acyclic relations, collapsing, deleting.

1 Introduction

The original version of this paper was written in 1995. To preserve that version, we have limited modifications to minor stylistic changes and to the addition of a paragraph that accounts for the new and related work during the years from 1996 to 2001. All citations of work during these five years use letters rather than numbers in the citation.

Problem and solution. The input to the *surface reconstruction problem* is a finite set of points scattered in three-dimensional Euclidean space. The general task is to find a surface passing through the points. There are of course many possible such surfaces, and we would want one that in some sense is most reasonable and best represents the way the input points are distributed. We allow for the case that some points lie off the surface inside the bounded volume. We propose a solution that provides structural information in terms of a mesh or complex connecting the

*Then: Department of Computer Science, University of Illinois at Urbana-Champaign, and Raindrop Geomagic, Champaign-Urbana, Illinois, USA. Now: Department of Computer Science, Duke University, Durham, and Raindrop Geomagic, Research Triangle Park, North Carolina, USA.

points. Section 2 will be more specific about what exactly we mean by this and what properties we expect from the mesh.

The first part of our solution consists of a description of the surface in geometric and topological terms. There are minimum distance functions and ideas from Morse theory turning these functions into vector fields and cell decompositions. For generic data sets, this description is unambiguous and completely determines the surface. The second part of our solution is an efficient algorithm that constructs the defined surface. The algorithm is based on Delaunay complexes and extracts a subcomplex through repeated collapsing. All ideas and results generalize to any arbitrary fixed number of dimensions. For reasons of specificity, the discussion in this paper is exclusively three-dimensional.

Work prior to 1995. The surface reconstruction problem has a long history. Most of the previous work assumes some kind of additional structure given along with the data points. A common assumption is that the points lie on curves defined by slicing a surface with a collection of parallel planes [9, 13]. The surface reconstruction is reduced to a sequence of steps, each connecting two curves in contiguous planes. Another common assumption is differentiability [11]. The surface is constructed from patches defining diffeomorphisms between $\mathbb{R}^2$ and local neighborhoods on the surface. Fairly dense point distributions are required to allow the reconstruction of tangents and normals.

We are interested in the *general* surface reconstruction problem that admits no assumption other than that the input consists of finitely many points in $\mathbb{R}^3$. At the time of writing this paper in 1995, we found only three pieces of work studying the general problem. In two cases, the surface or shape is obtained from the three-dimensional Delaunay complex of the input points. This is also the approach followed in this paper. Boissonnat [1] compromises the global nature of the approach by using local rules for removing simplices from the Delaunay complex. The resulting surfaces are somewhat unpredictable and not amenable to rational analysis. Edelsbrunner and Mücke [6] use distance relationships to identify certain subcomplexes of the Delaunay complex as alpha shapes of the given data set. For uneven densities, these shapes tend to either exhibit a lack of detail in dense regions or gaps and holes in sparse regions. Rather than subcomplexes of the Delaunay complex, Veltkamp [15] uses two-parameter neighborhood graphs to form surfaces from points in space. Depending on the choice of the parameters the graphs may exhibit self-intersections or poor shape representation.

Development after 1995. The algorithm described in this paper has been implemented in 1996 at Raindrop Geomagic, which successfully commercialized it as `geomagic Wrap`. It is also described in U. S. Patent No. 6,3777,865, which has issued on April 23, 2002.

The surface reconstruction problem has enjoyed increasing popularity over the last few years, both in computer graphics and in computational geometry. A number of essentially two-dimensional algorithms that rely primarily on density and smoothness assumptions of the data have been developed [C, E, F, J]. In parallel, the use of three-dimensional Delaunay complexes has been refined [A, B, D, I]. The focus of that work is the detailed study of Delaunay complexes for data sets that satisfy density and smoothness assumptions, and to exploit their special structure for surface reconstruction. Possibly surprisingly, neither development has come close to reproducing the ideas presented in this paper.

From a completely different angle, Robin Forman's development of a discrete Morse theory for simplicial complexes [H] is related to work in this paper. According to Forman, a *discrete Morse function* is a map f from the collection of simplices to the real numbers such that for every simplex τ the following two conditions hold:

(1) there is at most one face $\upsilon \leq \tau$ with $\dim \upsilon = \dim \tau - 1$ and $f(\upsilon) \geq f(\tau)$, and

(2) there is at most one coface $\sigma \geq \tau$ with $\dim \sigma = \dim \tau + 1$ and $f(\sigma) \leq f(\tau)$.

The theory developed in this paper uses relations that correspond to functions violating these conditions and thus does not seem to fit into Forman's framework. It would be interesting to elucidate the connection between the two approaches to a discrete Morse theory. Another recent development that resonates with the work in this paper is the introduction of persistent Betti numbers [G]. It relates to the discussion of granularity in Section 9, in which it is suggested to construct coarse-grained decompositions of the Delaunay complex by merging discrete stable manifolds, possibly by suppressing some of the less persistent critical points. Maybe the time has come to integrate all these ideas and to develop a hierarchical approach to surface reconstruction based on a more extensive use of algebraic structures developed in topology.

Outline. Section 2 displays sample results obtained with software implemented at Raindrop Geomagic. Section 3 reviews Delaunay complexes for finite point sets in $\mathbb{R}^3$. Section 4 reviews notions from Morse theory and constructs a family of Morse functions from local distance information. Section 5 derives an ordering principle for the Delaunay simplices from the Morse functions. Section 6 studies mechanisms to cluster simplices based on the ordering. Section 7 defines the basic surface construction as a sequence of collapses. Section 8 discusses generalizations of the basic construction with and without interactive surface modification. Section 9 mentions possible extensions of the presented results and related open questions.

2 Examples and Properties

We use notation and terminology from combinatorial topology [10] to describe the surfaces and the algorithm that constructs them. In a nut-shell, that algorithm starts with the Delaunay complex of the input set and constructs an acyclic partial order over its simplices. This order is motivated by a continuous flow field in which every point is attracted by the Voronoi vertex that is nearest in a weighted sense. The Voronoi vertex at infinity corresponds to a dummy simplex in the relation, and the reconstructed surface is obtained by sculpting away all predecessors of that dummy simplex. We begin by giving a few examples constructed by software implementing the algorithm.

Sample surfaces. In the simplest and possibly most common case, the surface constructed from a finite point set in $\mathbb{R}^3$ is connected like a sphere, as in Figure 1. With our approach, it is possible to modify the construction and to introduce tunnels, as illustrated in Figure 2. There are also cases, in which the surface cannot be naturally closed and remains connected like a disk, as in Figure 3. Finally, if the points are lined up, it is possible that the surface reconstructed by our algorithm degenerates to a curve, as in Figure 4.

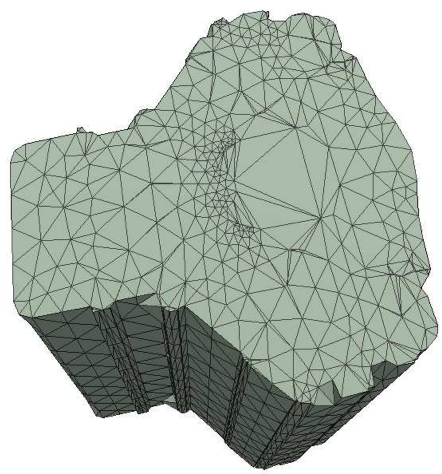

Figure 1: An engine block surface homeomorphic to a sphere.

In technical terms, the constructed surface is a simplicial complex of dimension 2 with the topology of a possibly pinched sphere. If forced by the distribution of the data points, the complex can be one- or zero-dimensional. We first introduce the relevant terminology and then describe the reconstructed surface in more detail.

Spaces and maps. All topological spaces in this paper are subsets of Euclidean space of some dimension k, denoted by $\mathbb{R}^k$. Without exception, we use the topology induced by the Euclidean metric in $\mathbb{R}^k$. The Euclidean distance between points x and y is denoted by $\|x - y\|$, and the *norm* of x is the distance from the origin, which is $\|x\| = \|x - 0\|$. Other than for Euclidean k-dimensional space, we need short notation for the k-dimensional sphere and the k-dimensional ball,

$$\begin{aligned} \mathbb{S}^k &= \{x \in \mathbb{R}^{k+1} \mid \|x\| = 1\}, \\ \mathbb{B}^k &= \{x \in \mathbb{R}^k \mid \|x\| \leq 1\}. \end{aligned}$$

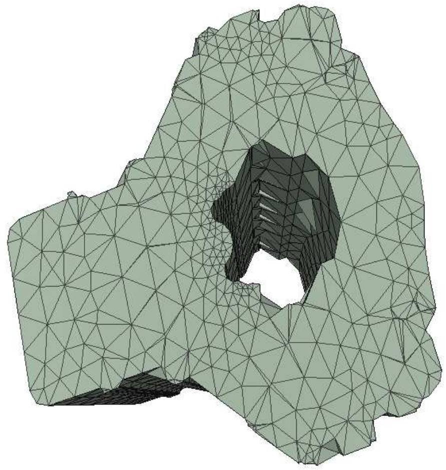

Figure 2: The engine block surface in Figure 1 after pushing open a tunnel.

We refer to $\mathbb{S}^k$ as the *k-sphere* and to $\mathbb{B}^k$ as the *k-ball*. For example, the 1-sphere is a circle, the 0-sphere is a pair of points, and the (-1)-sphere is the empty set. The 2-sphere is what we ordinarily call a sphere. The 2-ball is a closed disk, the 1-ball is a closed interval, and the 0-ball is a point.

Topological spaces are compared via continuous functions referred to as *maps*. A *homeomorphism*, $\beta : \mathbb{X} \to \mathbb{Y}$, is a continuous bijection with continuous inverse. The inverse of a homeomorphism is again a homeomorphism. $\mathbb{X}$ and $\mathbb{Y}$ are *homeomorphic* or *topologically equivalent*, denoted $\mathbb{X} \approx \mathbb{Y}$, if there is a homeomorphism between them. An *embedding* is an injective map $\iota : \mathbb{X} \to \mathbb{Y}$ whose restriction to the image, $\iota(\mathbb{X})$, is a homeomorphism.

A *homotopy* between two maps $a, b : \mathbb{X} \to \mathbb{Y}$ is a continuous function $F : \mathbb{X} \times [0, 1] \to \mathbb{Y}$ with $F(x, 0) = a(x)$ and $F(x, 1) = b(x)$ for all $x \in \mathbb{X}$. $\mathbb{X}$ and $\mathbb{Y}$ are *homotopy equivalent*, denoted by $\mathbb{X} \simeq \mathbb{Y}$, if there are maps $f : \mathbb{X} \to \mathbb{Y}$ and $g : \mathbb{Y} \to \mathbb{X}$ and homotopies between $g \circ f$ and the identity for $\mathbb{X}$ and between $f \circ g$ and the identity for $\mathbb{Y}$. Two spaces have the same *homotopy type* if they are homotopy equivalent. A space is *contractible* if it is homotopy equivalent to a point. Homotopy equivalence is weaker than topological equivalence, in the sense that $\mathbb{X} \approx \mathbb{Y}$ implies $\mathbb{X} \simeq \mathbb{Y}$.

Simplicial complexes. We use combinatorial structures to represent topological spaces in the computer. We begin by introducing the geometric elements that make up these structures. The convex hull of a finite collection of points U is denoted as $\operatorname{conv} U$. A *k-simplex*, σ, is the convex hull of $k+1$ affinely independent points. The *dimension* of σ is $\dim \sigma = k$. At most four points can be affinely independent in $\mathbb{R}^3$ and we have four types of simplices: *vertices* or 0-simplices, *edges* or 1-simplices,

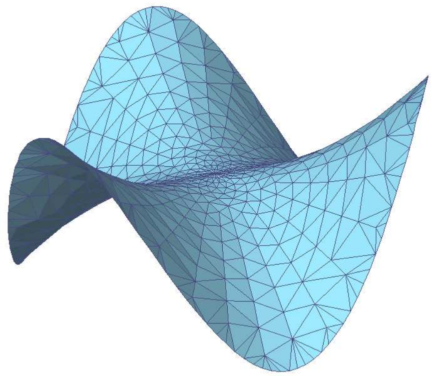

Figure 3: Points on a saddle surface triangulated to form a patch homeomorphic to a disk.

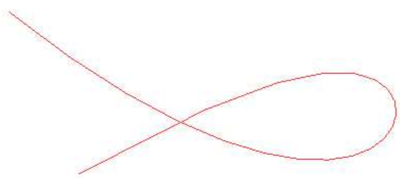

Figure 4: Points on the moment curve connected to form a curve homeomorphic to an interval.

triangles or 2-simplices, and *tetrahedra* or 3-simplices. A simplex $\tau = \operatorname{conv} T$ is a *face* of another simplex $\sigma = \operatorname{conv} U$, and σ is a *coface* of τ, if $T \subseteq U$. We denote this relationship by $\tau \leq \sigma$. The *boundary* of σ, $\operatorname{bd} \sigma$, is the union of all proper faces, and the *interior* is $\operatorname{int} \sigma = \sigma - \operatorname{bd} \sigma$. For example, the boundary of an edge consists of its two endpoints and the interior is the open edge, without endpoints. The boundary of a vertex is empty and the interior is the vertex itself.

A *simplicial complex*, $\mathcal{K}$, is a finite collection of simplices such that $\sigma \in \mathcal{K}$ and $\tau \leq \sigma$ implies $\tau \in \mathcal{K}$, and $\sigma, \sigma' \in \mathcal{K}$ implies that $\sigma \cap \sigma'$ is either empty or a face of both. The *dimension* of $\mathcal{K}$ is the maximum dimension of any of its simplices. A *principal simplex* has no proper coface in $\mathcal{K}$. For example, if $\dim \mathcal{K} = 3$ then every tetrahedron in $\mathcal{K}$ is a principal simplex. A *subcomplex* is a simplicial complex $\mathcal{L} \subseteq \mathcal{K}$. The *vertex set* of $\mathcal{K}$ is $\operatorname{Vert} \mathcal{K} = \{\sigma \in \mathcal{K} \mid \dim \sigma = 0\}$. $\mathcal{K}$ is *connected* if for every non-trivial partition $\operatorname{Vert} \mathcal{K} = V_1 \dot{\cup} V_2$ there is a simplex $\sigma \in \mathcal{K}$ that has vertices in V_1 and in V_2. The *underlying space* is $|\mathcal{K}| = \bigcup_{\sigma \in \mathcal{K}} \sigma$. The interiors of simplices partition the underlying space. In other words, for every $x \in \|\mathcal{K}\|$ there is a unique $\sigma \in \mathcal{K}$ with $x \in \operatorname{int} \sigma$.

Special subsets and subcomplexes. The *star* of a simplex τ in a simplicial complex $\mathcal{K}$ is the set of cofaces, the *closure* of a subset $L \subseteq \mathcal{K}$ is the smallest subcomplex that contains L, and the *link* of a simplex $\tau \in \mathcal{K}$ is the set of simplices in the closed star that are disjoint from τ:

$$\begin{aligned} \operatorname{St} \tau &= \{\sigma \in \mathcal{K} \mid \tau \leq \sigma\}, \\ \operatorname{Cl} L &= \{\tau \in \mathcal{K} \mid \tau \leq \sigma \in L\}, \\ \operatorname{Lk} \tau &= \{\sigma \in \operatorname{Cl} \operatorname{St} \tau \mid \sigma \cap \tau = \emptyset\}. \end{aligned}$$

In other words, the link consists of all simplices in the closed star of τ that do not belong to the (open) star of any face of τ. Links can be used to introduce combinatorial notions of interior and boundary. They are defined relative to the space that contains the complex, which in this paper is $\mathbb{R}^3$. The *interior* of $\mathcal{K}$ is the set of simplices, σ, whose links are homeomorphic to spheres of appropriate dimension, and the *boundary* consists of all other simplices:

$$\begin{aligned} \operatorname{Int} \mathcal{K} &= \{\sigma \in \mathcal{K} \mid \|\operatorname{Lk} \sigma\| \approx \mathbb{S}^{2-\dim \sigma}\}, \\ \operatorname{Bd} \mathcal{K} &= \mathcal{K} - \operatorname{Int} \mathcal{K}. \end{aligned}$$

In $\mathbb{R}^3$, the boundary of a simplicial complex consists of all simplices of dimension 2 or less that are not completely surrounded by tetrahedra.

Surface properties. Let $S \subseteq \mathbb{R}^3$ be finite. The solution to the surface reconstruction problem proposed in this paper is a simplicial complex, $\mathcal{W}$. Its underlying space is what we call a *pinched* 2*-sphere*, that is, $\|\mathcal{W}\|$ is the image of a map $\varphi : \mathbb{S}^2 \to \mathbb{R}^3$ and every neighborhood of φ contains an embedding of $\mathbb{S}^2$ in $\mathbb{R}^3$. Another way to express the latter condition is that for every real $\varepsilon > 0$ there is an embedding $\iota : \mathbb{S}^2 \to \mathbb{R}^3$ with $\|\varphi(x) - \iota(x)\| < \varepsilon$ for every $x \in \mathbb{S}^2$. As we will see in Section 6 — which contains the formal definition of $\mathcal{W}$ — not every topological type of a pinched 2-sphere can be realized by $\mathcal{W}$. Here we just list a few not necessarily independent properties of $\mathcal{W}$:

(P1) $\operatorname{Vert} \mathcal{W} \subseteq S$.

(P2) $\mathcal{W}$ is connected.

(P3) $\mathbb{R}^3 - \|\mathcal{W}\|$ consists of $c + 1 \geq 1$ open components exactly one of which is unbounded.

Let Ω be the unbounded component, and let $\mathcal{X} \supseteq \mathcal{W}$ be a complex triangulating the complement of Ω, that is, $\|\mathcal{X}\| = \mathbb{R}^3 - \Omega$. Our algorithm implicitly constructs such a complex $\mathcal{X}$ that satisfies the following again not necessarily independent properties:

(P4) $\operatorname{Vert} \mathcal{X} = S$.

(P5) If $c = 0$ then $\mathcal{X} = \mathcal{W}$.

(P6) $\|\mathcal{X}\|$ is contractible.

(P7) $\mathcal{W} = \mathrm{Bd}\,\mathcal{X}$.

(P8) $\mathcal{X}$ may have principal simplices of dimension less than 3.

We also consider variants of the basic construction yielding complexes $\mathcal{W}$ that violate property (P2) and complexes $\mathcal{X}$ that violate (P4) and (P6).

3 Delaunay Complexes

The complexes $\mathcal{W}$ and $\mathcal{X}$ of Section 2 are both constructed as subcomplexes of the Delaunay complex of the data set. This section introduces Delaunay complexes and mentions properties relevant to the discussions in this paper.

Voronoi cells and Delaunay simplices. Let S be a finite set in $\mathbb{R}^3$. The *Voronoi cell* of $p \in S$ is

$$V_p = \{x \in \mathbb{R}^3 \mid \|x - p\| \leq \|x - q\| \text{ for all } q \in S\}.$$

Let $V_T = \{V_p \mid p \in T\}$ for every $T \subseteq S$. Any two Voronoi cells have disjoint interiors and the collection of all Voronoi cells, V_S, covers the entire $\mathbb{R}^3$. Throughout this paper, we assume the generic case, in which no four points lie on a common plane and no five points lie on a common sphere. The algorithmic justification of this assumption is a simulated perturbation, as described in [5]. In the generic case, two Voronoi cells are either disjoint or they meet along a common two-dimensional face. Three cells either have no points in common or they meet along a common line segment or half-line. Four cells either have no points in common or they meet in a common point. Five or more cells have no points in common.

The intersection pattern among the Voronoi cells can be recorded using a simplicial complex. More specifically, the *Delaunay complex* of S, defined as

$$\mathrm{Del}\,S = \{\mathrm{conv}\,T \mid \bigcap V_T \neq \emptyset\},$$

is such a record. The non-degeneracy assumption implies that $\mathrm{Del}\,S$ is a simplicial complex in $\mathbb{R}^3$. Its underlying space is $\|\mathrm{Del}\,S\| = \mathrm{conv}\,S$, and its vertex set is $\mathrm{Vert}\,\mathrm{Del}\,S = S$. Whether or not a simplex belongs to $\mathrm{Del}\,S$ can be decided by a local geometric test. Call a sphere in $\mathbb{R}^3$ *empty* if all points of S lie on or outside the sphere.

FACT 1. $\sigma = \mathrm{conv}\,T \in \mathrm{Del}\,S$ iff there is an empty sphere that passes through all points of T.

If σ is a tetrahedron then $\mathrm{card}\,T = 4$ and there is a unique sphere $\Sigma = \Sigma(\sigma)$ passing through the four points. We call Σ the *orthosphere* of σ or T. $\mathrm{Del}\,S$ contains exactly all tetrahedra whose orthospheres are empty. The triangles, edges, and vertices in $\mathrm{Del}\,S$ are the faces of these tetrahedra.

Weighted points. In some circumstances, it is useful to assign real weights to the points in S. With the proper generalization of definitions, all results extend from the unweighted to the more general weighted case, in which S is a finite subset of $\mathbb{R}^3 \times \mathbb{R}$. It is convenient to be ambiguous about the meaning of a point p in S: it can either be the weighted point $p \in \mathbb{R}^3 \times \mathbb{R}$ or its projection to $\mathbb{R}^3$. In either case, the *weight* is denoted by $w_p \in \mathbb{R}$. The *weighted square distance* of a point $x \in \mathbb{R}^3$ from the weighted point p is $\pi_p(x) = \|x - p\|^2 - w_p$. Voronoi cells and Delaunay simplices can be defined as before, substituting weighted square distance for Euclidean distance. We still have $|\operatorname{Del} S| = \operatorname{conv} S$. It is possible that $\operatorname{Vert} \operatorname{Del} S$ is not equal to but rather a proper subset of S. Specifically, $p \in S$ is not a vertex of $\operatorname{Del} S$ iff its Voronoi cell is empty.

To extend the local criterion expressed by Fact 1, we need to generalize the notion of orthosphere. A sphere with center $z \in \mathbb{R}^3$ and radius r is *orthogonal* to a weighted point p if $\|p - z\|^2 = w_p + r^2$, and it is *further than orthogonal* if $\|p - z\|^2$ exceeds $w_p + r^2$. Here, $r^2 \in \mathbb{R}$ and it is convenient to choose r from the set of non-negative multiples of the real and the imaginary units, 1 and $\sqrt{-1}$. We call a sphere *empty* if it is orthogonal to or further than orthogonal from all weighted points in S.

FACT 2. $\sigma = \operatorname{conv} T \in \operatorname{Del} S$ iff there is an empty sphere orthogonal to all weighted points in T.

Again we assume non-degeneracy. In the weighted case, this means that every four weighted points have a unique sphere orthogonal to all of them, and no five points have such a sphere. The *orthosphere* of a tetrahedron $\sigma = \operatorname{conv} T$ is the sphere orthogonal to the four weighted points in T. $\operatorname{Del} S$ contains exactly all tetrahedra with empty orthosphere. The triangles, edges, and vertices in $\operatorname{Del} S$ are the faces of these tetrahedra.

Relative position. Call the non-empty intersection of $k+1$ Voronoi cells an *ℓ-cell*, for $0 \leq k \leq 3$ and $\ell = 3 - k$. An intersection of Voronoi cells $\nu = \bigcap V_T$ is an ℓ-cell iff $\sigma = \operatorname{conv} T$ is a k-simplex in $\operatorname{Del} S$. We are interested in the position of ν and σ relative to each other in space. Their affine hulls are orthogonal flats of complementary dimension, $\ell + k = 3$, that intersect at a point y. In the assumed generic case, y is either contained in the interior of ν or it lies outside ν. Similarly, y is either contained in $\operatorname{int} \sigma$ or $y \notin \sigma$. We distinguish four mutually exclusive cases:

(R1) $\operatorname{int} \nu \cap \operatorname{int} \sigma \neq \emptyset$,

(R2) $\operatorname{int} \nu \cap \operatorname{int} \sigma = \emptyset$ and $\operatorname{int} \nu \cap \operatorname{aff} \sigma \neq \emptyset$,

(R3) $\operatorname{int} \nu \cap \operatorname{int} \sigma = \emptyset$ and $\operatorname{aff} \nu \cap \operatorname{int} \sigma \neq \emptyset$,

(R4) $\operatorname{int} \nu \cap \operatorname{aff} \sigma = \operatorname{aff} \nu \cap \operatorname{int} \sigma = \emptyset$.

Figure 5 illustrates all cases for all values of k. For $k = 0$, only Cases (R1) and (R3) and for $k = 3$ only Cases (R1) and (R2) are possible. The ℓ-cell ν is the set of points $x \in \mathbb{R}^3$ for which the sphere with center x and radius $\pi_p(x)$, with $p \in T$,

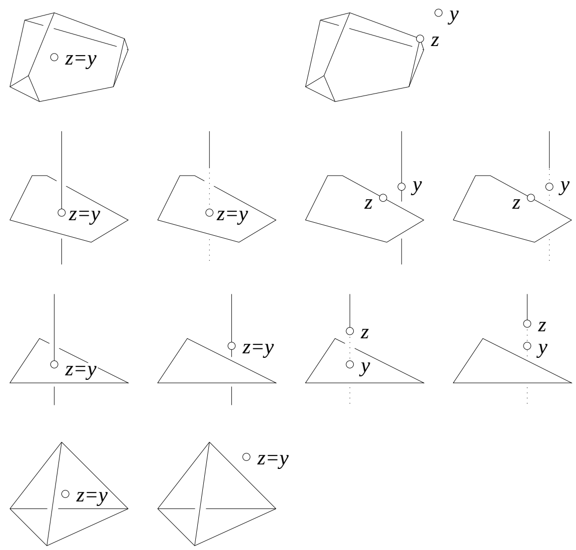

Figure 5: From left to right: Cases (R1), (R2), (R3), and (R4), and from top to bottom: $k = 0$, 1, 2, and 3. In each case, the center of the smallest empty sphere orthogonal to all $p \in T$ is z and the intersection of the two affi ne hulls is y.

is empty and orthogonal to all points in T. The smallest such sphere is centered at the point $z \in \nu$ closest to the point $y = \text{aff}\,\nu \cap \text{aff}\,\sigma$. This implies that for $k = 1$, the Cases (R2) and (R4) cannot occur unless the points are weighted. For $k = 2$, all four cases are possible even in the unweighted case.

4 Morse Functions

The complex $\mathcal{X}$ of Section 2 is a subcomplex of $\text{Del}\,S$ constructed by collapsing Delaunay simplices from the outside in. To decide which simplices to collapse and which not, we construct an acyclic relation among the Delaunay simplices, which is motivated by a particular family of Morse functions. This section constructs these functions after reviewing relevant concepts from Morse theory. The reader interested in a more complete account of that theory is referred to Milnor [14] or Wallace [16].

Vector fields and flow curves. We are interested in smooth functions $f : \mathbb{R}^3 \to \mathbb{R}$ that satisfy a few genericity assumptions. Smoothness means that f is continuous and infinitely often differentiable, but we will see later that this can be weakened to twice differentiable or even only once differentiable and almost everywhere twice differentiable. The *gradient* of f is a vector field $\nabla f : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ defined by

$$\nabla f(x) \quad = \quad \left(\frac{\partial f}{\partial x_1}(x), \frac{\partial f}{\partial x_2}(x), \frac{\partial f}{\partial x_3}(x) \right).$$

It is smooth because f is smooth. A point $y \in \mathbb{R}^3$ is *critical* if $\nabla f(y) = (0,0,0)$. The *Hessian* at y is the three-by-three matrix of partial second derivatives. A critical point is *non-degenerate* if the Hessian at that point has full rank. Non-degenerate critical points are necessarily isolated. A *Morse function* is a smooth function, f, whose critical points are non-degenerate. The fundamental Morse lemma asserts that for every critical point y there are local coordinates originating from y so that

$$f(x) \quad = \quad f(y) \pm x_1^2 \pm x_2^2 \pm x_3^2$$

for all x in a neighborhood of y. The number of minus signs is the *index* of y. For example, critical points of index 0 are local minima, and critical points of index 3 are local maxima. Critical points of index 1 and 2 are different types of saddle points.

The gradient of f defines a first-order differential equation. A solution is a maximal embedding $\iota : \mathbb{R} \rightarrow \mathbb{R}^3$ whose tangent vectors agree with the gradient of f. For each non-critical point $x \in \mathbb{R}^3$, there is a unique solution or *flow curve* ι_x that contains x, that is, $x = \iota_x(t)$ for some $t \in \mathbb{R}$. It follows that two flow curves are either the same or disjoint. The orientation of $\mathbb{R}$ from $-\infty$ to $+\infty$ imposes an orientation on the flow curve. It is convenient to compactify $\mathbb{R}^3$ to $\mathbb{S}^3$ by stipulating a critical point ω at infinity. Then every flow curve starts at a critical point and ends at a critical point. In this paper, ω will only be an endpoint of flow curves and we define its index to be 3. Let $C(f)$ be the collection of critical points, including ω. The *stable manifold* of $y \in C(f)$ is

$$M_y \quad = \quad \{y\} \cup \{x \in \mathbb{R}^3 \mid \iota_x(t) \rightarrow y \text{ as } t \rightarrow \infty\}.$$

If the index of y is j, then the stable manifold consists of y and a $(j-1)$-dimensional sphere of flow curves. For all $y \neq \omega$, M_y is the image of an injective map from $\mathbb{R}^j$ to $\mathbb{R}^3$. If M_y fails to be homeomorphic to $\mathbb{R}^j$ that is only because it is possible that flow curves ending at y share the same starting point. M_ω is the image of an injective map from $\mathbb{R}^3 - \{0\}$, the punctured three-dimensional space, to $\mathbb{R}^3$. The stable manifolds are mutually disjoint open sets that cover the entire $\mathbb{R}^3$. We call this the *complex of stable manifolds* and write

$$\mathrm{Sm}\, f \quad = \quad \{M_y \mid y \in C(f)\}.$$

The *j-cells* of $\mathrm{Sm}\, f$ are the stable manifolds of critical points of index j.

Distance from empty spheres. Given a finite set S of weighted or unweighted points in $\mathbb{R}^3$, we construct a real-valued function by considering empty spheres. Think of a sphere $\Sigma = (z, r)$ as a weighted point $(z, r^2) \in \mathbb{R}^3 \times \mathbb{R}$. With this interpretation, the weighted square distance of a point $x \in \mathbb{R}^3$ from Σ is well-defined as $\pi_\Sigma(x) = \|x - z\|^2 - r^2$. Consider the function that maps every point $x \in \mathbb{R}^3$ to the minimum weighted square distance from any empty sphere. For a point inside the convex hull of S, the minimum weighted square distance is defined by the orthosphere of a Delaunay tetrahedron. For a point outside the convex hull, the minimum weighted square distance is defined by an infinitely large sphere or hyperplane that supports the convex hull in a triangle. This sphere can be interpreted as the orthosphere of an infinitely large tetrahedron spanned by the triangle and a point at infinity.

The exact shape of this infinitely large tetrahedron can be obtained by constructing the Voronoi cells for the set of orthospheres, including the ones of infinite size. To eliminate the remaining ambiguity, we approximate each infinitely large sphere by a sphere of radius $\frac{1}{\delta}$ that is orthogonal to the same three weighted points. As δ goes to zero, some of the Voronoi cells do not change, some grow and eventually become unbounded, and some cells disappear to infinity. The first kind are the original Delaunay tetrahedra, and the second kind are the infinitely large tetrahedra defined by convex hull triangles. Together, the two types of tetrahedra cover the entire $\mathbb{R}^3$. Let D^3 be the set of tetrahedra of both types. We construct $g : \mathbb{R}^3 \to \mathbb{R}$ by defining $g(x) = \max\{-\pi_{\Sigma(\sigma)}(x) \mid \sigma \in D^3\}$. For a point x, the relevant orthosphere is defined by the tetrahedron that contains the point x.

FACT 3. If $x \in \sigma \in D^3$ then $g(x) = -\pi_{\Sigma(\sigma)}(x)$.

Note that $g(x) = +\infty$ if $x \notin \operatorname{conv} S$. This is a slight inconvenience in the subsequent discussion. All difficulties can be finessed by again approximating each infinitely large sphere by a sphere of radius $\frac{1}{\delta}$. To simplify the discussion, we do not explicate this approximation, but we do pretend that g is a continuous map from $\mathbb{R}^3$ to $\mathbb{R}$.

Smoothing. The function g is continuous but does not quite qualify as a Morse function because it is not everywhere smooth. Smooth functions can be derived from g using appropriate cut-off functions blending between adjacent Delaunay tetrahedra. Figure 6 illustrates the effect of smoothing on Delaunay edges in $\mathbb{R}^2$.

We construct an infinite family of smooth functions f_ε approximating g. Consider the graph of $g : \mathbb{R}^3 \to \mathbb{R}$, which is $G = \{(x, g(x)) \mid x \in \mathbb{R}^3\} \subseteq \mathbb{R}^4$. It is a three-dimensional manifold that consists of finitely many smooth patches, which fit together in a continuous but non-differentiable manner. Let $\varepsilon > 0$ be real, let $b_\varepsilon = \{u \in \mathbb{R}^4 \mid \|u\| \le \varepsilon\}$, and consider $G + b_\varepsilon = \{z + u \mid z \in G, u \in b_\varepsilon\}$. For sufficiently small positive ε, the lower boundary of $G + b_\varepsilon$ is the graph of a differentiable function

$$f_\varepsilon(x) = \min\{r \mid (x, r) \in G + b_\varepsilon\}.$$

The f_ε are not smooth in the sense of being infinitely often differentiable. Still, they are everywhere differentiable and almost everywhere twice differentiable, which

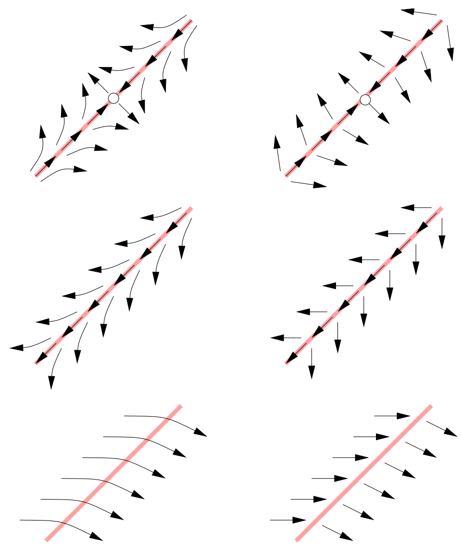

Figure 6: From top to bottom: the gradient of f_ε on the left and its limit, ∇g, on the right for a centered, a confident, an equivocal edge in $\mathbb{R}^2$. The centered and confident edges repel a flow curve unless it lies exactly on the edge. The equivocal edge is crossed by an interval of bending flow curves.

suffices for the purposes of this paper. Most importantly, the notions of gradient, critical point, and flow curve are defined. For example, outside the convex hull of S, g and therefore the f_ε are approximately infinitely steep and the flow curves go quickly to infinity. By assumption of non-degeneracy, the f_ε are twice differentiable at all critical points, and we can define indices and stable manifolds as before.

Limit construction. We use limit considerations to construct a vector field for g. For every point $x \in \mathbb{R}^3$, define $\nabla g(x) = \lim_{\varepsilon \to 0} \nabla f_\varepsilon(x)$. ∇g is a vector field albeit not continuous because g is not differentiable, as seen in Figure 6. Observe that the f_ε have identical critical points. In other words, $C(f_\varepsilon) = C(f_\delta)$ for any two sufficiently small $0 < \varepsilon < \delta$. The following relation between g and the f_ε is fairly straightforward to prove.

FACT 4. $y \in \mathbb{R}^3$ is a critical point of f_ε iff $\nabla g(y) = (0, 0, 0)$, and the index of y is j iff the Delaunay simplex σ that contains y in its interior has dimension j.

The vector field ∇g does not enjoy some of the nice properties of the ∇f_ε. In particular, ∇g is not continuous. We finesse some of the resulting difficulties with limit considerations. As an example, consider a non-critical point $x \in \mathbb{R}^3$. For each sufficiently small $\varepsilon > 0$, there is a unique flow curve $\iota_{x,\varepsilon}$ of ∇f_ε that contains x. Define the *limit curve* for x as $\lambda_x = \lim_{\varepsilon \to 0} \iota_{x,\varepsilon}$. The curve λ_x is a continuous though generally not a smooth embedding of $\mathbb{R}$ in $\mathbb{R}^3$. Indeed, λ_x is piecewise linear, and for each simplex σ, $\lambda_x \cap \operatorname{int}\sigma$ is either empty, a point, or a line segment. While two flow curves are either disjoint or the same, two limit curves can partially overlap. However, once they separate they stay apart. In other words, if $x \in \lambda_u \cap \lambda_v$ then the portions of λ_u and λ_v preceding x are the same. The reason for this is the repulsion of nearby flow curves by centered and confident simplices. Only equivocal simplices attract nearby flow curves, but these curves pass right through the simplex, without ambiguity or merging of curves.

Based on the definition of the λ_x, we can consider limits of stable manifolds and the complex they form. These limits provide the guiding principle motivating our surface construction method. At his moment, we recall that $\mathcal{W}$ refers to the surface and $\mathcal{X}$ refers to a triangulation of the portion of $\mathbb{R}^3$ bounded by $\mathcal{W}$.

INTUITION. In the limit, $\mathcal{X}$ triangulates the complement of the stable manifold of ω and $\mathcal{W}$ is the boundary of $\mathcal{X}$.

We will take small liberties in translating this intuition into an unambiguous construction, which will be given in Section 6.

5 Ordering Simplices

The flow and limit curves motivate the construction of an acyclic relation over the set of Delaunay simplices. The complexes $\mathcal{W}$ and $\mathcal{X}$ of Section 2 are then constructed by collapsing simplices from back to front in the relation.

Flow relation. Introduce a dummy simplex, ω, that represents the outside, or complement of $|\operatorname{Del} S|$. It replaces the collection of infinitely large tetrahedra introduced in Section 4. All tetrahedra in this collection have the same flow behavior and can be treated uniformly. We deliberately choose the same name, ω, for the dummy simplex and the dummy critical point, and this will not cause any confusion. By definition, the faces of ω are the convex hull faces, which are the simplices in $\operatorname{Bd}\operatorname{Del} S$. Let $D = \operatorname{Del} S \cup \{\omega\}$. The *flow relation*, $\prec \, \subseteq D \times D$, is constructed to mimic the behavior of the limit curves. Specifically, $\tau \prec \upsilon \prec \sigma$ if υ is a proper face of τ and of σ and there is a point $x \in \operatorname{int}\upsilon$ with λ_x passing from $\operatorname{int}\tau$ through x to $\operatorname{int}\sigma$. We pronounce this as τ precedes υ and υ precedes σ. The condition implies that every neighborhood of x contains a non-empty subset of $\lambda_x \cap \operatorname{int}\tau$ and a non-empty subset of $\lambda_x \cap \operatorname{int}\sigma$. We call τ a *predecessor* and σ a *successor* of υ. The sets of *descendents* and *ancestors* are

$$\begin{aligned} \operatorname{Des}\upsilon &= \{\upsilon\} \cup \bigcup_{\upsilon \prec \sigma} \operatorname{Des}\sigma, \\ \operatorname{Anc}\upsilon &= \{\upsilon\} \cup \bigcup_{\tau \prec \upsilon} \operatorname{Anc}\tau. \end{aligned}$$

It is convenient to study the flow relation by distinguishing three types of Delaunay simplices: centered, confident, and equivocal. These are related to the classification of Delaunay simplices introduced in Section 3 and illustrated in Figure 5. The three types are mutually exclusive and exhaust all possible Delaunay simplices. For each type, we specify predecessors and successors in terms of their local geometric properties.

Centered simplices. A simplex $\sigma \in D$ is *centered* if, for every $x \in \operatorname{int} \sigma$, the portion of λ_x succeeding x is contained in $\operatorname{int} \sigma$. Illustrations of a centered edge and a centered triangle can be seen in Figures 6 and 7. In words, limit curves enter but do not exit $\operatorname{int} \sigma$. In the case of $\sigma = \omega$, the limit curves diverge, and in the case

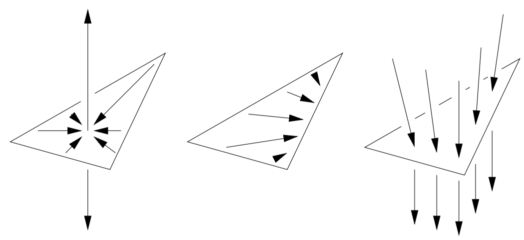

Figure 7: From left to right: a centered, a confident, an equivocal triangle in $\mathbb{R}^3$. We remember this terminology by thinking of a triangle that contains its own flow as confident. If on top of its own flow it also contains the limit point, we think of it as overly confident or (self-)centered. We think of a triangle that is too weak to contain its flow as equivocal.

of $\sigma \in \operatorname{Del} S$, they converge towards a critical point y in the interior of $\sigma = \operatorname{conv} T$. This point is also contained in the interior of the corresponding ℓ-cell $\nu = \bigcap V_T$. It follows that ν and σ fall into Case (R1), which is illustrated by the leftmost column in Figure 5. The index of y is the dimension of σ.

FACT 5. A Delaunay simplex $\sigma \in \operatorname{Del} S$ with dual Voronoi cell ν is centered iff $\operatorname{int} \nu \cap \operatorname{int} \sigma \neq \emptyset$. The intersection is a point $y \in \mathbb{R}^3$, y is critical for all f_ε, and the index of y is $\dim \sigma$.

Since the limit curves do not exit its interior, σ has no successors in the flow relation. All predecessors are faces, but in general not all faces are predecessors of σ.

Confident simplices. A simplex $\tau \in D$ is *confident* if it is not centered and for every $x \in \operatorname{int} \tau$ there is a sufficiently small neighborhood N of x with $\lambda_x \cap N \subseteq \operatorname{int} \tau$. Illustrations of a confident edge and a confident triangle can be seen in Figures 6 and 7. Confident simplices are quite similar to centered ones, in the sense that they would be centered if they covered a large enough part of their affine hull. In other words, the affine hull of $\tau = \operatorname{conv} T$ intersects the interior of the corresponding ℓ-cell $\nu = \bigcap V_T$. Thus, ν and τ fall into Case (R2), which is illustrated by the second column from the left in Figure 5. All tetrahedra that are not centered are confident.

FACT 6. A Delaunay simplex $\tau \in \operatorname{Del} S$ with dual Voronoi cell ν is confident iff $\operatorname{int} \nu \cap \operatorname{int} \tau = \emptyset$ and $\operatorname{int} \nu \cap \operatorname{aff} \tau \neq \emptyset$.

All predecessors and successors of a confident τ are faces of τ. To determine which faces are successors, consider the center of the smallest orthosphere of τ. This is the point $z = \operatorname{int} \nu \cap \operatorname{aff} \tau$, which lies outside $\operatorname{int} \tau$. Let $k = \dim \tau$ and consider each $(k-1)$-dimensional face ξ of τ. By assumption of general position, the affine hull of ξ does not contain z. Either $\operatorname{aff} \xi$ separates z and $\operatorname{int} \tau$ within $\operatorname{aff} \tau$ or z and $\operatorname{int} \tau$ lie on the same side of $\operatorname{aff} \xi$. A proper face υ of τ is a successor iff the affine hulls of all $(k-1)$-dimensional faces ξ of τ that contain υ separate z and $\operatorname{int} \tau$. Observe that there is a unique lowest-dimensional successor υ. It has the property that the successors of τ are precisely the proper faces of τ that are cofaces of υ. Every other proper face of τ is either a predecessor of τ or neither a successor nor a predecessor.

Equivocal simplices. A simplex $\upsilon \in D$ is *equivocal* if $\lambda_x \cap \operatorname{int} \upsilon = x$ for every point $x \in \operatorname{int} \upsilon$. Illustrations of an equivocal edge and an equivocal triangle can be seen in Figures 6 and 7. The center z of the smallest empty sphere orthogonal to all $p \in T$, with $\upsilon = \operatorname{conv} T$, lies outside the affine hull of υ. In other words, the affine hull of υ misses the interior of the corresponding ℓ-cell, $\nu = \bigcap V_T$. Thus ν and υ fall into Cases (R3) or (R4), which are illustrated by the right two columns of Figure 5.

FACT 7. A Delaunay simplex $\upsilon \in \operatorname{Del} S$ with dual Voronoi cell ν is equivocal iff $\operatorname{int} \nu \cap \operatorname{aff} \upsilon = \emptyset$.

All predecessors and successors are cofaces of υ. For example, an equivocal triangle has exactly one predecessor, a tetrahedral coface, and exactly one successor, the other tetrahedral coface. The second coface can also be ω. All predecessors and successors of an equivocal simplex are confident or centered. Symmetrically, all predecessors and successors of confident and centered simplices are equivocal. An equivocal edge or vertex can have an arbitrary number of successors, but there is always only one predecessor. This fact is important and deserves a proof.

CLAIM 8. Every equivocal simplex has exactly one predecessor.

PROOF. Let υ be equivocal and $x \in \operatorname{int} \upsilon$. Consider the collection of limit curves λ_u that pass through x. As mentioned earlier, all λ_u approach x from the same direction although they possibly leave x in different directions. Since before x all λ_u are the same, we consider only λ_x and in particular a sufficiently small portion of λ_x immediately preceding x. This portion is a line segment contained in the interior of a simplex $\tau = \operatorname{conv} U$. We have $\tau \prec \upsilon$ and τ is confident. Every limit curve that intersects $\operatorname{int} \tau$ does so in a portion that lies on a line passing through the center z of the smallest empty sphere orthogonal to all $p \in U$. It follows that for every point $x \in \operatorname{int} \upsilon$, a sufficiently small portion of λ_x immediately preceding x lies in the affine hull of υ and z and therefore in $\operatorname{int} \tau$. In other words, each x identifies the same simplex τ, which implies that τ is the only predecessor of υ. ◻

The unique predecessor τ of the equivocal υ can be determined through local geometric considerations. Recall the definitions of z and $y = \operatorname{aff} \nu \cap \operatorname{aff} \upsilon$. By Fact

7, we have $y \not\in \nu$, and since $z \in \nu$ is closest to y, it lies on the boundary. Let $U \subseteq S$ be maximal with $z \in \bigcap V_U$. We have $T \subseteq U$ and $T \neq U$ because $z \in \mathrm{bd}\,\nu$. The predecessor of υ is $\tau = \mathrm{conv}\,U$.

6 Clustering Simplices

A *sink* is a simplex $\sigma \in D$ without successor in the flow relation. By construction, the sinks are exactly the centered simplices together with ω. We use $\prec$ to define for each sink σ a set of simplices gravitating towards σ. Think of σ analogous to a critical point and of this set analogous to a stable manifold.

Acyclicity. We show that the flow relation is acyclic. This is plausible since the value of the function g strictly increases along limit curves. A *cycle* is a sequence of simplices $\sigma_1 \prec \sigma_2 \prec \ldots \prec \sigma_\ell$, with $\ell \geq 3$ and $\sigma_1 = \sigma_\ell$. The algorithm in Section 8 relies on the absence of cycles in the flow relation.

CLAIM 9. The relation $\prec$ is acyclic.

PROOF. Let $\sigma_i = \mathrm{conv}\,T \in \mathrm{Del}\,S$, and let $\Sigma_i = (z_i, r_i)$ be the smallest empty sphere orthogonal to all $p \in T$. Consider $\sigma_i \prec \sigma_j$ and note that σ_i cannot be centered since otherwise it has no successor. If σ_i is confident then σ_j is equivocal and we have $\Sigma_i = \Sigma_j$ and $\dim \sigma_i > \dim \sigma_j$. If σ_i is equivocal then σ_j is centered or confident. Hence, $z_i \neq z_j$ and by assumption of non-degeneracy we have $r_i^2 < r_j^2$. To prove a cycle cannot exist, we assign to each σ_i the pair $(r_i^2, -\dim \sigma_i)$. The pairs increase lexicographically along a chain, which implies the chain cannot come back to where it started. ◻

Ancestor sets. The analogy between stable manifolds and ancestor sets of centered simplices is generally correct but troubled by inconsistencies in the details. The source of the trouble are simplices with more than one successor. Their existence implies the possibility of non-disjoint ancestor sets. This is in contrast to stable manifolds of a smooth Morse function, which are pairwise disjoint, but not unlike the closures of stable manifolds, which can overlap. There are two types of simplices which may have more than one successor:

(S1) equivocal edges and vertices,

(S2) confident tetrahedra and triangles.

Type (S1) simplices relate to the pinching or flattening of stable manifolds that occurs in the limit. Type (S2) simplices are a result of the non-continuous dependence of the stable manifolds from the input points. There are no type (S2) simplices in the two-dimensional unweighted case, where the limit of the complex of stable manifolds is the Gabriel graph [12]. Nonetheless, type (S2) simplices appear already in the two-dimensional weighted and the three-dimensional unweighted cases.

Despite the possibility of type (S1) and (S2) simplices, ancestor sets retain the containment relation of the closures of stable manifolds. Let y and z be critical points of a differentiable map f_ε, for a sufficiently small $\varepsilon > 0$, and recall that M_y and M_z denote their stable manifolds, as defined in Section 4. Let σ and τ be the centered simplices with $y \in \operatorname{int}\sigma$ and $z \in \operatorname{int}\tau$, and recall that $\operatorname{Anc}\sigma$ and $\operatorname{Anc}\tau$ are their ancestor sets.

CLAIM 10. $M_z \subseteq \operatorname{cl} M_y$ implies $\operatorname{Anc}\tau \subseteq \operatorname{Cl}\operatorname{Anc}\sigma$.

PROOF. Assume first that there are no type (S2) simplices. Then the dimensions of the confident simplices along a chain of the flow relation cannot decrease. If follows that $\dim\tau$ is the maximum dimension of any simplex in the ancestor set of τ, $\dim\operatorname{Cl}\operatorname{Anc}\tau = \dim\tau$. The dimension of τ is also the index of z. The claimed subset relation follows because limit curves are approximated by the flow curves as ε goes to 0.

If D contains type (S2) simplices, the dimension of $\operatorname{Cl}\operatorname{Anc}\tau$ may exceed the dimension of τ. The claimed subset relation still holds because the simplices $\upsilon \in \operatorname{Anc}\tau$ with $\dim\upsilon > \dim\tau$ have descendents outside $\operatorname{Anc}\tau$ and in particular σ is a descendent of υ. ◻

The face-coface relation over the set of stable manifolds is transitive and induces a partial order over the collection of centered simplices. This relation will be used in Section 8.

Definition of $\mathcal{X}$ and $\mathcal{W}$. Ancestor sets seem slightly too large to faithfully represent stable manifolds. We introduce a more conservative notion that admits only simplices whose cofaces have descendent sets contained in ancestor sets. Let $C \subseteq D$ be the set of sinks, including ω. For a subset $B \subseteq C$, define its ancestor set as the union of ancestor sets of its members, $\operatorname{Anc} B = \bigcup_{\sigma\in B}\operatorname{Anc}\sigma$. The *conservative ancestor set* of B is

$$\begin{aligned}\operatorname{Cnc} B &= \operatorname{Int}\{\tau \in D \mid \operatorname{Des}\tau \subseteq \operatorname{Anc} B\}\\ &= \{\tau \in D \mid \operatorname{Des}\sigma \subseteq \operatorname{Anc} B \text{ for all } \sigma \in \operatorname{St}\tau\}.\end{aligned}$$

Observe that the sets $\operatorname{Cnc}\sigma = \operatorname{Cnc}\{\sigma\}$, $\sigma \in C$, do not necessarily cover D. Indeed, a simplex is not covered by any set $\operatorname{Cnc}\sigma$ if it belongs to more than one set $\operatorname{Anc}\sigma$. On the other hand, $\operatorname{Cnc} C = D$ which shows that $\operatorname{Cnc} B$ is generally not equal to the union of conservative ancestor sets of its members. In this paper, we are only interested in

$$\begin{aligned}\Omega &= \operatorname{Cnc}\omega\\ &= \operatorname{Int}\{\tau \in D \mid \operatorname{Des}\tau \subseteq \operatorname{Anc}\omega\}\\ &= D - \operatorname{Cl}\bigcup\operatorname{Anc}\sigma\\ &= D - \bigcup\operatorname{Cl}\operatorname{Anc}\sigma,\end{aligned}$$

where the union is taken over all sinks $\sigma \neq \omega$. The wrapping surface, $\mathcal{W}$, is the boundary of

$$\begin{aligned} \mathcal{X} &= D - \Omega \\ &= \bigcup \operatorname{Cl} \operatorname{Anc} \sigma, \end{aligned}$$

where the union is again taken over all sinks $\sigma \neq \omega$. We have $\Omega = \operatorname{Int} \Omega$, by definition, and therefore $\mathcal{X} = \operatorname{Cl} \mathcal{X}$. In words, $\mathcal{X}$ is a simplicial complex. We will see shortly that $|\mathcal{X}|$ is contractible. In summary, Ω and $\mathcal{X}$ have the topological properties suggested by the analogy between Ω and the stable manifold of ω.

7 Collapsing Simplices

In this section, we show how to construct the complex $\mathcal{X}$ by collapsing simplices of the Delaunay complex. We refer to this algorithm as the *basic construction*.

Collapses. Consider a simplicial complex $\mathcal{K}$, and let υ be a simplex with exactly one proper coface $\tau \in \mathcal{K}$. In this case, $\dim \upsilon = \dim \tau - 1$ and τ is a principal simplex in $\mathcal{K}$. The operation that removes τ and υ from $\mathcal{K}$ is called an *elementary collapse*, and we write $\mathcal{K} \searrow \mathcal{K}_1$, where $\mathcal{K}_1 = \mathcal{K} - \{\tau, \upsilon\}$. An elementary collapse maintains the homotopy type of the complex.

FACT 11. $|\mathcal{K}_1|$ is homotopy equivalent to $|\mathcal{K}|$.

The homotopy equivalence can be established by constructing a deformation retraction of $|\mathcal{K}|$ to $|\mathcal{K}_1|$. This is a homotopy $F : |\mathcal{K}| \times [0, 1] \to |\mathcal{K}|$ between the identity map on $|\mathcal{K}|$ and a map from $|\mathcal{K}|$ to $|\mathcal{K}_1|$ that keeps all points $x \in |\mathcal{K}_1|$ fixed for all $t \in [0, 1]$. Such a homotopy is indicated in Figure 8.

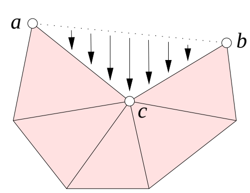

Figure 8: The elementary collapse removes the edge ab together with the triangle abc. The corresponding deformation retraction moves all points of abc parallel to the direction from the barycenter of ab to c.

An ℓ-simplex υ is *free* if there is a $k > \ell$ and a k-simplex $\tau \in \mathcal{K}$ such that all cofaces of υ are faces of τ. It follows that all cofaces of a free υ are free, except for τ, which is a principal simplex in $\mathcal{K}$. The operation that removes all cofaces of the free υ is called a (k, ℓ)*-collapse*. The number of simplices removed is $2m = 2^{k-\ell}$, and the (k, ℓ)-collapse can be written as a composition of m elementary collapses:

$$\mathcal{K} \searrow \mathcal{K}_1 \searrow \mathcal{K}_2 \searrow \ldots \searrow \mathcal{K}_m.$$

In $\mathbb{R}^3$, the (k, ℓ)-collapses satisfy $0 \leq \ell < k \leq 3$, and there are six different cases all illustrated in Figure 9. We use collapses to shrink a subcomplex $\mathcal{Y}$ of the Delaunay

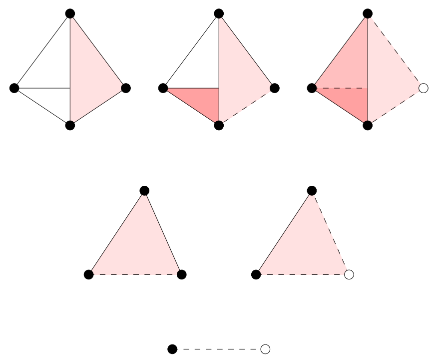

Figure 9: From left to right, top to bottom: collapsing a tetrahedron from a triangle, an edge, a vertex, collapsing a triangle from an edge, a vertex, and collapsing an edge from a vertex. In each case, the collapse removes the tetrahedron, the shaded triangles, the dashed edges, and the hollow vertices, if any.

complex. Call a pair $\upsilon \leq \tau$ *collapsible* if

(i) $\upsilon, \tau \in \mathcal{Y}$, υ is free, υ is equivocal, and

(ii) $\tau \prec \upsilon$, τ is the highest-dimensional coface of υ in $\mathcal{Y}$, υ is the lowest-dimensional successor of τ.

Observe that $\upsilon \in \mathrm{Bd}\,\mathcal{Y}$ because υ is free. Its coface τ may have several successors, all of which are free because they contain υ, which is free.

Correctness. A *constructive retraction* is an algorithm A that starts with the Delaunay complex and removes simplices by collapsing as long as there are collapsible pairs. Let $\mathcal{Y}_\mathsf{A}$ be the remaining subcomplex. We claim that every constructive retraction correctly constructs $\mathcal{X}$, no matter what sequence of collapses it chooses.

THEOREM. $\mathcal{Y}_\mathsf{A} = \mathcal{X}$ for every constructive retraction A.

PROOF. We prove $\mathcal{X} \subseteq \mathcal{Y} = \mathcal{Y}_\mathsf{A}$ by induction in the order of increasing descendent sets. Let $\xi \in D - \mathcal{Y}$. To show $\xi \in D - \mathcal{X}$, recall that $\xi \in \Omega = D - \mathcal{X}$ iff $\mathrm{Des}\,\eta \subseteq \mathrm{Anc}\,\omega$ for all cofaces η of ξ. Let $\upsilon \leq \tau$ be the pair whose collapse removes ξ from Del S, and note that $\upsilon \leq \xi \leq \tau$. We begin by proving $\mathrm{Des}\,\xi \subseteq \mathrm{Anc}\,\omega$. If $\xi \neq \tau$ then ξ is equivocal and all successors σ are cofaces that have already been removed. Then $\mathrm{Des}\,\sigma \subseteq \mathrm{Anc}\,\omega$ by induction and $\mathrm{Des}\,\xi \subseteq \mathrm{Anc}\,\omega$ follows. If $\xi = \tau$ then all successors σ are proper faces of τ and cofaces of υ. We just proved $\mathrm{Des}\,\sigma \subseteq \mathrm{Anc}\,\omega$ for all such σ and $\mathrm{Des}\,\xi \subseteq \mathrm{Anc}\,\omega$ again follows. Finally observe that every coface η of ξ has either already been removed or $\xi \leq \eta \leq \tau$. In both cases we have $\mathrm{Des}\,\eta \subseteq \mathrm{Anc}\,\omega$ and therefore $\xi \in \mathrm{Cnc}\,\omega = D - \mathcal{X}$.

We prove $\mathcal{Y} \subseteq \mathcal{X}$ by contradiction. Assume $\mathcal{Y} - \mathcal{X} \neq \emptyset$. For each simplex $\sigma = \operatorname{conv} T$, consider the pair $(r^2, -k)$, where $k = \dim \sigma$ and r is the radius of the smallest empty sphere orthogonal to all $p \in T$. By assumption of genericity, all pairs are different and, as shown in the proof of Claim 9, they lexicographically increase along chains in the flow relation. Let υ be the simplex in $\mathcal{Y} - \mathcal{X}$ with lexicographically largest pair. Since $\upsilon \in \Omega$, each successor σ of υ belongs to Ω. Furthermore, the pair of σ is lexicographically larger than that of υ. It follows that υ has no successor in $\mathcal{Y}$ and is therefore either a tetrahedron or equivocal. To contradict the first possibility, observe that for every tetrahedron υ with $\operatorname{Des} \upsilon \subseteq \operatorname{Anc} \omega$, there is an alternating sequence of tetrahedra and triangles,

$$\upsilon \prec \tau_0 \prec \upsilon_1 \prec \tau_1 \prec \ldots \prec \tau_j \prec \omega,$$

connecting υ to ω. So $\upsilon \notin \mathcal{Y}$ implies $\tau_0 \notin \mathcal{Y}$, contradicting the choice of υ. Second, consider the case in which υ is an equivocal simplex. Let τ be the predecessor of υ, which is unique and confident. The predecessor τ and its faces are the only cofaces of υ in $\mathcal{Y}$, since all others have larger r^2 value than υ. It follows that υ is free and $\upsilon \leq \tau$ is collapsible, which contradicts $\upsilon \in \mathcal{Y}$. ◻

The construction of $\mathcal{X}$ starts with $\operatorname{Del} S$, which is a contractible simplicial complex. The collapses maintain the homotopy type, which implies that $\| \mathcal{X} \|$ is indeed contractible, as claimed at the end of Section 6.

8 Deleting Simplices

A strength of the basic construction in Section 7 is that the wrapping surface is unique and its computation is fully automatic. A complementary weakness is the lack of variability in the result. This section generalizes the basic construction so the shape of the surface is influenced by the choice of additional parameters. The surface may wrap tighter around input points or develop holes and change its topology. We first describe a simplex removing operation that changes the homotopy type and then use this operation to modify the surface.

Discriminating by size. The idea is to collapse simplices not only from ω but more generally from all significant sinks. Recall that each sink, or centered simplex $\sigma = \operatorname{conv} T$, with $T \subseteq S$, corresponds to a critical point $y \in \operatorname{int} \sigma$. Call $|\sigma| = g(y)$ the *size* of σ, where g is the same as in Section 4. By definition of g, the sphere $\Sigma = (y, \sqrt{|\sigma|})$ is empty and the smallest sphere orthogonal to all $p \in T$. It is intuitively plausible that large size is indicative of space through which the wrapping surface may want to be pushed.

The non-degeneracy assumption on the input points implies that all sizes are different. Sort the sinks in order of decreasing size $|\sigma_0| > |\sigma_1| > \ldots > |\sigma_m|$, where

$\sigma_0 = \omega$ and $|\omega| = +\infty$. For each index $0 \le j \le m$, define

$$\begin{aligned}\mathcal{X}_j &= D - \mathrm{Cnc}\, B_j \\ &= \bigcup_{i=j+1}^{m} \mathrm{Cl}\,\mathrm{Anc}\,\sigma_i,\end{aligned}$$

where $B_j = \{\sigma_0, \sigma_1, \ldots, \sigma_j\}$ and $\mathrm{Cnc}\, B_j$ is the conservative ancestor set, as defined in Section 6. Define $\mathcal{W}_j = \mathrm{Bd}\,\mathcal{X}_j$. The $\mathcal{X}_j$ form a nested sequence of subcomplexes:

$$\mathcal{X} = \mathcal{X}_0 \supseteq \mathcal{X}_1 \supseteq \ldots \supseteq \mathcal{X}_m = \emptyset.$$

Correspondingly, the $\mathcal{W}_j$ form a nested sequence of wrapping surfaces.

An operation that removes a principal simplex σ from a complex $\mathcal{K}$ is called a *deletion*. In contrast to a collapse, a deletion alters the homotopy type of $|\mathcal{K}|$. A particular $\mathcal{X}_j$ is constructed from D by a succession of deletions and collapses. Each deletion is followed by collapses until no collapsible pairs remain. We refer to such an exhaustive sequence of collapses as a *retraction*. For example the basic construction computes $\mathcal{X} = \mathcal{X}_0$ from D by first deleting ω and then performing a retraction. The complex $\mathcal{X}_j$ is computed by repeating these two operations $j+1$ times, once each for $\sigma_0, \sigma_1, \ldots, \sigma_j$.

Local modifications. There is no reason other than convenience that requires a total order of the retractions. Indeed, it is possibly to perform retractions in any order consistent with the face-coface relation of stable manifolds. Recall that $C \subseteq D$ is the set of sinks, including ω. Let τ and σ be centered simplices and $z \in \mathrm{int}\,\tau$ and $y \in \mathrm{int}\,\sigma$ the corresponding critical points. The pair $\tau \vdash \sigma$ is in the *sink relation* $\vdash\ \subseteq C \times C$ if $M_z \subseteq \mathrm{cl}\, M_y$, as discussed in Claim 10. We call σ a *cosink* of τ and write $\mathrm{Cos}\,\tau$ for the set of cosinks, including τ. The relation is acyclic and transitive and therefore a partial order. It can be used to locally change or refine the wrapping surface. To describe how this may work, let B and B' be disjoint sets of centered simplices with $\mathcal{X} = D - \mathrm{Cnc}\, B$ and $B' \subseteq \mathcal{W} = \mathrm{Bd}\,\mathcal{X}$. We call the set of simplices in the conservative ancestor set of $B \cup B'$ that belong to $\mathcal{W}$ a *front*: $F = \mathcal{W} \cap \mathrm{Cnc}\,(B \cup B')$. Locality is understood in terms of F, that is, changes to the surface are triggered only from simplices in F. We exemplify this idea by setting a size threshold δ and removing simplices that have descendents σ with size exceeding δ. It suffices to consider sinks and we restrict attention to simplices τ in $F_\delta = \{\sigma \in F \cap C \mid |\sigma| > \delta\}$. To remove a simplex τ, we remove the entire conservative ancestor set of its cosinks. This is repeated for every $\tau \in F_\delta$. The local modification of $\mathcal{X}$ is completely specified by F and δ and creates $\mathcal{X}' = D - \mathrm{Cnc}\,(B \cup \mathrm{Cos}\, F_\delta)$. It is possible that $\mathcal{X}'$ contains a simplex $\upsilon \in F$ even though all centered descendents of υ and their cosinks have been removed. This is the case if υ has a coface ξ with at least one centered descendent remaining in $\mathcal{X}'$. The construction of $\mathcal{X}'$ is again reduced to a sequence of deletions and retractions: for each $\tau \in F_\delta$, we find all $\sigma \in \mathrm{Cos}\,\tau$ and repeatedly delete the σ without remaining cosinks. Each deletion is followed by a retraction. Similar to the basic construction, the sequence in which the $\tau \in F_\delta$ are removed is irrelevant, since any one results in the same $\mathcal{X}'$.

An interesting variant of the above described local deformation expands F_δ to include sinks υ that become part of $\mathcal{W}$ during the process. Only sinks υ with size $|\upsilon| > \delta$ are considered. This recursive construction amounts to replacing F_δ by G_δ, which is the maximal set of sinks $\upsilon \in \mathcal{X}$ with $|\upsilon| > \delta$ so that every component of $\mathrm{Cnc}\, G_\delta$ contains a simplex in F_δ.

9 Discussion

The solution to the surface reconstruction problem presented in this paper is based on discrete methods inspired by concepts in continuous mathematics. In particular, we construct subcomplexes of the Delaunay complex of a finite point set by collapsing and occasionally deleting simplices. Continuous mathematics enters in the form of Morse functions and their gradient fields, which constitute the rationale for the rules that decide when to collapse and when not. The remainder of this section briefly discusses possible extensions of the ideas in this paper and formulates open questions.

Adjusting granularity. The discrete version of the complex of stable manifolds can be interpreted as a coarse-grained view of the finer Delaunay complex. Each stable manifold is represented by a cluster of Delaunay simplices glued together by the flow relation. We can imagine an extension to a 1-parameter family of flow relations. The *granularity* parameter $\gamma \in \mathbb{R}$ controls the coarseness of the clustering. We aim at a parametrization with $\gamma = -\infty$, 0, and $+\infty$ corresponding to the Delaunay complex, the complex of stable manifolds, and $\{\mathbb{R}^3\}$. If $\gamma_1 < \gamma_2$, then the clustering for γ_1 should be a refinement of the clustering for γ_2. The present discussion conveniently ignores that ancestor sets representing stable manifolds can overlap and do not exactly partition the set of Delaunay simplices. Eventually, this set-theoretic inconvenience will have to be dealt with, possibly with concepts similar to conservative ancestor sets.

A mathematical formulation of granularity will have to be based on size and the comparison of sizes. Consider for example a centered triangle τ shared by tetrahedra σ_1 and σ_2. Then $|\tau| < |\sigma_1|$ and $|\tau| < |\sigma_2|$, and we suppose $|\sigma_1| < |\sigma_2|$. It is plausible to stipulate that for $\gamma > |\sigma_1| - |\tau|$, the triangle ought to change its behavior and act like an equivocal triangle with flow from smaller to larger size: $\sigma_1 \prec_\gamma \tau \prec_\gamma \sigma_2$. The intuition for this stipulation is to permit a limited amount of downward flow, namely from σ_1 to τ. The permitted amount is bounded from above by γ. The idea of limited downward flow can be generalized to simplices of all dimensions. For $\gamma > 0$, we cannot expect that the resulting cells are necessarily contractible. For negative γ, we get fewer pairs than in $\prec \; = \; \prec_0$ and therefore a finer partition of D than for $\gamma = 0$. Variants substituting $|\sigma|/|\tau|$ for $|\sigma| - |\tau|$ or cosinks for cofaces σ are conceivable.

Final remarks and open questions. An interesting variant of the basic construction maintains the wrapping surface $\mathcal{W} \subseteq \mathrm{Del}\, S$ while adding one point at a time to S. The Delaunay complex can be constructed in randomized time between constant and logarithmic per simplex [3, 7]. Is it possible to maintain $\mathcal{W}$ in about the same or possibly less time? More generally, we ask for an algorithm that maintains $\mathcal{W}$ through a sequence of point insertions, point deletions, point motions, and weight changes. An efficient such algorithm would be useful in conjunction with a fast algorithm for maintaining the Delaunay complex under such operations. The implementation of such an algorithm for three-dimensional Delaunay complexes is described in [8].

Given a viewing point $x \in \mathbb{R}^3$, a *depth-ordering* of the simplices in $\mathcal{W}$ is a linear extension of the in front-behind relation defined for x. Every Delaunay complex has a depth-ordering for every viewpoint in space [4], and since $\mathcal{W} \subseteq \mathrm{Del}\, S$, this is also true for $\mathcal{W}$. The depth-ordering opens up the possibility to use hidden-surface algorithms other than z-buffering to generate pleasing graphical representations.

The ideas presented in this paper can be extended to $\mathbb{R}^4$ and higher dimensions. It might be worthwhile to develop and implement such an extension, which could then be used in the analysis of point data beyond three dimensions. Such data is common in studies of time series, dynamical systems, and other areas of science and applied mathematics. The independence of the algorithm from assumptions about the data makes it an attractive approach to discovering structure in point data. The examples in Section 2 show that the algorithm has the ability to adapt to the dimension of the data, which is a useful feature in data exploration [2].

We conclude this paper with a question about the stability of the complex of stable manifolds. Small motion in the input data can cause critical points to appear or disappear. This causes non-continuous changes in the complex and possibly in the wrapping surface. It would be interesting to understand how exactly this lack of stability is related to the phenomenon of overlapping ancestor sets, or more precisely to the existence of type (S2) simplices, which are discussed in Section 6.

References

[1] J.-D. BOISSONNAT. Geometric structures for three-dimensional shape representation. *ACM Trans. Graphics* **3** (1984), 266–286.

[2] C. BREGLER AND S. M. OMOHUNDRO. Nonlinear manifold learning for visual speech recognition. *In* “Proc. 5th Internat. Conf. Comput. Vision, 1995”, 494-499.

[3] K. L. CLARKSON, K. MEHLHORN AND R. SEIDEL. Four results on randomized incremental constructions. *In* “Proc. 9th Ann. Sympos. Theoret. Aspects Comput. Sci., 1992”, 463–474, Lecture Notes in Comput. Sci. **577**, Springer-Verlag.

[4] H. EDELSBRUNNER. An acyclicity theorem for cell complexes in d dimensions. *Combinatorica* **10** (1990), 251–260.

[5] H. EDELSBRUNNER AND E. P. MÜCKE. Simulation of Simplicity: a technique to cope with degenerate cases in geometric algorithms. *ACM Trans. Graphics* **9** (1990), 66–104.

[6] H. EDELSBRUNNER AND E. P. MÜCKE. Three-dimensional alpha shapes. *ACM Trans. Graphics* **13** (1994), 43–72.

[7] H. EDELSBRUNNER AND N. R. SHAH. Incremental topological flipping works for regular triangulations. *In* “Proc. 8th Ann. Sympos. Comput. Geom., 1992”, 43–52.

[8] M. FACELLO. *Geometric Techniques for Molecular Shape Analysis.* Ph. D. Thesis, Dept. Comput. Sci., Univ. Illinois, Urbana, 1996.

[9] H. FUCHS, Z. KEDEM, AND S. P. USELTON. Optimal surface reconstruction from planar contours. *Comm. ACM* **20** (1977), 693–702.

[10] P. J. GIBLIN. *Graphs, Surfaces, and Homology.* Second edition, Chapman and Hall, London, 1981.

[11] H. HOPPE, T. DE ROSE, T. DUCHAMP, J. MCDONALD, AND W. STÜTZLE. Surface reconstruction from unorganized points. *Computer Graphics,* Proc. SIGGRAPH 1992, 71–78.

[12] D. W. MATULA AND R. R. SOKAL. Properties of Gabriel graphs relevant to geographic variation research and the clustering of points in the plane. *Geographic Analysis* **12** (1980), 205–222.

[13] D. MEYERS, S. SKINNER AND K. SLOAN. Surfaces from contours. *ACM Trans. Graphics* **11** (1992), 228–258.

[14] J. MILNOR. *Morse Theory.* Annals Math. Studies, Princeton Univ. Press, 1963.

[15] R. C. VELTKAMP. *Closed Object Boundaries from Scattered Points.* Springer-Verlag, Berlin, 1994.

[16] A. WALLACE. *Differential Topology. First Steps.* Benjamin, New York, 1968.

References

[A] N. AMENTA AND M. BERN. Surface reconstruction by Voronoi filtering. *Discrete Comput. Geom.* **22** (1999), 481–504.

[B] N. AMENTA, S. CHOI AND R. KOLLURI. The power crust, unions of balls, and the medial axis transform. *Comput. Geom. Theory Appl.* **19** (2001), 127–153.

[C] F. BERNARDINI, J. MITTLEMAN, H. RUSHMEIER, C. SILVA AND G. TAUBIN. The ball-pivoting algorithm for surface reconstruction. *IEEE Trans. Visual. Comput. Graphics* **5**(1999), 349–359.

[D] F. BERNARDINI AND C. L. BAJAJ. Sampling and reconstructing manifolds using α-shapes. *In* “Proc. 9th Canadian Conf. Comput. Geom., 1997”, 193–198.

[E] J.-D. BOISSONNAT AND F. CAZALS. Smooth surface reconstruction via natural neighbour interpolation of distance functions. *In* "Proc. 16th Ann. Sympos. Comput. Geom. 2000", 223–232.

[F] B. CURLESS AND M. LEVOY. A volumetric method for building complex models from range images. *Comput. Graphics,* Proc. SIGGRAPH 1996, 303–312.

[G] H. EDELSBRUNNER, D. LETSCHER AND A. ZOMORODIAN. Topological persistence and simplification. *Discrete Comput. Geom.*, to appear.

[H] R. FORMAN. Combinatorial differential topology and geometry. In *New Perspective in Geometric Combinatorics*, MSRI Publication **8**, 1999, 177–206.

[I] S. FUNKE AND E. A. RAMOS. Smooth surface reconstruction in near-linear time. *In* "Proc. 13th Ann. SIAM-ACM Sympos. Discrete Algorithms, 2002", 781–790.

[J] G. MEENAKSHISUNDARAM. *Theory and Practice of Sampling and Reconstruction for Manifolds with Boundaries.* Ph. D. Thesis, Dept. Comput. Sci., Univ. North Carolina, Chapel Hill, 2001.