Shape Reconstruction with Delaunay Complex

Herbert Edelsbrunner

Dept. Comput. Sci., Univ. Illinois at Urbana-Champaign, and
Raindrop Geomagic, Champaign, Illinois, USA

Abstract

The reconstruction of a shape or surface from a finite set of points is a practically significant and theoretically challenging problem. This paper presents a unified view of algorithmic solutions proposed in the computer science literature that are based on the Delaunay complex of the points.

Keywords. Computational geometry, combinatorial topology, morphology, Voronoi cells, Delaunay complex, α-shape, crust, $\mathcal{A}$-shape, wrap complex, neural network, surface triangulation, normalized mesh.

1 Introduction

This paper considers the problem of reconstructing a shape from a given finite set of points. Solutions based on the Delaunay complex of the set are surveyed and a unified view using restricted Delaunay complexes is developed

Problem description. Define a *shape* as a subset of Euclidean space. It inherits the topology of that space and can be viewed as a topological space itself. Given a finite set of points, $S \subseteq \mathbb{R}^d$, the *shape reconstruction problem* asks for a shape in $\mathbb{R}^d$ that best approximates S.

Without quantifying what it means that a shape approximates a point set, shape reconstruction remains a vague and primarily morphological problem. In spite of the importance of the problem there has been little success in phrasing it as an optimization problem. The main source of the difficulty is the fantastically rich variety of shape and form as apparent in nature around us [16, 39]. One way to cope with this difficulty is to limit the range of shapes. We call the result a *narrow* problem specification. An example is the requirement that the shape produced be a closed surface in $\mathbb{R}^3$, assuming $S \subseteq \mathbb{R}^3$. The trouble with this problem specification is that many sets $S \subseteq \mathbb{R}^3$ do not admit any reasonably approximation by a closed surface. As a consequence, any algorithm is limited to a subclass of input sets S, and it is difficult to characterize this subclass other than through success and failure of the algorithm. A *wide* problem specification permits the construction of any topological space. For example in $\mathbb{R}^3$, the shape can be a point, a curve, a surface, a solid, or any combination of these. The trouble with this specification is that an algorithm might not produce a closed surface even in cases where one approximating S would exist and the application would prefer one.

Ramifications. Versions of the shape reconstruction problem can be found in diverse areas of science and engineering. 2-dimensional versions arise in pattern recognition, image processing and computer vision [36]. For example, solutions to boundary reconstruction from images based on Delaunay complexes have been studied by Brandt and Algazi [11] and by Robinson et al. [38]. The most common 3-dimensional version of the problem is narrow and requires the reconstructed shape be a closed surface, see Lodha and Franke [31] for a survey of scattered point techniques for surfaces. The importance of this case stems from the fact that the closed surfaces are exactly the boundaries of the solids (3-manifolds with boundary). Given a closed surface, the solid can be physically created by modern 3D printing technology surveyed by Burns [13]. A particular algorithmic solution to the wide version of the 3-dimensional shape reconstruction problem are the α-shapes proposed by Edelsbrunner and Mücke [21]. They are dual to the space filling models of molecules and found extensive applications in molecular biology. The surface reconstruction problem in $\mathbb{R}^3$ has been generalized to the manifold learning problem in dimensions beyond 3 by Bregler and Omohundro [12]. This problem arises in the analysis of dynamical systems and of physical phenomena described by data of a fixed dimension greater than 3.

Reconstruction methods. This paper focuses on the *Delaunay approach* that reconstructs a shape from the Delaunay complex of the points. This is a simplicial complex that decomposes the convex hull of S by connecting the points with simplices of all possible dimensions. The shape is the underlying space of a subcomplex chosen from the Delaunay complex by some algorithm. Other approaches to the shape reconstruction problem are beyond the scope of this paper.

Most solutions described in the scientific literature follow either the Delaunay or one of three other approaches. The first other approach keeps the idea of taking subcomplexes and replaces the Delaunay by a different complex. Kirkpatrick and Radke [27] suggest the one-parameter family of β-skeletons to reconstruct the shape of a finite set in $\mathbb{R}^2$. The sphere-of-influence graph has been proposed by Toussaint, see [6], and used by Edelsbrunner, Rote and Welzl [22] to find shortest curves. The β-skeletons have been generalized by Veltkamp [41] and used for the reconstruction of surfaces in $\mathbb{R}^3$.

The second approach reconstructs surfaces in $\mathbb{R}^3$ from slices. Each slice represents the intersection of the shape with a plane by a collection of polygons in that plane. It is usually assumed that the slices are defined by a sequence of parallel planes. If two adjacent slices consist of a single polygon each then the reconstruction problem reduces to finding a cylindrical surface that connects the two polygons. Fuchs, Kedem and Uselton [25] describe a polynomial time algorithm for constructing minimum area and other optimal cylinders. In the general case each slice consists of several pairwise disjoint but possibly nested polygons. Solutions that first match and second connect the polygons are surveyed by Meyers, Skinner and Sloan [35]. Boissonnat and Geiger [10] combine the matching and connecting into one step using the Delaunay complex of the two slices.

The third approach takes advantage of the fact that the points in S used to specify the shape have otherwise no significance. S can be replaced by any set that leads to the same or a similar shape. Assuming a dense distribution along the hypothetical surface, Hoppe et al. [26] construct a signed distance function, $f : \mathbb{R}^3 \rightarrow \mathbb{R}$. The surface is defined as the zero-set, $f^{-1}(0)$, and constructed by the marching cube algorithm [32]. The limitation to sets S that are dense everywhere along the surface has been partially overcome by Curless and Levoy [15] who assume S is obtained by a scanner that provides, for each point, also the ray meeting the surface at that point.

Outline. Section 2 introduces the restricted Delaunay complex, which is the central notion in the unified view of solutions following the Delaunay approach. Section 3 discusses versions of the Delaunay approach that construct the restricting space from the data set. Section 4 considers version that assume the restricting space is given, either implicitly or explicitly. Section 5 concludes this paper.

2 Restricted Delaunay Complex

This section presents the definitions needed to unify and classify the proposed solutions following the Delaunay approach to shape reconstruction. The main concept is the so-called restricted Delaunay complex first defined in full generality by Edelsbrunner and Shah [23].

Voronoi cells. A finite set $S \subseteq \mathbb{R}^d$ induces a decomposition of the space into regions of influence. Specifically, let $\|x - p\|$ be the Euclidean distance between points $x, p \in \mathbb{R}^d$. The *Voronoi cell* of $p \in S$ is the set of points x whose distance from p is less than or equal to the distance from any other point in S:

$$V_p = \{x \in \mathbb{R}^d \mid \|x - p\| \leq \|x - q\|, q \in S\}.$$

Each Voronoi cell is a closed and possibly unbounded convex polyhedron, see Figure 1. The Voronoi cells meet

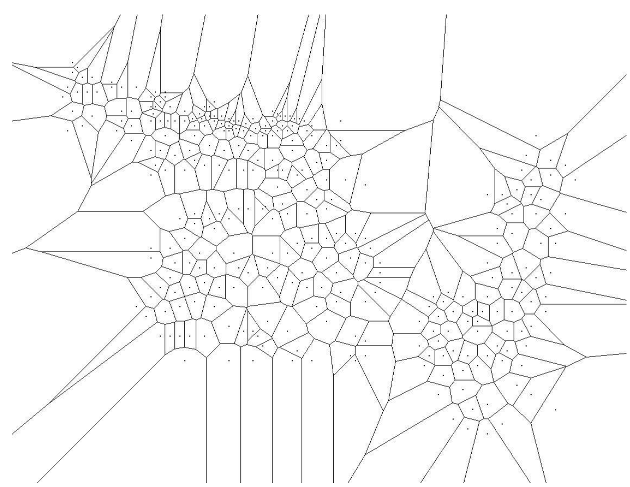

Figure 1: Decomposition of the plane by Voronoi cells of a finite set. The points are the locations of trees in the Allerton Park near Monticello, Illinois.

at most along common boundary faces, and together they cover the entire $\mathbb{R}^d$. The collection of Voronoi cells is denoted as $V_S = \{V_p \mid p \in S\}$.

In this paper we only consider Voronoi cells defined by finitely many (unweighted) points and the Euclidean metric. Refer to the survey article by Aurenhammer [5] for generalizations to points with weights, to infinite point sets, and to other metrics.

Nerve. The *nerve* of a finite collection of sets, A, is the set system (or set of collections) consisting of all subcollections with non-empty common intersection:

$$\mathrm{Nrv}\, A \quad = \quad \{X \subseteq A \mid \bigcap X \neq \emptyset\}.$$

It has been introduced by Alexandrov [1] as a tool to construct simplicial complexes. Observe that $Y \subseteq X$ and $X \in \mathrm{Nrv}\, A$ implies $Y \in \mathrm{Nrv}\, A$; this is the defining property of an *abstract simplicial complex.* To obtain an embedding we represent each set in A by a point in the Euclidean space of some dimension, e, and each collection $X \in \mathrm{Nrv}\, A$ by the convex hull of the corresponding points. Specifically, we find an injective function $\varphi : A \to \mathbb{R}^e$ so that

$$\mathrm{conv}\, \varphi(X) \cap \mathrm{conv}\, \varphi(Z) \quad = \quad \mathrm{conv}\, \varphi(X \cap Z)$$

holds for all $X, Z \in \mathrm{Nrv}\, A$. In words, $\varphi(X)$, $\varphi(Z)$, and $\varphi(X \cap Z)$ are finite sets of points, their convex hulls are simplices, and the intersection of the first two simplices is the simplex spanned by the intersection of the first two point sets. $\mathcal{K} = \{\mathrm{conv}\, \varphi(X) \mid X \in \mathrm{Nrv}\, A\}$ is a *simplicial complex* and $\mathcal{K}$ together with φ is a *geometric realization* of $\mathrm{Nrv}\, A$. The *underlying space* of $\mathcal{K}$ is the part of $\mathbb{R}^e$ covered by its simplices: $|\mathcal{K}| = \bigcup \mathcal{K}$.

A general position argument shows that there is always a geometric realization in dimension $e \geq 2m - 1$, where m is the maximum cardinality of any $X \in \mathrm{Nrv}\, A$, and there are examples that show $e = 2m - 1$ is sometimes necessary [24, 40]. For computational purposes it is important to keep e as small as possible, and for $A = V_S$ it turns out that $e = m - 1 = d$ is sufficient.

Delaunay complex. Recall that S is a finite set of points in $\mathbb{R}^d$ and V_S is the collection of Voronoi cells. We assume general position so that the common intersection of any k Voronoi cells is either empty or a convex polyhedron of dimension $d + 1 - k$. It follows that the collections X in the nerve of V_S have cardinality at most $d+1$. The *Delaunay complex* of S is the geometric realization of $\mathrm{Nrv}\, V_S$ defined by the injection $\varphi : V_S \to \mathbb{R}^d$ that maps every Voronoi cell to its generator, $\varphi(V_p) = p$:

$$\mathrm{Del}\, S \quad = \quad \{\mathrm{conv}\, \varphi(X) \mid X \in \mathrm{Nrv}\, V_S\},$$

see Figure 2. In other words, if two Voronoi cells share a common $(d-1)$-face then their generating points are connected by an edge, if three cells share a common $(d-2)$-face then their generators are connected by a triangle, etc.

General position is a convenient bua nyot a necessary assumption. Without this assumption we get cells that are not simplices. Specifically, the convex hull of k points is a cell in the Delaunay complex iff the corresponding k Voronoi cells have a non-empty common intersection not contained in any other Voronoi cell. In this paper we assume general position, which can be simulated computationally by a symbolic perturbation [20].

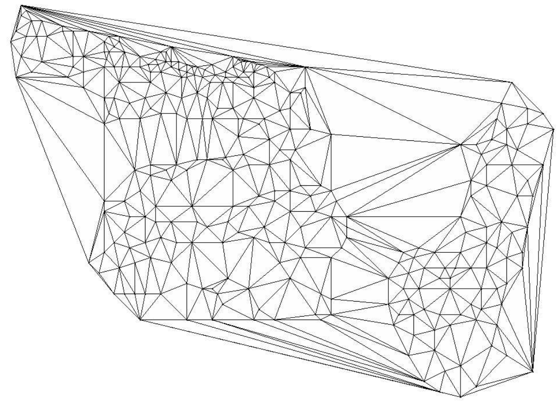

Figure 2: Delaunay complex corresponding to the decomposition of the plane into Voronoi cells shown in Figure 1.

Restricting Voronoi cells and Delaunay complex. Just as the Voronoi cells decompose $\mathbb{R}^d$, they decompose any topological subspace $\mathbb{X} \subseteq \mathbb{R}^d$. We call $V_p \cap \mathbb{X}$ the *restricted* Voronoi cell of p and consider the collection of all such cells: $V_{S,\mathbb{X}} = \{V_p \cap \mathbb{X} \mid p \in S\}$. The *restricted Delaunay complex* is the geometric realization in $\mathbb{R}^d$ of the nerve of the collection of restricted Voronoi cells:

$$\mathrm{Del}_{\mathbb{X}} S \quad = \quad \{\mathrm{conv}\, \varphi(Y) \mid Y \in \mathrm{Nrv}\, V_{S,\mathbb{X}}\},$$

see Figure 3 in the section on alpha shapes. Observe that the restricting space specifies a subcomplex of the Delaunay complex: $\mathrm{Del}_{\mathbb{X}} S \subseteq \mathrm{Del}\, S$.

The nerve theorem of Leray [29] implies that if all restricted Voronoi cells are contractible then $\mathbb{X}$ and the underlying space of the restricted Delaunay complex, $|\mathrm{Del}_{\mathbb{X}} S|$, are homotopy equivalent. This means that the two topological spaces are connected the same way: they can be geometrically different but they have the same kind and arrangement of holes. Edelsbrunner and Shah [23] prove that if $\mathbb{X}$ is a k-manifold with boundary then $\mathbb{X}$ and $|\mathrm{Del}_{\mathbb{X}} S|$ are homeomorphic if the restricted Voronoi cells satisfy the closed ball property:

(i) the common intersection of $\mathbb{X}$ and any $k + 1 - \ell$ Voronoi cells is either empty or a closed ℓ-ball, and

(ii) the common intersection of the boundary of $\mathbb{X}$ and any $k+1-\ell$ Voronoi cells is either empty or a closed $(\ell - 1)$-ball.

The closed ball property generalizes to a sufficient condition that implies homeomorphic reconstruction for general triangulable spaces $\mathbb{X}$.

3 Constructing the Restricting Space

All algorithmic solutions to shape reconstruction surveyed in this paper use restricted Delaunay complexes and only differ in how they arrive at the restricting space and how they treat it computationally. This section discusses solutions that generate the restricting space from the given data points.

Alpha shapes. In 1983, Edelsbrunner, Kirkpatrick and Seidel introduced the *α-shape* of a set $S \subseteq \mathbb{R}^2$ as the space generated by connecting point pairs that can be touched by an empty disk of radius α [19]. Specifically, points $p, q \in S$ are connected by a straight edge if there is a circle of radius α that passes through p and q, and all other points of S lie strictly outside the circle. The collection of edges decomposes $\mathbb{R}^2$ into *interior* regions that belong to the α-shape and *exterior* regions that constitute the background. The unbounded region is always exterior. An equivalent definition restricts the Delaunay complex of S with open disks of radius α centered at the points:

$$\mathbb{X} = \{x \in \mathbb{R}^2 \mid \|x - p\| < \alpha \text{ for some } p \in S\}.$$

The α-shape is the underlying space of $\mathcal{K}_\alpha = \mathrm{Del}_{\mathbb{X}} S$, see Figure 3. The *$\alpha$-complex*, $\mathcal{K}_\alpha$, triangulates the interior regions and thus clarifies the distinction between interior and exterior.

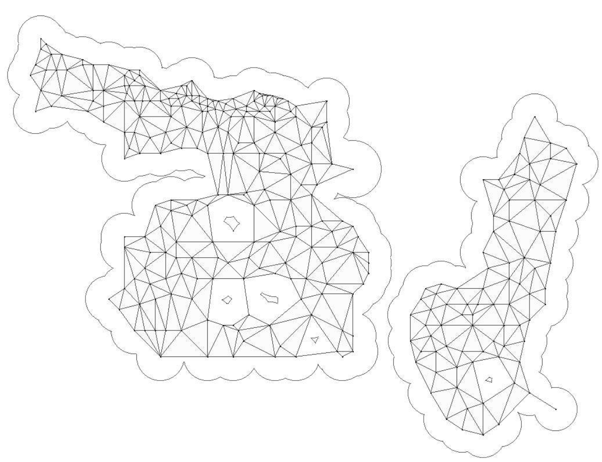

Figure 3: Each Voronoi cell in Figure 1 is restricted to within the open disk of radius α centered at the generating point. The result is a subcomplex of the Delaunay complex in Figure 2 that represents the shape at the resolution determined by α.

Each restricted Voronoi cell is the intersection of the original Voronoi cell with the disk of its generating point. Since the cell and the disk are both convex, the restricted cell is convex and therefore contractible. The nerve theorem implies that $\mathbb{X}$ and the α-shape are homotopy equivalent. This fact is reflected in Figure 3 where both spaces consist of 2 components, one with 4 holes and the other with 1 hole. The two spaces are not homeomorphic since $\mathcal{K}_\alpha$ contains 3 isolated edges, which under any retraction are the pinched images of locally 2-dimensional regions in $\mathbb{X}$. Each one of these edges violates condition (ii) of the closed ball property stated in section 2. The fact that restricted Voronoi cells are open, rather than closed as required by the closed ball property, is a minor difficulty that can be remedied by taking slightly smaller closed cells. Bernardini and Bajaj [8] reconstruct 2-dimensional shapes, and they prove that under some density requirements α-shapes correctly reconstruct uniformly sampled smooth curves.

Alpha shapes generalize to $\mathbb{R}^d$ by using open balls in the definition of the restricting space. Edelsbrunner and Mücke [21] discuss this construction in $\mathbb{R}^3$ and Bajaj, Bernardini and Xu [7] use it to reconstruct shapes and surfaces. The 3-dimensional case has also applications to computational biology where molecules are modeled as unions of spherical balls. Such models have been introduced by Lee and Richards [28] in 1971 and are commonly used to assess spatial properties of molecules such as volume, surface area, connectivity, shape, etc. Edelsbrunner [17] generalizes alpha shapes to points with weights in order to model molecules made up of atoms of varying size. That paper also contains inclusion-exclusion formulas that compute the volume and surface area of a molecule directly from the α-complex, without constructing the union of balls. These formulas have been used to measure molecules and their voids and pockets by Liang and collaborators, see e.g. [30].

Crust. In 1997, Amenta, Bern and Eppstein defined the *crust* of a set $S \subseteq \mathbb{R}^2$ as the subcomplex of $\mathrm{Del}\,(S \cup U)$ induced by S, where U is the set of vertices of the Voronoi cells defined by S [3]. In other words, a simplex in $\mathrm{Del}\,(S \cup U)$ belongs to the crust if all its vertices are points in S, see Figure 4. To reformulate the definition we consider the collection of Voronoi cells of $S \cup U$ and use the subset of cells generated by points in S to restrict the Delaunay complex of S:

$$\mathbb{X} = \mathrm{int}\,(\textstyle\bigcup_{p \in S} V_p).$$

Note that $\mathbb{X}$ is the set of points closer to S than to U, or equivalently it is the complement of the union of the Voronoi cells generated by points in U. The crust is the resulting restricted Delaunay complex: $\mathcal{C} = \mathrm{Del}_{\mathbb{X}} S$. No triangle in $\mathrm{Del}\,S$ can be in $\mathcal{C}$ because the corresponding intersection of the three Voronoi cells is a point in U, which necessarily lies outside $\mathbb{X}$.

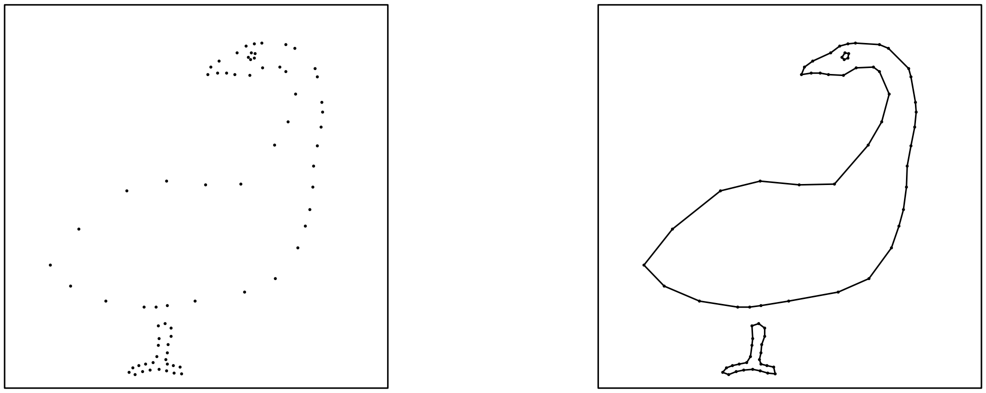

Figure 4: Point set and crust are courtesy of Nina Amenta at the University of Texas in Austin. The crust reconstructs the goose from a collection of points sampled from the outline.

The crust is suitable for the reconstruction of smooth boundary curves in the plane. The main result in [3] is a fairly modest condition on the sampling density under which the crust is guaranteed to reconstruct a smooth closed curve, γ. Define the *medial axis* of γ as the set of points $y \in \mathbb{R}^2$ with two or more closest points on γ, and for a point x let $f(x)$ be the distance to the medial axis. A finite set $S \subseteq \gamma$ is an *ε-sample* if every point $x \in \gamma$ is within distance $\varepsilon \cdot f(x)$ of some point in S. For $\varepsilon < 0.252$ the crust is guaranteed to contain an edge connecting points $p, q \in S$ iff they are contiguous along γ. This result justifies the definition of crust by the observation that the points in U approximate the medial axis of γ.

It is straightforward to extend the definition of crust to 3 and higher dimensions. However, already in $\mathbb{R}^3$ the Voronoi vertices of points sampled on a smooth surface no longer approximate the medial axis of that surface. The source of the trouble are slivers, which are Delaunay tetrahedra whose 4 vertices are almost cocircular. The center of the circumsphere belongs to U which implies that the sliver does not belong to the crust, but in many cases neither do the 4 triangles of the sliver. As a consequence, the crust develops holes or windows in the reconstructed surface. Amenta and Bern [2] cope with this difficulty by using only a subset of the points in U for the restricting space.

$\mathcal{A}$-shape. In 1997, Melkemi proposed a general family of shapes that includes α-shapes and the crust as special cases [34]. Let S be a finite set in $\mathbb{R}^2$. A member in this family is identified with the help of a second finite set $\mathcal{A} \subseteq \mathbb{R}^2$. The *$\mathcal{A}$-shape* of S is generated by drawing an edge connecting points $p, q \in S$ if there is a circle that passes through p, q, and a point $a \in \mathcal{A}$, and all other points of $S \cup \mathcal{A}$ lie strictly outside the circle. The crust is the special case where $\mathcal{A} = U$ is the collection of Voronoi vertices defined by S. The α-shape is the special case where $\mathcal{A}$ is the collection of points a on Voronoi edges that span empty circles of radius α with points in S. The reformulation of the definition is similar to the crust. The restricting space is the set of points closer to S than to $\mathcal{A}$:

$$\mathbb{X} \;=\; \operatorname{int}\left(\bigcup_{p \in S} V_p\right),$$

and the *$\mathcal{A}$-shape* is the boundary of the underlying space of the Delaunay complex restricted by $\mathbb{X}$.

The trouble with this definition is the high degree of freedom. $\mathcal{A}$ can be anything and it is not clear how to construct sets that bring out the shape of S best. To address this concern, Melkemi suggests a two-parameter family of point sets, $\mathcal{A} = \mathcal{A}(\alpha, t)$. The first parameter, $\alpha \geq 0$, controls the resolution and the second parameter, $t \in [0, 1]$, interpolates between the unweighted case and the case where points are weighted by the local density. To be specific, let $\delta(p)$ be the minimum Euclidean distance between $p \in S$ and any other point in S. For a given t, the *weighted distance* of a point $x \in \mathbb{R}^2$ from $p \in S$ is

$$\pi_{t,p}(x) \;=\; \|x - p\|^2 - t^2 \cdot \delta(p)^2,$$

which is the square length of a tangent line segment from x to the circle with center p and radius $t \cdot \delta(p)$. A point $a \in \mathbb{R}^2$ belongs to $\mathcal{A}(\alpha, t)$ if there are points $p, q \in S$ with $\alpha^2 = \pi_{t,p}(a) = \pi_{t,q}(a) < \pi_{t,r}(a)$ for all $r \in S - \{p, q\}$. The family $\mathcal{A}(\alpha, t)$ does in general not contain the particular sets that generate the α-shapes, but it would be easy to define a similar family that does.

Wrap complex. Commercial software produced at Raindrop Geomagic reconstructs the shape of a finite set $S \subseteq \mathbb{R}^3$ through an iteration that refines both the shape and the restricting space [37]. Let F map $\mathbb{X} \subseteq \mathbb{R}^d$ to the Delaunay complex restricted by $\mathbb{X}$, and let G map

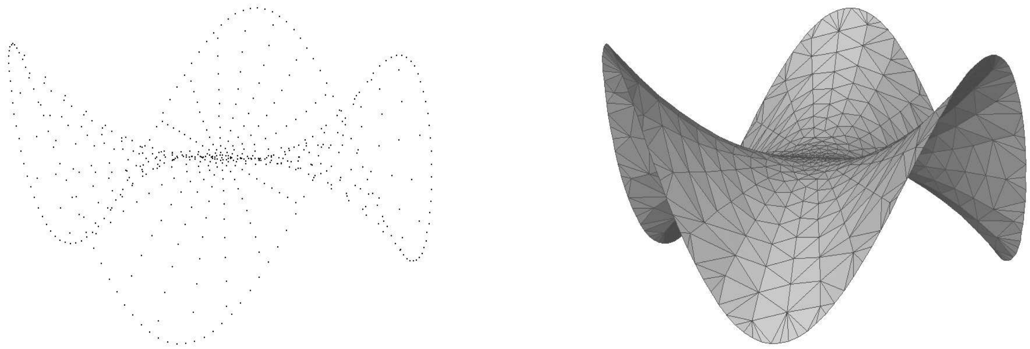

Figure 5: Wrap complex of a collection of points sampled on the surface of a monkey saddle.

a subcomplex $\mathcal{K} \subseteq \mathrm{Del}\, S$ to a topological subspace of $\mathbb{R}^d$. The composition maps a subcomplex of $\mathrm{Del}\, S$ to another such subcomplex, and we write $\mathcal{K} \preceq \mathcal{L}$ if $\mathcal{L} = F(G(\mathcal{K}))$. Edelsbrunner constructs G so the relation is acyclic [18]. It follows that the maximal elements in the relation are fixed points of $F \circ G$. Another property implied by the special choice of G is that the union of two fixed points is again a fixed point. This implies there exists a unique largest fixed point, which we call the *Wrap complex* of S, see Figure 5. It can be obtained by iterating $F \circ G$ starting with $\mathrm{Del}\, S = F(\mathbb{R}^d)$:

$$\mathrm{Del}\, S = \mathcal{X}_0, \mathcal{X}_1, \ldots, \mathcal{X}_j = \mathcal{X}_{j+1},$$

with $\mathcal{X}_{i+1} = F(G(\mathcal{X}_i))$. The software developed at Raindrop Geomagic offers the user convenient control that permits the transition between different fixed points, each representing a locally reasonable approximation of S.

4 Assuming the Restricting Space

This section discusses variants of the Delaunay approach to shape reconstruction that assume the restricting space is given, either implicitly or explicitly. The first two solutions make use of an oracle that returns a small bit of information about the restricting space. The third solution aims at reconstructing the restricted Delaunay complex without any information about the restricting space other than the sampled data points.

Neural network. In 1994, Martinetz and Schulten designed what they initially called the neural gas algorithm [33]. It constructs a neural network modeled as a 1-dimensional complex of nodes and edges that approximates a target space $\mathbb{X} \subseteq \mathbb{R}^d$. The approximation is achieved by sampling points from $\mathbb{X}$ and using them to locally adjust the nodes and connect them with edges. For example, if $\mathbb{X}$ is the state space of a dynamical system then each sample is a snapshot of that system during its evolution. The algorithm starts with a loose collection of nodes or points distributed in $\mathbb{R}^d$, see Figure 6. For each sample $x \in \mathbb{X}$, the positions of the nearby nodes are adjusted towards x and the *age* of every edge is increased by one. Furthermore, if the two nodes p and q closest to x are already connected by an edge, pq, then x is interpreted as further justification of that edge and its age is set back to 0. If pq is not yet in the network then it is now added with age 0. Edges whose age exceeds a certain threshold are removed from the network.

The connection to restricted Delaunay complexes arises from the fact that a sample point x with closest nodes p and q can be interpreted as evidence that the Voronoi cells of p and q share a common $(d-1)$-dimensional face that has a non-empty intersection with $\mathbb{X}$. It is therefore reasonable to expect that the neural gas algorithm approaches the 1-skeleton of the Delaunay complex of the nodes restricted by $\mathbb{X}$. The simple strategy of connecting the $k+1=2$ closest nodes does not extend to k-simplices for $k > 1$ because already the 3 closest nodes do not, in general, span a triangle in the Delaunay complex.

Surface triangulation. In 1993, Chew defined the *Delaunay triangulation* of a finite set S of points on a surface $\mathbb{X}$ in $\mathbb{R}^3$ by modifying the Euclidean empty disk criterion [14]. Specifically, the triangle formed by points $p, q, r \in S$ belongs to the surface Delaunay triangulation if there is a sphere K with center on $\mathbb{X}$ so that p, q, r lie on and all other points lie strictly outside K. Chew uses this definition in combination with a point placement mechanism to produce surface triangulations with good quality triangles, see Figure 7.

The connection to restricted Delaunay complexes should be obvious: the sphere K exists iff the 3-dimensional Voronoi cells of p, q, and r meet along an

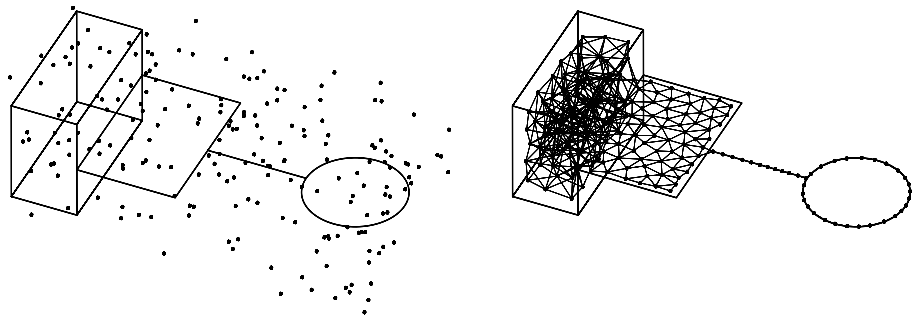

Figure 6: The pictures are courtesy of Klaus Schulten from the University of Illinois at Urbana-Champaign. The target space consists of a 3-dimensional box, a 2-dimensional rectangle and a ring with line segment. The left shows the starting configuration and the right shows the ending configurations after 40,000 steps of the neural gas algorithm.

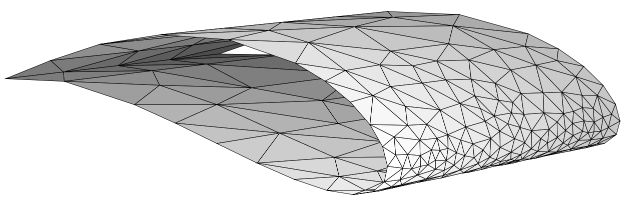

Figure 7: The triangulation is courtesy of Paul Chew at Cornell University. The surface is part of an airplane wing and it is triangulated using the modified empty disk criterion.

edge that has a non-empty common intersection with $\mathbb{X}$. In other words, Chew's surface Delaunay triangulation is the same as the 3-dimensional Delaunay complex restricted by that surface. This suggest that the closed ball property of section 2 be used as part of the point placement mechanism to guarantee the restricted complex is homeomorphic to the surface.

Normalized mesh. In 1997, Attali defined the *normalized mesh* of a finite set $S \subseteq \mathbb{X}$ as a complex that approximates the space $\mathbb{X}$ [4]. It contains the convex hull of a subset $T \subseteq S$ as a cell if there is a point $x \in \mathbb{X}$ equally far from all points in T and further from all other points: $\|x-p\| = \|x-q\| < \|x-r\|$ for all $p, q \in T$ and $r \in S-T$. Observe that for $\mathbb{X}$ a surface in $\mathbb{R}^3$ this is the same as Chew's surface Delaunay triangulation. In general, the normalized mesh is the same as the Delaunay complex restricted by $\mathbb{X}$.

The algorithmic problem tackled in [4] is the reconstruction of $\mathbb{X}$ through the construction of the normalized mesh. The problem is made difficult by assuming that $\mathbb{X}$ is not given at all, other than indirectly through the points in S. In two dimensions the strategy is to discriminate edges $pq \in \mathrm{Del}\, S$ by $\delta(pq)$ defined as the sum of angles opposite to pq in the two incident triangles. If only one triangle exists, the other opposite angle is set to 0. Under some smoothness conditions it is possible to prove that the normalized mesh consists exactly of all edges with small value of δ. Call $\mathbb{Y} \subseteq \mathbb{R}^2$ an *R-regular shape* if the circle passing through any three boundary points has radius greater than R. We assume that R is positive and $\mathbb{X}$ is the boundary of $\mathbb{Y}$. This implies that the curvature at any point $x \in \mathbb{X}$ is smaller than $\frac{1}{R}$. A finite set $S \subseteq \mathbb{X}$ is an *ε-sample* if every point $x \in \mathbb{X}$ is within distance $\varepsilon \cdot R$ of some point in S. The main result in [4] states that if $\varepsilon < \sin \frac{\pi}{8}$ then the normalized mesh consists exactly of all edges $pq \in \mathrm{Del}\, S$ with $\delta(pq) < \varepsilon \cdot R$. For similar reasons as mentioned in the discussion of the crust, this result does not generalize to $\mathbb{R}^3$ and Attali presents heuristics that patch up the holes in the partially reconstructed surfaces.

5 Conclusions

This paper unifies algorithmic solutions to the shape reconstruction problem following the Delaunay approach by identifying a common underlying concept: the restriction of the Delaunay complex by a topological space. The unification succeeds in all cases known to the author at this time, except for the work of Boissonnat [9] who suggests to sculpt a shape from the Delaunay complex by removing simplices from outside in. This idea is closely related to the crust, the wrap complex, and the normalized mesh, but in its general form it is not guided by any restricting space.

The seven solutions to shape reconstruction surveyed in this paper are classified according to their treatment of the restricting space. The four methods in section 3 *construct* the space from the data points, while the three methods in section 4 *assume* the space is given and cannot be altered by the algorithm. Another useful classification criterion for shape reconstruction is the one-dimensional scale from narrow to wide. The surface reconstruction methods ought to be classified as narrow, and examples are the crust, the surface Delaunay triangulation, and the normalized mesh. The other four methods show no bias for any particular type of shape and ought to be classified as wide.

References

[1] P. S. ALEXANDROV. Über den allgemeinen Dimensionsbegriff und seine Beziehungen zur elementaren geometrischen Anschauung. *Math. Ann.* **98** (1928), 617–635.

[2] N. AMENTA AND M. BERN. Surface reconstruction by Voronoi filtering. Manuscript, 1998.

[3] N. AMENTA, M. BERN AND D. EPPSTEIN. The crust and the β-skeleton: combinatorial curve reconstruction. *Graphical Models and Image Process.*, to appear.

[4] D. ATTALI. r-regular shape reconstruction from unorganized points. *In* "Proc. 13th Ann. Sympos. Comput. Geom., 1997", 248–253.

[5] F. AURENHAMMER. Voronoi diagrams – a study of a fundamental geometric data structure. *ACM Comput. Surveys* **23** (1991), 345–405.

[6] D. AVIS AND J. HORTON. Remarks on the sphere of influence graph. *In* "Proc. Conf. Discrete Geom. Convexity", J. E. Goodman et al. (eds.), Ann. New York Acad. Sci. **440** (1985), 323–327.

[7] C. L. BAJAJ, F. BERNARDINI AND G. XU. Automatic reconstruction of surfaces and scalar fields from 3D scans. *Comput. Graphics*, Proc. SIGGRAPH 1995, 109–118.

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

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

[10] J.-D. BOISSONNAT AND B. GEIGER. Three-dimensional reconstruction of complex shapes based on the Delaunay triangulation. *In* "Proc. Biomedical Image Process. Biomed. Visualization, 1993", 964–975.

[11] J. BRANDT AND V. R. ALGAZI. Continuous skeleton computation by Voronoi diagram. *Comput. Vision, Graphics, Image Process.* **55** (1992), 329–338.

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

[13] M. BURNS. *Automated Fabrication. Improving Productivity in Manufacturing.* Prentice Hall, Englewood Cliffs, New Jersey, 1993.

[14] L. P. CHEW. Guaranteed-quality mesh generation for curved surfaces. *In* "Proc. 9th Ann. Sympos. Comput. Geom., 1993", 274–280.

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

[16] W. D'ARCY THOMPSON. *Growth and Form.* Cambridge Univ. Press, 1917.

[17] H. EDELSBRUNNER. The union of balls and its dual shape. *Discrete Comput. Geom.* **13** (1995), 415–440.

[18] H. EDELSBRUNNER. Surface reconstruction by wrapping finite sets in space. Rept. rgi-tech-96-001, Raindrop Geomagic, Urbana, Illinois, 1996.

[19] H. EDELSBRUNNER, D. G. KIRKPATRICK AND R. SEIDEL. On the shape of a set of points in the plane. *IEEE Trans. Inform. Theory* **29** (1983), 551–559.

[20] 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.

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

[22] H. EDELSBRUNNER, G. ROTE AND E. WELZL. Testing the necklace condition for shortest tours and optimal factors in the plane. *Theoret. Comput. Sci.* **66** (1989), 157–180.

[23] H. EDELSBRUNNER AND N. R. SHAH. Triangulating topological spaces. *Internat. J. Comput. Geom. Appl.* **7** (1997), 365–378.

[24] A. FLORES. Über die Existenz n-dimensionaler Komplexe, die nicht in den $\mathbb{R}_{2n}$ topologisch einbettbar sind. *Ergeb. Math. Kolloq.* **5** (1932/33), 17–24.

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

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

[27] D. G. KIRKPATRICK AND J. D. RADKE. A framework for computational morphology. In *Computational Morphology*, G. Toussaint (ed.), Elsevier (1985), 217–248.

[28] B. LEE AND F. M. RICHARDS. The interpretation of protein structures: estimation of static accessibility. *J. Mol. Biol.* **55** (1971), 379–400.

[29] J. LERAY. Sur la forme des espaces topologiques et sur les point fixes des représentations. *J. Math. Pures Appl.* **24** (1945), 95–167.

[30] J. LIANG, H. EDELSBRUNNER AND C. WOODWARD. Anatomy of protein pockets and cavities: measurement of binding site geometry and implications for ligand design. Manuscript, 1997.

[31] S. K. LODHA AND R. FRANKE. Scattered data techniques for surfaces. In *Scientific Visualization: Methods and Applications*, G. M. Nielson, H. Hagen and F. Post (eds.), Springer-Verlag, Heidelberg, to appear.

[32] W. E. LORENSEN AND H. E. CLINE. Marching cubes: a high resolution 3D surface construction algorithm. *Comput. Graphics* **21**, Proc. SIGGRAPH 1987, 163–169.

[33] T. MARTINETZ AND K. SCHULTEN. Topology representing networks. *Neural Networks* **7** (1994), 507–522.

[34] M. MELKEMI. $\mathcal{A}$-shapes of a finite point set. *Correspondence in* "Proc. 13th Ann. Sympos. Comput. Geom., 1997", 367–369.

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

[36] Y.-L. O, A. TOET, D. FOSTER, H. J. A. M. HEIJMANS AND P. MEER (EDS.) *Shape in Picture. Mathematical Description of Shape in Grey-level images.* NATO ASI Series F: Computer and Systems Sciences **126**, Springer-Verlag, Berlin, 1994.

[37] RAINDROP GEOMAGIC, INC. `www.geomagic.com`.

[38] G. P. ROBINSON, A. C. F. COLCHESTER, L. D. GRIFFIN AND D. J. HAWKES. Integrated skeleton and boundary shape representation for medical image interpretation. *In* "Proc. European Conf. Comput. Vision, 1992", 725–729.

[39] R. THOM. *Structural Stability and Morphogenesis.* Addison Wesley, Reading, Massachusetts, 1989.

[40] E. R. VAN KAMPEN. Komplexe in euklidischen Räumen. *Abh. Math. Sem. Univ. Hamburg* **9** (1933), 72–78.

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