35 CURVE AND SURFACE RECONSTRUCTION

Tamal K. Dey

INTRODUCTION

The problem of reconstructing a shape from its sample appears in many scientific and engineering applications. Because of the variety in shapes and applications, many algorithms have been proposed over the last three decades, some of which exploit application-specific information and some of which are more general. We focus on techniques that apply to the general setting and have geometric and topological guarantees on the quality of reconstruction.

GLOSSARY

Simplex: A k-simplex in $\mathbb{R}^m$, $0 \leq k \leq m$, is the convex hull of $k+1$ affinely independent points in $\mathbb{R}^m$ where $0 \leq k \leq m$. The 0-, 1-, 2-, and 3-simplices are also called *vertices*, *edges*, *triangles*, and *tetrahedra* respectively.

Simplicial complex: A simplicial complex $\mathcal{K}$ is a collection of simplices with the conditions that, (i) all sub-simplices spanned by the vertices of a simplex in $\mathcal{K}$ are also in $\mathcal{K}$, and (ii) if $\sigma_1, \sigma_2 \in \mathcal{K}$ intersect, then $\sigma_1 \cap \sigma_2$ is a sub-simplex of both. The underlying space $|\mathcal{K}|$ of $\mathcal{K}$ is the set of all points in its simplices. (Cf. Chapter 15.)

Distance: Given two subsets $X, Y \subseteq \mathbb{R}^m$, the Euclidean distance between them is given by $d(X,Y) = \inf_{x\in X, y \in Y} \|x-y\|_2$. Additionally, $d(x,y)$ denotes the Euclidean distance between two points x and y in $\mathbb{R}^m$.

k-manifold: A k-manifold is a topological space where each point has a neighborhood homeomorphic to $\mathbb{R}^k$ or the halfspace $\mathbb{H}^k$. The points with $\mathbb{H}^k$ neighborhood constitute the boundary of the manifold.

Voronoi diagram: Given a point set $P \in \mathbb{R}^m$, a Voronoi cell V_p for each point $p \in P$ is defined as

$$V_p = \{x \in \mathbb{R}^m \mid d(x,p) \leq d(x,q), \forall q \in P\}\ .$$

The Voronoi diagram $\mathrm{Vor}\, P$ of P is the collection of all such Voronoi cells and their faces.

Delaunay triangulation: The Delaunay triangulation of a point set $P \in \mathbb{R}^m$ is a simplicial complex $\mathrm{Del}\, P$ such that a simplex with vertices $\{p_0, .., p_k\}$ is in $\mathrm{Del}\, P$ if and only if $\bigcap_{i=0,k} V_{p_i} \neq \emptyset$. (Cf. Chapter 27.)

Shape: A shape Σ is a subset of an Euclidean space.

Sample: A sample P of a shape Σ is a finite set of points from Σ.

Medial axis: The medial axis of a shape $\Sigma \in \mathbb{R}^m$ is the closure of the set of points in $\mathbb{R}^m$ that have more than one closest point in Σ. See Figure 35.1.1(a) for an illustration.

Preliminary version (July 19, 2017). To appear in the Handbook of Discrete and Computational Geometry,
J.E. Goodman, J. O'Rourke, and C. D. Tóth (editors), 3rd edition, CRC Press, Boca Raton, FL, 2017.

Local feature size: The local feature size for a shape $\Sigma \subseteq \mathbb{R}^m$ is a continuous function $f : \Sigma \to \mathbb{R}$ where $f(x)$ is the distance of $x \in \Sigma$ to the medial axis of Σ. See Figure 35.1.1(a).

ϵ-sample: A sample P of a shape Σ is an ϵ-sample if for each $x \in \Sigma$ there is a sample point $p \in P$ so that $d(p, x) \leq \epsilon f(x)$.

ϵ-local sample: A sample P of a shape Σ is ϵ-local if it is an ϵ-sample and for every pair of points $p, q \in P \times P$, $d(p, q) \geq \frac{\epsilon}{c}$ for some fixed constant $c \geq 1$.

ϵ-uniform sample: A sample P of a shape Σ is ϵ-uniform if for each $x \in \Sigma$ there is a sample point $p \in P$ so that $d(p, x) \leq \epsilon f_{min}$ where $f_{min} = \min\{f(x), x \in \Sigma\}$ and $\epsilon > 0$ is a constant.

35.1 CURVE RECONSTRUCTION

In its simplest form the reconstruction problem appeared in applications such as pattern recognition (Chapter 54), computer vision, and cluster analysis, where a curve in two dimensions is to be approximated from a set of sample points. In the 1980s several geometric graph constructions over a set of points in plane were discovered which reveal a pattern among the points. The influence graph of Toussaint [AH85], the β-skeleton of Kirkpatrick and Radke [KR85], the α-shapes of Edelsbrunner, Kirkpatrick, and Seidel [EKS83] are such graph constructions. Since then several algorithms have been proposed that reconstruct a curve from its sample with guarantees under some sampling assumption.

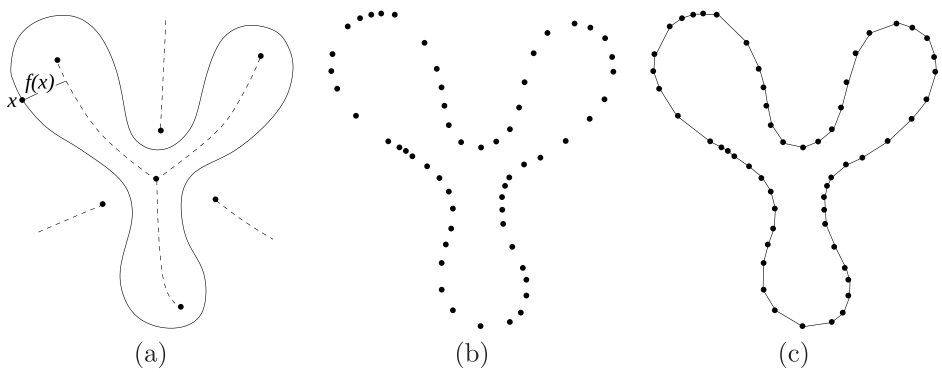

FIGURE 35.1.1
A smooth curve (solid), its medial axis (dashed) (a), sample (b), reconstruction (c).

GLOSSARY

Curve: A curve C in plane is the image of a function $p : [0, 1] \to \mathbb{R}^2$ where $p(t) = (x(t), y(t))$ for $t \in [0, 1]$ and $p[t] \neq p[t']$ for any $t \neq t'$ except possibly $t, t' \in \{0, 1\}$. It is *smooth* if p is differentiable and the derivative $\frac{d}{dt}p(t) = (\frac{dx(t)}{dt}, \frac{dy(t)}{dt})$ does not vanish.

Preliminary version (July 19, 2017). To appear in the Handbook of Discrete and Computational Geometry,
J.E. Goodman, J. O'Rourke, and C. D. Tóth (editors), 3rd edition, CRC Press, Boca Raton, FL, 2017.

Boundary: A curve C is said to have no boundary if $p[0] = p[1]$; otherwise, it is a curve with boundary.

Reconstruction: The reconstruction of C from its sample P is a geometric graph $G = (P, E)$ where an edge pq belongs to E if and only if p and q are adjacent sample points on C. See Figure 35.1.1.

Semiregular curve: One for which the left tangent and right tangent exist at each point of the curve, though they may be different.

UNIFORM SAMPLE

α-shapes: Edelsbrunner, Kirkpatrick, and Seidel [EKS83] introduced the concept of α-shape of a finite point set $P \subset \mathbb{R}^2$. It is the underlying space of a simplicial complex called the ***α-complex***. The α-complex of P is defined by all simplices with vertices in P that have an empty circumscribing disk of radius α. Bernardini and Bajaj [BB97] show that the α-shapes reconstruct curves from ϵ-uniform samples if ϵ is sufficiently small and α is chosen appropriately.

r-regular shapes: Attali considered r-regular shapes that are constructed using certain morphological operations with r as a parameter [Att98]. It turns out that these shapes are characterized by requiring that any circle passing through the points on the boundary has radius greater than r. A sample P from the boundary curve $C \subset \mathbb{R}^2$ of such a shape is called γ-dense if each point $x \in C$ has a sample point within γr distance. Let η_{pq} be the sum of the angles opposite to pq in the two incident Delaunay triangles at a Delaunay edge $pq \in \operatorname{Del} P$. The main result in [Att98] is that if $\gamma < \sin\frac{\pi}{8}$, Delaunay edges with $\eta_{pq} < \pi$ reconstruct C.

EMST: Figueiredo and Gomes [FG95] show that the Euclidean minimum spanning tree (EMST) reconstructs curves with boundaries when the sample is sufficiently dense. The sampling density condition that is used to prove this result is equivalent to that of ϵ-uniform sampling for an appropriate $\epsilon > 0$. Of course, EMST cannot reconstruct curves without boundaries and/or multiple components.

NONUNIFORM SAMPLE

Crust: Amenta, Bern, and Eppstein [ABE98] proposed the first algorithm called Crust to reconstruct a curve with guarantee from a sample that is not necessarily uniform. The algorithm operates in two phases. The first phase computes the Voronoi diagram of the sample points in P. Let V be the set of Voronoi vertices in this diagram. The second phase computes the Delaunay triangulation of the larger set $P \cup V$. The Delaunay edges that connect only sample points in this triangulation constitute the crust; see Figure 35.1.2.

The theoretical guarantee of the Crust algorithm is based on the notion of dense sampling that respects features of the sampled curve. The important concepts of local feature size and ϵ-sample were introduced by Amenta, Bern, and Eppstein [ABE98]. They prove:

THEOREM 35.1.1

For $\epsilon < \frac{1}{5}$, given an ϵ-sample P of a smooth curve $C \subset \mathbb{R}^2$ without boundary, the Crust reconstructs C from P.

Preliminary version (July 19, 2017). To appear in the Handbook of Discrete and Computational Geometry, J.E. Goodman, J. O'Rourke, and C. D. Tóth (editors), 3rd edition, CRC Press, Boca Raton, FL, 2017.

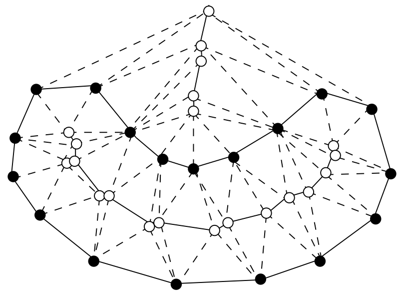

FIGURE 35.1.2
Crust edges (solid) among the Delaunay triangulation of a sample and their Voronoi vertices.

The two Voronoi diagram computations of the Crust are reduced to one by Gold and Snoeyink [GS01].

Nearest neighbor: After the introduction of the Crust, Dey and Kumar [DK99] proposed a curve reconstruction algorithm called NN-Crust based on nearest neighbors. They showed that all nearest neighbor edges that connect a point to its Euclidean nearest neighbor must be in the reconstruction if the input is $\frac{1}{3}$-sample. However, not all edges of the reconstruction are necessarily nearest neighbor edges. The remaining edges are characterized as follows. Let p be a sample point with only one nearest neighbor edge pq incident to it. Consider the halfplane with pq being an outward normal to its bounding line through p, and let r be the nearest to p among all sample points lying in this halfplane. Call pr the half-neighbor edge of p. Dey and Kumar show that all half-neighbor edges must also be in the reconstruction for a $\frac{1}{3}$-sample.

The algorithm first computes all nearest neighbor edges and then computes the half-neighbor edges to complete the reconstruction. Since all edges in the reconstruction must be a subset of Delaunay edges if the sample is sufficiently dense, all nearest neighbor and half-neighbor edges can be computed from the Delaunay triangulation. Thus, as Crust, this algorithm runs in $O(n \log n)$ time for a sample of n points.

THEOREM 35.1.2

For $\epsilon \leq \frac{1}{3}$, given an ϵ-sample P of a smooth curve $C \subset \mathbb{R}^2$ without boundary, NN-Crust reconstructs C from P.

NONSMOOTHNESS, BOUNDARIES

The crust and nearest neighbor algorithms assume that the sampled curve is smooth and has no boundary. Nonsmoothness and boundaries make reconstruction harder.

Traveling Salesman Path: Giesen [Gie00] considered a fairly large class of nonsmooth curves and showed that the Traveling Salesman Path (or Tour) reconstructs them from sufficiently dense samples. A semiregular curve C is ***benign*** if the angle between the two tangents at each point is less than π. Giesen proved the following:

THEOREM 35.1.3

For a benign curve $C \subset \mathbb{R}^2$, there exists an $\epsilon > 0$ so that if P is an ϵ-uniform sample of C, then C is reconstructed by the Traveling Salesman Path (or Tour) in case C has boundary (or no boundary).

The uniform sampling condition for the Traveling Salesman approach was later removed by Althaus and Mehlhorn [AM02], who also gave a polynomial-time algorithm to compute the Traveling Salesman Path (or Tour) in this special case of curve reconstruction. Obviously, the Traveling Salesman approach cannot handle curves with multiple components. Also, the sample points representing the boundary need to be known a priori to choose between path or tour.

Conservative Crust: In order to allow boundaries in curve reconstruction, it is essential that the sample points representing boundaries are detected. Dey, Mehlhorn, and Ramos presented such an algorithm, called the *conservative crust* [DMR00].

Any algorithm for handling curves with boundaries faces a dilemma when an input point set samples a curve without boundary densely and simultaneously samples another curve with boundary densely. This dilemma is resolved in conservative crust by a justification on the output. For any input point set P, the graph output by the algorithm is guaranteed to be the reconstruction of a smooth curve $C \subset \mathbb{R}^2$ possibly with boundary for which the input point set is a dense sample. The main idea of the algorithm is that an edge pq is chosen in the output only if there is a large enough ball centering the midpoint of pq which is empty of all Voronoi vertices in the Voronoi diagram of P. The rationale behind this choice is that these edges are small enough with respect to local feature size of C since the Voronoi vertices approximate its medial axis.

With a certain sampling condition tailored to handle nonsmooth curves, Funke and Ramos used conservative crust to reconstruct nonsmooth curves that may have boundaries [FR01].

SUMMARIZED RESULTS

The strengths and deficiencies of the discussed algorithms are summarized in Table 35.1.1.

TABLE 35.1.1 Curve reconstruction algorithms.

ALGORITHM	SAMPLE	SMOOTHNESS	BOUNDARY	COMPONENTS
α-shape	uniform	required	none	multiple
r-regular shape	uniform	required	none	multiple
EMST	uniform	required	exactly two	single
Crust	non-uniform	required	none	multiple
Nearest neighbor	non-uniform	required	none	multiple
Traveling Salesman	non-uniform	not required	must be known	single
Conservative crust	non-uniform	required	any number	multiple

Preliminary version (July 19, 2017). To appear in the Handbook of Discrete and Computational Geometry, J.E. Goodman, J. O'Rourke, and C. D. Tóth (editors), 3rd edition, CRC Press, Boca Raton, FL, 2017.

OPEN PROBLEM

All algorithms described above assume that the sampled curve does not cross itself. It is open to devise an algorithm that can reconstruct such curves under some reasonable sampling condition.

35.2 SURFACE RECONSTRUCTION

A number of surface reconstruction algorithms have been designed in different application fields. The problem appeared in medical imaging where a set of cross sections obtained via CAT scan or MRI needs to be joined with a surface. The points on the boundary of the cross sections are already joined by a polygonal curve and the output surface needs to join these curves in consecutive cross sections. A dynamic programming based solution for two such consecutive curves was first proposed by Fuchs, Kedem, and Uselton [FKU77]. A negative result by Gitlin, O'Rourke, and Subramanian [GOS96] shows that, in general, two polygonal curves cannot be joined by a nonself-intersecting surface with only those vertices; even deciding its feasibility is NP-hard. Several solutions with the addition of Steiner points have been proposed to overcome the problem, see [MSS92, BG93]. The most general version of the surface reconstruction problem does not assume any information about the input points other than their 3D coordinates, and requires a piecewise linear approximation of the sampled surface; see Figure 35.2.1. In the context of computer graphics and vision, this problem has been investigated intensely with emphasis on the practical effectiveness of the algorithms [BMR+99, Boi84, CL96, GCA13, GKS00, HDD+92]. In computational geometry, several algorithms have been designed based on Voronoi/Delaunay diagrams that have guarantees on geometric proximity (Hausdorff closeness) and topological equivalence (homeomorphism/isotopy). We focus mainly on them.

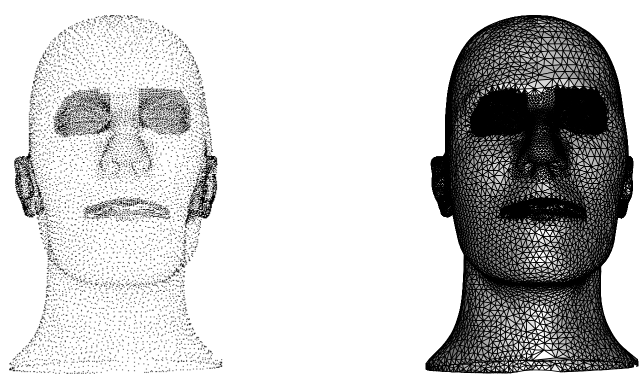

FIGURE 35.2.1
A point sample and the reconstructed surface.

Preliminary version (July 19, 2017). To appear in the Handbook of Discrete and Computational Geometry, J.E. Goodman, J. O'Rourke, and C. D. Tóth (editors), 3rd edition, CRC Press, Boca Raton, FL, 2017.

GLOSSARY

Surface: A surface $S \subset \mathbb{R}^3$ is a 2-manifold embedded in $\mathbb{R}^3$. Thus each point $p \in S$ has a neighborhood homeomorphic to $\mathbb{R}^2$ or halfplane $\mathbb{H}^2$. The points with neighborhoods homeomorphic to $\mathbb{H}^2$ constitute the boundary ∂S of S.

Smooth Surface: A surface $S \subset \mathbb{R}^3$ is smooth if for each point $p \in S$ there is a neighborhood $W \subseteq \mathbb{R}^3$ and a map $\pi : U \to W \cap S$ of an open set $U \subset \mathbb{R}^2$ onto $W \cap S$ so that

(i) π is differentiable,

(ii) π is a homeomorphism,

(iii) for each $q \in U$ the differential $d\pi_q$ is one-to-one.

A surface $S \subset \mathbb{R}^3$ with boundary ∂S is smooth if its interior $S \setminus \partial S$ is a smooth surface and ∂S is a smooth curve.

Smooth Closed Surface: We call a smooth surface $S \subset \mathbb{R}^3$ closed if it is compact and has no boundary.

Restricted Voronoi: Given a subspace $\mathbb{N} \subseteq \mathbb{R}^3$ and a point set $P \subseteq \mathbb{R}^3$, the restricted Voronoi diagram of P w.r.t $\mathbb{N}$ is $\mathrm{Vor}\, P|_{\mathbb{N}} = \{F \cap \mathbb{N} \mid F \in \mathrm{Vor}\, P\}$.

Restricted Delaunay: The dual of $\mathrm{Vor}\, P|_{\mathbb{N}}$ is called the restricted Delaunay triangulation $\mathrm{Del}\, P|_{\mathbb{N}}$ defined as

$$\mathrm{Del}\, P|_{\mathbb{N}} = \{\sigma \mid \sigma = \mathrm{Conv}\,\{p_0, ..., p_k\} \in \mathrm{Del}\, P \text{ where } (\textstyle\bigcap_{i=0,k} V_{p_i}) \cap \mathbb{N} \neq \emptyset\}.$$

Watertight surface: A 2-complex $\mathcal{K}$ embedded in $\mathbb{R}^3$ is called watertight if the underlying space $|\mathcal{K}|$ of $\mathcal{K}$ is the boundary of the closure of some 3-manifold in $\mathbb{R}^3$.

Hausdorff ϵ-close: A subset $X \subset \mathbb{R}^3$ is Hausdorff ϵ-close to a surface $S \subset \mathbb{R}^3$ if every point $y \in S$ has a point $x \in X$ with $d(x, y) \leq \epsilon f(y)$, and similarly every point $x \in X$ has a point $y \in S$ with $d(x, y) \leq \epsilon f(y)$.

Homeomorphism: Two topological spaces (e.g., surfaces) are homeomorphic if there is a continuous bijective map between them with continuous inverse.

Steiner points: The points used by an algorithm that are not part of the finite input point set are called Steiner points.

α-SHAPES

Generalization of α-shapes to 3D by Edelsbrunner and Mücke [EM94] can be used for surface reconstruction in case the sample is more or less uniform. An alternate definition of α-shapes in terms of the restricted Delaunay triangulation is more appropriate for surface reconstruction. Let $\mathbb{N}$ denote the space of all points covered by open balls of radius α around each sample point $p \in P$. The α-shape for P is the underlying space of the α-complex which is the restricted Delaunay triangulation $\mathrm{Del}\, P|_{\mathbb{N}}$; see Figure 35.2.2 below for an illustration in 2D. It is shown that the α-shape is always homotopy equivalent to $\mathbb{N}$. If P is a sample of a surface $S \subset \mathbb{R}^3$, the space $\mathbb{N}$ becomes homotopy equivalent to S if α is chosen appropriately and P

Preliminary version (July 19, 2017). To appear in the Handbook of Discrete and Computational Geometry,
J.E. Goodman, J. O'Rourke, and C. D. Tóth (editors), 3rd edition, CRC Press, Boca Raton, FL, 2017.

is sufficiently dense [EM94]. Therefore, by transitivity of homotopy equivalence, the α-shape is homotopy equivalent to S if α is appropriate and the sample P is sufficiently dense.

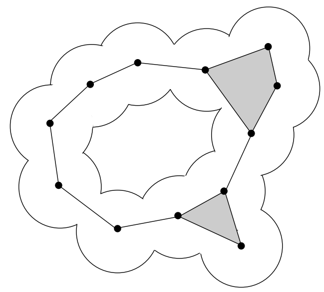

FIGURE 35.2.2
Alpha shape of a set of points in $\mathbb{R}^2$.

The major drawback of α-shapes is that they require a nearly uniform sample and an appropriate parameter α for reconstruction.

CRUST

The Crust algorithm for curve reconstruction was generalized for surface reconstruction by Amenta and Bern [AB99]. In case of curves in 2D, Voronoi vertices for a dense sample lie close to the medial axis. That is why a second Voronoi diagram with the input sample points together with the Voronoi vertices is used to separate the Delaunay edges that reconstruct the curve. Unfortunately, Voronoi vertices in 3D can lie arbitrarily close to the sampled surface. One can place four arbitrarily close points on a smooth surface which lie near the diametric plane of the sphere defined by them. This sphere can be made empty of any other input point and thus its center as a Voronoi vertex lies close to the surface. With this important observation Amenta and Bern forsake the idea of putting all Voronoi vertices in the second phase of crust and instead identify a subset of Voronoi vertices called *poles* that lie far away from the surface, and in fact close to the medial axis.

Let P be an ϵ-sample of a smooth closed surface $S \subset \mathbb{R}^3$. Let V_p be a Voronoi cell in the Voronoi diagram $\mathrm{Vor}\, P$. The farthest Voronoi vertex of V_p from p is called the positive pole of p. Call the vector from p to the positive pole the *pole vector* for p; this vector approximates the surface normal $\mathbf{n}_p$ at p. The Voronoi vertex of V_p that lies farthest from p in the opposite direction of the pole vector is called its negative pole. The opposite direction is specified by the condition that the vector from p to the negative pole must make an angle more than $\frac{\pi}{2}$ with the pole vector. Figure 35.2.3(a) illustrates these definitions. If V_p is unbounded, the positive pole is taken at infinity and the direction of the pole vector is taken as the average of all directions of the unbounded Voronoi edges in V_p.

The Crust algorithm in 3D proceeds as follows. First, it computes $\mathrm{Vor}\, P$ and then identifies the set of poles, say L. The Delaunay triangulation of the point set $P \cup L$ is computed and the set of Delaunay triangles, T, is filtered that have all three vertices only from P. This set of triangles almost approximates S but may not form a surface. Nevertheless, the set T includes all restricted Delaunay triangles in $\mathrm{Del}\, P|_S$. According to a result by Edelsbrunner and Shah [ES97], the underlying space $|\mathrm{Del}\, P|_S|$ of $\mathrm{Del}\, P|_S$ is homeomorphic to S if each Voronoi face satisfies a topological condition called the "closed ball property." Amenta and Bern show that if P is an ϵ-sample for $\epsilon \leq 0.06$, each Voronoi face in $\mathrm{Vor}\, P$ satisfies this property. This means that, if the triangles in $\mathrm{Del}\, P|_S$ could be extracted from T, we would have a surface homeomorphic to S. Unfortunately, it is impossible to detect the restricted Delaunay triangles of $\mathrm{Del}\, P|_S$ since S is unknown. However, the fact that T contains them is used in a manifold extraction step that computes a manifold out of T after a normal filtering step. This piecewise linear manifold surface is output by Crust. The Crust guarantees that the output surface lies very close to S.

THEOREM 35.2.1

For $\epsilon \leq 0.06$, given an ϵ-sample P of a smooth closed surface $S \subset \mathbb{R}^3$, the Crust algorithm produces a 2-complex that is Hausdorff ϵ-close to S.

Actually, the output of Crust is also homeomorphic to the sampled surface under the stated condition of the theorem above, a fact which was proved later in the context of the Cocone algorithm discussed next.

COCONE

The Cocone algorithm was developed by Amenta, Choi, Dey, and Leekha [ACDL02]. It simplified the reconstruction by Crust and enhanced its proof of correctness.

A ***cocone*** C_p for a sample point p is defined as the complement of the double cone with p as apex and the pole vector as axis and an opening angle of $\frac{3\pi}{4}$; see Figure 35.2.3(b). Because the pole vector at p approximates the surface normal $\mathbf{n}_p$, the cocone C_p (clipped within V_p) approximates a thin neighborhood around the tangent plane at p. For each point p, the algorithm then determines all Voronoi edges in V_p that are intersected by the cocone C_p. The dual Delaunay triangles of these Voronoi edges constitute the set of candidate triangles T.

It can be shown that the circumscribing circles of all candidate triangles are small [ACDL02]. Specifically, if $pqr \in T$ has circumradius r, then

(i) $r = O(\epsilon) f(x)$ where $f(x) = \min\{f(p), f(q), f(r)\}$.

It turns out that any triangle with such small circumradius must lie almost parallel to the surface, i.e., if $\mathbf{n}_{pqr}$ is the normal to a candidate triangle pqr, then

(ii) $\angle(\mathbf{n}_{pqr}, \mathbf{n}_x) = O(\epsilon)$ up to orientation where $x \in \{p, q, r\}$.

Also, it is proved that

(iii) T includes all restricted Delaunay triangles in $\mathrm{Del}\, P|_S$.

These three properties of the candidate triangles ensure that a *manifold extraction* step, as in the Crust algorithm, extracts a piecewise-linear surface which is homeomorphic to the original surface S.

Cocone uses a single Voronoi diagram as opposed to two in the Crust algorithm and also eliminates the normal filtering step. It provides the following guarantees.

Preliminary version (July 19, 2017). To appear in the Handbook of Discrete and Computational Geometry, J.E. Goodman, J. O'Rourke, and C. D. Tóth (editors), 3rd edition, CRC Press, Boca Raton, FL, 2017.

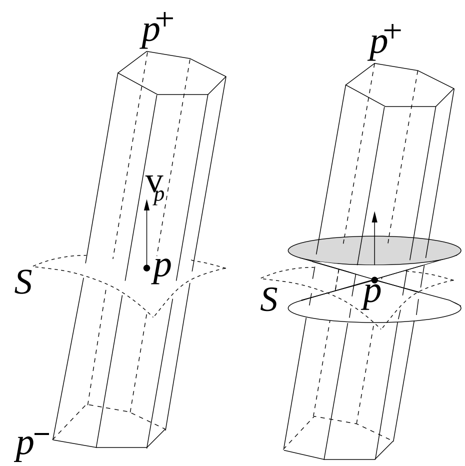

FIGURE 35.2.3
A long thin Voronoi cell V_p, the positive pole p^+, the pole vector $\mathbf{v}_p$ and the negative pole p^- (a), the cocone (b).

THEOREM 35.2.2

For $\epsilon \leq 0.06$, given a sample P of a smooth closed surface $S \subset \mathbb{R}^3$, the Cocone algorithm computes a Delaunay subcomplex $N \subseteq \mathrm{Del}\, P$ where $|N|$ is Hausdorff ϵ-close and is homeomorphic to S.

Actually, the homeomorphism property can be strengthened to *isotopy*, a stronger topological equivalence condition. Because of the Voronoi diagram computation, the Cocone runs in $O(n^2)$ time and space. Funke and Ramos [FR02] improved its complexity to $O(n \log n)$ though the resulting algorithm seems impractical. Cheng, Jin, and Lau [CJL17] simplified this approach making it more practical.

NATURAL NEIGHBOR

Boissonnat and Cazals [BC02] revisited the approach of Hoppe et al. [HDD$^+$92] by approximating the sampled surface as the zero set of a signed distance function. They used natural neighbors and an ϵ-sampling condition to provide output guarantees.

Given an input point set $P \subset \mathbb{R}^3$, the ***natural neighbors*** $N_{x,P}$ of a point $x \in \mathbb{R}^3$ are the Delaunay neighbors of x in $\mathrm{Del}\,(P \cup x)$. Letting $V(x)$ denote the Voronoi cell of x in $\mathrm{Vor}\,(P \cup x)$, this means

$$N_{x,P} = \{p \in P \,|\, V(x) \cap V_p \neq \emptyset\}.$$

Let $A(x,p)$ denote the volume stolen by x from V_p, i.e.,

$$A(x,p) = \mathrm{Vol}\,(V(x) \cap V_p).$$

The natural coordinate associated with a point p is a continuous function $\lambda_p : \mathbb{R}^3 \to \mathbb{R}$ where

$$\lambda_p(x) = \frac{A(x,p)}{\Sigma_{q \in P} A(x,q)}.$$

Preliminary version (July 19, 2017). To appear in the Handbook of Discrete and Computational Geometry, J.E. Goodman, J. O'Rourke, and C. D. Tóth (editors), 3rd edition, CRC Press, Boca Raton, FL, 2017.

Some of the interesting properties of λ_p are that it is continuously differentiable except at p, and any point $x \in \mathbb{R}^3$ is a convex combination of its natural neighbors: $\Sigma_{p \in N_{x,P}} \lambda_p(x)p = x$. Boissonnat and Cazals assume that each point p is equipped with a unit normal $\mathbf{n}_p$ which can either be computed via pole vectors, or is part of the input. A distance function $h_p : \mathbb{R}^3 \to \mathbb{R}$ for each point p is defined as $h_p(x) = (p-x) \cdot \mathbf{n}_p$. A global distance function $h : \mathbb{R}^3 \to \mathbb{R}$ is defined by interpolating these local distance functions with natural coordinates. Specifically,

$$h(x) = \Sigma_{p \in P} \lambda_p^{1+\delta}(x) h_p(x).$$

The δ term in the exponent is added to make h continuously differentiable. By definition, $h(x)$ locally approximates the signed distance from the tangent plane at each point $p \in P$ and, in particular, $h(p) = 0$.

Since h is continuously differentiable, $\widehat{S} = h^{-1}(0)$ is a smooth surface unless 0 is a critical value. A discrete approximation of $\widehat{S}$ can be computed from the Delaunay triangulation of P as follows. All Voronoi edges that intersect $\widehat{S}$ are computed via the sign of h at their two endpoints. The dual Delaunay triangles of these Voronoi edges constitute a piecewise linear approximation of $\widehat{S}$. If the input sample P is an ϵ-sample for sufficiently small ϵ, then a theorem similar to that for Cocone holds.

MORSE FLOW

Morse theory is concerned with the study of critical points of real-valued functions on manifolds. Although the original theory was developed for smooth manifolds, various extensions have been made to incorporate more general settings. The theory builds upon a notion of gradient of the Morse function involved. The critical points where the gradient vanishes are mainly of three types, local minima, local maxima, and saddles. The gradient vector field usually defines a flow that can be thought of as a mechanism for moving points along the steepest ascent. Some surface reconstruction algorithms build upon this concept of ***Morse flow***.

Given a point sample P of a smooth closed surface $S \subseteq \mathbb{R}^3$, consider the distance function $d : \mathbb{R}^3 \to \mathbb{R}$ where $d(x) := d(x, P)$ is the distance of x to the nearest sample point in P. This function is not differentiable everywhere. Still, one can define a flow for d. For every point $x \in \mathbb{R}^3$, let $A(x) \subseteq P$ be the set of sample points closest to x. The driver $r(x)$ for x is the closest point in the convex hull of $A(x)$. In the case where $r(x) = x$, we say x is critical and regular otherwise. The normalized gradient of d at a regular point x is defined as the unit vector in the direction $x - r(x)$. Notice that the gradient vanishes at a critical point x where $x = r(x)$. The flow induced by this vector field is a map $\phi : \mathbb{R}^3 \times \mathbb{R}^+ \to \mathbb{R}^3$ such that the right derivative of $\phi(x, t)$ at every point x with respect to time t equals the gradient vector. This flow defines flow curves (integral lines) along which points move toward the steepest ascent of d and arrive at critical points in the limit. See Figure 35.2.4 for an illustration in $\mathbb{R}^2$.

It follows from the definition that the critical points of d occur at the points where a Delaunay simplex in Del P intersects its dual Voronoi face in Vor P. Given a critical point c of d, the set of all points that flow into c constitute the *stable manifold* of c. The set of all stable manifolds partitions $\mathbb{R}^3$ into cells that form a cell complex together. Giesen and John [GJ08] named it as ***flow complex*** and studied its properties. The dimension of each stable manifold is the index of its associated critical point. Index-0 critical points are minima which are the Delaunay

Preliminary version (July 19, 2017). To appear in the Handbook of Discrete and Computational Geometry, J.E. Goodman, J. O'Rourke, and C. D. Tóth (editors), 3rd edition, CRC Press, Boca Raton, FL, 2017.

vertices, or equivalently the points in P. Index-1 critical points, also called the 1-saddles, are the intersections of Delaunay edges and their dual Voronoi facets. Index-2 critical points, also called the 2-saddles, are the intersections of Delaunay triangles and their dual Voronoi edges. Index-3 critical points are maxima which are a subset of Voronoi vertices.

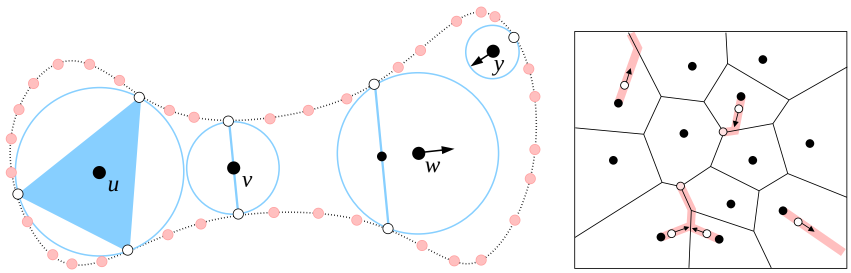

FIGURE 35.2.4
On left: Point sample from a curve, sets $A(x)$ are shown with hollow circles, driver for w is the smaller black circle, driver of y is the hollow circle, u and v are critical; On right: some flow curves for a point set in $\mathbb{R}^2$.

A surface reconstruction algorithm based on the flow complex was proposed by Dey, Giesen, Ramos, and Sadri [DGRS08]. They observe that the critical points of d separate into two groups, one near the surface S, called the *surface critical points*, and the other near the medial axis M, called the *medial axis critical points*. Let $S_\eta = \{x \in \mathbb{R}^3 \mid d(x, S) \leq \eta\}$ and $M_\eta = \{x \in \mathbb{R}^3 \mid d(x, M) = \eta\}$ denote the η-offset of S and M respectively. For any point $c \in \mathbb{R}^3$, let $\tilde{c}$ be its orthogonal projection on S, and $\check{c}$ be the point where the ray $\tilde{c}c$ intersects M first time. Let $\rho(\tilde{c}) = d(\tilde{c}, \check{c})$. An important result proved in [DGRS08] is:

THEOREM 35.2.3

For $\epsilon < \frac{1}{3}$, let P be an ϵ-sample of a smooth closed surface $S \subset \mathbb{R}^3$. Let c be any critical point of the distance function d. Then, either $c \in S_{\epsilon^2 f(\tilde{c})}$, or in $M_{2\epsilon\rho(\tilde{c})}$.

The algorithm in [DGRS08] first separates the medial axis critical points from the surface ones using an angle criterion. The union of the stable manifolds for the medial axis critical points separates further into two connected components, one for the outer medial axis critical points and the other for the inner medial axis critical points. These two connected components can be computed using a union-find data structure. The boundary of any one of these connected components is output as the reconstructed surface.

The main disadvantage of the flow complex based surface reconstruction is that the stable manifolds constituting this complex are not necessarily subcomplexes of the Delaunay complex. Consequently, its construction is more complicated. A Morse theory based reconstruction that sidesteps this difficulty is the Wrap algorithm of Edelsbrunner [Ede03]. A different distance function is used in Wrap. A Delaunay circumball $B(c, r)$ that circumscribes a Delaunay simplex can be treated as a weighted point $\hat{c} = (c, r)$. For any point $x \in \mathbb{R}^3$, one can define the weighted distance which is also called the *power distance* as $\pi(x, \hat{c}) = d(x, c)^2 - r^2$. For a point set P, let C denote the centers of the Delaunay balls for simplices in $\mathrm{Del}\, P$ and $\widehat{C}$ denote the corresponding weighted points. These Delaunay balls also include those of the infinite tetrahedra formed by the convex hull triangles and a point at infinity. Their centers are at infinity and their radii are infinite.

Define a distance function $g : \mathbb{R}^3 \to \mathbb{R}$ as $g(x) = \min_{\hat{c} \in \widehat{C}} \pi(x, \hat{c})$. Consider the Voronoi diagram of $\widehat{C}$ with the power distance metric. This diagram, also called the power diagram of $\widehat{C}$, denoted as $\mathrm{Pow}\,\widehat{C}$, coincides with the Delaunay triangulation $\mathrm{Del}\,P$ extended with the infinite tetrahedra. A point $x \in \mathbb{R}^3$ contained in a Delaunay tetrahedron $\sigma \in \mathrm{Del}\,P$ has distance $g(x) = \pi(x, \hat{c})$ where c is the Delaunay ball of σ. Analogous to d, one can define a flow for the distance function g whose critical points coincide with those of d. Edelsbrunner defined a *flow relation* among the Delaunay simplices using the flow for g. Let σ be a proper face of two simplices τ and ζ. We say $\tau < \sigma < \zeta$ if there is a point x in the interior of σ such that the flow curve through x proceeds from the interior of τ into the interior of ζ. The Wrap algorithm, starting from the infinite tetrahedra, collapses simplices according to the flow relation. It finds a simplex σ with a coface ζ where $\sigma < \zeta$ and ζ is the only coface adjacent to σ. The collapse modifies the current complex $\mathcal{K}$ to $\mathcal{K} \setminus \{\sigma, \zeta\}$, which is known to maintain a homotopy equivalence between the two complexes. The algorithm stops when it can no longer find a simplex to collapse. The output is a subcomplex of $\mathrm{Del}\,P$ and is necessarily homotopy equivalent to a 3-ball. The algorithm can be modified to create an output complex of higher genus by starting the collapse from other source tetrahedra. The boundary of this complex can be taken as the output approximating S. There is no guarantee for topological equivalence between the output complex and the surface S for the original Wrap algorithm. Ramos and Sadri [RS07] proposed a version of the Wrap algorithm that ensures this topological equivalence under a dense ϵ-sampling assumption.

WATERTIGHT SURFACES

Most of the surface reconstruction algorithms face a difficulty while dealing with undersampled surfaces and noise. While some heuristics such as in [DG03] can detect undersampling, it leaves holes in the surface near the vicinity of undersampling. Although this may be desirable for reconstructing surfaces with boundaries, many applications such as CAD designs require that the output surface be *watertight*, i.e., a surface that bounds a solid.

The natural neighbor algorithm of [BC02] can be adapted to guarantee a watertight surface. Recall that this algorithm approximates a surface $\widehat{S}$ implicitly defined by the zero set of a smooth map $h : \mathbb{R}^3 \to \mathbb{R}$. This surface is a smooth 2-manifold without boundary in $\mathbb{R}^3$. However, if the input sample P is not dense for this surface, the reconstructed output may not be watertight. Boissonnat and Cazals suggest to sample more points on $\widehat{S}$ to obtain a dense sample for $\widehat{S}$ and then reconstruct it from the new sample.

Amenta, Choi, and Kolluri [ACK01] use the crust approach to design the Power Crust algorithm to produce watertight surfaces. This algorithm first distinguishes the inner poles that lie inside the solid bounded by the sampled surface S from the outer poles that lie outside. A consistent orientation of the pole vectors is used to decide between inner and outer poles. To prevent outer poles at infinity, eight corners of a large box containing the sample are added. Let $\mathbb{B}_O$ and $\mathbb{B}_I$ denote the Delaunay balls centered at the outer and inner poles respectively. The union of Delaunay balls in $\mathbb{B}_I$ approximate the solid bounded by S. The union of Delaunay balls in $\mathbb{B}_O$ do not approximate the entire exterior of S although one of

Preliminary version (July 19, 2017). To appear in the Handbook of Discrete and Computational Geometry,
J.E. Goodman, J. O'Rourke, and C. D. Tóth (editors), 3rd edition, CRC Press, Boca Raton, FL, 2017.

its boundary components approximates S. The implication is that the cells in the power diagram Pow $(\mathbb{B}_O \cup \mathbb{B}_I)$ can be partitioned into two sets, with the boundary between approximating S. The facets in the power diagram Pow $(\mathbb{B}_O \cup \mathbb{B}_I)$ that separate cells generated by inner and outer poles form this boundary which is output by Power Crust.

Dey and Goswami [DG03] announced a watertight surface reconstructor called Tight Cocone. This algorithm first computes the surface with Cocone, which may leave some holes in the surface due to anomalies in sampling. A subsequent sculpting [Boi84] in the Delaunay triangulation of the input points recover triangles that fill the holes. Unlike Power Crust, Tight Cocone does not add Steiner points.

BOUNDARY

The algorithms described so far work on the assumption that the sampled surface has no boundary. The reconstruction of smooth surfaces with boundary is more difficult because the algorithm has to reconstruct the boundary from the sample as well. The algorithms such as Crust and Cocone cannot be extended to surfaces with boundary because they employ a manifold extraction step which iteratively prunes triangles with edges adjacent to a single triangle. Surfaces with non-empty boundaries necessarily contain such triangles in their reconstruction, and thus cannot withstand such a pruning step. Another difficulty arises on the theoretical front because the restricted Delaunay triangulation $\mathrm{Del}\, P|_S$ of a sample P for a surface S with boundary may not be homeomorphic to S no matter how dense P is [DLRW09]. This property is a crucial ingredient for proving the correctness of the Crust and Cocone algorithms. Specifically, the manifold extraction step draws upon this property. Dey, Li, Ramos, and Wenger [DLRW09] sidestepped this difficulty by replacing the prune-and-walk step with a peeling step.

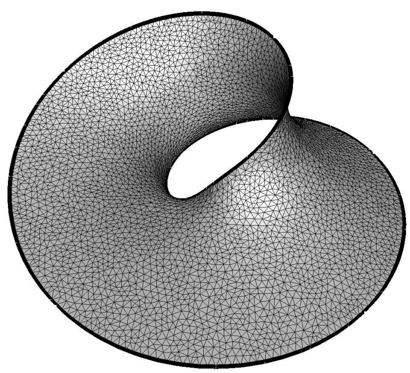

FIGURE 35.2.5
Reconstruction of a non-orientable surface with boundary (Möbius band) with Peel algorithm.

The Peel algorithm in [DLRW09] works as follows. It takes a point sample P and a positive real α as input. First, it computes the α-complex of P with the input parameter $\alpha > 0$. If P is dense, the Delaunay tetrahedra retained in this complex is proven to be "flat" meaning that they all have two non-adjacent edges,

Preliminary version (July 19, 2017). To appear in the Handbook of Discrete and Computational Geometry, J.E. Goodman, J. O'Rourke, and C. D. Tóth (editors), 3rd edition, CRC Press, Boca Raton, FL, 2017.

called flat edges, that subtend an internal angle close to π. A flat tetrahedron t is peelable if one of its flat edges, say e, is not adjacent to any tetrahedron other than t. A peeling of t means that the two triangles adjacent to e are removed along with e and t from the current complex. The algorithm successively finds peelable tetrahedra and peels them. It stops when there is no peelable tetrahedron and outputs the current complex. It is shown that the algorithm can always find a peelable tetrahedron as long as there is any tetrahedron at all. This means that the output is a set of triangles that are not adjacent to any tetrahedra. Dey et al. prove that the underlying space of this set of triangles is isotopic to S if P is sufficiently dense. Figure 35.2.5 shows an output surface computed by Peel.

The correctness of the Peel algorithm depends on the assumption that the input P is a globally uniform, that is, ϵ-uniform sample of S. A precise statement of the guarantee is given in the following theorem.

THEOREM 35.2.4

Let S be a smooth closed surface with boundary. For a sufficiently small $\epsilon > 0$ and $6\epsilon < \alpha \leq 6\epsilon + O(\epsilon)$, if P is an ϵ-uniform sample of S, Peel(P, α) produces a Delaunay sub-complex isotopic and Hausdorff ϵ-close to S.

NOISE

The input points are assumed to lie on the surface for all of the surface reconstruction algorithms discussed so far. In reality, the points can be a little off from the sampled surface. The noise introduced by these perturbations is often referred to as the Hausdorff noise because of the assumption that the Hausdorff distance between the surface S and its sample P is small. In the context of surface reconstruction, Dey and Goswami [DG06] first modeled this type of noise and presented an algorithm with provable guarantees.

A point set $P \subset \mathbb{R}^3$ is called an ϵ-noisy sample of a smooth closed surface $S \subset \mathbb{R}^3$ if conditions 1 and 2 below are satisfied. It is called an (ϵ, κ)-sample if additionally condition 3 is satisfied.

1. The orthogonal projection $\tilde{P}$ of P on S is an ϵ-sample of S.
2. $d(p, \tilde{p}) \leq \epsilon^2 f(\tilde{p})$.
3. For every point $p \in P$, its distance to its κ-th nearest point is at least $\epsilon f(\tilde{p})$.

The first two conditions say that the point sample is dense and sufficiently close to S. The third condition imposes some kind of relaxed version of local uniformity by considering the κ-th nearest neighbor instead of the nearest neighbor. Dey and Goswami [DG06] observe that, under the (ϵ, κ)-sampling condition, some of the Delaunay balls of tetrahedra in $\mathrm{Del}\, P$ remain almost as large as the Delaunay balls centering the poles in the noise-free case. The small Delaunay balls cluster near the sample points. See Figure 35.2.6. With this observation, they propose to separate out the 'big' Delaunay balls from the small ones by thresholding. Let $B(c, r)$ be a Delaunay ball of a tetrahedron $pqrs \in \mathrm{Del}\, P$ and let ℓ be the smallest among the distances of p, q, r, and s to their respective κ-th nearest neighbors. If $r > k\ell$ for a suitably chosen fixed constant k, the ball $B(c, r)$ is marked as big. After marking all such big Delaunay balls, the algorithm starts from any of the infinite tetrahedra and continues collecting any big Delaunay ball that has a

Preliminary version (July 19, 2017). To appear in the Handbook of Discrete and Computational Geometry, J.E. Goodman, J. O'Rourke, and C. D. Tóth (editors), 3rd edition, CRC Press, Boca Raton, FL, 2017.

positive power distance (intersects deeply) to any of the balls collected so far. As in the power crust algorithm, this process separates the inner big Delaunay balls from the outer ones. It is proved that, if the thresholds are chosen right and an appropriate sampling condition holds, the boundary ∂D of the union D of outer (or inner) big Delaunay balls is homeomorphic to S. To make the output a Delaunay subcomplex, the authors [DG06] suggest to collect the points $P' \subseteq P$ contained in the boundary of the big outer Delaunay balls and compute the restricted Delaunay triangulation of P' with respect to a smooth *skin surface* [Ede99] approximating ∂D. An even easier option which seems to work in practice is to take the restricted Delaunay triangulation $\mathrm{Del}\, P'|_{\partial D}$ which coincides with the boundary of the union of tetrahedra circumscribed by the big outer Delaunay balls.

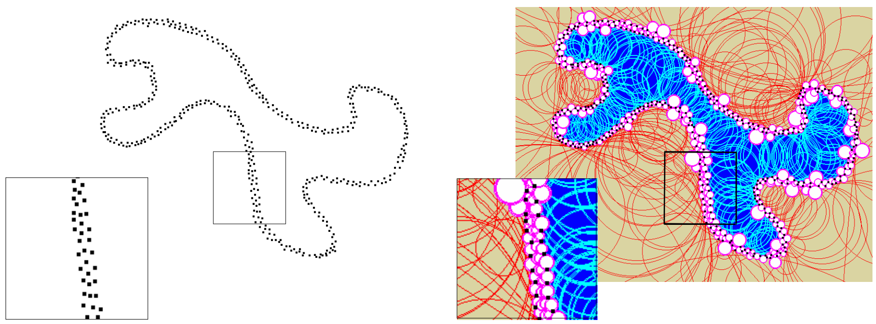

FIGURE 35.2.6
On left: Noisy point sample from a curve; on right: big and small Delaunay balls.

Mederos, Amenta, Vehlo, and Figueiredo [MAVF05] adapted the idea in [DG06] to the framework of power crust. They do not require the third uniformity condition, but instead rely on an input parameter c that needs to be chosen appropriately. Their algorithm starts by identifying the polar balls. A subset of the polar balls whose radii are larger than the chosen parameter c is selected. These "big" polar balls are further marked as inner and outer using the technique described before. Let $\mathbb{B}_O$ and $\mathbb{B}_I$ denote the set of these inner and outer polar balls respectively. Considering the polar balls as weighted points, the algorithm computes the power crust, that is, the facets in the power diagram of $\mathrm{Pow}\,(\mathbb{B}_O \cup \mathbb{B}_I)$ that separate a Voronoi cell corresponding to a ball in $\mathbb{B}_O$ from a cell corresponding to a ball in $\mathbb{B}_I$.

THEOREM 35.2.5

There exist $c > 0$, $\epsilon > 0$ so that if $P \subset \mathbb{R}^3$ is an ϵ-noisy sample of a smooth closed surface $S \subset \mathbb{R}^3$, the above algorithm with parameter c returns the power crust in $\mathrm{Pow}\,(\mathbb{B}_O \cup \mathbb{B}_I)$ *which is homeomorphic to S.*

SUMMARIZED RESULTS

The properties of the above discussed surface reconstruction algorithms are summarized in Table 35.2.1.

Preliminary version (July 19, 2017). To appear in the Handbook of Discrete and Computational Geometry,
J.E. Goodman, J. O'Rourke, and C. D. Tóth (editors), 3rd edition, CRC Press, Boca Raton, FL, 2017.

TABLE 35.2.1 Surface reconstruction algorithms.

ALGORITHM	SAMPLE	PROPERTIES	SOURCE
α-shape	uniform	α to be determined.	[EM94]
Crust	non-uniform	Theoretical guarantees from Voronoi structures, two Voronoi computations.	[AB99]
Cocone	non-uniform	Simplifies crust, single Voronoi computation with topological guarantee, detects undersampling.	[ACDL02] [DG03]
Natural Neighbor	non-uniform	Theoretical guarantees using Voronoi diagram and implicit functions.	[BC02]
Morse Flow	non-uniform	Draws upon the gradient flow of distance functions, flow complex introduces Steiner points, Wrap computes Delaunay subcomplex.	[DGRS08] [Ede03]
Power Crust	non-uniform	Watertight surface using power diagrams, introduces Steiner points.	[ACK01]
Tight Cocone	non-uniform	Watertight surface using Delaunay triangulation.	[DG03]
Peel	uniform	Theoretical guarantees for surfaces with boundaries.	[DLRW09]
Noise	uniform	Filters big Delaunay balls, assumes Hausdorff noise.	[DG06] [MAVF05]

OPEN PROBLEMS

All guarantees given by various surface reconstruction algorithms depend on the notion of dense sampling. Watertight surface algorithms can guarantee a surface without holes, but no theoretical guarantees exist under any type of undersampling. Also, non-smoothness and noise still need more investigation.

1. Design an algorithm that reconstructs nonsmooth surfaces under reasonable sampling conditions.

2. Design a surface reconstruction algorithm that handles noise more general than Hausdorff.

35.3 SHAPE RECONSTRUCTION

All algorithms discussed above are designed for reconstructing a shape of specific dimension from the samples. Thus, the curve reconstruction algorithms cannot handle samples from surfaces and the surface reconstruction algorithms cannot handle samples from curves. Therefore, if a sample is derived from shapes of mixed dimensions, i.e., both curves and surfaces in $\mathbb{R}^3$, none of the curve and surface reconstruction algorithms are adequate. General shape reconstruction algorithms should be able to handle any shape embedded in Euclidean spaces. However, this goal may be too ambitious, as it is not clear what would be a reasonable definition of dense samples for general shapes that are nonsmooth or nonmanifold. The

Preliminary version (July 19, 2017). To appear in the Handbook of Discrete and Computational Geometry, J.E. Goodman, J. O'Rourke, and C. D. Tóth (editors), 3rd edition, CRC Press, Boca Raton, FL, 2017.

ϵ-sampling condition would require infinite sampling in these cases. We therefore distinguish two cases: (i) smooth manifold reconstruction for which a computable sampling criterion can be defined, (ii) shape reconstruction for which it is currently unclear how a computable sampling condition could be defined to guarantee reconstruction. This leads to a different definition for the general shape reconstruction problem in the glossary below.

GLOSSARY

Shape reconstruction: Given a set of points $P \subseteq \mathbb{R}^m$, compute a shape that best approximates P.

Manifold reconstruction: Compute a piecewise-linear approximation to a manifold M, given a sample P of M.

MAIN RESULTS

Shape reconstruction: Not many algorithms are known to reconstruct shapes. The definition of α-shapes is general enough to be applicable to shape reconstruction. In Figure 35.2.2, the α-shape reconstructs a shape in $\mathbb{R}^2$ which is not a manifold. Similarly, it can reconstruct curves, surfaces and solids and their combinations in three dimensions. Melkemi [Mel97] proposed $\mathcal{A}$-shapes that can reconstruct shapes in $\mathbb{R}^2$. Its class of shapes includes α-shapes. Given a set of points P in $\mathbb{R}^2$, a member in this class of shapes is identified with another finite set $\mathcal{A} \subseteq \mathbb{R}^2$. The $\mathcal{A}$-shape of S is generated by edges that connect points $p, q \in P$ if there is a circle passing through p, q and a point in $\mathcal{A}$, and all other points in $P \cup \mathcal{A}$ lie outside the circle. The α-shape is a special case of $\mathcal{A}$-shapes where $\mathcal{A}$ is the set of all points on Voronoi edges that span empty circles with points in P. The crust is also a special case of $\mathcal{A}$-shape where $\mathcal{A}$ is the set of Voronoi vertices.

For points sampled from compact subsets of $\mathbb{R}^m$, Chazal, Cohen-Steiner, and Lieutier [CCL09] developed a sampling theory that guarantees a weaker topological equivalence. The output is guaranteed to be homotopy equivalent (instead of homeomorphic) to the sampled space when the input satisfies an appropriate sampling condition.

Manifold reconstruction: When the sample P is derived from a smooth manifold M embedded in some Euclidean space $\mathbb{R}^m$, the curse of dimensionality makes reconstruction harder. Furthermore, unlike in two and three dimensions, the restricted Delaunay triangulation $\mathrm{Del}\, P|_M$ may not have underlying space homeomorphic to M no matter how dense P is. As observed by Cheng, Dey, and Ramos [CDR05], the normal spaces of restricted Delaunay simplices in dimensions four and above can be arbitrarily oriented instead of aligning with those of M. This leads to the topological discrepancy between $\mathrm{Del}\, P|_M$ and M as observed by Boissonnat, Guibas, and Oudot [BGO09].

To overcome the difficulty with the topological discrepancy, Cheng, Dey, and Ramos [CDR05] proposed an algorithm that utilizes the restricted Delaunay triangulation of a weighted version $\widehat{P}$ of P. Using the concept of weight pumping in sliver exudations [CDE+00], they prove that $\mathrm{Del}\, \widehat{P}|_M$ becomes homeomorphic to M under appropriate weight assignments if P is sufficiently dense. Since it is possible for P to be dense for a curve and a surface simultaneously already in

Preliminary version (July 19, 2017). To appear in the Handbook of Discrete and Computational Geometry, J.E. Goodman, J. O'Rourke, and C. D. Tóth (editors), 3rd edition, CRC Press, Boca Raton, FL, 2017.

three dimensions [DGGZ03, BGO09], one needs some uniformity condition on P to disambiguate the multiple possibilities. Cheng et al. assumed P to be locally uniform, or ϵ-local. The algorithm in [CDR05] assigns some appropriate weights to P and then computes a Delaunay sub-complex of Del $\widehat{P}$ restricted to a space consisting of the union of cocone-like spaces around each point in P. This restricted complex is shown to be homeomorphic and Hausdorff close to M. Similar result was later obtained by Boissonnat, Guibas, and Oudot [BGO09] using the witness complex [CdS04] and weights.

The above algorithms compute complexes of dimensions equal to that of the embedding dimension $\mathbb{R}^m$. Consequently, they have complexity exponential in m. For example, it is known that Delaunay triangulation of n points in $\mathbb{R}^m$ has size $\Theta(n^{\lceil \frac{m}{2} \rceil})$. Usually, the intrinsic dimension k of the k-manifold M is small compared to m. Therefore, manifold reconstruction algorithms that replace the dependence on m with that on k have better time and size complexity. Boissonnat and Ghosh [BG14] achieve this goal. They compute the restricted Delaunay complex with respect to the approximate tangent spaces of M at each sample point. Each such individual tangent complex may not agree globally to form a triangulation of M. Nonetheless, Boissonnat and Ghosh [BG14] show that under appropriate weight assignment and sampling density, the collection of all tangent complexes over all sample points become consistent to form a reconstruction of M. The running time of this algorithm is exponential in k, but linear in m.

All known algorithms for reconstructing manifolds in high dimensions use data structures that are computationally intensive. A practical algorithm with geometric and topological guarantee is still elusive.

OPEN PROBLEMS

1. Design an algorithm that reconstructs k-manifolds in $\mathbb{R}^m$, $m \geq 4$, with guarantees and is practical.
2. Reconstruct shapes with guarantees.

35.4 SOURCES AND RELATED MATERIAL

BOOKS/SURVEYS

[Dey07]: Curve and surface reconstruction: Algorithms with mathematical analysis.

[Ede98]: Shape reconstruction with Delaunay complex.

[OR00]: Computational geometry column 38 (recent results on curve reconstruction).

[MSS92]: Surfaces from contours.

[MM98]: Interpolation and approximation of surfaces from 3D scattered data points.

Preliminary version (July 19, 2017). To appear in the Handbook of Discrete and Computational Geometry, J.E. Goodman, J. O'Rourke, and C. D. Tóth (editors), 3rd edition, CRC Press, Boca Raton, FL, 2017.

RELATED CHAPTERS

REFERENCES

[AM02] E. Althaus and K. Mehlhorn. Traveling salesman-based curve reconstruction in polynomial time. *SIAM J. Comput.*, 31: 27–66, 2002.

[AB99] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete Comput. Geom.*, 22:481–504, 1999.

[ACDL02] N. Amenta, S. Choi, T.K. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. *Internat. J. Comput. Geom. Appl.*, 12:125–141, 2002.

[ABE98] N. Amenta, M. Bern, and D. Eppstein. The crust and the β-skeleton: Combinatorial curve reconstruction. *Graphical Models and Image Processing*, 60:125–135, 1998.

[ACK01] N. Amenta, S. Choi, and R.K. Kolluri. The power crust, union of balls, and the medial axis transform. *Comput. Geom.*, 19:127–153, 2001.

[Att98] D. Attali. r-regular shape reconstruction from unorganized points. *Comput. Geom.*, 10:239–247, 1998.

[AH85] D. Avis and J. Horton. Remarks on the sphere of influence graph. *Ann. New York Acad. Sci.*, 440:323–327, 1985.

[BB97] F. Bernardini and C.L. Bajaj. Sampling and reconstructing manifolds using α-shapes. In *Proc. 9th Canad. Conf. Comput. Geom.*, pages 193–198, 1997.

[BMR+99] F. Bernardini, J. Mittleman, H. Rushmeier, C. Silva, and G. Taubin. The ball-pivoting algorithm for surface reconstruction. *IEEE Trans. Visual. Comput. Graphics*, 5:349–359, 1999.

[Boi84] J.-D. Boissonnat. Geometric structures for three-dimensional shape representation. *ACM Trans. Graphics*, 3:266–286, 1984.

[BC02] J.-D. Boissonnat and F. Cazals. Smooth surface reconstruction via natural neighbor interpolation of distance functions. *Comput. Geom.*, 22:185–203, 2002.

[BG93] J.-D. Boissonnat and B. Geiger. Three-dimensional reconstruction of complex shapes based on the Delaunay triangulation. In *Proc. Biomed. Image Process. Biomed. Visualization*, vol. 1905 of SPIE, pages 964–975, 1993.

[BG14] J.-D. Boissonnat and A. Ghosh. Manifold reconstruction using tangential Delaunay complexes. *Discrete Comput. Geom.*, 51:221–267, 2014.

[BGO09] J.-D. Boissonnat, L.J. Guibas, and S.Y. Oudot. Manifold reconstruction in arbitrary dimensions using witness complexes. *Discrete Comput. Geom.*, 42:37–70, 2009.

[CdS04] G. Carlsson and V. de Silva. Topological estimation using witness complexes. In *Proc. Eurographics Conf. Point-Based Graphics*, pages 157–166, ACM Press, 2004.

[CCL09] F. Chazal, D. Cohen-Steiner, and A. Lieutier. A sampling theory for compact sets in Euclidean space. *Discrete Comput. Geom.*, 41:461–479, 2009.

Preliminary version (July 19, 2017). To appear in the Handbook of Discrete and Computational Geometry,
J.E. Goodman, J. O'Rourke, and C. D. Tóth (editors), 3rd edition, CRC Press, Boca Raton, FL, 2017.

[CDE+00] S.-W. Cheng, T.K. Dey, H. Edelsbrunner, M.A. Facello, and S.-H. Teng. Sliver exudation. *J. ACM*, 47:883–904, 2000.

[CDR05] S.-W. Cheng, T.K. Dey, and E.A. Ramos. Manifold reconstruction from point samples. In *Proc. 16th ACM-SIAM Sympos. Discrete Algorithms*, pages 1018–1027, 2005.

[CJL17] S.-W. Cheng, J. Jin, and M.-K. Lau. A fast and simple surface reconstruction algorithm. *ACM Trans. Algorithms*, 13:27, 2017.

[CL96] B. Curless and M. Levoy. A volumetric method for building complex models from range images. In *Proc. 23rd Conf. Comput. Graphics Interactive Tech. (SIGGRAPH)*, pages 306–312, ACM Press, 1996.

[Dey07] T.K. Dey. *Curve and Surface Reconstruction: Algorithms with Mathematical Analysis.* Cambridge University Press, 2007.

[DG03] T.K. Dey and J. Giesen. Detecting undersampling in surface reconstruction. In B. Aronov, S. Basu, J. Pach, and M. Sharir, editors, *Discrete and Computational Geometry: The Goodman-Pollack Festschrift*, vol. 25 of *Algorithms Combin.*, pages 329–345, Springer, Berlin, 2003.

[DGGZ03] T.K. Dey, J. Giesen, S. Goswami, and W. Zhao. Shape dimension and approximation from samples. *Discrete Comput. Geom.*, 29:419–434, 2003.

[DGRS08] T.K. Dey, J. Giesen, E.A. Ramos, and B. Sadri. Critical points of the distance to an epsilon-sampling on a surface and flow-complex-based surface reconstruction. *Internat. J. Comput. Geom. Appl.*, 18:29–61 , 2008.

[DG03] T.K. Dey and S. Goswami. Tight cocone: A water-tight surface reconstructor. *J. Comput. Inf. Sci. Eng*, 3:302–307, 2003.

[DG06] T.K. Dey and S. Goswami. Provable surface reconstruction from noisy samples. *Comput. Geom.*, 35:124–141, 2006.

[DK99] T.K. Dey and P. Kumar. A simple provable curve reconstruction algorithm. In *Proc. 10th ACM-SIAM Sympos. Discrete Algorithms*, pages 893–894, 1999.

[DMR00] T.K. Dey, K. Mehlhorn, and E.A. Ramos. Curve reconstruction: Connecting dots with good reason. *Comput. Geom.*, 15:229–244, 2000.

[DLRW09] T.K. Dey, K. Li, E.A. Ramos, and R. Wenger. Isotopic reconstruction of surfaces with boundaries. *Computer Graphics Forum*, 28:1371–1382, 2009.

[Ede98] H. Edelsbrunner. Shape reconstruction with Delaunay complex. In *Proc. 3rd Latin American Sympos. Theoret. Inf.*, vol. 1380 of *LNCS*, pages 119–132, Springer, Berlin, 1998.

[Ede99] H. Edelsbrunner. Deformable smooth surface design. *Discrete Comput. Geom.*, 21:87–115, 1999.

[Ede03] H. Edelsbrunner. Surface reconstruction by wrapping finite point sets in space. In B. Aronov, S. Basu, J. Pach, and M. Sharir, editors, *Discrete and Computational Geometry: The Goodman-Pollack Festschrift*, vol. 25 of *Algorithms Comb.*, pages 379–404, Springer, Berlin, 2003.

[EKS83] H. Edelsbrunner, D.G. Kirkpatrick, and R. Seidel. On the shape of a set of points in the plane. *IEEE Trans. Inf. Theory*, 29:551–559, 1983.

[EM94] H. Edelsbrunner and E.P. Mücke. Three-dimensional alpha shapes. *ACM Trans. Graphics*, 13:43–72, 1994.

[ES97] H. Edelsbrunner and N.R. Shah. Triangulating topological spaces. *Internat. J. Comput. Geom. Appl.*, 7:365–378, 1997.

Preliminary version (July 19, 2017). To appear in the Handbook of Discrete and Computational Geometry, J.E. Goodman, J. O'Rourke, and C. D. Tóth (editors), 3rd edition, CRC Press, Boca Raton, FL, 2017.

[FG95] L.H. de Figueiredo and J. de Miranda Gomes. Computational morphology of curves. *Visual Computer*, 11:105–112, 1995.

[FKU77] H. Fuchs, Z.M. Kedem, and S.P. Uselton. Optimal surface reconstruction from planar contours. *Commun. ACM*, 20:693–702, 1977.

[FR01] S. Funke and E.A. Ramos. Reconstructing a collection of curves with corners and endpoints. In *Proc. 12th ACM-SIAM Sympos. Discrete Algorithms*, pages 344–353, 2001.

[FR02] S. Funke and E.A. Ramos. Smooth-surface reconstruction in near-linear time. In *13th ACM-SIAM Sympos. Discrete Algorithms*, pages 781–790, 2002.

[Gie00] J. Giesen. Curve reconstruction, the traveling salesman problem and Menger's theorem on length. *Discrete Comput. Geom.*, 24:577–603, 2000.

[GJ08] J. Giesen and M. John. The flow complex: A data structure for geometric modeling. *Comput. Geom.*, 39:178–190, 2008.

[GCA13] S. Giraudot, D. Cohen-Steiner, and P. Alliez. Noise-adaptive shape reconstruction from raw point sets. *Comput. Graphics Forum*, 32:229–238, 2013.

[GOS96] C. Gitlin, J. O'Rourke, and V. Subramanian. On reconstruction of polyhedra from slices. *Internat. J. Comput. Geom. Appl.*, 6:103–112, 1996.

[GS01] C.M. Gold and J. Snoeyink. Crust and anti-crust: A one-step boundary and skeleton extraction algorithm. *Algorithmica*, 30:144–163, 2001.

[GKS00] M. Gopi, S. Krishnan, and C.T. Silva. Surface reconstruction based on lower dimensional localized Delaunay triangulation. *Computer Graphics Forum*, 19:467–478, 2000.

[HDD$^+$92] H. Hoppe, T.D. DeRose, T. Duchamp, J. McDonald, and W. Stützle. Surface reconstruction from unorganized points. In *Proc. 19th Conf. Comput. Graphics Interactive Tech. (SIGGRAPH)*, pages 71–78, ACM Press, 1992.

[KR85] D.G. Kirkpatrick and J.D. Radke. A framework for computational morphology. In G.T. Toussaint, editor, *Computational Geometry*, vol. 2 of *Mach. Intell. Pattern Recog.*, pages 217–248, North-Holland, Amsterdam, 1985.

[MAVF05] B. Mederos, N. Amenta, L. Velho, and H. de Figueiredo. Surface reconstruction from noisy point clouds. In *Proc. 3rd Eurographics Sympos. Geom. Processing*, pages 53–62, 2005.

[Mel97] M. Melkemi. $\mathcal{A}$-shapes of a finite point set. Correspondence in *Proc. 13th Sympos. Comput. Geom.*, pages 367–369, ACM Press, 1997.

[MM98] R. Mencl and H. Müller. Interpolation and approximation of surfaces from three-dimensional scattered data points. In *State of the Art Reports (Eurographics)*, pages 51–67, 1998.

[MSS92] D. Meyers, S. Skinner, and K. Sloan. Surfaces from contours. *ACM Trans. Graphics*, 11:228–258, 1992.

[OR00] J. O'Rourke. Computational geometry column 38. *Internat. J. Comput. Geom. Appl.*, 10:221–223, 2000. Also in *SIGACT News*, 31:28–30, 2000.

[RS07] E.A. Ramos and B. Sadri. Geometric and topological guarantees for the WRAP reconstruction algorithm. In *Proc. 18th ACM-SIAM Sympos. Discrete Algorithms*, pages 1086–1095, 2007.

Preliminary version (July 19, 2017). To appear in the Handbook of Discrete and Computational Geometry, J.E. Goodman, J. O'Rourke, and C. D. Tóth (editors), 3rd edition, CRC Press, Boca Raton, FL, 2017.