

Isotopic Reconstruction of Surfaces with Boundaries

Tamal K. Dey¹ Kuiyu Li¹ Edgar A. Ramos² Rephael Wenger¹

¹The Ohio State University, Columbus OH, USA

²Escuela de Matemáticas, Universidad Nacional de Colombia, Sede Medellín, Medellín, Colombia

Abstract

We present an algorithm for the reconstruction of a surface with boundaries (including a non-orientable one) in three dimensions from a sufficiently dense sample. It is guaranteed that the output is isotopic to the unknown sampled surface. No previously known algorithm guarantees isotopic or homeomorphic reconstruction of surfaces with boundaries. Our algorithm is surprisingly simple. It ‘peels’ slivers greedily from an α -complex of a sample of the surface. No other post-processing is necessary. We provide several experimental results from an implementation of our basic algorithm and also a modified version of it.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Surface Reconstruction—

1. Introduction

This paper is concerned with the design of a provably correct algorithm that can reconstruct surfaces with *boundaries* in three dimensions. In the last decade a number of algorithms for the problem of surface reconstruction have been proposed [AB99, ACDL02, ACSTD07, BC00, FCOS05, HDD*92, JWS08, KBH06, OBA*03, PKKG03, WOK05]. Among them, the ones that come with theoretical guarantees [AB99, ACDL02, BC00, Dey07] assume that the sampled surface is smooth and *closed* (compact and no boundary). The proofs and the algorithms fail if boundaries are allowed.

For a provably correct reconstruction algorithm, typically one looks for guarantees on the geometric and topological similarities between the output and the sampled surface. Geometric guarantees generally mean small Hausdorff distance between the sampled surface $\Sigma \subset \mathbb{R}^3$ and the output. Topological guarantee, however, may be of homotopy equivalence, homeomorphism, or isotopy—listed in the increasing order of topological similarity. It is desirable that the output be isotopic to Σ . *Isotopy* means that there is a smooth deformation of \mathbb{R}^3 that brings the output to Σ while maintaining a homeomorphism between the two all the time.

The case of closed smooth surfaces is now well understood. There exist established algorithms and software that can reconstruct such surfaces with isotopy guarantee from point data [AB99, ACDL02] even if they are contaminated

with noise [Dey07]. However, the important case of *surfaces with boundaries* is less understood. The theory of compact reconstruction put forward by Chazal, Cohen-Steiner, and Lieutier [CCSL06] applies to such surfaces. However, this result only guarantees a homotopy equivalence between the sampled surface and the output. In fact, the algorithm which constructs α -complexes may return a single point for a disk. In this paper we show that a very simple algorithm can reconstruct all smooth surfaces $\Sigma \subset \mathbb{R}^3$, with or without boundary, while guaranteeing an isotopy and a small Hausdorff distance between the output and Σ . The algorithm computes an α -complex and then ‘peels’ tetrahedra from it greedily. No other post-processing is necessary. We argue mainly about the isotopy since Hausdorff distance claim follows from dense sampling.

For closed surface reconstructions, a fundamental property used for proofs is that a closed surface is homeomorphic (and isotopic) to the restricted Delaunay triangulation of a sufficiently dense point sample. (See Section 2.1 for a definition of the restricted Delaunay triangulation.) As shown in Section 4.1 and illustrated in Figure 3, this is not true for a surface Σ with boundary. To overcome this problem we use a collar extension of Σ . We show that a restricted Delaunay triangulation of a sufficiently dense point sample of Σ is isotopic to this collar extension. We emphasize that the collar is only needed for the proof of correctness, it does not need to be sampled or used in any way by the algorithm.

Reconstruction algorithms such as Crust and Co-

cone [AB99, ACDL02] identify candidate triangles for the reconstructed surface and then use a prune-and-walk approach to extract a manifold. Such an approach fails when boundaries are present, since the pruning repeatedly removes boundary triangles until no triangles are left. We present a ‘sliver peeling’ approach based on the following idea. We prove the existence of a *canonical peeling sequence* of tetrahedra from an α -complex induced by a *deformation retraction* of the associated union of α -balls. For a closed surface Σ , this sequence peels the α -complex to the restricted Delaunay triangulation which is known to be isotopic to Σ . Even though one cannot compute the canonical peeling sequence in absence of Σ , we show that a greedy peeling produces an output isotopic to that produced by the canonical peeling sequence. For surfaces with boundaries, the argument remains essentially the same except that now one has to argue with the restricted Delaunay triangulation of a collar extension.

Another problem that arises in reconstructing surfaces with boundaries is the existence of non-unique solutions. For instance, it can be impossible to distinguish between the point sample of a sphere with a small disk removed and a sphere which happens to not have any sample points in a small disk. We use an input parameter α which allows us to distinguish between boundary regions and small regions which have no sample points.

One noteworthy feature of our algorithm is that it can reconstruct non-orientable surfaces (necessarily open ones as they are embedded in \mathbb{R}^3). Since there is no global orientation of normals for such surfaces, it is difficult to extract such a surface from a complex by any consistent walk. The technique of sliver ‘peeling’ resolves this problem.

Our algorithm works nicely for dense uniform samples in practice as the theory predicts. However, it faces problems when the sample is not uniform. We extend our basic ‘peeling’ algorithm to handle non-uniform samples and show experimental results for this modification. Our algorithm is not geared to handle noise in data.

2. Preliminaries

We set up some basic notations and concepts that are needed to describe the reconstruction algorithm.

2.1. Complexes, feature size, and sampling

Voronoi and Delaunay Complexes: For a point set $P \subset \mathbb{R}^3$, we use $\text{Vor } P$ and $\text{Del } P$ to denote the Voronoi and Delaunay complexes of P . The Voronoi cell of p is V_p . The Voronoi face dual to a Delaunay simplex $\sigma \in \text{Del } P$ is denoted V_σ . The *restricted Delaunay complex* of P with respect to a topological space $\mathbb{X} \subseteq \mathbb{R}^3$ is $\text{Del}_{\mathbb{X}}(P) = \{\sigma \in \text{Del } P : V_\sigma \cap \mathbb{X} \neq \emptyset\}$.

Alpha-Complex: For $\alpha > 0$, the α -ball of p , denoted $B(p, \alpha)$, is the closed ball centered at p and of radius α . The

union of α -balls for a set P is

$$\mathcal{B}_\alpha(P) = \bigcup_{p \in P} B(p, \alpha).$$

The α -complex $\mathcal{C}_\alpha(P)$ of P is the restricted Delaunay complex of P with respect to $\mathcal{B}_\alpha(P)$. That is,

$$\mathcal{C}_\alpha(P) = \{\sigma \in \text{Del } P : V_\sigma \cap \mathcal{B}_\alpha(P) \neq \emptyset\}.$$

It is well-known that $\mathcal{B}_\alpha(P)$ and $\mathcal{C}_\alpha(P)$ have the same homotopy type [Ede95]. If P is a sufficiently dense sample from a shape Σ , it is known that $\mathcal{B}_\alpha(P)$ and $\mathcal{C}_\alpha(P)$ have topological similarity to Σ for appropriate α [CCSL06, CSEH05, NSW08]. In particular, $\mathcal{B}_\alpha(P)$ and $\mathcal{C}_\alpha(P)$ become homotopy equivalent to Σ when the sample is sufficiently dense [CCSL06, NSW08], a fact we need for our results.

Local feature size: We consider a smooth surface Σ with smooth boundaries $\partial\Sigma$. The surface normal n_x to Σ at x is then well-defined in the surface interior and also on the boundary by taking the *limit*. At a boundary point $y \in \partial\Sigma$, there is also an open half-circle of normal directions with the two orientations of the surface normal n_y as the *limit points*.

The *local feature size function* is a measure of the local level of detail. For a closed surface, it is defined as the distance, $d(x, \mathcal{M})$, to the medial axis \mathcal{M} of the surface [AB99]. For a smooth surface with smooth boundaries, this is still well-defined but does not capture the required minimum sampling (to see this, consider for example a flat disk). We consider \mathcal{M}' , the medial axis of the boundary $\partial\Sigma$ separately and define:

$$\text{lfs}(x) = \min\{d(x, \mathcal{M}), (d(x, \mathcal{M}') + d(x, \partial\Sigma)) / 2\}.$$

Note that the \mathcal{M}' can intersect the surface Σ and so $d(x, \mathcal{M}')$ may be zero. However, if $d(x, \mathcal{M}')$ is zero, then x is ‘‘far’’ from $\partial\Sigma$ and so $\text{lfs}(x)$ is still non-zero.

With the above definition of lfs , Proposition 1 extends three essential properties of closed surfaces to surfaces with boundaries. For a triangle t we use n_t to denote its normal. Unless specified otherwise we measure angles between vectors by the acute angle made by their supporting lines. The notation $O(\epsilon)$ means a value that is less than $c\epsilon$ for some constant $c > 0$ when ϵ is sufficiently small.

Proposition 1 Let $\epsilon \leq 1$ and $x, y, z \in \Sigma$. If $\|x - y\| = O(\epsilon)\text{lfs}(x)$ and the circumradius of $t = \triangle xyz$ is at most $O(\epsilon)\text{lfs}(x)$, then:

- (i) (Normal variation is small) $\angle n_x n_y = O(\epsilon)$,
- (ii) (Short edges are close to tangent) $\angle n_x \vec{xy} \geq \pi/2 - O(\epsilon)$,
- (iii) (Small triangles are almost parallel to the surface) $\angle n_t n_x = O(\epsilon)$.

Essentially the same proof [AB99, ACDL02] applies here as for the same properties for closed surfaces once (i) is established. The proof of (i) can be carried out along the lines presented in [CDR05]. Although the definition of lfs in [CDR05] differs from the one used here, for a point $x \in \Sigma$,

$\text{lfs}(x)$ remains smaller than the medial balls touching Σ at x which is essential for the proof to go through. See appendix for the proof.

Sampling condition: The standard sampling condition requires that for any $x \in \Sigma$, $B(x, \text{lfs}(x))$ contains a sample [AB99]. In the presence of boundaries, this is not sufficient for correct reconstruction. There is no unique tightening of the condition. We use a global measure $\rho = \inf_{x \in \Sigma} \text{lfs}(x)$ and define P to be an ε -sample of Σ if $B(x, \varepsilon\rho)$ contains a point in P for any $x \in \Sigma$. We do not require the boundaries to be sampled by points lying on them. Requiring such a sampling of the boundaries would make the problem easier but less realistic.

2.2. Peeling tetrahedra

Our idea is to compute a subcomplex \mathcal{K} of an alpha complex $\mathcal{C}_\alpha(P)$ for sufficiently small α and then peel tetrahedra from it to arrive at a triangulated surface. The existence and computation of the peeling order are justified by a deformation retraction of the ball union $\mathcal{B}_\alpha(P)$. Before describing this retraction, we elaborate on tetrahedra peeling and observe some of its key properties.

Let $\sigma \in \mathcal{K}$ be a tetrahedron with an edge e and t_1, t_2 be two triangles of σ incident to e . We say σ is *peelable* by e if no triangles other than t_1 and t_2 are incident to e in \mathcal{K} . A new complex \mathcal{K}' is obtained by removing the collection of simplices $\{\sigma, t_1, t_2, e\}$ from \mathcal{K} . In that case we say \mathcal{K}' is obtained by *peeling* e from \mathcal{K} and write $\mathcal{K} \xrightarrow{e} \mathcal{K}'$. A peeling takes out two triangles of a tetrahedron while leaving the other two. If \mathcal{K}' is obtained from \mathcal{K} by peeling a sequence edges e_1, e_2, \dots, e_k , we write $\mathcal{K} \xrightarrow{e_1, e_2, \dots, e_k} \mathcal{K}'$. A peeling does not delete any vertex and therefore the vertex set of \mathcal{K} and \mathcal{K}' remains the same.

We need one more definition before stating our generic results on peeling. Two edge sequences $\{e_i\}$ and $\{g_i\}$ by which a complex \mathcal{K} is peeled are called *compatible* if for all pair of edges $e \in \{e_i\}$ and $g \in \{g_i\}$ where e and g peel the same tetrahedron in \mathcal{K} , either $e = g$ or e and g are vertex disjoint. Essentially two compatible peeling sequences peel a tetrahedron either at the same edge or at two vertex disjoint edges if both of them peel it. Two such sequences may differ in length. Figure 1 illustrates how compatibility ensures isotopy. The peeling of vertex disjoint edges e_1 and e_3 in sequences $\{e_2, e_1\}$ and $\{e_3, e_2\}$ leads to isotopy.

Our main observation is Proposition 2 below which says that if a simplicial 3-complex is peeled by two compatible sequences both of which remove *all* tetrahedra, the resulting 2-complexes are necessarily isotopic. Notice that we do not require any extra condition on the input 3-complex for this proposition. Its proof by induction uses Proposition 3 whose proof appears in the appendix.




Figure 1: The sequences $\{e_2, e_1\}$ and $\{e_3, e_2\}$ are compatible, but $\{e_4, e_1\}$ is not compatible with either of them.

Proposition 2 Suppose \mathcal{K}_0 is a simplicial 3-complex with $k \geq 0$ tetrahedra. If $\mathcal{K}_0 \xrightarrow{e_1, e_2, \dots, e_k} \mathcal{K}_1$ and $\mathcal{K}_0 \xrightarrow{g_1, g_2, \dots, g_k} \mathcal{K}_2$ for two compatible sequences of edges $\{e_1, e_2, \dots, e_k\}$ and $\{g_1, g_2, \dots, g_k\}$, then \mathcal{K}_1 and \mathcal{K}_2 are isotopic.

Proof We prove by induction on k . For $k = 1$, \mathcal{K}_0 has a single tetrahedron which is peelable by e_1 and g_1 . If $e_1 = g_1$, $\mathcal{K}_1 = \mathcal{K}_2$. If $e_1 \neq g_1$, they are opposite edges of the single tetrahedron in \mathcal{K}_0 . Clearly, peeling by e_1 and g_1 gives two isotopic complexes. We now assume that for any $k' < k$ the proposition is true and verify it for k . Let $\mathcal{K}_0 \xrightarrow{g_1} \mathcal{K}'$. Since $\{g_1\}$ and $\{e_1, e_2, \dots, e_k\}$ are compatible, by Proposition 3 below there is a sequence $\{e'_1, e'_2, \dots, e'_{k-1}\}$ of edges compatible with $\{e_1, \dots, e_k\}$ so that $\mathcal{K}' \xrightarrow{e'_1, e'_2, \dots, e'_{k-1}} \mathcal{K}_3$ where \mathcal{K}_3 is isotopic to \mathcal{K}_1 . By assumption $\mathcal{K}' \xrightarrow{g_2, g_3, \dots, g_k} \mathcal{K}_2$, and also the sequences $\{g_2, g_3, \dots, g_k\}$ and $\{e'_1, e'_2, \dots, e'_{k-1}\}$ are compatible since both of them are compatible with $\{e_1, e_2, \dots, e_k\}$. Then, by inductive hypothesis \mathcal{K}_2 and \mathcal{K}_3 are isotopic which implies that \mathcal{K}_2 and \mathcal{K}_1 are isotopic. \square

Proposition 3 Suppose \mathcal{K}_0 is a simplicial 3-complex with $\mathcal{K}_0 \xrightarrow{e_1, \dots, e_k} \mathcal{K}_1$ for edges e_1, \dots, e_k . Let e be an edge compatible with e_1, \dots, e_k so that $\mathcal{K}_0 \xrightarrow{e} \mathcal{K}'$ and e peels a tetrahedron that is peeled by some edge in the sequence e_1, \dots, e_k . Then, there is a sequence e'_1, \dots, e'_{k-1} compatible with e_1, \dots, e_k so that $\mathcal{K}' \xrightarrow{e'_1, \dots, e'_{k-1}} \mathcal{K}_2$ and \mathcal{K}_1 is isotopic to \mathcal{K}_2 .

3. Closed surface

For simplicity first we assume that the sampled surface has no boundary and then describe the modifications needed to accommodate boundaries. Proposition 2 motivates the following approach:

- a. Consider an α -complex \mathcal{K} spanning the input point set P .

- b. Prove that there is a *canonical edge sequence* which peel all tetrahedra from \mathcal{K} and the resulting complex is the restricted Delaunay triangulation $\text{Del}_\Sigma(P)$.
- c. Unfortunately, this canonical sequence is impossible to compute in absence of Σ . So, algorithmically find another peeling sequence which peels all tetrahedra from \mathcal{K} .
- d. Prove that the peeling sequence adopted by the algorithm is compatible with the canonical sequence. Appeal to Proposition 2 to claim that the resulting complex is isotopic to $\text{Del}_\Sigma(P)$.

3.1. Canonical edge sequence

To show the existence of a canonical edge sequence, we define a deformation retraction $\text{ret}(x, t)$ on \mathbb{R}^3 which retracts $\mathcal{B}_\alpha(P)$ to the surface Σ . For any non medial axis point $x \in \mathbb{R}^3$, let \bar{x} be its closest point on Σ . Define

$$\text{ret}(x, t) = \bar{x} + (1 - t) \cdot (x - \bar{x}) \text{ for } t \in [0, 1].$$

If α is sufficiently small, this deformation retraction is a witness to the homotopy equivalence between \mathcal{B}_α and Σ . This deformation retraction is reminiscent of the flow function used in [Ede03, GJ02]. However, a key difference is that we consider distances to the surface Σ instead of distances to the sample P .

We are interested in the interpretation of the retraction $\text{ret}(x, t)$ in the context of sub-complexes of $\text{Del}P$. We define

$$\mathcal{B}_{\alpha,t}(P) = \{\text{ret}(x, t) : x \in \mathcal{B}_\alpha(P)\} \text{ and } \mathcal{C}_{\alpha,t}(P) = \text{Del}_{\mathcal{B}_{\alpha,t}(P)}(P).$$

Observe that $\mathcal{B}_\alpha(P) = \mathcal{B}_{\alpha,0}(P)$ retracts to $\Sigma = \mathcal{B}_{\alpha,1}(P)$. In the dual α -complex $\mathcal{C}_\alpha(P) = \mathcal{C}_{\alpha,0}(P)$ retracts to the restricted Delaunay triangulation $\text{Del}_\Sigma(P) = \mathcal{C}_{\alpha,1}(P)$. Let $\sigma_1, \dots, \sigma_n$ denote the sequence of simplices removed from $\mathcal{C}_{\alpha,t}(P)$ as $\mathcal{C}_\alpha(P) = \mathcal{C}_{\alpha,0}(P)$ transforms to $\text{Del}_\Sigma(P) = \mathcal{C}_{\alpha,1}(P)$. We call $\sigma_1, \dots, \sigma_n$ the *canonical simplex sequence* for $\mathcal{C}_\alpha(P)$. We argue that the subsequence of edges in this simplex sequence indeed peels all simplices in the sequence and thus form a canonical edge sequence we are looking for.

Let us look at how $\partial\mathcal{B}_{\alpha,t}(P)$ sweeps over the Voronoi diagram. When $\partial\mathcal{B}_{\alpha,t}(P)$ sweeps over a Voronoi vertex, edge, or a facet completely, their dual simplices get removed from the restricted triangulation $\mathcal{C}_{\alpha,t}(P)$. If all Voronoi edges and facets intersecting $\mathcal{B}_{\alpha,t}(P)$ intersect its boundary $\partial\mathcal{B}_{\alpha,t}(P)$ transversally, they go out of $\mathcal{B}_{\alpha,t}(P)$ only when a Voronoi vertex is swept over. This means $\mathcal{C}_{\alpha,t}(P)$ only changes when a tetrahedron is removed. Now let us look at how $\partial\mathcal{B}_{\alpha,t}(P)$ sweeps over a Voronoi vertex v . In Figure 2, the vertex v is about to be swept when the dual tetrahedron σ is peeled. It is proved in Proposition 6 that σ is a *sliver* tetrahedron (see section 3.3 for definition) if α is small. In particular, this means two of the Voronoi edges incident to v make large angle close to π with \vec{v} and the other two make small angle close to 0 with it. As a result when v is swept over, exactly two of the Voronoi edges, say e_1 and e_2 , incident to v leave $\mathcal{B}_{\alpha,t}(P)$. Interpreting dually, the tetrahedron σ is peeled from $\mathcal{C}_{\alpha,t}(P)$




Figure 2: Sweeping a Voronoi vertex.

at the edge, say e , dual to the Voronoi facet containing e_1 and e_2 . The subsequence of edges like e from the canonical simplex sequence form the desired canonical edge sequence.

The above argument dwells on two facts: (i) Voronoi facets and edges intersect $\partial\mathcal{B}_{\alpha,t}(P)$ transversally, and (ii) all tetrahedra in $\mathcal{C}_\alpha(P)$ are slivers. We show these two facts.

3.2. Transversal intersections

To talk about transversal intersections between Voronoi edges, facets and $\partial\mathcal{B}_{\alpha,t}(P)$, we need to define a notion of normals to $\partial\mathcal{B}_{\alpha,t}(P)$ since it is not necessarily a smooth surface. Consider $\partial\mathcal{B}_\alpha(P) = \partial\mathcal{B}_{\alpha,0}(P)$. This is the boundary of a spherical polyhedron. The boundaries of the facets in this spherical polyhedron constitute non-smooth regions in $\partial\mathcal{B}_\alpha(P)$. A point $x(t) \in \partial\mathcal{B}_{\alpha,t}(P)$ is given by $x(t) = \bar{x} + (1 - t)(x - \bar{x})$ where $x \in \partial\mathcal{B}_\alpha(P)$. It follows that $x(t)$ is a smooth point in $\partial\mathcal{B}_{\alpha,t}(P)$ if x is smooth in $\partial\mathcal{B}_\alpha(P)$. For a non-smooth point $x(t) \in \partial\mathcal{B}_{\alpha,t}(P)$, consider the set of normals $N_{x(t)}$ that are limit points of the normals to the smooth points in any neighborhood of $x(t)$. The normal cone at $x(t)$ is given by the convex hull $\text{Conv}N_{x(t)}$.

We first establish a bound on the angles between normals to $\partial\mathcal{B}_\alpha(P)$ at smooth points and the normals of Σ and then extend the bound for all points in $\partial\mathcal{B}_{\alpha,t}(P)$, $t \in [0, 1]$. Proposition 4 is proved in the appendix. Recall that P is assumed to be an ϵ -sample of Σ .

Proposition 4 Let $\alpha = \kappa\epsilon\rho$ and $\kappa\epsilon \leq 1/2$. Let x be any point in $\partial\mathcal{B}_{\alpha,t}(P)$. The normal of $\partial\mathcal{B}_{\alpha,t}(P)$ at x and the normal to Σ at \bar{x} form an angle of at most $4\sin^{-1}(2/\kappa)$.

Let us now consider a point x on a Voronoi facet or edge that $\partial\mathcal{B}_{\alpha,t}(P)$ is about to sweep. Let $6\epsilon\rho \leq \alpha \leq 6\epsilon\rho + O(\epsilon\rho)$. For this choice of α , we have $\kappa = 6$ in Proposition 4 which assures that a normal to a point x on $\partial\mathcal{B}_{\alpha,t}(P)$ makes an angle of at most $4\sin^{-1}(1/3) \approx 80^\circ$ with the surface normal at \bar{x} . By Proposition 1, the Voronoi edges and facets intersecting $\mathcal{B}_{\alpha,t}(P)$ make an angle of $O(\epsilon)$ with the surface normals. A standard calculation reveals that they intersect $\partial\mathcal{B}_{\alpha,t}(P)$ with an angle of at most $4\sin^{-1}(1/3) + O(\epsilon)$ which is smaller than $\frac{\pi}{2}$ if ϵ is sufficiently small. Therefore, Voronoi edges and facets cannot intersect $\partial\mathcal{B}_{\alpha,t}(P)$ tangentially implying that all peelings are associated with a

Voronoi vertex going out of $\mathcal{B}_{\alpha,t}(P)$. Thus, a canonical simplex sequence $\sigma_1, \dots, \sigma_n$ induces a canonical edge sequence at which tetrahedra are peeled which effectively removes all of $\{\sigma_1, \dots, \sigma_n\}$. We have:

Proposition 5 For $6\rho\varepsilon \leq \alpha \leq 6\varepsilon\rho + O(\varepsilon\rho)$, let $\mathcal{K} = \mathcal{C}_\alpha(P)$. There is a canonical edge sequence $\{e_1, e_2, \dots, e_k\}$ where $\mathcal{K} \xrightarrow{e_1, e_2, \dots, e_k} \text{Del}_\Sigma(P)$.

3.3. Slivers and top-down peeling

Now we show that all tetrahedra in $\mathcal{C}_\alpha(P)$ are 2-2-flat. We also refer to them as *slivers* although strictly speaking they may not be slivers by standard definition in meshing. A tetrahedron σ is called ε -flat if all normals to its triangles differ by $O(\varepsilon)$ angle. For a sufficiently small ε , an ε -flat tetrahedron can only be in two configurations with respect to the internal dihedral angles at its edges. These internal dihedral angles are either $\pi - O(\varepsilon)$ or $O(\varepsilon)$. If two angles are $\pi - O(\varepsilon)$ and the other four are $O(\varepsilon)$, we say it is a 2-2-flat tetrahedron (corresponds to slivers.) Otherwise, when three angles are $\pi - O(\varepsilon)$ and the other three are $O(\varepsilon)$, we say it is a 1-3-flat tetrahedron. No other configuration is possible for an ε -flat tetrahedron.

Proposition 6 Let $\alpha = O(\varepsilon)\rho$. For $t \in [0, 1]$ let $\sigma \in \mathcal{C}_{\alpha,t}(P)$. (i) If σ is a triangle, $\angle n_\sigma n_p = O(\varepsilon)$ where p is any vertex of σ . (ii) If σ is a tetrahedron, it is 2-2-flat.

Proof (i) Since the dual Voronoi edge is intersected by $\mathcal{B}_{\alpha,t}(P)$, the circumradius of σ is at most $\alpha = O(\varepsilon)\rho$. The claim follows from Proposition 1(iii). (ii) First observe that σ is an ε -flat tetrahedron due to (i). If σ were 1-3-flat, there is a vertex p of σ which subtends a solid angle close to 2π . The line of the normal n_p intersects the opposite triangle τ of σ , for otherwise n_p has to be almost parallel to τ contradicting Proposition 1(iii). We now apply the argument of Lemma 11 of [ACDL02] to claim that this is impossible. \square

Let us now go back to the scenario when a Voronoi vertex v is about to be swept over. The edge e at which the dual tetrahedron of v is peeled, subtends a large dihedral angle $\pi - O(\varepsilon)$. We use this observation to ensure a compatible edge sequence in the algorithm. The following definition helps. We say a 2-2-flat tetrahedron is peeled *top-down* if it is peeled at an edge subtending a large dihedral angle of $\pi - O(\varepsilon)$. We also say an edge sequence e_1, e_2, \dots, e_k is *top-down* if all tetrahedra are peeled top-down by the sequence. Since any top-down edge sequence peels a 2-2-flat tetrahedron only at edges that subtend large dihedral angles, two such sequences must be compatible.

Proposition 7 For a complex containing only 2-2-flat tetrahedra, any two top-down peeling sequences are compatible.

3.4. Algorithm

We already indicated that we cannot determine the canonical sequence because Σ is not known. But, thanks to Proposition 2, if we find any edge sequence for \mathcal{K} that is compatible

with the canonical one, we obtain a 2-complex isotopic to $\text{Del}_\Sigma(P)$. Our algorithm below finds such a sequence.

PEEL(P, α)

1. Compute the α -complex $\mathcal{K} := \mathcal{C}_\alpha(P)$;
2. While there is a peelable tetrahedron in \mathcal{K} , peel it top-down;
3. Output the resulting 2-complex.

Theorem 1 For $6\varepsilon\rho \leq \alpha \leq 6\varepsilon\rho + O(\varepsilon\rho)$ the algorithm PEEL computes a subcomplex of $\mathcal{C}_\alpha(P)$ that is a manifold triangulation isotopic to Σ , and is within a small Hausdorff distance.

Proof First, we argue that the algorithm peels all tetrahedra from $\mathcal{C}_\alpha(P)$. By Proposition 5 there is a canonical edge sequence e_1, e_2, \dots, e_k so that $\mathcal{C}_\alpha(P) \xrightarrow{e_1, e_2, \dots, e_k} \text{Del}_\Sigma(P)$. Suppose that the algorithm gets “stuck” with tetrahedra $\sigma_1, \dots, \sigma_s$ none of which is peelable, and let σ_i be the tetrahedron among these that first appears in the canonical sequence of peeling. Since σ_i is first, the canonical sequence finds σ_i peelable at a moment when all of the other $\sigma_j, j \neq i$, are still present. This is a contradiction and our algorithm should also find σ_i peelable.

Let e'_1, e'_2, \dots, e'_k be the edge sequence by which the algorithm peels $\mathcal{C}_\alpha(P)$, that is, $\mathcal{C}_\alpha(P) \xrightarrow{e'_1, e'_2, \dots, e'_k} T$ where T is the output complex. In the canonical sequence e_1, e_2, \dots, e_k , each peel is top-down. Each peel conducted by the algorithm is also top-down. Therefore, the sequences $\{e_i\}$ and $\{e'_i\}$ are compatible (Proposition 7). By Proposition 2, T is isotopic to $\text{Del}_\Sigma(P)$. The claim about Hausdorff distance follows from standard calculations as T consists of triangles with small circumradii, see [AB99, Dey07]. \square

4. Surface with boundaries

We show that the algorithm PEEL computes a surface isotopic to Σ even if Σ has a non-empty boundary. A first difficulty to overcome is that a surface with boundaries do not necessarily admit a restricted Delaunay triangulation which is homeomorphic to the surface no matter how dense the sample is. We use a collar extension to counter this difficulty.

4.1. Collar extension

Recall that we do not require the boundary curves of the surface themselves be well-sampled. Under this condition, for a good sample, $\text{Del}_\Sigma(P)$ may not be an isotopic or homeomorphic reconstruction of Σ (though it is homotopy equivalent to it). See the picture on left in Figure 3. The restricted Delaunay triangulation there is not a proper manifold (there is a pinching at the boundary). This can be repaired though by considering the Delaunay triangulation restricted to the surface with added *collar*. Let x be any point in the boundary $\partial\Sigma$. Consider the tangent fiber γ_x at x that is perpendicular to $\partial\Sigma$ at x . A segment of this fiber of length μ is denoted γ_x^μ . Let




Figure 3: Left: The restricted Delaunay triangulation in the ‘almost’ planar configuration has a non-manifold feature in its boundary. Right: The collar is obtained by extending from every boundary point x a “fiber” in the tangent plane at x .

$\Sigma_\mu = \Sigma \cup \{\gamma_x^\mu\}_{x \in \partial\Sigma}$. See the picture on right in Figure 3. It is a standard result in differential topology that every smooth surface can be extended with a collar while maintaining an isotopy if the extension length μ is sufficiently small.

We need to investigate how much the normal changes within the collar compared to the normal at the original boundary. Note that the normal can indeed change along an extension fiber. The following proposition is proved in the appendix.

Proposition 8 Let x be a boundary point and γ_x^μ its extension fiber. For $x' \in \gamma_x^\mu \subset \Sigma_\mu$ with $\|x - x'\| \leq \mu$ we have $\tan \angle n_x n_{x'} = O(\mu/\rho)$.

Finally, we want to show that the Delaunay triangulation restricted to Σ_μ is a good canonical reconstruction when $\mu = O(\rho\varepsilon)$ is chosen appropriately. For a sufficiently small μ , the bound on the normal in the collar implies that Voronoi edges and facets intersect Σ_μ almost orthogonally in the collar (which is already known for Σ). We also need a similar fact for the boundary $\partial\Sigma_\mu$.

Proposition 9 Let $\mu = \kappa\varepsilon\rho$. Let F be a Voronoi facet that intersects $\partial\Sigma_\mu$ at x . If τ_x is the tangent to $\partial\Sigma_\mu$ at x and n_F is the normal to F , then $\angle \tau_x, n_F \leq O(\kappa\varepsilon) + \cos^{-1} \frac{\kappa}{1+\kappa}$ where ε is sufficiently small.

Standard arguments [Dey07] using the near orthogonality between the surface and Voronoi faces provide the following result (see appendix for a proof).

Theorem 2 For $2\varepsilon\rho \leq \mu \leq 2\varepsilon\rho + O(\varepsilon\rho)$, $\text{Del}_{\Sigma_\mu} P$ is a manifold triangulation isotopic to Σ_μ when ε is sufficiently small.

4.2. Boundary and peeling

We follow PEEL to remove tetrahedra from $C_\alpha(P)$. Assume that $6\varepsilon\rho \leq \alpha \leq 6\varepsilon\rho + O(\varepsilon\rho)$ as before. One problem we now face is that not all points in $\mathcal{B}_\alpha(P)$ project in the interior of Σ by the closest point map. Observe that the proof of Proposition 4 still remains valid for surfaces with boundaries. Therefore, at a point $x \in \mathcal{B}_{\alpha,t}(P)$ on a Voronoi face f , the segment

$x\bar{x}$ makes small angle with a normal to $\partial\mathcal{B}_{\alpha,t}(P)$. But, unlike the closed surface case, $x\bar{x}$ may not be almost parallel to f because if \bar{x} is on $\partial\Sigma$, the angle $\angle x\bar{x}, n_{\bar{x}}$ could be large (recall the definition of surface normals at boundary points). We cannot claim that a Voronoi face in $\mathcal{B}_{\alpha,t}(P)$ is never tangent to $\partial\mathcal{B}_{\alpha,t}(P)$.

As a remedy we consider extending Σ by a small amount so that all points in $\mathcal{B}_\alpha(P)$ project in the interior of the extended Σ . This requires a slightly larger balls in the union to contain the extended surface inside. This larger union sweeps over all Voronoi faces it intersects, but we only focus on the subset of the Voronoi faces which intersect $\mathcal{B}_\alpha(P)$.

Proposition 10 Any point in $\mathcal{B}_\alpha(P)$ projects to an interior point in $\Sigma_{3\alpha}$.

Proof Let x be any point in $\mathcal{B}_\alpha(P)$. If x projects normally to the interior of Σ , it does so for any Σ_μ if μ is small. Consider x projecting to a point \bar{x} in $\partial\Sigma$. The closest point of x on the surface Σ_μ , $\mu > 2\alpha$, cannot be further away than 2α from \bar{x} . It follows that x projects to the interior of Σ_μ for $\mu \geq 3\alpha$. \square

Consider $\mathcal{B}_{4\alpha}(P)$. Observe that any point in $\Sigma_{3\alpha}$ is within a distance of $\varepsilon\rho + 3\alpha \leq 4\alpha$ from a sample point. It means that $\mathcal{B}_{4\alpha}(P)$ contains $\Sigma_{3\alpha}$. We consider the retraction of $\mathcal{B}_{4\alpha}(P)$ to $\Sigma_{3\alpha}$. This retraction defines a canonical sequence of simplex deletions reducing $C_{4\alpha}(P)$ to $\text{Del}_{\Sigma_{3\alpha}}(P)$. We are interested in the restriction of this sequence to the simplices of $C_\alpha(P)$. Because of Proposition 10, we can argue similarly as in the closed surface case that a Voronoi face in $\mathcal{B}_\alpha(P)$ is never tangent to $\partial\mathcal{B}_{4\alpha,t}(P)$ if $6\varepsilon\rho \leq \alpha \leq 6\varepsilon\rho + O(\varepsilon\rho)$. Therefore the canonical sequence of deleted simplices in $C_{4\alpha}(P)$ indeed induces an edge sequence at which all tetrahedra in $C_\alpha(P)$ are peeled top-down. This is the canonical sequence of peeling we consider for $C_\alpha(P)$.

There is one more difference from the closed surface case that we need to address. In the closed surface case $C_\alpha(P)$ retracts to $\text{Del}_\Sigma(P)$. Here, since we are considering the retraction of $\mathcal{B}_{4\alpha}(P)$, we cannot claim that the induced peeling of $C_\alpha(P)$ will provide $\text{Del}_\Sigma(P)$.

Let \mathcal{K}^* be the complex produced by the canonical peeling of $C_\alpha(P)$ induced by the canonical sequence of simplex deletions in $C_{4\alpha}(P)$. First notice that since $C_{4\alpha}(P)$ retracts to $\text{Del}_{\Sigma_{3\alpha}}$ and \mathcal{K}^* is obtained by restricting this retraction to $C_\alpha(P)$, we have $\mathcal{K}^* \subseteq \text{Del}_{\Sigma_{3\alpha}}(P)$. Also, since $\text{Del}_{\Sigma_{2\varepsilon\rho}}(P) \subseteq \text{Del}_{\Sigma_{3\alpha}}(P)$ and $C_\alpha(P)$ contains $\text{Del}_{\Sigma_{2\varepsilon\rho}}(P)$ for $\alpha > 3\varepsilon\rho$, we have $\text{Del}_{\Sigma_{2\varepsilon\rho}}(P) \subseteq \mathcal{K}^*$. In essence, $\text{Del}_{\Sigma_{2\varepsilon\rho}}(P) \subseteq \mathcal{K}^* \subseteq \text{Del}_{\Sigma_{3\alpha}}$, that is, \mathcal{K}^* is sandwiched between two triangulations both of which are isotopic to Σ . This leads to the main result of this paper.

Theorem 3 Let P be an ε -sample of a smooth compact surface with boundary. For sufficiently small $\varepsilon > 0$ and $6\varepsilon\rho < \alpha \leq 6\varepsilon\rho + O(\varepsilon\rho)$, PEEL(P, α) produces a 2-complex T isotopic to Σ .

Proof The algorithm PEEL computes $C_\alpha(P)$ in step 2. Consider the complex \mathcal{K}^* produced by the canonical peeling in

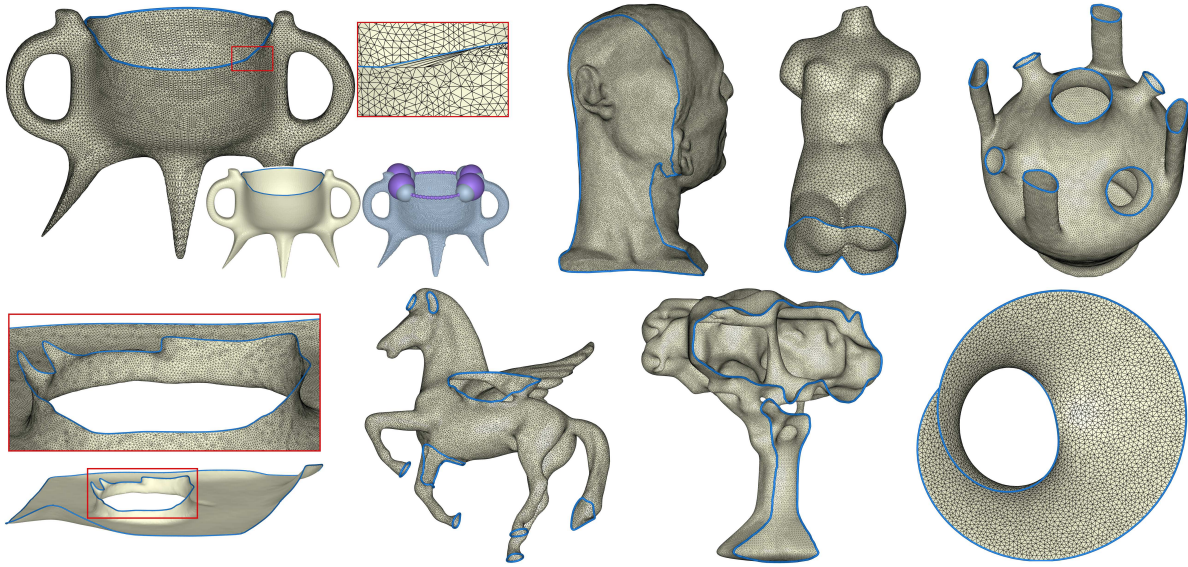


Figure 4: From top to bottom, left to right: Reconstruction of Ding, Nicolo, Venus, Botijo, Volcano, Pegasus, Tree, and Möbius band from uniform samples.

$C_\alpha(P)$ induced by the canonical peeling of $C_{4\alpha}(P)$. First we argue that if \mathcal{K}^* is isotopic to Σ , our algorithm produces a complex that is isotopic to Σ as well.

Let e_1, e_2, \dots, e_k be the canonical sequence of edges where $C_\alpha(P) \xrightarrow{e_1, e_2, \dots, e_k} \mathcal{K}^*$. Since all Voronoi vertices dual to the tetrahedra in $C_\alpha(P)$ project normally to the interior of $\Sigma_{3\alpha}$, the canonical peeling is top-down. We can apply the same argument as in Theorem 1 to claim that the algorithm finds a peeling sequence which reduces $C_\alpha(P)$ to a 2-complex T where \mathcal{K}^* and T are isotopic. Hence T is isotopic to Σ if \mathcal{K}^* is isotopic to Σ .

We argue that \mathcal{K}^* is indeed isotopic to Σ . For convenience we write $R_1 = \text{Del}_{\Sigma_{2\alpha}}(P)$ and $R_2 = \text{Del}_{\Sigma_{3\alpha}}$. We have $R_1 \subseteq \mathcal{K}^* \subseteq R_2$. By theorem 2, R_1 and R_2 are isotopic to Σ and thus are isotopic to each other. Also, \mathcal{K}^* is homotopy equivalent to $C_\alpha(P)$ since each peel maintains a homotopy equivalence between the complexes before and after the peel and $C_\alpha(P)$ is homotopy equivalent to Σ to begin with [CCSL06].

First we observe a property of the complex that represents the difference between R_1 and R_2 . Formally, let $W = \text{cl}(R_2 - R_1)$. The space $|W|$ cannot contain any handle or Möbius band since then R_1 will have different topology than R_2 . Next we prove that \mathcal{K}^* is a manifold.

We introduce the following definition for convenience. For a vertex v in a complex \mathcal{K} , let T_v denote its star. If the underlying space $|T_v|$ is a topological disk, and v is in the interior of $|T_v|$, we say v is *complete* in \mathcal{K} . If $|T_v|$ is a topological disk but v lies on the boundary, we say v has a half-disk neighborhood in \mathcal{K} . Observe that each complete vertex of R_1 remains complete in \mathcal{K}^* and R_2 . Also, the vertex set of R_1 , R_2 and \mathcal{K}^* is same. These two facts imply that each boundary in \mathcal{K}^*

and R_2 is generated from the boundaries of R_1 by possibly attaching triangles among boundary vertices. Since connecting vertices across different boundaries would require triangles with edges larger than 4α , boundaries of \mathcal{K}^* and R_2 are generated by attaching triangles that connect boundary vertices in the same boundary of R_1 .

Now consider an incomplete vertex v of \mathcal{K}^* . Since v is a boundary vertex in R_1 , we can assume that v is in a boundary b in R_1 . If v does not have a half-disk neighborhood in \mathcal{K}^* , it is incident to at least two boundaries in \mathcal{K}^* which is generated from b . But then \mathcal{K}^* would contain more boundaries than R_1 or R_2 since each boundary of R_2 also provides at least one boundary in \mathcal{K}^* . This is because an incomplete vertex in R_2 also remains incomplete in \mathcal{K}^* . Since the difference $\text{cl}(\mathcal{K}^* - R_1) \subseteq W$ has no handles and Möbius bands, \mathcal{K}^* has same number of handles and Möbius bands as in R_1 . Then, \mathcal{K}^* cannot have different number of boundaries from R_1 since \mathcal{K}^* and R_1 are homotopy equivalent. It follows that \mathcal{K}^* is a manifold since it has vertices which are either complete or have a half-disk neighborhood.

Now we examine the space of $\text{cl}(\mathcal{K}^* - R_1)$ more closely. A boundary, say b' in \mathcal{K}^* is generated from a boundary b in R_1 by possibly attaching triangles between vertices of b . Furthermore, a single boundary b in R_1 generates exactly one boundary in \mathcal{K}^* . The space bounded by b and b' in \mathcal{K}^* does not have any handle, boundary, or Möbius strip. This means b and b' bounds a cylinder which is possibly pinched at common vertices of b and b' . There is a natural deformation retraction of \mathcal{K}^* to R_1 defined by the obvious deformation retraction of b'_i s to b_i s along these cylinders. This deforma-

tion retraction is a witness to an isotopy between \mathcal{K}^* and R_1 . Since R_1 is isotopic to Σ , we have the claimed result. \square

5. Extensions and experiments

We implemented PEEL and some results are shown in Figure 4. When the sample is uniformly dense, PEEL produces correct results as these examples show. The output is a 2-manifold without any artifact. We can even reconstruct non-orientable surfaces such as the Möbius band example in Figure 4. We chose the α to be a factor (six times) of the largest nearest neighbor distance among the given points. Since all these examples have almost uniform sampling, this choice worked nicely confirming our theory.

In practice, however, the data are often non-uniform. We made adjustments to our algorithm to handle non-uniform data. We emphasize that if the data is a non-uniform sample of a surface with a boundary, there is a theoretical hurdle to reconstruct it provably. It stems from the fact that the same sample can be dense for two topologically different surfaces if boundaries are allowed. Therefore, there is no unique correct reconstruction. We employ a heuristic to handle non-uniform samples of surfaces with boundaries.

If the input is non-uniform, there may be no global α for which an α -complex may contain an isotopic triangulation to the sampled surface. Since nearest neighbor distances may vary widely if non-uniform density is allowed, an α fitting local density cannot be estimated from them. Instead we propose to estimate the local density at each input point and then take a ball around the point whose size respects the estimated density. As before, we consider the union of balls \mathcal{U} and peel tetrahedra from the restricted Delaunay complex $\text{Del}_{\mathcal{U}}(P)$. Notice that if the balls are not too large or too small with respect to local feature sizes, the Voronoi elements intersecting \mathcal{U} will do so transversally and therefore a deformation retraction of \mathcal{U} would induce a canonical edge sequence in the restricted complex.

First, we determine the set R of all sample points near the boundary by the modified Cocone algorithm [DG01] (We use the software COCONE). A parameter in COCONE regulates which points are detected as boundary. Then, for each point p we compute a *width* w_p as follows. For each non-boundary point $p \in P - R$, the width w_p is the *cocone width* [DG01]. Roughly speaking, the cocone width estimates the largest distance of p to its local Voronoi neighbors on the surface (cocone neighbors) and thus tracks the local sparsity/density of the sample. For a boundary point $p \in R$, we take the largest width of its non-boundary cocone neighbors as w_p . Then,

$$\mathcal{U} = \bigcup_{p \in P} B(p, w_p).$$

In our implementation, for efficiency, we do not consider all simplices in $\text{Del}_{\mathcal{U}}(P)$. Instead a subcomplex consisting

of simplices whose dual Voronoi elements intersect the cocones [ACDL02, DG01] is considered. The rationale is that this subcomplex contains the restricted Delaunay triangulation of the surface and at the same time inherits a canonical edge sequence from the retraction of $\text{Del}_{\mathcal{U}}(P)$. Figure 5 shows the results of this extension.

6. Concluding remarks

Our results on sliver peeling may be of independent interest. By getting rid of the prune-and-walk step, sliver peeling makes the reconstruction process more parallelizable. For the same reason, it also allows reconstruction of non-orientable surfaces which is not possible with any of the existing algorithms.

The question of handling non-uniform samples with theoretical guarantees remains open. Our approach in section 5 may very well have provable guarantees for locally uniform samples. However, without any local uniformity, it is difficult to estimate the local density of the sample which is required to estimate the ball sizes. Ideas from [FR02] may be used to convert a non-uniform sample to a locally uniform sample and then apply the method in section 5. Our approach does not handle noisy samples, an important case which remains open.

Acknowledgments: We acknowledge AIM@SHAPE database for models and the NSF grant CCF-0635008 for financial support.

References

- [AB99] AMENTA N., BERN M.: Surface reconstruction by voronoi filtering. *Discr. Comput. Geom.* 22 (1999), 481–504.
- [ACDL02] AMENTA N., CHOI S., DEY T. K., LEEKHA N.: A simple algorithm for homeomorphic surface reconstruction. *Internat. J. Comput. Geom. & Applications* 12 (2002), 125–141.
- [ACSTD07] ALLIEZ P., COHEN-STEINER D., TONG Y., DESBRUN M.: Voronoi-based variational reconstruction of unoriented point sets. In *Proc. Sympos. Geom. Processing* (2007), pp. 39–48.
- [BC00] BOISSONNAT J. D., CAZALS F.: Smooth surface reconstruction via natural neighbor interpolation of distance functions. In *Proc. 16th Ann. Sympos. Comput. Geom.* (2000), pp. 223–232.
- [CCSL06] CHAZAL F., COHEN-STEINER D., LIEUTIER A.: A sampling theory for compacts in euclidean space. In *Proc. 22nd Ann. Sympos. Comput. Geom.* (2006), pp. 319–326.
- [CDR05] CHENG S.-W., DEY T. K., RAMOS E. A.: Manifold reconstruction from point samples. In *Proc. ACM-SIAM Sympos. Discr. Algorithms* (2005), pp. 1018–1027.
- [CSEH05] COHEN-STEINER D., EDELSBRUNNER H., HARER J.: Stability of persistence diagram. In *Proc. 21st Ann. Sympos. Comput. Geom.* (2005), pp. 263–271.
- [Dey07] DEY T. K.: *Curve and surface reconstruction : algorithms with mathematical analysis*. Cambridge University Press, New York, 2007.
- [DG01] DEY T. K., GIESEN J.: Detecting undersampling in surface reconstruction. In *Proc. 17th Ann. Sympos. Comput. Geom.* (2001), pp. 257–263.




Figure 5: From top to bottom, left to right: Reconstruction of Terrain, Manneq, SGP09, and Foot from non-uniform samples. A parameter in Cocone regulates which points are detected as boundary. The points bordering the letters in SGP09 and also three spots in Foot bottom are not detected as boundary points. Tightening this parameter would detect them as boundary points but small holes in relatively undersampled places would appear.

- [Ede95] EDELSBRUNNER H.: The union of balls and its dual shape. *Discr. Comput. Geom.* 13 (1995), 415–440.
- [Ede03] EDELSBRUNNER H.: Surface reconstruction by wrapping finite point sets in space. In *Ricky Pollack and Eli Goodman Festschrift*, Aronov B., Basu S., Pach J., Sharir M., (Eds.). Springer Verlag, New York, 2003, pp. 379–404.
- [ES94] EDELSBRUNNER H., SHAH N.: Triangulating topological spaces. In *Proc. 10th Ann. Sympos. Comput. Geom.* (1994), pp. 285–292.
- [FCOS05] FLEISHMAN S., COHEN-OR D., SILVA C. T.: Robust moving least-squares fitting with sharp features. In *Proc. ACM SIGGRAPH* (2005), pp. 544–552.
- [FR02] FUNKE S., RAMOS E. A.: Smooth-surface reconstruction in near-linear time. In *Proc. 13th ACM-SIAM Sympos. Discr. Algorithms* (2002), pp. 781–790.
- [GJ02] GIESEN J., JOHN M.: Surface reconstruction based on a dynamical system. In *Proc. Eurographics* (2002), pp. 363–371.
- [HDD*92] HOPPE H., DE ROSE T., DUCHAMP T., McDONALD J., STUETZLE W.: Surface reconstruction from unorganized points. In *Proc. ACM SIGGRAPH* (1992), pp. 71–78.
- [JWS08] JENKE P., WAND M., STRASSER W.: Patch-graph reconstruction for piecewise smooth surfaces. In *Proc. Vision, Modeling and Visualization* (2008), pp. 3–12.
- [KBH06] KAZHDAN M., BOLITHO M., HOPPE H.: Poisson surface reconstruction. In *Proc. Sympos. Geom. Processing* (2006), pp. 61–70.
- [NSW08] NIYOGI P., SMALE S., WEINBERGER S.: Finding the homology of sub-manifolds with high confidence from random samples. *Discr. Comput. Geom.* 39 (2008), 419–441.
- [OBA*03] OHTAKE Y., BELYAEV A., ALEXA M., TURK G., SEIDEL H.-P.: Multi-level partition of unity implicits. In *Proc. ACM SIGGRAPH* (2003), pp. 463–470.
- [PKKG03] PAULY M., KEISER R., KOBELT L. P., GROSS M.: Shape modeling with point-sampled geometry. In *Proc. ACM SIGGRAPH* (2003), pp. 641–650.
- [WOK05] WANG J., OLIVEIRA M. M., KAUFMAN A. E.: Reconstructing manifold and non-manifold surfaces from point clouds. In *Proc. IEEE Visualization* (2005), pp. 415–422.

Appendix A: Proof of Proposition 1

Proof The proof for the closed surface case in [CDR05] takes the segment xy , considers the closest point projection γ of xy onto Σ and then (a) puts an upper bound on the length of γ by a multiple of the length of xy , and (b) bounds the normal variation on γ . In the argument for (a), the crucial point is that for $p \in xy$, the segment pp' from p to its projection p' in Σ is normal to Σ at p' and so there is an empty tangent medial ball B at p' whose center is in the ray from p' in the direction of p and its radius is $\text{lfs}(p')$. The difference here is that p' may lie on $\partial\Sigma$ and so pp' does not need to be normal to the surface, but to its bounding curve; still the same conclusion about the tangent ball holds. Using this and arguing as in the Figure 6 leads to the conclusion that $\text{length}(\gamma) \leq 2\|x - y\|$. The argument for (b) is the same as for closed surfaces. \square




Figure 6: The closest point q' to q in Σ is in $B' = B(q, \|q - p'\|)$ but outside B . From this follows that $\|q - q'\| \leq 2\|p - p'\|$ for sufficiently small ε .

Appendix B: Proof of proposition 3

Proof For the proof by induction on k we add to the conclusion that \mathcal{K}_1 and \mathcal{K}_2 contain the same set of tetrahedra, if any at all, and that the isotopy can be chosen to keep $\mathcal{K}_1 \cap \mathcal{K}_2$ fixed. When $k = 1$, e and e_1 peels the same tetrahedron, say τ . In this case $\mathcal{K}' = \mathcal{K}_2$. If $e = e_1$, we have $\mathcal{K}_1 = \mathcal{K}'$ and there is nothing to prove. If $e \neq e_1$, let \mathcal{K}_1 have the triangles t_1, t_2 of τ and \mathcal{K}' have the triangles t_3, t_4 of τ . Since e and e_1 are compatible and hence vertex disjoint, the two sets of triangles are disjoint. Hence $t_1 \cup t_2$ can be deformed to $t_3 \cup t_4$ with an isotopy that can keep all other points of \mathcal{K}_1 fixed. In essence we have an isotopy between \mathcal{K}_1 and $\mathcal{K}' = \mathcal{K}_2$. Clearly, \mathcal{K}_1 and $\mathcal{K}' = \mathcal{K}_2$ contain same set of tetrahedra since both of them are obtained from \mathcal{K}_0 by peeling τ .

Now consider $k > 1$. Let $\mathcal{K}_0 \xrightarrow{e_1, \dots, e_{k-1}} \mathcal{K}'_1$. First assume that e peels a tetrahedron that is not peeled by e_k . Then it is peeled by an edge in the sequence e_1, \dots, e_{k-1} . We can apply the inductive hypothesis by which there is a sequence e'_1, \dots, e'_{k-2}

compatible with e_1, \dots, e_{k-1} so that $\mathcal{K}' \xrightarrow{e'_1, \dots, e'_{k-2}} \mathcal{K}'_2$ and \mathcal{K}'_1 is isotopic to \mathcal{K}'_2 with the same set of tetrahedra. Since the isotopy keeps $\mathcal{K}'_1 \cap \mathcal{K}'_2$ fixed, the tetrahedron peeled by e_k in \mathcal{K}'_1 is also peelable at e_k in \mathcal{K}'_2 . Peeling e_k from \mathcal{K}'_1 produces \mathcal{K}_1 . Peeling the same from \mathcal{K}'_2 produces \mathcal{K}_2 . We have \mathcal{K}_1 and \mathcal{K}_2 isotopic with the same set of tetrahedra and the isotopy keeps $\mathcal{K}_2 \cap \mathcal{K}_1$ fixed. So, the sequence $e'_1 \dots e'_{k-2} e_k$ satisfy the inductive hypothesis.

Next consider the case when e peels the tetrahedron which is peeled by e_k . We claim that e_1, \dots, e_{k-1} peel the same sequence of tetrahedra in \mathcal{K}' as in \mathcal{K}_0 . If not, let e_j be the first in the sequence which cannot satisfy this claim. Let e_j peel τ in \mathcal{K}_0 . The only reason why e_j cannot peel τ in \mathcal{K}' is that it is still incident to more than two triangles. But, that is impossible since e_1, \dots, e_{j-1} peeled same set of tetrahedra in both and \mathcal{K}' has a subset of triangles that are in \mathcal{K}_0 . Let \mathcal{K}_2 be such that $\mathcal{K}' \xrightarrow{e_1, \dots, e_{k-1}} \mathcal{K}_2$.

Obviously, $\mathcal{K}_0 \xrightarrow{e, e_1, \dots, e_{k-1}} \mathcal{K}_2$. Also notice that $\mathcal{K}_0 \xrightarrow{e_1, \dots, e_{k-1}, e} \mathcal{K}_2$. We also have $\mathcal{K}_0 \xrightarrow{e_1, \dots, e_{k-1}, e_k} \mathcal{K}_1$. So, the difference between \mathcal{K}_1 and \mathcal{K}_2 are made by e_k and e respectively. We have argued for similar situation in the case when $k = 1$. This argument implies that \mathcal{K}_2 and \mathcal{K}_1 have same set of tetrahedra and are isotopic where the isotopy keeps $\mathcal{K}_1 \cap \mathcal{K}_2$ fixed. We


have $e'_1, \dots, e'_{k-1} = e_1, \dots, e_{k-1}$ satisfying the inductive hypothesis. \square

Appendix C: Proof of Proposition 4

We first prove the following proposition.

Proposition 11 Let $\alpha = \kappa\varepsilon\rho$ and $\kappa\varepsilon \leq 1/2$. Let x be a smooth point in $\partial\mathcal{B}_\alpha(P)$. The angle between the normal of $\partial\mathcal{B}_\alpha(P)$ at x and the normal of Σ at \tilde{x} is at most $2\sin^{-1}(2/\kappa)$.

Proof Consider a point x in the interior facet of $\partial\mathcal{B}_\alpha(P)$, lying on the boundary of $B_\alpha(p)$ for $p \in P$, and its closest point \tilde{x} on Σ . For convenience write $\beta = \varepsilon\rho$. So $\alpha = \kappa\beta$. Because $\kappa\varepsilon \leq 1/2$ then $\alpha \leq \rho/2$. We want to first bound the angle $\angle px\tilde{x}$. The sample q closest to \tilde{x} lies in $B_\beta(\tilde{x})$ but outside of the medial balls $B = B(c, R)$ and $B' = B(c', R)$ tangent to Σ at \tilde{x} , where $R \equiv \rho$. See the figure. So p must lie in $\text{cl}(B(x, \|q - x\|) - B - B')$. Under this situation, $\angle px\tilde{x}$ is maximized by the angle $\theta = \angle vx\tilde{x}$ as shown in the figure. Considering the medial ball B' , we obtain



$$\angle vx\tilde{x} = \cos^{-1}(\beta/2R)$$

Applying law of sines to triangle $\Delta ux\tilde{x}$ we have

$$\sin \angle ux\tilde{x} = \frac{\beta}{\alpha} \sqrt{1 + \left(\frac{\beta}{2R}\right)^2}$$

and a second application to the same triangle gives

$$h = \alpha \left(\sqrt{1 - \left(\frac{\beta}{\alpha}\right)^2} \left(1 - \left(\frac{\beta}{2R}\right)^2\right) - \frac{\beta}{\alpha} \cdot \frac{\beta}{2R} \right).$$

Using $\alpha \leq R/2$, this implies the following bounds for h/α

$$1 - \frac{1}{\kappa^2} \left(1 + \frac{\alpha}{2R}\right) \leq \frac{h}{\alpha} \leq 1.$$

Then, applying the law of cosines to Δcvx , noting that $\angle vxc = \pi - \theta$ and using $\sin^2 \theta/2 = (1 - \cos \theta)/2$, we have

$$\sin(\theta/2) = \sqrt{\frac{1}{2} \cdot \left(1 - \frac{h}{\alpha}\right) \cdot \left(1 + \frac{\alpha + h}{2(R - h)}\right)}$$

Using the bounds for h/α above and $\alpha \leq R/2$, we obtain

$$\sin(\theta/2) \leq \sqrt{1 - \frac{h}{\alpha}} \leq \frac{1}{\kappa} \sqrt{1 + \frac{\alpha}{2R}} \leq \frac{2}{\kappa}$$

and so

$$\theta \leq 2 \sin^{-1} \frac{2}{\kappa}$$

Now, note that $\angle px\bar{x}$ is equal to $\angle \bar{p}\bar{x}, n_{\bar{x}}$, so we have found that

$$\angle \bar{p}\bar{x}, n_{\bar{x}} \leq 2 \sin^{-1} \frac{2}{\kappa}.$$

□

Proof[Proof of Proposition 4.] Let $x = x(t)$ and $\theta(x(t))$ denote this angle. First assume that $x(t)$ is a smooth point in $\partial\mathcal{B}_{\alpha,t}(P)$. We already observed that $x(t)$ remains smooth in $\partial\mathcal{B}_{\alpha,t}(P)$ for all $t \in [0, 1]$. Let $W(t) \subset \partial\mathcal{B}_{\alpha,t}(P)$ be a sufficiently small smooth neighborhood of $x(t)$. Consider the map $f : W(0) \rightarrow W(t)$ given by $f(w) = v(w) + (1-t)(\iota(w) - v(w))$ where $v : W(0) \rightarrow \Sigma$ is the restriction of the closest point map to $W(0)$ and $\iota : \partial\mathcal{B}_{\alpha}(P) \rightarrow \partial\mathcal{B}_{\alpha,t}(P)$ is the identity. The derivative Df maps the tangent spaces $TW(0)$ of $W(0)$ to the tangent spaces $TW(t)$ of $W(t)$ and is given by $Df = Dv + (1-t)(Dt - Dv)$. Specifically, for a tangent vector $\tau(0) \in TW(0)$ at $x(0)$ we have $\tau(t)$ at $x(t)$ where $\tau(t) = Df_{x(0)}(\tau(0)) = \tau_{\bar{x}} + (1-t)(\tau(0) - \tau_{\bar{x}})$ and $\tau_{\bar{x}} = Dv_{x(0)}(\tau(0))$ is a tangent to Σ at \bar{x} . We see that tangent spaces are linearly interpolated by t between $TW(0)$ at $x(0)$ and $T\Sigma$ at \bar{x} . Hence the normals to $\partial\mathcal{B}_{\alpha,t}(P)$ at smooth points are also interpolated linearly by t . It follows that $\theta(x(t))$ is also interpolated by t between $\theta(x(0))$ and $\theta(x(1))$. Proposition 11 provides $\theta(x(0)) \leq 2 \sin^{-1}(2/\kappa)$. Clearly $\theta(x(1)) = 0$. Thus, we have $\theta(x(t)) \leq \theta(x(0)) \leq 2 \sin^{-1}(2/\kappa)$.

If $x(t)$ is not a smooth point, a normal in $ConvN_{x(t)}$ cannot make more than $2 \sin^{-1}(2/\kappa)$ angle with any normal on the boundary of $ConvN_{x(t)}$ since all these normals make at most $2 \sin^{-1}(2/\kappa)$ angle with the same surface normal $n_{\bar{x}}$. It follows that any normal at $x(t)$ makes at most $4 \sin^{-1}(2/\kappa)$ angle with $n_{\bar{x}}$. □


Appendix D: Proof of Proposition 8

Proof Let b be the boundary curve of Σ containing x . The tangent plane at x' contains the fiber γ_x and so $n_{x'}$ lies in a plane orthogonal to γ_x . Thus, to determine $n_{x'}$ one needs to investigate what the slope of the surface is when moving in the direction orthogonal to γ_x when seen as a function defined on the tangent plane T_x (the tangent plane at x). To be precise, say $x = (0, 0, 0)$ and $x' = (\mu, 0, 0)$, and consider another $y' = (\mu, \Delta, \delta\Delta)$ in the collar close to x' whose closest point in Σ is y , that is, $y' \in \gamma_y$. Thus, we have that $\tan \angle n_x n_{x'} = \delta$ when $\Delta \rightarrow 0$. We assume $\delta \geq 0$. Let $B = B(c, R)$ be the tangent ball to b at x with center in γ_x and radius $R = \text{lfs}(x)$. Since $x \in \Sigma$ then y lies in the cup $C = \text{cl}(B' - B)$ where $B' = B(y', r)$ where $r = \|y' - x\| =$

$\sqrt{\mu^2 + \Delta^2 + v^2\Delta^2}$ (y cannot be further away from y' than x). See top figure. Let $z \in C$ be the furthest point in C from x . From the geometry we find that

$$\begin{aligned} r' = \|z - x\| &= 2R \sin(\angle xcz/2) \\ &= 2R \sin \arctan \frac{\Delta\sqrt{1+\delta^2}}{\sqrt{\Delta^2(1+\delta^2) + (R-\mu)^2}} \\ &= \frac{2\Delta\sqrt{1+\delta^2}}{1-\mu/R} + o(\Delta). \end{aligned}$$

Within $B'' = B(x, r')$, all the medial balls around b at x (all of



radius $\text{lfs}(x)$) constraint the location of y to a cone-like shape C' as illustrated in the middle figure. Let v be the top most point of C' (with highest z coordinate). A simple calculation shows that the height of v is (see bottom figure) $h = \frac{r'^2}{2R}$. Note that h is second order in Δ while $\delta\Delta$ (the height of y') is first order in Δ . By Proposition 1, for x and y on $b \subseteq \Sigma$ we have the bound

$$\angle n_x n_y \leq c \frac{\|x - y\|}{R},$$

where c is a constant. Now, writing $y = (y_1, y_2, y_3)$, we know $y_2 \leq r'$ and $|y_1|, |y_3| \leq \|y\|^2/2R$ (as in the expression for h above). Since $\angle n_x n_y = \angle x x' y y'$, then

$$\tan \angle n_x n_y = \frac{\sqrt{(y_2 - \Delta)^2 + (\delta\Delta - y_3)^2}}{\mu - y_1}$$

and so putting this together with the upper bound we obtain

$$c \frac{\sqrt{y_1^2 + y_2^2 + y_3^2}}{R} \geq \tan^{-1} \frac{\sqrt{(y_2 - \Delta)^2 + (\delta\Delta - y_3)^2}}{\mu - y_1}.$$

>From here

$$c \frac{\sqrt{y_1^2 + y_2^2 + y_3^2}}{R} \geq \tan^{-1} \frac{\max\{|y_2 - \Delta|, |\delta\Delta - y_3|\}}{\mu - y_1}.$$

First, taking this inequality with the term $|y_2 - \Delta|$ in the max, as $\Delta \rightarrow 0$, we obtain (note that y_1, y_3 are second order in Δ and so they can be neglected)

$$\frac{\Delta}{1 + c\mu/R} \leq y_2 \leq \frac{\Delta}{1 - c\mu/R}.$$

So

$$y_2 \leq \min \left\{ \frac{2\sqrt{1 + \delta^2}}{1 - \mu/R}, \frac{1}{1 - c\mu/R} \right\} \Delta.$$

Next, taking that inequality with the term $|\delta\Delta - y_3|$ in the max, as $\Delta \rightarrow 0$, we obtain

$$\delta \leq \frac{c\mu}{R} \min \left\{ \frac{2\sqrt{1 + \delta^2}}{1 - \mu/R}, \frac{1}{1 - c\mu/R} \right\} = O(\mu/R).$$

□

Appendix E: Proof of Proposition 9

Proof Let F be formed by samples p, q . Consider the empty ball $B(x, \|p - x\|)$. The closest point \bar{x} of x in Σ is on $\partial\Sigma$ by our collar extension and \bar{x} is in B . Since \bar{x} has a sample within $\varepsilon\rho$, the radius of B , $\|p - x\|$, is at most $\mu + \varepsilon\rho$.

Since B intersects Σ and the radius of B is $O(\varepsilon\rho)$, it intersects Σ in a topological ball (apply the argument of Lemma 1.1 [Dey07]). The points $p, q \in \Sigma$ forming the Voronoi facet F lie on its opposite sides. It means that the plane of F intersects Σ and since it intersects it almost orthogonally, it intersects $\partial\Sigma$ as well.

The tangents of $\partial\Sigma$ and $\partial\Sigma_\mu$ at points \bar{x} and x respectively are parallel by construction. Furthermore, the tangents at all points of $\partial\Sigma$ within B vary by $O(\kappa\varepsilon)$ angle since any two such points are within $O(\kappa\varepsilon\rho)$ distance. So, we measure the angle between the normal n_F and the tangent $\tau_{\bar{x}}$ at \bar{x} where the plane of F intersects $\partial\Sigma$. This angle is within $O(\kappa\varepsilon)$ of the angle $\angle \tau_{\bar{x}}, n_F$ which we want to bound.

Consider the disk at which the plane containing $x\bar{x}$ and $\tau_{\bar{x}}$ intersects B . See Figure 7. Clearly, $\angle \tau_{\bar{x}}, n_F$ is no larger than the angle θ as shown in the figure.




Figure 7: Proposition 9

We obtain

$$\cos \theta = \frac{\mu}{\mu + \varepsilon\rho} = \frac{\kappa}{1 + \kappa}.$$

Therefore, $\angle \tau_{\bar{x}}, n_F \leq \angle \tau_{\bar{x}}, n_F + O(\kappa\varepsilon) = O(\kappa\varepsilon) + \cos^{-1}(\frac{\kappa}{1+\kappa})$ as claimed. □

Appendix F: Proof of Theorem 2

Proof First we verify that it is a manifold triangulation homeomorphic to Σ_μ . We need to show that each of the restricted Voronoi face satisfies the *topological ball property* [ES94]: for $V_\tau \in \text{Vor}P$, $V_\tau \cap \Sigma_\mu$ and $V_\tau \cap \partial\Sigma_\mu$ are topological balls of the appropriate dimension. The proof for the closed surfaces applies here when there is no intersection with the boundary. All that is essential is that Voronoi faces that intersect Σ_μ do so almost orthogonally [Dey07], which is also the case here inside the collar according to Proposition 8. So it is only necessary to verify the ball property on the boundary $\partial\Sigma_\mu$. By Proposition 9, $\partial\Sigma_\mu$ intersects a Voronoi facet at an angle $O(\kappa\varepsilon) + \cos^{-1} 2/3$ for $\kappa \geq 2$. This angle is no more than 50° for sufficiently small ε . This fixed bound is enough to carry out the standard proofs from the literature [Dey07] to establish that $\partial\Sigma_\mu$ intersects Voronoi facets and cells with topological ball property. Therefore, $\text{Del}_{\Sigma_\mu}(P)$ is homeomorphic to Σ_μ and hence to Σ as Σ_μ and Σ are isotopic by our construction.

We argue that $\text{Del}_{\Sigma_\mu}(P)$ is indeed isotopic to Σ . Consider $\Sigma_{4\mu}$. We project $\text{Del}_{\Sigma_\mu}(P)$ by the closest point map to $\Sigma_{4\mu}$. Since the circumradii of the triangles in $\text{Del}_{\Sigma_\mu}(P)$ are at most 2μ , it can be shown that the closest point map $v : \text{Del}_{\Sigma_\mu}(P) \rightarrow \Sigma_{4\mu}$ is injective [ACDL02]. The map also induces an isotopy between $\text{Del}_{\Sigma_\mu}(P)$ and its image $T = v(\text{Del}_{\Sigma_\mu}(P)) \subseteq \Sigma_{4\mu}$. If we show that T and Σ are isotopic, we are done.

Let us examine the difference between T and Σ both of which are contained in $\Sigma_{4\mu}$. Formally we define $W = (\Sigma - T)$. The space $|W|$ cannot have any handle, boundary, or Möbius bands since Σ does not have them and according to surface classification theorem Σ and $\Sigma_{4\mu}$ should have same number of them. Therefore, $\text{cl}(W)$ is a set of cylinders possibly pinched at the points where $\partial\Sigma$ and ∂T intersect. A natural deformation retraction from ∂T over these cylinders to $\partial\Sigma$ is a witness to an isotopy between T and Σ . □