Reconstruction of Surfaces with Sharp Features

James McClain

Abstract

The question of how to reconstruct a surface from an unorganized collection of point samples is of wide importance, with applications ranging from medicine to manufacturing. There are many algorithms which provide homeomorphic reconstructions, but this work provides that guarantee in the presence of edges, corners, and other sharp features. We take the approach of constructing the relative neighborhood graph of the point samples and then we show that under appropriate conditions, the union of the interiors of certain small cycles in the graph resembles the surface sufficiently to provide a homeomorphic reconstruction.

1 Introduction

The scanning and digitization of physical objects has numerous applications in areas of both industrial and academic interest. For many of those applications, having a representation in the computer of the scanned surface, not just a collection of points on the surface, is necessary. That requires surface reconstruction. Some applications include: reverse engineering, manufacturing, medicine, and computer graphics [13, 9]. In addition, affordable 3D-printing technology means that this capability is of increasing interest to ordinary consumers, as well as manufacturers.

The two major challenges that surface reconstruction algorithms face are noisy data and the reconstruction of surfaces with sharp features such as cones, corners, and sharp edges. This article addresses the second of those challenges. In this article we present a deterministic algorithm that is guaranteed to provide homeomorphic reconstruction of a wide class of surfaces, including those with the previously-mentioned traits. By "deterministic" we mean that the algorithm's guarantees are firm, rather than statistical.

1.1 Problem Statement

In this article we give the name $\mathfrak{M}$ to the set two-dimensional, closed (boundaryless), finite-area surfaces embedded in $\mathbb{R}^3$ that are reconstructible by our algorithm. The *input surface* $\mathcal{M}$ is a member of $\mathfrak{M}$ and the *sample set* $\mathcal{S} \subset \mathcal{M}$ is a finite samples set from the surface which conforms to the reconstruction algorithm's sample conditions. The task is to find a polygonal surface $\mathcal{P}$ such that $\mathcal{S} \subset \mathcal{P}$ and $\mathcal{P}$ and $\mathcal{M} \in \mathfrak{M}$ are homeomorphic. $\mathcal{P}$ is called the *output surface*.

1.2 Notation and Terminology

The k-neighborhood $\mathcal{N}(u;k) \subset \mathcal{S}$ of $u \in \mathcal{S}$ consists of the k members of the sample set that are closest to u terms of Euclidean distance.

The open ball $\mathcal{B}_X(u,r)$ is the set of all points less than distance r from the point u. The space X in which the distance is defined could be the ambient space $\mathbb{R}^3$ or the surface $\mathcal{M}$. If r is the distance between a and b, then the *lune* of those two points[23] is $\text{Lune}_X(p,q) := \mathcal{B}_X(a,r) \bigcap \mathcal{B}_X(b,r)$, the intersection of two open spheres with radius r centered at a and b, respectively.

The *relative neighborhood graph*[23, 19] is the graph that results from connecting every pair of vertices $a, b \in V$ where $\text{Lune}_X(a,b) \bigcap V = \emptyset$. In this article, we are interested in the relative neighborhood graphs $\text{RNG}_{\mathcal{M}}(\mathcal{S})$ and $\text{RNG}_{\mathbb{R}^3}(\mathcal{S})$. Distances are defined geodesically on $\mathcal{M}$ in the first and in the Euclidean fashion in the second.

The *Urquhart graph*[24, 25] is computed by starting with the intrinsic Delaunay triangulation and deleting the longest edge in every Delaunay triangle from the graph. In our notation, the Urquhart graph on the sample set on the surface is $\text{UG}_{\mathcal{M}}(\mathcal{S})$.

Symbol	Meaning
D	A topological disk
$\mathfrak{M}$	The class of surfaces reconstructible by Algorithm 1
$\mathcal{M}$	The input surface, $\mathcal{M} \in \mathfrak{M}$
$\mathcal{S}$	The sample set, $\mathcal{S} \subset \mathcal{M}$
$\mathcal{P}$	The polygonal output surface, $\mathcal{S} \subset \mathcal{P}$
$\mathcal{P}_{\text{UG}}$	A perfect reconstruction of $\mathcal{M}$ on $\mathcal{S}$ (see §3.0.1)
$\mathcal{N}(u;k)$	The k points in $\mathcal{S}$ closest to u in $\mathbb{R}^3$
$\mathcal{B}_X(u,r)$	$\{y \mid \mathrm{d}(u,y) < r\}$ where d is the distance in X
$\text{Lune}_X(a,b)$	$\mathcal{B}_X(a,\mathrm{d}(a,b)) \bigcap \mathcal{B}_X(b,\mathrm{d}(a,b))$ where d is the distance in X
$\text{RNG}_X(Y)$	The relative neighborhood graph of Y in the metric of X
$\text{Conv}_{\mathbb{R}^3}\{Y\}$	The convex hull of Y in $\mathbb{R}^3$
$\text{DT}_{\mathcal{M}}(Y)$	The Delaunay triangulation of Y on $\mathcal{M}$
$\text{UG}_{\mathcal{M}}(Y)$	The Urquhart graph of Y on $\mathcal{M}$

Table 1: Notation used in this article.

The *intrinsic Delaunay triangulation* of the sample set, $\text{DT}_{\mathcal{M}}(\mathcal{S})$, is computed by considering all $a, b, c \in \mathcal{S}$ and adding $\triangle abc$ to the triangulation if and only if there is a geodesic circle on $\mathcal{M}$ with those three points on its boundary and no sample points in its interior. An alternative construction is to add the edge ab to the graph if and only if there is some geodesic circle on $\mathcal{M}$ with those two points on its boundary and no points in its interior. In this article, we assume for convenience that every cocircular formation of sample points on the surface consists of no more than three samples.

The relative neighborhood graph and the Urquhart graph are not always the same[22], but they do coincide when the vertices are placed with sufficient uniformity. Both of these graphs arose from the pattern recognition community where they are considered to "resemble" the data. That resemblance is what motivates our use of theme here.

An *isometric cycle* in a graph[10, Chapter 21] is a cycle such that between any two vertices in the cycle, here exists a shortest path (relative to the whole graph) that only makes use of edges in that same cycle. When speaking of isometric cycles, it is necessary to have in mind a metric on the graph so that the notion of distance on it is meaningful. For the purpose of defining isometric cycles, in this article we consider all edges in a graph to have unit length, regardless of the Euclidean or geodesic distance between the endpoints.

The *symmetric difference* of two cycles is computed by taking all of the edges that appear in one cycle or the other, but not both. A *cycle basis* of a graph is a collection of simple cycles which is capable of generating any cycle in the graph via the symmetric difference operation. A 2-*basis* of a graph is a cycle basis in which every edge appears exactly twice. For example, the set of bounded faces of a planar graph is a 2-basis for that graph, an intuitive idea that will reappear later in this article.

A *sharp point* on $\mathcal{M}$ is an isolated point at which the local feature size[1] is zero. A *sharp edge* is defined to be a simple, connected curve $\gamma(t) \subset \mathcal{M}$ that is the interface between three or fewer locally-flat areas. Each sharp edge is of measure zero on $\mathcal{M}$, and as with sharp points, $\text{LFS}(\gamma(\text{t})) = 0$ at every t in the domain of the curve. We define a *corner* as the common meeting point of a constant number of sharp edges. A *smooth point* on $\mathcal{M}$ is one that is neither a sharp point, nor a corner, nor on a sharp edge.

We make frequent use of the words like *trivial* and *non-trivial* with respect to loops on surfaces. In this article, those words are used in the homological sense: a loop on a surface is trivial if it can be continuously deformed into a point on the surface, and non-trivial otherwise. We also refer frequently to *homeomorphisms* between spaces. A homeomorphism is a bicontinuous bijection, and if A and B are homeomorphic, we use that notation $A \cong B$. Please see [14] for more on these topics.

[1] The concept of local feature size is defined and discussed in Section 2.

1.3 Organization

This article is organized as follows. In Section 2 we give a brief survey of previous work in this area. Section 3 introduces the sample conditions required by our reconstruction algorithm. A description of the algorithm is given in Section 4. Analysis of the algorithm is given in Section 5. In Section 6 we present the results of various experiments and the article concludes in Section 7.

2 Previous Work

The surface reconstruction algorithm described in this article is a deterministic algorithm which can reconstruct surfaces with sharp features. For purposes of comparison, we are mostly interested in two main streams of previous work: the *deterministic surface reconstruction algorithms*, such as the work of Amenta and Bern and various *meshing algorithms*, some of which can handle sharp surfaces. There are also some previous deterministic algorithms that can handle surfaces with sharp features, and we will compare our algorithm to those, as well.

At the root of the deterministic algorithms and intersecting the meshing algorithms is the work presented in [11]. That article by Edelsbrunner and Shah introduces the idea of the closed ball property. In the context of a boundaryless 2-manifold embedded in $\mathbb{R}^3$ (the setting for this article) the closed ball property can be described as follows.

A sample set $\mathcal{S}$ has *the closed ball property* if for every $T \subset \mathcal{S}$ of size $|T| = 3 - \ell$, the intersection of $\mathcal{M}$ with the Voronoi cells of $u \in T$ is either empty or is a closed ℓ-ball. If that property is possessed by the sample set, then the intersection of the Voronoi cells in $\mathbb{R}^3$ of $\mathcal{S}$ with $\mathcal{M}$ gives a set of triangles whose union is homeomorphic to $\mathcal{M}$.

An example of an algorithm which makes use of the closed ball property is given by Amenta and Bern in [1]. There, they show that sample conditions based on the *medial axis* (the closure of the set of points in $\mathbb{R}^3$ with more than one nearest-neighbor in $\mathcal{M}$) can be used to generate a sample set that has the closed ball property. That is typically done by considering the *local feature size* of a point $u \in \mathcal{M}$, the distance from u to the medial axis, and requiring inversely-proportionally more sample points around u when the local feature size is smaller. This idea was implemented by Amenta, Bern, and Kamvysselis in [2] and carried further in [3]. Because the medial axis can actually touch the surface when sharp features are present, these types algorithms are per force unable to provably reconstruct surfaces with sharp features (because no finite sample is adequate).

There are at least two meshing algorithm which have guarantees for surfaces containing sharp features. One of these algorithms was given in [7] by Chazal and Lieutier and another is [8] by Cheng, Dey, and Ramos. Meshing and surface reconstruction are closely related, but they differ in that meshing problems frequently assume that the original surface is "known" and that the task is to create a sample set from which topological invariants can be recovered. This can be contrasted to surface reconstruction where the only thing that is know about the original surface is the sample set (which presumably meets some sample conditions) and the goal is recover the unknown surface.

We are also aware of two previous deterministic algorithms which come with topological guarantees for the reconstruction of sharp surfaces.

One of those algorithms is [5] by Boissonnat and Oudot. That algorithm places a restriction on the difference in the height of two points compared to their distance apart when viewed relative to a particular planar projection. Because there is a fixed limit on this factor, there are many shapes, such as a cone with a sharp angle, that violate this requirement.

The second algorithm is [6] by Chazal, Cohen-Steiner, and Lieutier. That article gives a sampling theory for handling noisy sample sets and even those taken from sharp shapes. In that paper it is shown how to create a reconstruction which is isotopic to the input. The work presented in that article has a different goal than what is presented here: the output of their reconstruction process is an offset – a set that contains the surface – rather than a surface itself.

The work that we describe in this article takes a strongly graph-centric approach to surface reconstruction. The works of Mencl [16] and Mencl and Müller [17, 18] are an example of a previous graph-oriented approaches. In those works, the Euclidean minimum spanning tree of the sample set is computed, extended from a tree into a graph that has loops, then various features in the graph (on the surface) are recognized and handled according to particular sets of rules. Like the present article, these works are concerned with being able to reconstruct surfaces with sharp features and non-convex areas. Our work can be contrasted in that it provides theoretical guarantees of the reconstruction quality. Our algorithm is also somewhat simpler, in that our algorithm behaves the same in every neighborhood rather than applying different rules to different features.

Another example of a graph-based approach is given in [21] by Tewari, et al. In that work, the authors assume that the surface belongs to a particular class of shapes (those that are homeomorphic to a donut) and use that foreknowledge to compute a parametrization of the sample set which is used to triangulate the sample set. This is achieved by computing the k-nearest neighbor graph of the sample set; under appropriate assumptions about the sample set, homologically-trivial isometric loops are known to be of relatively small circumference whereas non-trivial loops are relatively large. This allows the authors to infer where the hole in the surface is and proceed from there. Our work does not assume any particular genus for the original surface and also remains graph-oriented throughout rather than using an initial graph to compute an explicit parametrization.

Although not technically a surface reconstruction algorithm, a well-known method for parametrization and embedding of data is Isomap[20]. The goal of that algorithm is non-linear dimensionality reduction while to the extent possible preserving distances.

The first step of the Isomap algorithm is similar to that of [16], [21] and indeed to our own algorithm: the initial step is to compute a Euclidean graph which is hopefully of low local distortion with respect to the true metric. In the case of [16] the graph is an augmented minimum spanning tree, in [21] and Isomap it is a Euclidean k-nearest neighbor graph, and in our case it is the relative neighborhood graph.

Isomap and [21] maintain their correspondence in the second step by computing an embedding with explicit coordinates. For Isomap this is done by computing a distance matrix and projecting it into a lower-dimensional space and in Tewari, et al. it is done by solving harmonic equations on the surface. The work of Mencl and our work diverge here, the former seeking to recognize features and the latter mapping combinatorial features in the relative neighborhood graph to features on the surface.

3 Sample Conditions

A sample set $\mathcal{S}$ is considered to be adequate if there is a constant integer k and a small constant ε so that conditions listed below are met.

The requirements are given notationally first, then described in detail in the three subsections that follow. For the discussion below, we define $\mathcal{M}_u := \text{Conv}_{\mathbb{R}^3}\{\mathcal{N}(u;k)\} \bigcap \mathcal{M}$. $\mathcal{M}_u$ is the intersection of the convex hull of the k-neighborhood of u with the surface, and from a mental-picture standpoint can be thought of the as the neighborhood of u on $\mathcal{M}$. We also make use of the object $\mathcal{P}_{\text{UG}}$ which is not defined until Section 3.0.1. For right now, it can be thought of as a polygonal object whose faces are the face cycles of the Urquhart graph $\text{UG}_{\mathcal{M}}(\mathcal{S})$ on the surface $\text{DT}_{\mathcal{M}}(\mathcal{S})$.

The sample conditions are as follows.

I **Possibility.** $\bigcup_{\triangle \in \text{DT}_{\mathcal{M}}(\mathcal{S})} \triangle = \mathcal{M}$ and $\mathcal{P}_{\text{UG}} \cong \mathcal{M}$. ($\mathcal{P}_{\text{UG}}$ is defined in Section 3.0.1.)

II **Density.**

- i For every $u \in \mathcal{S}$, $\mathcal{M}_u \cong D$.
- ii If every $x \in \mathcal{M}_u$ is a smooth point, then $\text{d}_{\mathcal{M}}(y,z) < (1+\varepsilon)\|y-z\|_{\mathbb{R}^3}, \forall y,z \in \mathcal{M}_u$.
- iii There exists a basis $\{\square_i\}$ for the subspace of cycles in $\text{UG}_{\mathcal{M}}(\mathcal{S})$ that are trivial on $\mathcal{M}$ such that $\square \in \{\square_i\} \implies \square \subset \mathcal{N}(u;k)$ for some u.
- iv If $\hexagon$ is an isometric cycle on $\text{UG}_{\mathcal{M}}(\mathcal{S})$ that is nontrivial on $\mathcal{M}$, then $|\hexagon| > 33k$.

III **Uniformity.**

- i The degree of $u \in \mathcal{S}$ is no more than 33 in the graph $\text{DT}_{\mathcal{M}}(\mathcal{S})$.
- ii Every $\triangle abc \in \text{DT}_{\mathcal{M}}(\mathcal{S})$ has a unique longest-edge, and the indices a, b, c are chosen so that ac is that edge. The edge lengths are bounded so that

$$\frac{\text{d}_{\mathcal{M}}(a,b)}{\text{d}_{\mathcal{M}}(a,c)} < \frac{\sqrt{3}(1-\varepsilon)}{2}$$

and

$$\frac{\mathrm{d}_{\mathcal{M}}(b,c)}{\mathrm{d}_{\mathcal{M}}(a,c)} < \frac{\sqrt{3}(1-\varepsilon)}{2}$$

are both true.

In the three conditions above and the three subsections below, we sometimes move rhetorically back and forth between cycles in $\mathrm{UG}_{\mathcal{M}}(\mathcal{M})$ and loops on $\mathcal{M}$. It is possible to define a mapping $f(\cdot)$ between every simple isometric cycle $\hexagon$ of adjacent sample points $a_0a_1a_2\cdots a_{n-1}a_0$ in $\mathrm{UG}_{\mathcal{M}}(\mathcal{M})$ and a collection of loops $\bigcirc_m$ that results from every possible choice of geodesics connecting a_i to $a_{(i+1 \mod n)}$ on $\mathcal{M}$. As part of the conditions above, every loop in $f(\hexagon)$ on $\mathcal{M}$ belongs to the same homotopy equivalence class.[2] That allows us to speak about $f(\hexagon)$ in the singular, and talk about cycles as if they were on $\mathcal{M}$.

3.0.1 Possibility

Condition I requires that $\mathcal{S}$ is adequate, in principle, to allow a homeomorphic reconstruction of $\mathcal{M}$. The goal of our algorithm is to produce a reconstructuion, so it is necessary that that union of the intrinsic Delaunay triangles of the sample set $\mathcal{S}$ (the gold standard reconstruction) is equal to the original surface. In addition to that, when the intrinsic Delaunay triangulation is transformed into the Urquhart graph by removing the longest edge of every triangle, we want the faces generated thereby to provide a homeomorphic reconstruction.

If the basis consisting of small isometric cycles on $\mathrm{UG}_{\mathcal{M}}(\mathcal{S})$ is assigned to the variable $\{\square_i\}$ on line 3 of Algorithm 1, proceeding from there gives a homeomorphic reconstruction $\mathcal{P}_{\mathrm{UG}} \cong \mathcal{M}$. If such an oracular process is not able to provide a homeomorphic reconstruction, then we will not attempt to do so with an algorithm.

3.0.2 Density

The first two parts of Condition II are jointly comparable to the medial-axis sample conditions found in some of the works referenced in Section 2. In the event that $\mathcal{M}_u$ does not contain a sharp point, Condition II-ii says how dense the sample must be in a smooth area: dense enough to keep the Euclidean and geometric distances between points in $\mathcal{M}_u$ within a factor of $(1+\varepsilon)$ of each other. The requirement that $\mathcal{M}_u \cong D$ in Condition II-i makes sure that the sample set is sufficiently dense to ensure that there is always some distance between distinct parts of the surface. Even when the surface is very flat, if there is another region close by (if the local feature size is low despite low curvature), the sample set must be dense enough to keep neighborhoods on one part of the surface well away from those on the other part.

The last two conditions concern the sizes of homologically trivial and non-trivial isometric cycles. Those conditions require that the sample set to be such that any holes can be readily distinguished from face cycles. Condition II-iii requires that a basis for the trivial cycles on $\mathcal{M}$ can be found inside of the neighborhoods $\mathcal{N}(u;k)$, that the elements of the cycle basis are all fairly small. Condition II-iv requires that cycles that are nontrivial on $\mathcal{M}$ are of fairly large circumference.

These requirement can be compared to naturally-arising facts about isometric cycle sizes in k-nearest neighbor graphs described in [12] and applied in [21]. In those works, the differences in length between trivial and non-trivial isometric cycles arise naturally from a sufficiently dense sample set on a smooth surface. Since we are concerned with non-smooth surfaces, we cannot use those results, but it is realistic to expect this phenomenon anyway.

3.0.3 Uniformity

Condition III-i puts a bound on the degree of each vertex in the Delaunay triangulation. If the sample set is to be described as "uniform", we would not expect that any sample point to meet more than a constant number of Delaunay triangles.

Condition III-ii puts bounds on edge lengths of triangles in $\mathrm{DT}_{\mathcal{M}}(\mathcal{S})$. Those bounds hold those triangles away from being equilateral. That might be seem to be the opposite of uniformity, but we have the following mental picture in mind.

We imagine that the samples a and b lie on a scan line generated by the laser scanner, that the scan line goes from down to up, and that c lies to the right of b on an adjacent and approximately parallel scan line. Under that scenario, the geodesic segment ab connects b with the point above it on its same scan line, the edge bc connects b with the sample c that

[2] In other words, one can imagine that Conditions II-iii and II-iv are respectively quantified over all $f(\square)$ and all $f(\hexagon)$. Alternatively, one can imagine that $\mathcal{M}_u$ is equal to the convex closure of $\mathrm{Conv}_{\mathbb{R}^3}\{\mathcal{N}(u;k)\} \bigcap \mathcal{M}$.

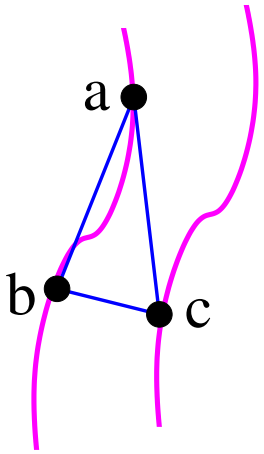

Figure 1: The scanlines are shown in magenta, the sample points are in black, and the Delaunay triangle is shown in blue. The angle $\angle abc$ is close to $\frac{\pi}{2}$.

is closest to it on the scan line to its right, and the long edge ac connects the point to the right of b with the point above b. Here, "uniformity" means that $\mathcal{S}$ is closer to a square lattice on $\mathcal{M}$ than a hexagonal one (the latter of which would suggest equilateral triangles). Please see Figure 1 for an illustration.

When the sample set is uniform in the sense described above, the Urquhart graph and the relative neighborhood graph coincide. We can prove that statement by divding it into Lemmas 1 and 2 and recombining the lemmas into Theorem 1. The first lemma shows that under the uniformity assumption given in Condition III-ii, the relative neighborhood graph contains the Urquhart graph.

Lemma 1. *If $\mathcal{S}$ obeys Condition III-ii, then* $\mathrm{UG}_{\mathcal{M}}(\mathcal{S}) \subseteq \mathrm{RNG}_{\mathcal{M}}(\mathcal{S})$ *(the subset symbol is with respect to the edges present in the graphs).*

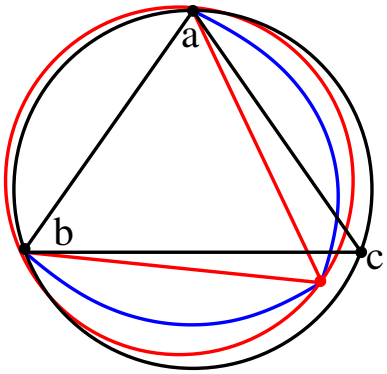

Figure 2: $\triangle abc$ and $\bigcirc$ are shown in black. The equilateral triangle, and $\circ$ are shown in red. The right half of the intrinsic lune of a and b is shown in blue.

Proof. We will assume that $ab \in \mathrm{UG}_{\mathcal{M}}(\mathcal{S})$ and show that $ab \in \mathrm{RNG}_{\mathcal{M}}(\mathcal{S})$. With that assumption, ab is by definition a Delaunay edge, but not the longest edge of either intrinsic Delaunay triangle that it meets. Because ab meets two such triangles, there are two intrinsic Delaunay circumcircles containing a and b whose interiors contain no points in $\mathcal{S}$.

Consider the Delaunay triangle with vertices abc where c is on the right side of the geodesic that goes through a and b on $\mathcal{M}$. If we change scale so that $\mathrm{d}_{\mathcal{M}}(a, c) = 2$, then the radius of the circumcircle of the Delaunay triangle is at least one.

Algorithm 1 The reconstruction algorithm.

Require: $\mathcal{S}, k$
1: RNG $\leftarrow$ LocalizedRNG($\mathcal{S}, k$)
2: $\{\pentagon_i\} \leftarrow$ TrivialIsometricCycles(RNG, k)
3: $\{\square_i\} \leftarrow$ Backtracking($\emptyset, \{\pentagon_i\}, k$)
4: **return** Triangulate($\{\square_i\}$)

(The length of ac bounds the diameter of the circumcircle from below and the radius is half of that.) We call the cirumcircle of the Delaunay triangle $\bigcirc$.

Next, we draw an equilateral triangle where two of its vertices are a and b and the third vertex is on the same side of ab as c. This equilateral triangle is named $\triangle$ and its circumcircle is called $\circ$. We know that to the right of ab, $\circ$ is strictly contained within $\bigcirc$ because we have supposed in Condition III-ii that $\mathrm{d}_{\mathcal{M}}(a,b) < \sqrt{3}$, which implies that the radius of $\circ$ is strictly less than one. (For the circumcircle of an equilateral triangle, that radius is $\frac{\sqrt{3}}{3}$ times the length of one of the sides.)

The half of $\mathrm{Lune}_{\mathcal{M}}(a,b)$ to the right of ab is composed of two circular segments which begin at a and b (respectively) and which terminate at the third vertex of $\triangle$. Those circular segments are on the boundaries of respective geodesic circles centered at a and b with radii of $\sqrt{3}$. That means that the half of $\mathrm{Lune}_{\mathcal{M}}(a,b)$ to the right of ab is contained within $\circ$.

We concluded earlier that the $\circ$ is contained within $\bigcirc$. The latter encompasses no points in $\mathcal{S}$ by definition, so that $\circ$ and the right half of the lune do not either. This same argument can be applied to the Delaunay triangle on the left side of ab. Based on that, we can conclude that $ab \in \mathrm{RNG}_{\mathcal{M}}(\mathcal{S})$. Please see Figure 2 for a picture. □

The next lemma shows the converse of Lemma 2.

Lemma 2. *If $\mathcal{S}$ obeys Condition III-ii, then* $\mathrm{RNG}_{\mathcal{M}}(\mathcal{S}) \subseteq \mathrm{UG}_{\mathcal{M}}(\mathcal{S})$.

Proof. We will prove that $ac \notin \mathrm{UG}_{\mathcal{M}}(\mathcal{S}) \implies ac \notin \mathrm{RNG}_{\mathcal{M}}(\mathcal{S})$. We start by citing [4], wherein it is shown that $\mathrm{RNG}_{\mathcal{M}}(\mathcal{S}) \subseteq \mathrm{DT}_{\mathcal{M}}(\mathcal{S})$. That means that all edges ac which are not members of some Delaunay triangle are automatically excluded from $\mathrm{RNG}_{\mathcal{M}}(\mathcal{S})$. The only thing left to do is show that the longest edge ac of some Delaunay triangle $\triangle abc$ is not found in the relative neighborhood graph.

We will rescale so that $\mathrm{d}_{\mathcal{M}}(a,c) = 2$. Using Condition III-ii, it can be seen that $\mathrm{d}_{\mathcal{M}}(a,b) < \sqrt{3} < 2$ and the same for $\mathrm{d}_{\mathcal{M}}(b,c)$. The sample b is evidently in the $\mathrm{Lune}_{\mathcal{M}}(a,c)$ of a and c so that $ac \notin \mathrm{RNG}_{\mathcal{M}}(\mathcal{S})$. □

With the two previous lemmas proven, the following theorem can be stated without further proof.

Theorem 1. *If $\mathcal{S}$ obeys Condition III-ii, then* $\mathrm{RNG}_{\mathcal{M}}(\mathcal{S}) = \mathrm{UG}_{\mathcal{M}}(\mathcal{S})$.

This theorem is helpful because it allows us to refer interchangeably to the intrinsic Urquhart graph and the intrinsic relative neighborhood graph when $\mathcal{S}$ obeys the conditions and when $\mathcal{M}$ is smooth (or even on smooth portions of a not-entirely-smooth $\mathcal{M}$).

4 Algorithm

The reconstruction algorithm is fairly simple. Line 1 of Algorithm 1 computes a localized version of the Euclidean relative neighborhood graph and assigns it to the variable RNG. The graph is localized in the sense that it is the union of the $\mathbb{R}^3$ relative neighborhood graph of every neighborhood, but via Conditions I and II-iii that is the same as the standard relative neighborhood graph for our purposes. Next, the trivial isometric cycles in that graph are found on line 2 and assigned to the variable $\{\pentagon_i\}$. On line 3, a backtracking strategy is used to compute a 2-basis for the graph (a closed collection of faces), and that is assigned to the variable $\{\square_i\}$. Finally, on line 4 a triangulation of the faces into a piecewise linear complex is returned.

As stated before in Section 3.0.1, the intrinsic Delaunay triangulation of the sample set on the surface is gold-standard to which any reconstruction algorithm aspires. Since $\mathcal{M}$ is unknown, this ideal can never be achieved, but we try to get as close as possible. In Section 1.2, we mentioned that intrinsic Urquhart graph is derived from the intrinsic Delaunay triangulation

Algorithm 2 The localized relative neigborhood graph.

function LocalizedRNG($\mathcal{S}, k$)
  $G \leftarrow \emptyset$ ▷ An initially-empty graph.
  for all $u \in \mathcal{S}$ **do**
    $N \leftarrow \mathcal{N}(u; k)$ ▷ Compute the neighborhood of u.
    $G' \leftarrow$ RelativeNeighborhoodGraph(N) ▷ Compute the R.N.G. of the neighborhood.
    $G \leftarrow G \bigcup G'$ ▷ Add that to the graph.
  end for
  return G
end function

Algorithm 3 Find potential face cycles.

function TrivialIsometricCycles(RNG, k)
  $\{\pentagon_i\} \leftarrow \emptyset$ ▷ An initially-empty set of small isometric cycles.
  for all $u \in \mathcal{S}$ **do** ▷ For every sample point u ...
    for all $\pentagon \in$ IsometricCyclesAt(u, RNG) **do** ▷ ... for every isometric cycle in RNG meeting u ...
      if $\pentagon \subset \mathcal{N}(u; 33k)$ for some u **then** ▷ ... if that cycle is small add it.
        $\{\pentagon_i\} \leftarrow \{\pentagon_i\} \bigcup \pentagon$
      end if
    end for
  end for
  return $\{\pentagon_i\}$
end function

by removing the longest edge of every Delaunay triangle. Theorem 1 allows us to identify the intrinsic Urquhart graph with the intrinsic relative neighborhood graph. We show later, in Theorem 2, that intrinsic and Euclidean relative neighborhood graphs coincide when $\mathcal{S}$ obeys the sample conditions and $\mathcal{M}$ is smooth. We then show in Theorems 3 and 4 that the similarity is sufficient to allow reconstruction even when $\mathcal{M}$ has sharp points and sharp edges.

From MacLane's Planarity criterion[15], we know that a graph is planar if and only if it has a 2-basis. In fact, the elements of that basis can be identified with the faces of the graph, which are in-turn small isometric cycles. Since $\mathrm{UG}_{\mathcal{M}}(\mathcal{S})$ is (locally) planar by Condition I, we can compute a reconstruction of the surface by looking for a 2-basis among the small isometric cycles of RNG. (Right now that is just a hope, but it will be proven later.) Condition II-iii guarantees that all of the cycles in the basis of $\mathrm{UG}_{\mathcal{M}}(\mathcal{S})$ (and also of RNG) can be found in k-neighborhoods. The trivial isometric cycles in RNG are found by the subroutine shown in Algorithm 3 (the are known to be trivial by condition II-i).

The backtracking subroutine (shown in Algorithm 4) works in the following way. The basis $\{\square_i\}$ (the collection of face cycles) is initially empty. A new cycle $\pentagon$ is chosen to add to that collection subject to the constraints that i) each of its edges of $\pentagon$ meets the existing basis zero or one times and ii) at least one of its edges of $\pentagon$ meets the existing basis once.

That ensures that the basis uses each of its edges zero, one, or two times (we call that *the backtracking invariant*). Let $\mathcal{P}'$ be the surface implied by a closed set of faces $\{\square_i\}$. Before $\{\square_i\}$ is returned to the top-level caller, that basis is checked to make sure that $\mathrm{Conv}_{\mathbb{R}^3}\{\mathcal{N}(u; k)\} \bigcap \mathcal{P}' \cong D$ for every $u \in \mathcal{S}$. That is *the backtracking success condition* and is found on on line 6. Consider a set of cycles that is maximal with respect to those criterion; if the cycles are viewed as faces, then their union is a closed surface.

After a valid solution is found, the faces are triangulated on line 4 of Algorithm 1. Triangulation is done by finding a collection of triangles associated with each cycle that is closed (in the sense used above) except at the boundary of the collection, which is the cycle. Those collections of triangles are chosen so that the interior of each triangle in the complete set of triangles is disjoint from that of all others.

Algorithm 4 Find a closed, manifold collection of faces.

▷ The set $\{\square_i\}$ is the collect of face cycles.
▷ The set $\{\pentagon_i\}$ is the collection of trivial isometric cyclces.
function Backtracking($\{\square_i\}, \{\pentagon_i\}, k$)
    candidate-p $\leftarrow \lambda(\pentagon, \text{TestInvariant}(\pentagon \bigcup \{\square_i\}))$ ▷ Function to see if $\pentagon \bigcup \{\square_i\}$ meets the invariant.
    $\{\pentagon_i'\} \leftarrow$ Filter(candidate-p, $\{\pentagon_i\}$) ▷ Collection of possible new face cycles.
    if $\{\pentagon_i'\} = \emptyset$ **then** ▷ If there are no candidates either ...
        if SuccessCondition($\{\square_i\}$) **then** ▷ ... immediately return a solution to the top-level caller...
            nonlocal-return $\{\square_i\}$
        else ▷ ... or terminate this branch of the search.
            return $\emptyset$
        end if
    else ▷ If there are candidates, try them.
        for all $\pentagon \in \{\pentagon_i'\}$ **do**
            backtracking($\{\square_i\} \bigcup \{\pentagon\}, \{\pentagon_i\} - \{\pentagon\}$)
        end for
    end if
end function

5 Analysis

We will start the analysis by showing that Algorithm 1 successfully reconstructs smooth surfaces. That is done in Theorem 2 via its two principal parts Lemmas 4 and 5. Before addressing Lemma 4, we need to put a bound on how close two adjacent samples in the intrinsic Delaunay triangulation can be.

Lemma 3. *When ε is sufficiently small, $\dfrac{\mathrm{d}_{\mathcal{M}}(x,y)}{\mathrm{d}_{\mathcal{M}}(x,z)} > \delta > 5\varepsilon$ for all $x \in \mathcal{S}$ and any $y, z \in \mathcal{S}$ adjacent to x in* $\mathrm{DT}_{\mathcal{M}}(\mathcal{S})$.

Proof. By Condition I, if we superimpose $\mathrm{DT}_{\mathcal{M}}(\mathcal{S})$ onto $\mathcal{M}$, there is a fan of Delaunay triangles around x. According to Condition III-i, the degree of that vertex in $\mathrm{DT}_{\mathcal{M}}(\mathcal{S})$ is bounded by 33. Condition III-ii implies that the edge-length-ratio of any two edges of any Delaunay triangle is bounded by $\frac{2-\sqrt{3}}{\sqrt{3}} > \frac{3}{20}$. Since the ratio $\frac{\mathrm{d}_{\mathcal{M}}(x,y)}{\mathrm{d}_{\mathcal{M}}(x,z')}$ is greater than a constant when $\triangle xyz' \in \mathrm{DT}_{\mathcal{M}}(\mathcal{S})$, $\frac{\mathrm{d}_{\mathcal{M}}(x,y)}{\mathrm{d}_{\mathcal{M}}(x,z)} \leq 5\varepsilon$ would require a fan around x composed of at least $\Omega\big(\log \frac{1}{\varepsilon}\big)$ triangles. Since ε is small, that quantity overwhelms the constant degree of x, which is a contradiction. At the same time, δ can be no less than $\frac{3}{20}$ raised to the power of the greatest degree of any vertex. □

Note that the ε in Lemma 3 is the same one that appears in the sample conditions. The purpose of this lemma is to establish a bound on the relative lengths of the Delaunay edges incident on x. That bound is given in terms of ε for later convenience.

With Lemma 3 in place, it is now possible to move on to Lemmas 4 and 5. The next lemma is analogous to Lemma 1, except that it deals with the $\mathbb{R}^3$ relative neighborhood graph instead of the intrinsic one.

Lemma 4. *If $\mathcal{M}$ is smooth and $\mathcal{S}$ obeys the sample conditions given in Section 3, then* $\mathrm{UG}_{\mathcal{M}}(\mathcal{S}) \subseteq \mathrm{RNG}_{\mathbb{R}^3}(\mathcal{S})$.

Proof. We will show that $ab \in \mathrm{UG}_{\mathcal{M}}(\mathcal{S}) \implies ab \in \mathrm{RNG}_{\mathbb{R}^3}(\mathcal{S})$. As in Lemma 1, we will assume that $\triangle abc \in \mathrm{DT}_{\mathcal{M}}(\mathcal{S})$ and that ac is the long edge of that triangle. However, this time we will notice that Condition III-ii allows us to scale so that $\mathrm{d}_{\mathcal{M}}(a,c) = \frac{2}{(1-\varepsilon)}$ (instead of just 2) and $\mathrm{d}_{\mathcal{M}}(a,b) < \sqrt{3}$.

We will consider a circumcircle around the Delaunay triangle. That circumcircle has radius at least $\frac{1}{(1-\varepsilon)}$ and we will call it $\bigcirc$. We will also consider the equilateral triangle with edge ab and its third vertex to the right (respectively, left) of ab. That triangle is called $\triangle$.

Since $\mathcal{M}$ has no sharp points, the difference between the geometric and Euclidean distances between any two points in $\mathcal{M}_u$ is bounded by a factor of $(1+\varepsilon)$ in Condition II-ii. For purposes of shorthand, we will define a new object $\mathcal{M}_{a,b} := \mathrm{Lune}_{\mathbb{R}^3}(a,b) \bigcap \mathcal{M}$. Even taking distortion into account, we know that the half of $\mathcal{M}_{a,b}$ to the right (left) of ab

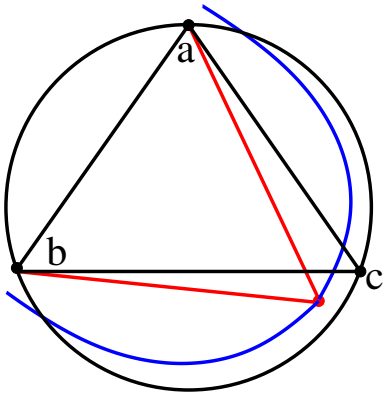

Figure 3: This figure is similar to Figure 3. The right half of $\mathcal{M}_{a,b}$ is shown in blue.

fits between two circular segments i) with their respective centers at a and b, ii) with radius $< \sqrt{3}(1+\varepsilon)$, iii) which go from respective points both within $4\varepsilon \mathrm{d}_{\mathcal{M}}(a,b)$ units (in the worst case) of b or a, and iv) which terminate at the third point of $\triangle$.

With that, the only places where the right half of $\mathcal{M}_{a,b}$ might not be contained in $\bigcirc$ are two small areas within $4\varepsilon \mathrm{d}_{\mathcal{M}}(a,b)$ units of a and b, $respectively$.

Consider $d \in \mathcal{S}$ near a so that $d \in \mathcal{M}_{a,b}$, and $d \notin \bigcirc$. The 1-nearest neighbor graph is a subgraph of the Delaunay triangulation, so if d is distinct from a, then Lemma 3 says that $\mathrm{d}_{\mathcal{M}}(d,a) > 5\varepsilon \mathrm{d}_{\mathcal{M}}(a,b)$. However, $5\varepsilon \mathrm{d}_{\mathcal{M}}(a,b)$ is beyond the $4\varepsilon \mathrm{d}_{\mathcal{M}}(a,b)$ bound just established, so claiming that d is distinct from a leads to a contradiction. With that, we conclude that the portion of $\mathcal{M}_{a,b}$ to the right (left) of ab does not intersect with $\mathcal{S}$. That implies that $ab \in \mathrm{RNG}_{\mathbb{R}^3}(\mathcal{S})$. Please see Figure 3 for a picture. □

Now we will show the converse of Lemma 5, which says that the set of edges in the Euclidean relative neighborhood graph is contained within that of the intrinsic Urquhart graph when the sample set is taken according to the sample conditions.

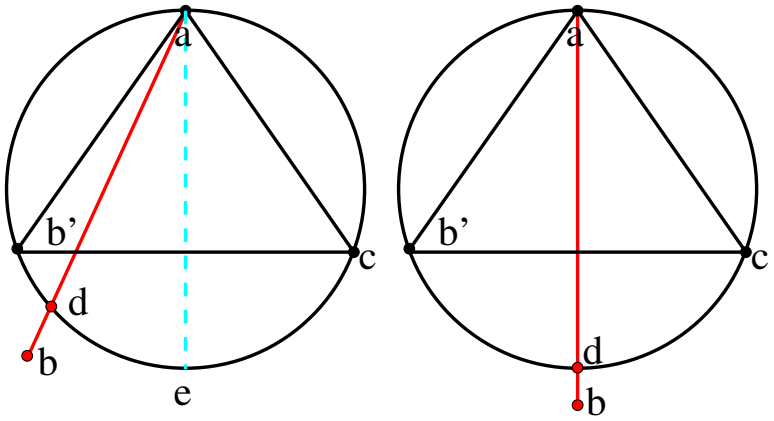

(a) (b)

Figure 4: The first subfigure shows the worst-case (greatest) distance between a and b' as compared to the distance between a and d. The second subfigure show the worst-case distance between b and b' relative to the distance between a and d.

Lemma 5. *If $\mathcal{M}$ is smooth and $\mathcal{S}$ obeys the sample conditions given in Section 3 (with sufficiently small ε), then* $\mathrm{RNG}_{\mathbb{R}^3}(\mathcal{S}) \subseteq \mathrm{UG}_{\mathcal{M}}(\mathcal{S})$.

Proof. We will show that $ab \notin \text{UG}_{\mathcal{M}}(\mathcal{S}) \implies ab \notin \text{RNG}_{\mathbb{R}^3}(\mathcal{S})$. If we assume that $ab \notin \text{UG}_{\mathcal{M}}(\mathcal{S})$, then that edge must either be the longest edge in some $\triangle abc \in \text{DT}_{\mathcal{M}}(\mathcal{S})$ or alternatively it is not an edge in any Delaunay triangle in $\text{DT}_{\mathcal{M}}(\mathcal{S})$.

The first of those cases leads to the following reasoning. If ab is the longest edge in an intrinsic Delaunay triangle, and if $\text{d}_{\mathcal{M}}(a,b)$ is scaled to length 2, then by Conditions II-i and III-ii, the distance to c from both a and b is $< \sqrt{3}(1-\varepsilon)$. At worst, $\|a-b\|_{\mathbb{R}^3} \geq \frac{2}{(1+\varepsilon)}$ so that c is well inside of $\text{Lune}_{\mathbb{R}^3}(a,b)$.

We will now deal with the situation where ab is not part of any intrinsic Delaunay triangle. We know that every point $\mathcal{S}$ is surrounded by a fan of Delaunay triangles on $\mathcal{M}$. We will focus on a Delaunay triangle $\triangle ab'c$ and assume without loss of generality that the geodesic segment between a and b intersects the geodesic circumcircle of $\triangle ab'c$ on the circular segment between b' and c at point d. The point d must exist, otherwise b would be inside of the circumcircle of $\triangle ab'c$, contradicting its status as a Delaunay triangle.

Let ae be a diameter of the circumcircle. Once again without loss of generality, we will assume that a geodesic going between a and d separates b' from e. If $b = b'$ then $ab = ab'$, contradicting the assumption that ab is not part of any Delaunay triangle. If $d = b'$ then b' is on a geodesic between a and b, so that even modulo distortion, $b' \in \text{Lune}_{\mathbb{R}^3}(a,b)$.

Therefore, the only case left to deal with is that of $b' \neq d$ and $b' \neq b$, so that b is outside of the circumcircle on a segment sandwiched between ab' and the diameter ae. Since b and b' are distinct, Lemma 3 says that $\text{d}_{\mathcal{M}}(b,b') > \delta \text{d}_{\mathcal{M}}(a,b')$. (That is because the 1-nearest neighbor graph is a subgraph of the Delaunay triangulation, so if b is less than that distance from b' then there must be a sample adjacent to b' in the Delaunay triangulation that is less than that distance.) Because $\frac{\text{d}_{\mathcal{M}}(b,b')}{\text{d}_{\mathcal{M}}(a,b')}$ is bounded away from zero, the angle $\angle b'ab$ is bounded away from zero, as well. That gives

$$\text{d}_{\mathcal{M}}(a,b) > \text{d}_{\mathcal{M}}(a,b') \implies \|a-b\|_{\mathbb{R}^3} > \|a-b'\|_{\mathbb{R}^3}$$

because ε is small. Please see Figure 4a for an illustration.

We also have

$$\|a-b\|_{\mathbb{R}^3} > \|b-b'\|_{\mathbb{R}^3}$$

because the worst-case ad (from the perspective of wanting $\|a-b\|_{\mathbb{R}^3} > \|b-b'\|_{\mathbb{R}^3}$) is if ad is a diameter of the circumcircle of $\triangle ab'c$. Please see Figure 4b for an illustration.

All of the above-stated facts combine to imply that $b' \in \text{Lune}_{\mathbb{R}^3}(a,b)$ so that $ab \notin \text{RNG}_{\mathbb{R}^3}(\mathcal{S})$. □

We now have the following theorem concerning surfaces with sharp edges.

Theorem 2. *If $\mathcal{M}$ is smooth, $\mathcal{S}$ obeys the sample conditions given in Section 3 (with sufficiently small ε), and $\mathcal{P}$ is the output from Algorithm 1, then $\mathcal{P} \cong \mathcal{M}$.*

Proof. It is assumed in Condition II-iii that there is a basis $\{\square_i\}$ for the subspace of cycles on $\text{UG}_{\mathcal{M}}(\mathcal{S})$ that are trivial on $\mathcal{M}$ and that basis has the property that every $\square \in \{\square_i\}$ is contained in the k-neighborhood of some u. Lemmas 4 and 5 combine to give $\text{RNG} = \text{RNG}_{\mathbb{R}^3}(\mathcal{S}) = \text{UG}_{\mathcal{M}}(\mathcal{S})$, so such a basis exists for RNG (the variable from line 1 of Algorithm 1). By Condition II-iii that basis must be a subset of the set held in variable $\{\hexagon_i\}$ (which is assigned on line 2). The basis is then discovered by the backtracking process on line 3. From there that point, the output of the algorithm is as if $\{\square_i\}$ was set equal to the basis for the trivial cycles in $\text{UG}_{\mathcal{M}}(\mathcal{M})$, so that $\mathcal{P} = \mathcal{P}_{\text{UG}} \cong \mathcal{M}$. □

Having established parity with existing methods, we will now move on to sharp features. In Section 1.1 we were vague about the extent of the set $\mathfrak{M}$, the class of manifolds homeomorphically reconstructible by the algorithm. We will now flesh that out.

Theorem 3. *Let $\mathcal{S}$ be a sample set that follows the conditions given in Section 3 (with sufficiently small ε) and let $\mathcal{M}$ be a surface that is smooth except for one sharp edge $\gamma(t)$. If*

- *The tangent $\gamma'(t)$ exists for every t in the domain of $\gamma(t)$, so that every $\mathcal{M}_u \bigcap \gamma(t)$ is either a straight, connected segment or the empty set,*
- *The curve $\gamma(t) \subset \mathcal{M}$ contains all of the points on $\mathcal{M}$ at which the local feature size is zero,*

then $\mathcal{P} \cong \mathcal{M}$.

Proof. As shown in Theorem 2, $\mathcal{P}$ coincides with the perfect reconstruction in the parts of the surface covered by smooth neighborhoods, so the task is to show that the portions of the surface that are in the vicinity of $\gamma(t)$ are also properly handled. To do that, it is necessary to show that

1. The backtracking subroutine in Algorithm 4 successfully terminates,
2. No "extra holes" exist in $\mathcal{P}$ relative to $\mathcal{M}$,
3. There are no "missing holes" in $\mathcal{P}$ relative to $\mathcal{M}$

when $\mathcal{S}$, $\mathcal{M}$, and $\gamma(t)$ are as described in the hypothesis. The first of these guarantees that $\mathcal{P}$ is a manifold and the second and third make sure that $\mathcal{P} \cong \mathcal{M}$.

Issue 1 can be addressed by noting that $\gamma(t)$ is completely covered by isometric cycles in RNG that are contractible on $\mathcal{M}$.

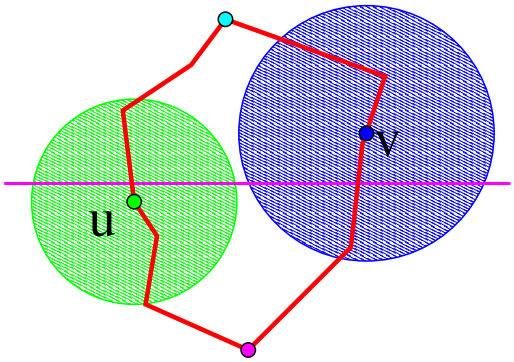

Figure 5: The horizontal magenta line is $\gamma(t)$. $\mathcal{M}_u$ is shown in green and $\mathcal{M}_v$ is shown in blue. The cycle that covers a portion of $\gamma(t)$ is shown in red.

As in Figure 5, we will imagine that the surface is positioned before us so that $\gamma(t)$ is horizontal. Let $u, v \in \mathcal{S}$ be two sample points such that i) $\mathcal{N}(u;k)$ has samples above and below $\gamma(t)$, ii) $\mathcal{N}(v;k)$ does too, iii) $\mathcal{M}_u \bigcap \mathcal{M}_v = \emptyset$, and iv) $v \in \mathcal{N}(u;2k)$. The two neighborhoods both cover portions of $\gamma(t)$ and are nearby, but do not overlap.

It is known that the Euclidean relative neighborhood graph of any set of points is a supergraph of the minimum spanning tree of the same points[19]. Along with that, it can be seen that Algorithm 2 computes the relative neighborhood graphs of $\mathcal{N}(u;k)$ and $\mathcal{N}(v;k)$ and includes them as subgraphs of RNG. Because of those two facts, and because the minimum spanning tree is connected, RNG has paths of edges which cross $\gamma(t)$ from above to below in both neighborhoods. $\mathcal{M}$ is flat above and below $\gamma(t)$, so the respective beginning and ending points of those paths connect above and below the sharp edge to form a cycle. All of the cycles formed in this way are of size less than $33k$ and will therefore be found by Algorithm 3.

The fact that the cycles are in the collection of isometric small isometric cycles, $\{\pentagon_i\}$, means that they are available for inclusion in the 2-basis, $\{\square_i\}$, by Algorithm 4. In fact, this construction shows that there are sufficient isometric cycles in RNG to provide a 2-basis for every $\mathcal{N}(u;k)$ that is compatible with the nearby smooth neighborhoods so that the *success condition* of the backtracking algorithm can be satisfied. All of that implies that the presence of $\gamma(t)$ is not sufficient to prevent Algorithm 4 from finding a closed collection of faces, so that subroutine terminates successfully.

Issue 2 is whether $\mathcal{P}$ has any "extra holes" or not. An extra hole exists if and only if there is some isometric cycle in the edge set of $\mathcal{P}$ that is not contractible on $\mathcal{P}$ but is on $\mathcal{M}$. (Here, we are using the same correspondence between cycles and loops that we discussed in Section 3.) Actually, a new hole implies two types of new nontrivial isometric cycles: those around the circumference of the hole and those that go through it. We will first consider a cycle of the latter type. $\hexagon_{\text{around}}$ is a cycle that is of minimal length among circumferential cycles.

If we assume for contradiction that $\mathcal{P}$ does in fact have an extra hole, then $\hexagon$ exists. We have previously established Theorem 2, which guarantees that the reconstruction must be correct in smooth areas. That means that it must be possible to cover $\hexagon$ with neighborhoods $\mathcal{N}(u;k)$ where $\mathcal{M}_u \bigcap \gamma(t) \neq \emptyset$, because $\gamma(t)$ is the exclusive cause of errors in this situation.

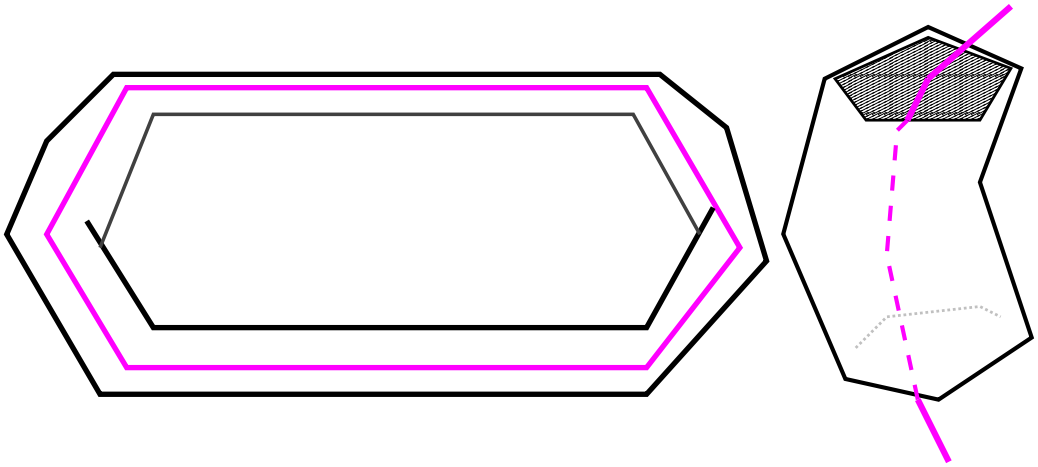

Figure 6: In both subfigures, $\gamma(t)$ is represented by a magenta line. The two subfigures show the two situations possible when an extra hole is created in $\mathcal{P}$ due to a sharp edge. In the first picture, the circumference of the new hole fits entirely in one sharp neighborhood. In the second picture, only a portion of the new hole fits into one neighborhood. One of these situations must obtain. In both cases, the neighborhood is homeomorphic to a donut instead of to a disk, in violation of Conditon II-i.

Since the surface is closed and $\mathcal{M}_v$ smooth implies $\text{Conv}_{\mathbb{R}^3}\{\mathcal{N}(v;k)\} \bigcap \mathcal{P} \cong D$ (due to the fact that smooth neighborhoods are already known to be correctly reconstructed), it must be the case that there is some neighborhood where $\mathcal{M}_u \bigcap \gamma(t) \neq \emptyset$ and $\text{Conv}_{\mathbb{R}^3}\{\mathcal{N}(u;k)\} \bigcap \mathcal{P}$ is either homeomorphic to a donut or a tube. However, that violates explicit check made in the *success condition* of the backtracking subroutine (Algorithm 4) which ensures that $\text{Conv}_{\mathbb{R}^3}\{\mathcal{N}(u;k)\} \bigcap \mathcal{P} \cong D$. Please see Figure 6 for illustrations.

Issue 3 is whether or not the output surface $\mathcal{P}$ has a "missing hole". There is a missing hole if and only if there is an isometric cycle in the edges of $\mathcal{P}$ that goes around the circumference of the missing hole, so that it is nontrivial on $\mathcal{M}$ but contractible on $\mathcal{P}$. We will give the name $\hexagon$ to some cycle that is of minimal length among those answering that description.

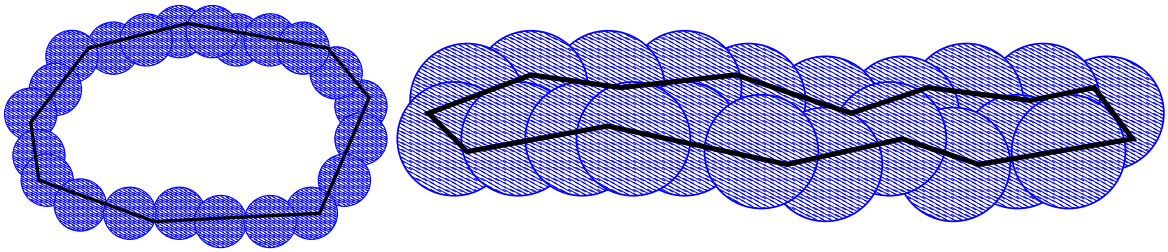

Figure 7: In both subfigures, the black loop is the circumference of a hole that is present in $\mathcal{M}$ but missing in $\mathcal{P}$. The first subfigure shows a wide loop, where the chords necessary to close the loop cannot reach across it because the neighborhoods are too small. The second subfigure shows a skinny hole in which some neighborhoods are not homeomorphic to a disk, in violation of Condition II-i. One of these two situations must obtain.

By Condition II-iv, $|\hexagon| > 33k$. The closedness of the output surface $\mathcal{P}$ implies the existence long chords across $\hexagon$ in RNG because there are either no samples in the hole, or if there are samples they belong to another part of the surface. However, those long chords are impossible because edges in RNG only exist within k-neighborhoods and $\mathcal{M}_u \cong D$. Please see Figure 7 for pictures. □

So we have shown that the presence of a sharp edge does not interfere with the quality of the reconstruction. Now we will deal with the issue of isolated sharp points and corners. This issue can be addressed with almost no additional analysis.

Theorem 4. *Let $\mathcal{M}$ be as described in Theorem 3 with the addition a single isolated sharp point or corner. If $\mathcal{S}$ is as described in Section 3 (with ε is sufficiently small), then $\mathcal{P} \cong \mathcal{M}$.*

Proof. Every neighborhood except those that host the isolated sharp point or the corner are correctly reconstructed by Theorem 2. Let u be the sample closest to the isolated sharp point or corner. By Theorem 2, a valid collection of face cycles $\{\square_i\}$ is present in RNG, except possibly in the area touched by $\mathcal{N}(u; k)$. Regardless of whether that area is correctly reconstructed or not, it is surrounded by an isometric loop of circumference less than $33k$, which in the worst case can be made part of the basis $\{\square_i\}$. □

Theorem 5. *Let $\mathcal{S}$ be a sample set that follows the conditions given in Section 3 (with ε sufficiently small). If $\mathcal{M}$ is a surface that is smooth except that*

- *There are a constant number of sharp edges $\gamma_i(t)$ which are isolated from one another except at a constant number of corner points which are isolated from one another,*
- *There are a constant number of isolated sharp points,*

then $\mathcal{P} \cong \mathcal{M}$.

Proof. The smooth parts of the surface are guaranteed to be correctly reconstructed by Theorem 2. The portions of the sharp edges that are isolated from other sharp features are guaranteed to be correctly reconstructed by Theorem 3. The isolated sharp points and corners are guaranteed to be correctly reconstructed by Theorem 4. □

Theorem 5 completes the work except for one minor defect. Based on what we have stated, the reconstruction $\mathcal{P}$ is not guaranteed to use all of the sample points $u \in \mathcal{M}$, even though $\mathcal{P} \cong \mathcal{M}$. This can be easily remedied by attaching any stray points to $\mathcal{P}$ with cones (this has to be done in such a way that the interiors of all triangles are kept disjoint).

6 Experiments

We have created a prototype implementation of the algorithm, and this section documents experiments done on various datasets with the prototype. The implementation is substantially similar to Algorithm 1 and its subroutines, except that various heuristic simplifications were used in various places. For example, instead of explicitly computing a 2-basis of small isometric cycles, we compute the collection of small isometric cycles and remove those that are found to go through or be inside of the consensus surface. Bringing the heuristic portions of the prototype into closer conformance with Section 4 would only improve the quality of the output (though, perhaps at the expense of greater complexity).

No effort has been made to clean up or remove defects from the output of the prototype. The results shown in Table 2 and Figures 8, 9, and 10 are raw output from the algorithm, including some flaws which could have been corrected with simple heuristics. This style of presentation may be in contrast to other work in this area, but we have chosen it in the interest of simplicity and transparency. By decoupling the separate issues of surface reconstruction and post-processing, we hope to make it easier for the reader to understand and judge the quality of the work. To that end, we have also made the prototype code freely available.[3]

We used datasets from three different sources to evaluate the prototype. The first five datasets shown in Table 2 are sample sets on surfaces with sharp features, including sharp non-convex features, that were custom-generated for this project. All of the five datasets were designed to meet the sample conditions given in Section 3, and are meant to establish a baseline for the algorithm's performance. They were all reconstructed perfectly, as expected.[4]

The next five datasets are from various universities and were found on the AIM@SHAPE Project Shape Repository[5]. The *fandisk* dataset is courtesy of MPII and the rest are from INRIA. These datasets are not known to conform to the sample conditions, and therefore we do not expect all of them to be correctly reconstructed. In spite of this, *fandisk* was perfectly reconstructed.

[3] `https://github.com/jamesmcclain/sr33`

[4] Some of the cube-shaped surfaces appear in the renderings to have flaws at some of the corners, but that appearance is due to the fact that there are no samples at those corners. These datasets are distributed with the prototype program.

[5] `http://shapes.aim-at-shape.net/`

Dataset Name	k	Open Edges	Overloaded Edges	Total Edges	Rendering
L	33	316 (all on the boundary)	0	15041	Figure 8a
Cone	33	0	0	9252	Figure 8b
Conebox	33	0	0	20592	Figure 8c
Cube	33	0	0	21258	Figure 8d
Cutcube	33	0	0	21261	Figure 8e
Fandisk	33	0	0	35883	Figure 9a
Foot	107	69	14	76488	Figure 9b
Niccolò	107	1069 (most on the boundary)	27	72481	Figure 9c
Caesar	107	1043 (most on the boundary)	54	72790	Figure 9d
Eros	107	1150 (most on the boundary)	81	72841	Figure 9e
Dragon	42	638	13	67298	Figure 10a
Stanford Bunny	42	1081 (many on the boundary)	40	106572	Figure 10b
Happy Buddha	42	1991	7	92183	Figure 10c

Table 2: Summary of the experiments.

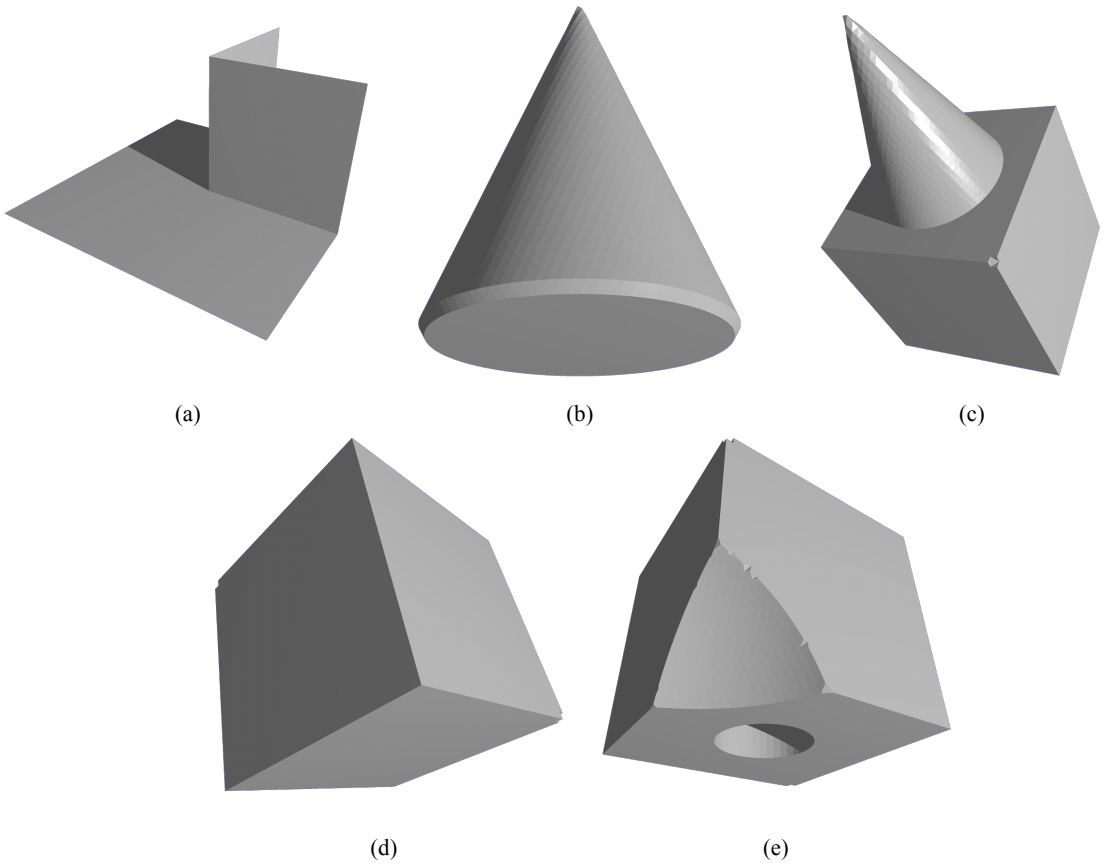

Figure 8: Surfaces with sharp features and square-lattice sample patterns.

The other four reconstructions do exhibit small flaws. The *Foot* dataset, for example, has 69 open edges (edges that meet only one triangle) and 14 overloaded edges (edges that meet more than two triangles). These 83 flawed edges constitute about 0.11% of the 76,488 edges in the reconstruction. The *Eros* dataset was the most difficult among this group, with 1.69% potentially flawed edges. It should be noted that *Niccolò*, *Caesar*, and *Eros* all have boundaries, so a significant number of the open edges are due to that.

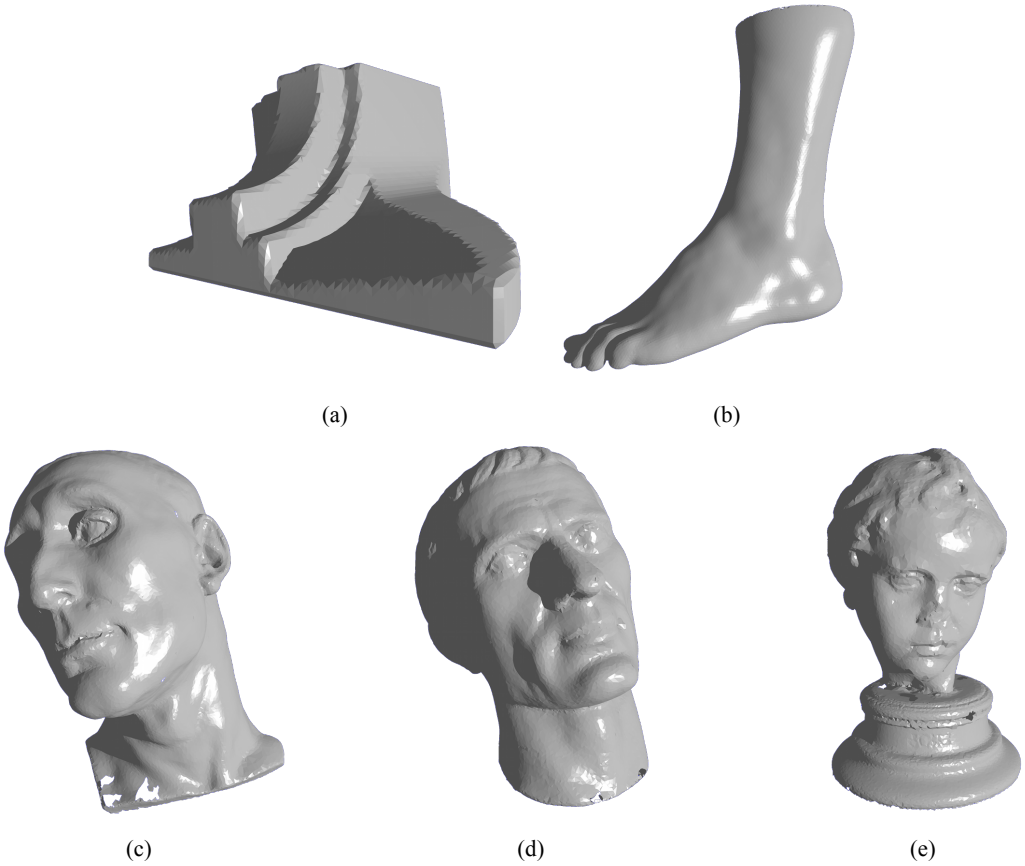

Figure 9: The *fandisk* dataset is from the MPII, the rest came from INRIA.

The last three datasets that we tested came from The Stanford 3D Scanning Repository[6]. The *Dragon* dataset had 0.97% flawed edges. *Stanford Bunny* had about 1% potentially-flawed edges (this model has a boundary, two holes on the bottom). The *Happy Buddha* statue was the most challenging and had about 2.1% flawed edges.

7 Conclusions

In this article, we presented an algorithm (Section 4) and associated sample conditions (Section 3) that guarantee the homeomorphic reconstruction of surfaces, even those with certain sharp features. The started point for the algorithm is the observation made in the pattern-recognition community years ago that the Urquhart graph tends to informally "resemble" the data that they are computed over. When we combined that observation with the similarity between the Urquhart graph and intrinsic Delaunay triangulation, we had a good starting point.

From there, we made use of the similarity between the intrinsic Urquhart graph the intrinsic relative neighborhood graph. These two objects are in fact identical when the set of sample points is uniform enough (Theorem 1). We then made yet another step and noticed the similarity between the intrinsic and Euclidean relative neighborhood graphs when the

[6] `http://graphics.stanford.edu/data/3Dscanrep/`

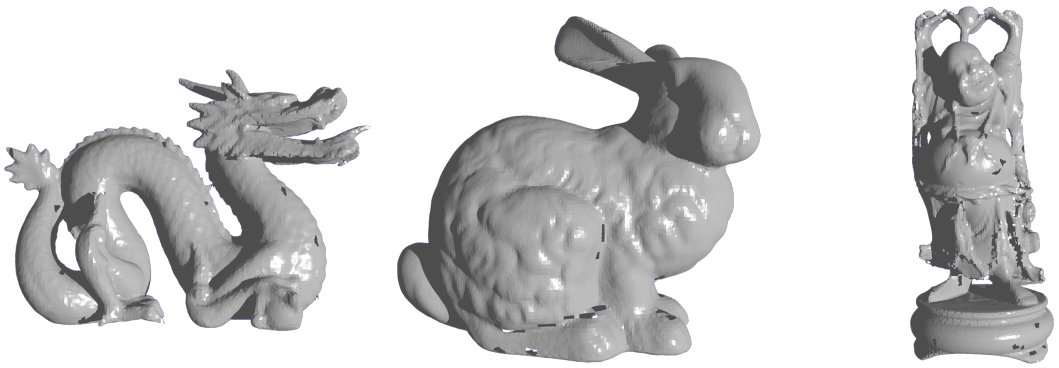

(a) (b) (c)

Figure 10: Datasets from Stanford.

sample set is both uniform and dense (Theorems 2, 3, 4, and 5). That last fact provides the basis for a workable algorithm, because the Euclidean relative neighborhood graph is accessible to us, whereas any intrinsic graph is not.

We demonstrated analytically in Section 5 that the algorithm works and gave experimental evidence in Section 6. When the sample set is known to obey the sample conditions, our prototype implementation of the algorithm produces perfect results, as expected (see Figure 8). Even when the sample set does not necessary conform to the sample conditions, the algorithm is still able to do well (see Figures 9, 10).

Acknowledgments

We would like to thank Neeraj Kashyap for providing valuable thoughts and perspective. We would also like to acknowledge the AIM@SHAPE Project Shape Repository, the Max Planck Institute for Informatics (MPII), the Institut National de Recherche en Informatique et en Automatique (INRIA), and the Stanford Computer Graphics Laboratory for their provision of some of the datasets used in this project.

References

[1] N. Amenta and M. Bern. Surface reconstruction by voronoi filtering. In *Proceedings of the fourteenth annual symposium on Computational geometry*, SCG '98, pages 39–48, New York, NY, USA, 1998. ACM.

[2] N. Amenta, M. Bern, and M. Kamvysselis. A new voronoi-based surface reconstruction algorithm. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '98, pages 415–421, New York, NY, USA, 1998. ACM.

[3] N. Amenta, S. Choi, T. K. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. In *Proceedings of the sixteenth annual symposium on Computational geometry*, SCG '00, pages 213–222, New York, NY, USA, 2000. ACM.

[4] D. V. Andrade and L. H. de Figueiredo. Good approximations for the relative neighbourhood graph. In *CCCG*, pages 25–28, 2001.

[5] J.-D. Boissonnat and S. Oudot. Provably good sampling and meshing of lipschitz surfaces. In *Proceedings of the twenty-second annual symposium on Computational geometry*, SCG '06, pages 337–346, New York, NY, USA, 2006. ACM.

[6] F. Chazal, D. Cohen-Steiner, and A. Lieutier. A sampling theory for compact sets in euclidean space. In *Proceedings of the twenty-second annual symposium on Computational geometry*, SCG '06, pages 319–326, New York, NY, USA, 2006. ACM.

[7] F. Chazal and A. Lieutier. Weak feature size and persistent homology: computing homology of solids in $\mathbb{R}^n$ from noisy data samples. In *Proceedings of the twenty-first annual symposium on Computational geometry*, SCG '05, pages 255–262, New York, NY, USA, 2005. ACM.

[8] S.-W. Cheng, T. K. Dey, and E. A. Ramos. Delaunay refinement for piecewise smooth complexes. *Discrete Comput. Geom.*, 43:121–166, December 2009.

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

[10] M. Deza and M. Laurent. *Geometry of cuts and metrics*, volume 15. Springer, 1997.

[11] H. Edelsbrunner and N. R. Shah. Triangulating topological spaces. In *Proceedings of the tenth annual symposium on Computational geometry*, SCG '94, pages 285–292, New York, NY, USA, 1994. ACM.

[12] C. Gotsman, K. Kaligosi, K. Mehlhorn, D. Michail, and E. Pyrga. Cycle bases of graphs and sampled manifolds. *Computer Aided Geometric Design*, 24(8):464–480, 2007.

[13] M. Gross and H. Pfister. *Point-Based Graphics (The Morgan Kaufmann Series in Computer Graphics)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2007.

[14] A. Hatcher. *Algebraic topology*. Cambridge University Press, 2002.

[15] S. Mac Lane. *A combinatorial condition for planar graphs*. Seminarium Matemat., 1936.

[16] R. Mencl. A graph–based approach to surface reconstruction. In *Computer Graphics Forum*, volume 14, pages 445–456. Wiley Online Library, 1995.

[17] R. Mencl and H. Müller. Interpolation and approximation of surfaces from three-dimensional scattered data points. In *Scientific Visualization Conference, 1997*, pages 223–223. IEEE, 1997.

[18] R. Mencl and H. Müller. Graph-based surface reconstruction using structures in scattered point sets. In *Computer Graphics International, 1998. Proceedings*, pages 298–311. IEEE, 1998.

[19] K. J. Supowit. The relative neighborhood graph, with an application to minimum spanning trees. *Journal of the ACM (JACM)*, 30(3):428–448, 1983.

[20] J. B. Tenenbaum, V. De Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.

[21] G. Tewari, C. Gotsman, and S. J. Gortler. Meshing genus-1 point clouds using discrete one-forms. *Computers & Graphics*, 30(6):917–926, 2006.

[22] G. Toussaint. Comment: Algorithms for computing relative neighbourhood graph. *Electronics Letters*, 16(22):860–860, 1980.

[23] G. T. Toussaint. The relative neighbourhood graph of a finite planar set. *Pattern recognition*, 12(4):261–268, 1980.

[24] R. Urquhart. Algorithms for computation of relative neighbourhood graph. *Electronics Letters*, 16(14):556–557, 1980.

[25] R. Urquhart. Graph theoretical clustering based on limited neighbourhood sets. *Pattern recognition*, 15(3):173–187, 1982.