A survey of Delaunay structures for surface representation

Article · February 2009

CITATIONS
8

READS
523

4 authors, including:

Ramsay Dyer
Tahsis Centre for Applied Geometry
35 PUBLICATIONS **665** CITATIONS

SEE PROFILE

Hao Zhang
Hebei University of Technology
39 PUBLICATIONS **2,278** CITATIONS

SEE PROFILE

A survey of Delaunay structures for surface representation

Ramsay Dyer , Hao Zhang, and Torsten Möller*
GrUVi Lab, School of Computing Science, Simon Fraser University, Canada

January 16, 2009

Abstract

The Delaunay triangulation characterizes a natural neighbour relation amongst points distributed in a Euclidean space. In this survey we examine extensions of the Delaunay paradigm that have been used to define triangle meshes for representing smooth surfaces embedded in three dimensional Euclidean space.

Progress in this area has stemmed primarily from work done in surface reconstruction and surface meshing. However, our focus is not on the surface reconstruction or meshing algorithms themselves, but rather on the structures that they aim to produce. In particular we concentrate on three distinct Delaunay structures which differ according to the metric involved in their definition: the metric of the ambient Euclidean space; the intrinsic metric of the original surface; or the intrinsic metric of the mesh itself. We study the similarities and distinctions between these objects; their strengths and weaknesses both as theoretical tools and as practical data structures.

The topic of this survey lies within the realm of geometry processing, a field of study generally associated with computer graphics. However, it is hoped that this survey will be of interest not just to those who study computer graphics, but to anybody whose research touches on a need to represent non-Euclidean geometry with Euclidean simplices.

*Email: rhdyer, haoz, torsten@cs.sfu.ca

Contents

1 Introduction

The goal of this report is to outline the main work that has been done in developing and understanding the properties of Delaunay structures as representations of the surfaces of objects in $\mathbb{R}^3$.

In Section 2 we review Voronoi diagrams and Delaunay triangulations as they pertain to Euclidean geometry, outlining their essential properties and applications.

Then we move on in Section 3 to the application of these concepts to surfaces in $\mathbb{R}^3$. There are several ways to extend the ideas that are so well developed in the Euclidean domain. Our focus is on three specific Delaunay structures and their differences and similarities. We then examine how well these structures perform as discrete approximations to an initial surface (Section 4).

Most of the theory that has been developed in this regard has arisen in the context of surface meshing and reconstruction. In Section 5 we review the literature in this area with particular focus on the use and emergence of Delaunay structures.

We conclude with a discussion of the work that needs to be done in order to reap the potential benefits of the newest of these three Delaunay structures.

2 Background

Voronoi diagrams and Delaunay triangulations are among the most fundamental structures in the computational geometry toolbox. The Voronoi diagram gets its name from the Russian mathematician Georges Voronoi [Vor07, Vor08], although it has appeared under many different names as it has been independently discovered in many different fields. In fact there is evidence that René Descartes was using Voronoi diagrams early in the seventeenth century [Ede01]. The Delaunay triangulation gets its name from another Russian mathematician. Boris Delaunay [Del34] produced the first systematic modern development of the theory of the structure that is dual to the Voronoi diagram.

These structures are discussed in any computational geometry textbook, see [dBvKOS98] or [Ede01] for example. Comprehensive surveys of the uses, properties and variations of Voronoi diagrams in particular can be found in [AK00] and [OBSC00].

In this section we review the fundamental properties and uses of the Voronoi diagram and the Delaunay triangulation as well as the relationship between them. For the most part, we will confine our attention to Euclidean domains of two or three dimensions, and we focus particularly on properties that will have relevance when we discuss these structures in the context of surface representation.

2.1 Voronoi diagrams

The Voronoi diagram is easy to describe and, via a duality relationship, it facilitates the description of the Delaunay triangulation. Given a set P of n points in $\mathbb{R}^d$, the Voronoi diagram partitions $\mathbb{R}^d$ into n cells: one cell is associated with each point in P. For $p \in P$, we denote the associated Voronoi cell by $V(p)$. The extent of $V(p)$ is simply the entire

region of $\mathbb{R}^d$ whose distance to P is realized by the distance to p. That is, the set of points that is at least as close to p as it is to any other $q \in P$. Formally, we have[1]

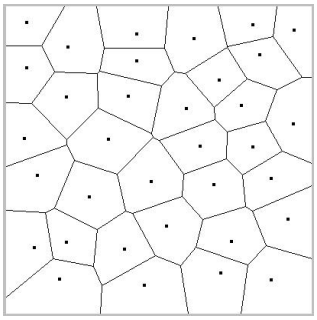

A Voronoi diagram

Definition 2.1 (Voronoi diagram) The *Voronoi cell* of $p \in P$ is defined by

$$V(p) = \{x \in \mathbb{R}^d \mid d_{\mathbb{R}^3}(p, x) \leq d_{\mathbb{R}^3}(q, x), \forall q \in P\}. \tag{1}$$

The set of Voronoi cells forms a covering of S called the *Voronoi diagram* of P.

The Voronoi diagram gives a very natural definition of the neighbours of a point $p \in P$: q and p are *Voronoi neighbours* if their Voronoi cells have a nonempty intersection. This is the relationship that will be exploited to define the Delaunay triangulation in Section 2.2.

The Voronoi cells are convex polygons in $\mathbb{R}^2$, and in higher dimensions they are convex polytopes. Indeed, $V(p)$ can be constructed as the intersection of the $n-1$ half spaces each of which contains p and is bounded by the orthogonal bisector of $[p, q]$ for some $q \in P$.

Note that we have defined the Voronoi cells to be closed; neighbouring Voronoi cells have a nonempty intersection. The intersection of $d+1$ or more Voronoi cells is either empty or a single point, called a *Voronoi vertex*. A Voronoi vertex v is equidistant from the elements of P whose Voronoi cells define it. Thus if $v = \cap_{i=0}^{d} V(p_i)$, then the p_i all lie on a common hypersphere centred at v. For a random set of points $P \subset \mathbb{R}^d$, the chances of more than $d+1$ points lying on a common hypersphere is vanishingly small [Ede01]. The set P is said to be in *general position* if the intersection of more than $d+1$ Voronoi cells is always empty. It is common practice to assume that P is in general position, and in fact this is not a big constraint to impose. If P is not in general position, an arbitrarily small perturbation of the positions of its violating elements is sufficient to bring it into general position.

In the computational geometry literature, the elements of P are often referred to as *sites*. For our purposes, we prefer to consider the elements of P to be *samples*. In the context of geometry processing, the elements of P will be the discrete sample points that represent the geometry (i.e., positional information) of the underlying continuous surface. More generally, a sample set may provide a discrete representation of any signal (or function) defined on the domain.

Suppose then that to each $p_i \in P$ there is associated a value $f_i \in \mathbb{R}$ and we interpret the f_i as the value of some unknown continuous function $f : \mathbb{R}^d \to \mathbb{R}$. The Voronoi diagram naturally provides an interpolation scheme that allows us to define a function $\tilde{f} : \mathbb{R}^d \to \mathbb{R}$ that coincides with f on P. For $x \in V(p_i)$, we simply define $\tilde{f}(x)$ to be f_i. Although crude and primitive, this interpolation scheme, called *nearest neighbour interpolation*, is used in some applications and it provides the basis for more sophisticated interpolation and approximation schemes.

A more sophisticated interpolation scheme based on the Voronoi diagram was introduced by Sibson [Sib81]. In this scheme, refered to as *natural neighbour interpolation*, the value

[1] Image from http://www.cs.wustl.edu/~pless/546/lectures/l16.html

of $\tilde{f}(x)$ is determined by considering the Voronoi cell $V(x)$ of $P \cup \{x\}$. The value of f at x is then given as a weighted sum of the values of f at the sample points p, with the weights given by the portion of the original Voronoi cell of p that is occupied by the new Voronoi cell $V(x)$. Thus

$$\tilde{f}(x) = \sum_{p \in P} \frac{|V(x; \{x\} \cup P) \cap V(p; P)|}{|V(p; P)|} f(p),$$

where $V(q; Q)$ denotes the Voronoi cell of q in the Voronoi diagram of Q, and $|V(q; Q)|$ represents its volume. Sibson demonstrated that this interpolation scheme is smooth everywhere except at the sample points.

Nearest neighbour interpolation is simple, but it is not even continuous. Natural neighbour interpolation is smooth almost everywhere, but it lacks simplicity. A happy medium is obtained with piecewise linear interpolation between the sample points. This involves constructing a triangulation of the sample points and indeed the Voronoi diagram imposes a natural triangulation on the sample set. This triangulation, the Delaunay triangulation, embodies the paradigm around which this report is based.

2.2 Delaunay triangulations

2.2.1 Triangulations

A common approach to discretizing a domain is to use a triangulation. A set of samples P is created and these samples are connected by linear elements (line segments, triangles, tetrahedra, etc.) to form a covering of the domain that provides a convenient framework for interpolation and numerical computations. A simplicial complex is a well established structure that allows us to define a triangulation in a unified way for all dimensions. The definition we present here is standard to most textbooks on algebraic topology, see Munkres [Mun84] for example.

The elements that are the building blocks of a triangulation are called simplices. A set of points $X = \{p_0, p_1, \ldots, p_m\} \subset \mathbb{R}^d$ is *affinely independent* if the vectors $\{p_i - p_0\}_{i \in [1,\ldots,m]}$ are linearly independent. Given such an affinely independent set, the *m-simplex*, σ, spanned by $p_0, \ldots, p_m$ is the set of points $x \in \mathbb{R}^d$ such that

$$x = \sum_{i=0}^{m} t_i p_i \qquad \text{with } \sum_{i=0}^{m} t_i = 1 \text{ and } t_i \geq 0 \; \forall i.$$

Observe that an m-simplex is an m-dimensional object. Any subset of the $Y \subset X$, with $h = |Y| - 1$, defines an h-simplex $\tau \subset \sigma$ that is called an *h-face* of σ. The 0-faces of σ are the p_i themselves and they are called *vertices*. The 1-faces are line segments connecting the vertices and they are refered to as *edges*. Simplices of dimension 2 and 3 are triangles and tetrahedra respectively.

An object σ may be similarly defined even if X is not affinely independent. In this case σ will not be a truly m-dimensional object, and we say that it is a *degenerate simplex*.

A *simplicial complex* $K \subset \mathbb{R}^d$ is a collection of simplices such that

1. If $\sigma \subset K$, then every face $\tau \subset \sigma$ also belongs to K

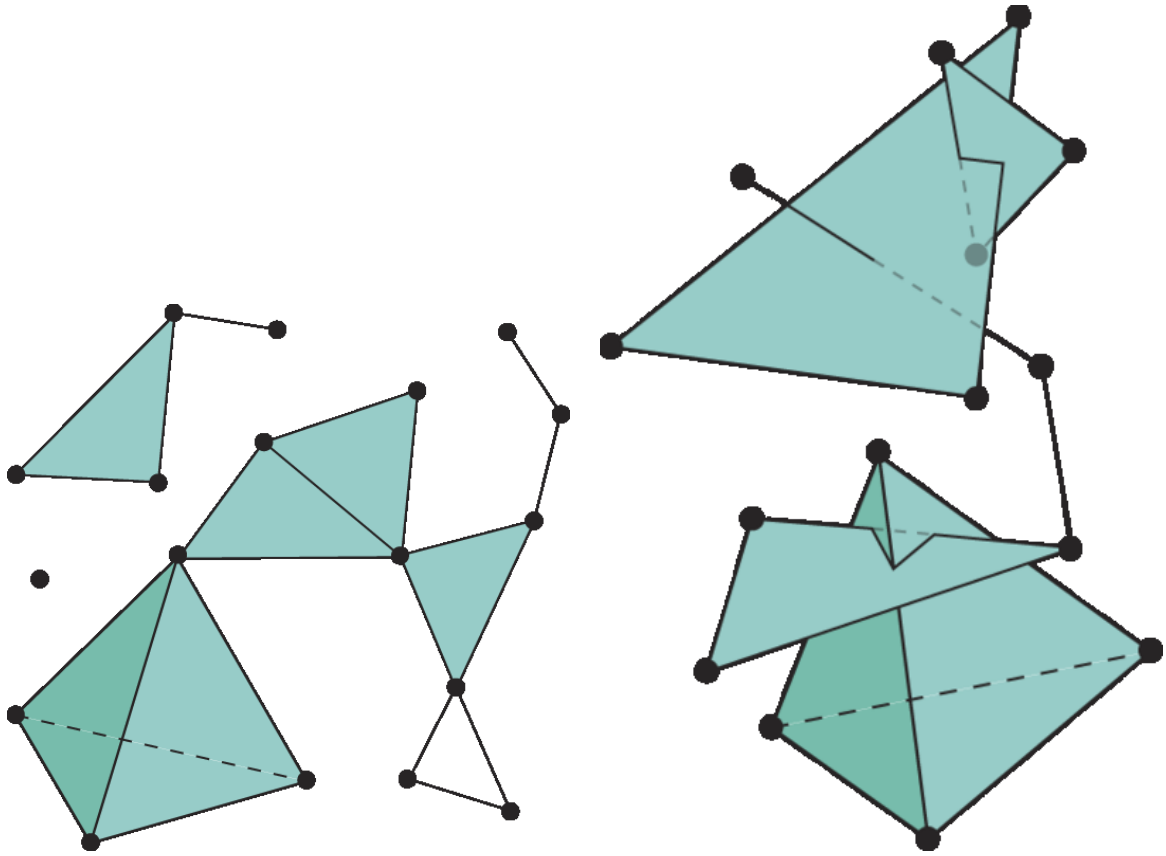

(a) (b)

Figure 1: (a) A simplicial complex. (b) Not a simplicial complex

2. The intersection of any two simplices is a face of each of them.

For the last point, note that the empty set is a face of all simplices. Figure 1 gives examples of a simplical complex and a collection of simplices that is not a simplicial complex[2]

A *triangulation* of a set of points $P \subset \mathbb{R}^d$ is a simplicial complex K whose vertices are exactly P and such that the union of all the simplices in K forms the convex hull of P.

To define a triangulation of a domain $D \subset \mathbb{R}^d$ we need to decompose D into a collection of simplices. However, such a decomposition can only be possible if D itself is a polytope, i.e., it has a piecewise linear boundary. If this is not the case, then D must be approximated by some polytope $\tilde{D}$ which can then be triangulated: a triangulation of $\tilde{D}$ is a simplicial complex K whose vertices include the vertices of $\tilde{D}$ and such that the union of simplices of K coincides with $\tilde{D}$.

2.2.2 Introducing Delaunay triangulations

For P in general position, the Delaunay triangulation of $P \in \mathbb{R}^d$ is the dual of the Voronoi diagram of $P \in \mathbb{R}^d$. In the planar setting, the duality relationship is as follows: To each

[2]images from Wikipedia: http://en.wikipedia.org/wiki/Simplicial_complex

Voronoi vertex c we associate a *Delaunay triangle*, t whose vertices are the three samples which define c. An edge $e = [p, q]$ of t is dual to the Voronoi edge $V(p) \cap V(q)$. The vertices of the Delaunay triangulation are the sample points, and they are dual to the corresponding Voronoi cells in the Voronoi diagram.

Note that if c is a Voronoi vertex defined by samples $p, q, s \in P$, then those samples lie on a circle, C, centred at c and no samples lie in the interior of the disk bounded by C. The circle C is the circumcircle of the Delaunay triangle $t = pqs$. If P is in general position, then each edge of t can be contained in a disk that has the endpoints of the edge as the only two samples on its boundary and that has no samples in its interior. This empty circumcircle property is the standard way to formally define Delaunay triangulations.

Definition 2.2 Let $P \subset \mathbb{R}^d$ be a finite set of points. A *Delaunay triangulation* of P is a triangulation, K, such that each $\sigma \in K$ has a d-ball that has the vertices of σ on its boundary and no elements of P contained in its interior.

If P is in general position, the Delaunay triangulation is known to be unique. If P is not in general position, the Delaunay triangulation is not uniquely defined. For example, in the plane if four points lie on the boundary of a circle, there are two ways to triangulate their convex hull: both fit the definition of a Delaunay triangulation. Points that are not in general position present a technical annoyance, but as discussed in Section 2.1, demanding that P be in general position is not a significant constraint.

We will assume that P is in general position.

2.2.3 Properties of Delaunay triangulations

In the planar setting, the Delaunay triangulation owes its popularity in part to its ease of definition and construction, but also to the many properties it enjoys. In the plane, the Delaunay triangulation can be constructed with $\mathcal{O}(n \log n)$ operations; in the general case of d dimensions $\mathcal{O}(n \log n + n^{\lceil d/2 \rceil})$ operations are required [dBvKOS98].

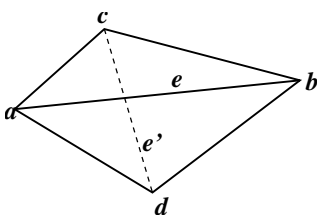

An early, yet still practical algorithm for producing a planar Delaunay triangulation was introduced by Lawson [Law77]. This algorithm takes an arbitrary triangulation of P as input, and produces a Delaunay triangulation by *edge flipping*. An edge flip replaces an edge $e = [a, b]$ that is adjacent to triangles $[a, b, c]$ and $[b, a, d]$ with the edge $e' = [c, d]$ that is the other diagonal of the quadrilateral $[a, d, b, c]$. We refer to e' as the *opposing edge* to e. The quadrilateral $[a, d, b, c]$ must be convex for the triangulation to remain valid.

We say that $e' = [c, d]$ is *locally Delaunay* if a is not contained in the circumcircle of $[b, c, d]$, or equivalently, if b is not contained in the circumcircle of $[c, a, d]$. There is a convenient characterization of a locally Delaunay edge e': the sum of the angles at the opposing vertices, a, and b, does not exceed π [BS07]. With this observation it is easy to show that if edge e is not locally Delaunay (NLD), then $[a, d, b, c]$ is a convex quadrilateral and the opposing edge, e', will be locally Delaunay.

Lawson's algorithm consists of repeatedly selecting an NLD edge and flipping it. Such a flip is called a *Delaunay flip*. If a nondecreasing sequence $n_{\mathcal{T}}$ is constructed consisting of

all the face angles in the triangulation $\mathcal{T}$, then it is shown that a Delaunay flip produces a triangulation $\mathcal{T}'$ in which the associated sequence $n_{\mathcal{T}'}$ is lexicographically smaller than $n_{\mathcal{T}}$ [Ede01]. Thus the edge flipping algorithm terminates with the triangulation that has the lexicographically minimal sequence of angles.

Sibson [Sib78] showed that the resulting triangulation is in fact the unique Delaunay triangulation for samples in general position. Thus if *all* the edges in a triangulation are locally Delaunay, then the triangulation will be a Delaunay triangulation. Note, however, that an edge that is locally Delaunay in a general triangulation will not necessarily belong to the Delaunay triangulation. The property of being a locally Delaunay edge is not as strong as the property of being a Delaunay edge. Indeed the number of NLD edges in a triangulation is not necessarily a good indicater of how many edge flips are required to achieve a Delaunay triangulation. It is possible to construct a configuration in which there is only one NLD edge, but none of the triangles are Delaunay (see figure) [GR04].

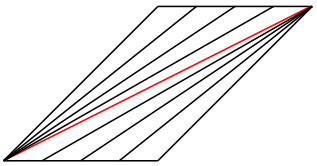

The red edge is the only NLD edge, but none of the triangles are Delaunay.

It turns out that there are several properties, in addition to face angles, that monotonically decrease (or increase) with Delaunay flips. Any such property is optimized (minimized or maximized) by the planar Delaunay triangulation and can be used to demonstrate termination of the edge swapping algorithm. If the property is something that is defined on each triangle, then the optimization can be expressed in terms of a lexicographic ordering, as described for the triangle face angles.

Consider again the case where each sample x_i has an associated value f_i that is considered to be the value of some continuous function defined over the domain. Let $f_{\mathcal{T}}$ be the piecewise linear function that interpolates the data over triangulation $\mathcal{T}$. Rippa [Rip90] showed that, regardless of the data $\{f_i\}$, the planar Delaunay triangulation minimizes the *Dirichlet energy* of $f_{\mathcal{T}}$:

$$E(f_{\mathcal{T}}) = \int_D \|\nabla f_{\mathcal{T}}\|^2 \, da = \sum_{t \in \mathcal{T}} \|\nabla f_{\mathcal{T}}|_t\|^2 \, a_t,$$

where a_t is the area of triangle t. In other words, the graph of $f_{\mathcal{T}}$ has "minimal roughness" when $\mathcal{T}$ is a Delaunay triangulation.

Another property that is minimized by the planar Delaunay triangulation is the *harmonic index* [Mus97]. The harmonic index of a triangle, t, with sides of length a, b and c, is given by $h(t) = \frac{a^2+b^2+c^2}{a_t}$, and the harmonic index of a triangulation, T, is given by the sum of the harmonic indices of all the triangles. The harmonic index of t can also be expressed as four times the sum of the cotangents of the angles of t ([BS06]). Note that the harmonic index of a triangle t is minimized if t is equilateral.

The planar Delaunay triangulation also minimizes the maximum circumradius of the triangles [DS89]. As will be discussed in Section 4, this is a property that is interesting in the context of surface approximation.

All these properties mentioned have been demonstrated for the planar Delaunay triangulation by exploiting the edge flipping algorithm. However the edge flipping algorithm does not easily extend to higher dimensions, and those extensions that exist are known to not work if started from an arbitrary initial triangulation [ES92]. Thus for most of these

properties, it is not known whether or not they extend to a property that is optimized by the higher dimensional Delaunay triangulation.

In some cases, a property that is optimized by the planar Delaunay triangulation is known to be not optimized by higher dimensional Delaunay triangulations. An example is the mean inradius of the triangles. This is maximized by the Delaunay triangulation in 2D, but not the higher dimensions [Lam94].

By using a completely different technique of proof, Rajan [Raj94, Raj91] showed that for all dimensions the Delaunay triangulation minimizes the radius of the maximum smallest enclosing circle. Note that, for obtuse triangles in 2D, for example, this radius is smaller than the circumradius. Rajan's proof does not exploit the flip algorithm, and in fact the Delaunay triangulation does *not* minimize the smallest enclosing circles in a lexicographic sense.

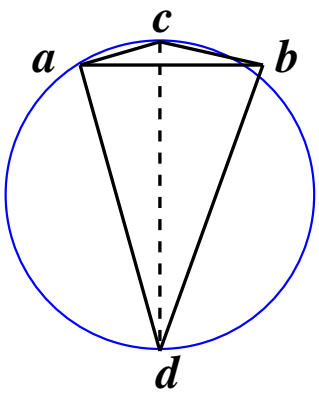

The NLD edge $[a, b]$ is shorter than the opposing Delaunay edge $[c, d]$.

It is worth mentioning a couple of properties that are *not* generally possessed by the Delaunay triangulation. Contrary to the erroneous assertion in [Lei99][Lemma 24], the Delaunay triangulation does not minimize the total edge length. It is not hard to construct a Delaunay edge flip that increases the length of an edge. The problem of constructing a triangulation that minimizes the total edge length is called the *minimal weight triangulation problem*, and it has recently been shown to be NP-hard [MR06].

The Delaunay triangulation also does not minimize the maximum angle. Although triangulations which possess this property can be desirable for scientific computing applications, computing such triangulations appears more difficult than computing Delaunay triangulations [Ede01][p.12].

2.2.4 Meshing and Delaunay refinement

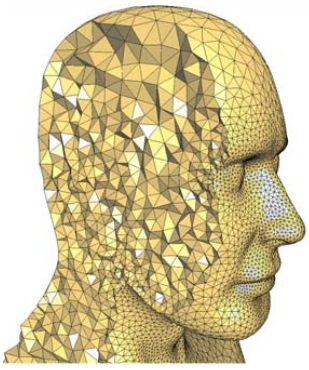
A cut-away view of a meshed 3D Euclidean domain

The process of producing a triangle mesh to represent a given domain is called *meshing*. In subsequent sections we will focus on the problem of meshing the surface of a three dimensional object, however the problem of meshing a Euclidean domain is a huge topic in its own right, and remains an area of active research. Many of the tools and insights discovered in this latter context are beginning to find their way into the field of surface meshing.

A *Euclidean domain* D is a subset of Euclidean space (e.g., $D \subset \mathbb{R}^2$ or $D \subset \mathbb{R}^3$) that is equal to the topological closure of its interior. In particular, the topological interior of D is non-empty. So for example, the solid head of Max Plank[3] is a Euclidean domain, D. The surface of the head is the boundary of D and it is not a Euclidean domain.

The problem of meshing Euclidean domains arises in scientific computing; a PDE is to be solved numerically over the given domain. In this case it is important that the elements

[3]Image from Aliez et al. "Variational Tetrahedral Meshing" SIGGRAPH 2005

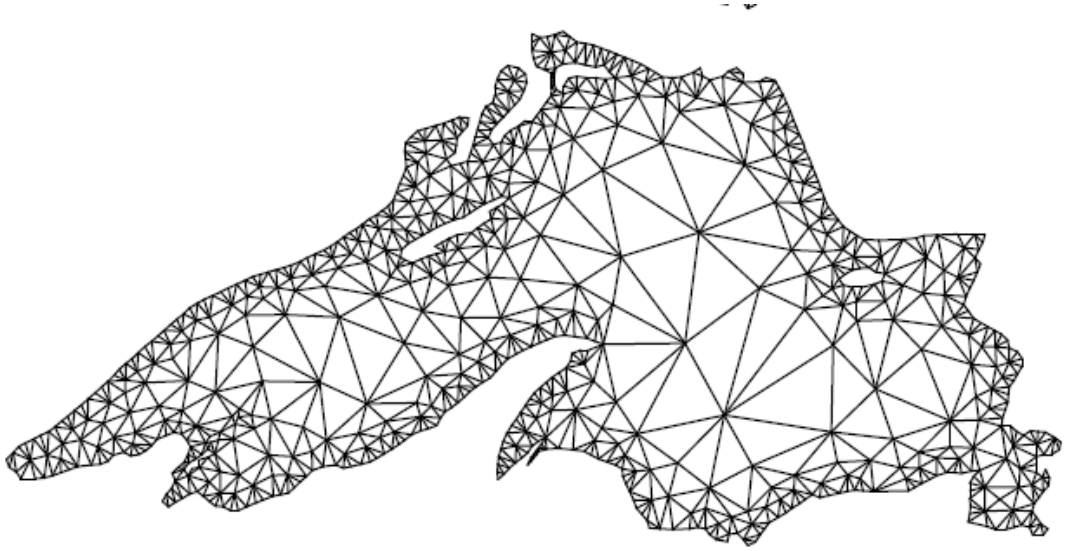

Figure 2: A meshing of Lake Superior, a Euclidean domain. The sample density decreases away from the boundary.

of the mesh, the triangle faces if it is a 2D domain, or the tetrahedra if it is a 3D domain, are of good quality. That is, they must be close to regular (maximally symmetric).

The geometry of the problem comes from the boundary of the domain. A mesh must be produced that uses as few vertices as possible given constraints on element size and quality. In general there will be smaller elements (more samples) near the boundary of the domain than near the centre (see Figure 2)[4].

The dominant paradigm for meshing Euclidean domains is *Delaunay refinement*. Delaunay refinement is a method for improving the element quality of a given Delaunay mesh by strategically inserting new vertices. Focusing on the 2D case, element quality may be measured by the *aspect ratio* of the triangles: the ratio between the circumradius and the length of the shortest edge, for example. A Delaunay refinement algorithm will choose the element with the worst (largest) aspect ratio and insert a new sample at its circumcentre. The new sample is then incorporated into the existing mesh and edge flips are performed to maintain the Delaunay property.

A good survey of Delaunay refinement algorithms can be found in the works of Johnathan Shewchuk [She97, She99]. The essential idea of Delaunay refinement is attributed to Rupert [Rup93, Rup95], although a Delaunay refinement algorithm for meshing curved surfaces was independently proposed by Chew about the same time [Che93]. There are many variations on the Delaunay refinement paradigm, a notable one being the longest side bisection algorithm of Rivara [RI97], where new vertices are inserted on existing edges, rather than at circumcentres.

In some meshing applications, geometry may be imposed by constraints additional to the boundary of the domain. A *constrained Delaunay triangulation* of a collection of points and

[4]Image from [She97]

line segments is a triangle mesh that has the points as vertices and the segments as edges and is in some sense as close to being Delaunay as possible, given the constraints on the fixed edges. Chew [Che89] described how to compute a constrained Delaunay triangulation from a given arrangement.

A related concept is a *conforming Delaunay triangulation*. In this case, we demand that the final triangulation be Delaunay, however the initial constraint line segments can be "split" by the insertion of new vertices. Edelsbrunner and Tan [ET93] demonstrated an algorithm which, given an initial configuration of line segments and vertices, will compute a conforming Delaunay triangulation with an asymptotically optimal (minimal) number of inserted vertices.

Guibas et al. [GR04] considered the problem of maintaining a Delaunay mesh of a domain that is deforming over time. The paper makes a number of interesting observations about the fragility of the Delaunay triangulation; a small deformation of the point set may demand a large change in the Delaunay triangulation.

In some meshing applications, the sample points are not static, but change location slightly with time. Guibas et al. [GR04] analyzed various algorithms for maintaining Delaunay triangulations under small perturbations of the vertices. One notable figure in that paper illustrates a configuration in which only one of the edges of the triangulation is NLD, yet none of the edges are Delaunay edges. This emphasizes the fact that the number of edges that are not locally Delaunay is not necessarily a good measure of how "close" a given triangulation is to being Delaunay.

Finally, it should be mentioned that recent work in meshing Euclidean domains has employed alternate metrics to the Euclidean metric. In particular, "geodesic distances" (or "inner distances") measured by lower bounds of lengths of paths *within* the domain, have proven useful [GGOW08]. In this case the term "Euclidean domain" is arguably no longer appropriate.

2.3 Generalizations and related structures

The Voronoi-Delaunay paradigm is central to the field of computational geometry, and it has spawned many variations and related structures. The basic Delaunay structures for surface representation discussed in Section 3 represent extensions of this paradigm to non-Euclidean domains. Before turning to that section, we review here a few structures that are fundamental and related to the Voronoi-Delaunay structures introduced above.

One natural generalization of the Delaunay triangulation is the *weighted Delaunay triangulation*. It is sometimes referred to as a *regular triangulation*, but this latter terminology has been applied also to quite different concepts (c.f. [FSBS06]) and so we will avoid it.

In a weighted Delaunay triangulation, each vertex p is assigned a weight w_p. The weighted Delaunay triangulation can then be described as the dual of the *power diagram* of $P_w = \{(p, w_p)\}$. The power diagram is analogous to the Voronoi diagram, but the distance associated with each sample is dependent on its weight. The *power distance* from (p, w_p) to $x \in D$ is given by $\pi_p(x) = d_{\mathbb{R}^3}(x, p)^2 - w_p^2$. Then the *power cell* associated with (p, w_p) is given by

$$V_w(p) = \{x \in D \mid \pi_p(x) \leq \pi_q(x) \ \ \forall q \in P_w\}.$$

It can be shown that these power cells are convex polyhedra and the power diagram is the collection of all the power cells. The weighted Delaunay triangulation is the dual of the power diagram in the same way that the usual Delaunay triangulation is the dual of the Voronoi diagram. In fact, if all the weights are zero, the power diagram is exactly the Voronoi diagram described in Definition 2.1, and the corresponding weighted Delaunay triangulation is just the ordinary Delaunay triangulation.

One geometric interpretation of the power diagram is obtained by considering w_p to be the radius of a sphere centred at p. The power distance to x is then the square of the length of the line segment from x to a tangent point on the sphere. More details can be found in [Dey07] or [Ede01].

An interesting power diagram can be constructed as follows: Let $P \subset \mathbb{R}^d$ and let V be the set of Voronoi vertices in the Voronoi diagram of P. Assign to each Voronoi vertex, v, a weight, w_v, equal to the radius of the circumsphere of the Delaunay simplex dual to v. Then the power diagram of $V_w = \{(v, w_v)\}$ is exactly the Delaunay triangulation of P [Dey07][§10.4].

The power diagram represents a distortion of the Voronoi diagram which maintains the convex polyhedron property of the cells. Another natural generalization of the Voronoi diagram of Definition 2.1, is obtained by replacing the Euclidean metric $d_{\mathbb{R}^3}$ with an alternate metric. In this case convexity of the Voronoi cells is no longer assured. In fact for general metrics many desirable properties of the Voronoi cells can be lost, and even defining the Voronoi diagram can be a challenge. For example, the Manhattan (L_1) metric in the plane can yield Voronoi cell bisectors that are regions and not curves [KW88]. Klein and Wood [KW88] developed a set of axioms that, when satisfied by a metric, assure that the Voronoi diagram in the plane is well behaved. Aurenhammer and Klein [AK00] refer to this class of metrics as *nice metrics*. In particular, nice metrics yield connected Voronoi cells with boundaries composed of a finite set of Voronoi cell bisectors that are simple curves.

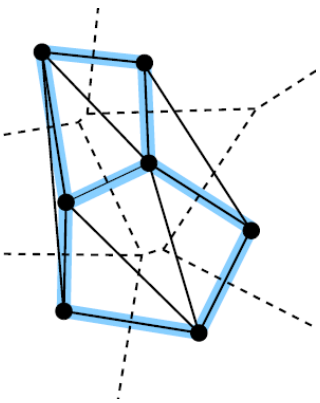

The Gabriel graph consists of those Delaunay edges that intersect their dual Voronoi edges.

An nice example of Delaunay meshes constructed using an alternate metric is in the work of Labelle and Shewchuk [LS03], in which they studied anisotropic metrics. As they remark, the setting they consider is (at least locally) equivalent to the case involving the intrinsic metric of a surface, as will be discussed in Section 3.2.2. However, for computational reasons they consider an approximation to the Voronoi diagram of the Riemannian metric, and in so doing they lose some of its nice properties. In particular, the Voronoi cells may not be connected.

Another structure built on the Delaunay theme that will be of interest to us is the *Gabriel complex*. This is actually a substructure of the usual Delaunay triangulation[5]. Recall that the Delaunay triangulation consists of those simplices that can be contained in a ball whose interior is empty of samples. In dimension d, the Gabriel complex is the simplicial complex which consists of those $(d-1)$-simplices for which the circumscribing ball (i.e., the smallest ball that has all the vertices of the simplex on its boundary) is empty, together with the faces of

[5]Image from [GJ02]

these simplices. In 2D the Gabriel complex is called the Gabriel graph.

A description of the Gabriel graph was published in 1969 [GS69], and it is a well established data structure. However, the higher dimensional analogue that we have described here has a much shorter history. The Gabriel complex in three dimensions was introduced by Petitjean and Boyer [PB01].

It should be noted that there is a potential for confusion because it is also natural to define a Gabriel graph in higher dimensions: it is the simplicial complex consisting of those edges (1-simplices) for which the diametric ball is empty. In contrast, the Gabriel complex in $\mathbb{R}^d$ consists of $(d-1)$-simplices and their faces. So for example, in $\mathbb{R}^3$, the Gabriel complex will consist of triangles and their edges, however these edges need not belong to the Gabriel *graph*. So we must be careful when refering to "Gabriel edges". In this report we are concerned only with the Gabriel complex, and will not consider the Gabriel graph, except in the planar case where the two concepts coincide.

3 Delaunay structures for surfaces

In the previous section we studied Delaunay triangulations of Euclidean domains. In the 2D setting for example, the domain to be triangulated is planar. If geometry enters into the problem, it is due to the boundary of the domain. Thus meshing and refinement algorithms that use adaptive sampling will produce a higher sampling density in the neighbourhood of small features in the domain boundary, and as a rule, the sampling density will decrease as the distance to the boundary increases, as demonstrated in Figure 2.

We now turn our attention to non-Euclidean domains. In particular our interest is in the 2D surfaces that bound objects in 3D space. In this exposition we will assume that the surface $S \subset \mathbb{R}^3$ is a smooth (C^∞) manifold surface that is compact, orientable and has no boundary.

We begin in Section 3.1 by introducing extentions of triangulations to surfaces (i.e., non-euclidean domains). For our purposes the appropriate structure with which to represent (approximate) S is a manifold triangle mesh, which is introduced in Section 3.1.2. The Delaunay paradigm is used to define three distinct mesh structues which are introduced in Section 3.2. These three structures are the primary focus of this survey. In Section 3.3 we develop the characterization of these Delaunay structures as combinatorial duals to a Voronoi diagram corresponding to a particular choice of metric on S. This provides a unifying framework in which to compare their differences and similarities. Finally, in Section 3.4 we discuss related structures which employ the Delaunay paradigm but are not the focus of this survey.

3.1 Triangulations and triangle meshes for surfaces

We used a simplicial complex to define a triangulation of a Euclidean domain. In order to triangulate a non-Euclidean domain with Euclidean simplices, we either need to warp the simplices to fit the domain, or distort the domain into a piecewise linear representation. The difference is really just a matter of perspective, and for geometry processing the latter is more natural.

3.1.1 Surface triangulations

A *triangulation of* S is a simplicial complex, K, together with a homeomorphism $h : K \to S$. We can think of this as drawing triangles on S (via the images of the 1-simplices in K). If a set of samples $P \subset S$ is given, then a triangulation of P on S is a triangulation $h : K \to S$ such that the vertices of K are mapped bijectively onto P. It can be shown that the surfaces we consider always admit a triangulation, however it is not true that any set of samples $P \subset S$ is sufficient to be the vertex set of a triangulation. The topology of the surface will demand a minimum number of samples.

The *geodesic distance* between two points $p, q \in S$, denoted $d_S(p, q)$, is the length of the shortest path between them. Such a shortest path always exists, and it is called a *geodesic*. If S is viewed as a Riemannian manifold, the metric d_S arises naturally from its Riemannian metric tensor and as such does not depend on the embedding in $\mathbb{R}^3$ for its definition. (Although the embedding can be used to define the Riemannian metric tensor, this tensor remains well defined in the absense of any embedding.) We call this metric the *intrinsic* metric on S and likewise all properties which depend only on this metric for their definition, are deemed intrinsic.

A *geodesic disk* centred at $c \in S$ with radius r is the set $B_S(c; r) = \{x \in S \mid d_S(c, x) < r\}$. In general a geodesic disk need not even be a topological disk (it can wrap around on itself), but for any $c \in S$, there is an r small enough to ensure that $B_S(c; r)$ is a topological disk.

For a sufficiently dense sample set P, the *intrinsic Delaunay triangulation* (iDt) of $P \subset S$ can be defined analogously to the Euclidean Delaunay triangulation. We demand that each triangle be circumscribed by a unique geodesic ball that is empty of all sample points. Sampling criteria that guarantee that the iDt is well defined were proposed by Leibon and Letscher [LL00, Lei99]. A loose estimate on the actual number of samples this implies was published in [OI03]. For surfaces, the sampling criteria has recently been relaxed [DZM08].

The empty ball criterion defines the connectivity of the Delaunay triangulation, however even with a sufficient sampling, and assuming general position, this does not define a unique triangulation according to our definition. In other words, the combinatorial structure of the simplicial complex K is dictated by the empty ball property, but the homeomorphism $h : K \to S$ is not fully constrained. We desire the edges of the triangles, i.e., the images of the 1-simplices in K, be described by well defined curves in S. It turns out that the density requirements that ensure a well defined topology for the iDt are also sufficient to ensure that there is a unique geodesic of minimal length between each pair of neighbouring samples in P. A triangulation of S whose edges are composed of geodesics is called a *geodesic triangulation*[6].

The Voronoi diagram induced by the intrinsic metric, with Voronoi cells $\mathcal{V}(p) = \{x \in S \mid d_S(x, p) \leq d_S(x, q)\ \forall q \in P\}$, is the *intrinsic Voronoi diagram* (iVd) of P on S. The iDt can be seen to be the dual of the iVd, just as in the Euclidean case: Delaunay triangles are dual to Voronoi vertices and Delaunay edges have dual Voronoi edges. However, unlike the iDt, the iVd is well defined regardless of the sampling density.

[6] An alternative to geodesic curves is mentioned in [LL00]. They propose to employ "middle planes". This approach has appeal in that it is amenable to extensions to higher dimensions, however the details of their proposal were never published.

3.1.2 Triangle meshes

The goal is to employ the Delaunay paradigm to produce a discrete (piecewise linear) structure M that will serve as an approximation to S. Thus we are interested in constructing a simplicial complex K that defines a triangulation of S. However, it will be convenient to separate the geometry from the combinatorial topology that is inherent in a simplicial complex as defined in Section 2.

An *abstract simplicial complex* $\mathcal{K}$ is a collection of finite nonempty sets such that if $A \in \mathcal{K}$, then so is every nonempty subset of A. An element A of $\mathcal{K}$ is called a *simplex* of $\mathcal{K}$. The dimension of a simplex is one less than the number of elements it contains.

Following [PSS01][7], we define a *triangle mesh*, M, as a pair $M = (X, \mathcal{K})$, where X is a set of n point positions in $\mathbb{R}^3$ and $\mathcal{K}$ is an abstract simplicial complex composed of three types of subsets of $\{1, \ldots, n\}$: vertices $\{i\}$, edges $\{i, j\}$ and faces $\{i, j, k\}$.

The abstract simplices of M, together with the associated geometric position information in X, define (possibly degenerate) Euclidean simplices in $\mathbb{R}^3$, as defined in Section 2. The union of these Euclidean simplices is the *geometric realization* of M. The reason for separating geometry and connectivity information is that we can allow self intersections of the mesh surface. Thus the collection of simplices that form the geometric realization of M may not be a simplicial complex as defined in Section 2. We identify vertices, edges and faces with their geometric realizations.

The *star* of a vertex i is the set of faces that are incident to i. A *manifold triangle mesh* is a triangle mesh such that the star of every vertex is homeomorphic to a closed disk. We will focus on manifold triangles meshes *without boundary*, which means that every vertex lies in the topological interior of the disk defined by its star.

Ideally we wish to produce a manifold triangle mesh M that is a good approximation to S. Exactly what is meant by a good approximation will be discussed in detail in Section 4. This problem, of producing a mesh to represent a given surface is known as *meshing* and has seen considerable attention. In Section 5 we will survey the literature in meshing and related applications.

3.2 Three Delaunay structures

In this section we examine means of employing the Delaunay paradigm to define meshes that represent $S \subset \mathbb{R}^3$. In this setting we will present three distinct Delaunay structures. These are triangle meshes which are defined via the Delaunay paradigm.

In Section 3.2.1 we discuss the restricted Delaunay triangulation (rDt) which results from the criterion that each mesh triangle face t be contained in a (3D) Euclidean ball centred on S and such that the vertices of t are on its boundary and no sample points of P are in its interior. In Section 3.2.2 we look at the mesh that is combinatorially equivalent to the iDt. This mesh, dubbed the iDt-mesh, has triangle faces whose vertices belong to a triangle in the iDt: they are defined by empty geodesic disks. Then in Section 3.2.3 we explore Delaunay meshes whose faces are defined by empty geodesic disks on the surface of the mesh itself.

[7]This paper is also the one in which the phrase "digital geometry processing" was coined.

One way to view these structures in a unifying framework is to consider each of them to be the result of a different choice of metric on S. The metric yields a Voronoi diagram on S and the corresponding mesh structure is seen as the mesh which is a combinatorial dual to this Voronoi diagram. For the rDt the metric is the Euclidean distance whereas for the iDt-mesh we employ the intrinsic metric of S. For Delaunay meshes the metric employed properly belongs to the mesh itself, however if the mesh is a good representation of S we are able to port this metric to S itself. This perspective of viewing the different Delaunay structures as manifestations of different choices of metric is discussed further in Section 3.3.

3.2.1 Restricted Delaunay triangulation (rDt)

Of the three main Delaunay structures we present in this section, the restricted Delaunay triangulation (rDt) has been the most extensively studied in the geometry processing community. The rDt can be defined as the mesh that is dual to the *restricted Voronoi diagram* (rVd). The rVd was formally introduced by Edelsbrunner and Shah [ES94], where they gave conditions on that guarantee that the rDt is homeomorphic to the original surface, S. The rDt itself appeared in an earlier meshing algorithm by Chew [Che93], but no explicit reference to the rVd was made in that work. Instead, the faces of the mesh were defined directly in terms of empty Euclidean spheres. There is an unfortunate nomenclature conflict here. Despite its name, the restricted Delaunay triangulation is a mesh, not a triangulation as discussed in Section 3.1.1. However, the name is well established.

In the rVd, the Voronoi cell of a point $p \in P$, is given by $V(p) \cap S$, where $V(p)$ is the Voronoi cell of p in the standard Euclidean Voronoi diagram of P in $\mathbb{R}^3$. Thus the rVd is well named: it is the 3D Euclidean Voronoi diagram restricted to S. Equivalently, we can view it as the Voronoi diagram resulting from choosing the Euclidean distance as a metric on S, thus the restricted Voronoi cell of $p \in S$ is given by $V|_S(p) = \{x \in S \mid d_{\mathbb{R}^3}(x,p) \leq d_{\mathbb{R}^3}(x,q)\ \forall q \in P\}$.

The rDt is the mesh that is obtained when Voronoi neighbours in the rVd are connected by mesh edges; the faces of the rDt are Euclidean simplices dual to the Voronoi vertices in the rVd. The rDt can be equivalently characterized as consisting of those triangle faces of the 3D Delaunay tetrahedralization for which the dual Voronoi edge intersects S. This latter characterization leads to the property, exploited by Chew [Che93], that the faces of the rDt are contained in empty Euclidean spheres centred on S.

Since the initial works of Chew and Edelsbrunner and Shah, many meshing and surface reconstruction algorithms have been based upon an the rDt and an analysis of the rVd. In particular, the first provably correct surface reconstruction algorithms [AB98, ACDL00] exploited the work of Edelsbrunner and Shah. We will review these and other algorithms in Section 4 and also in section 5.

3.2.2 Intrinsic Delaunay triangulation mesh (iDt-mesh)

Using the Euclidean distance $d_{\mathbb{R}^3}(p,q)$ to measure the distance between two points $p, q \in S$ is simple and convenient in practice. However, there are advantages to using the metric of geodesic distance: $d_S(p,q)$.

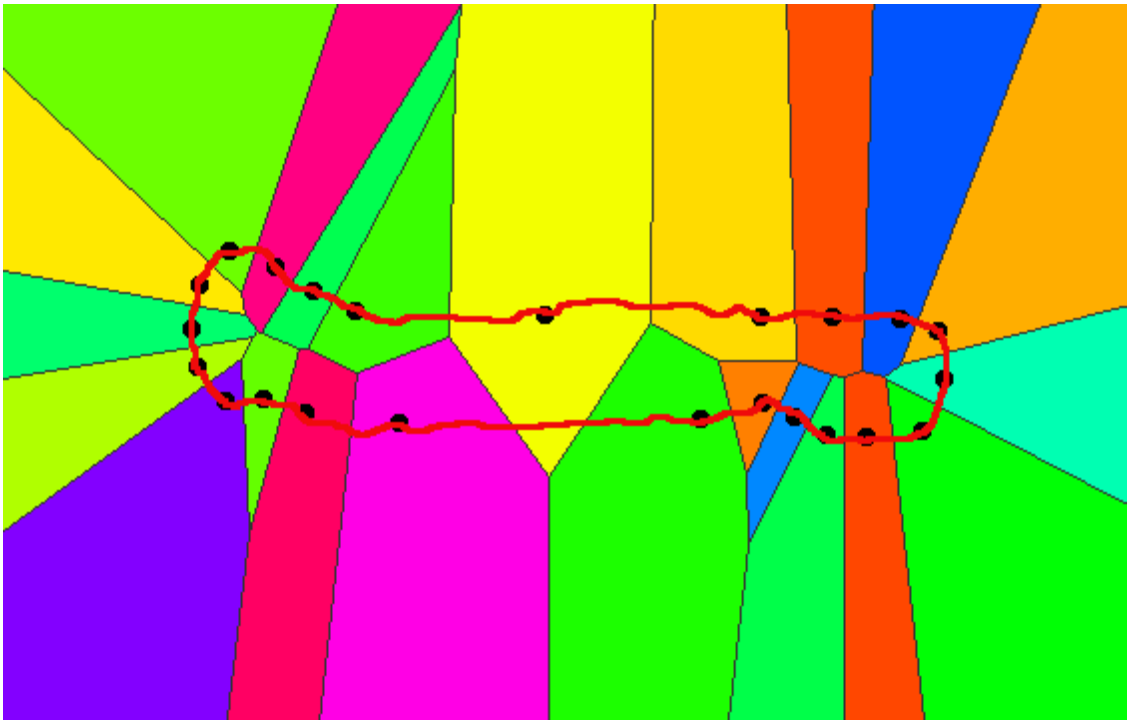

Figure 3: In a restricted Voronoi diagram the Voronoi cells need not be connected, as shown by the restriction to the red curve of the yellow Voronoi cell near the centre of the figure.

We have the relation $d_{\mathbb{R}^3}(p,q) \leq d_S(p,q)$. The difference between $d_{\mathbb{R}^3}(p,q)$ and $d_S(p,q)$ is most pronounced when p and q lie close in the ambient space but far on the surface, as would be the case if we chose p and q to lie at the centre, but on opposite sides on the surface of a pancake, for example. (See Figure 3.)[8]

It can be shown that the Voronoi cells of the intrinsic Voronoi diagram are connected in the topological sense. The argument is exactly as given by Aurenhammer and Klein [AK00] in the context of nice metrics. This is in contrast to the Voronoi cells of the rVd, which may contain more than one connected component. An example may be easily constructed using the pancake mentioned above. Thus the intrinsic metric is in a sense "nicer" than the Euclidean metric in the context of surfaces embedded in $\mathbb{R}^3$.

Another argument in favour of the intrinsic metric is that we are liberated from dependence on the ambient space. This argument is weak in the context of computer graphics where an attempt to divorce the surface from the surrounding space appears to be an exercise in senseless abstraction. However, for many other applications, especially when the surface lies in a higher dimensional space, the intrinsic metric is the only sensible choice. This is because to work in the high dimensional ambient space renders one subject to "the curse of dimensionality": some computations become intractable in a high dimensional setting. Finally, an analysis based on the intrinsic metric allows us to tap into a wealth of results in Riemannian geometry, and in some cases stronger statements can be made in this context

[8]Figure composed with the aid of Chew's Voronoi diagram applet: http://www.cs.cornell.edu/home-/chew/Delaunay.html

than have been made with more traditional extrinsic methods [DZM08].

The mesh that is dual to the iVd is called the *iDt-mesh* and for sufficient sampling it will be a manifold triangle mesh. The iDt-mesh is the mesh with vertices P and whose connectivity is that of the intrinsic Delaunay triangulation discussed in Section 3.1.1.

3.2.3 Delaunay meshes

The rDt and the iDt-mesh have both been well studied and proven themselves to be important structures in theory and practice. However, they share a trait that can be problematic in practice: for their definition, both the rDt and the iDt-mesh depend in an essential way on the underlying surface S that is being approximated. This is an obvious difficulty in surface reconstruction, for example, where S is known only by the samples P, and some assumptions on the density of P and on the regularity of S. In this case, when one has constructed a mesh M whose vertices are P, one would like to make claims of the fitness of M as a representation of S. If it could be demonstrated that M is an rDt or an iDt-mesh, then an appeal could be made to the known qualities of these structures as surface approximations (results which we will discuss in Section 4). Unfortunately, such a direct approach is impossible since S is unknown.

In general, given a mesh M, there are no known algorithms which enable one to decide if M is an rDt or an iDt-mesh of some smooth surface S. It would be nice to have a family of meshes, $\mathcal{F}$, which display good approximation properties and for which the inclusion of a given mesh, M, in $\mathcal{F}$ can be determined algorithmically based on the properties of M alone.

In a notable work in this direction Petitjean and Boyer [PB01] introduced the notion of *regular interpolants*. They defined these by means of the *discrete medial axis* of a mesh M. This is a subset of the 3D Voronoi diagram of the vertices of M: It is the union of all of the Voronoi edges, facets and vertices which are not dual to a simplex of M. They then defined the *local thickness* of M at a vertex p, which is essentially a measure of the distance from p to the discrete medial axis. Then M is a regular interpolant if for each $p \in M$ the local thickness is larger than the circumradius of each triangle in M that is incident to p. Although the definition is attractive in its simplicity, it suffers a fatal flaw. The problem is the existence of *sliver tetrahedra* in the 3D Voronoi diagram.

It was pointed out by Amenta and Bern [AB98] that sliver tetrahedra have dual Voronoi vertices that may be arbitrarily close to the original surface S. Sliver tetrahedra are those whose vertices are distributed near the equator of their circumsphere. They are flat and have a small volume, but no short edges (See Figure 4)[9]. Such tetrahedra cannot be eliminated by simply increasing the sampling density and as a consequence, it follows that the discrete medial axis may also be arbitrarily close to S, regardless of the sampling density. Thus regular interpolants are too restrictive; we are not guaranteed to achieve a regular interpolant for any given density based sampling criteria.

Recently an alternative approach has begun to emerge in the form of yet another Delaunay structure. In 2005 Bobenko and Springborn presented a formal study and uniqueness proof of for the intrinsic Delaunay triangulation of the vertices of a piecewise flat surface [BS05]

[9]These figures are from Kolluri et al. "Spectral reconstruction from noisy point clouds" (SGP 2004), and Amenta et al. "A new Voronoi based surface reconstruction algorithm" (SIGGRAPH 1998) respectively.

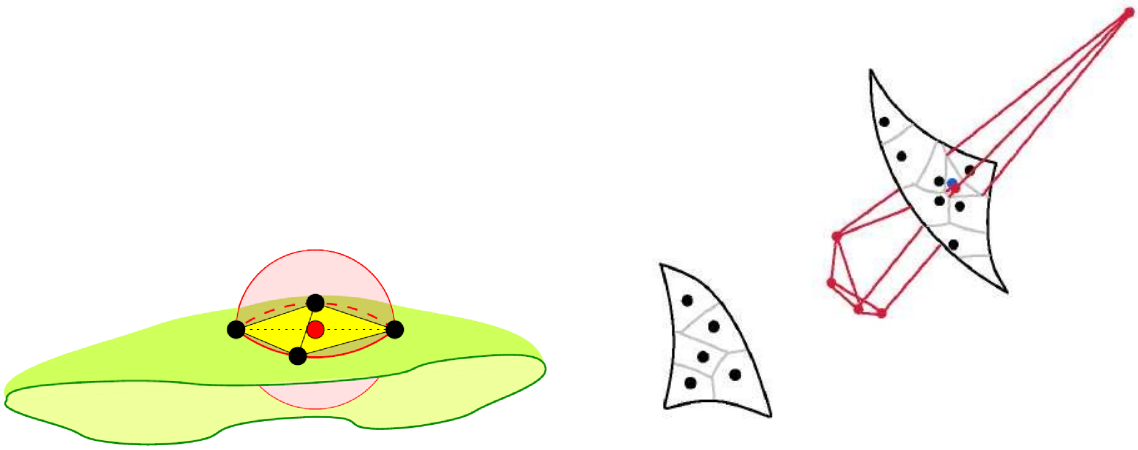

(a) (b)

Figure 4: A sliver tetrahedron can have its circumcentre (Voronoi vertex) arbitrarily close to the surface (a). The discrete medial axis includes only those Voronoi edges and facets that are not dual to faces or edges of the mesh, but it includes *all* the Voronoi vertices. In (b) the Voronoi cell of the blue sample is drawn in red and the restricted Voronoi cells are drawn in grey on the surface. Because the blue sample belongs to a sliver tetrahedron, at least two of the edges of its 3D Voronoi cell will come close to the surface: they are the two edges that are dual to the two facets of the sliver tetrahedron that don't belong to the surface mesh.

(formally published in [BS07]). Although the structure had been studied previously, [Riv90], [ILTC01], Bobenko and Springborn's work was the first to recognize its potential utility in geometry processing. For our purposes we may interpret a piecewise flat (pwf) surface to be a manifold triangle mesh, M. The class of pwf surfaces is actually more general than this, and a comprehensive study of the geometry of these objects can be found in the treatise by Aleksandrov and Zalgaller [AZ67].

The sample set P is the vertices of M. The definition of the iDt of M is based on the idea of empty circumdisks and is not unlike the definition of the geodesic iDt of P on a smooth surface S. The principle difference is that by extending the concept of a triangulation of the surface, the iDt of the vertex set of a pwf surface is well defined without imposing restrictions on the vertex set P. In particular, it is not required that the vertices or edges of a triangle be distinct; this allows the triangulation to contain "loop" edges that terminate at the same vertex at each end. Also, two vertices may have more than one edge between them.

We say the triangulation of M is *proper* if it lacks these anomalous features. In other words it defines an abstract simplicial complex. Figure 3.2.3 gives an example of a Delaunay triangulation that is not proper. As discussed in Section 4.1, the indications are that the iDt of M will be proper if M is a "good" representation of a sufficently well sampled smooth surface S. However, that assertion has not been rigorously quantified.

An algorithm and data structure to generate the iDt of a pwf surface was presented by

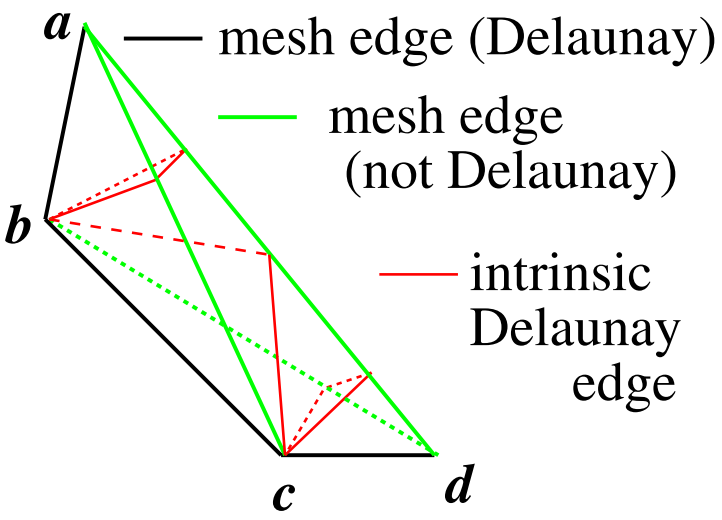

Figure 5: An example of a non-proper intrinsic Delaunay triangulation of a pwf surface. The surface here is a tetrahedron defined by the triangle faces $[a, b, c]$, $[b, c, d]$, $[a, c, d]$, and $[a, d, b]$. The intrinsic Delaunay triangulation is defined by the red and black geodesics. The black geodesics coincide with the original mesh edges, but the red geodesics traverse faces that define the tetrahedron. It is composed of "triangles" $[a, b, b]$ which has a loop edge and shares an edge with itself; $[c, d, c]$ which has the same characteristics; $[b, c, b]$, and $[c, c, b]$ which share two edges and each have a loop edge.

Fisher et al. [FSBS06]. The principle interest in this structure is the nice properties it lends to the discrete Laplace-Beltrami operator described by the cotan formula [BS05, FSBS06]. The disadvantage is that a separate, somewhat complicated, data structure must be maintained to describe the iDt, in addition to the triangulation inherent in the mesh itself. The obvious solution to this problem is to seek a mesh representation M such that the iDt of M coincides with its inherent triangulation. Such meshes are called *Delaunay meshes* and it has been shown [DZM07a] that they can be reliably constructed from an initial mesh by an edge flipping algorithm. If the initial mesh is too coarse, new vertices may need to be added, and in fact, at the cost of adding more vertices, the geometric distortion from the initial mesh can be made to vanish.

An interesting observation is that flipping a mesh edge to make it locally Delaunay reduces (or at least never increases) the surface area of the mesh. Unlike the rDt and the iDt-mesh, the Delaunay mesh associated with a sample set P is not unique, even if the surface topology is constrained to conform with a reference surface. Also, as will be discussed in Section 4, regardless of how densely P was sampled on a smooth surface S, some Delaunay meshes on P can be very poor representations of S.

In summary, as a data structure the Delaunay mesh is not as well understood as the more mature rDt and iDt-mesh. There are no known existence results for Delaunay meshes on a sample set $P \subset S$, and, in contrast to the rDt and the iDt-mesh, it is known to be not unique. However, intuition dictates that if P represents a sufficiently dense sampling, there will exist a "nice Delaunay mesh" that possesses the desirable approximation guarantees of the rDt and iDt-mesh. If such a Delaunay mesh can be characterized by local properties of

the mesh itself, such as the dihedral angles across edges, it would have the advantage over Delaunay structures that make explicit reference to S in that it can be explicitly verified in the absence of S.

3.3 The Delaunay structure associated with a metric

As has been mentioned, the three Delaunay structures we have discussed so far in this section can be seen in a unifying framework if we consider them to be manifestations of particular choices of metric on S. Indeed, given an arbitrary metric $d : S \times S \to \mathbb{R}$, the associated Voronoi diagram of P on S is naturally defined by replacing the Euclidean metric $d_{\mathbb{R}^3}$ with d in Definition 2.1. Thus the Voronoi cell associated with $p \in P$ is $\{x \in S \,|\, d(x,p) \leq d(x,q), \forall q \in P\}$. The associated Delaunay structure is then the dual of this Voronoi diagram.

To obtain the rDt, the Euclidean distance is used: $d = d_{\mathbb{R}^3}$ and for the iDt-mesh the intrinsic (geodesic) metric of the surface is employed: $d = d_S$. Although the Euclidean distance is perhaps more commonly used, it can be argued that the intrinsic metric is "nicer" because it always leads to topologically connected Voronoi cells, as was mentioned in Section 3.2.2. However, if P is sufficiently dense on S then the Voronoi cells of the rVd will be connected.

For a Delaunay mesh, M, under certain conditions, we can also view it as a Delaunay structure resulting from a choice of metric on S. Define a *conforming homeomorphism* as a homeomorphism $h : M \to S$ that leaves the samples in P invariant: $h|_P = \mathrm{id}_{\mathbb{R}^3}$. A mesh M is a *conforming mesh* for S if it admits a conforming homeomorphism. The metric of a conforming Delaunay mesh can be used on S: for $x, y \in S$, $d(x,y) = d_M(h^{-1}(x), h^{-1}(y))$. We can then view M as the dual mesh to the Voronoi diagram of its vertices on S with respect to this metric. Such a characterization of a Delaunay mesh is somewhat circular, but it may provide a convenient framework for analysis. An important task will be to identify local characteristics that guarantee that a Delaunay mesh conforms to some well sampled surface S.

This highlights another area that has not been studied in depth: how does the Voronoi diagram of P on S change when the metric is perturbed? Specifically, it would be interesting to identify under which conditions the metric can be perturbed while ensuring that the dual mesh remains homeomorphic to S. For example, are conditions which are sufficient to guarantee that the rDt is homeomorphic to S also sufficient to guarantee that the iDt-mesh is homeomorphic to surf?

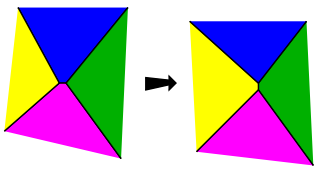

Although all three Delaunay structures, with the appropriate qualifications, are expected to provide good representations of S, it is known that they can all be combinatorially distinct structures, regardless of the sampling density [DZM07b]. The differences result from samples that are "almost not in general position". As the sampling density increases, the differences between the associated metrics over the span of a Voronoi cell will decrease. However, there will always be small differences and the metric perturbation near two Voronoi vertices that are very close in one metric may result in a Voronoi edge disappearing in another metric. Such discrepencies are expected to arise in the neighbourhood of sliver tetrahedra in the 3D Delaunay triangulation for example.

3.4 Alternate and related structures

The Delaunay paradigm has proven itself extremely useful for surface representation. However other triangle mesh families have also shown their utility. We will review some of these alternative structures in this subsection. Most of them can be seen to be generalizations of (or at least closely related to) one or more of the Delaunay structures discussed above.

In Section 3.3 we saw that each of the three Delaunay structures, the rDt, the iDt-mesh, and the Delaunay mesh, could in a sense be interpreted as resulting from a specific choice of metric for S. Indeed the range of possible Delaunay-type structures is at least as broad as the range of choices of metric that can be imposed upon S.

In adaptive sampling theory, the density of the sample set P on S is governed by some sizing function, which typically specifies a higher sampling density in regions of higher curvature. Various sizing functions are discussed in Section 4. In some contexts it is natural to want to modulate the intrinsic distance function on S by an appropriate sizing function. A nice example of where exactly such an approach is used to produce a Delaunay structure is in the farthest point sampling algorithm presented by Peyré and Cohen [PC03, PC06]. In this algorithm the Voronoi diagram of P with respect to the modified metric is maintained while iteratively adding samples at points on S that are farthest from the current P with respect to this metric. The final mesh that is produced is the dual to the resulting Voronoi diagram.

Approaching the modified metric from a more discrete perspective, Glickenstein [Gli05] studied *weighted triangulations* of piecewise flat surfaces. The corresponding Delaunay structure is a *regular triangulation*, which is the dual of the *power diagram* and these are defined just as in the Euclidean setting. Such structures have been studied extensively in Euclidean domains [AK00], [Ede01], but the systematic study of them in the context of surface meshes is in its infancy.

The Gabriel complex introduced in Section 2 plays a role in surface representation. For a well sampled smooth planar simple curve, the Gabriel graph will contain the polygonal reconstruction of the curve, i.e. the unique polygon that connects the samples in the order specified by the curve. It is assumed that for a sufficiently well sampled surface, the Gabriel complex of the samples will contain a mesh that is homeomorphic to the surface, however this assumption has not been formally demonstrated in the literature (although with an additional uniformity assumption on the samples, Cheng and Dey [CD07] were able to prove the existence of a mesh all of whose faces satisfy the Gabriel property). We refer to a mesh whose faces all have the Gabriel property as a *Gabriel mesh*. It is easy to show that a Gabriel mesh must be a Delaunay mesh, as well as being a substructure of the 3D Delaunay tetrahedralization, so insights into one should inform our understanding of the other.

The flow complex introduced by Giesen et al. [GJ02, GJ03] represents another structure that has close ties with the Voronoi diagram and the Delaunay triangulation. It is a triangle mesh whose vertices include the original samples as well as additional points which are certain critical points of the distance function defined from the original sample set.

Finally, there has been recent interest in imposing a stronger criterion than the Delaunay condition on a triangle mesh. Specifically, one may demand that no triangle in the mesh contains an obtuse angle. Such a *non-obtuse mesh* is necessarily a Delaunay mesh, however in

general a Delaunay mesh will contain obtuse angles. For a given sample set P, a non obtuse mesh interpolating P will not exist in general, however approximation algorithms exist which will produce a non-obtuse mesh with small perturbations of the original samples [LZ06]. Such meshes are convenient for certain algorithms, such as the fast marching algorithm of Kimmel and Sethian [KS98], which become much more complicated in the presence of obtuse angles.

The idea of a non-obtuse mesh extends to higher dimensions where the resulting simplicial complexes are referred to as *well centred* meshes. This is an allusion to the fact that the circumcentre of each simplex must reside within the simplex itself. Algorithms for producing such structures have recently been proposed [VHGR08]. In this context the interest in well centred meshes stems from the theory of discrete exterior calculus, where the theory is simplified if such a mesh can be assumed. However, Glickenstein [Gli05] has shown that a weaker condition suffices to supply this simplification, and for two dimensional domains it is enough that the mesh be a Delaunay mesh.

4 Surface approximation theory

Having reviewed the principle triangle mesh structures that have been employed to represent surfaces, we now turn our attention to the question of what constitutes a good representation of a surface. Of course the answer to such a question depends very much upon the intended application of the structure.

The principle concern in many, if not most applications is the accuracy of the geometric approximation. The quality of geometric approximation can be evaluated based upon a three level hierarchy:

1. topological consistency
2. pointwise approximation
3. normal approximation

At the first level, we require for topological consistency that the approximation M be homeomorphic to the original surface S. In other words we demand that M and S both have the same genus. In this sense the first level is of a different nature from level (2) and (3) in that it is either attained or it is not. The latter geometric criteria can be viewed as convergence issues where the quality of the approximation is expected to improve with increased sampling density.

Normally, in order to evaluate the quality of the pointwise and normal approximation of M, one assumes that the criterion of topological consistency has been met. This then allows for the definition of a homeomorphism, a bijective, bicontinuous mapping: $\xi : M \to S$. Then the pointwise approximation error of M can be evaluated as some measure of the difference between x and $\xi(x)$ for each $x \in M$. Likewise, the quality of the normal approximation is judged by a comparison of the normal to M at x with the normal to S at $\xi(x)$.

A family of meshes will be useful as representations of the geometry of S if it can be shown that they will converge to S in both position and normal as the number of vertices is increased. The two types of convergence need to be evaluated independently. A classic

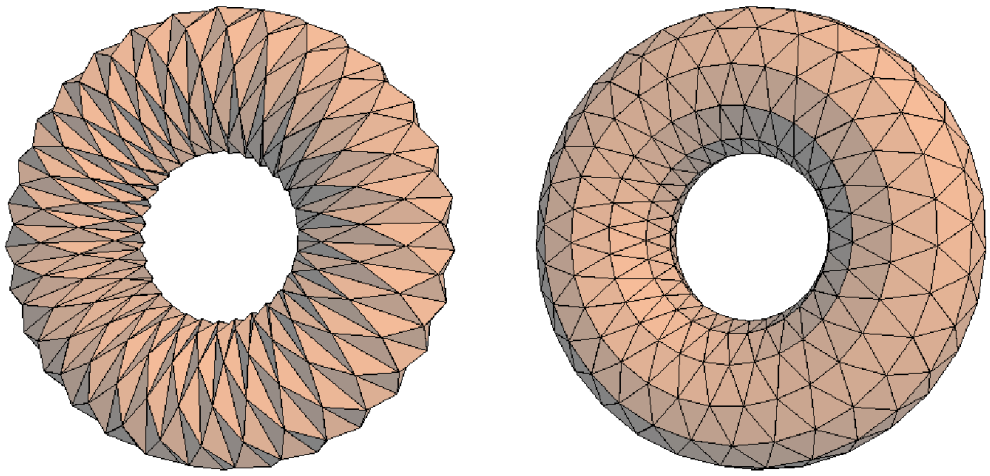

(a) (b)

Figure 6: (a)A Schwarz torus. The Schwarz lantern is a triangle mesh of a cylinder that does not display normal convergence as the sampling increases. Here we have used the same tessellation pattern on a torus. (b) A Delaunay mesh on the same vertex set. The triangle normals are more consistent with the expected surface normals.

example of a family of meshes that converges pointwise to a cylinder, but which display no normal convergence (hence no convergence in surface area), is the *Schwarz lantern* [HPW06] (see Figure 6).

The problem of creating a triangle mesh that is a good geometric representation of S is called *meshing*. The problem involves not only sampling S, but also deciding how those samples are connected together to form the triangle mesh. The Delaunay paradigm has been exploited in both these aspects of meshing.

The connectivity of the samples will influence the geometric accuracy of the mesh approximation. However, as we discuss in Section 4.3, another criterion that is influenced by connectivity is element quality: how close are the triangle faces to equilateral triangles. Optimizing for geometric accuracy alone may lead to meshes with poor element quality. This is undesirable in some applications.

4.1 Sampling and topological consistency

Perhaps the most fundamental issue that must be addressed when meshing or reconstructing a surface is determining the distribution of samples necessary to ensure that the final mesh is topologically equivalent to the original surface. Ideally the sampling density is specified by an adaptive sampling criterion. Such criteria generally impose restrictions on the sampling density based on local curvature properties as well as semi-local properties relating to some notion of the distance to "the other side" of the surface. We refer to functions that can be used to modulate the sampling density in this way as *sizing functions*. These functions take positive values which can be thought of as having the units of distance. Thus at each point on the surface, a sizing function specifies a radius within which a certain proportion of

representative samples is expected. The square of such a function can, for example, be used to define a weighted area measure for governing stochastic sampling.

Perhaps the best known sizing function is the *local feature size* (lfs), introduced by Amenta and Bern [AB98]. At a point $x \in S$, the lfs at x, denoted $\rho_f(x)$ is given by the Euclidean distance from x to the medial axis of S. The *medial axis* of S is the closure of the set of points in $\mathbb{R}^3$ whose distance to S is realized by more than one distinct point in S. Equivalently, it is the closure of the set of centres of the *medial balls*: Euclidean balls which contain no points of the surface in their interior and are maximal in the sense that they are contained in no other Euclidean ball which possesses this property.

The local feature size becomes smaller in the presence of higher curvature. This is a property that is expected in any sizing function. Another property possessed by the lfs is that it is *Lipschitz continuous*, specifically it is a 1-Lipschitz function, which means that for any $p, q \in S$, $|\rho_f(p) - \rho_f(q)| \leq d_{\mathbb{R}^3}(p, q)$. This is especially convenient for analysis because it enables us to bound $\rho_f(x)$ in a neighbourhood of a point p if we have a bound on $\rho_f(p)$.

In terms of the lfs, a typical sampling criterion would demand that for any $x \in S$ there be a $p \in P$ such that $p \in B_{\mathbb{R}^3}(x; \epsilon\rho_f(x))$, the Euclidean ball of radius $\epsilon\rho_f(x)$ centred at x. That is, no point x on the surface is farther than $\epsilon\rho_f(x)$ from a sample point, where $0 < \epsilon \leq 1$. A sample set that satisfies this criteria is called a *lfs ϵ-sample* set for S. Amenta and Bern [AB98] demonstrated that the rDt will be homeomorphic to S provided that P fulfills such a sampling criteria, with $\epsilon < 0.1$.

Their proof relies upon a fundamental result by Edelsbrunner and Shah [ES94] in which it was demonstrated that the rDt would be homeomorphic to S provided that the rVd satisfied the *closed ball property*. This property demands that each Voronoi cell be a topological disk (a closed 2-ball), that two neighbouring Voronoi cells share no more than a single Voronoi edge (i.e. that the intersection of two Voronoi cells is either empty or a topological 1-ball), and that the intersection of three Voronoi cells, if non-empty, be a single Voronoi vertex (topological 0-ball). The result applies to higher dimensional objects; the definition of the closed ball property extends in an obvious manner.

The closed ball property can be viewed as a qualitative measure of the fitness of a sample set for capturing the topology of S. It has been observed that if the iVd of the vertices of a mesh, M, satisfies the closed ball property, then the iDt of M will be a proper triangulation [DZM07b]. This is the observation behind the as yet unquantified assertion that a necessary condition for a mesh to be a good representation of a smooth surface is that it have a proper iDt.

Observing that the closed ball property applies as a topological consistency criterion not just to the rVd, but to any Voronoi diagram defined on S, Dyer et al. [DZM08] demonstrated that the iDt-mesh is homeomorphic to S if no point $x \in S$ is farther than $0.2\rho_f(x)$ from the nearest sample point. This result was obtained via intrinsic sizing functions introduced by Leibon and Letscher [LL00]. In the latter work, more stringent sampling criteria were developed to guarantee a well defined iDt of S.

In the surface reconstruction and meshing literature there have appeared many other works that presented algorithms which produced meshes provably homeomorphic to the target surface. Often these meshes are closely related to a Delaunay structure. We will review the important examples of such algorithms in Section 5.

Finally, we mention that there are other notions of topological consistency. Besides demanding that M be homeomorphic to S, one may require that they are *ambient isotopic*. This is a stronger requirement than homeomorphism. It essentially requires that the ambient space can be continuously deformed to bring M to coincide with S, but without inducing any self-intersections along the way. As an example, a tube tied in a trefoil knot[10] is homeomorphic to a torus, but the two are not ambient isotopic. Amenta et al. [APR03] showed that if M and S are sufficiently close that a homeomorphism $\xi : M \rightarrow S$ is given by the projection to the closest point on S, then they are ambient isotopic.

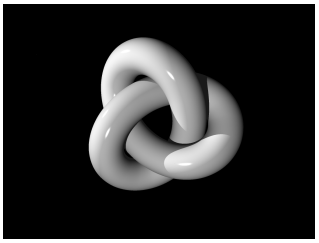

The trefoil knot is homeomorphic, but not ambient isotopic to a torus.

In light of that result, it is seen that most sampling criteria which guarantee that M is homeomorphic to S, e.g. [AB98], [DZM08], are sufficient to guarantee an ambient isotopy as well.

Another notion of topological consistency is the concept of homotopy equivalence. This concept is of little interest in the context of manifold surface approximation, but it becomes much more relevant when constructing a mesh to represent a non-manifold object. In this context, homotopy equivalence is a weaker criterion than a homeomorphism. Recent work by Gao et al. [GGOW08] provides guarantees of homotopy equivalence in the meshing of arbitrary planar domains.

Also, it should be pointed out that the three point hierarchy presented in the introduction to this section is subject to debate. Specifically, it can be argued that topological consistency is not a prerequisite to a good geometric approximation. For closed surfaces, the topology is completely characterized by the genus: the number of handles (or holes) in the surface. From a topological perspective, all holes have equal importance, but from the point of view of geometry this is simply not true. For many applications it may not be important to preserve tiny handles. The notion of *topological persistence* was developed [ELZ00] as a way of capturing the relative size and importance of handles. Of course if topological equivalence is sacrificed, no homeomorphism between the original surface and its approximation is possible, so subsequent geometric accuracy analysis becomes more problematic.

4.2 Convergence and geometric accuracy

Assuming the criterion of topological consistency has been met, we turn our attention to the geometric fidelity of M as a representation of S. In other words we are interested in evaluating the pointwise positional accuracy of M and the accuracy of the face normals of the triangles of M as representatives of corresponding surface normals on S.

4.2.1 Pointwise convergence and the projection mapping

Even if the set P of vertices of M all lie on S itself, there may be points of M that are far from S if the density of P is not high enough with respect to the local feature size. If $d_{\mathbb{R}^3}(x, S)$ is the (minimum) distance from $x \in M$ to S, then one natural measure of the pointwise accuracy of M is $\sup_{x \in M} d_{\mathbb{R}^3}(x, S)$. This is called the *Hausdorff distance* from M

[10] Image from http://commons.wikimedia.org/wiki/Image:Trefoil_knot_arb.png

to S. Notice that this distance is not symmetric: The Hausdorff distance from S to M will be different in general. The *symmetric Hausdorff distance* between M and S is the maximum of these two numbers. If we have a family of meshes M_n whose vertices, P are samples of S with $|P| = n$, then if the symmetric Hausdorff distance between M_n and S goes to zero as n goes to infinity, we say the family $\{M_n\}$ displays *pointwise convergence*.

If M is pointwise sufficiently close to S, then there is a natural mapping $M \to S$ that embodies the concept of "distance to the closest point" implied in the above discussion. If $m \in \mathbb{R}^3$ is not a point on the medial axis of S, then there is a unique point $p \in S$ such that $d_{\mathbb{R}^3}(m, p) = d_{\mathbb{R}^3}(m, S)$. The point m must lie on the line generated by the normal to S at p, so p can be viewed as the orthogonal projection of m onto S. For a set $U \subset \mathbb{R}^3$, the projection mapping $\xi : U \to S$, takes each m in U to its closest point in S. If $\xi|_M$ is injective, then it defines a homeomorphism of M onto S. Standard techniques in geometric accuracy analysis exploit this homeomorphism.

For any smooth surface S, we can construct a neighbourhood $U \supset S$ upon which the projection ξ is well defined. For a point $p \in S$, the *local reach* at p is the radius of the smallest of the two medial balls at p. The local reach is a continuous function $\rho_R : S \to \mathbb{R}^+$ [Fed59]. For each $p \in S$, let I_p denote the open interval on the normal line through p, centred at p and with length $2\rho_R(p)$. By construction, any point $m \in I_p$ is closer to p than to any other point in S, thus for $p \neq q$, $I_p \cap I_q = \emptyset$. The set $U_{\rho_R} = \cup_{p \in S} I_p$ is called a *tubular neighbourhood* of S, and the projection ξ onto S is well defined on U_{ρ_R}. In fact, ξ is well defined on any neighbourhood of S that does not contain any point of the medial axis.

4.2.2 Normal convergence

Establishing pointwise convergence is not sufficient to guarantee that M_n will be a good representation of S for sufficiently high n. It is crucial that one also demonstrates normal convergence. To evaluate how effectively M approximates the normals of S one needs a means of corresponding points on M with points on S. Typically the projection mapping is used, so we are interested in the size of $\|\hat{\boldsymbol{n}}(x) - \hat{\boldsymbol{n}}(\xi(x))\|$. If the supremum of this number over all $x \in t$ for all triangles t in M_n goes to zero as n goes to infinity, we say that $\{M_n\}$ displays *normal convergence*.

Note that this criterion implies that the difference in the normals of two adjacent triangles must go to zero as n goes to infinity. The normal vector at a vertex in a triangle mesh does not need to be defined. If $p \in M_n$ for all n, then the normals to triangles that have p as a vertex must converge to $\hat{\boldsymbol{n}}(p)$, the normal to S at p. In computer graphics, defining a normal vector at the vertex of a mesh is important. Most schemes for defining a normal at a vertex involve a weighted sum of the face normals of the adjacent triangles. Recent work indicates that weighting the face normals by the face angle at the vertex yields convenient properties [BA05]. In the emerging formalism of the discrete exterior calculus, the formulation of a consistent scheme for defining tangent planes at vertices is considered an open problem [DHLM03]. Any sane definition of a normal vector at the vertices will converge to the normal of the surface at that point if the meshes display normal convergence as we have defined it.

It has been shown [MT04], [HPW06] that normal convergence is equivalent to convergence of surface area and to the convergence of the intrinsic distance functions of the M_n to that

of S.

4.2.3 Convergence properties of Delaunay structures

One of the attractive properties of both the rDt and the iDt-mesh is that these families both display pointwise and normal convergence. This was demonstrated for the former in [AB98], and more recently for the latter in [DLYG06, DLJ+07]. It should be mentioned that the bounds for the normal error that appeared in [AB98] as well as in many subsequent works, e.g., [ACDL00, Dey07], rested on a "normal variation lemma" which bounds the difference in the normals between two points whose separation is bounded by the lfs. A recent erratum, [AD], has corrected the proof of that lemma and improved its bound. Thus the normal error bounds mentioned in affected previous works can be improved.

For general meshes, increasing sampling density alone is not sufficient to ensure a decrease in the normal error, as the Schwarz lantern will attest. However, if the circumradius of the triangles can be controlled, then so too can be the normal error. For general meshes, the normal error depends linearly on the largest circumradius of the triangles [MT04]. This highlights one of the reasons for the success of the Delaunay paradigm in surface meshing: Delaunay triangulations favour triangles with small circumradius. Actually, that statement merits further exploration. While it is known that the 2D planar Delaunay triangulation minimizes the maximum circumradius of the triangles [DS89], the corresponding claims for the structures presented in Section 3 have not been established. Recent work by Cheng and Dey [CD07] shows that, with mild constraints on the dihedral angles and some uniformity and density constraints on the sampling, a mesh that lexicographically minimizes the circumradii of the triangles will be a Gabriel mesh. In particular, a physical edge flip of an edge that is NLD with respect to the intrinsic metric of the mesh has the effect of reducing the maximum circumradius of the triangles involved. Thus some sort of circumradii optimality result has been obtained for Delaunay meshes at least. We expect that the uniformity constraints on the sampling can be dropped.

That all three Delaunay structures constrain the triangle circumradius with respect to the sampling density can be seen intuitively if we take the approach of considering the dual Voronoi diagram on S that results from the appropriate choice of metric d. Suppose a sampling density is specified in terms of a *sampling radius*: a function $\rho(x) : S \to \mathbb{R}^+$ such that for any point $x \in S$ there must be a $p \in P$ with $d(x, p) < \rho(x)$. If the sampling radius is sufficiently small, the point z on S that is farthest from any sample point will be a Voronoi vertex. If $d = d_M$, the metric of a conforming Delaunay mesh, then the distance of z to the closest sample will be exactly the circumradius of the triangle face t dual to z. It follows that t will have the largest circumradius of all triangles in M and that this circumradius is necessarily smaller than $\rho(z)$. Of coures, no sampling density would be specified in terms of the metric of a mesh that has yet to be built, but the analysis will hold as long as the metric of the mesh can be bounded from above in terms of the metric used to define the sampling density.

Likewise, for the rDt and the iDt-mesh, the above argument will not hold exactly, but the distance from z to the closest sample will differ from the circumradius of t by only a small amount, corresponding to the distortion of the metric as we go from S to M. In either

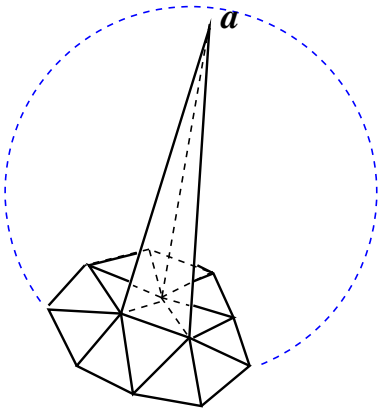

Figure 7: Without some kind of regularity condition, a Delaunay mesh is not necessarily a good representation of the surface from which its vertices are sampled. The mesh fragment shown here is necessarily Delaunay since it has no obtuse face angles. However if the samples represent a well sampled smooth surface, the vertex a should not be associated with the neighbours that it has. For example, a could be a point on the antipodal side of a well sampled sphere.

case, we are able to bound the circumradius of the triangle with respect to the sampling radius.

It is worth being more explicit on this point. Convenient bounds on the normal error and the positional error for a Euclidian triangle t whose vertices belong to S are given in [DLYG06][Theorem 2]. Specifically, if the vertices of t are contained in $B_S(c;r)$, a geodesic disk of radius r and such that the maximal curvature is bounded by κ in $B_S(c;r)$ and $r < 1/(4\kappa)$, then for any point $m \in t$, we have $\|\hat{\boldsymbol{n}}(\xi(m)) - \hat{\boldsymbol{n}}_t\| \leq 4.5\kappa r$ and also $d_{\mathbb{R}^3}(m, \xi(m)) \leq 9\kappa r^2$. Now if we have a sampling radius that demands $r < \epsilon\rho_f(c)$, for a sufficiently small ϵ, then using the Lipschitz continuity of the lfs the normal error is bounded by $\mathcal{O}(\epsilon)$ and the positional error is bounded by $\mathcal{O}(\rho_f(c)\epsilon^2)$. It is possible to obtain the same bounds if we are using the Euclidean metric instead of the intrinsic metric [DZM08][Lemma 9]. Thus these bounds hold for both the rDt and the iDt-mesh. We expect that it will also be possible to establish a corresponding result for a class of conforming Delaunay meshes. The principle difficulty will be characterizing, and establishing existence of such Delaunay meshes.

The specification of a subclass of Delaunay meshes is important here. Indeed, without the conforming qualifier, the definition of a Delaunay mesh is too broad to guarantee a useful surface representation. Even when the vertices of M densely sample S, M may be Delaunay but not admit a conforming homeomorphism. A simple example can be constructed from a densely sampled sphere. Construct a Delaunay mesh, but choose a sample point at the north pole to be connected only to neighbours in the neighbourhood of the south pole. Thus M resembles a sphere but with a pyramidal spike penetrating its axis. Yet M may still be a Delaunay mesh (see Figure 7).

4.3 Element quality

There is another quality of M that is often sought in addition to, or even in preference to geometric fidelity. We often want to ensure that M contains no "bad triangles". This is the goal of good element quality. What exactly is meant by a bad triangle is the subject of a sizable paper by Shewchuk [She02]. There are many measures for judging the quality of a triangle, perhaps the best known is the ratio of the shortest edge to the circumradius, a number that is generally desired to be as large as possible. Good triangle quality is important for applications which discretize partial differential equations on S [She02]. Computing eigenfunctions of the Laplace-Beltrami operator, a computation that is drawing increasing interest in the geometry processing community [ZvKD08], can be seen to fall into this category.

Although there may be no clear consensus of exactly what defines a bad triangle, it is generally agreed that a triangle that is near to being equilateral is good. In general the Delaunay structures provide good element quality, but for a fixed sample budget, the best geometric representations of S are composed of triangles that are far from equilateral. In the presence of anisotropic curvature, such as on the surface of a cylinder for example, long thin triangles oriented in the direction of low curvature provide a better geometric approximation than may be obtained with triangles that are closer to being equilateral.

A nice demonstration of this phenomenon is presented in works that seek to optimize triangle meshes by edge flipping algorithms geared towards minimizing some discrete curvature measure [DHKL01] [vDA95]. Although these works appeal primarily to the visual quality of the results and lack quantitative analysis, they are supported by more quantitative works: In [DLR90] the best planar mesh to represent a known, possibly anisotropic, function $f : D \subset \mathbb{R}^2 \to \mathbb{R}$ is sought, and in [AKTvD00] experiments are performed to assess the geometric quality of the meshes produced by optimizing the quality measures discussed in [vDA95]. The meshes produced are visually appealing, but they contain many long skinny triangles. Interestingly, edge flipping algorithms that minimize the *Willmore energy*, which is, roughly speaking, the integral of the square of the mean curvature, tend to produce nice triangle quality at the expense of mesh smoothness [ABR06].

Thus it appears that the goal of geometric fidelity is at odds with the goal of good element quality. In particular, the Delaunay structures discussed in Section 3, despite the established convergence results for the rDt and the iDt-mesh, are not generally the best meshes for representing the geometry of S. This fact coupled with the demonstrated practical and theoretical attributes of Delaunay structures should motivate future work into Delaunay structures based on anisotropic metrics which reflect the principle curvatures.

Returning again to the question of what makes a bad element, we consider the issue from the point of view of geometric accuracy. There is evidence which supports the assertion that triangles with very large angles may be more problematic than triangles with very small angles. In the context of normal convergence, Morvan and Thibert [MT04] identified the *rightness* of a triangle t as an important measure of element quality. The rightness of t is simply the sine of the largest angle. The importance of rightness arises in the context of normal convergence of meshes approximating a smooth surface: the normal error is minimized when the minimum rightness of the triangles is maximized. This result translates into the linear dependence of the normal error on the circumradius by the formula $\rho_t = \frac{\eta_t}{2\gamma_t}$, where for

triangle t, ρ_t is its circumradius, η_t is the length of its longest edge and γ_t is its rightness.

5 Surface meshing and reconstruction

Most results pertaining to surface representation by Delaunay structures have arisen in the context of surface meshing and surface reconstruction. These are two related application domains that both depend on surface sampling theory and the geometric accuracy of triangle meshes.

In surface meshing, the input is a surface S and one must produce a set, P, of samples of S, and a mesh, M, whose vertices are P, such that M meets given geometric accuracy requirements.

In surface reconstruction, the input is the set of samples P and, aside from some regularity assumptions (i.e. that it was a smooth surface, or at least a Lipschitz surface), the surface S is unknown. Again one wishes to construct a mesh that adequately represents S.

The work in these domains is extensive, and there appeared recently a good survey on the subject of Delaunay structures in surface reconstruction [CG06]. Also, a recent book on surface reconstruction [Dey07], gives details and analysis on state of the art surface reconstruction algorithms. The Delaunay paradigm is prominent in most of the algorithms discussed there.

In this section we give an overview of the main results with an emphasis on how they relate to Delaunay structures in general and the Delaunay meshes described in Section 3.2.3 in particular.

5.1 Reconstruction algorithms

5.1.1 Quality guarantees

Many surface reconstruction algorithms provide guarantees on the quality of the output surface if specific sampling density assumptions are met. The first such algorithm was the Crust algorithm developed by Amenta and Bern [AB98]. This was a seminal paper in which several important concepts were introduced including the local feature size and the *pole* of a sample point. The pole of a sample p is the Voronoi vertex of $V(p)$ that is farthest from p. The pole enables an estimate of the normal vector at p. The mesh M was extracted from the Delaunay tetrahedralization of the poles and the original samples together; only triangles whose vertices are all samples are considered while constructing M. The algorithm and the analysis was subsequently simplified with the Co-cone algorithm [ACDL00]. A later development was the Power Crust [ACK01] which creates a piecewise linear surface via the consideration of medial balls centred at the poles.

The 3D Delaunay tetrahedralization and the rDt played prominently in the design and analysis of these algorithms, and it is fair to say that the scaffold provided by the Delaunay tetrahedralization has played a role in most reconstruction and meshing algorithms that have provable guarantees on their output quality.

5.1.2 Sample density considerations

Any surface reconstruction algorithm which guarantees a topologically correct output, must make some assumptions about the density of the samples that are used as input. Typically, the sample set is assumed to be a lfs ϵ-sample set, for some $\epsilon \approx 0.1$. An important observation made by Dey and Giesen [DG01] is that regions of undersampling in the input sample points can be detected without direct knowledge of the initial surface S. Their work was based on an analysis of the shape of the Voronoi cells in the 3D Voronoi diagram. Petitjean and Boyer [PB01] set out with a similar agenda to characterize a good sample set via the properties of certain interpolating meshes. However, as noted in Section 3, their scheme turned out to be too restrictive because it did not allow sample sets which contain sliver tetrahedra in their Delaunay tetrahedralization.

5.1.3 Surface area minimization

Among the early publications dealing with the surface reconstruction problem was a short paper by O'Rourke [O'R81] in which local edge swapping is performed to reduce the surface area of an initial mesh. The assumption was that the mesh of minimal surface area connecting the vertices P would be one of the best representations of S amongst all the possible ways of connecting P. It was mentioned that attaining the true minimal surface area mesh would probably not be computationally practical, since the lower dimensional analogous problem of finding a minimal length polyhedron connecting a set of points in the plane is essentially the traveling salesman problem, which is known to be NP complete.

However, the traveling salesman problem is NP-complete for a generic set of points, but points which are densely sampled from some smooth curve have properties which reflect their special origin. In a series of papers initiated by Giesen [Gie99a, Gie99b, AM01] it was demonstrated that the traveling salesman perspective yields a powerful algorithm for producing a polygonal reconstruction of the original planar curve. These algorithms are able to perform reliably on much weaker regularity assumptions than are required by Voronoi-based curve reconstruction algorithms.

Thus it may be interesting to revisit the idea of polyhedra of minimal surface area in the context of surface reconstruction. A further motivation for such an investigation is the strong ties to the Delaunay mesh as indicated by the observation [DZM07a] that performing a physical edge swap on an NLD edge reduces the surface area of the mesh.

5.1.4 The flow complex and wrap algorithms

There have been works which present algorithms which are metaphorically referred to as shrink wrapping algorithms, and these give the impression of being related to surface area minimization. Among these were later works by Giesen et al. [GJ02, GJ03] which introduced the *flow complex*. The flow complex is a piecewise linear structure whose vertices are the critical points of the distance function as measured from the sample point set, as well as the sample points themselves. The precise description of the flow structure is quite involved. It is neither a substructure of the 3D Delaunay tetrahedralization nor of the Voronoi diagram, but it can be described completely in terms of them both. The wrap algorithm that was

developed independently by Edelsbrunner [Ede04] is closely related to the flow complex, but it employs a power distance from Voronoi vertices. It can be seen as the dual of the flow complex, but the critical points are the same [RS07]. In contrast to the mesh constructed by the flow complex, the output of the wrap algorithm is a subcomplex of the 3D Delaunay tetrahedralization. Another notable shrink wrapping algorithm with this property was presented by Chaine [Cha03]. The algorithm is described as a convection algorithm and is perhaps the most worthy of the shrink wrapping classification.

Although the shrink wrapping algorithms give the impression of surface area minimization, just as their metaphorical moniker implies, no explicit claims of surface area minimization are made. These algorithms have not been analyzed from this perspective. A reconstruction algorithm that does produce a polyhedron of minimal surface area was introduced by Althaus and Fink [AF02]. However the algorithm applies only to data sampled on parallel *planar contours*. A contour on S is defined as the curve of intersection between S and a plane in $\mathbb{R}^3$. The contours parallel if their corresponding planes are parallel.

5.1.5 Local region growing

Another class of surface reconstruction algorithms is based on the idea of local region growing. The Delaunay paradigm has appeared here in work by Gopi et al. [GKS00], where samples in the vicinity of interest are projected onto an estimated tangent plane, and the 2D Delaunay triangulation of the result guides the connectivity of the mesh under construction. A similar approach was employed by Chen et al. [CB97] who also do a local parametrization but to establish the connectivity they use ellipses reflecting the parametric distortion of the Delaunay circles, thus approximating an iDt-mesh more than an rDt. There are many other algorithms that are based on local region growing. One that was recently presented [KA08] is worth noting because to form an umbrella of triangles around each point, triangles with the smallest circumradius are selected.

Some region growing algorithms do come with quality guarantees [LT06, KA08], however since these algorithms tend to work with k-nearest neighbours around a point, they typically require a uniformity constraint on the sample set. For example, P is a lfs (ϵ, δ)-sample set if it is a lfs ϵ-sample set and for any $p, q \in P$, $d_{\mathbb{R}^3}(p, q) \geq \delta \rho_f(p)$ for $\frac{\epsilon}{2} \leq \delta < \epsilon < 1$ [LT06].

5.1.6 Witness complex

A new approach to the rDt has recently emerged in the form of the *witness complex*. This construction provides for the identification of the rDt of $P \subset S$ when S itself is known only in the form of a much denser set of samples $\tilde{P}$, of which P is a subset. Elements of the set P, the vertices of the rDt, are referred to as *landmarks*. A recent work by Attali et al. [AEM07] paved the way for the use of witness complexes in surface reconstruction.

5.1.7 Alpha, Beta and Gabriel complexes

The relevance of the Gabriel complex has been recognized in surface reconstruction [AGJ00], [PB01], [Cha03]. In general the Gabriel complex is expected to contain a good surface rep-

resentation for a sufficiently dense sample set, however this has not been formally demonstrated, except for the case of a uniform sample distribution [CD07].

A related complex is the *α-shape* [EM94]. This consists of those triangles that admit an empty ball of radius α on their dual Voronoi edge. Surface reconstruction algorithms which are based on the α-shape generally need to assume a uniform sampling density. The *β-skeleton* [KR85] was exploited for curve reconstruction [ABE98], but it did not easily generalize for the purposes of surface reconstruction [AB98]. The β-skeleton consists of those edges in the planar Delaunay triangulation which possess empty disks that intersect at the endpoints of the edge and with radius of β times the length of the edge.

5.1.8 Implicit surface construction

Many surface reconstruction algorithms compute a scalar field, i.e., a function $h : \mathbb{R}^3 \to \mathbb{R}$ such that the surface is approximated by a level set. This is called an *implicit surface* representation. This has the advantage that an analytic representation of the surface may be obtained. Algorithms designed to handle noisy points sets tend to be of this type.

One such algorithm that exploited the Voronoi diagram was presented by Boissonnat and Cazals [BC00, BC01]. They constructed a signed distance function with the aid of the natural neighbour interpolation discussed in Section 2.

5.2 Surface meshing and remeshing

Related to surface reconstruction is the problem of meshing a surface. In this case the surface S is known and a mesh is constructed by sampling S. Commonly the surface to be meshed is an implicit surface such as is produced by some reconstruction algorithms.

5.2.1 Delaunay refinement

The seminal meshing paper was Chew's Delaunay refinement algorithm [Che93]. This is the same Delaunay refinement paradigm that is used for planar meshes, but it is the restricted Delaunay triangulation of S that is being refined. Virtually all subsequent meshing papers admit to being a refinement of Chew's algorithms. Dey et al. have developed variations for remeshing existing meshes [DLR05] as well as for meshing isosurfaces [DL07]. For smooth surfaces there are algorithms which guarantee a topologically correct and geometrically faithful output [BO03, BO05, CDRR04], and similar guarantees now also apply to Lipschitz surfaces [BO06] which have a much weaker smoothness constraint.

5.2.2 Use of the intrinsic Voronoi diagram

Although there are no surface reconstruction algorithms based on the theory of the intrinsic Voronoi diagram, the farthest point sampling algorithm published by Peyré and Cohen [PC03, PC06] and also by Moenning and Dodgson [MD03] present a remeshing method based upon the iVd. This algorithm computes an approximation of the iVd, and the resulting mesh is its dual. In other words this algorithm produces an approximation to the iDt-mesh.

5.3 Alternatives to the Delaunay paradigm

Our focus is on Delaunay structures, however there are many approaches to surface reconstruction which do not employ Delaunay methods. These works come primarily from the computer graphics community and they usually yeild an implicit surface representation. A seminal work in this vein was produced by Hoppe [HDD+92], where a signed distance function was constructed using estimates of the tangent planes at the data points. An analytic representation of a scalar field whose zero level set is an approximation to the surface can be obtained using radial basis functions [CBC+01].

A popular method for defining a surface from a point cloud is the via the moving least squares projection operator [ABCO+01] which defines the surface locally according to the nearby samples. The resulting surface is refered to as an MLS surface. It has been shown that the MLS surface can be represented an implicit surface [AK04].

Other surface reconstruction methods rely on numerically solving partial differential equations. Works in this category include the work of Kazhdan et al. [KBH06], where a scalar field is constructed by solving a Poisson equation. The input to this algorithm is a set of *oriented point samples*: samples that have a surface normal vector associated with them. Another PDE based method evolves a scalar field according to an energy functional that drives its zero level set towards the sample points [ZOF01]. This later technique belongs to a general framework called *level set methods*.

These give a taste for the range of different surface reconstruction methods that have appeared. We emphasise that we have mentioned only a tiny fraction of the work that has been published. The primary purpose of this report is to survey those techniques which can or have supplied insight into the theory of Delaunay structures for surface representation.

In the theory of surface meshing, there has also been much work that isn't Delaunay-based. In particular there is a body of abstract theoretical work from the mathematics community that has yet to be exploited in geometry processing. Surface triangulation has been approached from the point of view of metric spaces by Saucan [Sau04] and also Clarkson [Cla06]. Rivin [Riv90], focused on hyperbolic spaces, and Chavel explores triangulations in abstract Riemannian manifolds [Cha06][§IV.4]. These works surely contain deep contributions that can be exploited in geometry processing if the results can be recognized and coaxed from their abstract environment. Rivin's work in particular appears to have possible relevance to the theory of Delaunay meshes.

6 Discussion

Of the three principle Delaunay structures we have discussed, the rDt, the iDt-mesh and the Delaunay mesh, the rDt is by far the most established and best understood. In practice it is much easier to work with Euclidean distances rather than geodesic distances on some smooth surface. By definition the rDt is a substructure of the 3D Delaunay tetrahedralization. For surface reconstruction algorithms in particular, this is an often exploited property. For a sufficiently well sampled surface the iDt-mesh and nice Delaunay meshes are also expected to possess this property, but this has yet to be formally demonstrated.

In many ways analysis of such things as approximation error is more naturally performed

when geodesic distances are considered. Recently there has been increased interest in the intrinsic viewpoint [DLYG06], [GGOW08], [DZM08]. The intrinsic viewpoint is of greater interest in those domains, such as manifold learning, where working in the high dimensional ambient space is impractical, but even in geometry processing for graphics applications, the strength of this viewpoint is beginning to be recognized.

The study of Delaunay meshes is in its infancy. One of the appealing characteristics of these meshes is that, in contrast to the rDt and the iDt-mesh, they do not rely on a reference surface S for their definition. When given a mesh M, one can determine with a quick check of the angles whether or not it is a Delaunay mesh. In contrast, there is no way to tell if it is a rDt or an iDt-mesh, unless the reference surface S is also supplied.

However, in practice M is supposed to represent some reference surface. The definition of a Delaunay mesh is so loose that even if its vertices form a dense sampling P of S, there is no guarantee that the Delaunay mesh M interpolating them is a good representation of S. We need some way of ensuring that M will admit a conforming homeomorphism onto a surface S for which P forms a sufficiently dense sample set.

Assuming that P does represent a sufficiently dense sample set, it is expected that some constraints on the dihedral angles and edge lengths will ensure a sane Delaunay mesh. If such criteria could be established it would provide a means of certifying the correctness of the output of a reconstruction algorithm. Ideally a simple verification of local properties of the mesh will ensure that it is a good representation of a smooth surface for which its vertices form a sufficiently dense sample set.

References

[AB98] Nina Amenta and Marshall W. Bern. Surface reconstruction by Voronoi filtering. In *Symp. Comp. Geom.*, pages 39–48, 1998.

[ABCO+01] Marc Alexa, Johannes Behr, Daniel Cohen-Or, Shachar Fleishman, David Levin, and Claudio T. Silva. Point set surfaces. In *IEEE Visualization*, 2001.

[ABE98] Nina Amenta, Marshall Bern, and David Eppstein. The crust and the beta-skeleton: combinatorial curve reconstruction. *Graphical Models and Image Processing*, 60(2):125–135, 1998.

[ABR06] Lyuba Alboul, Willie Brink, and Marcos Rodrigues. Mesh optimisation based on Willmore energy. In *Euro. Workshop on Comp. Geom.*, pages 133–136, 2006.

[ACDL00] Nina Amenta, Sunghee Choi, Tamal K. Dey, and Naveen Leekha. A Simple Algorithm for Homeomorphic Surface Reconstruction. In *Symp. Comp. Geom.*, pages 213–222, 2000.

[ACK01] Nina Amenta, Sunghee Choi, and Ravi Krishna Kolluri. The power crust. In *Sixth ACM Symposium on Solid Modeling and Applications*, pages 249–260, 2001.

[AD] N. Amenta and T. Dey. Normal variation for adaptive feature size. Erratum.

[AEM07] Dominique Attali, Herbert Edelsbrunner, and Yuriy Mileyko. Weak witnesses for Delaunay triangulations of submanifolds. In *Symp. Solid and Physical Modeling*, pages 143–150, 2007.

[AF02] Ernst Althaus and Christian Fink. A polyhedral approach to surface reconstruction from planar contours. In *Integer Programming and Combinatorial Optimization*, pages 258–272, 2002.

[AGJ00] Udo Adamy, Joachim Giesen, and Matthias John. New techniques for topologically correct surface reconstruction. In *IEEE Visualization*, 2000.

[AK00] Franz Aurenhammer and Rolf Klein. Voronoi diagrams. In J. Sack and G. Urrutia, editors, *Handbook of Computational Geometry*, pages 201–290. Elsevier Science Publishing, 2000.

[AK04] Nina Amenta and Yong Joo Kil. Defining point-set surfaces. *ACM Trans. Graph.*, 23(3):264–270, 2004.

[AKTvD00] Lyuba Alboul, Gertjan Kloosterman, Cornelis Traas, and Ruud van Damme. Best data-dependent triangulations. *J. Comput. Appl. Math.*, 119(1-2):1–12, 2000.

[AM01] Ernst Althaus and Kurt Mehlhorn. Traveling salesman-based curve reconstruction in polynomial time. *SIAM J. Comput.*, 31(1):27–66, 2001.

[APR03] Nina Amenta, Thomas J. Peters, and Er Russell. Computational topology: ambient isotopic approximation of 2-manifolds. *Theoretical Computer Science*, 305:3–15, 2003.

[AZ67] A. D. Aleksandrov and V. A. Zalgaller. *Intrinsic Geometry of Surfaces*, volume 15 of *Transactions of Mathematical Monographs*. AMS, 1967.

[BA05] J. Andreas Baerentzen and Henrik Aanaes. Signed distance computation using the angle weighted pseudonormal. *IEEE Transactions on Visualization and Computer Graphics*, 11(3):243–253, 2005.

[BC00] Jean-Daniel Boissonnat and Frédéric Cazals. Smooth surface reconstruction via natural neighbour interpolation of distance functions. In *Symp. Comp. Geom.*, pages 223–232, 2000.

[BC01] Jean-Daniel Boissonnat and Frédéric Cazals. Natural neighbour coordinates of points on a surface. *Computational Geometry Theory and Applications*, 19(2):155–173, 2001.

[BO03] Jean-Daniel Boissonnat and Steve Oudot. Provably good surface sampling and approximation. In *Symp. Geometry Processing*, pages 9–18, 2003.

[BO05] Jean-Daniel Boissonnat and Steve Oudot. Provably good sampling and meshing of surfaces. *Graphical Models*, 67(5):405–451, 2005.

[BO06] Jean-Daniel Boissonnat and Steve Oudot. Provably good sampling and meshing of lipschitz surfaces. In *Symp. Comp. Geom.*, pages 337–346, 2006.

[BS05] Alexander I. Bobenko and Boris A. Springborn. A discrete Laplace-Beltrami operator for simplicial surfaces. arXiv:math.DG/0503219 v1, 2005.

[BS06] Alexander I. Bobenko and Boris A. Springborn. A discrete Laplace-Beltrami operator for simplicial surfaces. arXiv:math.DG/0503219 v3, 2006.

[BS07] Alexander I. Bobenko and Boris A. Springborn. A discrete Laplace-Beltrami operator for simplicial surfaces. *Discrete and Computational Geometry*, 38(4):740–756, 2007.

[CB97] H. Chen and J. Bishop. Delaunay triangulation for curved surfaces. In *Meshing Roundtable*, pages 115–127, 1997.

[CBC+01] J. C. Carr, R. K. Beatson, J. B. Cherrie, T. J. Mitchell, W. R. Fright, B. C. McCallum, and T. R. Evans. Reconstruction and representation of 3d objects with radial basis functions. In *ACM SIGGRAPH*, 2001.

[CD07] Siu-Wing Cheng and Tamal K. Dey. Delaunay edge flips in dense surface triangulations, 2007. arXiv.org:0712.1959v1.

[CDRR04] Siu-Wing Cheng, Tamal K. Dey, Edgar A. Ramos, and Tathagata Ray. Sampling and meshing a surface with guaranteed topology and geometry. In *Symp. Comp. Geom.*, pages 280–289, 2004.

[CG06] Frédéric Cazals and Joachim Giesen. Delaunay triangulation based surface reconstruction. In Jean-Daniel Boissonnat and Monique Teillaud, editors, *Effective Computational Geometry for Curves and Surfaces*. Springer, 2006.

[Cha03] Raphaëlle Chaine. A geometric convection approach of 3-D reconstruction. In *Symp. Geometry Processing*, pages 218–229, 2003.

[Cha06] Isaac Chavel. *Riemannian Geometry, A modern introduction*. Cambridge, 2nd edition, 2006.

[Che89] L. Paul Chew. Constrained Delaunay triangulations. *Algorithmica*, 4(1):97–108, 1989.

[Che93] L. Paul Chew. Guaranteed-quality mesh generation for curved surfaces. In *Symp. Comp. Geom.*, pages 274–280, 1993.

[Cla06] Kenneth L. Clarkson. Building triangulations using epsilon-nets. In *Proceedings of the 38th ACM Symposium on Theory of Computing (STOC'06)*. ACM Press, 2006.

[dBvKOS98] Mark de Berg, Marc van Kreveld, Mark Overmars, and Otfried Schwarzkopf. *Computational Geometry. Algorithms and Applications*. Springer-Verlag, 1998.

[Del34] B. Delaunay. Sur la sphère vide. *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvennyka Nauk*, 7:793–800, 1934.

[Dey07] Tamal Dey. *Curve and Surface Reconstruction; Algorithms with Mathematical Analysis*. Cambridge University Press, 2007.

[DG01] Tamal K. Dey and Joachim Giesen. Detecting undersampling in surface reconstruction. In *Symp. Comp. Geom.*, pages 257–263, 2001.

[DHKL01] Nira Dyn, Kai Hormann, Sun-Jeong Kim, and David Levin. Optimizing 3D triangulations using discrete curvature analysis. In *Mathematical Methods for Curves and Surfaces*, pages 135–146. Vanderbilt University, 2001.

[DHLM03] M. Desbrun, A. N. Hirani, M. Leok, and J. E. Marsden. Discrete exterior calculus. (preprint, arXiv:math.DG/0508341), 2003.

[DL07] Tamal K. Dey and Joshua A. Levine. Delaunay meshing of isosurfaces. In *Shape Modelling International*, pages 241–250, 2007.

[DLJ+07] Junfei Dai, Wei Luo, Miao Jin, Wei Zeng, Ying He, Shing-Tung Yau, and Xianfeng Gu. Geometric accuracy analysis for discrete surface approximation. *Comput. Aided Geom. Des.*, 24(6):323–338, 2007.

[DLR90] Nira Dyn, David Levin, and Samuel Rippa. Data dependent triangulations for piecewise linear interpolation. *IMA Journal of Numerical Analysis*, 10:137–154, 1990.

[DLR05] Tamal K. Dey, Gang Li, and Tathagata Ray. Polygonal surface remeshing with Delaunay refinement. In *Meshing Roundtable*, pages 343–361, 2005.

[DLYG06] Junfei Dai, Wei Luo, Shing-Tung Yau, and Xianfeng Gu. Geometric accuracy analysis for discrete surface approximation. In *Geometric Modeling and Processing*, pages 59–72, 2006.

[DS89] E. F. D'Azevedo and R. B. Simpson. On optimal interpolation triangle incidences. *SIAM J. Sci. Statist. Comput.*, 10(6):1063–1075, 1989.

[DZM07a] R. Dyer, H. Zhang, and T. Möller. Delaunay mesh construction. In *Symp. Geometry Processing*, pages 271–282, 2007.

[DZM07b] R. Dyer, H. Zhang, and T. Möller. Voronoi-Delaunay duality and Delaunay meshes. In *Symp. Solid and Physical Modeling*, pages 415–420, 2007.

[DZM08] R. Dyer, H. Zhang, and T. Möller. Surface sampling and the intrinsic Voronoi diagram. *Computer Graphics Forum (Special Issue of Symposium on Geometry Processing 2008)*, 27(5):1393–1402, 2008.

[Ede01] H. Edelsbrunner. *Geometry and Topology for Mesh Generation*. Cambridge, 2001.

[Ede04] H. Edelsbrunner. Surface reconstruction by wrapping finite point sets in space. *Discrete and Computational Geometry*, 32:231–244, 2004.

[ELZ00] Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. Topological persistence and simplification. In *IEEE Symposium on Foundations of Computer Science*, pages 454–463, 2000.

[EM94] Herbert Edelsbrunner and Ernst P. Mücke. Three-dimensional alpha shapes. *ACM Trans. Graph.*, 13(1):43–72, 1994.

[ES92] H. Edelsbrunner and N. R. Shah. Incremental topological flipping works for regular triangulations. In *Symp. Comp. Geom.*, pages 43–52, 1992.

[ES94] Herbert Edelsbrunner and Nimish R. Shah. Triangulating topological spaces. In *Symp. Comp. Geom.*, pages 285–292, 1994.

[ET93] Herbert Edelsbrunner and Tiow Seng Tan. An upper bound for conforming Delaunay triangulations. *Discrete and Computational Geometry*, 10(2):197–213, 1993.

[Fed59] Herbert Federer. Curvature measures. *Trans. Amer. Math. Soc.*, 93:418–491, 1959.

[FSBS06] Matthew Fisher, Boris Springborn, Alexander I. Bobenko, and Peter Schröder. An algorithm for the construction of intrinsic Delaunay triangulations with applications to digital geometry processing. In *SIGGRAPH Courses*, pages 69–74, 2006.

[GGOW08] Jie Gao, Leonidas J. Guibas, Steve Y. Oudot, and Yue Wang. Geodesic Delaunay triangulation and witness complex in the plane. In *Symp. on Discrete Algorithms*, pages 571–580, 2008.

[Gie99a] Joachim Giesen. Curve reconstruction in arbitrary dimension and the traveling salesman problem. *Lecture Notes in Computer Science*, 1568:164–176, 1999.

[Gie99b] Joachim Giesen. Curve reconstruction, the traveling salesman problem and Menger's theorem on length. In *Symp. Comp. Geom.*, pages 207–216, 1999.

[GJ02] Joachim Giesen and Matthias John. Surface reconstruction based on a dynamical system. *Computer graphics forum*, 21:363–371, 2002.

[GJ03] Joachim Giesen and Matthias John. The flow complex: a data structure for geometric modeling. In *Symp. on Discrete Algorithms*, pages 285–294, 2003.

[GKS00] M. Gopi, S. Krishnan, and C. T. Silva. Surface reconstruction based on lower dimensional localized Delaunay triangulation. In M. Gross and F. R. A. Hopgood, editors, *Computer Graphics Forum (Eurographics)*, volume 19(3), 2000.

[Gli05] David Glickenstein. Geometric triangulations and discrete Laplacians on manifolds, 2005. arxiv:math.MG/0508188v1.

[GR04] Leonidas Guibas and Daniel Russel. An empirical comparison of techniques for updating Delaunay triangulations. In *Symp. Comp. Geom.*, pages 170–179, 2004.

[GS69] K. Ruben Gabriel and Robert R. Sokal. A new statistical approach to geographic variation analysis. *Systematic Zoology*, 18(3):259–278, 1969.

[HDD+92] Hugues Hoppe, Tony DeRose, Tom Duchamp, John McDonald, and Werner Stuetzle. Surface reconstruction from unorganized points. In *ACM SIGGRAPH*, pages 71–78, 1992.

[HPW06] Klaus Hildebrandt, Konrad Polthier, and Max Wardetzky. On the convergence of metric and geometric properties of polyhedral surfaces. *Geometriae Dedicata*, 123(1):89–112, 2006.

[ILTC01] C. Indermitte, Th. Liebling, M. Troyanov, and H. Clémençon. Voronoi diagrams on piecewise flat surfaces and an application to biological growth. *Theor. Comput. Sci.*, 263(1-2):268–274, 2001.

[KA08] Y. J. Kil and N. Amenta. GPU-assisted surface reconstruction on locally-uniform samples. In *Meshing Roundtable*, pages ??–??, 2008.

[KBH06] M. Kazhdan, M. Bolitho, and H. Hoppe. Poisson surface reconstruction. In *Symp. Geometry Processing*, pages 61–70, 2006.

[KR85] D. G. Kirkpatrick and J. D. Radke. A framework for computational morphology. In G. T. Toussaint, editor, *Computational Geometry*, pages 217–248. 1985.

[KS98] R. Kimmel and J. A. Sethian. Computing geodesic paths on manifolds. In *Proceedings of National Academy of Sciences*, volume 95(15), pages 8431–8435, 1998.

[KW88] Rolf Klein and Derick Wood. Voronoi diagrams based on general metrics in the plane. In *STACS '88: Proceedings of the 5th Annual Symposium on Theoretical Aspects of Computer Science*, pages 281–291. Springer-Verlag, 1988.

[Lam94] Timothy Lambert. The Delaunay triangulation maximizes the mean inradius. In *Proc. 6th Canadian Conf. on Computational Geometry*, pages 201–206, 1994.

[Law77] C. L. Lawson. Software for C^1 surface interpolation. In J. R. Rice, editor, *Math. Software III*, pages 161–194. Academic Press, New York, 1977.

[Lei99] Gregory Leibon. *Random Delaunay triangulations, the Thurston-Andreev theorem, and metric uniformization.* PhD thesis, UCSD, 1999. arXiv:math/0011016v1.

[LL00] Greg Leibon and David Letscher. Delaunay triangulations and Voronoi diagrams for Riemannian manifolds. In *Symp. Comp. Geom.*, pages 341–349, 2000.

[LS03] Francois Labelle and Jonathan Richard Shewchuk. Anisotropic voronoi diagrams and guaranteed-quality anisotropic mesh generation. In *Symp. Comp. Geom.*, pages 191–200, 2003.

[LT06] Chi-Wan Lim and Tiow-Seng Tan. Surface reconstruction by layer peeling. *The Visual Computer*, 22(9):593–603, 2006.

[LZ06] J. Li and H. Zhang. Nonobtuse remeshing and decimation. In *Symp. Geometry Processing*, pages 235–238, 2006.

[MD03] Carsten Moenning and Neil Dodgson. Fast marching farthest point sampling for point clouds and implicit surfaces. Technical Report 565, University of Cambridge Computer Laboratory, 2003.

[MR06] Wolfgang Mulzer and Günter Rote. Minimum weight triangulation is NP-hard. In *Symp. Comp. Geom.*, pages 1–10, 2006.

[MT04] Jean-Marie Morvan and Boris Thibert. Approximation of the normal vector field and the area of a smooth surface. *Discrete & Computational Geometry*, 32(3):383–400, 2004.

[Mun84] James R. Munkres. *Elements of Algebraic Topology.* Addison-Wesley, 1984.

[Mus97] Oleg R. Musin. Properties of the Delaunay triangulation. In *Symp. Comp. Geom.*, pages 424–426, 1997.

[OBSC00] Atsuyuki Okabe, Barry Boots, Kokichi Sugihara, and Sung Nok Chiu. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams.* John Wiley, second edition, 2000.

[OI03] Kensuke Onishi and Jin-ichi Itoh. Estimation of the necessary number of points in Riemannian Voronoi diagram. In *Proc. 15th Canad. Conf. Comput. Geom.*, pages 19–24, 2003.

[O'R81] Joseph O'Rourke. Polyhedra of minimal area as 3D object models. In *Intl. Joint Conf. on AI*, pages 664–666, 1981.

[PB01] Sylvain Petitjean and Edmond Boyer. Regular and non-regular point sets: Properties and reconstruction. *Computational Geometry: Theory and Applications*, 19(2-3):101–126, 2001.

[PC03] Gabriel Peyré and Laurent Cohen. In *Proceedings IEEE Workshop on Variational and Level Set Methods*, pages 33–40, 2003.

[PC06] Gabriel Peyré and Laurent Cohen. Geodesic remeshing using front propagation. *International Journal on Computer Vision*, 69(1):145–156, 2006.

[PSS01] Emil Praun, Wim Sweldens, and Peter Schröder. Consistent mesh parameterizations. In *ACM SIGGRAPH*, pages 179–184, 2001.

[Raj91] V. T. Rajan. Optimality of the Delaunay triangulation in $\mathbb{R}^d$. In *Symp. Comp. Geom.*, pages 357–363, 1991.

[Raj94] V. T. Rajan. Optimality of the Delaunay triangulation in $\mathbb{R}^d$. *Discrete and Computational Geometry*, 12:189,202, 1994.

[RI97] María-Cecilia Rivara and Patricio Inostroza. Using longest-side bisection techniques for the automatic refinement of Delaunay triangulations. *Int. J. Num. Meth. in Eng.*, 40:581–597, 1997.

[Rip90] S. Rippa. Minimal roughness property of the Delaunay triangulation. *Comput. Aided Geom. Des.*, 7(6):489–497, 1990.

[Riv90] Igor Rivin. Euclidean structures on simplicial surfaces and hyperbolic volume. *Annals of Mathematics*, 139:553–580, 1990.

[RS07] Edgar A. Ramos and Bardia Sadri. Geometric and topological guarantees for the wrap reconstruction algorithm. In *Symp. on Discrete Algorithms*, pages 1086–1095, 2007.

[Rup93] Jim Ruppert. A new and simple algorithm for quality 2-dimensional mesh generation. In *SODA '93: Proceedings of the fourth annual ACM-SIAM symposium on discrete algorithms*, pages 83–92, Philadelphia, PA, USA, 1993. Society for Industrial and Applied Mathematics.

[Rup95] J. Ruppert. A Delaunay refinement algorithm for quality 2-dimensional mesh generation. *J. of Algorithms*, 18(3):548–585, 1995.

[Sau04] Emil Saucan. Surface triangulation – the metric approach. arXiv:cs.GR/0401023 v1, 2004.

[She97] Jonathan Richard Shewchuk. *Delaunay Refinement Mesh Generation*. PhD thesis, School of Computer Science, Carnegie Mellon University, 1997. Technical Report CMU-CS-97-137.

[She99] Jonathan Richard Shewchuk. Lecture notes on Delaunay mesh generation, 1999.

[She02] Jonathan Richard Shewchuk. What is a good linear finite element? - interpolation, conditioning, anisotropy, and quality measures. http://www.cs.berkeley.edu/~jrs/papers/elemj.pdf (last viewed 2008), 2002.

[Sib78] R. Sibson. Locally equiangular triangulations. *Computer Journal*, 21:243–245, 1978.

[Sib81] R. Sibson. A brief description of natural neighbour interpolation. In V. Barnet, editor, *Interpreting multivariate data*, pages 21–36. John Wiley & sons, 1981.

[vDA95] R. van Damme and L. Alboul. Tight triangulations. In *Mathematical Methods for Curves and Surfaces*, pages 517–526. Vanderbilt University Press, 1995.

[VHGR08] Evan Vanderzee, Anil N. Hirani, Damrong Guoy, and Edgar Ramos. Well-centered triangulation. Technical Report UIUCDCS-R-2008-2936, Department of Computer Science, University of Illinois at Urbana-Champaign, February 2008. Also on arXiv: arXiv:0802.2108v1 [cs.CG].

[Vor07] G. Voronoi. Nouvelles applications des paramètres continus à la théorie des formes quadratiques. *J. Reine Angew. Math.*, 133:97–178, 1907.

[Vor08] G. Voronoi. Nouvelles applications des paramètres continus à la théorie des formes quadratiques. *J. Reine Angew. Math.*, 134:198–287, 1908.

[ZOF01] Hong-Kai Zhao, Stanley Osher, and Ronald Fedkiw. Fast surface reconstruction using the level set method. In *1st IEEE Workshop on Variational and Level Set Methods*, pages 194–202, Vancouver, 2001. 8th ICCV.

[ZvKD08] Hao Zhang, Oliver van Kaick, and Ramsay Dyer. Spectral mesh processing. *Computer Graphics Forum*, pages ??–??, 2008.

Index