

Provably Good Sampling and Meshing of Surfaces

Jean-Daniel Boissonnat and Steve Oudot*

{Jean-Daniel.Boissonnat,Steve.Oudot}@sophia.inria.fr
INRIA, 2004 route des lucioles, 06902 Sophia-Antipolis, France.

August 4, 2008

Abstract

The notion of ε -sample, introduced by Amenta and Bern, has proven to be a key concept in the theory of sampled surfaces. Of particular interest is the fact that, if E is an ε -sample of a C^2 -continuous surface S for a sufficiently small ε , then the Delaunay triangulation of E restricted to S is a good approximation of S , both in a topological and in a geometric sense. Hence, if one can construct an ε -sample, one also gets a good approximation of the surface. Moreover, correct reconstruction is ensured by various algorithms.

In this paper, we introduce the notion of loose ε -sample. We show that the set of loose ε -samples contains and is asymptotically identical to the set of ε -samples. The main advantage of loose ε -samples over ε -samples is that they are easier to check and to construct. We also present a simple algorithm that constructs provably good surface samples and meshes. Given a C^2 -continuous surface S without boundary, the algorithm generates a sparse ε -sample E and at the same time a triangulated surface $\text{Del}_{|S}(E)$. The triangulated surface has the same topological type as S , is close to S for the Hausdorff distance and can provide good approximations of normals, areas and curvatures. A notable feature of the algorithm is that the surface needs only to be known through an oracle that, given a line segment, detects whether the segment intersects the surface and, in the affirmative, returns the intersection points. This makes the algorithm useful in a wide variety of contexts and for a large class of surfaces.

Key-words: Surface mesh generation, ε -sampling, surface approximation, restricted Delaunay triangulation, mesh refinement

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Curve, surface, solid, and object representations

*Corresponding author. Fax: +33-492-387-643.

1 Introduction

Meshing and reconstructing surfaces are two fundamental problems in geometry processing. In surface reconstruction, a finite set of points E on a surface S is given and one wants to compute a good approximation of S from E . This is of course only possible if E is a good sample of S in some sense. In surface mesh generation, the problem is somehow opposite. A surface S is known and we want to compute a triangulated surface that suitably approximates S . Clearly, the vertices of the triangulated surface have to sample correctly S . Hence, in both applications and also in many others, including the new arena of point set surfaces [1, 2], the notion of good sample is crucial.

The notion of ε -sample, introduced by Amenta and Bern [3], has proven to be a key concept in the theory of sampled surfaces. Roughly, an ε -sample E of a surface S is a (non necessarily uniform) point set that is sufficiently dense with respect to the *distance to the medial axis* of S – see Section 2. Of particular interest is the fact that if E is an ε -sample of a C^2 -continuous surface S for a sufficiently small ε , the Delaunay triangulation of E restricted to S , $\text{Del}_S(E)$, is a good approximation of S , both in a topological and in a geometric sense – see Section 2 for more details. Hence, given an ε -sample of a surface, it is easy to get a good approximation of the surface.

This result (and variants of it) plays a central role in the analysis of all surface reconstruction algorithms that offer theoretical guarantees [8]. In particular, if E is an ε -sample of a surface S for a sufficiently small ε , these algorithms can reconstruct a surface that has the same topology type as S and is close to S .

One drawback of the concept of ε -sample is the fact that it is difficult to check whether a sample is an ε -sample of a given surface, and even more difficult to construct a (preferably sparse) ε -sample of a given surface. This is due to the fact that a direct application of the definition of an ε -sample leads to complicated operations like cutting the surface with balls.

In this paper, we introduce the notion of loose ε -sample. The set of loose ε -samples contains and is asymptotically identical to the set of ε -samples. The main advantage of loose ε -samples over ε -samples is that they are easier to check and to construct. Indeed, checking that a sample is a loose ε -sample reduces to checking whether a finite number of spheres are small enough with respect to the distance from their centers to the medial axis of the surface.

We also present a construction algorithm which derives from Chew’s surface meshing algorithm [17]. Given a C^2 -continuous surface S without boundary, the algorithm generates a sparse ε -sample E and at the same time a triangulated surface $\text{Del}_S(E)$. The triangulated surface has the same topological type as S , is close to S for the Hausdorff distance and can provide good approximations of normals, areas and curvatures. A notable feature of the algorithm is that the surface needs only to be known through an oracle that, given a line segment, detects whether the segment intersects the surface and, in the affirmative, returns the intersection points. This makes the algorithm useful in a wide variety of contexts and for a large class of surfaces.

The paper is organized as follows. After the recall of several useful concepts and the introduction of the notion of loose ε -sample (Section 2), the paper is divided in two parts:

- The first part deals with the structural properties of loose ε -samples. In Section 3, we present some of their local properties that are used in Section 4 to establish our main results. We prove that, for sufficiently small ε , $\text{Del}_{|S}(E)$ is a 2-manifold without boundary that is ambient isotopic to S and whose Hausdorff distance to S is $O(\varepsilon^2)$. We also prove that S is covered by the so-called surface Delaunay balls, and that loose ε -samples are $\varepsilon(1 + 8.5 \varepsilon)$ -samples. In Section 5, we bound the size of loose ε -samples.
- The second part of the paper presents our surface mesh generator, as an application of the theoretical results introduced in the first part. We describe our algorithm in Sections 6 and 7, and we perform its analysis in Section 8. In Section 9, we present several improvements that can be made to the algorithm, either to speed it up or to enhance the quality of its output. Section 10 deals with the computation of the distance to the medial axis of the surface, which plays a central role in the context of ε -sampling. Finally, in Section 11 we present our implementation as well as some experimental results.

Readers mainly interested in the algorithm can skip the first part of the paper in a first reading. For completeness, we recall a few well-known results in appendix. In the body of the paper, references to these results begin with an “A” (Theorem A.1, Lemma A.2 etc.).

2 Definitions and preliminary observations

In the paper, S denotes a compact, orientable, C^2 -continuous surface without boundary. S will be called a smooth closed surface for short. By $\vec{n}(p)$ we denote the surface normal at point $p \in S$, and by $T(p)$ the plane tangent to S at p .

Our analysis uses the fact that locally a smooth closed surface is the graph of a function. More precisely, given an orthonormal frame (O, x, y, z) of \mathbb{R}^3 , a subset of \mathbb{R}^3 is said to be *xy-monotone* if it is the graph of a function of the two variables x and y . A *terrain* is a surface that is *xy-monotone* in some frame (O, x, y, z) of \mathbb{R}^3 . Similarly, given an orthonormal frame (O, x, y) of \mathbb{R}^2 , a subset of \mathbb{R}^2 is said to be *x-monotone* if it is the graph of a function of variable x .

2.1 Restricted Delaunay triangulation

In the paper, E denotes a finite point sample of S and $\text{Del}(E)$ the 3-dimensional Delaunay triangulation of E . By $\mathcal{V}(E)$ we denote the set of all edges of the Voronoi diagram of E .

We call Delaunay triangulation of E *restricted to S* , and we note $\text{Del}_{|S}(E)$, the sub-complex of $\text{Del}(E)$ that consists of the facets of $\text{Del}(E)$ whose dual Voronoi edges intersect S – see figure 11 (color section). An edge or vertex of $\text{Del}(E)$ belongs to $\text{Del}_{|S}(E)$ if it is incident to at least one facet of $\text{Del}_{|S}(E)$. Notice that we depart from the usual definition [17, 22] and do not consider vertices and edges with no incident facet of $\text{Del}_{|S}(E)$. See Figure 1 for an illustration.

A facet (resp. edge, vertex) of $\text{Del}_{|S}(E)$ is called a *restricted Delaunay facet* (resp. *restricted Delaunay edge*, *restricted Delaunay vertex*). For a restricted Delaunay facet f , we call *surface*

Delaunay ball of f any ball circumscribing f centered at some point of $S \cap f^*$, where f^* is the Voronoi edge dual to f . We call *surface Delaunay patch* the intersection of a surface Delaunay ball with S . Notice that the centers of the surface Delaunay balls are precisely the intersection points of S and $\mathcal{V}(E)$.

2.2 Loose ε -samples and ε -samples

The *medial axis* of S , denoted by M , is the topological closure of the set of points of \mathbb{R}^3 that have more than one nearest neighbour in S .

For a point $x \in \mathbb{R}^3$, we call *distance to the medial axis* at x , and write $d_M(x)$, the Euclidean distance from x to the medial axis of S .

As noticed by Amenta and Bern [3], d_M is 1-Lipschitz, i.e. $|d_M(x) - d_M(y)| \leq \|x - y\|$.

We define $d_M^{\text{inf}} = \inf \{d_M(x), x \in S\}$ and $d_M^{\text{sup}} = \sup \{d_M(x), x \in S\}$. Since S is a smooth closed surface, both d_M^{inf} and d_M^{sup} are finite positive constants.

We borrow from Amenta and Bern [3] the notion of ε -sample, defined below. In the whole paper, $B(c, r)$ denotes the ball of center c and radius r .

Definition 2.1 E is an ε -sample of S if $\forall x \in S, E \cap B(x, \varepsilon d_M(x)) \neq \emptyset$.

For sufficiently small values of ε , ε -samples enjoy many beautiful properties. We recall the most important ones in our context.

- *Normals*: the angle between the normal to a facet f of $\text{Del}_{|S}(E)$ and the normal to S at the vertices of f is $O(\varepsilon)$ [3].
- *Area*: the area of $\text{Del}_{|S}(E)$ approximates the area of S [28].
- *Curvatures*: the curvature tensor of S can be estimated from $\text{Del}_{|S}(E)$ [18].
- *Homeomorphism*: $\text{Del}_{|S}(E)$ is homeomorphic to S [3].
- *Hausdorff distance*: the Hausdorff distance between S and $\text{Del}_{|S}(E)$ is $O(\varepsilon^2)$ (Theorem 4.5).
- *Reconstruction*: several algorithms can reconstruct from E a surface that is homeomorphic [3, 4, 8, 19] or even ambient isotopic [5] to S .

We will show that these properties hold for loose ε -samples as well.

Definition 2.2 E is a loose ε -sample of S if the two following conditions are verified:

1. $\forall x \in S \cap \mathcal{V}(E), E \cap B(x, \varepsilon d_M(x)) \neq \emptyset$
2. $\text{Del}_{|S}(E)$ has vertices on all the connected components of S

Since the centers of the surface Delaunay balls are precisely the intersection points of S with the Voronoi edges, Condition 1 of Definition 2.2 is verified if and only if every surface Delaunay ball $B(c, r)$ has a radius of at most $\varepsilon d_M(c)$.

Observe that Condition 1 alone is not sufficient to control the density of E . Indeed, according to our definition of the restricted Delaunay triangulation, a point of E is a vertex of $\text{Del}_{|S}(E)$ only if at least one edge of the boundary of its Voronoi cell intersects S . It follows that some of the points of E may not be vertices of $\text{Del}_{|S}(E)$. In some situations (see Figure 1 for an example), $\text{Del}_{|S}(E)$ may even be empty, in which case Condition 1 is trivially verified for any value of ε .




Figure 1: A case when $\text{Del}_{|S}(E)$ is empty: the four points of E are placed on a torus, such that the Voronoi edges pass through the hole.

Loose ε -samples and ε -samples are closely related but not identical concepts. The next lemma follows from Definitions 2.1 and 2.2, and from Theorem 2 of [3].

Lemma 2.3 *If E is an ε -sample of S , with $\varepsilon < 0.1$, then it is also a loose ε -sample of S .*

Proof Clearly, Condition 1 of Definition 2.2 is automatically verified when E is an ε -sample of S . To show that Condition 2 is also verified, we use Theorem 2 of [3], which states that, when $\varepsilon < 0.1$, every k -face of $\text{Vor}(E)$ either does not intersect S or intersects S along a $(k - 1)$ -topological ball. Hence, the Voronoi cell p^* of any point $p \in E$ intersects S along a topological disk whose boundary lies in the boundary of p^* since S has no boundary. Moreover, the boundary of $S \cap p^*$ cannot lie inside a single facet of the boundary of p^* , since otherwise this facet would intersect S along a topological circle, and not along a topological arc. It follows that at least one edge of the boundary of p^* intersects S , which means that p is a vertex of $\text{Del}_{|S}(E)$. Moreover, since $\varepsilon < 1$, E has some points on all the connected components of S . Condition 2 of Definition 2.2 is therefore fulfilled. \square

The converse of Lemma 2.3 is true asymptotically, as we will see in Section 4.3 (Corollary 4.10).

Since adding points to an ε -sample results in another ε -sample, we cannot hope for an upper bound on ε -samples without making some additional assumptions. The same observation can also be made for loose ε -samples since ε -samples are loose ε -samples, by Lemma 2.3. This motivates the following definition.

Definition 2.4 *Given a positive constant κ and a positive 1-Lipschitz function ϕ , a point sample E of S is said to be (κ, ϕ) -sparse if $\forall x \in E$, $\text{dist}(x, E \setminus \{x\}) \geq \kappa \phi(x)$.*

Notice that (κ, ϕ) -sparse samples are not assumed to be ε -samples nor loose ε -samples. In Section 5.2 (Theorem 5.4), we will give an upper bound on the size of (κ, ϕ) -sparse samples, which will lead to a sufficient (local) condition for loose ε -samples and ε -samples to have an optimal size.

2.3 Other notations

The following constants are used in the paper:

- $\varepsilon_0 \approx 0.091$ is the smallest positive value of ε such that $\frac{2\varepsilon}{1-8\varepsilon} + \arcsin \frac{\varepsilon}{1-\varepsilon} \geq \frac{\pi}{4}$.
- $\varepsilon_1 \approx 0.096$ is the smallest positive value of ε such that $\frac{2\varepsilon}{1-7\varepsilon} + \arcsin \frac{\varepsilon\sqrt{3}}{1-\varepsilon} \geq \frac{\pi}{4}$.
- $\varepsilon_2 = \frac{\pi}{4+9\pi} \approx 0.097$.
- $\varepsilon_3 \approx 0.12$ is the smallest positive value of ε such that $\frac{2\varepsilon}{1-4\varepsilon} + \arcsin \frac{\varepsilon\sqrt{3}}{1-\varepsilon} \geq \frac{\pi}{4}$.
- $\varepsilon_4 \approx 0.17$ is the smallest positive value of ε such that $\frac{\varepsilon}{1-5\varepsilon} + \arcsin \frac{\varepsilon\sqrt{3}}{1-\varepsilon} \geq \frac{\pi}{2}$.
- $\varepsilon_5 \approx 0.065$ is the smallest positive value of ε such that $\varepsilon(1 + 8.5\varepsilon) \geq 0.1$.

We also use the notation (\vec{u}, \vec{v}) to denote the modulus of the angle (measured in $[-\pi, \pi]$) between vectors \vec{u} and \vec{v} of \mathbb{R}^3 , and $\vec{u} \cdot \vec{v}$ to denote their dot-product.

Part I

Properties of loose ε -samples

3 Local properties of loose ε -samples

In this section, we prove that surface Delaunay balls of sufficiently small radii keep important properties of planar disks. In particular, we show that they intersect S along topological disks whose boundaries pairwise intersect in at most two points (Proposition 3.10).

3.1 Technical lemmas

In this paragraph, we introduce two technical results (Lemmas 3.5 and 3.6) that will be useful in the remainder of the paper. Their proofs rely on the following statements by Amenta and Bern [3] on one hand, Mederos, Velho and Figueiredo [27] on the other hand.

Lemma 3.1 *Let f be a facet of $\text{Del}_S(E)$. Assume that some surface Delaunay ball $B(c, r)$ of f is such that $r \leq \rho d_M(c)$, with $\rho < \frac{1}{7}$. Let a be a vertex of f . If a has an inner angle of at least $\pi/3$, then the smaller angle between the line normal to f and the normal to S at a is at most $\arcsin \frac{\rho\sqrt{3}}{1-\rho}$. Otherwise, the smaller angle between the line normal to f and the normal to S at a is at most $\frac{2\rho}{1-7\rho} + \arcsin \frac{\rho\sqrt{3}}{1-\rho}$.*

Proof By assumption, the radius of $B(c, r)$ is $r \leq \rho d_M(c)$, which is at most $\rho (d_M(a) + \|c - a\|)$ since d_M is 1-Lipschitz. Since $B(c, r)$ circumscribes f , we have $\|c - a\| = r$. Thus, $r \leq \frac{\rho}{1-\rho} d_M(a)$. It follows that the proof of Lemma 7 of [3] holds here. \square

Lemma 3.2 *For any two points p and q on S with $\|p - q\| \leq \rho d_M(p)$, the smaller angle between the line segment pq and the surface normal at p is at least $\frac{\pi}{2} - \arcsin \frac{\rho}{2}$.*

Lemma 3.3 *For any two points p and q on S with $\|p - q\| \leq \rho \min \{d_M(p), d_M(q)\}$, for any $\rho < \frac{1}{3}$, the modulus of the angle between the normals to S at p and at q is at most $\frac{\rho}{1-3\rho}$.*

Lemma 3.4 *Let f and f' be two facets of $\text{Del}_S(E)$ with a common edge. If f and f' are circumscribed respectively by surface Delaunay balls $B(c, r)$ and $B(c', r')$, such that $r < \varepsilon_3 d_M(c)$ and $r' < \varepsilon_3 d_M(c')$, then the dihedral angle formed by f and f' is greater than $\frac{\pi}{2}$.*

Proof This result is a consequence of Theorem 1 of [27], which states that the dihedral angle formed by f and f' is at least $\pi - 2 \left(\frac{2\varepsilon_3}{1-4\varepsilon_3} + \arcsin \frac{\varepsilon_3\sqrt{3}}{1-\varepsilon_3} \right)$, which is greater than $\frac{\pi}{2}$. \square

Using the above statements, we can now prove our technical results.

Lemma 3.5 *Let c and c' be two points of S such that $\|c - c'\| \leq \varepsilon (d_M(c) + d_M(c'))$, where $\varepsilon < \frac{1}{8}$. There exists a vector \vec{v} orthogonal to $\vec{cc'}$, such that the angle between \vec{v} and the normal to S at any point of $S \cap B(c, 2\varepsilon d_M(c))$ is at most $\frac{2\varepsilon}{1-8\varepsilon} + \arcsin \frac{\varepsilon}{1-\varepsilon}$. Hence, if $\varepsilon \leq \varepsilon_0$, this angle is at most $\frac{\pi}{4}$.*

Proof Let $B^+ = B(c, 2\varepsilon d_M(c))$. We have

$$\forall x \in B^+ \cap S, \|x - c\| \leq 2\varepsilon d_M(c) \quad (1)$$

thus

$$\begin{aligned} \forall x \in B^+ \cap S, d_M(x) &\geq d_M(c) - \|x - c\| \\ &\geq (1 - 2\varepsilon) d_M(c) \end{aligned} \quad (2)$$

(1) and (2) give

$$\forall x \in B^+ \cap S, \|x - c\| \leq \frac{2\varepsilon}{1-2\varepsilon} \min \{d_M(c), d_M(x)\}$$

which implies, according to Lemma 3.3,

$$\forall x \in B^+ \cap S, (\vec{n}(x), \vec{n}(c)) \leq \frac{2\varepsilon/(1-2\varepsilon)}{1-6\varepsilon/(1-2\varepsilon)} = \frac{2\varepsilon}{1-8\varepsilon}$$

By hypothesis, $\|c - c'\| \leq \varepsilon d_M(c) + \varepsilon d_M(c') \leq 2\varepsilon d_M(c) + \varepsilon \|c - c'\|$, hence $\|c - c'\| \leq \frac{2\varepsilon}{1-\varepsilon} d_M(c)$. Lemma 3.2 then tells that

$$\min \{(\vec{cc'}, \vec{n}(c)), (\vec{c'c}, \vec{n}(c))\} \geq \frac{\pi}{2} - \arcsin \frac{\varepsilon}{1-\varepsilon} \quad (3)$$

Inside plane $(c, \vec{cc'}, \vec{n}(c))$, let \vec{v} be the unitary vector that is orthogonal to $\vec{cc'}$ and has a positive dot-product with $\vec{n}(c)$. According to (3) we have $(\vec{n}(c), \vec{v}) \leq \arcsin \frac{\varepsilon}{1-\varepsilon}$. Thus,

$$\begin{aligned} \forall x \in S \cap B^+, (\vec{n}(x), \vec{v}) &\leq (\vec{n}(x), \vec{n}(c)) + (\vec{n}(c), \vec{v}) \\ &\leq \frac{2\varepsilon}{1-8\varepsilon} + \arcsin \frac{\varepsilon}{1-\varepsilon} \end{aligned}$$

□

Lemma 3.6 *An edge of $\text{Vor}(E)$ cannot intersect S in more than one point x such that $\text{dist}(x, E) < \frac{1}{7} d_M(x)$. Moreover, in such a point the intersection is transversal.*

Proof Let $\varepsilon < \frac{1}{7}$. Let e be an edge of $\text{Vor}(E)$ and f its dual Delaunay facet. We assume that e intersects S in at least one point x such that $\text{dist}(x, E) \leq \varepsilon d_M(x)$. By definition, x is the center of a surface Delaunay ball of f . Let a be the vertex of f that has the largest inner angle. The value of this angle is at least $\frac{\pi}{3}$, hence Lemma 3.1 says that

$$(\vec{n}(a), \vec{n}_f) \leq \arcsin \frac{\varepsilon\sqrt{3}}{1-\varepsilon} \quad (4)$$

where \vec{n}_f denotes the unitary vector orthogonal to f that makes the smaller angle with $\vec{n}(a)$. Let B_a be the ball $B(a, \frac{\varepsilon}{1-\varepsilon} d_M(a))$. For any point x as above, we have

$$\begin{aligned} \|x - a\| = \text{dist}(x, E) &\leq \varepsilon d_M(x) \\ &\leq \varepsilon (d_M(a) + \|x - a\|) \end{aligned}$$

Hence, $\|x - a\| \leq \frac{\varepsilon}{1-\varepsilon} d_M(a)$. In other words, x lies in B_a . We now show that $B_a \cap S$ is a terrain over the plane Π_f that supports f . For any $y \in B_a \cap S$, we have

$$\|y - a\| \leq \frac{\varepsilon}{1-\varepsilon} d_M(a) \leq \frac{\varepsilon}{1-\varepsilon} (d_M(y) + \|y - a\|)$$

which implies $\|y - a\| \leq \frac{\varepsilon}{1-2\varepsilon} d_M(y)$. According to Lemma 3.3, we then have

$$(\vec{n}(y), \vec{n}(a)) \leq \frac{\varepsilon/(1-2\varepsilon)}{1-3\varepsilon/(1-2\varepsilon)} = \frac{\varepsilon}{1-5\varepsilon} \quad (5)$$

(4) and (5) give

$$(\vec{n}(y), \vec{n}_f) \leq (\vec{n}(y), \vec{n}(a)) + (\vec{n}(a), \vec{n}_f) \leq \frac{\varepsilon}{1-5\varepsilon} + \arcsin \frac{\varepsilon\sqrt{3}}{1-\varepsilon}$$

which is less than $\frac{\pi}{2}$ since $\varepsilon < \frac{1}{7} < \varepsilon_4$. Thus, by Lemma A.4, $B_a \cap S$ is a terrain over Π_f . Since e is orthogonal to Π_f , it cannot intersect $B_a \cap S$ more than once, nor tangentially. The result follows since, as shown above, every point $x \in e \cap S$ such that $\text{dist}(x, E) \leq \varepsilon d_M(x)$ lies in B_a . \square

3.2 Topological disks and terrains

Lemma 3.7 ([7]) *Let B be a ball that intersects S . If the intersection is not a topological disk, then B contains a point of the medial axis of S . As a consequence, if E is a loose ε -sample, with $\varepsilon < 1$, then surface Delaunay patches are topological disks.*

Lemma 3.8 *If E is a loose ε -sample, with $\varepsilon < \varepsilon_2$, then, for every surface Delaunay ball $B = B(c, r)$, for any point $x \in S \cap B$, $S \cap B(x, 2r)$ is a topological disk and a terrain over $T(x)$.*

Proof Since E is a loose ε -sample, we have $r \leq \varepsilon d_M(c) \leq \varepsilon(d_M(x) + \|x - c\|) \leq \varepsilon(d_M(x) + r)$, that is, $r \leq \frac{\varepsilon}{1-\varepsilon} d_M(x)$. Thus, $2r < d_M(x)$ since $\varepsilon < \varepsilon_2 < \frac{1}{3}$. According to Lemma 3.7, $S \cap B(x, 2r)$ is thus a topological disk.

In addition, since $r \leq \frac{\varepsilon}{1-\varepsilon} d_M(x)$, we have

$$\begin{aligned} \forall y \in B(x, 2r), \|y - x\| &\leq 2r \leq \frac{2\varepsilon}{1-\varepsilon} d_M(x) \\ &\leq \frac{2\varepsilon}{1-\varepsilon} (d_M(y) + \|y - x\|) \end{aligned}$$

which implies that $\|y - x\| \leq \frac{2\varepsilon}{1-3\varepsilon} d_M(y)$. It follows that $\|y - x\| \leq \rho \min\{d_M(y), d_M(x)\}$, with $\rho = \frac{2\varepsilon}{1-3\varepsilon}$. Thus, according to Lemma 3.3,

$$\forall y \in B(x, 2r), (\vec{n}(y), \vec{n}(x)) \leq \frac{\rho}{1-3\rho} = \frac{2\varepsilon}{1-9\varepsilon}$$

Since $\varepsilon < \varepsilon_2$, we have $\frac{2\varepsilon}{1-9\varepsilon} < \frac{\pi}{2}$, hence $S \cap B(x, 2r)$ is a terrain by Lemma A.4. \square

3.3 Pseudo-disks

Definition 3.9 *Topological disks are pseudo-disks if they pairwise intersect along topological disks (that may be empty or reduced to a point) and if their boundaries pairwise intersect in at most two points.*

Observe that the boundaries of two pseudo-disks either do not intersect, or intersect in one point tangentially, or intersect in two points transversally.

The aim of this section is to prove the following proposition:

Proposition 3.10 *If E is a loose ε -sample, with $\varepsilon \leq \varepsilon_0$, then surface Delaunay patches are pseudo-disks.*

To assist the reader, the proof is organized in a hierarchical way, with several lemmas and claims with independent proofs. These proofs can be skipped in a first reading.

Proof of Proposition 3.10

Let $B = B(c, r)$ and $B' = B(c', r')$ be two surface Delaunay balls. According to Lemma 3.7, $D = B \cap S$ and $D' = B' \cap S$ are topological disks, since $\varepsilon \leq \varepsilon_0 < 1$. Their boundaries C and C' are topological circles. Let us assume that balls B and B' intersect, the other case being trivial. Notice that none of them can be contained in the other one, since they are Delaunay balls. Thus, their bounding spheres ∂B and $\partial B'$ also intersect. Let Γ be the circle $\partial B \cap \partial B'$, ρ its radius ($\rho < \min\{r, r'\}$) and P its supporting plane. We define $\Delta = B \cap P$ and notice that $\Gamma = \partial\Delta$. Since S is a closed surface, we have $C \subset \partial B$ and $C' \subset \partial B'$, which implies that

$$C \cap C' \subseteq S \cap \Gamma \tag{6}$$

Let $B^+ = B(c, 2r)$. Since $\|c - c'\| \leq r + r' \leq \varepsilon d_M(c) + \varepsilon d_M(c')$, with $\varepsilon \leq \varepsilon_0$, Lemma 3.5 says that there exists a vector \vec{v} orthogonal to $\vec{cc'}$ such that

$$\forall x \in S \cap B^+, (\vec{n}(x), \vec{v}) \leq \frac{\pi}{4} \tag{7}$$

Let us choose in \mathbb{R}^3 a reference frame with origin c , y -axis directed along $\vec{c'c}$, and z -axis directed along \vec{v} . We call L_l and L_r the two lines of P , parallel to the z -axis, that are tangent to Γ . The region of P bounded by L_l and L_r is called G (see Figure 2). In the following, ξ denotes $S \cap B^+ \cap G$.




Figure 2: Definitions of G and G'

Lemma 3.11 ξ is a connected x -monotone arc.

Proof According to (7), we have $\forall x \in S \cap B^+$, $(\vec{n}(x), \vec{v}) \leq \frac{\pi}{4}$. Thus, by Lemma A.4, $B^+ \cap S$ is xy -monotone, which implies that ξ is x -monotone. Moreover, according to Lemma A.5, $B^+ \cap S$ lies outside the cone of apex $c \in S$, of vertical axis and of half-angle $\frac{\pi}{4}$. The equation of the cone in our frame is $z^2 = x^2 + y^2$. It intersects P along two hyperbolic arcs of equations $z = \pm\sqrt{x^2 + d^2}$, where $d \leq r$ is the distance from c to P . Consider the subregion G' of G that is bounded vertically by the two hyperbolic arcs (see Figure 2). Since $S \cap B^+$ lies outside the cone, ξ is included in G' .

The points of G' that are farthest from c are the points $(\pm\rho, -d, \pm\sqrt{\rho^2 + d^2})$. Their distance to c is

$$\sqrt{2(\rho^2 + d^2)} < 2r$$

In other words, $G' \subset \text{int}(B^+)$. It follows that ξ is included in $\text{int}(B^+)$ and cannot intersect ∂B^+ . Its endpoints must then lie on the vertical lines L_r and L_l . But there can be only one endpoint per vertical line, since ξ is x -monotone. Hence, ξ has at most two endpoints and is thus connected. \square

Lemma 3.12 $|S \cap \Gamma| \leq 2$.

Proof Let us assume for a contradiction that $|S \cap \Gamma| > 2$. First, we show that there exists a point where the curvature of ξ is high and hence the distance to the medial axis M is small. Then we work out a contradiction with the fact that E is a loose ε -sample, with $\varepsilon \leq \varepsilon_0$.

Claim 3.12.1 *There exists a point q at which the curvature of ξ is at least $\frac{1}{\rho}$.*

Proof We made the assumption that $|S \cap \Gamma| > 2$. Since $\Gamma \subset G$ and $\Gamma \subset B^+$, ξ also intersects Γ more than twice. And since ξ is connected by Lemma 3.11, there is a subarc ab of ξ that lies

outside Δ and whose endpoints a and b lie on Γ . This subarc may be reduced to a point ($a = b$), since ξ may be tangent to Γ . But in this case, in the vicinity of a , ξ is locally included in Δ and tangent to Γ at a . Thus, its curvature at a is at least $\frac{1}{\rho}$, which proves the claim with $q = a$. So now we assume that arc ab of ξ is not reduced to a point. Since ξ is x -monotone by Lemma 3.11, a and b lie on the same half of Γ , upper half or lower half (say upper half). Thus, the smaller arc of Γ that joins a and b is also x -monotone. Then, by Lemma A.3, there is a point q of arc ab of ξ at which the curvature of ξ is at least $\frac{1}{\rho}$, which proves the claim. \square

Claim 3.12.2 $d_M(q) \leq \rho\sqrt{2}$.

Proof Let $\vec{n}_\xi(q)$ be the normal to planar curve ξ at point q . By (7), $\vec{n}(q)$ is not orthogonal to P , thus $\vec{n}_\xi(q)$ is oriented along the projection of $\vec{n}(q)$ onto P . Hence, by Lemma A.2, we have $(\vec{n}(q), \vec{n}_\xi(q)) \leq (\vec{n}(q), \vec{v})$ which is at most $\frac{\pi}{4}$ by (7). According to Theorem A.1, we then have at q

$$\mathbb{I}(\xi', \xi') \geq \cos \frac{\pi}{4} \|\xi''\|$$

ξ' is the unit tangent vector of ξ at q and $\|\xi''\|$ is the curvature of ξ at q , which is more than $\frac{1}{\rho}$ according to Claim 3.12.1. So, at q we have

$$\mathbb{I}(\xi', \xi') \geq \frac{1}{\rho\sqrt{2}} \quad (8)$$

Recall that \mathbb{I} is a symmetric bilinear form, thus it can be diagonalized in an orthonormal frame, and its eigenvalues are the minimum and maximum curvatures of S at q . Let us call these values $\kappa_{\min}(q)$ and $\kappa_{\max}(q)$ respectively. Since ξ' is a unit vector, we have $\mathbb{I}(\xi', \xi') \leq \max\{|\kappa_{\min}(q)|, |\kappa_{\max}(q)|\}$. It follows, according to (8), that $\max\{|\kappa_{\min}(q)|, |\kappa_{\max}(q)|\} \geq \frac{1}{\rho\sqrt{2}}$, or, equivalently, that the minimal radius of curvature of S at q is at most $\rho\sqrt{2}$. The claim follows. \square

The end of the proof of Lemma 3.12 is now immediate. We have

$$\begin{aligned} d_M(c) &\leq d_M(q) + \|c - q\| \\ &\leq \rho\sqrt{2} + 2r \\ &\leq r(\sqrt{2} + 2) \end{aligned}$$

So, the radius of ball B is at least $\frac{1}{\sqrt{2}+2} d_M(c)$, which contradicts the assumption that E is a loose ε -sample, with $\varepsilon \leq \varepsilon_0 < \frac{1}{\sqrt{2}+2}$. This ends the proof of Lemma 3.12. \square

From Lemma 3.12, it immediately follows that $|C \cap C'| \leq 2$, by (6).

Lemma 3.13 $S \cap \Delta$ is not reduced to two points.

Proof Let us assume that S intersects Δ in two points exactly, say a and b . Then, the subarc of ξ that joins points a and b lies outside Δ . It follows, by the same reasoning as in the proof of

Claim 3.12.1, that there exists some point q of ξ at which the curvature of ξ is at least $\frac{1}{\rho}$. It follows by Claim 3.12.2 that $d_M(q) \leq \rho\sqrt{2}$, which leads to a contradiction, as in the end of the proof of Lemma 3.12. \square

We can now complete the proof of Proposition 3.10. It follows from Lemmas 3.12 and 3.13 that D and D' intersect along a topological disk. The result is clear if $D \subseteq D'$ or if $D' \subseteq D$. Otherwise, we have $|C \cap C'| \leq 2$, by Lemma 3.12. If $|C \cap C'| = 0$, then $D \cap D'$ is empty. If $|C \cap C'| = 1$, then $D \cap D'$ is reduced to a point. If $|C \cap C'| = 2$, then $D \cap D'$ is either a topological disk or equal to $C \cap C'$. But if $D \cap D' = C \cap C'$, then $S \cap \Delta = C \cap C'$ since $C \cap C' \subseteq S \cap \Delta \subseteq D \cap D'$. This contradicts Lemma 3.13. Hence, $D \cap D'$ is not equal to $C \cap C'$ and is therefore a topological disk. This ends the proof of Proposition 3.10. \square

4 Global properties of loose ε -samples

In this section, E is a loose ε -sample of S , with $\varepsilon \leq \varepsilon_0$. Using the results of the previous section, we prove that $\text{Del}_S(E)$ is a manifold without boundary (Theorem 4.3), ambient isotopic to S (Theorem 4.4), at Hausdorff distance $O(\varepsilon^2)$ from S (Theorem 4.5). From the latter we deduce that E is an $\varepsilon(1 + 8.5\varepsilon)$ -sample of S (Corollary 4.10). We also prove that the surface Delaunay balls cover S (Theorem 4.12).

4.1 Manifold

We first prove that every edge of $\text{Del}_S(E)$ is incident to exactly two facets of $\text{Del}_S(E)$. We then prove that every vertex of $\text{Del}_S(E)$ has only one *umbrella*. An umbrella of a vertex v is a subset of facets of $\text{Del}_S(E)$ incident to v whose adjacency graph is a cycle.

Since E is a loose ε -sample of S , with $\varepsilon \leq \varepsilon_0$, every point $x \in S \cap \mathcal{V}(E)$ is at distance less than $\frac{1}{7} d_M(x)$ from E . Thus, by Lemma 3.6, the dual of any facet f of $\text{Del}_S(E)$ intersects S only once. We denote by $B_f = (c_f, r_f)$ the only surface Delaunay ball that circumscribes f . The surface Delaunay patch of f , $S \cap B_f$, is denoted by D_f . Let $C_f = \partial D_f$.

Proposition 4.1 *Every edge of $\text{Del}_S(E)$ is incident to exactly two facets of $\text{Del}_S(E)$.*

Proof Let e be an edge of $\text{Del}_S(E)$. We denote by e^* the Voronoi face dual to e . Since S has no boundary, its intersection with $\text{Aff}(e^*)$ (the affine hull of e^*) is a collection of simple closed curves, none of which intersects the boundary ∂e^* of e^* tangentially, by Lemma 3.6. Thus, by the Jordan curve theorem, each component of $S \cap \text{Aff}(e^*)$ intersects ∂e^* in an even number of points. It follows that S intersects ∂e^* in an even number of points. Moreover, by Lemma 3.6, each edge of ∂e^* is intersected at most once by S . Thus, S intersects an even number of edges of ∂e^* , and e is incident to an even number of restricted Delaunay facets.

In addition, two restricted Delaunay facets incident to e make a dihedral angle greater than $\frac{\pi}{2}$, by lemma 3.4. It follows that e may be incident to at most three restricted Delaunay facets.

In conclusion, the number of facets of $\text{Del}_{|S}(E)$ incident to e is even, at least one and at most three. The result follows. \square

It follows from the above proposition that the restricted Delaunay facets incident to a vertex of $\text{Del}_{|S}(E)$ form a set of umbrellas.

Proposition 4.2 *Every vertex of $\text{Del}_{|S}(E)$ has exactly one umbrella.*

Proof Let v be a vertex of $\text{Del}_{|S}(E)$. Let $F(v)$ be the set of all facets of $\text{Del}_{|S}(E)$ that are incident to v . Let f_v be the facet of $F(v)$ that has the surface Delaunay ball of largest radius. We call r_{f_v} this radius. Then $B(v, 2r_{f_v})$ contains the surface Delaunay balls of all facets of $F(v)$. Moreover, by lemma 3.8, $S \cap B(v, 2r_{f_v})$ is a topological disk and a terrain over $T(v)$, the tangent plane of S at v . Hence, the orthogonal projection of $S \cap B(v, 2r_{f_v})$ onto $T(v)$ is one-one and therefore preserves several topological properties, such as pseudo-disks or the fact that the interiors of surface Delaunay patches are empty of points of E . For simplicity of notations, we identify objects with their projection onto $T(v)$. Let $F_1(v)$ be one of the umbrellas formed by the facets of $F(v)$. We call $U_1(v)$ the union of the facets of $F_1(v)$, and $R_1(v)$ the union of the surface Delaunay patches associated with the facets of $F_1(v)$.

Claim 4.2.1 $v \in \text{int}(R_1(v))$.

Proof If $v \in \text{int}(U_1(v))$, then it is clear that $v \in \text{int}(R_1(v))$, since surface Delaunay patches are pseudo-disks, by proposition 3.10. Therefore, it suffices to prove that $v \in \text{int}(U_1(v))$. Let us assume the contrary. Let $[vu]$ be an edge of the boundary of $U_1(v)$ that is incident to v . $[vu]$ is incident to two facets of $F_1(v)$, say (v, u, w) and (v, u, w') . These facets both lie on the same side of $[vu]$, which is a boundary edge of $U_1(v)$. Since surface Delaunay patches are pseudo-disks by proposition 3.10, either w is included in the interior of the surface Delaunay patch of (v, u, w') or w' is included in the interior of the surface Delaunay patch of (v, u, w) , which contradicts the fact that these patches are empty of points of E . This ends the proof of the claim. \square

We will now prove that $F_1(v)$ is the only umbrella of v . We assume for a contradiction that there exists a restricted Delaunay facet $f = (v, u, w) \notin F_1(v)$ that is incident to v . Vertices u and w lie outside $\text{int}(R_1(v))$, whereas v lies inside $\text{int}(R_1(v))$, by Claim 4.2.1. It follows that C_f intersects the boundary of $R_1(v)$, at some point z that lies on the boundary of the surface Delaunay patch of some facet $f' = (v, u', w')$ of $F_1(v)$. By proposition 3.10, C_f and $C_{f'}$ intersect at points v and z only. By the same proposition, open arcs (v, u') and (v, w') of $C_{f'}$ are included in $\text{int}(R_1(v))$. Since $v \in \text{int}(R_1(v))$, z lies on arc (u', w') of $C_{f'}$. If $z \neq u'$ and $z \neq w'$, then u' and w' lie on different sides of C_f , hence one of them lies in $\text{int}(D_f)$, which violates the Delaunay property. Otherwise (say $z = u'$), w' must lie outside D_f . In this case, consider the facet $f'' = (v, u', w'')$ of $F_1(v)$ that is incident to f' through edge $[u', w']$. By proposition 3.10, C_f intersects arc (u', w'') of $C_{f''}$ at point u' only, thus u' and w'' lie on different sides of C_f . Hence, either u' or w'' lies in $\text{int}(D_f)$, which contradicts the fact that $\text{int}(D_f)$ is empty of points of E . \square

The next theorem follows from Propositions 4.1 and 4.2.

Theorem 4.3 *Let S be a smooth closed surface and E a loose ε -sample of S . If $\varepsilon \leq \varepsilon_0 \approx 0.091$, then $\text{Del}_{|S}(E)$ is a 2-manifold without boundary.*

Since $\text{Del}_{|S}(E)$ is a closed 2-manifold embedded in \mathbb{R}^3 , we can orient the normals of its facets consistently. For instance, they can be chosen so as to point to the unbounded component of $\mathbb{R}^3 \setminus \text{Del}_{|S}(E)$.

4.2 Homeomorphism and ambient isotopy

Let $\pi : \mathbb{R}^3 \rightarrow S$ map each point of \mathbb{R}^3 to the closest point of S . In [4], the authors have shown that the restriction of π to a 2-simplicial complex W whose vertices lie on S is a homeomorphism between W and S , provided that:

H0 W is a manifold without boundary.

H1 W has vertices on all the connected components of S .

H2 The angle between the oriented normals of two facets of W sharing a vertex is lower than $\frac{\pi}{2}$.

H3 (SMALL TRIANGLE CONDITION) every facet f of W has a surface Delaunay ball of radius at most $0.113 \min \{d_M(v), v \text{ vertex of } f\}$.

H4 (FLAT TRIANGLE CONDITION) for every facet f of W , the line normal to f makes an angle of at most 0.375 radians with $\vec{n}(v)$, where v is the vertex of f of largest inner angle.

We will show that, since E is a loose ε -sample of S , with $\varepsilon \leq \varepsilon_0$, Assertions H0-H4 are verified by $W = \text{Del}_{|S}(E)$. H0 has already been stated for $\text{Del}_{|S}(E)$ in Theorem 4.3. H1 is guaranteed by Condition 2 of Definition 2.2.

Proof of H2

Let v be a vertex of $\text{Del}_{|S}(E)$ and let $F(v)$ be the umbrella of v . By Lemma 3.1, the smaller angle between $\vec{n}(v)$ and the line normal to any facet of $F(v)$ is at most $\frac{2\varepsilon}{1-7\varepsilon} + \arcsin \frac{\varepsilon\sqrt{3}}{1-\varepsilon}$, which is less than $\frac{\pi}{4}$ since $\varepsilon < \varepsilon_1$. It follows that the angle between $\vec{n}(v)$ and the oriented normal of the facet is less than $\frac{\pi}{4}$ or greater than $\frac{3\pi}{4}$. Moreover, any two consecutive facets in the umbrella of v make a dihedral angle greater than $\frac{\pi}{2}$, by Lemma 3.4, thus the angles between $\vec{n}(v)$ and the oriented normals of the facets of $F(v)$ are all less than $\frac{\pi}{4}$, or they are all greater than $\frac{3\pi}{4}$. It follows that the angle between the oriented normals of any two facets of $F(v)$ is less than $\frac{\pi}{2}$. \square

Proof of H3

Since E is a loose ε -sample, every facet f of $\text{Del}_{|S}(E)$ has a surface Delaunay ball $B_f = B(c_f, r_f)$ of radius $r_f \leq \varepsilon d_M(c_f)$. Let v be any vertex of f . We have $d_M(c_f) \leq d_M(v) + \|v - c_f\| \leq d_M(v) + r_f$, thus $r_f \leq \frac{\varepsilon}{1-\varepsilon} d_M(v)$, which is less than $0.113 d_M(v)$ since $\varepsilon < 0.1$. \square

Proof of H4

Let $f \in \text{Del}_{|S}(E)$ and let v be the vertex of f of largest inner angle. By Lemma 3.1, $\vec{n}(v)$ and

the line normal to f make an angle of at most $\arcsin \frac{\varepsilon\sqrt{3}}{1-\varepsilon}$, which is less than 0.375 radians since $\varepsilon < 0.17$. \square

We can now state the main result of this section.

Theorem 4.4 *Let S be a smooth closed surface and E a loose ε -sample of S , with $\varepsilon \leq \varepsilon_0 \approx 0.091$. The restriction of the mapping π to $\text{Del}_{|S}(E)$ is a homeomorphism between $\text{Del}_{|S}(E)$ and S . Moreover, $\text{Del}_{|S}(E)$ and S are ambient isotopic.*

Proof The fact that $\pi|_{\text{Del}_{|S}(E)}$ is a homeomorphism follows directly from H0-H4 and from Theorem 19 of [4]. Moreover, since the SMALL TRIANGLE CONDITION (H3) is verified by the facets of $\text{Del}_{|S}(E)$, Lemma 12 of [4] says that $\forall x \in \text{Del}_{|S}(E)$, $\|x - \pi(x)\| < 0.165 d_M(\pi(x))$. It follows that $\text{Del}_{|S}(E)$ and S are ambient isotopic, by Theorem 9 of [5]. \square

4.3 Hausdorff distance

Theorem 4.5 *Let S be a smooth closed surface and E a loose ε -sample of S , with $\varepsilon \leq \varepsilon_0 \approx 0.091$. The Hausdorff distance between S and $\text{Del}_{|S}(E)$ is at most $4.5 \varepsilon^2 d_M^{\text{sup}}$.*

The idea is to bound the distance from $\text{Del}_{|S}(E)$ to S , and then to use the surjectivity of π to prove that the bound also holds for the distance from S to $\text{Del}_{|S}(E)$.

Lemma 4.6 *Let $c \in S$. For any point $x \in S$ at distance at most $\varepsilon d_M(c)$ from c , the distance from x to $T(c)$ is at most $\frac{1}{2}\varepsilon^2 d_M(c)$.*

Proof Let B_1 and B_2 be the two balls of radius $d_M(c)$, tangent to S at c . Their interiors cannot intersect S and therefore do not contain x . Let x' be the intersection point other than c of the segment $[c, x]$ with the boundary of $B_1 \cup B_2$. Let h be the distance of x to $T(c)$ and θ the angle between \vec{cx} and $T(c)$. We have

$$\|c - x'\| = 2d_M(c) \sin \theta \leq \|c - x\| \leq \varepsilon d_M(c)$$

Therefore, $\sin \theta \leq \frac{\varepsilon}{2}$ and $h = \|c - x\| \sin \theta \leq \frac{1}{2}\varepsilon^2 d_M(c)$. \square

Lemma 4.7 *Let $c \in S$ and let y be a point of $T(c)$ at distance at most $\varepsilon d_M(c)$ from c . The distance of y to S is at most $4\varepsilon^2 d_M(c)$.*

Proof Let z be the point of S closest to y , t the projection of z onto $T(c)$ and $\phi = \angle yzt$, which is also the angle between the normals to S in c and in z . We have

$$\|c - z\| \leq \|c - y\| + \|y - z\| \leq 2\|c - y\| \leq 2\varepsilon d_M(c) \tag{9}$$

It then follows from Lemma 4.6 that $\|z-t\| \leq 2\varepsilon^2 d_M(c)$. Moreover, since $d_M(c) \leq d_M(z) + \|c-z\|$, (9) implies that $\|c-z\| \leq \frac{2\varepsilon}{1-2\varepsilon} d_M(z)$. It follows from Lemma 3.3 that $\phi \leq \frac{2\varepsilon}{1-8\varepsilon}$. Since $\varepsilon < 0.1$, we have $\frac{2\varepsilon}{1-8\varepsilon} \leq 1$, thus $\phi \leq 1$. It follows that $\frac{1}{\cos \phi} \leq \frac{1}{1-\frac{\phi^2}{2}} \leq 1 + \phi^2$, from which we deduce

$$\begin{aligned} \|y-z\| &= \frac{\|z-t\|}{\cos \phi} \leq 2\varepsilon^2 d_M(c) \left(1 + \left(\frac{2\varepsilon}{1-8\varepsilon}\right)^2\right) \\ &\leq 4\varepsilon^2 d_M(c) \end{aligned}$$

□

Using Lemmas 4.6 and 4.7, we can bound the distance from $\text{Del}_S(E)$ to S .

Proposition 4.8 *Every point $x \in \text{Del}_S(E)$ is at distance at most $4.5 \varepsilon^2 d_M(c) \leq 4.5 \varepsilon^2 d_M^{\text{sup}}$ from S , where c is the center of the surface Delaunay ball of the facet that contains x .*

Proof Let $x \in \text{Del}_S(E)$. Let f be a facet of $\text{Del}_S(E)$ on which x lies, and let $B(c, r)$ be the surface Delaunay ball of f . Let x' be the orthogonal projection of x onto $T(c)$. We have $\|x-c\| \leq r$, which is at most $\varepsilon d_M(c)$ since E is a loose ε -sample. Thus, by Lemma 4.6, the distance from x to $T(c)$ is at most $\frac{1}{2}\varepsilon^2 d_M(c)$. Hence, $\|x-x'\| \leq \frac{1}{2}\varepsilon^2 d_M(c)$. In addition, we have $\|x'-c\| \leq \|x-c\| \leq \varepsilon d_M(c)$. Thus, by Lemma 4.7, the distance from x' to S is at most $4\varepsilon^2 d_M(c)$. It follows that the distance from x to S is at most $4.5 \varepsilon^2 d_M(c) \leq 4.5 \varepsilon^2 d_M^{\text{sup}}$. □

We can now bound the distance from S to $\text{Del}_S(E)$, which completes the proof of Theorem 4.5.

Proposition 4.9 *Every point $x \in S$ is at distance at most $\min\{4.5 \varepsilon^2 d_M^{\text{sup}}, 5.3 \varepsilon^2 d_M(x)\}$ from $\text{Del}_S(E)$.*

Proof Let $x \in S$. Since the restriction of π to $\text{Del}_S(E)$ is surjective, we have $\pi_{|\text{Del}_S(E)}^{-1}(x) \neq \emptyset$. Let $x' \in \pi_{|\text{Del}_S(E)}^{-1}(x)$. According to Proposition 4.8, $\|x-x'\| \leq 4.5 \varepsilon^2 d_M(c) \leq 4.5 \varepsilon^2 d_M^{\text{sup}}$, where c is the center of the surface Delaunay ball of a facet that contains x' .

In addition, we have $\|x'-c\| \leq \varepsilon d_M(c)$, since E is a loose ε -sample. Thus, $\|x-c\| \leq (\varepsilon + 4.5 \varepsilon^2) d_M(c) \leq (\varepsilon + 4.5 \varepsilon^2)(d_M(x) + \|x-c\|)$. It follows that $\|x-c\| \leq \frac{\varepsilon + 4.5 \varepsilon^2}{1 - \varepsilon - 4.5 \varepsilon^2} d_M(x)$, which is at most $0.17 d_M(x)$ since $\varepsilon < 0.1$. Hence,

$$\|x-x'\| \leq 4.5 \varepsilon^2 d_M(c) \leq 4.5 \varepsilon^2 (d_M(x) + \|x-c\|) \leq 5.3 \varepsilon^2 d_M(x)$$

□

By Lemma 2.3, we know that ε -samples are loose ε -samples, for a sufficiently small value of ε . The converse is not true, but the following corollary shows that loose ε -samples are close to being ε -samples.

Corollary 4.10 *Let S be a smooth closed surface and E a loose ε -sample of S , with $\varepsilon \leq \varepsilon_0 \approx 0.091$. E is an $\varepsilon(1 + 8.5\varepsilon)$ -sample of S .*

Proof By Proposition 4.9, any point $x \in S$ is at distance at most $5.3 \varepsilon^2 d_M(x)$ from $\text{Del}_{|S}(E)$. Let x' be the point of $\text{Del}_{|S}(E)$ closest to x , and f a facet of $\text{Del}_{|S}(E)$ that contains x' . We call c the center of the surface Delaunay ball of f , and c' the center of the circumcircle of f . Let v be the vertex of f closest to x' . Since x' belongs to f , we have $\|x' - v\| \leq \|c' - v\| \leq \|c - v\|$. Moreover, $\|c - v\| \leq \varepsilon d_M(c) \leq \varepsilon(d_M(v) + \|c - v\|)$, hence $\|c - v\| \leq \frac{\varepsilon}{1-\varepsilon} d_M(v)$. Thus,

$$\begin{aligned} \|x - v\| &\leq \|x - x'\| + \|x' - v\| \\ &\leq 5.3 \varepsilon^2 d_M(x) + \frac{\varepsilon}{1-\varepsilon} d_M(v) \\ &\leq 5.3 \varepsilon^2 d_M(x) + \frac{\varepsilon}{1-\varepsilon} (d_M(x) + \|x - v\|) \end{aligned}$$

It follows that

$$\|x - v\| \leq \frac{5.3(1-\varepsilon)}{1-2\varepsilon} \varepsilon^2 d_M(x) + \frac{\varepsilon}{1-2\varepsilon} d_M(x)$$

Since $\varepsilon < 0.1$, we have $\frac{5.3(1-\varepsilon)}{1-2\varepsilon} \leq 6$ and $\frac{1}{1-2\varepsilon} \leq 1 + 2.5 \varepsilon$, thus

$$\|x - v\| \leq 6 \varepsilon^2 d_M(x) + \varepsilon(1 + 2.5 \varepsilon) d_M(x)$$

□

4.4 Covering

Let $\bigcup_{f \in \text{Del}_{|S}(E)} B_f$ (or $\bigcup_f B_f$, for short) denote the union of the surface Delaunay balls.

Let f_0 be a facet of $\text{Del}_{|S}(E)$. Our goal is to prove that $C_{f_0} \subset \text{int}\left(\bigcup_f B_f\right)$. In fact, we will prove a slightly more precise result, stated as Lemma 4.11.

Let $F(f_0)$ be the set of all facets of $\text{Del}_{|S}(E)$ that share a vertex with f_0 , including f_0 itself. Since $\varepsilon \leq \varepsilon_0$, $\text{Del}_{|S}(E)$ is a manifold without boundary, by Theorem 4.3. Hence, $F(f_0) \setminus \{f_0\}$ contains one facet of $\text{Del}_{|S}(E)$ adjacent to f_0 through each edge of f_0 . We define $R(f_0)$ as the union of all surface Delaunay patches associated with facets of $F(f_0)$.

Lemma 4.11 $C_{f_0} \subseteq \text{int}(R(f_0))$.

Proof Let u, v and w be the vertices of f_0 . We call f_{uv}, f_{vw} and f_{wu} the three facets of $F(f_0)$ that are incident to f_0 through edges uv, vw and wu respectively. By Proposition 3.10, arcs uv, vw and wu of C_{f_0} are included in $D_{f_{uv}}, D_{f_{vw}}$ and $D_{f_{wu}}$ respectively, and only their endpoints may possibly lie on $C_{f_{uv}}, C_{f_{vw}}$ or $C_{f_{wu}}$. Thus, the three arcs are included in the interior of $R(f_0)$, except for their endpoints which may possibly lie on the boundary of $R(f_0)$.

We claim that u, v and w also belong to $\text{int}(R(f_0))$. Let $F(u)$ be the umbrella of facets of $\text{Del}_{|S}(E)$ incident to u , and $R(u)$ the union of the surface Delaunay patches of all facets of $F(u)$. Notice that $R(u) \subseteq R(f_0)$, since $F(u) \subseteq F(f_0)$. Let $B(c, r)$ be the ball of biggest radius, among the surface Delaunay balls of the facets of $F(u)$. For every point $x \in R(u)$, we have $\|x - u\| \leq 2r$, hence

$R(u)$ is included in $S \cap B(u, 2r)$, which is a terrain over $T(u)$, by Lemma 3.8. It follows that the surface Delaunay patches of the facets of $F(u)$, which are pseudo-disks in \mathbb{R}^3 , are also pseudo-disks in projection onto $T(u)$. Then, the projection of u lies inside the projection of $\text{int}(R(u))$, by Claim 4.2.1. It follows that $u \in \text{int}(R(u))$, since $R(u)$ is a terrain over $T(u)$. Thus, $u \in \text{int}(R(f_0))$. The reasoning holds for v and w as well, which completes the proof of the lemma. \square

Theorem 4.12 *Let S be a smooth closed surface and E a loose ε -sample of S , with $\varepsilon \leq \varepsilon_0 \approx 0.091$. S is included in $\bigcup_f B_f$.*

Proof By Lemma 4.11, the union of all surface Delaunay patches has no boundary. Thus, S does not intersect the boundary of $\bigcup_f B_f$. Moreover, since $\text{Del}_S(E)$ has vertices on all the connected components of S , $\bigcup_f B_f$ intersects all the connected components of S . It follows that S cannot exit $\bigcup_f B_f$. \square

Recall that our definition of $\text{Del}_S(E)$ excludes edges and vertices with no incident restricted Delaunay facet. Hence there might exist points of E that are not vertices of $\text{Del}_S(E)$. In fact, this cannot happen, as stated in the following corollary of Theorem 4.12.

Corollary 4.13 *Let S be a smooth closed surface and let E be a loose ε -sample of S , with $\varepsilon \leq \varepsilon_0 \approx 0.091$. Every point of E is a vertex of $\text{Del}_S(E)$.*

Proof Let p be a point of E . By Theorem 4.12, $S \subseteq \bigcup_f B_f$, thus p belongs to the surface Delaunay ball $B_f = B(c, r)$ of some facet f of $\text{Del}_S(E)$. Let u, v and w be the vertices of f . Since B_f is a Delaunay ball, p, u, v and w belong to its bounding sphere. If $p \in \{u, v, w\}$, then f is incident to p , thus p is a vertex of $\text{Del}_S(E)$. Otherwise, (p, u, v, w) is a Delaunay tetrahedron, whose dual Voronoi vertex is $c \in S$. Then, every facet of (p, u, v, w) is a restricted Delaunay facet, and here again p is a vertex of $\text{Del}_S(E)$. \square

5 Size of loose ε -samples

5.1 Lower bound

Erickson [24] has shown that $\Omega\left(\frac{\mu(S)}{\varepsilon^2}\right)$, with $\mu(S) = \iint_S \frac{dx}{d_M^2(x)}$, is a lower bound on the number of points of any ε -sample of S , with $\varepsilon < \frac{1}{5}$. This bound holds for loose ε -samples as well, by Corollary 4.10. However, in the following we rewrite Erickson's proof in the case of loose ε -samples directly and improve on the constant.

Theorem 5.1 *Let S be a smooth closed surface and let E be a loose ε -sample of S , with $\varepsilon \leq \varepsilon_0 \approx 0.091$. We have $|E| \geq 2(1 - g) + \frac{2}{5\pi} \frac{\mu(S)}{\varepsilon^2}$, where g is the genus of S .*

Proof By Theorem 4.12, we have $S \subseteq \bigcup_{f \in \text{Del}_{|S}(E)} B_f$. Thus,

$$\mu(S) = \iint_S \frac{dx}{d_M^2(x)} \leq \sum_{f \in \text{Del}_{|S}(E)} \iint_{D_f} \frac{dx}{d_M^2(x)} \quad (10)$$

Moreover, since E is a loose ε -sample, we have $\forall f \in \text{Del}_{|S}(E), \forall x \in D_f, \|x - c_f\| \leq \varepsilon d_M(c_f)$. It follows that $\|x - c_f\| \leq \frac{\varepsilon}{1-\varepsilon} d_M(x)$ and that $d_M(x) \geq (1-\varepsilon)d_M(c_f)$, since d_M is 1-Lipschitz. Thus,

$$\forall f \in \text{Del}_{|S}(E), \iint_{D_f} \frac{dx}{d_M^2(x)} \leq \frac{\text{Area}(D_f)}{(1-\varepsilon)^2 d_M^2(c_f)} \quad (11)$$

Since $\forall x \in D_f, \|x - c_f\| \leq \varepsilon d_M(c_f)$ and $\|x - c_f\| \leq \frac{\varepsilon}{1-\varepsilon} d_M(x)$, by Lemma 3.3 we have $\forall x \in D_f, (\vec{n}(x), \vec{n}(c_f)) \leq \frac{\varepsilon}{1-4\varepsilon}$, which is less than $\frac{\pi}{2}$ since $\varepsilon \leq \varepsilon_0 < \frac{\pi}{2+4\pi}$. Thus, by Lemma A.4, D_f is a terrain over $T(c_f)$, the plane tangent to S at c_f . We can then bound the area of D_f by projecting it orthogonally onto $T(c_f)$. Let us call proj the orthogonal projection onto $T(c_f)$. Since $\text{proj}(D_f)$ is included in the disk of radius $\varepsilon d_M(c_f)$ centered at c_f , we have

$$\text{Area}(D_f) \leq \frac{\text{Area}(\text{proj}(D_f))}{\min_{x \in D_f} \cos(\vec{n}(x), \vec{n}(c_f))} \leq \frac{\pi \varepsilon^2 d_M^2(c_f)}{\cos \frac{\varepsilon}{1-4\varepsilon}} \quad (12)$$

It follows from (10), (11) and (12) that

$$\mu(S) = \iint_S \frac{dx}{d_M^2(x)} \leq \frac{\pi \varepsilon^2}{(1-\varepsilon)^2 \cos \frac{\varepsilon}{1-4\varepsilon}} m$$

where m is the number of facets of $\text{Del}_{|S}(E)$. According to Theorem 4.4, $\text{Del}_{|S}(E)$ is homeomorphic to S , thus the number of vertices of $\text{Del}_{|S}(E)$ is $2(1-g) + \frac{m}{2}$, by Euler's formula. Hence, $|E| \geq 2(1-g) + \frac{m}{2} \geq 2(1-g) + \frac{1}{2\pi} (1-\varepsilon)^2 \cos \frac{\varepsilon}{1-4\varepsilon} \frac{\mu(S)}{\varepsilon^2}$, which is at least $2(1-g) + \frac{2}{5\pi} \frac{\mu(S)}{\varepsilon^2}$ since $\varepsilon \leq \varepsilon_0$. \square

5.2 Upper bound

In this paragraph, we give an upper bound on the size of (κ, ϕ) -sparse samples, which *a priori* are not assumed to be ε -samples nor loose ε -samples. Let E be a (κ, ϕ) -sparse sample of S . According to Definition 2.4, E is *a fortiori* (κ', ϕ') -sparse, for any $\kappa' \leq \kappa$ and any $\phi' \leq \phi$. Therefore, we may assume for convenience that $\kappa < 2$ and $\phi \leq d_M$.

For every point $x \in E$, we define B_x as the open ball centered in x of radius $\frac{\kappa}{2} \phi(x)$. Since E is (κ, ϕ) -sparse, $\frac{\kappa}{2} \phi(x)$ is smaller than $\frac{1}{2} \text{dist}(x, E \setminus \{x\})$, which implies that the balls $(B_x)_{x \in E}$ are pairwise disjoint.


Lemma 5.2 *For every point $x \in E$, B_x is included in the Voronoi cell of x in $\text{Vor}(E)$.*

Proof Let z be a point of B_x . We have $\|z - x\| \leq \frac{\kappa}{2} \phi(x)$. Since E is (κ, ϕ) -sparse, every point y of $E \setminus \{x\}$ is at distance at least $\kappa \phi(x)$ from x . Thus, $\|y - x\| \geq 2\|z - x\|$, which implies that $\|z - x\| \leq \|z - y\|$, by the triangle inequality. \square

Before bounding the size of E , we prove that every ball B_x intersects S along a topological disk of large size.

Lemma 5.3 *For every point $x \in E$, we have $\text{Area}(S \cap B_x) \geq \frac{3}{16} \pi \kappa^2 \phi^2(x)$.*

Proof By definition, the radius of B_x is $\frac{\kappa}{2} \phi(x)$, which is less than $d_M(x)$ since $\kappa < 2$ and $\phi \leq d_M$. Thus, $B_x \cap M = \emptyset$. It follows that $S \cap B_x$ is a topological disk, by Lemma 3.7. It follows also that S does not intersect the open balls $B(y, \frac{\kappa}{2} d_M(x))$ and $B(z, \frac{\kappa}{2} d_M(x))$, where y and z are the intersection points of the normal of S in x with the bounding sphere of B_x . Hence, $S \cap B_x$ lies outside $B(y, \frac{\kappa}{2} d_M(x)) \cup B(z, \frac{\kappa}{2} d_M(x))$, which contains $B(y, \frac{\kappa}{2} \phi(x)) \cup B(z, \frac{\kappa}{2} \phi(x))$.



Moreover, since S has no boundary, the boundary of $S \cap B_x$ lies on the bounding sphere of B_x . Thus, $\text{proj}(S \cap B_x)$, the orthogonal projection of $S \cap B_x$ onto $T(x)$, contains the projection of $B_x \cap B(y, \frac{\kappa}{2} \phi(x))$, which is a disk of radius $\frac{\kappa\sqrt{3}}{4} \phi(x)$.

Hence,

$$\begin{aligned} \text{Area}(S \cap B_x) &\geq \text{Area}(\text{proj}(S \cap B_x)) \\ &\geq \frac{3}{16} \pi \kappa^2 \phi^2(x) \end{aligned}$$

\square

Theorem 5.4 *Let S be a smooth closed surface and let E be a (κ, ϕ) -sparse point sample of S . We have $|E| \leq \frac{64}{3\pi} \frac{1}{\kappa^2} \iint_S \frac{dx}{\phi^2(x)}$.*

Proof We proceed as in the case of planar meshes [23, 30] and bound the integral over S of $1/\phi^2(t)$. Since $S \cap B_x \subseteq S$ for every point $x \in E$, we have

$$\iint_S \frac{dt}{\phi^2(t)} \geq \iint_{\bigcup_{x \in E} (B_x \cap S)} \frac{dt}{\phi^2(t)}$$

Moreover, the balls $(B_x)_{x \in E}$ are pairwise disjoint, thus

$$\iint \bigcup_{x \in E} (B_x \cap S) \frac{dt}{\phi^2(t)} = \sum_{x \in E} \iint_{(B_x \cap S)} \frac{dt}{\phi^2(t)}$$

In addition, since ϕ is 1-Lipschitz, $\forall y \in B_x$, $\phi(y) \leq \phi(x) + \|y - x\| \leq (1 + \frac{\kappa}{2}) \phi(x)$. It follows that

$$\sum_{x \in E} \iint_{(B_x \cap S)} \frac{dy}{\phi^2(y)} \geq \sum_{x \in E} \frac{\text{Area}(B_x \cap S)}{(1 + \frac{\kappa}{2})^2 \phi^2(x)}$$

Since $\text{Area}(B_x \cap S) \geq \frac{3}{16} \pi \kappa^2 \phi^2(x)$ by Lemma 5.3, we have

$$\sum_{x \in E} \frac{\text{Area}(B_x \cap S)}{(1 + \frac{\kappa}{2})^2 \phi^2(x)} \geq \sum_{x \in E} \frac{\frac{3}{16} \pi \kappa^2 \phi^2(x)}{(1 + \frac{\kappa}{2})^2 \phi^2(x)} = \frac{3}{16 (1 + \frac{\kappa}{2})^2} \pi \kappa^2 |E|$$

which is greater than $\frac{3}{64} \pi \kappa^2 |E|$ since $\kappa < 2$. The result follows. \square

It follows from Theorem 5.4 that every loose ε -sample of S that is (κ, d_M) -sparse, for some $\kappa = \Omega(\varepsilon)$, has a size $O\left(\frac{\mu(S)}{\varepsilon^2}\right)$, which is optimal in view of the bound of Theorem 5.1. This gives a sufficient (local) condition for a loose ε -sample to have an optimal size.

Part II

Application to surface sampling and meshing

We are given a smooth closed surface S . From the previous sections (specifically from Theorems 4.4 and 4.5), we know that, if we can construct a loose ε -sample E of S , for an $\varepsilon \leq \varepsilon_0$, then $\text{Del}_{|S}(E)$ will be a good approximation of S , both in a topological and in a geometric sense. In Sections 6 and 7, we present an algorithm that produces such a sample E and computes its restricted Delaunay triangulation $\text{Del}_{|S}(E)$. Its analysis is done in Section 8. In Section 9, we present several improvements that can be made to the algorithm, either to speed it up or to enhance the quality of its output. Section 10 deals with the computation of the distance to the medial axis of the surface, which plays a central role in the context of ε -sampling. Finally, in Section 11 we present our implementation as well as some experimental results.

6 Our Algorithm

The algorithm is greedy and derives from Chew’s surface meshing algorithm [17]. It takes as input the surface S and it is parameterized by some user-defined function $\sigma : S \rightarrow \mathbb{R}$ that is positive and 1-Lipschitz. The algorithm starts with a small initial point sample E and, at each iteration, it inserts a new point of S into E and updates $\text{Del}_{|S}(E)$. Each point inserted into E is the center of a *bad* surface Delaunay ball, *i.e.* a surface Delaunay ball $B(c, r)$ such that $r > \sigma(c)$. The algorithm stops when there are no more bad surface Delaunay balls, which will eventually happen since σ does not vanish on S , as shown in Section 8, Lemma 8.1.

The surface is known only through an *oracle* that, given a line segment s , computes all the points of $s \cap S$, which generically are finitely many. $\text{Del}_{|S}(E)$ is stored as a subcomplex of $\text{Del}(E)$ and computed by detecting the intersections of the Voronoi edges with S , thanks to the oracle. At each step of the algorithm, only the part of $\text{Vor}(E)$ that has changed after the point insertion is tested. As for the bad surface Delaunay balls, they are stored in a list L .

The initial point sample is constructed in such a way that the output point set verifies Condition 2 of Definition 2.2, *i.e.* that its restricted Delaunay triangulation has vertices on all the connected components of S . The details of the construction are given in Section 7. Once the initial point sample has been constructed, the algorithm computes $\text{Del}_{|S}(E)$ and stores the bad surface Delaunay balls in L . Then, the algorithm executes the following loop:

```

while  $L$  is not empty {
  take an element  $B(c, r)$  from  $L$ ;
  insert  $c$  into  $E$  and update  $\text{Del}(E)$ ;
  update  $\text{Del}_S(E)$  by testing all the Voronoi edges that have changed or appeared:
    delete from  $\text{Del}_S(E)$  the Delaunay facets whose dual Voronoi edges no longer intersect  $S$ ;
    add to  $\text{Del}_S(E)$  the new Delaunay facets whose dual Voronoi edges intersect  $S$ ;
  update  $L$  by
    deleting all the elements of  $L$  which are no longer bad surface Delaunay balls;
    adding all the new surface Delaunay balls that are bad;
}

```

The algorithm stops at the end of the loop and returns E as well as $\text{Del}_S(E)$.

7 Construction of the initial point sample

As explained in Section 6, our purpose here is to construct, at the beginning of the algorithm, an initial point sample such that the output point set is guaranteed to verify Condition 2 of Definition 2.2. For clarity, E_i and E_o will denote respectively the initial point sample and the output point set.

Let us assume that $\text{Del}_S(E_i)$ contains a facet f_i circumscribed by a surface Delaunay ball $B_i = B(c_i, r_i)$ such that $r_i \leq \frac{1}{3}\sigma(c_i)$. Such a facet will be called a *persistent facet*. Persistent facets are interesting in our context because they share a nice persistence property, illustrated in Figure 3 and stated in the following lemma.

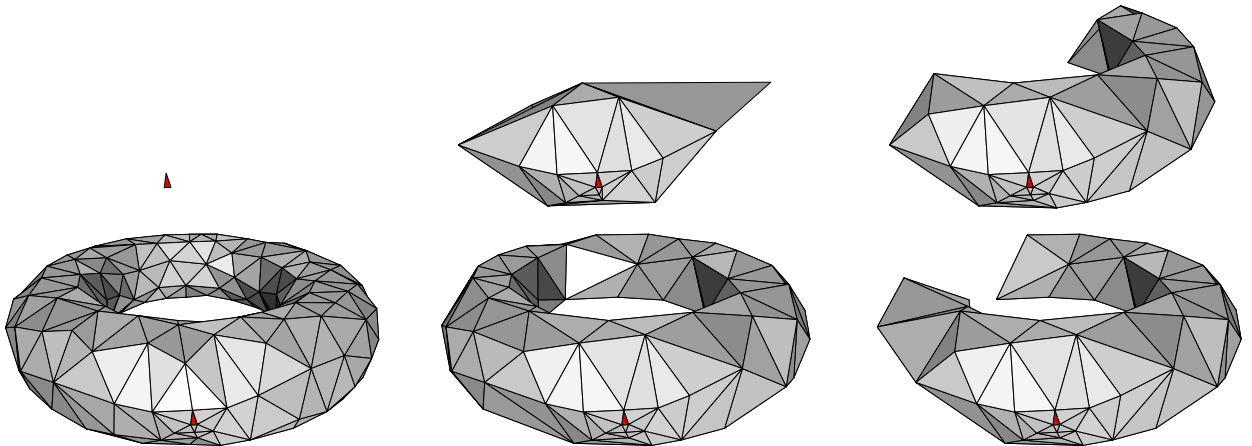


Figure 3: Meshing of a torus: the persistent facet (shown in the upper-left corner) remains in the restricted Delaunay triangulation throughout the process. In compensation, the final mesh has smaller triangles in its vicinity.

Lemma 7.1 *Every persistent facet remains a restricted Delaunay facet throughout the course of the algorithm. In other words, all persistent facets are facets of $\text{Del}_S(E_o)$.*

Proof Let f_i be a persistent facet and $B_i = B(c_i, r_i)$ a surface Delaunay ball circumscribing f_i such that $r_i \leq \frac{1}{3}\sigma(c_i)$. Assume that, at the end of the algorithm, $f_i \notin \text{Del}_{|S}(E_o)$. This implies that, at some stage, the algorithm inserts some point x in the interior of B_i . Hence, $\|x - c_i\| < \frac{1}{3}\sigma(c_i)$, which gives $\sigma(x) \geq \sigma(c_i) - \|x - c_i\| > \frac{2}{3}\sigma(c_i)$, since σ is 1-Lipschitz. Let v be one of the vertices of f_i . Since v and x both lie in B_i , we have $\|v - x\| \leq 2r_i \leq \frac{2}{3}\sigma(c_i) < \sigma(x)$. Now, since x is inserted by the algorithm, it is the center of some bad surface Delaunay ball B , whose radius is greater than $\sigma(x)$. It follows that v belongs to the interior of B , which contradicts the fact that the latter is a Delaunay ball. \square

To guarantee that E_o verifies Condition 2 of Definition 2.2, it suffices to construct E_i in such a way that each connected component of S contains at least one vertex of a persistent facet. Here is how we proceed: we pick up at least one point on each connected component of S , and we insert the collected points in a set called E'_i . For each point $x \in E'_i$, we consider a ball B_x centered at x of radius less than $\min\{\frac{1}{6}\text{dist}(x, E'_i \setminus \{x\}), d_M(x), \frac{1}{6}\sigma(x)\}$. By shooting rays repeatedly inside B_x , we pick up three points (u_x, v_x, w_x) of $S \cap B_x$ such that the line orthogonal to the plane $\text{Aff}(u_x, v_x, w_x)$ (affine hull of u_x, v_x, w_x) and passing through the center of the circumcircle of (u_x, v_x, w_x) intersects $S \cap B_x$. We insert u_x, v_x and w_x in E_i , which will be the initial point sample of the algorithm. After doing this for each point of E'_i , we have $|E_i| = 3|E'_i|$.

Lemma 7.2 *For any $x \in E'_i$, the points (u_x, v_x, w_x) picked up from $B_x \cap S$ lie on the same connected component of S as x and form a persistent facet of $\text{Del}(E_i)$.*

Proof As explained in the construction of E_i , the line orthogonal to the plane $\text{Aff}(u_x, v_x, w_x)$ and passing through the center of the circumcircle of (u_x, v_x, w_x) intersects $S \cap B_x$. Let x' be a point of intersection, and $B_{x'}$ the ball centered at x' that circumscribes (u_x, v_x, w_x) . Since x' and u_x both lie in B_x , the radius of $B_{x'}$ is at most twice the radius r_x of B_x .

In the same manner, for any $y \in E'_i \setminus \{x\}$, the radius of $B_{y'}$ is at most $2r_y$. Since r_x and r_y are both less than $\frac{1}{6}\|x - y\|$, and since $\|x' - x\| \leq r_x$ and $\|y' - y\| \leq r_y$, we have $\|x' - y'\| > \frac{2}{3}\|x - y\|$, which is greater than the sum of the radii of $B_{x'}$ and $B_{y'}$. It follows that $B_{x'} \cap B_{y'} = \emptyset$. Since this is true for any $y \in E'_i$, $B_{x'}$ is a Delaunay ball of E_i because every point of E_i belongs to some B_y , $y \in E'_i$.

In addition, the radius of $B_{x'}$ is at most $2r_x$, which is less than $\frac{1}{3}\sigma(x)$, thus (u_x, v_x, w_x) is a persistent facet of $\text{Del}(E_i)$. Moreover, u_x, v_x and w_x lie on the same connected component of S as x , since they belong to B_x , whose radius is less than $d_M(x)$. \square

Since this is true for every $x \in E'_i$, and since E'_i intersects all the connected components of S , E_i contains at least three points per connected component of S that form a persistent facet. It follows, by Lemma 7.1, that E_o verifies Condition 2 of Definition 2.2.

8 Termination and output guarantees

Lemma 8.1 *The algorithm terminates.*

Proof Notice that, since σ is Lipschitz, it is continuous. Moreover, σ does not vanish on S , which is compact; thus, there exists some constant $\sigma_0 > 0$ such that $\sigma(x) \geq \sigma_0 \forall x \in S$.

Since the construction of the initial point set takes a finite amount of time, the only thing to prove here is that the main loop of the algorithm terminates. At each iteration, the center of some bad surface Delaunay ball $B(c, r)$ is inserted. At this time, the distance from c to E is r , which is greater than $\sigma(c) \geq \sigma_0$ since $B(c, r)$ is bad. Therefore, the distance between any two points inserted during the main loop is at least $\sigma_0 > 0$, and, as a consequence, the open balls of radius $\frac{\sigma_0}{2}$, centered at the points inserted during the main loop, are pairwise disjoint. Since S is compact, there can be only a finite number of such balls. Thus, a finite number of points are inserted during the main loop of the algorithm, which terminates since it inserts one point at each iteration. \square

Lemma 8.2 *If $\sigma \leq \varepsilon d_M$, then, upon termination of the algorithm, E is a loose ε -sample of S .*

Proof When the algorithm stops, all surface Delaunay balls are good. Since a surface Delaunay ball $B(c, r)$ is good if $r \leq \sigma(c) \leq \varepsilon d_M(c)$, E verifies Condition 1 of Definition 2.2.

In addition, according to Section 7, once the initial point set has been constructed, every connected component of S contains at least one vertex of a persistent facet. Therefore, upon termination, $\text{Del}_S(E)$ has at least one vertex on each connected component of S , by Lemma 7.1, and hence E verifies Condition 2 of Definition 2.2. \square

It follows from Theorem 4.4, Theorem 4.5 and Lemma 8.2, that, if $\sigma < \varepsilon_0 d_M \approx 0.091 d_M$, then the algorithm outputs a good approximation of S , in terms of topology and geometry.

From now on, we assume that E'_i , the point set from which the initial point sample is constructed (see Section 7), contains at most c points per connected component of S , where c is a constant that does not depend on S .

Lemma 8.3 *If $\sigma \leq d_M$, then the size of the output point sample is $O\left(\iint_S \frac{dx}{\sigma^2(x)}\right)$.*

Proof For convenience, we call E_i the initial point sample constructed at the beginning of the algorithm, and E_o the output point sample. Since the algorithm is greedy, we have $E_i \subseteq E_o$. $E_o \setminus E_i$ is the set of the points that are inserted during the main loop of the algorithm. To bound the size of E_o , it suffices to bound independently the sizes of E_i and of $E_o \setminus E_i$.

As explained in Section 7, we have $|E_i| = 3 |E'_i|$. Moreover, by assumption, $|E'_i| \leq c |S|$, where $|S|$ is the number of connected components of S . Thus, $|E_i| \leq 3c |S|$.

Claim 8.3.1 $|S| \leq \frac{32}{3\pi} \mu(S)$, where $\mu(S) = \iint_S \frac{dx}{d_M^2(x)}$.

Proof For each connected component S_i of S , we define δ_i as the diameter of S_i . Let $a_i, b_i \in S_i$ such that $\|a_i - b_i\| = \delta_i$. Let $E_S = \bigcup_i \{a_i, b_i\}$. We have $|E_S| = 2 |S|$.

Since the balls $B(a_i, \delta_i)$ and $B(b_i, \delta_i)$ both contain S_i , we have $\|a_i - b_i\| = \delta_i \geq \max\{d_M(a_i), d_M(b_i)\}$, by Lemma 3.7. Moreover, for any $j \neq i$, $\|a_i - a_j\| \geq d_M(a_i)$ since a_i and a_j lie on different connected

components of S . It follows that E_S is (κ, d_M) -sparse, with $\kappa = 1$. Hence, $|E_S| \leq \frac{64}{3\pi} \mu(S)$, by Theorem 5.4. This proves the claim. \square

It follows from Claim 8.3.1 that $|E_i| \leq \frac{32c}{\pi} \iint_S \frac{dx}{d_M^2(x)} = O\left(\iint_S \frac{dx}{\sigma^2(x)}\right)$.

We now bound the size of $E_o \setminus E_i$. Every point x inserted during the main loop of the algorithm is at distance at least $\sigma(x)$ from E at the time when it is inserted, since it is the center of some bad surface Delaunay ball. It follows that, for any two points x and y of $E_o \setminus E_i$, we have $\|x - y\| \geq \sigma(x)$ or $\|x - y\| \geq \sigma(y)$, depending on whether x is inserted last or not. In both cases, we have $\|x - y\| \geq \frac{1}{2} \sigma(x)$, since σ is 1-Lipschitz. Thus, $E_o \setminus E_i$ is (κ, σ) -sparse, with $\kappa = \frac{1}{2}$. It follows that $|E_o \setminus E_i| = O\left(\iint_S \frac{dx}{\sigma^2(x)}\right)$, by Theorem 5.4. This ends the proof of Lemma 8.3. \square

The next result follows from Lemmas 8.2 and 8.3. It gives an upper bound on the size of the output point sample, which is optimal in view of Theorem 5.1.

Lemma 8.4 *If $\sigma = \varepsilon d_M$, with $\varepsilon \leq 1$, then, upon termination of the algorithm, E is a loose ε -sample of S of size $O\left(\frac{\mu(S)}{\varepsilon^2}\right)$, where $\mu(S) = \iint_S \frac{dx}{d_M^2(x)}$.*

9 Improvements

In this section, we introduce several modifications that can be made to the algorithm to enhance the aspect ratio of the facets of the output mesh (Section 9.1) or to speed up the process (Section 9.2). We also present a way to avoid the use of persistent facets in the construction of the initial point sample (Section 9.3). This latter change comes with no guarantee on the size of the output point sample, but still with guarantees on the termination of the algorithm and on the topological and geometric properties of its output mesh.

9.1 Removing the skinny facets

Once an ε -sample E_ε of S has been obtained (or is given), one can remove the skinny facets from the mesh, simply by running the algorithm with E_ε as the initial point sample and by using a new definition of a bad surface Delaunay ball. From now on and until the end of Section 9.1, a surface Delaunay ball is said *bad* if the restricted Delaunay facet it circumscribes is skinny, *i.e.* one of its inner angles is less than a user-defined value θ . The output point sample \bar{E}_ε of the algorithm contains E_ε and hence is an ε -sample of S . Moreover, all the facets of $\text{Del}_S(\bar{E}_\varepsilon)$ are well-shaped, *i.e.* no angle is less than θ . Some results are reported in Figure 9.

For every point x of \bar{E}_ε , we call $r(x)$ the distance from x to the point sample E (maintained by the algorithm) right before its insertion. If $x \in E_\varepsilon$, then $r(x) = \text{dist}(x, E_\varepsilon \setminus \{x\})$. In addition, we call *local feature size of x* , or simply $\text{lfs}(x)$, the distance from x to its second nearest neighbor in E_ε . If $x \in E_\varepsilon$, then we have $\text{lfs}(x) = \text{dist}(x, E_\varepsilon \setminus \{x\}) = r(x)$. It is proved in [30] that lfs is 1-Lipschitz and does not vanish anywhere. From now on, we assume that $\theta < \frac{\pi}{6}$.

Lemma 9.1 $\forall x \in \bar{E}_\varepsilon$, we have $r(x) \geq \frac{\beta-1}{\beta} \text{lfs}(x)$, where $\beta = \frac{1}{2 \sin \theta} > 1$.

Proof We reason by induction. We first prove that the result holds for every point of E_ε , then we show that it holds also for every point inserted by the algorithm.

- Let $x \in E_\varepsilon$. By definition we have $r(x) = \text{lfs}(x)$, which is greater than $\frac{\beta-1}{\beta} \text{lfs}(x)$ since $\beta > 1$.
- Let $x \in \bar{E}_\varepsilon \setminus E_\varepsilon$. When x is inserted in E , it is the center of a ball of $\text{Del}_S(E)$ that circumscribes a facet f whose smallest inner angle θ_f is less than θ . Let e be the smallest edge of f , and v the vertex of e that was inserted last (if both vertices of e belong to E_ε , then, without loss of generality, we choose v to be either one of them). A quick computation shows that the sine of the smallest inner angle θ_f of f (which is opposite to e) is greater than half the ratio between $|e|$ and the circumradius of f . Therefore, we have $\sin \theta \geq \frac{1}{2} \frac{r(v)}{r(x)}$, since $|e| \geq r(v)$ and since the circumradius of f is at most $r(x)$. It follows that $r(x) \geq \beta r(v)$.

Now, according to the induction hypothesis, we have $r(v) \geq \frac{\beta-1}{\beta} \text{lfs}(v)$, which is at least $\frac{\beta-1}{\beta} (\text{lfs}(x) - \|x - v\|)$ since lfs is 1-Lipschitz. Hence, $r(x) \geq \beta \frac{\beta-1}{\beta} (\text{lfs}(x) - r(x))$, i.e. $r(x) \geq \frac{\beta-1}{\beta} \text{lfs}(x)$. \square

Corollary 9.2 \bar{E}_ε is $\left(\frac{\beta-1}{2\beta-1}, \text{lfs}\right)$ -sparse.

Proof Let x and y be two points of \bar{E}_ε . If x has been inserted after y by the algorithm, or if x and y both belong to E_ε , then we have $\|x - y\| \geq r(x)$, which is at least $\frac{\beta-1}{\beta} \text{lfs}(x)$, by Lemma 9.1. Thus $\|x - y\| \geq \frac{\beta-1}{2\beta-1} \text{lfs}(x)$, since $\beta > 1$. Otherwise, y has been inserted after x by the algorithm, hence $\|x - y\| \geq r(y)$, which is at least $\frac{\beta-1}{\beta} \text{lfs}(y)$, by Lemma 9.1. Since lfs is 1-Lipschitz, we have $\frac{\beta-1}{\beta} \text{lfs}(y) \geq \frac{\beta-1}{\beta} (\text{lfs}(x) - \|x - y\|)$. Thus, $\|x - y\| \geq \frac{\beta-1}{2\beta-1} \text{lfs}(x)$. The result follows, since this is true for any pair (x, y) of points of \bar{E}_ε . \square

It follows from Corollary 9.2 and Theorem 5.4 that¹

$$|\bar{E}_\varepsilon| = O\left(\frac{1}{(\beta-1)^2} \iint_S \frac{dx}{\text{lfs}^2(x)}\right) \quad (13)$$

In particular, the algorithm terminates, since lfs does not vanish anywhere and since the algorithm inserts one point per iteration. Moreover, if E_ε is a sparse ε -sample of S , then by definition we have $\forall x \in E_\varepsilon$, $\text{lfs}(x) = \Omega(\varepsilon)d_M(x)$, which implies that $\forall x \in S$, $\text{lfs}(x) = \Omega(\varepsilon)d_M(x)$. It follows that $|\bar{E}_\varepsilon| = O\left(\frac{1}{(\beta-1)^2} \frac{\mu(S)}{\varepsilon^2}\right)$, by (13). This quantity is $O\left(\frac{1}{(\beta-1)^2} |E_\varepsilon|\right)$, by Theorem 5.1, hence the algorithm for skinny facets increases the amount of points by a factor of $O\left(\frac{1}{(\beta-1)^2}\right)$.

9.2 Speeding up the oracle

As stated in Section 6, the algorithm relies on an oracle that is able to compute all the intersection points between a line segment and the surface. This oracle is called the *intersection oracle* in the

¹Notice that Theorem 5.4 assumes that $\text{lfs} \leq d_M$, which is the case here since E_ε is an ε -sample.

sequel. In practice, *e.g.* for implicit surfaces, computing all the intersection points can be quite time-consuming. We can reduce the requirements for the oracle as follows: given a line segment s , the oracle is now supposed to be only able to:

- compute the parity of the number of transversal intersections between s and S
- find one point of $s \cap S$ when s intersects S transversally an odd number of times.

On several types of surfaces, in particular implicit surfaces, the two above operations are much easier to perform than the former ones – see Section 11.1 for more details. However, the new oracle knows that a given Voronoi edge e intersects S only when e intersects S transversally an odd number of times. Such an edge is called *bipolar*, because its two endpoints lie in different connected components of $\mathbb{R}^3 \setminus S$. Its dual Delaunay facet is called a *bipolar Delaunay facet*. The subcomplex of $\text{Del}_S(E)$ made of the bipolar Delaunay facets is called the *bipolar Delaunay triangulation* and denoted by $\text{Del}_S^b(E)$. The new oracle, which discriminates only the bipolar elements among the set of all Voronoi edges, is called *bipolar oracle*, as opposed to the intersection oracle which detects all restricted Delaunay edges. A Delaunay ball centered at some point computed by the bipolar oracle is called a *bipolar ball*.

To build the initial point sample, we perform the construction described in Section 7, which provides us with one persistent facet per connected component of S . However, some of the persistent facets may not be bipolar at this stage. Let f be one of them. We insert in E the centers of all the surface Delaunay balls of f , except the ones whose center c and radius r verify $r \leq \frac{1}{3} \sigma(c)$. We do this for all the persistent facets, and we take the result as the initial point sample of the algorithm.

Since $\text{Del}_S^b(E)$ is a subcomplex of $\text{Del}_S(E)$, the proof of Lemma 8.1 holds when the algorithm is run with the bipolar oracle instead of the intersection oracle. Hence, the algorithm terminates.

Lemma 9.3 *If $\sigma \leq \varepsilon_0 d_M$, then, upon termination of the algorithm, $\text{Del}_S^b(E)$ has the topological and geometric properties stated in Section 4 for $\text{Del}_S(E)$.*

The proof relies on some results of Sections 3 and 4. Hence, it can be skipped in a first reading.

Proof It suffices to check that $\text{Del}_S^b(E)$ verifies Assertions H0-H4 of Section 4.2.

Proof of H0

Since bipolar balls are surface Delaunay balls and since they have small radii upon termination of the algorithm, all the results of Section 3 that are stated for surface Delaunay balls and restricted Delaunay facets hold here for bipolar balls and bipolar facets.

The proof of Proposition 4.1 does not hold since it relies on the observation that, when all the surface Delaunay balls are small, a Voronoi edge cannot intersect S more than once, which is not true when only bipolar balls are small. However, it suffices to change the sentences *Moreover, by Lemma 3.6, each edge of ∂e^* is intersected at most once by S . Thus, S intersects an even number of edges of ∂e^** , by the following one: *As a consequence, the number of edges of ∂e^* that intersect*

S transversally an odd number of times is even, to make the proof of Proposition 4.1 hold with $\text{Del}_{|S}^b(E)$.

In addition, the proof of Proposition 4.2 relies only on the results of Section 3 and on Proposition 4.1, thus it holds here with $\text{Del}_{|S}^b(E)$. It follows that Theorem 4.3 is true with $\text{Del}_{|S}^b(E)$, which means that the latter verifies Assertion (H0). \square

Proof of H1

Since $\sigma < \frac{1}{7} d_M$, Lemma 3.6 says that every persistent facet f is circumscribed by exactly one ball $B_f = B(c_f, r_f)$ such that $r_f \leq \frac{1}{3} \sigma(c_f)$, and that every other ball $B(c, r)$ circumscribing f is such that $r > \sigma(c)$. Hence, after the construction of the initial point sample, f is circumscribed by at most one surface Delaunay ball, namely B_f . Moreover, every point inserted in E during the construction of the initial point set is farther than $\sigma(x)$ from E . Thus, by Lemma 7.1, B_f remains a surface Delaunay ball after the construction of the initial point sample. As a consequence, f is bipolar since B_f is the only surface Delaunay ball of f .

Throughout the rest of the course of the algorithm, every point x inserted in E is farther than $\sigma(x)$ from E , hence B_f remains a surface Delaunay ball, by Lemma 7.1. Therefore, f remains bipolar since B_f is the only surface Delaunay ball of f .

Since there is one persistent facet on each connected component of S , $\text{Del}_{|S}^b(E)$ verifies (H1) upon termination of the algorithm. \square

Proof of H2, H3 and H4

The proofs of (H2), (H3) and (H4) detailed in Section 4.2 hold here for $\text{Del}_{|S}^b(E)$ since they rely exclusively on (H0) and on Lemmas 3.1 and 3.4. \square

So, $\text{Del}_{|S}^b(E)$ verifies Assertions H0-H4 of Section 4.2 upon termination of the algorithm. As a consequence, the properties stated for $\text{Del}_{|S}(E)$ in Section 4 hold here for $\text{Del}_{|S}^b(E)$, which ends the proof of Lemma 9.3. \square

Among the results of Section 4, Corollary 4.10 says that E is a $\varepsilon(1 + 8.5 \varepsilon)$ -sample of S . If $\varepsilon < \varepsilon_5 \approx 0.065$, then $\varepsilon(1 + 8.5 \varepsilon) < 0.1$ and hence $\text{Del}_{|S}(E)$ is homeomorphic to S , by Theorem 2 of [3]. This implies that $\text{Del}_{|S}^b(E) = \text{Del}_{|S}(E)$, since $\text{Del}_{|S}^b(E)$ and $\text{Del}_{|S}(E)$ are homeomorphic triangulated surfaces without boundary, and since $\text{Del}_{|S}^b(E)$ is a subcomplex of $\text{Del}_{|S}(E)$.

In conclusion, replacing the intersection oracle by the bipolar oracle implies that the algorithm works with $\text{Del}_{|S}^b(E)$ instead of $\text{Del}_{|S}(E)$. If $\sigma \leq \varepsilon_0 d_M \approx 0.091 d_M$, then upon termination of the algorithm $\text{Del}_{|S}^b(E)$ has all the properties stated in Section 4 for $\text{Del}_{|S}(E)$. If $\sigma < \varepsilon_5 d_M \approx 0.065 d_M$, then $\text{Del}_{|S}(E)$ and $\text{Del}_{|S}^b(E)$ are equal. As explained in Section 11.1 and illustrated in Table 1, the bipolar oracle is much more efficient than the intersection oracle when run on implicit surfaces.

9.3 Getting rid of persistent facets

By definition, the radius of the surface Delaunay ball of a persistent facet is three times as small as the radius of a standard surface Delaunay ball. Therefore, the output mesh of the algorithm has smaller triangles in the vicinity of persistent facets, which may not be satisfactory. An illustration of this phenomenon is given in Figure 3.

To avoid the construction of persistent facets during the initialization phase of the algorithm, we use the set E'_i described in Section 7 as the initial point sample, instead of E_i , and, inside the main loop of the algorithm, we use the elements of E'_i as control points to check whether $\text{Del}_{|S}(E)$ has vertices on all the connected components of S (Condition 2 of Definition 2.2). Specifically, if all the points of E'_i are vertices of $\text{Del}_{|S}(E)$, then $\text{Del}_{|S}(E)$ verifies Condition 2 of Definition 2.2 since E'_i intersects all the connected components of S . Otherwise, we cannot decide whether $\text{Del}_{|S}(E)$ verifies that condition or not, because we do not know on which connected components of S the vertices of $\text{Del}_{|S}(E)$ lie. Hence, we insert additional points of S in E and go on running the main loop of the algorithm until all surface Delaunay balls are good and all the points in E'_i have become vertices of $\text{Del}_{|S}(E)$. The additional points of S are computed by casting rays in random directions from the points of E'_i that are not yet vertices of $\text{Del}_{|S}(E)$. Here is the modified version of the main loop:

```

while  $L$  is not empty or some point  $x$  of  $E'_i$  is not a vertex of  $\text{Del}_{|S}(E)$  {
  if  $L \neq \emptyset$  { // proceed as before
    take an element  $B(c, r)$  from  $L$ ;
    insert  $c$  into  $E$  and update  $\text{Del}(E)$ ;
    update  $\text{Del}_{|S}(E)$ ;
    update  $L$ ;
  }
  else { //  $x \in E'_i$  is not a vertex of  $\text{Del}_{|S}(E)$ 
    choose a random vector  $\vec{v}$ ;
    insert in  $E$  the first point of intersection of ray  $(x, \vec{v})$  with  $S$  (if it exists);
    insert in  $E$  the first point of intersection of ray  $(x, -\vec{v})$  with  $S$  (if it exists);
    update  $\text{Del}_{|S}(E)$ ;
    update  $L$ ;
  }
}

```

Lemma 9.4 shows that this process terminates. The proof relies on the fact that, in the “worst-case” scenario, the algorithm creates persistent facets on all the connected components of S , which implies that $\text{Del}_{|S}(E)$ eventually verifies Condition 2 of Definition 2.2.

Lemma 9.4 *The modified algorithm terminates, provided that $\sigma \leq \varepsilon_0 d_M$.*

Proof Since the initial point sample, E'_i , is clearly computed in a finite amount of time, all we have to prove is that the main loop of the algorithm terminates. Let E_a denote the set of all points inserted by the main loop. Let E_c be the set of all inserted centers of bad surface Delaunay balls.

Since σ is continuous and does not vanish on S , which is compact, E_c is finite by the same packing argument as in the proof of Lemma 8.1. Let us bound the size of $E_a \setminus E_c$, which will complete the proof of the lemma since the main loop of the algorithm inserts one point per iteration.

Let $x \in E'_i$. We call S_x the connected component of S on which x lies. On one side of $T(x)$, all vectors point towards the interior of the object bounded by S . Therefore, any ray cast on this side of $T(x)$ from x is such that its first intersection point with S lies on S_x . Moreover, its first transversal intersection point with S lies on S_x as well. Hence, whatever oracle is used, by casting such a ray we find a point of S_x . It follows that, if at some stage of its execution the algorithm detects that x is not a vertex of $\text{Del}_S(E)$, then it inserts at least one point of S_x in E , since the random vector \vec{v} it chooses does not belong to $T(x)$ generically. After a finite number of iterations of the main loop, either x has become and will remain a vertex of $\text{Del}_S(E)$, or there are enough points of S_x in E to guarantee that at least three of them are close enough to one another to form a persistent facet, since S_x is compact. In both cases, $\text{Del}_S(E)$ has a vertex on S_x for the rest of the course of the algorithm. Since this is true for every $x \in E'_i$, and since E'_i intersects all the connected components of S , E verifies Condition 2 of Definition 2.2 after a finite time. At this stage, E also verifies Condition 1 of Definition 2.2 with $\varepsilon = \varepsilon_0$, since the centers of bad surface Delaunay balls are inserted first and since we took $\sigma \leq \varepsilon_0 d_M$. Therefore, all the points of E (in particular, those of E'_i) are vertices of $\text{Del}_S(E)$ by Corollary 4.13. Hence, $E_a \setminus E_c$ is finite. \square

Although we have no explicit upper bound on the size of the output, in practice it turns out that this version of the algorithm generates sparse samples with no persistent facet – see Figures 4 and 6. As an extreme, we could replace the initialization step by just computing a few random points on S by shooting along random lines of \mathbb{R}^3 . As reported in [10], this simple procedure satisfies Condition 2 of Definition 2.2 with a probability that increases dramatically with the number of initial random points.

10 Choice of function σ

Although the algorithm is conceptually simple, the theoretical guarantees hold only if σ is at most ε times d_M , the distance to the medial axis of S . For some surfaces, *e.g.* skin surfaces [15], d_M is equal to the minimum radius of curvature and can therefore be estimated locally. However, in most cases, d_M depends on the global shape of the surface and is difficult to compute. This makes the search for optimal ε -samples difficult.

A much easier quest is to compute a uniform ε -sample of S . This can be achieved by taking for σ a constant positive function $\sigma_0 \leq \varepsilon d_M^{\text{inf}}$. The only prerequisite is to compute d_M^{inf} or any smaller positive value, which is a simpler issue than computing d_M at each center of surface Delaunay ball, and which can be done once for all at the beginning of the algorithm. The way one can estimate d_M^{inf} depends highly on the nature of the surface. Strategies are discussed in Section 11, for several types of surfaces.

Once a uniform ε -sample \bar{E} has been constructed, it can be used to estimate d_M . In [12], Chazal and Lieutier introduced the notion of λ -medial axis. The λ -medial axis of \bar{E} , $M_\lambda(\bar{E})$, is a subset of $\text{Vor}(\bar{E})$. Chazal and Lieutier proved that, for values of λ greater than $10\sqrt{3/2} \text{diam}(\bar{E})^{3/4} \sqrt{\varepsilon}$,

$M_\lambda(\bar{E})$ approximates M within a semi-Hausdorff distance of $O(\varepsilon)$. It follows that d_{M_λ} , the distance to $M_\lambda(\bar{E})$, is a good estimate of d_M . Moreover, given $x \in S$, $d_{M_\lambda}(x)$ is easy to compute. Therefore, in a second stage, if one wishes to generate a sparse loose ε -sample of S , one can restart the meshing algorithm from the beginning, taking $\sigma = \varepsilon' d_{M_\lambda}$, for some ε' slightly smaller than ε . According to Lemma 8.4, the result is a size-optimal loose ε -sample of S , provided that E'_i , the point set from which the initial point sample of the algorithm is constructed (see Section 7), contains a constant number of points per connected component of S . To construct E'_i , one can simply pick up one vertex on each connected component of $\text{Del}_S(\bar{E})$. Since $\text{Del}_S(\bar{E})$ is homeomorphic to S and close to S for the Hausdorff distance, E'_i will contain exactly one point of each connected component of S .

11 Implementation and results

This section presents our implementation of the algorithm, which works with various types of surfaces: implicit surfaces, level sets in 3D-images, point set surfaces and polyhedra. It is written in C^{++} and uses the CGAL library [32], which provides a data structure for representing and manipulating the restricted Delaunay triangulation as a subcomplex of the 3D Delaunay triangulation.

We adopted the two-passes strategy described in Section 10, using the version of the algorithm described in Section 9.3, which does not use persistent facets. This way, we minimized the prerequisites on the knowledge of the surface, which are the following ones:

- P1** we can pick up at least one point from each connected component of S .
- P2** we know (or we can compute) a positive constant less than d_M^{inf} .
- P3** we can implement at least one of the two versions of the oracle.

These prerequisites are discussed in the sequel, for several types of surfaces. The purpose is to build implementations to fulfill (P1), (P2) and (P3). Notice that (P3) is used throughout the algorithm, whereas (P1) is involved only in the initialization phase and (P2) is a simple precalculation. Therefore, algorithmic issues arise mainly from (P3).

This section also presents some experimental results which show that the algorithm works well on a wide variety of input surfaces, including some piecewise smooth surfaces, for which our theoretical results do not hold since the distance to the medial axis vanishes at singularities.

11.1 Implicit surfaces

In this paragraph, S is a level set (say the zero-set) of some function f whose expression is given explicitly. In other words, we have $S = f^{-1}(\{0\})$. We assume that f is C^2 -continuous and that its gradient does not vanish on S . Then, by the implicit function theorem, S is a C^2 -differentiable 2-manifold.

P1 We compute the points of S that have a horizontal tangent plane. Each connected component of S has at least two such points, since S has no boundary. These points are the critical points of f with respect to the height function, *i.e.* the solutions of the following system:

$$\begin{cases} f(x, y, z) = 0 \\ \frac{\partial f}{\partial x}(x, y, z) = 0 \\ \frac{\partial f}{\partial y}(x, y, z) = 0 \end{cases}$$

which is generically zero-dimensional. If f is a polynomial, then the system is algebraic and can be solved by various means. Our approach consists in computing the generalized normal form modulo the ideal generated by the three polynomials of the system, and then finding the roots from eigencomputation – this method was developed in [31] and implemented in C^{++} as part of the SYNAPS library [33], which we use in our implementation. If f is not a polynomial but still continuous, then we compute the solutions of the system using interval arithmetic. Notice that this computation may be quite complicated but is invoked only once, during the initialization phase of the algorithm.

P2 Since S is compact, there exists a point $p \in S$ such that $d_M(p) = d_M^{\text{inf}}$. Let c be the point of M closest to p . We have $\text{dist}(c, S) = \|c - p\|$, since $\|c - p\| = d_M^{\text{inf}}$. Hence, p is a nearest neighbour of c on S . If c belongs to the boundary of M , then $\|c - p\|$ equals the minimum radius of curvature of S at p . Otherwise, c has another nearest neighbour on S , say q . We have $\|c - q\| = \|c - p\| = d_M^{\text{inf}}$, thus $\|c - q\| = d_M(q)$. It follows that the balls $B(p, \|c - p\|)$ and $B(q, \|c - q\|)$ are both tangent to M in c . Hence, p , c and q are collinear, since c belongs to the relative interior of M . This implies that c is the midpoint of the line segment $[p, q]$, since $p \neq q$. So, the ball $B(c, \|c - p\|)$ is tangent to S in two diametral points.

To sum up, computing d_M^{inf} can be done by finding:

1. the point of S at which the smallest radius of curvature is minimal, which reduces to solving some low-dimensional optimization problem over S .
2. the smallest sphere bitangent to S with diametral contact points, which reduces to finding the smallest real positive root of some zero-dimensional system.

These two issues can be solved using the same tools as for prerequisite (P1).

P3 The bipolar oracle is quite easy to implement for implicit surfaces. Indeed, the parity of the number of transversal intersections between S and a given line segment $[a, b]$ is given by the signs of f at a and b : if the signs are equal, then the number of transversal intersections is even, otherwise it is odd. In the latter case, we can find a transversal intersection point by binary search. The computation time is $O(1)$ for the bipolar test, and $O\left(\log \frac{\|a-b\|}{\Delta}\right)$ for finding an intersection point within a precision of $\Delta > 0$. When f is a polynomial, we use Descartes' rule instead of evaluating the signs of f at a and b , which reduces the computation time.

The intersection oracle is also simple to implement, and we did it by means of a divide-and-conquer strategy. However, computing all the intersection points of S and segment $[a, b]$ within a precision of $\Delta > 0$ takes $\Theta\left(\frac{\|a-b\|}{\Delta}\right)$ time. Hence, for implicit surfaces, the intersection oracle is far less efficient than the bipolar oracle. Timings are reported in Table 1.




Figure 4: Results on smooth closed algebraic surfaces, with $\sigma = 0.09 d_M$.

Experimental results Results on smooth closed algebraic surfaces are reported in Figure 4. The top line shows the inputs, the bottom line shows the outputs. From left to right, we have a torus, a genus-three surface of degree 4, called “chair”, and a genus-five surface of degree 4, called “tanglecube”. These surfaces have been meshed taking $0.09 d_M$ for function σ . As predicted by our theoretical results, we obtained good topological and geometric approximations of the surfaces.

Figures 12 (color section) and 5 show the results of the algorithm respectively on the standard left trefoil knot and on a more intricate knot with high self-entanglement. In both cases, the result of the algorithm is ambient isotopic to the original algebraic model, as predicted by Theorem 4.4. Notice that every knot is represented in “sausage” format, *i.e.* as the boundary of a thickening of some knotted curve. Figure 12 (left) shows that such a curve can be defined as the intersection of two algebraic surfaces, $f_1 = 0$ and $f_2 = 0$, which are the images through a stereographic projection of two 2-manifolds embedded in the unit sphere of \mathbb{R}^4 , as explained in [11]. One possible thickening of the curve $f_1 = f_2 = 0$ is $f_1^2 + f_2^2 < \delta$, whose boundary $f_1^2 + f_2^2 = \delta$ is a smooth closed surface, for sufficiently small δ . Notice that this thickening does not have a constant radius, as one can observe in Figure 12 (right).




Figure 5: An intricate knot, in “sausage” format, meshed with $\sigma = 0.09 d_M$.

Some results on piecewise smooth algebraic surfaces are presented in Figure 6. Since the distance to the medial axis vanishes at singularities, our theoretical results do not hold here. Therefore, to generate each mesh, we used a constant function σ , whose value was chosen according to the desired resolution of the output mesh, with respect to the diameter of the surface. The left column of Figure 6 contains the input models, the center column the output meshes, and the right column some close-ups of both objects. From top to bottom, we have:

- Barth’s octic surface, of degree 8 (according to its name), made of eight pillows placed at the vertices of a cube and connected along the edges of the cube by means of two singular points. Hence, there are 24 singular points in total. We took $\sigma = 0.03$ to mesh this surface, whose diameter is 4.
- a degree-six surface called “heart”, with two pinch points, one at the top and the other at the bottom. We took $\sigma = 0.05$ to mesh this surface, whose diameter is 4.
- Klein’s bottle, which is not a manifold when immersed in 3-space. Here, we took $\sigma = 0.1$, while the diameter of the surface is 8.

Figure 6 illustrates two things. First, the algorithm terminates on all kinds of compact surfaces, in accordance with the statement of Lemma 8.1. Second, the algorithm does a good job far from the singularities. We have added to our implementation a patch that checks whether $\text{Del}_{|S}(E)$ is a manifold and that goes on refining the mesh in the negative. This patch can be toggled by the user, which allows him to force the algorithm to generate a manifold. This option has been activated for the “heart” and Barth’s octic. The corresponding meshes are manifolds. Notice that, as a consequence, the singular points of Barth’s octic are missing in our triangulated version.

Timings for the above algebraic models are reported in Table 1, normalized with respect to the 3D incremental Delaunay triangulation algorithm. The precalculation of d_M^{inf} is not taken into account here. Moreover, since the meshing process works in two passes, timings are reported only for the

Surface	Intersection oracle (uses $\text{Del}_S(E)$)							Bipolar oracle (uses $\text{Del}_S^b(E)$)						
	Engine		Oracle			Total time / Del.	Output size (# pts)	Engine		Oracle			Total time / Del.	Output size (# pts)
	Time/Del.	%	# calls	Time/Del.	%			Time/Del.	%	# calls	Time/Del.	%		
Sphere	0.64	14.6	8,929	3.8	85.4	4.4	445	0.66	87.5	3,226	0.1	12.5	0.76	445
Ellipsoid	2.5	6.3	7,410	37	93.7	39.5	380	2.75	91.7	8,170	0.25	8.3	3	380
Torus	0.29	0.2	58K	158	99.8	158	1,268	0.31	88.7	63K	0.39	11.3	0.7	1,307
Chair	2.67	0.5	260K	562	99.5	565	6,619	2.7	89.4	255K	0.32	10.6	3.02	6,461
Tanglecube	2.77	2.9	168K	118	97.1	121	4,225	2.57	90.8	165K	0.27	9.2	2.84	4,242
Trefoil knot	4.18	1.6	545K	261	98.4	265	8,329	4.18	93.2	592K	0.31	6.8	4.49	8,317
Intricate knot	2.56	2.3	5.2M	110	97.7	113	133K	2.31	89.9	7M	0.26	10.1	2.57	148K
Barth's octic	2.31	1.7	547K	132	98.3	134	13,928	2.35	83.8	603K	0.45	16.2	2.8	14,168
Heart	2.24	7	252K	29.6	93	31.9	8,445	2.15	91.8	263K	0.19	8.2	2.34	8,539
Klein's bottle	0.15	3.5	101K	41.2	96.5	41.5	3,424	1.58	92	102K	0.14	8	1.72	3,445

Table 1: Timings (normalized w.r.t. 3D incremental Delaunay triangulation algorithm) and output size for various algebraic surfaces.




Figure 6: Results on non-smooth or non-manifold algebraic surfaces, with $\sigma = \sigma_0 > 0$.

second pass. We tried the two versions of the oracle: intersection oracle and bipolar oracle. For each version, we separated the time spent in the oracle from the time spent in the rest of the program (the *engine*).

The first observation is that the ratio between the timings of the two versions of the oracle is quite large. In the table, it ranges from several units for small models (*e.g.* 4 for the sphere), to several dozens for huge models (*e.g.* 40 for the intricate knot). As explained previously, this phenomenon

is due to the algorithmic structure of the oracle, which makes the running time of its standard version linear w.r.t. $1/\Delta$, where Δ is a user-defined precision threshold (10^{-6} here), while the time complexity of its bipolar version is logarithmic w.r.t. $1/\Delta$.

The second observation is that the timings of the engine are comparable to the timings of the incremental Delaunay triangulation algorithm (the ratio is never greater than a few units, and sometimes even less than 1). This observation is not surprising since both algorithms are similar. The trade-off between the performances of these algorithms is due to the oracle. With the bipolar oracle, the trade-off is small and the performances of both algorithms are comparable.

11.2 Level sets in 3D images

In this paragraph, S is a grey level in a 3D greyscaled image. Formally, it is a level set of some potential function f whose expression is *not* given explicitly. However, we can retrieve the value of $f(x)$ for any $x \in \mathbb{R}^3$ by interpolation. Our implementation uses the YAV++ library [34], which provides us with basic tools for manipulating 3D greyscaled images.

P1 We compute the points of S that have a horizontal tangent plane, by sweeping a horizontal plane vertically across the image. Since S has no boundary, this gives us at least two points per connected component of S .

P2 We use a thinning algorithm [13] to compute a discrete approximation \tilde{M} of the medial axis of S . An estimate of d_M can be easily computed from \tilde{M} , which allows us to mesh S in one pass only.

P3 We proceed as in the case of an implicit surface.

Experimental results Figure 7 shows some results of the algorithm on 3D greyscaled images. The input images are courtesy of Olivier Clatz. The first one has size $256^3 \approx 16.7M$ voxels, the second one $50 \times 60 \times 60 = 180K$ voxels. The first mesh has about 25K vertices, which is far less than the size of the output of the Marching Cubes algorithm (several millions of vertices). The advantage of our mesher over the Marching Cubes algorithm is that the user can specify through function σ , at which level of detail he wants to work. The second mesh has approximately 100K vertices, which is comparable to the size of the input image. The reason for this is that the surface is quite complicated, with a small d_M^{inf} , hence generating an ε -sample requires a lot of points.

11.3 Point set surfaces

Here, S is known only through an unorganized finite point set E . The points of E are assumed to be close to S , though they may not lie on S exactly. We use our algorithm to generate a new point sample. The original one is used only for fulfilling prerequisites (P1), (P2) and (P3).




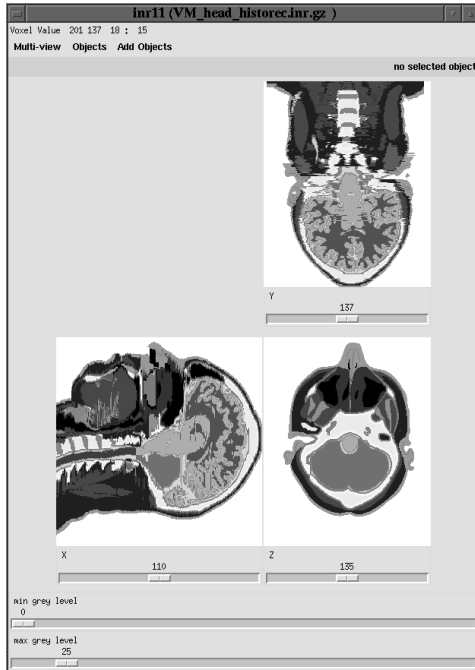


Figure 7: Results on 3D greyscaled images, with $\sigma = 0.09 d_M$.

P3 Since S is known only through E , it is impossible to compute new points of S . Therefore, to fulfill prerequisite (P3), we have to derive from E some local approximation \tilde{S} of S . Then, our algorithm will sample \tilde{S} instead of S . Such a strategy has been proposed in [1, 8, 25]. In [8],

\tilde{S} is a smooth closed surface, and our results imply that the meshing algorithm outputs a good approximation of \tilde{S} . In [1], \tilde{S} is also a smooth closed surface, provided that E is a *good* sample of S in some sense. Notice that we can choose \tilde{S} freely. In particular, \tilde{S} does not need to pass through the points of E .

P1 If E is a *good* sample of S , in the sense of [1] or in the sense of [20], then we can compute an approximation of the tangent plane of S in each point of E . We pick up the points of E whose tangent plane is almost horizontal. Since S has no boundary, this gives us at least two points per connected component of S .

P2 As explained in Section 10, we can approximate the distance to the medial axis of \tilde{S} with the distance to the λ -medial axis of E . Using this estimate, we can sample \tilde{S} in one pass only, without computing d_M^{inf} .

Experimental results Figure 8 shows the results of the algorithm on a noisy point set of $42K$ points – upper-left corner. In the upper-right corner is the output of our algorithm when used with the interpolating surface of [7]. In the bottom-right corner is the result of our algorithm when used with the approximating surface of [1] – shown in the bottom-left corner. Both outputs have the right topology type. However, the interpolating surface leads to a wrong approximation of the normals of the original surface², due to the presence of noise. This is not the case with the approximating surface.

11.4 Polyhedra

It is trivial to fulfill (P1) when S is a polyhedron. Moreover, we can forget about (P2) because, as a choice for σ , we take a constant function that has no relationship with d_M , since S contains singular points. As for (P3), it can be satisfied by means of a naive procedure which, given a line segment s , checks the intersection of s with each facet of S . The segment-surface-intersection test is then performed in linear time, with respect to the number of facets of S . Our implementation of the oracle uses some code developed by M. Samozino from an original idea by P. Bhattacharya [6]. This code uses an octree to speed-up the segment-surface-intersection test.

Experimental results Figure 10 shows some results on polyhedra. One can notice that the algorithm does a good job far from the sharp edges. When all dihedral angles of S are larger than $\frac{\pi}{2}$, the algorithm behaves as if the surface were smooth. It is the case of the horse model, from which the algorithm generated a manifold with the right topology type. However, in the vicinity of sharp edges, the output mesh may have cracks or holes, or it may even not be a manifold. An illustration is given in the close-up of the octopus, in which the mesh is shown with Gouraud shading. The consequence of the non-orientability of the mesh is that dark stains appear in places where the normal is oriented in the wrong direction. In these places, the surface is not locally a manifold.

²Ondulations are clearly visible in the upper-right corner of Figure 8.

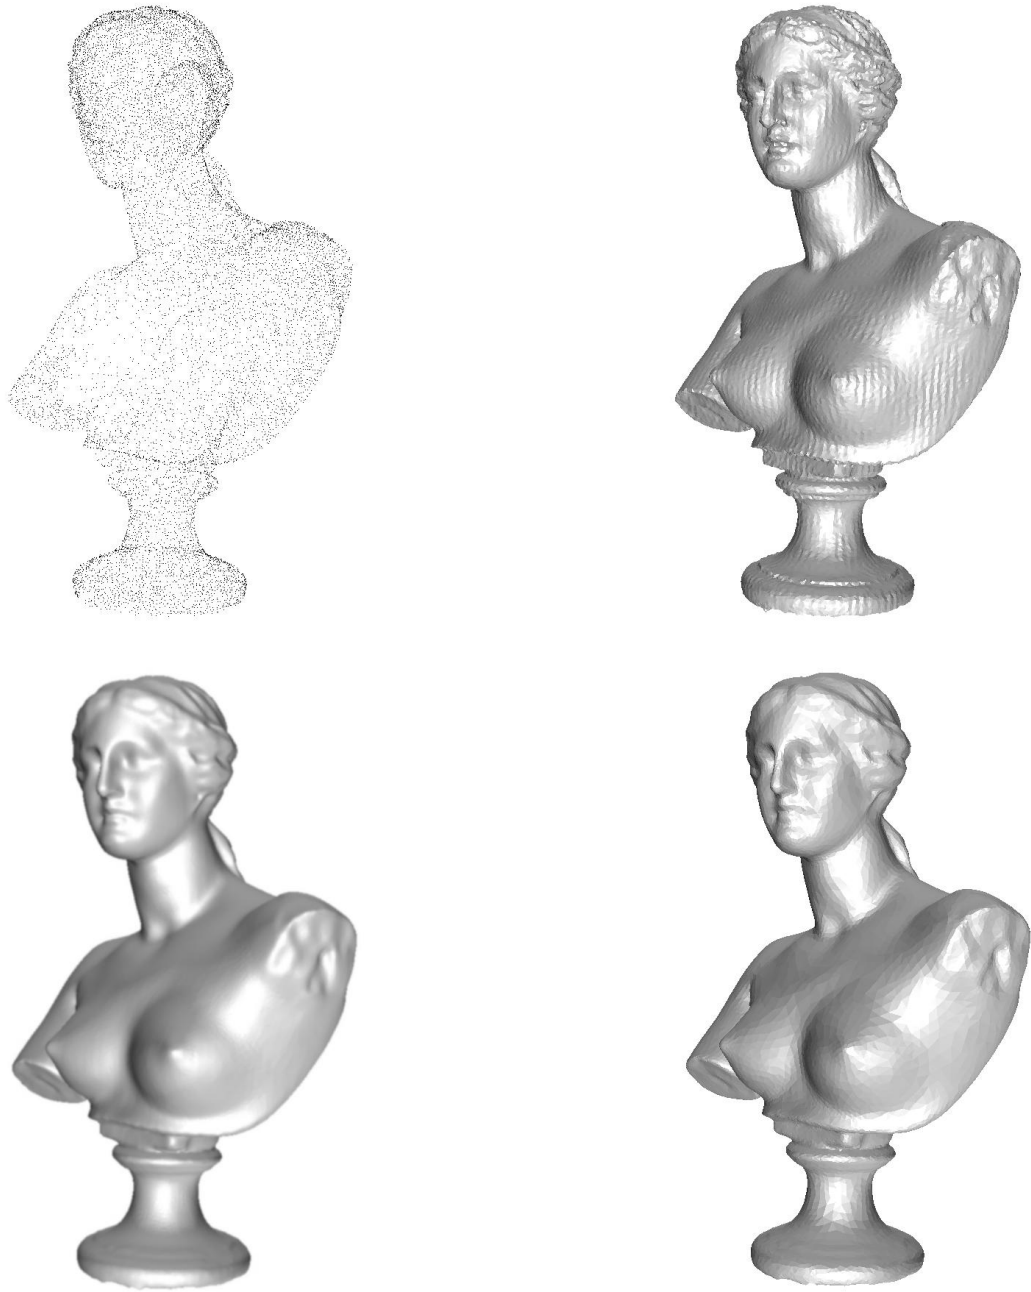


Figure 8: Results on a point set surface, with $\sigma = 0.09 d_M$.

Figure 9 shows the result of the skinny facet removal procedure on the triceratops model. As predicted in Section 9.1, the procedure terminated and removed all triangles with angles less than $\frac{\pi}{6}$.

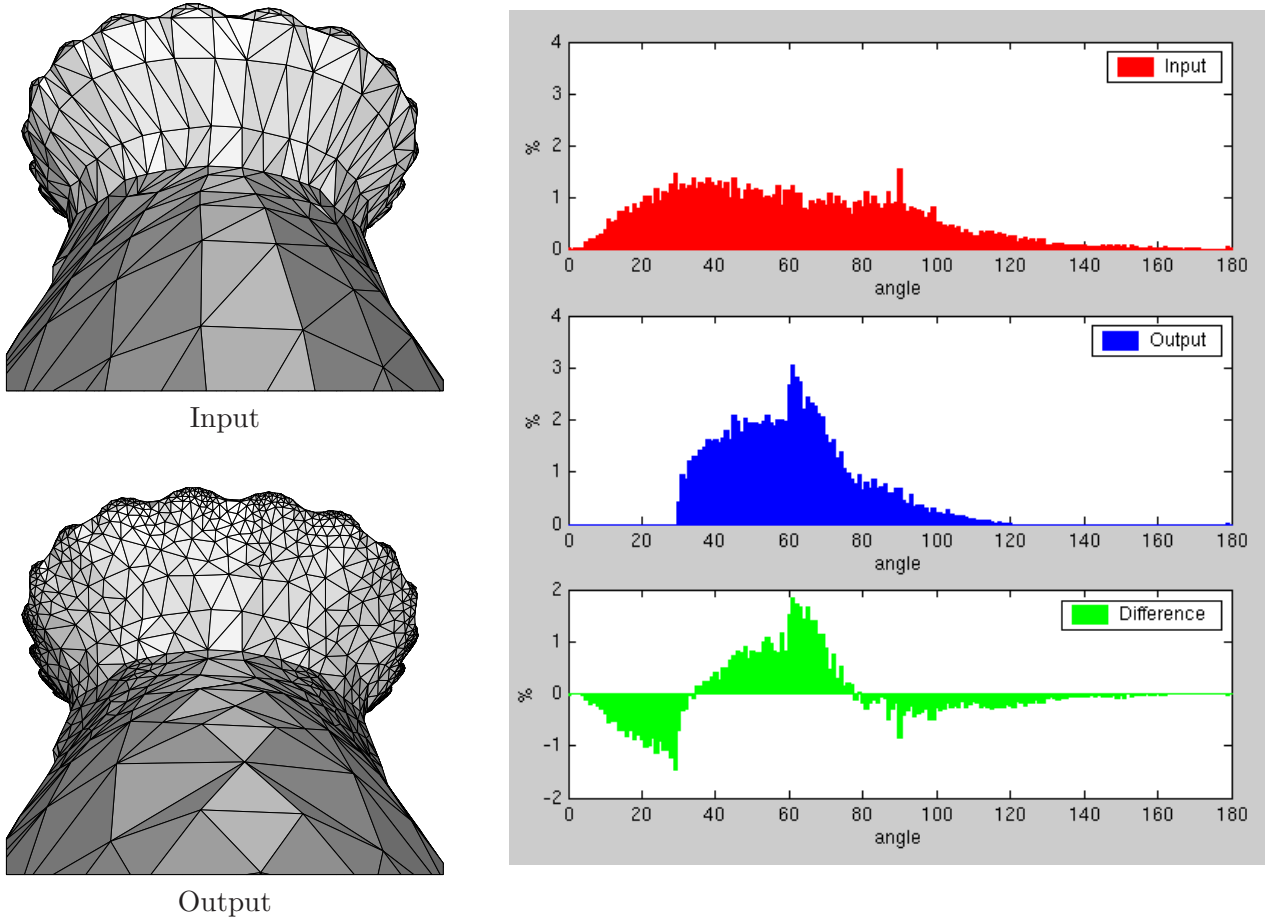


Figure 9: Removing the skinny facets (triceratops model).

12 Conclusion

We have introduced a new notion of surface sample, the so-called loose ε -samples. We have shown that loose ε -samples are $\varepsilon(1 + 8.5\varepsilon)$ -samples and share the main properties of ε -samples. Checking if a sample E of a surface S is a loose ε -sample reduces to comparing the radii of the surface Delaunay balls with the distances of their centers to the medial axis of S . Hence we obtain a new sufficient condition for sampling a surface with topological and geometric properties. This condition is similar in spirit to other sampling conditions [3, 4, 22]. An important advantage of our condition is that it leads to a simple and provably correct algorithm to sample and mesh surfaces, which has several advantages over the Marching Cubes algorithm [16, 26]: topological guarantees, approximation of normals, facets with bounded aspect ratio, optimal size.

This paper has only considered the case of smooth closed surfaces. We plan to extend our work to surfaces with boundaries and to piecewise smooth surfaces. Experimental results have shown that the algorithm is robust and can produce good geometric approximations of surfaces with




Figure 10: Results on polyhedra, with $\sigma = \sigma_0 > 0$.

singularities.

Our approach can be used for curves in any dimension: extending the proofs of this paper is not difficult, and in fact the proofs are simpler. Further research is needed to extend this work to manifolds of dimension more than one embedded in spaces of higher dimensions.

Since the submission of the conference version of this paper, several other certified algorithms have been proposed to mesh smooth closed implicit surfaces. Plantinga and Vegter [29] proposed

a variant of the Marching Cubes algorithm with topological guarantees. However, they have no optimality result, and no guarantee on the aspect ratio of the facets. Cheng et al. [14] proposed another Delaunay refinement algorithm for meshing implicit surfaces. Their algorithm assumes to be able to compute critical points of height functions defined over S or over intersections of S with planes, which seems to be a hard task except for simple cases (when the implicit function is a polynomial of low degree). Boissonnat et al. [9] proposed a meshing algorithm that offers topological guarantees based on some results from Morse theory. Their method requires to compute the indices of critical points and has not been implemented yet.

Acknowledgment

This work has been partially supported by the IST Programme of the EU as a Shared-cost RTD (FET Open) Project under Contract No IST-2000-26473 (ECG - Effective Computational Geometry for Curves and Surfaces).

References

- [1] Anders Adamson and Marc Alexa. Approximating and Intersecting Surfaces from Points. *Proc. 1st Annu. Sympos. on Geometry Processing*, pp 230-239, 2003.
- [2] Marc Alexa, Johannes Behr, Daniel Cohen-Or, Shachar Fleishman, David Levin and Claudio T. Silva. Point Set Surfaces. *Proc. 12th IEEE Visualization*, 2001.
- [3] N. Amenta and M. Bern. Surface Reconstruction by Voronoi Filtering. *Discrete and Computational Geometry*, 22:481-504, 1999.
- [4] N. Amenta, S. Choi, T. K. Dey and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. *Internat. Journal of Comput. Geom. and Applications*, 12:(1-2), pp 125-141, 2002.
- [5] N. Amenta, T.J. Peters and A.C. Russell. Computational Topology: Ambient Isotopic Approximation of 2-Manifolds. *Theoretical Computer Science*, 305:1-3, pp 3-15, 2003.
- [6] P. Bhattacharya. *Efficient Neighbor Finding Algorithms in Quadtree and Octree*. Master's thesis, Indian Institute of Technology, Kampur, 2001.
- [7] J-D Boissonnat and F. Cazals. Natural Neighbour Coordinates of Points on a Surface. *Comp. Geom. Theory and Appl.*, 155-174, 2001.
- [8] J-D Boissonnat and F. Cazals. Smooth Surface Reconstruction via Natural Neighbour Interpolation of Distance Functions. *Comp. Geom. Theory and Appl.*, 185-203, 2002.
- [9] J-D Boissonnat, D. Cohen-Steiner and G. Vegter. Isotopic Implicit Surface Meshing. *Proc. 36th annu. Sympos. on Theory of computation*, pp 301-309, 2004.

- [10] J-D Boissonnat and S. Oudot. Provably Good Surface Sampling and Approximation. *Proc. 1st Annu. Sympos. on Geometry Processing*, pp 9-19, 2003.
- [11] E. Brieskorn and H. Knörrer. *Plane Algebraic Curves*. Birkhäuser, Basel Boston Stuttgart, 1986
- [12] F. Chazal, A. Lieutier. The “ λ -medial axis”. *Graphical Models*, 2005. To appear.
- [13] Y-S Chen, W-H Hsu. Parallel Thinning Algorithm for Binary Digital Patterns. *Handbook of Pattern Recognition and Computer Vision*, World Scientific, pp 458-461, 1993.
- [14] S-W Cheng, T. K. Dey, E. A. Ramos and T. Ray. Sampling and meshing a surface with guaranteed topology and geometry. *Proc. 20th Annu. Sympos. on Comput. Geom.*, pp 280-289, 2004.
- [15] H-L Cheng, T. K. Dey, H. Edelsbrunner and J. Sullivan. Dynamic skin triangulation. *Discrete and Computational Geometry*, 25(4):525-568, 2001.
- [16] E. V. Chernyaev. Marching Cubes 33: Construction of Topologically Correct Isosurfaces. Technical report CERN CN 95-17, 1995.
- [17] L. P. Chew. Guaranteed-Quality Mesh Generation for Curved Surfaces. *Proc. 9th Annu. ACM Sympos. Comput. Geom.*, pp 274-280, 1993.
- [18] D. Cohen-Steiner and J-M Morvan. Restricted Delaunay Triangulations and Normal Cycle. *Proc. 19th Annu. ACM Sympos. Comput. Geom.*, 2003.
- [19] T. K. Dey and S. Goswami. Tight Cocone: A water tight surface reconstructor. *Proc. 8th ACM Sympos. Solid Modeling Appl.*, pp 127-134, 2003.
- [20] T. K. Dey and S. Goswami. Provable surface reconstruction from noisy samples. *Proc. 20th Annu. Sympos. Comput. Geom.*, pp 330–339, 2004.
- [21] M. P. Do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice Hall, 1976, p. 142.
- [22] H. Edelsbrunner, N.R. Shah. Triangulating topological spaces. *International Journal of Computational Geometry and Applications*, Vol. 7, No. 4 (1997), pp 365-378.
- [23] H. Edelsbrunner. *Geometry and Topology for Mesh Generation*. Cambridge University Press, pp 39-43, 2001.
- [24] J. Erickson. Nice point sets can have nasty Delaunay Triangulations. *Discrete and Computational Geometry*, 30(1):109-132, 2003.
- [25] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald and W. Stuetzle. Mesh Optimization. *Proc. SIGGRAPH*, pp 19-26, 1993.
- [26] W. E. Lorensen and H. E. Cline. Marching Cubes: a High-Resolution 3D Surface Construction Algorithm. *Proc. SIGGRAPH*, 1987.
- [27] B. Mederos, L. Velho and L. H. de Figueiredo. Moving Least Squares Multiresolution Surface Approximation. *Proc. SIBGRAPI*, 2003.

- [28] J-M Morvan and B. Thibert. On the Approximation of the Area of a Surface. INRIA RR-4375, 2002.
- [29] S. PLantinga and G. Vegter. Isotopic Approximation of Implicit Curves and Surfaces. *Proc. 2nd Sympos. on Geometry Processing*, pp 251-260, 2004.
- [30] J. Ruppert. A Delaunay Refinement Algorithm for Quality 2-Dimensional Mesh Generation. *Journal of Algorithms*, May 1995.
- [31] P. Trébuchet. *Vers une résolution stable et rapide des équations algébriques*. Doctoral thesis, University of Paris 6, 2002.
<http://www.inria.fr/rrrt/tu-1067.html>
- [32] <http://www.cgal.org/index.html>
- [33] <http://www-sop.inria.fr/galaad/logiciels/synaps/index.html>
- [34] <http://www-sop.inria.fr/epidaure/personnel/Herve.Delingette/yav++.html>

Color figures




Figure 11: Example of a Delaunay triangulation (in blue) restricted to a curve (in black).





Figure 12: The standard left trefoil knot, in “sausage” format, meshed with $\sigma = 0.09 d_M$.

Appendix: various lemmas

This appendix recalls a few well-known results.

Theorem A.1 [*Meusnier's theorem*]

Let S be a C^2 -continuous surface, m be a point of S , and ξ be a curve drawn on S that passes through m . We parameterize ξ by arclength. Then,

$$\|\xi''\| = \frac{\mathbf{II}(\xi', \xi')}{\cos \theta}$$

where ξ' and ξ'' are respectively the first and second derivatives of ξ at point m , \mathbf{II} is the second fundamental form of S at m , and θ is the angle between the normal to S and the normal to ξ at m .

A proof of this theorem can be found in [21].

Lemma A.2 Let \vec{v} be a vector of Euclidean space \mathbb{R}^3 , and Π a vectorial plane that is not orthogonal to \vec{v} . Let \vec{v}^π denote the orthogonal projection of \vec{v} onto Π . For any vector \vec{u} of $\Pi \setminus \{\vec{0}\}$, we have $(\vec{v}, \vec{u}) \geq (\vec{v}, \vec{v}^\pi)$.

Proof Notice that $(\vec{v} - \vec{v}^\pi)$ is orthogonal to Π . Thus, for all $\vec{x} \in \Pi$, $\vec{v} \cdot \vec{x} = \vec{v}^\pi \cdot \vec{x}$. Consider $C = \{\vec{x} \in \Pi \mid \|\vec{x}\| = \|\vec{v}^\pi\|\}$. For all $\vec{u} \in C$, we have $\vec{v}^\pi \cdot (\vec{u} - \vec{v}^\pi) \leq 0$. Thus, $\vec{v} \cdot \vec{u} = \vec{v} \cdot \vec{v}^\pi + \vec{v} \cdot (\vec{u} - \vec{v}^\pi) = \vec{v} \cdot \vec{v}^\pi + \vec{v}^\pi \cdot (\vec{u} - \vec{v}^\pi) \leq \vec{v} \cdot \vec{v}^\pi$, which means that $\cos(\vec{v}, \vec{u}) \leq \cos(\vec{v}, \vec{v}^\pi)$ since $\|\vec{u}\| = \|\vec{v}^\pi\|$. It follows that $(\vec{v}, \vec{u}) \geq (\vec{v}, \vec{v}^\pi)$.

Now, let \vec{u} be any non-zero vector of Π . There exists a unique vector $\vec{u}^c \in C$ such that $(\vec{u}, \vec{u}^c) = 0$. It follows that $(\vec{u}, \vec{v}) = (\vec{u}^c, \vec{v}) \geq (\vec{v}, \vec{v}^\pi)$. \square

Lemma A.3 Let f and g be two univariate functions of class C^2 . Let x_a and x_b ($x_a < x_b$) be two reals such that

- (i) $f(x_a) = g(x_a)$ and $f(x_b) = g(x_b)$
- (ii) $\forall x \in [x_a, x_b]$, $f(x) \geq g(x)$
- (iii) $\forall x \in]x_a, x_b[$, $g''(x) \leq 0$

Then there exists a real $x_c \in]x_a, x_b[$ such that $f''(x_c) \leq g''(x_c) \leq 0$.

Proof By $f'(x_a)$ and $g'(x_a)$ we denote the right derivatives of f and g at x_a . Idem, by $f'(x_b)$ and $g'(x_b)$ we denote the left derivatives of f and g at x_b . (i) and (ii) imply that $f'(x_a) \geq g'(x_a)$, since otherwise there would exist a neighbourhood V_a of x_a such that $\forall x \in V_a \setminus \{x_a\}$, $\frac{f(x) - f(x_a)}{x - x_a} < \frac{g(x) - g(x_a)}{x - x_a}$, which would give that $f < g$ on $V_a \setminus \{x_a\}$, which contradicts (ii). Idem, we have $f'(x_b) \leq g'(x_b)$.

Now, Taylor-Lagrange formula (at first order), applied to function $(f'-g')$, tells that there exists a real $x_c \in]x_a, x_b[$ such that $(f' - g')'(x_c) = \frac{(f' - g')(x_b) - (f' - g')(x_a)}{x_b - x_a}$, which is negative since $f'(x_a) - g'(x_a) \geq 0$ and $f'(x_b) - g'(x_b) \leq 0$. It follows that $f''(x_c) - g''(x_c) \leq 0$. \square

Lemma A.4 *Let S be a compact surface without boundary, embedded in \mathbb{R}^3 , and let \vec{v} be a vector. We choose an orthonormal frame (O, x, y, z) such that \vec{v} is oriented along the $[0, z)$ direction. Let Ω be a convex subset of \mathbb{R}^3 , such that $\forall x \in S \cap \Omega$, $(\vec{n}(x), \vec{v}) < \frac{\pi}{2}$, where $\vec{n}(x)$ denotes the normal to S at x . Then $S \cap \Omega$ is xy -monotone.*

Proof We assume without loss of generality that S is oriented such that its normal points outwards. Let us assume for a contradiction that there exists a point (x_0, y_0) of plane (O, x, y) such that the vertical line d passing through (x_0, y_0) intersects $S \cap \Omega$ at least twice. Let (x_0, y_0, z_1) and (x_0, y_0, z_2) be two points of intersection that are consecutive along d . If there are not two such points, then this means that d intersects $S \cap \Omega$ along a segment (which is a degenerate case), and at each point of this segment the normal to S is orthogonal to d , and thus has zero dot-product with \vec{v} , which contradicts the hypothesis of the lemma. So now we assume that points (x_0, y_0, z_1) and (x_0, y_0, z_2) do exist. By definition, they are consecutive among the points of $S \cap \Omega \cap d$. Since Ω is convex, $\Omega \cap d$ is a segment of d , hence points (x_0, y_0, z_1) and (x_0, y_0, z_2) are also consecutive among the points of $S \cap d$. Thus, the open segment of d that joins them is included in one component of $\mathbb{R}^3 \setminus S$. It follows that $\vec{n}(x_0, y_0, z_1)$ or $\vec{n}(x_0, y_0, z_2)$ has a negative or zero dot-product with \vec{v} , since the normal of S always points outwards. This contradicts the hypothesis of the lemma. \square

Lemma A.5 *Let S be a compact surface without boundary, embedded in \mathbb{R}^3 , and let \vec{v} be a vector. We choose an orthonormal frame (O, x, y, z) such that \vec{v} is oriented along the $[0, z)$ direction. Let B be a ball centered at point $c \in S$, such that $\forall x \in B \cap S$, $(\vec{n}(x), \vec{v}) \leq \frac{\pi}{4}$. Then $S \cap B$ lies outside the cone K of apex c , of vertical axis and of half-angle $\frac{\pi}{4}$.*

Proof Let proj be the vertical projection onto plane (O, x, y) . Since $B \cap S$ is xy -monotone by Lemma A.4, the projection of $B \cap S$ is one-one. Let us assume for a contradiction that there exists a point $c' \in B \cap S$ that lies inside K . Let P be the vertical plane that passes through c and c' . It intersects $B \cap S$ along a set of simple arcs, since $B \cap S$ is xy -monotone. We consider the segment that joins $\text{proj}(c)$ and $\text{proj}(c')$:

1. if it is included in $\text{proj}(B \cap S)$, then c and c' belong to the same connected arc of $P \cap B \cap S$. The problem becomes then two dimensional: inside plane P , K is a cone of apex c , of vertical axis and of semi-angle $\frac{\pi}{4}$, and c and c' belong to a connected arc that is the graph of a function f , and whose normal makes an angle lower than $\frac{\pi}{4}$ with the vertical direction (which means that $|f'| \leq 1$). Let (O, t, z) be an orthonormal frame of P . We call c_t and c'_t the t -coordinates of c and c' . We have $|f(c'_t) - f(c_t)| = \left| \int_{c_t}^{c'_t} f'(t) dt \right| \leq \left| \int_{c_t}^{c'_t} |f'(t)| dt \right| \leq |c'_t - c_t|$, which means that c' does not belong to K , which contradicts the assumption.

2. if the segment that joins $\text{proj}(c)$ and $\text{proj}(c')$ is not entirely included in $\text{proj}(B \cap S)$, then we have to find another point that satisfies all the assumptions of case 1. We call c'_1 and c'_2 the points of ∂B that have same (x, y) coordinates as c' , and we assume without loss of generality that

c'_1 lies above c'_2 . Let p_1 and p_2 be the upper and lower poles of ∂B . We consider the meridian m of ∂B that passes through p_1 , c'_1 , c'_2 and p_2 . The smaller arc α_1 of m that joins p_1 and c'_1 , and the smaller arc α_2 of m that joins p_2 and c'_2 , project themselves onto the segment that joins $\text{proj}(c) = \text{proj}(p_1) = \text{proj}(p_2)$ and $\text{proj}(c') = \text{proj}(c'_1) = \text{proj}(c'_2)$. Since this segment is not entirely included in $\text{proj}(B \cap S)$, it intersects the boundary of $\text{proj}(B \cap S)$. Let c^* be the point of intersection that is closest to $\text{proj}(c)$. Since S has no boundary, the boundary of $B \cap S$ belongs to ∂B , thus the point of $B \cap S$ that projects onto c^* lies on α_1 or α_2 . We call $\text{proj}^{-1}(c^*)$ this point. Since $c' \in K$ by assumption, c'_1 and c'_2 also belong to K . Thus, all the points of α_1 belong to K since they are closer than c'_1 to the vertical axis of the cone, while their z-coordinate is bigger than that of c'_1 . Idem, $\alpha_2 \subset K$. Thus, $\text{proj}^{-1}(c^*)$ belongs to K . Moreover, since c^* is the point of intersection that is closest to $\text{proj}(c)$, segment $[\text{proj}(c), c^*]$ is included in $\text{proj}(B \cap S)$. So, $\text{proj}^{-1}(c^*)$ verifies all the assumptions of case 1., and thus leads to the same contradiction. \square