

Surface sampling and the intrinsic Voronoi diagram

Ramsay Dyer, Hao Zhang, and Torsten Möller[†]

Abstract

We develop adaptive sampling criteria which guarantee a topologically faithful mesh and demonstrate an improvement and simplification over earlier results, albeit restricted to 2D surfaces. These sampling criteria are based on functions defined by intrinsic properties of the surface: the strong convexity radius and the injectivity radius. We establish inequalities that relate these functions to the local feature size, thus enabling a comparison between the demands of the intrinsic sampling criteria and those based on Euclidean distances and the medial axis.

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

1. Introduction

A standard digital geometry representation of the surface of a 3D object is a triangle mesh. The process of producing a mesh from a known surface is called *meshing*. The vertices of the mesh are sample points lying on the given surface. Thus part of the meshing process involves *sampling*: placing sample points on the surface. A fundamental problem is to determine the sampling density required for the mesh to meet given accuracy requirements. At the most basic level one demands topological consistency: the mesh is homeomorphic to the surface in question. Geometric criteria which demand accurate positional and normal approximations usually assume that topological consistency has been met.

It is the fundamental question of topological consistency that we address in this paper. This matter has certainly been studied previously, mainly from two different perspectives: the *intrinsic* and the *extrinsic*. On the one hand the surface can be viewed as an abstract Riemannian manifold, possessing an intrinsic metric. No reference is made to the ambient embedding space. A notable work in this vein is that of Leibon and Letscher [LL00], where the notions of injectivity radius and strong convexity radius, well established in the Riemannian geometry literature, were brought to bear on the problem of surface sampling. However, these ideas have not received a lot of attention in the geometry processing community. Thus the intrinsic viewpoint has not yet proven itself to have practical utility for surfaces in 3D.

The more common approach to surface sampling involves measuring distances in the ambient Euclidean space. Since the seminal work of Amenta and Bern [AB98], sampling density on a smooth surface has typically been specified in terms of the *local feature size*: the distance to the medial axis. Their original sampling criterion is built upon the foundation laid down by Edelsbrunner and Shah [ES94], who defined a set of Voronoi diagram properties to characterize the sample sets that yield a homeomorphic mesh. Amenta and Bern then quantified a sampling density to ensure the needed Voronoi diagram properties.

Contribution: We take the intrinsic approach, utilizing concepts previously exploited by Leibon and Letscher, to define sampling criteria that ensure the required properties in the Voronoi diagram. Our contribution to the existing theoretical framework on surface sampling is twofold:

- We obtain a significant relaxation in the sampling requirements over those obtained by Leibon and Letscher [LL00, Lei99], and simplify the exposition.
- We establish inequalities relating local feature size to the strong convexity radius and the injectivity radius, bridging a gap between the extrinsic and intrinsic views on sampling criteria. In particular, we are now able to say what extrinsic criterion is needed to ensure that a sample set satisfies a given intrinsic criterion. This opens the door to exploiting intrinsic analysis even when extrinsic methods are employed. It is shown that the extrinsic criterion that is needed to satisfy our intrinsic sampling condition is weaker than existing extrinsic sampling criteria for topological consistency.

[†] GrUVi Lab, School of Computing Science, Simon Fraser University, BC, Canada. Email: rhdyer, haoz, torsten@cs.sfu.ca

Related work: The literature on surface sampling is vast. We mention here only recent works and those that have directly influenced our current work. The work of Boissonnat and Oudot [BO05], with its graceful exploitation of the Voronoi-Delaunay duality in surface sampling, was an inspiration. A clear distinction between meshing and surface reconstruction was made in that work. Cheng et al. [CDRR04] present a meshing algorithm with guaranteed topological consistency which avoids explicit reference to the local feature size. The algorithm strives for a sparse sampling, but there is no concise expression for the final sampling density. Gao et al. [GGOW08] use an intrinsic sizing function, the homotopy feature size (hfs), to mesh planar domains. The hfs appears to be closely related to the injectivity radius discussed here.

Morvan and Thibert [MT04], as well as Hildebrandt et al. [HPW06], have done a nice quantitative analysis of the metric distortion between a smooth surface and a mesh that approximates it. Although our work focuses more on the qualitative local behaviour of geodesics, we did exploit these works when it was necessary to bound geodesic lengths to allow a comparison between intrinsic and extrinsic sampling criteria. Dai et al. [DLYG06] have also recently produced a geometric accuracy analysis. Their analysis is from the intrinsic viewpoint, using Leibon and Letscher's work [LL00] as the topological correctness foundation. The sampling conditions required by their main theorem (Thm 3), involve a minimum amongst terms representing Leibon and Letscher's criterion and (larger) scaled extrinsic sizing functions. Thus it can be both relaxed and simplified in light of our work.

Paper organization: The main body of the paper is organized into three sections. In section 2, we define the intrinsic Voronoi diagram and strong convexity and conclude with Theorem 1, a new observation on how strong convexity can enforce desirable topological properties in the intrinsic Voronoi diagram. Section 3 reviews some sizing functions that are related to variable sampling density requirements. Our new sampling density requirements and the estimates on the intrinsic sizing functions with respect to the local feature size are presented in Section 4, where we go on to synthesize these results and enable a direct comparison between extrinsic and intrinsic sampling criteria.

Notation: Throughout this paper, S will refer to a smooth, compact, connected surface without boundary isometrically embedded in \mathbb{R}^3 , and P will refer to a set of sample points on S . Given two points p and q on S , $d_{\mathbb{R}^3}(p, q)$ and $d_S(p, q)$ denote respectively the Euclidean and geodesic distances between p and q . Likewise $B_S(x; r) = \{y \in S \mid d_S(x, y) < r\}$ denotes the open geodesic disk centred at x , and $B_{\mathbb{R}^3}(x; r)$ is an open Euclidean ball. A closed disk/ball is indicated with a bar: $\bar{B}_S(x; r)$. All geodesics and space curves are arc-length parameterized. The length of a curve γ is $\ell(\gamma)$ and we abuse the notation by identifying γ with its image. Thus for example, $z \in \gamma$ means $z \in \gamma([0, \ell(\gamma)])$.

2. The intrinsic Voronoi diagram

Geodesic curves: The notion of geodesic curves on S is essential in our exposition. We use the following defining characteristic: A *geodesic* on S is a curve $\gamma \subset S$ that, when viewed as a space curve $\gamma \subset \mathbb{R}^3$, has its curvature vector parallel to the normal vector on S at all corresponding points where the curvature of γ does not vanish. In particular, the curvature of γ at $x \in \gamma$ is bounded by the largest magnitude of the principle curvatures of S at x .

Compact surfaces are geodesically complete; any two points $p, q \in S$ can be connected by a smooth curve γ of minimal length: $\ell(\gamma) = d_S(p, q)$, and this curve is a geodesic (the Hopf-Rinow theorem [dC92]). Also, for any $z \in \gamma$, the portion of γ between p and z is also a minimal geodesic between p and z [dC92]. However, a minimal geodesic is not necessarily unique. If two points are sufficiently close together we can be assured of a unique minimal geodesic between them [dC92].

The intrinsic Voronoi diagram of a sample set P on S is naturally defined without imposing restrictions on P .

Definition 1 (Voronoi diagram) The *Voronoi cell* of $p \in P$ is defined by $\mathcal{V}(p) = \{x \in S \mid d_S(p, x) \leq d_S(q, x), \forall q \in P\}$. The set of Voronoi cells forms a covering of S called the *intrinsic Voronoi diagram* (iVd) of P on S . Replacing d_S with $d_{\mathbb{R}^3}$, we obtain the *restricted Voronoi diagram* (rVd) on S .

2.1. The closed ball property

Edelsbrunner and Shah. [ES94] introduced the *closed ball property* to describe those restricted Voronoi diagrams on S whose dual, the *restricted Delaunay triangulation* (rDt), is a simplicial complex homeomorphic to S . For a surface without boundary, the property expresses three conditions:

1. each Voronoi cell is a closed topological disk (2-ball);
2. the intersection of two Voronoi cells is either empty or a closed topological 1-ball: a *single Voronoi edge*;
3. the intersection of three Voronoi cells is either empty or a single point (0-ball): a *single Voronoi vertex*.

General position: The closed ball property implicitly imposes a *general position* condition on the sample points in that the intersection of more than three Voronoi cells must necessarily be empty (otherwise there would be two Voronoi cells whose intersection was a single point). No sampling criteria based on sample density alone can guarantee that the general position property is satisfied, but in principle and practice such failures can be fixed with an arbitrarily small perturbation; we will henceforth ignore this technicality.

Generality: Edelsbrunner and Shah. [ES94] showed that the rDt is homeomorphic to S when the rVd satisfies the closed ball property. Their proof looks at the topological properties of the dual complex resulting from a Voronoi diagram satisfying this property. An examination of that proof

reveals that it does not rely on the specific metric employed to generate the Voronoi diagram. In particular, we can apply the result to meshes constructed from the iVd: the *iDt-mesh* is a triangle mesh whose edges connect samples that are Voronoi neighbours (i.e., that share a Voronoi edge). To every Voronoi vertex in the iVd, there corresponds a face in the iDt-mesh. If the iVd satisfies the closed ball property, the iDt-mesh is a manifold mesh homeomorphic to S .

Note the distinction between the iDt-mesh and the iDt itself: The *intrinsic Delaunay triangulation* (iDt) of P on S is a triangulation of S formed by connecting, with minimal geodesics, samples that are Voronoi neighbours. Although a iVd satisfying the closed ball property guarantees a homeomorphic manifold iDt-mesh, whose edges are straight line segments in \mathbb{R}^3 , it does *not* guarantee that the iDt itself is well defined. In particular, there may be more than one minimal geodesic connecting a given pair of sample points. However, if the iVd also has the property that there is a unique minimal geodesic between samples that are Voronoi neighbours, then the iDt exists and it is combinatorially equivalent to the iDt-mesh.

Redundancy of condition 3: We have assumed that S is a single component. We further assume that there are at least four distinct samples in P . In this case, the third condition of the closed ball property becomes redundant. Indeed, suppose that the first two closed ball conditions are satisfied but the third is violated. Thus we have three Voronoi cells $\mathcal{V}(p)$, $\mathcal{V}(q)$ and $\mathcal{V}(r)$ whose intersection includes two distinct points a and b . Then $\mathcal{V}(p) \cap \mathcal{V}(q)$ will be a Voronoi edge, e , following condition 2, and a and b must necessarily be its endpoints. Likewise, $\mathcal{V}(p) \cap \mathcal{V}(r)$ will yield another distinct Voronoi edge, e' , also with a and b as endpoints. Since $\mathcal{V}(p)$ is a topological disk (condition 1), e and e' make up its entire boundary. It follows that p has no Voronoi neighbours other than q and r . Arguing similarly for q and r , we conclude that there are only three samples on S .

Definition 2 (Well formed Voronoi diagram) A (restricted or intrinsic) Voronoi diagram is *well formed* if it consists of at least four samples and satisfies the first two closed ball conditions.

2.2. Strong convexity

We seek a sampling criterion that guarantees a well formed Voronoi diagram and thus a valid homeomorphic iDt-mesh. To this end it is useful to examine the notion of convexity of sets on a surface. In the planar setting a set A is convex if a line segment connecting any two points in A lies in A . On a surface, lines are replaced by geodesics.

There are several ways to extend the notion of convexity to sets on a surface. We follow Chavel [Cha06]:

Definition 3 (Strongly convex set on surface) A set $A \subset S$ is *strongly convex* if for every $p, q \in A$,

1. there is a *unique minimal* geodesic γ in S connecting p and q ;
2. γ lies entirely within A ;
3. no other geodesic connecting p and q lies within A .


The set A is *convex* if it satisfies only the first two conditions above, but we are interested in strong convexity. Note that there are many non-equivalent definitions of convexity and strong convexity in the Riemannian geometry literature, so care must be taken when referring to other works.

Note that the intersection of strongly convex sets is strongly convex. Such sets are also contractible [dC92], which implies the following useful observation:

Lemma 1 A strongly convex set is simply connected.

2.3. The closed ball property via strong convexity

As the sampling density increases we expect the Voronoi cells to more closely exhibit the characteristics of those in a planar Voronoi diagram. One notable characteristic of a Voronoi cell in the plane is that it is convex. Voronoi cells on a surface cannot share this property, however, because Voronoi edges are not necessarily geodesics. To see this, let



p, q lie on a plane, in which case their Voronoi boundary is a straight line (a geodesic).

Now introduce a small bump near p , between p and q , that is far from the original Voronoi boundary, as shown in the figure. This will distort the originally straight boundary so that it is no longer a geodesic. For a general surface, regardless of the sampling density, we cannot expect Voronoi edges to be geodesics.

Therefore, we cannot demand convexity from the Voronoi cells on S . However, a sufficient sampling density will ensure that the Voronoi cells are contained in strongly convex neighbourhoods, which turns out to be a useful criterion.

Theorem 1 (Strong convexity and well formed iVd) If $\forall p \in P$ there exists a strongly convex set $U_p \subset S$ with $\mathcal{V}(p) \subset U_p$, then the iVd of $P \subset S$ is well formed.

The proof of Theorem 1 relies on the following lemma.

Lemma 2 Suppose that $\Omega \subset S$ is a union of Voronoi cells that is bounded by only two Voronoi cells, $\mathcal{V}(p), \mathcal{V}(q) \subset S \setminus \Omega$ with $\mathcal{V}(p) \cap \mathcal{V}(q) \cap \Omega \neq \emptyset$. Then Ω contains a geodesic γ that cannot be contained in a strongly convex set.

Proof. Let a be a Voronoi vertex in $\mathcal{V}(p) \cap \mathcal{V}(q) \cap \Omega$, and let $s \in P$ be a sample in Ω with $a \in \mathcal{V}(s)$. Let γ_{as} be a minimal geodesic between a and s . Then we must have $\gamma_{as} \subset \mathcal{V}(s)$, since otherwise, $\exists w \in \gamma_{as}$ and $u \in P$ with $u \neq s$ and $w \in \mathcal{V}(u)$ but $w \notin \mathcal{V}(s)$. It follows that $d_S(a, u) \leq d_S(a, w) + d_S(w, u) < d_S(a, w) + d_S(w, s) = d_S(a, s)$, contradicting $a \in \mathcal{V}(s)$.

We extend γ_{as} through s until it exits Ω at t . Let $\gamma = \gamma_{at}$ denote the resulting geodesic. Clearly, $\gamma \in \Omega$. Without loss of




Figure 1: Ω is a region enclosed by the boundaries of $\mathcal{V}(p)$ and $\mathcal{V}(q)$. (a) The geodesic γ is at least as long as the green curve and cannot be contained in any strongly convex set. (b) The minimal geodesic α must be contained in U_p . If α does not cross γ , then either U_q or U_p contains a loop of minimal geodesics (red+green curves) that encompasses γ . Assuming it is U_p , as shown, then U_p cannot be simply connected and not contain γ . (c) If α does cross γ , then U_p (or U_q) will contain two geodesics between a and z .

generality, let t lie on the boundary of $\mathcal{V}(p)$; see Figure 1(a). Note that if γ did not exit Ω , it would exceed the diameter of S and be too long to reside in a strongly convex set.

Now note that both a and t lie on the boundary of $\mathcal{V}(p)$ and, by our choice of s , $d_S(a, p) = d_S(a, s)$, since a is a Voronoi vertex. Also $d_S(t, p) \leq d_S(t, s)$. Thus γ is at least as long as the path between a and t comprised of a minimal geodesic between a and p and a minimal geodesic between p and t , shown as the green curve in Figure 1(a). It follows that γ cannot be a unique minimal geodesic between a and t and so cannot be contained in any strongly convex set. \square

Proof of Theorem 1: Suppose that $\mathcal{V}(p)$ is not a topological disk. By Lemma 1, U_p is simply connected. Thus a homotopically nontrivial loop in $\mathcal{V}(p)$ is homotopically trivial in U_p , and so U_p must contain a region Ω that is exterior to but bounded on all sides by $\mathcal{V}(p)$. By an argument identical to the proof of Lemma 2 we see that such an Ω cannot be contained in a strongly convex set. Therefore $\mathcal{V}(p)$ must be a topological disk.

Now suppose that $\mathcal{V}(p)$ and $\mathcal{V}(q)$ meet at more than one distinct Voronoi edge. These cells then bound a region Ω as described in Lemma 2. Consider the strongly convex neighbourhoods U_p and U_q . Since $\mathcal{V}(p) \subset U_p$ and $\mathcal{V}(q) \subset U_q$, the intersection $U_p \cap U_q$ must contain all Voronoi edges in $\mathcal{V}(p) \cap \mathcal{V}(q)$. Also, since $U_p \cap U_q$ is strongly convex the minimal geodesic, α , between points a and b on distinct Voronoi edges (see Figure 1(b)) must lie in $U_p \cap U_q$.

Let γ be the geodesic in Ω that was constructed in Lemma 2. Suppose that α does not cross γ . Then γ must be contained in the region bounded by a loop of minimal geodesics (a geodesic triangle in fact) involving α and either p or q . Assume it is p ; refer to Figure 1(b). Now U_p must contain α , the minimal geodesic between a and p , and the one between p and b . However, by Lemma 2, U_p cannot contain γ . We arrive at a contradiction to Lemma 1.

Thus α must cross γ . Let z be the first such intersection that is encountered on a traversal of γ starting at point a and assume that this portion of γ between a and z is contained in a region bounded by a loop of minimal geodesics involving p (green curves) and α (red curve); see Figure 1(c). But now there are two geodesics between a and z ; U_p cannot contain this portion of γ , but it is forced to if it is to contain α and remain simply connected. Again a contradiction. \square

3. Sizing functions for surface sampling

Adaptive sampling criteria on surfaces generally impose restrictions on the sampling density based on local curvature properties as well as semi-local properties relating to some notion of the distance to “the other side” of the surface. We refer to functions that can be used to modulate the sampling density in this way as *sizing functions*. These functions take positive values which can be thought of as having the units of distance. Thus at each point on the surface, a sizing function specifies a radius within which a certain proportion of representative samples is expected. The square of such a function can, for example, be used to define a weighted area measure for governing stochastic sampling.

We first present a natural hierarchy of familiar sizing functions that have extrinsic definitions and then discuss the properties of some less familiar intrinsic sizing functions. In Section 4.2 we find inequalities that bridge the gap between representatives of these two sets of sizing functions.

3.1. A natural hierarchy of extrinsic sizing functions

We will make use of the following definition:

Definition 4 (medial ball and medial axis) Given a closed set $C \subset \mathbb{R}^3$, e.g., a surface, a *medial ball* is an open ball $B \in \mathbb{R}^3 \setminus C$ that is maximal with respect to inclusion (i.e., no other open ball can contain B). If $p \in C$ is contained in the closure of B , we say B is a medial ball at p . The *medial axis* of C is the closure of the set of centres of all the medial balls.

Local feature size: The *local feature size* (lfs) at a point $x \in S$, denoted by $\rho_f(x)$, is the distance from x to the medial axis of S . It was introduced by Amenta and Bern [AB98] and one would be hard pressed to find a subsequent paper on surface sampling that does not mention it. For this reason we will go to some lengths to compare our sampling criteria to those that are expressed in terms of the lfs.

The lfs is particularly convenient to work with because it enjoys the property of Lipschitz continuity,

$$|\rho_f(x) - \rho_f(z)| \leq d_{\mathbb{R}^3}(x, z), \quad (1)$$

allowing us to bound $\rho_f(x)$ in terms of a nearby $\rho_f(z)$.

Local reach: Each point $x \in S$ is associated with two medial balls, one on each side of the surface; one of them may have infinite radius. These balls are tangent to S at x . The radius, $\rho_R(x)$, of the smaller of the two medial balls at x is called the *local reach* at x . It is the distance from x to the medial axis along a direction normal to S at x . It was introduced by Federer [Fed59] where it was observed that it is a continuous function on S . It is bounded below by the lfs for all $x \in S$.




Figure 2: lfs vs. local reach.

More commonly encountered than the local reach is its global lower bound. The *reach* of S is $\rho_R = \inf_{p \in S} \rho_R(p)$. The reach is important because it specifies the maximal constant thickness of a tubular neighbourhood of S for which the natural (orthogonal) projection onto S is well defined. Specifically, if $U_{\rho_R} \subset \mathbb{R}^3$ is given by $\{m \in \mathbb{R}^3 \mid d_{\mathbb{R}^3}(m, S) < \rho_R\}$, then the projection $\xi : U_{\rho_R} \rightarrow S$ can be defined by specifying $\xi(m)$ to be the unique point on S that is closest to m .

Maximal curvature: The *maximal curvature* at $x \in S$ is the maximum of the absolute values of the two principal curvatures: $\kappa(p) = \max\{|\kappa_1(p)|, |\kappa_2(p)|\}$. The associated sizing function, the *maximal curvature radius*, is given by the radius of the associated osculating sphere: $\rho_\kappa(p) = 1/\kappa(p)$.

Since the smallest medial ball at x can never exceed the size of the osculating sphere at x , we have $\rho_R(x) \leq \rho_\kappa(x)$. Although the maximal curvature is continuous, it can vanish. Thus $\rho_\kappa(x)$ is not bounded nor even well defined everywhere. However, for its principle employment as an upper bound on the lfs, we may interpret $\rho_\kappa(x) = \infty$ at those points where the maximal curvature vanishes.

Gaussian curvature: The *Gaussian curvature* is the product of the principle curvatures: $G(x) = \kappa_1(x)\kappa_2(x)$. The *Gaussian curvature radius*, defined by $\rho_G(x) = 1/\sqrt{G(x)}$, is the associated sizing function. Like the maximal curvature radius, the Gaussian curvature radius suffers from a problem of definition when the Gaussian curvature is non-positive. Read $\rho_G(x) = \infty$ when $G(x) \leq 0$. We have $\rho_\kappa(x) \leq \rho_G(x)$.

The hierarchy: For all $x \in S$ we have the relations

$$\rho_f(x) \leq \rho_R(x) \leq \rho_\kappa(x) \leq \rho_G(x). \quad (2)$$

Of the four sizing functions represented here, the latter two are useful primarily for theoretical analysis; their definition is too local to be used as a sizing function for a sampling criterion. Unlike the other three, the Gaussian curvature is truly intrinsic, but it fits comfortably into this hierarchy with the extrinsic functions.

3.2. Intrinsic sizing functions

We now present intrinsic sizing functions which embody standard concepts in Riemannian geometry. An introduc-

tory textbook on Riemannian geometry, e.g., [dC92], may be consulted for further details on the statements made here.

3.2.1. Strong convexity radius

On the plane, convexity has to do with the shape of a set, but on a surface, the strong convexity condition also limits the size of the set. For example a geodesic disk is not strongly convex in general. Consider a geodesic disk on a cylinder. If the radius of the disk exceeds one quarter the circumference of the cylinder, then there will be points on the disk whose shortest connecting geodesic leaves the disk.

Definition 5 (Strong convexity radius) The *strong convexity radius* (*scr*) at a point $x \in S$ is defined as

$$\rho_{sc}(x) = \sup\{\rho \mid B_S(x; r) \text{ is strongly convex } \forall r < \rho\}.$$

It can be shown [dC92] that for any $x \in S$, $\rho_{sc}(x) > 0$. The scr is an intrinsic quantity, and in general, if x is in a region of high curvature, $\rho_{sc}(x)$ will be small. However, Gaussian curvature alone is not sufficient to characterize the scr. Consider again the example of a cylinder. The scr at any point will be no greater than 1/4 the circumference of the cylinder, but the Gaussian curvature radius is unbounded.

There are no continuity results for the strong convexity radius. However, it is worth mentioning that Klingenberg [Kli95][1.9.9] has a definition of strong convexity in which a further axiom is imposed: A set A is strongly convex in Klingenberg's sense if it satisfies the axioms of Definition 3 and further has the property that any geodesic disk B contained in A is also convex. It easily follows that the resulting scr is 1-Lipschitz. It is shown that the scr is always positive even with this additional axiom.

3.2.2. Injectivity radius

An explanation of the injectivity radius requires a brief description of the exponential map. We denote by $T_x S$ the tangent plane of S at x . The *exponential map* at x is a smooth mapping $\exp_x : T_x S \rightarrow S$ that takes $X \in T_x S$ to the point $\gamma_X(\|X\|) \in S$, where γ_X is the geodesic emanating from x with tangent vector $X/\|X\|$. Restricted to a small enough disk in $T_x S$, \exp_x is a diffeomorphism onto its image [dC92].

Definition 6 (Injectivity radius) The *injectivity radius* at $x \in S$ is the supremum of the radii for which \exp_x is injective:

$$\rho_i(x) = \sup\{\rho \mid \exp_x \text{ is injective on } B_{T_x S}(0; r) \forall r < \rho\}.$$

The function $\rho_i(x)$ is continuous on S [Cha06]. One of the most useful properties of $\rho_i(x)$ follows from the definition: if $d_S(x, p) < \rho_i(x)$, then there is a unique minimal geodesic γ between x and p and it will be the only geodesic between x and p that is contained in $B_S(x; \rho_i(x))$. By the first axiom of Definition 3, the radius of a strongly convex disk cannot exceed the injectivity radius of the centre. Thus $\rho_{sc}(x) \leq \rho_i(x)$.

Note that the image under \exp_x of any disk D centred at 0 in T_xS is exactly a geodesic disk, and that if the radius of D is less than $\rho_i(x)$, then the image of D will be a topological disk. However, it is incorrect to say that $\rho_i(x)$ is the largest radius for which the geodesic disk centred at x is an embedded topological disk. Indeed, as the radius increases, there are two ways that the exponential map can fail to be injective. One is if the disk wraps around and merges with itself to create nontrivial topology.

But the other situation that can occur is that the Jacobian of \exp_x may become degenerate. Suppose this happens at a point z . The concentric geodesic circles centred at x will be smooth provided their radius is less than $\rho_i(x)$, how-




Figure 3: Two ways the exponential map can fail to be injective.

ever the circle through z may have a cusp. If γ is the minimal geodesic joining x to z we say that z is a *conjugate point* to x along γ (and vice versa: it is reciprocal). For all points y on γ between x and z , the minimal geodesic between x and y will be a portion of γ . However, for all points on the extension of γ past z , γ will not be a portion of the minimizing geodesic [Cha06]. Thus $z \notin B_S(x; \rho_i(x))$. More generally, if an open neighbourhood V of x is the bijective image of $\exp_x|_U$ for some $U \subset T_xS$, then V contains no conjugate points and hence, by the inverse function theorem $\exp_x|_U$ is a diffeomorphism onto V .

A *geodesic loop* is a geodesic that starts and ends at the same point. A *closed geodesic* (sometimes called a *periodic geodesic*) is a geodesic loop γ whose tangent vectors agree at its endpoints: $\dot{\gamma}(0) = \dot{\gamma}(\ell(\gamma))$. Returning now to the case where a change in the topology of the image of \exp_x occurs, it can be shown that there will be a geodesic loop starting and ending at x . The midpoint of this loop is the closest point to x at which \exp_x fails to be injective [Cha06].

If we extend a geodesic γ from x , there will be a closest point $z \in \gamma$ beyond which γ is no longer a minimizing geodesic. The set of all such points is called the *cut locus* of x and it is compact [dC92]. The assertions made above are summarized by a theorem [Cha06][III.2.4, p.118] due to Klingenberg:

Lemma 3 (Klingenberg) If q is the point on the cut locus of x that is closest to x , then q is either

- (i) conjugate to x along a minimal geodesic connecting them, or,
- (ii) the midpoint of a geodesic loop starting and ending at x .

3.2.3. Intrinsic sampling radius

There are no sampling criteria based exclusively on the injectivity radius. We give a sampling criterion, Corollary 1, that is based on the scr alone, but by combining the two sizing functions we are able to relax that criterion.

Definition 7 (Intrinsic sampling radius) The *intrinsic sampling radius* at $x \in S$ is given by

$$\rho_m(x) = \min \left\{ \rho_{sc}(x), \frac{1}{2}\rho_i(x) \right\}.$$

Clearly, $\rho_m(x) \leq \rho_{sc}(x) \leq \rho_i(x)$. In Section 4.2, we establish a relationship between these quantities and the lfs.

4. Relating extrinsic and intrinsic sampling criteria

In this section we examine the relationship between intrinsic sampling criteria based on the scr or the intrinsic sampling radius and extrinsic ones based on the lfs. We show that for any $\epsilon > 0$, there exists an $\epsilon_f > 0$ such that any sample set P that satisfies the extrinsic criterion:

$$\forall x \in S \quad \exists p \in P \text{ such that } p \in B_{\mathbb{R}^3}(x; \epsilon_f \rho_f(x)), \quad (3)$$

will also satisfy the corresponding intrinsic criterion:

$$\forall x \in S \quad \exists p \in P \text{ such that } p \in B_S(x; \epsilon \rho_m(x)). \quad (4)$$

We begin by observing in Section 4.1 that the strong convexity condition of Theorem 1 immediately yields a sampling criterion of the form (4), employing $\rho_{sc}(x)$ rather than $\rho_m(x)$ as the sampling radius. Subsequently we obtain a weaker criterion using $\rho_m(x)$. We then proceed to find an explicit extrinsic criterion (3) that is sufficient to meet this. A first step in this direction is to form an estimate on the intrinsic sampling radius based on the lfs. In particular, we seek a constant C such that

$$\rho_m(x) \geq C \rho_f(x) \text{ for any } x \in S. \quad (5)$$

Both $\rho_f(x)$ and $\rho_m(x)$ become smaller as the local maximal curvature becomes larger. However, $\rho_f(x)$ also becomes smaller when geodesically distant points of the surface become close in the ambient space — a property that $\rho_m(x)$ does not possess. Thus although the lfs may be bounded above by the intrinsic sampling radius, we cannot hope to find a constant that would allow us to reverse the inequality in Equation (5). For this same reason, without further qualifications, we can never guarantee that any intrinsic sampling (4) will satisfy a given extrinsic criterion (3).

In Section 4.2, we obtain an estimate for the constant C in Equation (5). In equation (3), distances to the sample set P are measured in the ambient space, whereas the intrinsic conditions (4) are specified with respect to geodesic distances on the surface, which are larger in general. So the next step is to put an upper bound on the geodesic distance between a point $x \in S$ and a nearby point $p \in P \subset S$ in terms of the Euclidean distance between them. This is done in Section 4.3. Finally, in Section 4.4 we develop an explicit relationship between the ϵ of Equation (4) and the ϵ_f of Equation (3) and we use this relationship to compare the sampling criteria derived in Section 4.1 with more familiar extrinsic sampling criteria for meshing and surface reconstruction.

4.1. Intrinsic sampling criteria

Theorem 1 ensures that if the Voronoi cells can be contained in strongly convex neighbourhoods, then a valid homeomorphic iDt-mesh can be constructed. We can use the strong convexity radius to define a sampling density that will guarantee the conditions of Theorem 1.

Corollary 1 If $\forall x \in S$ there exists a $p \in P$ such that $p \in B_S(x; \frac{1}{2}\rho_{sc}(x))$, then the iVd of P on S is well formed.

Proof. For any $\mathcal{V}(p)$, choose $x \in \mathcal{V}(p)$ that is at a maximal geodesic distance from p . The triangle inequality yields $\mathcal{V}(p) \subset B_S(x; \rho_{sc}(x))$ and Theorem 1 applies. \square

Note that this sampling condition ensures more. If we choose x to be a Voronoi vertex, then the associated samples must all lie within $B_S(x; \frac{1}{2}\rho_{sc}(x))$ and so there are unique minimal geodesics between them. Thus Corollary 1 can be strengthened to apply to intrinsic Delaunay triangulations:

Corollary 2 If $\forall x \in S$ there exists a $p \in P$ such that $p \in B_S(x; \frac{1}{2}\rho_{sc}(x))$, then the iDt of P on S exists.

By contrast, Theorem 1 itself is *not* apriory sufficient to ensure that the iDt itself exists. Thus the above sampling condition is stronger than the condition imposed by Theorem 1. Compared to Corollary 2, the sampling criterion of [LL00] is more complicated. Also it requires at least that there be a $p \in P$ such that $p \in B_S(x; \frac{1}{5}\rho_{sc}(x))$.

4.1.1. A weaker criterion

One observation that comes up in the demonstration of Corollary 2 is that if we have a sampling criterion that demands only that a sample lies within the strong convexity radius of any point on S , then we are guaranteed that there will be a unique minimal geodesic between any two samples that are Voronoi neighbours. So the question arises, in this case where the ϵ has been doubled from $\frac{1}{2}$ to 1: Is the iVd still well formed? This would imply the existence of the iDt.

We have not obtained an affirmative answer to this question. However, we are able to ensure that the iVd is well formed provided no point on the surface is as far as its intrinsic sampling radius from the nearest sample point. To facilitate this result, we borrow some terminology from Boissonat and Oudot [BO05]:

Definition 8 (pseudo-disks) A family $\{B_i\}$ of topological disks on S are *pseudo-disks* if for any two distinct disks B_i and B_j , their boundaries either do not intersect, or they intersect tangentially at a single point, or they intersect transversely at exactly two points.


Qualitatively pseudo-disks intersect each other in the manner expected of Euclidean disks:

Lemma 4 (Three circles) If B_1, B_2 , and B_3 are pseudo-disks whose boundaries all intersect at p and q , then one of the disks is contained in the union of the other two.

Proof. It is sufficient to show that the boundary of one of the disks is contained in the union of the other two disks.

Let $C_i = \partial B_i$. Each of these circles is composed of two arcs joining p and q .

Choose a consistent orientation on the circles and consider the three arcs emanating from p . One of these arcs must be inside one of the two other disks, but outside of the other. Suppose this arc belongs to C_2 (blue in figure), and that it is inside B_1 and outside B_3 . Since the intersections are transversal, the other arc on C_2 must be outside B_1 , but inside B_3 . Thus C_2 is contained in $B_1 \cup B_3$. \square



The following lemma is an improvement and simplification of its namesake in [Lei99]:

Lemma 5 (Small circle intersection) For $x \in S$ and $r < \rho_m(x)$, the disks $B_S(x; r)$ are pseudo-disks.

Proof. Let $C_x = \partial B_S(x; r_x)$ and $C_y = \partial B_S(y; r_y)$ be geodesic circles with $r_x < \rho_m(x)$ and $r_y < \rho_m(y)$ and $x \neq y$. If C_x and C_y intersect tangentially at z , then by the Gauss lemma [dC92][p.69], the minimal geodesics γ_{xz} and γ_{yz} connecting x and y with z must have parallel tangent vectors at z . It follows that x, y and z all lie on a common geodesic γ .

If $\gamma'_{xz}(r_x) = -\gamma'_{yz}(r_y)$, then $\ell(\gamma) = r_x + r_y < \frac{1}{2}(\rho_i(x) + \rho_i(y))$. It follows that either $x \in B_S(y; \rho_i(y))$, or $y \in B_S(x; \rho_i(x))$ and γ is the unique minimal geodesic connecting x and y . If there were another intersection at $w \neq z$, then $\alpha = \gamma_{xw} \cup \gamma_{yw}$ would be a path between x and y with $\ell(\alpha) = r_x + r_y$, contradicting the minimality of γ .

If on the other hand $\gamma'_{xz}(r_x) = \gamma'_{yz}(r_y)$, then either $x \in \gamma_{yz}$ or $y \in \gamma_{xz}$. Assume the former. Then $r_y = r_x + d_S(x, y)$. Now if there is another intersection at w , this same equality must apply and we conclude that $x \in \gamma_{yw}$. But a radial geodesic of length r_y from y can only meet C_y once, and assuming $d_S(x, y) > 0$, there is only one such geodesic that contains x . Therefore we must have $w = z$. Thus if C_x and C_y intersect tangentially at z , there can be no other points of intersection.

Now suppose that C_x and C_y intersect transversely at z . Then $D = B_S(x; r_x) \cap B_S(y; r_y) \neq \emptyset$. Because it is strongly convex, D must be a single connected component bounded by an arc of C_x and an arc of C_y . It follows that C_x and C_y must intersect transversely at another point distinct from z , and that there can be no further transverse intersections. \square

Lemmas 4 and 5 provide an obstruction to neighbouring Voronoi cells sharing more than two Voronoi vertices. This yields our main sampling result:

Theorem 2 The iDt of P on S exists if

$$\forall x \in S \quad \exists p \in P \text{ such that } p \in B_S(x; \rho_m(x)).$$

Proof. The sampling condition implies that there is a unique minimal geodesic between samples that are Voronoi neighbours. It remains to prove that the iVd is well formed.

Since p lies within the injectivity radius of each $x \in \mathcal{V}(p)$, there is an open neighbourhood V of $\mathcal{V}(p)$ that is

the diffeomorphic image under \exp_p of some $U \subset T_p S$ (see Section 3.2.2). This implies that $\mathcal{V}(p)$ is contractible (use $\exp_p \circ t \cdot \exp_p^{-1}$, $t \in [0, 1]$). Thus $\mathcal{V}(p)$ is a topological disk.

It remains to show that $\mathcal{V}(p)$ and $\mathcal{V}(q)$ cannot share more than a single Voronoi edge. If this were the case, we would have a region Ω , as in Figure 1(a), that is bounded completely by $\mathcal{V}(p)$ and $\mathcal{V}(q)$. Suppose that the Voronoi vertices a and b were both on the boundary of $\mathcal{V}(s) \subset \Omega$. Consider the geodesic circles centred at a and b and with radius $d_S(a, s) < \rho_m(a)$ and $d_S(b, s) < \rho_m(b)$, respectively. Since a and b are Voronoi vertices, these two circles would have to intersect at p and q in addition to s , contradicting Lemma 5.

Suppose then that the Voronoi vertices a and b are on the boundaries of $\mathcal{V}(s) \subset \Omega$ and $\mathcal{V}(v) \subset \Omega$, respectively. By hypothesis, $\mathcal{V}(p)$ and $\mathcal{V}(q)$ share more than a single Voronoi edge. Therefore there are at least four Voronoi vertices in $\mathcal{V}(p) \cap \mathcal{V}(q)$. Let c be such a Voronoi vertex, distinct from a and b . Consider the three disks centred at these Voronoi vertices and with radii such that their respective three closest samples lie on the boundary. These three disks are pseudo-disks, and their boundaries intersect at p and q . Thus by Lemma 4 one of these disks, say B_a , is contained in the union of the other two. However, B_a has a third sample, s , on its boundary, contradicting the fact that all three disks must have empty interiors. Thus the iVd must be well formed. \square

Since the injectivity radius is never smaller than the scr, we are assured that the conditions imposed by Theorem 2 are at least as weak as those demanded by Corollary 2. For the comparison with the lfs which we develop next, the result of Theorem 2 is twice as good.

4.2. A lfs estimate on the intrinsic sampling radius

In this section we arrive at Theorem 4, one of the main results of this paper. By producing a relationship (10) between the intrinsic sampling radius and the lfs, it opens the door for comparing intrinsic and extrinsic sampling criteria. We exploit curvature bounds established in the Riemannian geometry literature. Two facts, Lemmas 6 and 7, which respectively give insight into case (i) and (ii) of the Klingenberg Lemma 3, enable us to get an estimate on the injectivity radius. With the addition of a result from Chavel, Lemma 8, we obtain an estimate on the scr.

To tackle case (i) of Lemma 3, a theorem [Cha06][II.6.3, p.86] attributed to Morse and Schönberg states that if q is conjugate to p along a geodesic γ and the Gaussian curvature along γ is bounded above by G , then $\ell(\gamma) \geq \pi/\sqrt{G}$. In other words, $\ell(\gamma) \geq \pi \inf_{z \in \gamma} \rho_G(z)$, and since γ is a compact set, there will be a point that attains the bound. Thus we can state the theorem in a more convenient form:

Lemma 6 (Morse, Schönberg) If a geodesic γ connecting p to q contains a point conjugate to one of its endpoints, or if

p and q are conjugate along γ , then

$$\exists z \in \gamma \text{ such that } \rho_G(z) \leq \frac{\ell(\gamma)}{\pi}.$$

For case (ii) of Lemma 3 there is an extension to open curves of a famous theorem by Fenchel that we can exploit. Let γ be a smooth space curve from p to q . Let v be the vector from p to q in \mathbb{R}^3 . Denote by α and β the angles $\angle(\gamma'(0), v)$ and $\angle(\gamma'(\ell(\gamma)), v)$ respectively, and let $k_\gamma(t) = \|\gamma''(t)\|$ be the curvature of γ at $\gamma(t)$. Then the inequality, which is referred to in [Top06][p. 56] as the Fenchel-Reshetnyak inequality, states that

$$\int_\gamma k_\gamma(t) dt \geq \alpha + \beta. \tag{6}$$

The Fenchel-Reshetnyak inequality applies to non-closed curves, however, if we have $\gamma(0) = \gamma(\ell(\gamma)) = p$, then we can break γ into two pieces and obtain a curvature bound by applying Equation (6) to each piece.

Choose a point q on γ and let γ_1 be the portion of γ from p to q and let γ_2 be the remaining portion from q back to p . Denote the associated initial and final angles by α_1, β_1 and α_2, β_2 respectively. Then $\beta_1 + \alpha_2 = \pi$ and

$$\begin{aligned} \int_\gamma k_\gamma(t) dt &= \int_{\gamma_1} k_{\gamma_1}(t) dt + \int_{\gamma_2} k_{\gamma_2}(t) dt \geq \alpha_1 + \beta_1 + \alpha_2 + \beta_2 \\ &\geq (\pi - \angle(\gamma'(0), \gamma'(\ell(\gamma)))) + \pi. \end{aligned}$$

All we can assert about $(\pi - \angle(\gamma'(0), \gamma'(\ell(\gamma))))$ is that it is positive. Thus for an arbitrary loop curve γ we have that the total curvature is greater than π . It follows that there must be some point on γ where the curvature is greater than $\pi/\ell(\gamma)$. Applied to geodesics, this gives us the following:

Lemma 7 (Fenchel, Reshetnyak) On a geodesic loop γ ,

$$\exists z \in \gamma \text{ such that } \rho_\kappa(z) < \frac{\ell(\gamma)}{\pi}.$$

We obtain an estimate on the injectivity radius:

Theorem 3 (Injectivity radius estimate) For all $x \in S$

$$\rho_i(x) \geq \left(\frac{\pi}{2 + \pi} \right) \rho_f(x).$$

Proof. In the first case of Lemma 3, we have from Lemma 6 and Equation (2) that there exists a point z in $\bar{B}_S(x; \rho_i(x))$ with $\rho_f(z) \leq \rho_G(z) \leq \rho_i(x)/\pi$. Using the Lipschitz continuity of lfs (1), we have $\rho_f(x) \leq \rho_f(z) + d_S(x, z) \leq \frac{\rho_i(x)}{\pi} + \rho_i(x)$, and so $\rho_i(x) \geq \left(\frac{\pi}{1 + \pi} \right) \rho_f(x)$.

In the second case, we have by Lemma 7, a z in $\bar{B}_S(x; \rho_i(x))$ with $\rho_f(z) \leq \rho_\kappa(z) \leq 2\rho_i(x)/\pi$, and Lipschitz continuity yields $\rho_i(x) \geq \left(\frac{\pi}{2 + \pi} \right) \rho_f(x)$. Taking the smaller of the two bounds gives our estimate on $\rho_i(x)$. \square

Our estimate on the scr is based on a theorem in [Cha06][IX.6.1, p.404] which gives a global lower bound

for the scr in terms of global bounds on the injectivity radius and the Gaussian curvature. The following lemma is extracted from Chavel's proof:

Lemma 8 (Chavel) Let xpq be a geodesic triangle consisting of geodesics γ_1, γ_2 and γ_3 connecting xp, pq and qx respectively. Suppose there is a constant $r_0 > 0$, such that $\rho_G(z) \geq r_0$ on $B_S(x; \rho_i(x))$. Suppose also that $\sum_{i=1}^3 \ell(\gamma_i) < \min\{2\rho_i(x), 2\pi r_0\}$. Then

$$\gamma_2 \subset B_S(x; \rho),$$

where $\rho = \max\{d_S(x, p), d_S(x, q)\}$.

This result, together with Lemmas 6 and 7, yields the main result of this section:

Theorem 4 (scr estimate) For all $x \in S$,

$$\rho_{sc}(x) \geq \left(\frac{\pi}{4+3\pi}\right) \rho_f(x).$$

Proof. Consider the geodesic disk of radius r centred at x . There are three ways in which $B_S(x; r)$ can fail to be strongly convex: There exist $p, q \in B_S(x; r)$ such that either

- (i) the minimizer γ connecting p and q is not unique, or
- (ii) in addition to γ , there is another geodesic α connecting p and q and contained in $B_S(x; r)$, or
- (iii) γ is not contained in $B_S(x; r)$.

Case (i) cannot happen if $d_S(p, q) < \max\{\rho_i(p), \rho_i(q)\}$, and possibility (ii) is eliminated if $B_S(x; r) \subset B_S(p; \rho_i(p))$ for any $p \in B_S(x; r)$. Thus we eliminate the first two cases if we ensure that $\rho_i(p) \geq 2r$ for all $p \in B_S(x; r)$.

If $\rho_i(p) < 2r$, then by the Klingenberg Lemma 3 either

- (a) there is a $z \in B_S(p; 2r) \subset B_S(x; 3r)$ that is conjugate to p along a minimizing geodesic, or
- (b) there is a geodesic loop in $B_S(p; 2r)$.

In case (a), the Morse-Schönberg Lemma 6 gives us a $z \in B_S(x; 3r)$ with $\rho_f(z) \leq \frac{2r}{\pi}$. The Lipschitz continuity of lfs yields $\rho_f(x) \leq \rho_f(z) + d_S(x, z) \leq \frac{2r}{\pi} + 3r$. Thus

$$r \geq \left(\frac{\pi}{2+3\pi}\right) \rho_f(x). \quad (7)$$

In case (b), the Fenchel-Reshetnyak Lemma 7 yields a $z \in B_S(x; 3r)$ with $\rho_f(z) < \frac{4r}{\pi}$. Again using Lipschitz continuity to bring the lfs bound to x , we obtain

$$r > \left(\frac{\pi}{4+3\pi}\right) \rho_f(x). \quad (8)$$

Thus if we ensure that r is smaller than the bounds (7) and (8) then cases (i) and (ii) cannot happen and we need only consider case (iii). For this we turn to Chavel's Lemma 8.

Consider the geodesic triangle xpq consisting of minimal geodesics, with notation as in Lemma 8. By hypothesis we now have $\rho_i(x) \geq 2r$, so $\sum_{i=1}^3 \ell(\gamma_i) < 4r \leq 2\rho_i(x)$. Thus the conditions of Lemma 8 are satisfied unless there is a $z \in$

$B_S(x; 2r)$ with $2\pi\rho_G(z) < 4r$. This would imply $\rho_f(z) < \frac{2r}{\pi}$, and the Lipschitz shuffle to x yields

$$r > \frac{1}{2} \left(\frac{\pi}{1+\pi}\right) \rho_f(x). \quad (9)$$

By (8), the smaller of the three estimates, we have that $B_S(x; r)$ is strongly convex whenever $r \leq \left(\frac{\pi}{4+3\pi}\right) \rho_f(x)$, and we obtain the theorem by the definition of the scr. \square

Since the constant in the scr bound is less than half of that of the injectivity bound, we can use it also as a bound on the intrinsic sampling radius. Thus for all $x \in S$

$$\rho_m(x) \geq \left(\frac{\pi}{4+3\pi}\right) \rho_f(x). \quad (10)$$

4.3. Bounding geodesic lengths

Since Euclidean distances between two points never exceed the geodesic distances, $p \in B_S(x; \rho)$ implies $p \in B_{\mathbb{R}^3}(x; \rho)$. However, we need to make claims about the containment of points within geodesic disks, given their presence within a Euclidean ball. Estimates on $d_S(x, p)$ relative to $d_{\mathbb{R}^3}(x, p)$, for $x, p \in S$ sufficiently close, are provided in the works of [MT04] and [HPW06]. We follow the terminology and notation of the former.

We exploit the projection mapping, $\xi: U_{\rho_R} \rightarrow S$ discussed in Section 3.1. The *relative curvature*, $\omega(m)$, at a point $m \in U_{\rho_R}$ is defined as

$$\omega(m) = \frac{d_{\mathbb{R}^3}(m, \xi(m))}{\rho_\kappa(\xi(m))}.$$

From the definition of U_{ρ_R} , $d_{\mathbb{R}^3}(m, \xi(m)) \leq \rho_R(\xi(m))$, so Equation (2) gives $\omega(m) \leq 1$.

Suppose that $p \in B_{\mathbb{R}^3}(x; \varepsilon_f \rho_f(x))$ and let $I =]x, p[$ be the open Euclidean line segment between x and p , and let $\omega = \sup_{m \in I} \omega(m)$. Then according to [MT04]:

$$\ell(\xi(I)) \leq \frac{1}{1-\omega} d_{\mathbb{R}^3}(x, p). \quad (11)$$

If γ is a minimal geodesic between x and p , then $d_S(x, p) = \ell(\gamma) \leq \ell(\xi(I))$ gives us the needed bound.

Since ξ takes m to the *closest* point on S , we have, for all $m \in I$, $d_{\mathbb{R}^3}(m, \xi(m)) \leq \frac{1}{2} d_{\mathbb{R}^3}(x, p) \leq \frac{1}{2} \varepsilon_f \rho_f(x)$. For the denominator of ω we have $\rho_\kappa(\xi(m)) \geq \rho_f(\xi(m))$, and by the Lipschitz continuity of lfs (1), $\rho_f(\xi(m)) \geq (1 - \varepsilon_f) \rho_f(x)$. Thus $\omega \leq \frac{\varepsilon_f}{2(1-\varepsilon_f)}$. For the estimate to be usable, we need $\omega \leq 1$, so we demand $\varepsilon_f < 2/3$.

Plugging this estimate into (11) together with $d_{\mathbb{R}^3}(x, p) < \varepsilon_f \rho_f(x)$ yields the needed bound on the geodesic length:

Lemma 9 If $p \in S \cap B_{\mathbb{R}^3}(x; \varepsilon_f \rho_f(x))$, with $\varepsilon_f < 2/3$, then $p \in B_S(x; \tilde{\varepsilon} \rho_f(x))$ for

$$\tilde{\varepsilon} \leq \frac{\varepsilon_f(1-\varepsilon_f)}{1-\frac{3}{2}\varepsilon_f}.$$

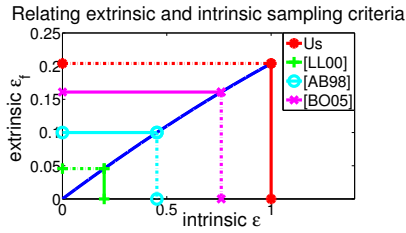


Figure 4: The ϵ for intrinsic sampling (4) is on the horizontal axis. The required ϵ_f for equation (3) is on the vertical axis. Sampling criteria from other works are compared.

4.4. Extrinsic criteria meeting intrinsic demands

Equipped with Equation (10) and Lemma 9 we determine that if P satisfies equation (3), then it will also satisfy equation (4) provided that

$$\frac{\epsilon_f(1 - \epsilon_f)}{1 - \frac{3}{2}\epsilon_f} \leq \epsilon \frac{\pi}{4 + 3\pi}.$$

Putting $C = \epsilon \frac{\pi}{4 + 3\pi}$, we get

$$\epsilon_f^2 - (1 + \frac{3}{2}C)\epsilon_f + C \geq 0,$$

an inequality that will be satisfied whenever ϵ_f is smaller than the smaller of the two positive roots.

For our sampling criteria of Theorem 2, $\epsilon = 1$, yielding $C \approx 0.234$ and $\epsilon_f \leq 0.204$ is required. This compares well with existing lfs sampling requirements for topological consistency. For example, $\epsilon_f \leq 0.1$ is required in [AB98]. In [BO05], a loose ϵ -sample is required to have $\epsilon \leq 0.091$. A loose ϵ -sample only requires samples to lie within $B_{\mathbb{R}^3}(c; \epsilon \rho_f(c))$ when c is a vertex of the rVd. According to Corollary 4.10 of that work, such a sampling will be an ϵ_f -sampling in the sense of Equation (3) for $\epsilon_f \approx 0.161$. These comparisons are summarized in Figure 4.

5. Conclusions

Through an analysis of the iVd, we improved upon the sampling criteria of [LL00, Lei99]. By deriving inequalities relating the injectivity radius and the strong convexity radius to the local feature size, we have enabled comparison between sampling criteria in the intrinsic and extrinsic domains.

The indication is that sharper bounds may result from an intrinsic analysis even if an algorithm is based on an extrinsic model. However, we have not demonstrated that the new sampling conditions are sufficient to ensure that the iDt-mesh is a substructure of the 3D Delaunay tetrahedralization. Since good convergence properties of the iDt-mesh have been demonstrated [DLYG06], this would put it on a more or less equal footing with the rDt. It is known that the rDt and the iDt-mesh are not necessarily combinatorially equivalent, regardless of sampling density [DZM07].

The estimates on the injectivity radius and the scr in Section 4.2 can probably be tightened. The constant in Equation (10) may be improved from $\frac{\pi}{4+3\pi}$ to $\frac{1}{2} \left(\frac{\pi}{1+\pi} \right)$ if it can be shown that: if distinct geodesics α and γ connect p and q , then there is a point $z \in \alpha \cup \gamma$ with $\rho_R(z) \leq \frac{\ell(\alpha) + \ell(\gamma)}{2\pi}$. Furthermore, we would no longer rely on Lemmas 6, 7, or 3.

Although the statements of Theorems 1 and 2 make sense in higher dimensions, their proofs are not easily extended. However, extending the proofs of Theorems 3 and 4 may require no more than a reworking of Section 3.1.

References

- [AB98] AMENTA N., BERN M. W.: Surface Reconstruction by Voronoi Filtering. In *Symp. Comp. Geom.* (1998), pp. 39–48.
- [BO05] BOISSONNAT J.-D., OUDOT S.: Provably good sampling and meshing of surfaces. *Graphical Models* 67, 5 (2005), 405–451.
- [CDRR04] CHENG S.-W., DEY T. K., RAMOS E. A., RAY T.: Sampling and meshing a surface with guaranteed topology and geometry. In *Symp. Comp. Geom.* (2004), pp. 280–289.
- [Cha06] CHAVEL I.: *Riemannian Geometry, A modern introduction*, 2nd ed. Cambridge, 2006.
- [dC92] DO CARMO M. P.: *Riemannian Geometry*. Birkhäuser, 1992.
- [DLYG06] DAI J., LUO W., YAU S.-T., GU X.: Geometric accuracy analysis for discrete surface approximation. In *Geometric Modeling and Processing* (2006), pp. 59–72.
- [DZM07] DYER R., ZHANG H., MÖLLER T.: Voronoi-Delaunay duality and Delaunay meshes. In *Symp. Solid and Physical Modeling* (2007), pp. 415–420.
- [ES94] EDELSBRUNNER H., SHAH N. R.: Triangulating topological spaces. In *Symp. Comp. Geom.* (1994), pp. 285–292.
- [Fed59] FEDERER H.: Curvature measures. *Trans. Amer. Math. Soc.* 93 (1959), 418–491.
- [GGOW08] GAO J., GUIBAS L. J., OUDOT S. Y., WANG Y.: Geodesic delaunay triangulation and witness complex in the plane. In *Symp. on Discrete Algorithms* (2008), pp. 571–580.
- [HPW06] HILDEBRANDT K., POLTHIER K., WARDETZKY M.: On the Convergence of Metric and Geometric Properties of Polyhedral Surfaces. *Geometriae Dedicata* 123, 1 (2006), 89–112.
- [Kli95] KLINGENBERG W. P. A.: *Riemannian Geometry*, second ed. Walter de Gruyter, 1995.
- [Lei99] LEIBON G.: *Random Delaunay triangulations, the Thurston-Andreev theorem, and metric uniformization*. PhD thesis, UCSD, 1999. arXiv:math/0011016v1.
- [LL00] LEIBON G., LETSCHER D.: Delaunay Triangulations and Voronoi Diagrams for Riemannian Manifolds. In *Symp. Comp. Geom.* (2000), pp. 341–349.
- [MT04] MORVAN J.-M., THIBERT B.: Approximation of the normal vector field and the area of a smooth surface. *Discrete & Computational Geometry* 32, 3 (2004), 383–400.
- [Top06] TOPONOGOV V. A.: *Differential Geometry of Curves and Surfaces*. Birkhäuser, 2006.