Project number IST-25582

CGL

Computational Geometric Learning

Delaunay-type structures for manifolds I: Stability

STREP

Information Society Technologies

Period covered:	November 1, 2011–October 31, 2012
Date of preparation:	June 15, 2012
Start date of project:	November 1, 2010
Duration:	3 years
Project coordinator name:	Joachim Giesen (FSU)
Project coordinator organisation:	Friedrich-Schiller-Universität Jena Jena, Germany

Contents

Delaunay-type structures for manifolds I: Stability

Jean-Daniel Boissonnat, Ramsay Dyer, and Arijit Ghosh*

June 15, 2012

Abstract

We introduce a parametrized notion of genericity for Delaunay triangulations which, in particular, implies that the Delaunay simplices of δ-generic point sets are thick. Equipped with this notion, we study the stability of Delaunay triangulations under perturbations of the metric and of the vertex positions. We then show that, for any sufficiently regular submanifold of Euclidean space, and appropriate ϵ and δ, any sample set which meets a localized δ-generic ϵ-dense sampling criteria yields a manifold intrinsic Delaunay complex which is equal to the restricted Delaunay complex.

1 Introduction

One of the central properties of Delaunay complexes, which was demonstrated when they were introduced [Del34], is that under very mild assumptions they triangulate Euclidean space. This paper addresses issues that arise when the Delaunay paradigm is employed for triangulating non-Euclidean manifolds whose dimension may exceed two.

For a submanifold of Euclidean space, the restricted Delaunay complex [ES97], which is defined by the ambient metric restricted to the submanifold, was employed by Cheng et al. [CDR05] as the basis for a triangulation. However, it was found that sampling density alone was insufficient to ensure a triangulation, and manipulations of the complex were employed.

In an earlier work, Leibon and Letscher [LL00] announced sampling density conditions which would ensure that the Delaunay complex defined by the intrinsic metric of the manifold was a triangulation. In fact, as shown in Appendix A, the stated result is incorrect: sampling density alone is insufficient to guarantee an intrinsic Delaunay triangulation (see Theorem A.3). Topological defects can arise when the vertices lie too close to a degenerate or "quasi-cospherical" configuration.

We address this problem with the introduction of a parameterized notion of genericity for Delaunay complexes. Our interest in the intrinsic Delaunay complex stems from its close relationship with other Delaunay-like structures that have been proposed in the context of non-homogeneous metrics. For example, anisotropic Voronoi diagrams [LS03] and anisotropic Delaunay triangulations emerge as natural structures when we want to mesh a domain of $\mathbb{R}^m$ while respecting a given metric tensor field.

*This research has been partially supported by the 7th Framework Programme for Research of the European Commission, under FET-Open grant number 255827 (CGL Computational Geometry Learning).

This paper builds over preliminary results on anisotropic Delaunay meshes [BWY11] and manifold reconstruction using the tangential Delaunay complex [BG11]. The central idea in both cases is to define Euclidean Delaunay triangulations locally and to glue these local triangulations together by removing inconsistencies between them. We view the inconsistencies as arising from instability in the Delaunay triangulations, and provide explicit bounds on the stability with respect to the genericity parameter.

The stability of Delaunay triangulations has not previously been studied in this way. Related work can be found in the context of kinetic data structures [AGG$^+$10] or in the context of robust computation [BS04], and in particular, the concept of protection we introduce in Section 3 is embodied in the guarded insphere predicate which has been employed in a controlled perturbation algorithm for 2D Delaunay triangulation [FKMS05].

Although we make no explicit reference to Voronoi diagrams, the Delaunay complexes we study can be equivalently defined as the nerve of the Voronoi diagram associated with the metric under consideration. We provide criteria for ensuring that the Delaunay complex is a manifold without explicit requirements on the properties of the Voronoi diagram [ES97], as has often been done in related work [LL00, LS03, CDR05, DZM08, CG12].

After presenting background material in Section 2, we introduce the concept of δ-generic point sets for Euclidean Delaunay triangulations in Section 3. We show that Delaunay simplices of δ-generic point sets are thick; they satisfy a quality bound. Then in Section 4 we quantify how δ leads to robustness in the Delaunay triangulation when either the points or the metric are perturbed. The primary challenge is bounding the displacement of simplex circumcentres. Finally, we use these results in Section 5 to demonstrate conditions under which the intrinsic Delaunay complex and the restricted Delaunay complex coincide and are manifold.

2 Background

Within the context of the standard m-dimensional Euclidean space $\mathbb{R}^m$, when distances are determined by the standard norm, $\|\cdot\|$, we use the following conventions. The distance between a point p and a set $X \subset \mathbb{R}^m$, is the infimum of the distances between p and the points of X, and is denoted $d_{\mathbb{R}^m}(p, X)$. We refer to the distance between two points a and b as $\|b - a\|$ or $d_{\mathbb{R}^m}(a, b)$ as convenient. A ball $B_{\mathbb{R}^m}(c; r) = \{x \mid \|x - c\| < r\}$ is open, and $\overline{B}_{\mathbb{R}^m}(c; r)$ is its topological closure. Generally, we denote the topological closure of a set X by $\overline{X}$, the interior by $\operatorname{int} X$, and the boundary by ∂X. The convex hull is denoted $\operatorname{conv}(X)$, and the affine hull is $\operatorname{aff}(X)$.

We will make use of other metrics besides the Euclidean one. A generic metric is denoted d, and the associated open and closed balls are $B(c; r)$, and $\overline{B}(c; r)$. If a specific metric is intended, it will be indicated by a subscript, for example in Section 5 we introduce $d_{\mathcal{M}}$, the intrinsic metric on a manifold $\mathcal{M}$, which has associated balls $B_{\mathcal{M}}(c; r)$.

If A is a $k \times j$ matrix, we denote its i^{th} singular value by $s_i(A)$. We use the operator norm $\|A\| = s_1(A) = \sup_{\|x\|=1} \|Ax\|$, and employ the following standard observation:

Lemma 2.1 If $\eta > 0$ is an upper bound on the norms of the columns of A, then $\|A\| \leq \sqrt{j}\eta$.

We will also be interested in obtaining a lower bound on the smallest singular value, for which the following observation is useful:

Lemma 2.2 If A is a $k \times j$ matrix of rank $j \leq k$, then the *pseudo inverse* $A^{\dagger} = (A^{\mathsf{T}}A)^{-1}A^{\mathsf{T}}$ is the unique left inverse of A whose kernel is the orthogonal complement of the column space of A. Furthermore,

$$s_i(A^{\dagger}) = s_{j-i+1}(A)^{-1}.$$

If U and V are vector subspaces of $\mathbb{R}^m$, with $\dim U \leq \dim V$, the *angle* between them is defined by

$$\cos \angle(U, V) = \inf_{u \in U} \sup_{v \in V} \frac{u^{\mathsf{T}} v}{\|u\| \, \|v\|}.$$

This is the largest principal angle between U and V. The angle between affine subspaces K and H is defined as the angle between the corresponding parallel vector subspaces.

2.1 Sampling parameters and perturbations

The structures of interest will be built from a finite set $\mathsf{P} \subset \mathbb{R}^m$, which we consider to be a set of *sample points*. If $D \subset \mathbb{R}^m$ is a bounded set, then P is an *ϵ-sample set* for D if $d_{\mathbb{R}^m}(x, \mathsf{P}) < \epsilon$ for all $x \in \overline{D}$. We say that ϵ is a *sampling radius* for D satisfied by P. If no domain D is specified, we say P is an ϵ-sample set if $d_{\mathbb{R}^m}(x, \mathsf{P} \cup \partial \mathrm{conv}(\mathsf{P})) < \epsilon$ for all $x \in \mathrm{conv}(\mathsf{P})$. Equivalently, P is an ϵ-sample set if it satisfies a sampling radius ϵ for

$$D_{\epsilon}(\mathsf{P}) = \{x \in \mathrm{conv}(\mathsf{P}) \,|\, d_{\mathbb{R}^m}(x, \partial \mathrm{conv}(\mathsf{P})) \geq \epsilon\}.$$

The set P is *λ-sparse* if $d_{\mathbb{R}^m}(p, q) > \lambda$ for all $p, q \in \mathsf{P}$. We usually assume that the sparsity of a ϵ-sample set is proportional to ϵ, thus: $\lambda = \mu_0 \epsilon$.

We consider a perturbation of the points $\mathsf{P} \subset \mathbb{R}^m$ given by a function $\zeta : \mathsf{P} \to \mathbb{R}^m$. If ζ is such that $d_{\mathbb{R}^m}(p, \zeta(p)) \leq \rho$, we say that ζ is a *ρ-perturbation*. As a notational convenience, we frequently define $\tilde{\mathsf{P}} = \zeta(\mathsf{P})$, and let $\tilde{p}$ represent $\zeta(p) \in \tilde{\mathsf{P}}$. We will only be considering ρ-perturbations where ρ is less than half the sparsity of P, so $\zeta : \mathsf{P} \to \tilde{\mathsf{P}}$ is a bijection.

Points in P which are not on the boundary of $\mathrm{conv}(\mathsf{P})$ are *interior points* of P.

Lemma 2.3 Suppose P is an ϵ-sample set, and $\zeta : \mathsf{P} \to \tilde{\mathsf{P}}$ is a ρ-perturbation with $2\rho \leq \epsilon$. If point $p \in \mathsf{P}$ satisfies $d_{\mathbb{R}^m}(p, \partial \mathrm{conv}(\mathsf{P})) \geq 3\epsilon$, then $\tilde{p} = \zeta(p)$ is an interior point of $\tilde{\mathsf{P}}$.

Proof Let $S = \partial B$ be the bounding sphere for $B = B_{\mathbb{R}^m}(\tilde{p}; 3\epsilon/2)$. Then $d_{\mathbb{R}^m}(p, S) \geq \epsilon$ and for any $x \in S$, $d_{\mathbb{R}^m}(x, \partial \mathrm{conv}(\mathsf{P})) \geq \epsilon$. Thus the sampling assumption ensures that for every point $x \in S$, there is a point $q \in \mathsf{P}$ with $p \neq q$ and $d_{\mathbb{R}^m}(x, q) < \epsilon$. It follows that $d_{\mathbb{R}^m}(x, \zeta(q)) < 3\epsilon/2$, and thus that $\tilde{p}$ is not the closest point in $\tilde{\mathsf{P}}$ for any point on S.

Thus $\tilde{p}$ cannot belong to $\partial \mathrm{conv}(\tilde{\mathsf{P}})$. Indeed if $\tilde{p} \in \partial \mathrm{conv}(\tilde{\mathsf{P}})$, then take a unit vector v in an outward direction orthogonal to a closed half-space supporting $\mathrm{conv}(\tilde{\mathsf{P}})$ at $\tilde{p}$. The ray from $\tilde{p}$ defined by v must intersect S at some point y, and $\tilde{p}$ would be the closest point in $\tilde{\mathsf{P}}$ to y, a contradiction. □

2.2 Simplices

Given a set of $j+1$ points $\{p_0, \dots, p_j\} \subset \mathsf{P} \subset \mathbb{R}^m$, a (geometric) *$j$-simplex* $\sigma = [p_0, \dots, p_j]$ is defined by the convex hull: $\sigma = \operatorname{conv}(\{p_0, \dots, p_j\})$. The points p_i are the *vertices* of σ. Any subset $\{p_{i_0}, \dots, p_{i_k}\}$ of $\{p_0, \dots, p_j\}$ defines a k-simplex τ which we call a *face* of σ. We write $\tau \leq \sigma$ if τ is a face of σ, and $\tau < \sigma$ if τ is a *proper face* of σ, i.e., if the vertices of τ are a proper subset of the vertices of σ.

The *boundary* of σ, is the union of its proper faces: $\partial\sigma = \bigcup_{\tau<\sigma} \tau$. In general this is distinct from the topological boundary defined above, but we denote it with the same symbol. The *interior* of σ is $\operatorname{int}\sigma = \sigma \setminus \partial\sigma$. Again this is generally different from the topological interior. Other geometric properties of σ include its diameter (its longest edge), $\Delta(\sigma)$, and its shortest edge, $L(\sigma)$.

For any vertex $p \in \sigma$, the *face oppposite* p is the face determined by the other vertices of σ, and is denoted σ_p. If τ is a j-simplex, and p is not a vertex of τ, we may construct a $(j+1)$-simplex $\sigma = p * \tau$, called the *join* of p and τ. It is the simplex defined by p and the vertices of τ, i.e., $\tau = \sigma_p$.

Our definition of a simplex has made an important departure from standard convention: we do not demand that the vertices of a simplex be affinely independent. A j-simplex σ is a *degenerate simplex* if $\dim \operatorname{aff}(\sigma) < j$. If we wish to emphasise that a simplex is a j-simplex, we write j as a superscript: σ^j; but this always refers to the *combinatorial* dimension of the simplex.

If σ is non-degenerate, then it has a *circumcentre*, $C(\sigma)$, which is the centre of the smallest circumscribing ball for σ. The radius of this ball is the *circumradius* of σ, denoted $R(\sigma)$. In the context of the Euclidean Delaunay complexes we will work with, the degenerate simplices we may encounter also have these properties. The ratio of the circumradius to the shortest edge is denoted $\Phi(\sigma) = R(\sigma)/L(\sigma)$. We will make use of the affine space $N(\sigma)$ composed of the centres of the balls that circumscribe σ. This space is orthogonal to $\operatorname{aff}(\sigma)$ and intersects it at the circumcentre of σ. Its dimension is $m - \dim \operatorname{aff}(\sigma)$.

The *altitude* of p in σ is $D(p,\sigma) = d_{\mathbb{R}^m}(p, \operatorname{aff}(\sigma_p))$. A poorly-shaped simplex can be characterized by the existence of a relatively small altitude. The *thickness* of a j-simplex σ is the dimensionless quantity

$$\Upsilon(\sigma) = \begin{cases} 1 & \text{if } j = 0 \\ \min_{p\in\sigma} \frac{D(p,\sigma)}{j\Delta(\sigma)} & \text{otherwise.} \end{cases}$$

We say that σ is Υ_0-thick, if $\Upsilon(\sigma) \geq \Upsilon_0$. If σ is Υ_0-thick, then so are all of its faces. Indeed if $\tau \leq \sigma$, then the smallest altitude in τ cannot be smaller than that of σ, and also $\Delta(\tau) \leq \Delta(\sigma)$.

Other parameters such as the volume [Whi57], or the radius of the largest contained ball centred at the barycentre [Mun68], can be used to characterize simplex quality. We find a direct bound on the altitudes to be more convenient, in part due to the following consequence of Lemma 2.2:

Lemma 2.4 (Thickness and singular value) Let $\sigma = [p_0, \dots, p_j]$ be a non-degenerate j-simplex in $\mathbb{R}^m$, with $j > 0$, and let P be the $m \times j$ matrix whose i^{th} column is $p_i - p_0$. Then

$$s_j(P) \geq \sqrt{j}\,\Upsilon(\sigma)\Delta(\sigma).$$

Proof Let w_i^{T} be the i^{th} row of $P^{\dagger}$. Then w_i belongs to the column space of P, and it is orthogonal to all $(p_{i'} - p_0)$ for $i' \neq i$. Let $u_i = w_i / \|w_i\|$. It follows from the definition that $u_i^\mathsf{T}(p_i - p_0) = D(p_i, \sigma)$. Thus $w_i = D(p_i, \sigma)^{-1} u_i$. Since $s_i(A^\mathsf{T}) = s_i(A)$ for any matrix A, we have

$$s_1(P^\dagger) \leq \sqrt{j} \max_{1 \leq i \leq j} D(p_i, \sigma)^{-1},$$

by Lemma 2.1. Thus Lemma 2.2 yields

$$s_j(P) \geq \frac{1}{\sqrt{j}} \min_{1 \leq i \leq j} D(p_i, \sigma) = \sqrt{j}\Upsilon(\sigma)\Delta(\sigma).$$

□

The proof of Lemma 2.4 shows that the pseudoinverse of P has a natural geometric interpretation in terms of the altitudes of σ, and thus the altitudes provide a convenient lower bound on $s_j(P)$. By Lemma 2.1, $s_1(P) \leq \sqrt{j}\Delta(\sigma)$, and thus $\Upsilon(\sigma) \leq \frac{s_j(P)}{s_1(P)}$. In other words, $\Upsilon(\sigma)^{-1}$ provides a convenient upper bound on the *condition number* of P.

Whitney [Whi57, p. 127] proved that the affine hull of a thick simplex makes a small angle with any hyperplane which lies near all the vertices of the simplex. Employing Lemma 2.4 in the proof of Whitney's Lemma yields a simpler proof and a sharper result:

Lemma 2.5 (Whitney angle bound) Suppose σ is a j-simplex whose vertices all lie within a distance η from a hyperplane, $H \subset \mathbb{R}^m$. Then

$$\sin \angle(\operatorname{aff}(\sigma), H) \leq \frac{2\eta}{\Upsilon(\sigma)\Delta(\sigma)}.$$

Proof Suppose $\sigma = [p_0, \dots, p_j]$. Choose p_0 as the origin of $\mathbb{R}^m$, and let $U \subset \mathbb{R}^m$ be the vector subspace defined by $\operatorname{aff}(\sigma)$. Let W be the $(m-1)$-dimensional subspace parallel to H, and let $\pi : \mathbb{R}^m \to W$ be the orthogonal projection onto W.

For any unit vector $u \in U$, $\sin \angle(\operatorname{aff}(\sigma), H) = \sin \angle(U, W) \leq \|u - \pi u\|$. Since the vectors $v_i = (p_i - p_0)$, $i \in \{1, \dots, j\}$ form a basis for U, we may write $u = Pa$, where P is the $m \times j$ matrix whose i^{th} column is v_i, and $a \in \mathbb{R}^j$ is the vector of coefficients. Then, defining $X = P - \pi P$, we get

$$\|u - \pi u\| = \|Xa\| \leq \|X\| \, \|a\| .$$

Since $d_{\mathbb{R}^m}(p_i, H) \leq \eta$ for all $0 \leq i \leq j$, W is at a distance less than η from H, and $\|v_i - \pi v_i\| \leq 2\eta$. It follows then from Lemma 2.1 that

$$\|X\| \leq 2\sqrt{j}\eta.$$

Observing that $1 = \|u\| = \|Pa\| \geq s_j(P) \, \|a\|$, we find

$$\|a\| \leq \frac{1}{s_j(P)},$$

and the result follows from Lemma 2.4. □

2.3 Complexes

Given a finite set P, an *abstract simplicial complex* is a set of subsets $K \subset 2^{\mathsf{P}}$ such that if $\sigma \in K$, then every subset of σ is also in K. The Delaunay complexes we study are abstract simplicial complexes, but their simplices carry a canonical geometry induced from the inclusion map $\iota : \mathsf{P} \hookrightarrow \mathbb{R}^m$. (We assume ι is injective on P, and so do not distinguish between P and $\iota(\mathsf{P})$.) To each abstract simplex $\sigma \in K$, we have an associated geometric simplex $\mathrm{conv}(\iota(\sigma))$, and normally when we write $\sigma \in K$, we are referring to this geometric object. Occasionally, when it is convenient to emphasise a distinction, we will write $\iota(\sigma)$ instead of σ.

Thus we view such a K as a set of simplices in $\mathbb{R}^m$, and we refer to it as a *complex*, but it is not generally a (geometric) simplicial complex. A geometric *simplicial complex* is a finite collection G of non-degenerate simplices in $\mathbb{R}^N$ such that if $\sigma \in G$, then all of the faces of σ also belong to G, and if $\sigma, \tilde{\sigma} \in G$ and $\tau = \sigma \cap \tau \neq \emptyset$, then $\tau \leq \sigma$ and $\tau \leq \tilde{\sigma}$. An abstract simplicial complex is defined from a geometric simplicial complex in an obvious way. A *geometric realization* of an abstract simplicial complex K is a geometric simplicial complex whose associated abstract simplicial complex may be identified with K. A geometric realization always exists for any complex. Details can be found in algebraic topology textbooks; the book by Munkres [Mun84] for example.

The *carrier* of an abstract complex K is the underlying topological space $|K|$, associated with a geometric realization of K. Thus if G is a geometric realization of K, then $|K| = \bigcup_{\sigma \in G} \sigma$. For our complexes, the inclusion map ι induces a continous map $\iota : |K| \to \mathbb{R}^m$, defined by barycentric interpolation on each simplex. If this map is injective, we say that K is *embedded*. In this case ι also defines a geometric realization of K, and we may identify the carrier of K with the image of ι.

A subset $K' \subset K$ is a *subcomplex* of K if it is also a complex. The *star* of a subcomplex $K' \subseteq K$ is the subcomplex generated by the simplices incident to K'. I.e., it is all the simplices that share a face with a simplex of K', plus all the faces of such simplices. This is a departure from a common usage of this same term in the topology literature. The star of K' is denoted $\mathrm{star}(K')$ when there is no risk of ambiguity, otherwise we also specify the parent complex, as in $\mathrm{star}(K'; K)$.

A *triangulation* of $\mathsf{P} \subset \mathbb{R}^m$ is an embedded complex K with vertices P such that $|K| = \mathrm{conv}(\mathsf{P})$. A complex K is a *j-manifold complex* if the star of every vertex is isomorphic to the star of a triangulation of $\mathbb{R}^j$. In order to exploit the local nature of the definition of a manifold complex, it is convenient to have a local notion of triangulation for the star of a vertex in K, even if the whole of K is not a triangulation of its vertices:

Definition 2.6 (Triangulation at a point) A complex K is a *triangulation at* $p \in \mathbb{R}^m$ if:

- p is a vertex of K.
- $\mathrm{star}(p)$ is embedded.
- p lies in $\mathrm{int}\,|\mathrm{star}(p)|$.
- For all $\tau \in K$, and $\sigma \in \mathrm{star}(p)$, if $\mathrm{int}\,\tau \cap \sigma \neq \emptyset$, then $\tau \in \mathrm{star}(p)$.

If σ is a simplex with vertices in P, then any map $\zeta : \mathsf{P} \to \tilde{\mathsf{P}} \subset \mathbb{R}^m$ defines a simplex $\zeta(\sigma)$ whose verticies in $\tilde{\mathsf{P}}$ are the images of vertices of σ. If K is a complex on P, and $\tilde{K}$ is a complex on $\tilde{\mathsf{P}}$, then ζ induces a *simplicial map* $K \to \tilde{K}$ if $\zeta(\sigma) \in \tilde{K}$ for every $\sigma \in K$. We denote this map by the same symbol, ζ. We are interested in the case when ζ is an *isomorphism*, which means it establishes a bijection between K and $\tilde{K}$. We then say that K and $\tilde{K}$ are *isomorphic*, and write $K \cong \tilde{K}$, or $K \overset{\zeta}{\cong} \tilde{K}$ if we wish to emphasise that the correspondence is given by ζ.

A simplicial map $\zeta : K \to \tilde{K}$ defines a continuous map $\zeta : |K| \to \left|\tilde{K}\right|$, by barycentric interpolation on each simplex $\sigma \in K$. We observe the following local version of a standard result:

Lemma 2.7 Suppose K is a complex with vertices $\mathsf{P} \subset \mathbb{R}^m$, and $\tilde{K}$ a complex with vertices $\tilde{\mathsf{P}} \subset \mathbb{R}^m$. Suppose also that K is a triangulation at $p \in \mathsf{P}$, and that $\zeta : \mathsf{P} \to \tilde{\mathsf{P}}$ induces an injective simplicial map $\mathrm{star}(p) \to \mathrm{star}(\zeta(p))$. If $\tilde{K}$ is a triangulation at $\zeta(p)$, then

$$\zeta(\mathrm{star}(p)) = \mathrm{star}(\zeta(p)).$$

Proof Since $\mathrm{star}(p)$ is embedded, ζ defines a continuous map $\zeta : |\mathrm{star}(p)| \to |\mathrm{star}(\zeta(p))|$ that is injective on each simplex. Since $\mathrm{star}(\zeta(p))$ is also embedded, this continuous map is injective on $|\mathrm{star}(p)|$. Since K is a triangulation at p, there is an open ball B centred at p such that $B \subset \mathrm{int}\,|\mathrm{star}(p)|$. Then $\zeta|_B : \mathbb{R}^m \supset B \to \zeta(B) \subset \mathbb{R}^m$ is a homeomorphism by Brouwer's invariance of domain [Dug66, Ch. XVII].

We need to show that $\mathrm{star}(\zeta(p)) \subseteq \zeta(\mathrm{star}(p))$. Suppose $\sigma \in \mathrm{star}(\zeta(p))$ and $\zeta(p)$ is a vertex of σ. Then, since σ is not degenerate, there is an $x \in B \cap \mathrm{int}\,\sigma$, and from the above argument, x also lies in some simplex $\tilde{\tau} \in \zeta(\mathrm{star}(p))$. Since $\mathrm{star}(\zeta(p))$ is embedded, $\tilde{\tau} \cap \sigma$ is a face of σ, but since x is in the interior of σ, it must be that $\tilde{\tau} = \sigma$. Thus $\sigma \in \zeta(\mathrm{star}(p))$.

If $\tau \in \mathrm{star}(\zeta(p))$, then there is some $\sigma \in \mathrm{star}(\zeta(p))$ such that $\zeta(p)$ is a vertex of σ and $\tau \leq \sigma$. Since $\sigma \in \zeta(\mathrm{star}(p))$, we also have $\tau \in \zeta(\mathrm{star}(p))$, by the definition of a simplicial map. □

3 Parameterized genericity

In this section we examine the Delaunay complex of $\mathsf{P} \subset \mathbb{R}^m$, taking the view that poorly-shaped simplices arise from almost degenerate configurations of points. We introduce the concept of a protected Delaunay ball, which leads to a parameterized definition of genericity. We then show that a lower bound on the protection of the maximal simplices yields a lower bound on their thickness.

3.1 The Delaunay complex

An *empty ball* is one that contains no point from P.

Definition 3.1 (Delaunay complex) A *Delaunay ball* is a maximal empty ball. Specifically, $B = B_{\mathbb{R}^m}(x; r)$ is a Delaunay ball if any empty ball centred at x is contained in B. A simplex σ is a *Delaunay simplex*, if there exists some Delaunay ball B such that the vertices of σ belong to $\partial B \cap \mathsf{P}$. The *Delaunay complex* is the set of Delaunay simplices, and is denoted $\mathrm{Del}(\mathsf{P})$.

The Delaunay complex has the combinatorial structure of an abstract simplicial complex, but $\mathrm{Del}(\mathsf{P})$ is embedded only when P satisfies appropriate genericity requirements, as discussed in Section 3.2. Otherwise, $\mathrm{Del}(\mathsf{P})$ contains degenerate simplices. We make here some observations that are not dependent on assumptions of genericity.

The union of the Delaunay simplices is $\mathrm{conv}(\mathsf{P})$. A simplex $\sigma \in \mathrm{Del}(\mathsf{P})$ is a *boundary simplex* if all its vertices lie on $\partial \mathrm{conv}(P)$. We observe

Lemma 3.2 (Maximal simplices) If $\mathrm{aff}(\mathsf{P}) = \mathbb{R}^m$, then every Delaunay j-simplex, σ, is a face of a Delaunay simplex σ' with $\dim \mathrm{aff}(\sigma') = m$. In particular, if $j \leq m$, then σ is a face of a Delaunay m-simplex. If σ is not a boundary simplex, and $\dim \operatorname{aff}(\sigma) < m$, then there are at least two Delaunay $(j+1)$-simplices that have σ as a face.

Proof Suppose $\dim \operatorname{aff}(\sigma) < m$. Let $B = B_{\mathbb{R}^m}(c; r)$ be a Delaunay ball for σ. Let ℓ be the line through c and $C(\sigma)$. If $c = C(\sigma)$, let ℓ be any line through c and orthogonal to $\operatorname{aff}(\sigma)$. There must be a point $\hat{c} \in \ell$ such that the circumscribing ball for σ centred at $\hat{c}$ is not empty. If this were not the case, we would have $\operatorname{aff}(\sigma) = \operatorname{aff}(\mathsf{P})$, and thus $\dim \operatorname{aff}(\mathsf{P}) < m$. It follows then (from the continuity of the radius of the circumballs parameterized by ℓ), that there is a point $c' \in [c, \hat{c}]$ that is the centre of a Delaunay ball for a simplex σ' that has σ as a proper face. The first assertion follows.

The second assertion follows from the same argument, and the observation that if σ is not on the boundary of $\operatorname{conv}(\mathsf{P})$, then there must be non-empty balls centred on ℓ at either side of c. If $p \in \mathsf{P} \setminus \operatorname{aff}(\sigma)$ is on the boundary of an empty ball centred at one side of c, by the intersection properties of spheres, it cannot be on the boundary of an empty ball centred on the other side of c. Thus there must be at least two distinct Delaunay $(k+1)$-simplices that share σ as a face. $\square$

Lemma 3.2 gives rise to the following observation, which plays an important role in Section 3.3, where we argue that protecting the Delaunay m-simplices yields a thickness bound on the simplices.

Lemma 3.3 (Separation) If $\tau \in \operatorname{Del}(\mathsf{P})$ is a j-simplex that is not a boundary simplex, and $q \in \mathsf{P} \setminus \tau$, then there is a Delaunay m-simplex σ^m which has τ as a face, but does not include q.

Proof Assume $j < m$, for otherwise there is nothing to prove. If $\sigma = q * \tau$ is not Delaunay, the assertion follows from the first part of Lemma 3.2. Assume σ is Delaunay and let $\tilde{\sigma}^m$ be a Delaunay m-simplex that has σ as a face. Thus $\tilde{\sigma}^m = q * \sigma^{m-1}$ for some Delaunay $(m-1)$-simplex, σ^{m-1}. Since $\tau \leq \sigma^{m-1}$ does not belong to the boundary of $\operatorname{conv}(\mathsf{P})$, neither does σ^{m-1}, so by the second part of Lemma 3.2, there is another Delaunay m-simplex σ^m that has σ^{m-1} (and therefore τ) as a face. Since σ^m is distinct from $\tilde{\sigma}^m$, it does not have q as a vertex. $\square$

3.1.1 The Delaunay complex in other metrics

We will also consider the Delaunay complex defined with respect to a metric d on $\mathbb{R}^m$ which differs from the Euclidean one. Specifically, if $\mathsf{P} \subset U \subset \mathbb{R}^m$ and $d : U \times U \to \mathbb{R}$ is a metric, then we define the Delaunay complex $\operatorname{Del}_d(\mathsf{P})$ with respect to the metric d.

The definitions are exactly analogous to the Euclidean case: A Delaunay ball is a maximal empty ball $B(x; r)$ in the metric d. The resulting Delaunay complex $\operatorname{Del}_d(\mathsf{P})$ consists of all the simplices which are circumscribed by a Delaunay ball with respect to the metric d. The simplices of $\operatorname{Del}_d(\mathsf{P})$ are, possibly degenerate, geometric simplices in $\mathbb{R}^m$. As for $\operatorname{Del}(\mathsf{P})$, $\operatorname{Del}_d(\mathsf{P})$ has the combinatorial structure of an abstract simplicial complex, but unlike $\operatorname{Del}(\mathsf{P})$, $\operatorname{Del}_d(\mathsf{P})$ may fail to be embedded even when there are no degenerate simplices.

3.2 Protection

A Delaunay simplex σ is *δ-protected* if it has a Delaunay ball B such that $d_{\mathbb{R}^m}(q, \partial B) > \delta$ for all $q \in \mathsf{P} \setminus \sigma$. We say that B is a δ-protected Delaunay ball for σ. If $\tau < \sigma$, then B is also a Delaunay ball for τ, but it cannot be a δ-protected Delaunay ball for τ. We say that σ is *protected* to mean that it is δ-protected for some unspecified $\delta > 0$.

Definition 3.4 (δ-generic) A finite set of points $\mathsf{P} \subset \mathbb{R}^m$ is *δ-generic* if all the Delaunay m-simplices are δ-protected. The set P is simply *generic* if it is δ-generic for some unspecified $\delta > 0$.

In his seminal work, Delaunay [Del34] demonstrated that if there is no empty ball with $m + 2$ points from P on its boundary, then $\mathrm{Del}(\mathsf{P})$ is realized as a simplicial complex in $\mathbb{R}^m$. In other words it is a triangulation, the *Delaunay triangulation*. If P is generic according to Definition 3.4, then Delaunay's criterion will be met. This is obvious if there are no degenerate m-simplices, and Definition 3.4 ensures that a degenerate m-simplex cannot exist in $\mathrm{Del}(\mathsf{P})$: If σ^m is degenerate, then by Lemma 3.2, there is a simplex σ with $\mathrm{aff}(\sigma) = \mathbb{R}^m$, and $\sigma^m < \sigma$. An affinely independent set of $m + 1$ vertices from σ defines a non degenerate m-simplex $\tilde{\sigma}^m$, and since its unique circumball is also a Delaunay ball for σ, it cannot be protected, a contradiction.

In particular, if P is generic if and only if there are no Delaunay simplices with dimension higher than m. We can say more. There are no degenerate Delaunay simplices. This can be inferred directly from Delaunay's result [Del34], but is also easily established from Lemma 3.2. In Section 3.3 we will quantify this observation with a bound on the thickness of the Delaunay simplices.

The δ-generic assumption means that all the Delaunay m-simplices are δ-protected, but the lower dimensional Delaunay do not necessarily enjoy this level of protection. The fact that there are no degenerate Delaunay simplices implies that all the simplices of all dimensions are $\tilde{\delta}$-protected for some $\tilde{\delta} > 0$. A lower bound on this number is established in Appendix B, but for the current work we have no need to consider the protection on the lower dimensional simplices.

3.2.1 Deep interior simplices

Delaunay avoided boundary complications by assuming a periodic point set, but the point sets that we will work with in Section 5 come from local patches of a well-sampled compact manifold without boundary. Periodic boundary conditions are not appropriate for our setting, but this is not a problem because Delaunay's argument applies locally:

Lemma 3.5 (Local Delaunay triangulation) If $p \in \mathsf{P}$ is an interior point, and the Delaunay m-simplices incident to p are protected, then $\mathrm{Del}(\mathsf{P})$ is a triangulation at p.

Proof We first show that p is surrounded by incident simplices, i.e. that $p \in \mathrm{int}\, \iota(|\mathrm{star}(p)|)$. (Recall that $\iota(|\mathrm{star}(p)|)$ is the image of $|\mathrm{star}(p)|$ in $\mathbb{R}^m$ from the map induced by the inclusion $\iota : \mathsf{P} \to \mathbb{R}^m$.) By Lemma 3.2, the set of m-simplices incident to p is not empty. Also, there are no degenerate simplices incident to p, since Lemma 3.2 implies that such a simplex would be a face of a k-simplex τ with $\dim \mathrm{aff}(\tau) = m$ and $k > m$. Then τ would have a non-degenerate m-simplex $\tilde{\tau}$

as a face, and $\tilde{\tau}$ would not be protected, since it shares its circumball with τ, a contradiction.

Let r be the minimum of the altitudes of p in each of the simplices incident to p. Let $S = \partial B$, be the bounding sphere for $B = B_{\mathbb{R}^m}(p; r/2)$. Then S intersects every simplex incident to p, and every simplex in $\mathrm{star}(p)$ that it intersects is incident to p. We claim that $S \subset \iota(|\mathrm{star}(p)|)$.

Suppose to the contrary that $x \in S \setminus \iota(|\mathrm{star}(p)|)$. Since $\iota(|\mathrm{star}(p)|)$ is closed, there is a neighbourhood of x which does not intersect $\iota(|\mathrm{star}(p)|)$. Consider $C = \partial\iota(|\mathrm{star}(p)|) \cap S$. Since only a finite collection of simplices contribute to C, and C separates two neighbourhoods on S, C must have codimension 1 in S, i.e., $\dim C = m - 2$. This means that there must be an $m - 1$ simplex σ^{m-1} that contributes to C. But this would mean that σ^{m-1} does not have an m-simplex on one side, contradicting Lemma 3.2. Thus $p \in \mathrm{int}\,\iota(|\mathrm{star}(p)|)$.

We now show that $\mathrm{star}(p)$ is embedded, i.e., that if $\sigma, \tau \in \mathrm{star}(p)$ intersect, then they intersect in a common face. By Lemma 3.2, we may assume that σ and τ are m-simplices. Let B_1 and B_2 be the Delaunay balls for σ and τ, and let H be the $(m-1)$-flat defined by $\mathrm{aff}(\partial B_1 \cap \partial B_2)$. Since B_1 and B_2 are empty balls, if $\sigma \neq \tau$, H separates the interiors of σ and τ, and thus they must intersect in H, i.e., at the common face defined by the vertices in $\partial B_1 \cap \partial B_2$.

Finally we show that if $\mathrm{int}\,\tau \cap \iota(|\mathrm{star}(p)|) \neq \emptyset$, for some $\tau \in \mathrm{Del}(\mathsf{P})$, then $\tau \in \mathrm{star}(p)$. Suppose $x \in \mathrm{int}\,\tau \cap \sigma$ for $\sigma \in \mathrm{star}(p)$. We may assume that σ is an m-simplex. Then consider the Delaunay balls B_1 for σ and B_2 for τ, and H defined as above. Since σ is protected, x can lie in the interior of σ only if $\tau = \sigma$. Otherwise, $x \in H$, and therefore all vertices of τ lie in H, and τ is a face of σ. $\square$

We wish to consider the properties of Delaunay triangulations in regions which are comfortably in the interior of $\mathrm{conv}(\mathsf{P})$, and avoid the complications that arise as we approach the boundary of the point set. We introduce some language to facilitate this.

If none of the vertices of σ lie on $\partial\mathrm{conv}(P)$, then it is an *interior simplex*. We wish to identify a subcomplex of the interior simplices of $\mathrm{Del}(\mathsf{P})$ consisting of those simplices whose neighbour simplices are also all interior simplices with small circumradius. An interior simplex near the boundary of $\mathrm{conv}(\mathsf{P})$ does not necessarily have its circumradius constrained by the sampling radius. However, we have the following:

Lemma 3.6 If P is an ϵ-sample set, and $\sigma \in \mathrm{Del}(\mathsf{P})$ has a vertex p such that $d_{\mathbb{R}^m}(p, \partial\mathrm{conv}(\mathsf{P})) \geq 2\epsilon$, then $R(\sigma) < \epsilon$ and σ is an interior simplex.

Proof Let $B_{\mathbb{R}^m}(c; r)$ be a Delaunay ball for σ. We will show $r < \epsilon$. Suppose to the contrary. Let x be the point on $[c, p]$ such that $d_{\mathbb{R}^m}(p, x) = \epsilon$. Then p is the closest point in P to x, and so the sampling criteria imply that $d_{\mathbb{R}^m}(x, \partial\mathrm{conv}(\mathsf{P})) < \epsilon$. But then $d_{\mathbb{R}^m}(p, \partial\mathrm{conv}(\mathsf{P})) \leq d_{\mathbb{R}^m}(p, x) + d_{\mathbb{R}^m}(x, \partial\mathrm{conv}(\mathsf{P})) < 2\epsilon$, contradicting the hypothesis on p.

Thus $r < \epsilon$, and it follows that σ is an interior simplex because if $q \in \sigma$, then $d_{\mathbb{R}^m}(p, q) \leq 2r < d_{\mathbb{R}^m}(p, \partial\mathrm{conv}(\mathsf{P}))$. $\square$

This suggests the following:

Definition 3.7 (Safe interior points) Suppose $\mathsf{P} \subset \mathbb{R}^m$ is an ϵ-sample set. The subset $\mathsf{P}_I \subset \mathsf{P}$ consisting of all $p \in \mathsf{P}$ with $d_{\mathbb{R}^m}(p, \partial\mathrm{conv}(\mathsf{P})) \geq 4\epsilon$ is the set of *safe interior points*.

By Lemma 3.6, all the simplices that include a safe interior point, as well as all the neighbours of such simplices, will have a small circumradius. For technical reasons it is inconvenient to demand that *all* the Delaunay m-simplices be δ-protected. We focus instead on a subset defined with respect to a set of safe interior points:

Definition 3.8 (δ-generic for P_J) The set $\mathsf{P} \subset \mathbb{R}^m$ is δ*-generic for* P_J if $\mathsf{P}_J \subseteq \mathsf{P}_I$ and all the m-simplices in $\mathrm{star}(\mathrm{star}(\mathsf{P}_J; \mathrm{Del}(\mathsf{P})))$ are δ-protected. The *deep interior simplices* are the simplices in $\mathrm{star}(\mathsf{P}_J; \mathrm{Del}(\mathsf{P}))$.

Thus the deep interior simplices are determined by our choice of $\mathsf{P}_J \subseteq \mathsf{P}_I$, and our protection requirements ensure that all the m-simplices that share a face with a deep interior simplex are δ-protected and have a small circumradius.

3.3 Thickness from protection

Our goal here is to demonstrate that the deep interior simplices on a δ-generic point set are Υ_0-thick. If $\delta = \nu_0 \epsilon$, for some constant ν_0, then we obtain a constant Υ_0 which depends only on ν_0. The key observation is that together with Lemma 3.3, protection imposes constraints on all the Delaunay simplices; they cannot be too close to being degenerate. In the particular case that $j = 0$, Lemma 3.3 immediately implies that the vertices of the deep interior simplices are δ-sparse:

Lemma 3.9 (Sparsity from protection) If P is δ-generic for P_J, then $L(\sigma) > \delta$ for any deep interior simplex σ.

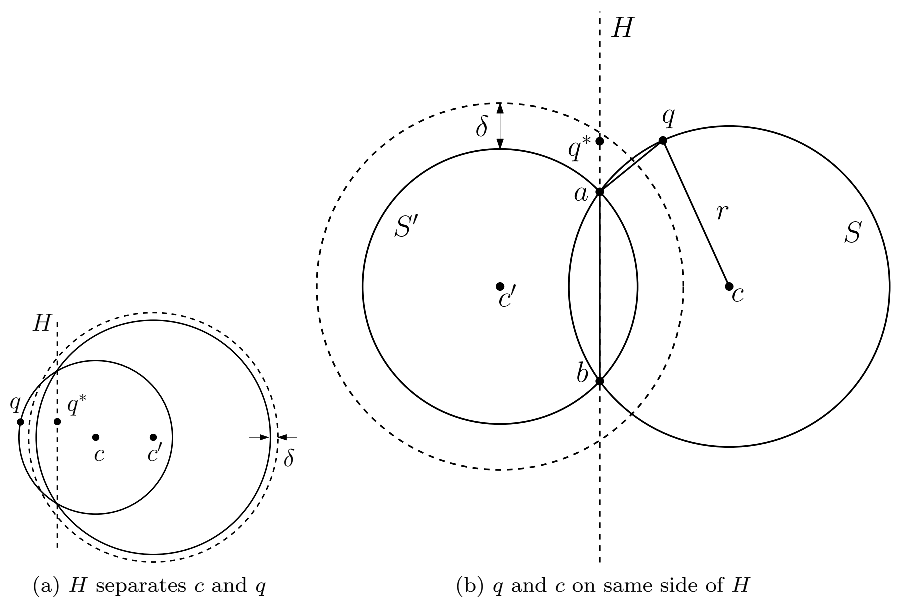

Figure 1: Diagram for Lemma 3.10. (a) When H separates q and c then $d_{\mathbb{R}^d}(q, q^*) > \delta$. (b) Otherwise, a lower bound on the distance between q and its projection q^* on H is obtained by an upper bound on the angle $\angle qab$.

Lemma 3.10 Suppose that $B = B_{\mathbb{R}^m}(c; r)$ is a Delaunay ball for $\sigma = q * \tau$ with $r < \epsilon$ and that $L(\tau) \geq \lambda$ for some $\lambda \leq \epsilon$. Suppose also that $\tau \leq \sigma'$ and that σ is not a face of σ'.

If B' is a δ-protected Delaunay ball for σ', and $H = \operatorname{aff}(\partial B \cap \partial B')$, then

$$d_{\mathbb{R}^m}(q, H) > \frac{\sqrt{3}\delta}{4\epsilon}(\lambda + \delta).$$

It follows that, if P is δ-generic for P_J, with sampling radius ϵ, and τ is a deep interior simplex, then

$$D(q, \sigma) > \frac{\sqrt{3}\delta^2}{2\epsilon}.$$

Proof Let $B' = B_{\mathbb{R}^m}(c'; r')$ be the δ-protected Delaunay ball for σ'. Our geometry will be performed in the plane, Q, defined by c, c', and q. This plane is orthogonal to the $(m-1)$-flat H, and it follows that the distance $d_{\mathbb{R}^m}(q, H)$ is realized by a segment in the plane Q: the projection, q^*, of q onto H lies in Q, and $d_{\mathbb{R}^m}(q, H) = d_{\mathbb{R}^m}(q, q^*)$.

If H separates q from c, then $\partial B'$ separates q from q^*, and $d_{\mathbb{R}^m}(q, q^*) > d_{\mathbb{R}^m}(q, \partial B') > \delta$, since B' is δ-protected (Figure 1(a)). The lemma then follows since λ and δ are each no larger than ϵ. Thus assume that q and c lie on the same side of H, as shown in Figure 1(b). Let $S' = Q \cap \partial B'$, and $S = Q \cap \partial B$, and let a and b be the points of intersection $S' \cap S$. Thus $H \cap Q$ is the line through a and b.

We will bound $d_{\mathbb{R}^m}(q, q^*)$ by finding an upper bound on the angle $\gamma = \angle qab$. This is the same as the standard calculation for upper-bounding the angles in a triangle with bounded circumradius to shortest edge ratio. Without loss of generality, we may assume that $\gamma \geq \angle qba$, and we will assume that $\gamma \geq \pi/2$ since otherwise $q* \in B'$ and thus $d_{\mathbb{R}^m}(q, q*) > \delta$ and the lemma is again trivially satisified.

Since $d_{\mathbb{R}^m}(a, q) > \delta$, we have $d_{\mathbb{R}^m}(q, q^*) = d_{\mathbb{R}^m}(a, q) \sin\gamma > \delta \sin\gamma$. Also note that $d(a, b) \geq 2R(\tau) \geq L(\tau) \geq \lambda$. Let $\alpha = \angle qac$. Then $\cos\alpha = \frac{d_{\mathbb{R}^m}(a,q)}{2r} \geq \frac{\delta}{2\epsilon}$. Similarly, with $\beta = \angle cab$, we have $\cos\beta \geq \frac{\lambda}{2\epsilon}$. Thus since $\gamma = \alpha + \beta \geq \pi/2$, we have

$$\begin{aligned}
d_{\mathbb{R}^m}(q, q^*) &> \delta \sin\left(\arccos\frac{\delta}{2\epsilon} + \arccos\frac{\lambda}{2\epsilon}\right) \\
&\geq \delta\left(\frac{\lambda}{2\epsilon}\sin\left(\arccos\frac{\delta}{2\epsilon}\right) + \frac{\delta}{2\epsilon}\sin\left(\arccos\frac{\lambda}{2\epsilon}\right)\right) \\
&\geq \frac{\sqrt{3}\delta}{4\epsilon}(\lambda + \delta),
\end{aligned}$$

where the last inequality follows from $\lambda \leq \epsilon$ and $\delta \leq \epsilon$.

Since $\operatorname{aff}(\tau) \subset H$, it follows that $D(q, \sigma) \geq d_{\mathbb{R}^m}(q, H)$, and if P is δ-generic for P_J, then $\lambda \geq \delta$, and Lemma 3.3 ensures that there is a δ-protected σ' that contains τ but not q. □

We thus obtain a bound on the thickness of the deep interior simplices when P is δ-generic. Since Lemma 3.10 yields a lower bound of $\frac{\sqrt{3}\delta^2}{2\epsilon}$ on the altitudes of the deep interior simplices, and since $\Delta(\sigma) \leq 2\epsilon$, we have that $\Upsilon(\sigma) \geq \frac{\sqrt{3}\delta^2}{4\epsilon^2}$ for all deep interior σ. If $\delta = \nu_0\epsilon$, we obtain a constant thickness bound.

Theorem 3.11 (Thickness from protection) If $\mathsf{P} \subset \mathbb{R}^m$ is δ-generic for P_J with $\delta = \nu_0\epsilon$, where ϵ is a sampling radius for P, then the deep interior simplices are Υ_0-thick, with

$$\Upsilon_0 = \frac{\sqrt{3}\nu_0^2}{4}.$$

4 Delaunay stability

We find upper bounds on the magnitude of a perturbation for which a protected Delaunay ball remains a Delaunay ball. We consider both perturbations of the sample points in Euclidean space, and perturbations of the metric itself. The primary technical challenge is bounding the effect of a perturbation on the circumcentre of an m-simplex. We then find the relationship between the perturbation parameter ρ and the protection parameter δ which ensures that a δ-protected Delaunay simplex will remain a Delaunay simplex.

4.1 Perturbations and circumcentres

As expected, a bound on the displacement of the circumcentre requires a bound on the thickness of the simplex.

4.1.1 Almost circumcentres

If σ is thick, a point whose distances to the vertices of σ are all almost the same, will lie close to $N(\sigma)$.

Lemma 4.1 If $\sigma = [p_0, ..., p_k] \subset \mathbb{R}^m$ is a non-degenerate k-simplex, and $x \in \mathbb{R}^m$ is such that

$$\left| \|p_i - x\|^2 - \|p_j - x\|^2 \right| \leq \xi^2 \quad \text{for all } i, j \in [0, ..., k], \tag{1}$$

then there is a $c \in N(\sigma)$ such that $\|c - x\| \leq \eta$, where

$$\eta = \frac{\xi^2}{2\Upsilon(\sigma)\Delta(\sigma)}.$$

In particular, if σ is an m-simplex then $x \in \overline{B}_{\mathbb{R}^m}(C(\sigma); \eta)$.

If the inequalities in Equations (1) are made strict, then the conclusions may also be stated with strict inequalities.

Proof First suppose $k = m$. The circumcentre of σ is given by the linear equations $\|C(\sigma) - p_i\|^2 = \|C(\sigma) - p_0\|^2$, or

$$(p_i - p_0)^{\mathsf{T}} C(\sigma) = \frac{1}{2}(\|p_i\|^2 - \|p_0\|^2).$$

Letting b be the vector whose i^{th} component is defined by the right hand side of the equation, and letting P be the $m \times m$ matrix, whose i^{th} column is $(p_i - p_0)$, we write the equations in matrix form:

$$P^{\mathsf{T}} C(\sigma) = b. \tag{2}$$

Without loss of generality, assume p_0 is the vertex that minimizes the distance to x. Then, defining x_a to be the vector whose i^{th} component is $\frac{1}{2}(\|p_i - x\|^2 - \|p_0 - x\|^2)$, we have $\|p_i - x\|^2 = \|p_0 - x\|^2 + 2(x_a)_i$, and we find

$$P^{\mathsf{T}} x = b - x_a. \tag{3}$$

From Equations (2) and (3) we have

$$\|C(\sigma) - x\| = \left\|P^{-T} x_a\right\| \leq \left\|P^{-1}\right\| \|x_a\|.$$

From Equation (1), it follows that $\|x_a\| \leq \frac{\sqrt{m}\xi^2}{2}$, and from Lemmas 2.2 and 2.4 we have $\left\|P^{-1}\right\| \leq (\sqrt{m}\Upsilon(\sigma)\Delta(\sigma))^{-1}$, and thus the result holds for full dimensional simplices.

If σ is a k-simplex with $k \leq m$, then we consider $\hat{x}$, the orthogonal projection of x into $\mathrm{aff}(\sigma)$. We observe that $\hat{x}$ also must satisfy Equation (1), and we conclude from the above argument that $\|C(\sigma) - \hat{x}\| \leq \eta$. Then letting $c = C(\sigma) + (x - \hat{x})$ we have that $c \in N(\sigma)$ and $\|c - x\| \leq \eta$. $\square$

It will be convenient to have a name for a point that is almost equidistant to the vertices of a simplex:

Definition 4.2 A ξ-*centre* for a simplex $\sigma = [p_0, \dots, p_k] \subset \mathbb{R}^m$ is a point x that satisfies

$$\big| \|p_i - x\| - \|p_j - x\| \big| \leq \xi \qquad \text{for all } i, j \leq k. \tag{4}$$

With a bound on the distance from x to the vertices of σ, Lemma 4.1 can be transformed into a bound on the distance from a ξ-centre to the closest true centre in $N(\sigma)$:

Lemma 4.3 If $\sigma = [p_0, \dots, p_k] \subset \mathbb{R}^m$ is non-degenerate, and x is a ξ-centre such that

$$\|p_i - x\| < \tilde{\epsilon} \qquad \text{for all } i, j \leq k,$$

then there exists a $c \in N(\sigma)$ such that $\|x - c\| < \eta$, where

$$\eta = \frac{\tilde{\epsilon}\xi}{\Upsilon(\sigma)\Delta(\sigma)}.$$

In particular, if σ is an m-simplex, then $x \in B_{\mathbb{R}^m}(C(\sigma); \eta)$.

Proof Let $R = \max_i \|p_i - x\|$ and $r = \min_i \|p_i - x\|$. Then

$$R^2 - r^2 = (R + r)(R - r) < 2\tilde{\epsilon}(R - r) \leq 2\tilde{\epsilon}\xi,$$

and the result then follows from Lemma 4.1. $\square$

4.1.2 Circumcentres and metric perturbations

We will show here that for an Υ_0-thick m-simplex σ in $\mathbb{R}^m$, and a metric d that is close to $d_{\mathbb{R}^m}$, there will be a metric circumcentre c near $C(\sigma)$. We require the metric d to be continuous in the topology defined by $d_{\mathbb{R}^m}$. Henceforth, whenever we refer to a *perturbation of the Euclidean metric*, this topological compatibility will always be assumed.

The proof is a topological argument based on considering a mapping into $\mathbb{R}^m$ of a small ball around the circumcentre of σ. The mapping is based on the metric and is such that circumcentres get mapped to the origin. In the mapping associated to the Euclidean metric, points that get mapped close to the origin are ξ-centres, and since the ξ-centres are in the interior of the ball, the boundary of the ball does not get mapped close to the origin. A small perturbation of the metric yields a small perturbation in the mapping, and so we can argue that there is a homotopy between the mapping associated with the Euclidean metric and the one associated to the perturbed metric, such that no point on the boundary of the ball ever gets mapped to the origin. A consideration of the degree of the mapping allows us to conclude that the ball must contain a circumcentre for the perturbed metric.

We will demonstrate the following:

Lemma 4.4 (Circumcentres: metric perturbation) Let $U \subset \mathbb{R}^m$, and let $d : U \times U \to \mathbb{R}$ be a continuous metric with respect to the topology defined by $d_{\mathbb{R}^m}$, and such that for any $x, y \in U$ with $d_{\mathbb{R}^m}(x,y) < 2\epsilon$, we have $|d(x,y) - d_{\mathbb{R}^m}(x,y)| \leq \rho$, with

$$\rho \leq \frac{\Upsilon_0 \mu_0 \epsilon}{8}.$$

If $\sigma = [p_0, \ldots, p_m] \subset U$ is an Υ_0-thick m-simplex with $R(\sigma) < \epsilon$, and $L(\sigma) \geq \mu_0 \epsilon$, and such that $d_{\mathbb{R}^m}(p_i, \partial U) \geq 2\epsilon$, then there is a point

$$c \in B = B_{\mathbb{R}^m}(C(\sigma); \eta) \quad \text{ with } \eta = \frac{8\rho}{\Upsilon_0 \mu_0},$$

and such that $d(c, p_i) = d(c, p_j)$ for all $p_i, p_j \in \sigma$.

In order to prove Lemma 4.4, we will use a particular case of Lemma 4.3:

Lemma 4.5 Suppose σ is an Υ_0-thick m-simplex such that $L(\sigma) \geq \mu_0 \epsilon$. If x is a ξ-centre for σ with $d_{\mathbb{R}^m}(x,p) < 2\epsilon$ for all $p \in \sigma$, then $x \in B_{\mathbb{R}^m}(C(\sigma); \eta)$, where $\eta = \frac{2\xi}{\Upsilon_0 \mu_0}$.

Let $B = B_{\mathbb{R}^m}(C(\sigma); \eta)$ be the open ball which contains the ξ-centres for σ. We will show that if $\xi = 4\rho$, then a circumcentre c for σ with respect to d will also lie in B. However, we make no claim that c is unique. Note that $\overline{B} \subset U$.

Consider the function $f_e : \overline{B} \to \mathbb{R}^m$ given by

$$f_e(x) = (d_{\mathbb{R}^m}(x, p_1) - d_{\mathbb{R}^m}(x, p_0), \ldots, d_{\mathbb{R}^m}(x, p_m) - d_{\mathbb{R}^m}(x, p_0))^{\mathsf{T}}. \tag{5}$$

Observe that f_e maps the circumcentre of σ, and only the circumcentre, to the origin: $f_e^{-1}(0) = \{C(\sigma)\}$.

We construct a similar function for the metric d,

$$f(x) = (d(x, p_1) - d(x, p_0), \ldots, d(x, p_m) - d(x, p_0))^{\mathsf{T}}, \tag{6}$$

and we will show that there must be a $c \in f^{-1}(0) \subset B$. We first show that there is a homotopy between f and f_e such that the image of $\partial\overline{B}$ never touches the origin:

Lemma 4.6 Under the hypotheses of Lemma 4.4, if $\xi = 4\rho \leq \frac{\Upsilon_0 \mu_0 \epsilon}{2}$, then there is a homotopy $F : \overline{B} \times [0,1] \to \mathbb{R}^m$ with $F(x,t) \neq 0$ for all $x \in \partial\overline{B}$ and $t \in [0,1]$.

Proof We define the homotopy $F : \overline{B} \times [0,1] \to \mathbb{R}^m$ by $F(x,t) = (1-t)f_e(x) + tf(x)$. By the bounds on ξ and $R(\sigma)$, for every $x \in \overline{B}$, and $p \in \sigma$, we have $d_{\mathbb{R}^m}(x,p) \leq \frac{2\xi}{\Upsilon_0 \mu_0} + R(\sigma) < 2\epsilon$. Thus it follows from Lemma 4.5 that $x \in \partial\overline{B}$ cannot be a ξ-centre.

It is convenient to consider the max norm on $\mathbb{R}^m$ defined by the largest magnitude of the components: $\|f_e(x)\|_\infty = \max_i |f_e(x)_i|$. (This gives us a better bound than working with the standard Euclidean norm.) If $\|f_e(y)\|_\infty \leq \frac{\xi}{2}$, then y must be a ξ-centre. Indeed, we would have $|\|p_i - y\| - \|p_j - y\|| \leq |\|p_i - y\| - \|p_0 - y\|| + |\|p_0 - y\| - \|p_j - y\|| \leq \frac{\xi}{2} + \frac{\xi}{2} = \xi$ for all i and j. Thus, since $x \in \partial\overline{B}$ is not a ξ-centre, we must have $\|f_e(x)\|_\infty > \frac{\xi}{2}$.

Also, from the hypothesis on d, we have $\|f_e(x) - f(x)\|_\infty \leq 2\rho = \frac{\xi}{2}$, for all $x \in \partial\overline{B}$. It follows that $\|F(x,t)\|_\infty \geq \|f_e(x)\|_\infty - t\,\|f(x) - f_e(x)\|_\infty > 0$ when $x \in \partial\overline{B}$. $\square$

We will need the following observation:

Lemma 4.7 The origin is a regular value for the function f_e defined in Equation (5).

Proof Choose a coordinate system such that $C(\sigma) \in B$ is the origin. We show by a direct calculation that $\det J(f_e)_0 \neq 0$, where $J(f_e)_0$ is the Jacobian matrix representing the derivative of f_e in our coordinate system.

Let $p_i = (p_{i0}, \ldots, p_{im})^\mathsf{T}$ for all $p_i \in \{p_0, \ldots, p_m\}$. For $x = (x_1, \ldots, x_m)^\mathsf{T} \in \mathbb{R}^m$, let $f_e(x) = (f_0(x), \ldots, f_m(x))^\mathsf{T}$, where

$$f_i(x) = \|p_i - x\| - \|p_0 - x\| = \sqrt{\sum_{k=1}^{m}(p_{ik} - x_k)^2} - \sqrt{\sum_{k=1}^{m}(p_{0k} - x_k)^2}\,.$$

We find

$$\left.\frac{\partial f_i}{\partial x_j}\right|_0 = \frac{p_{0j} - p_{ij}}{R(\sigma)},$$

and thus

$$J(f)_0 = -\frac{1}{R(\sigma)} P^\mathsf{T}, \tag{7}$$

where as usual P is the matrix whose columns are $p_i - p_0$. Since $\mathrm{vol}(\sigma^m) = \frac{|\det(P)|}{m!}$, Equation (7) implies

$$|\det J(f)_0| = \frac{m!\ \mathrm{vol}(\sigma^m)}{R(\sigma)^m}.$$

Thus since $f_e^{-1}(0) = \{0\}$, 0 is a regular value for f_e provided σ is non-degenerate. $\square$

Lemma 4.4 now follows from a consideration of the degree of the mappings f and f_e relative to zero. The *degree* of a smooth map $g : \overline{B} \to \mathbb{R}^m$ at a regular point $p \in g(B)$ is defined by

$$\deg(g, p, B) = \sum_{x \in g^{-1}(p)} \operatorname{sign} \det J(g)_x\,,$$

where $J(g)_x$ is the Jacobian matrix of g at x. The exposition by Dancer [Dan00] is a good reference for the degree of maps from manifolds with boundary. It is shown that the definition of $\deg(g, p, B)$ extends to continuous maps g that are not necessarily differentiable. If $h : \overline{B} \to \mathbb{R}^m$ is homotopic to g by a homotopy $H : \overline{B} \times [0,1] \to \mathbb{R}^m$ such that $H(x,t) \neq p$ for all $t \in [0,1]$, and $x \in \partial B$, then $\deg(g,p,B) = \deg(h,p,B)$.

Since $f_e^{-1}(0) = \{C(\sigma)\}$, it follows from Lemma 4.7 that $\deg(f_e, 0, B) = \pm 1$. Then Lemma 4.6 implies $\deg(f, 0, B) = \deg(f_e, 0, B)$, and since this is nonzero, it must be that $f^{-1}(0) \neq \emptyset$. The demonstration of Lemma 4.4 is complete.

4.1.3 Circumcentres and point perturbations

The exact same argument as was used to demonstrate Lemma 4.4 can be used to show that an m-simplex $\tilde{\sigma} = [\tilde{p}_0, \ldots, \tilde{p}_m]$ whose vertices lie close to a thick m-simplex σ, will also have a circumcentre that lies close to $C(\sigma)$. We replace the function f defined in Equation (6) by the function

$$\tilde{f}(x) = (d_{\mathbb{R}^m}(x, \tilde{p}_1) - d_{\mathbb{R}^m}(x, \tilde{p}_0), \ldots, d_{\mathbb{R}^m}(x, \tilde{p}_m) - d_{\mathbb{R}^m}(x, \tilde{p}_0))^\mathsf{T},$$

and the same argument goes through. We obtain:

Lemma 4.8 (Circumcentres: point perturbation) Suppose that $\sigma = [p_0, \dots, p_m]$ is an Υ_0-thick m-simplex with $R(\sigma) < \epsilon$ and $L(\sigma) \geq \mu_0 \epsilon$. Suppose also that $\tilde{\sigma} = [\tilde{p}_0, \dots, \tilde{p}_m]$ is such that $\|\tilde{p}_i - p_i\| \leq \rho$ for all $i \in [0, \dots, m]$. If

$$\rho \leq \frac{\Upsilon_0 \mu_0 \epsilon}{8}, \qquad \text{then} \qquad d_{\mathbb{R}^m}(C(\tilde{\sigma}), C(\sigma)) < \frac{8\rho}{\Upsilon_0 \mu_0}.$$

4.2 Perturbations and protection

Suppose $\zeta : \mathsf{P} \to \tilde{\mathsf{P}}$ is a ρ-perturbation. If σ is a δ-protected m-simplex in $\mathrm{Del}(\mathsf{P})$, then we want an upper bound on ρ that will ensure that $\tilde{\sigma} = \zeta(\sigma)$ is protected in $\mathrm{Del}(\tilde{\mathsf{P}})$. The following definition will be convenient:

Definition 4.9 (Secure simplex) A simplex $\sigma \in \mathrm{Del}(\mathsf{P})$ is *secure* if it is a δ-protected m-simplex that is Υ_0-thick and satisfies $R(\sigma) < \epsilon$ and $L(\sigma) \geq \mu_0 \epsilon$.

Our stability results apply to subcomplexes of secure simplices, the definition of which employs multiple parameters. For deep interior simplices Lemma 3.9 and Theorem 3.11 allow us to consolidate some of these parameters with the ratio δ/ϵ:

Lemma 4.10 (Deep interior simplices are secure) If P satisfies a sampling radius ϵ and is δ-generic for P_J, with $\delta = \nu_0 \epsilon$, then the deep interior m-simplices are secure with $\mu_0 = \nu_0$ and $\Upsilon_0 = \frac{\sqrt{3}\nu_0^2}{4}$.

Lemma 4.11 (Protection and point perturbation) Suppose that $\mathsf{P} \subset \mathbb{R}^m$ and $\sigma \in \mathrm{Del}(\mathsf{P})$ is secure. If $\zeta : \mathsf{P} \to \tilde{\mathsf{P}}$ is a ρ-perturbation with

$$\rho \leq \frac{\Upsilon_0 \mu_0}{18} \delta,$$

then $\zeta(\sigma) = \tilde{\sigma} \in \mathrm{Del}(\tilde{\mathsf{P}})$ and has a $(\delta - \frac{18}{\Upsilon_0 \mu_0}\rho)$-protected Delaunay ball.

Proof Let $B = B_{\mathbb{R}^m}(c; r)$ be the δ-protected Delaunay ball for $\sigma \in \mathrm{Del}(\mathsf{P})$, and let $\tilde{B} = B_{\mathbb{R}^m}(\tilde{c}; \tilde{r})$ be the circumball for the corresponding perturbed simplex $\tilde{\sigma}$. We wish to establish a bound on ρ that will ensure that $\tilde{B}$ is protected with respect to $\tilde{\mathsf{P}}$.

Let $q \in \mathsf{P}$ be a point not in σ. We need to ensure that the corresponding $\tilde{q}$ lies outside the closure of $\tilde{B}$, i.e., that $d_{\mathbb{R}^m}(\tilde{q}, \tilde{c}) > \tilde{r}$.

Since $\delta \leq \epsilon$, the hypothesis of Lemma 4.8 is satisfied by ρ, and we have $d_{\mathbb{R}^m}(\tilde{c}, c) < \eta\rho$, where $\eta = \frac{8}{\Upsilon_0 \mu_0}$. Thus for $p \in \sigma$ and corresponding $\tilde{p} \in \tilde{\sigma}$ we have

$$\begin{aligned} \tilde{r} &\leq d_{\mathbb{R}^m}(c, p) + d_{\mathbb{R}^m}(c, \tilde{c}) + d_{\mathbb{R}^m}(p, \tilde{p}) \\ &< r + (\eta + 1)\rho. \end{aligned}$$

Also

$$\begin{aligned} d_{\mathbb{R}^m}(\tilde{q}, \tilde{c}) &\geq d_{\mathbb{R}^m}(q, c) - d_{\mathbb{R}^m}(\tilde{c}, c) - d_{\mathbb{R}^m}(\tilde{q}, q) \\ &> r + \delta - \rho(\eta + 1). \end{aligned}$$

Therefore $\tilde{q}$ will be outside of the closure of $\tilde{B}$ provided $r + \delta - \rho(\eta+1) \geq r + (1+\eta)\rho$, i.e., when $\delta \geq 2(\eta + 1)\rho$. The result follows from the definition of η and the observation that μ_0 and Υ_0 are each no larger than one. $\square$

A similar argument yields a bound on the metric perturbation that will ensure the Delaunay balls for the m-simplices remain protected:

Lemma 4.12 (Protection and metric perturbation) Suppose $U \subset \mathbb{R}^m$ contains conv(P) and $d : U \times U \to \mathbb{R}$ is a metric such that $|d_{\mathbb{R}^m}(x, y) - d(x, y)| \leq \rho$ for all $x, y \in U$. Suppose also that $\sigma \in \mathrm{Del}(\mathsf{P})$ is secure. If

$$\rho \leq \frac{\Upsilon_0 \mu_0}{20} \delta,$$

and $d_{\mathbb{R}^m}(p, \partial U) \geq 2\epsilon$ for every vertex $p \in \sigma$, then σ also belongs to $\mathrm{Del}_d(\mathsf{P})$, and has a $(\delta - \frac{20}{\Upsilon_0 \mu_0}\rho)$-protected Delaunay ball in the metric d.

Proof Let $B = B_{\mathbb{R}^m}(c; r)$ be the Euclidean δ-protected Delaunay ball for $\sigma \in \mathrm{Del}(\mathsf{P})$, and let $\tilde{B} = B_{\mathbb{R}^m}(\tilde{c}; \tilde{r})$ be a circumball for σ in the metric d. We wish to establish a bound on ρ that will ensure that $\tilde{B}$ is protected with respect to d.

Let $q \in \mathsf{P}$ be a point not in σ. We need to ensure that $d(q, \tilde{c}) > \tilde{r}$. Since $\delta \leq \epsilon$, the hypothesis ensures that $\rho \leq \frac{\Upsilon_0 \lambda}{8}$, and so Lemma 4.4 yields $d_{\mathbb{R}^m}(\tilde{c}, c) < \eta\rho$, where $\eta = \frac{8}{\Upsilon_0 \mu_0}$. Thus for $p \in \sigma$

$$\begin{aligned} \tilde{r} &\leq d(c, p) + d(c, \tilde{c}) \\ &< (r + \rho) + (\eta\rho + \rho) \\ &= r + (\eta + 2)\rho, \end{aligned}$$

and

$$\begin{aligned} d(q, \tilde{c}) &\geq d(q, c) - d(\tilde{c}, c) \\ &> r + \delta - (\eta + 2)\rho. \end{aligned}$$

Thus $\tilde{q}$ will be outside of the closure of $\tilde{B}$ provided $r + \delta - (\eta + 2)\rho \geq r + (\eta + 2)\rho$, i.e., when

$$\delta \geq 2(\eta + 2)\rho.$$

The result follows from the definition of η and the observation that μ_0 and Υ_0 are each no larger than one. ☐

4.3 Perturbations and Delaunay stability

The results of Section 4.2 translate into stability results for Delaunay triangulations. In the case of point perturbations in Euclidean space, the connectivity of the Delaunay triangulation cannot change as long as the simplices corresponding to the initial m-simplices remain protected. This is a direct consequence of Delaunay's original result [Del34], but we explicitly lay out the argument.

In the case of metric perturbation, we can no longer take for granted that the Delaunay complex cannot change its connectivity if the m-simplices remain protected. This is because we are no longer guaranteed that the Delaunay complex will be a triangulation. Using the consequences of the point-perturbation result, we establish bounds that ensure that the Delaunay complex in the perturbed metric will be the same as the original Delaunay triangulation.

4.3.1 Point perturbations

Lemma 4.11 establishes bounds on a ρ-perturbation $\zeta : \mathsf{P} \to \tilde{\mathsf{P}}$ which will guarantee that $\zeta(\mathrm{star}(\mathsf{P}_J)) \subseteq \mathrm{Del}(\tilde{P})$. Since Lemma 4.11 also guarantees that, if ρ is small enough, the vertices of $\mathrm{star}(\zeta(\mathsf{P}_J); \mathrm{Del}(\tilde{\mathsf{P}}))$ will be generic, we know from Delaunay's Lemma 3.5 that $\mathrm{star}(\zeta(\mathsf{P}_J); \mathrm{Del}(\tilde{\mathsf{P}}))$ will be a triangulation at each $\tilde{p} \in \zeta(\mathsf{P}_J)$, and thus $\mathrm{star}(\zeta(\mathsf{P}_J)) = \zeta(\mathrm{star}(\mathsf{P}_J))$ follows from Lemma 2.7. Explicitly, we have:

Lemma 4.13 Suppose $\mathsf{P} \subset \mathbb{R}^m$ is a generic sample set, and $\zeta : \mathsf{P} \to \tilde{\mathsf{P}}$ is a perturbation such that $\zeta(\mathrm{star}(\mathsf{P}_J; \mathrm{Del}(\mathsf{P}))) \subseteq \mathrm{star}(\zeta(\mathsf{P}_J); \mathrm{Del}(\tilde{\mathsf{P}}))$, and every $\tilde{\sigma} \in \zeta(\mathrm{star}(\mathsf{P}_J))$ is protected in $\mathrm{Del}(\tilde{\mathsf{P}})$. Then $\zeta(\mathrm{star}(\mathsf{P}_J)) = \mathrm{star}(\zeta(\mathsf{P}_J))$.

Thus, considering Lemma 4.11, we have the following stability theorem for Delaunay triangulations of δ-protected points:

Theorem 4.14 (Stability under point perturbation) Suppose $\mathsf{P} \subset \mathbb{R}^m$ and $\mathsf{P}_J \subseteq \mathsf{P}$ is such that every m-simplex in $\mathrm{star}(\mathsf{P}_J)$ is secure. If $\zeta : \mathsf{P} \to \tilde{\mathsf{P}}$ is a ρ-perturbation, with

$$\rho \leq \frac{\Upsilon_0 \mu_0}{18} \delta$$

then

$$\mathrm{star}(\mathsf{P}_J; \mathrm{Del}(\mathsf{P})) \overset{\zeta}{\cong} \mathrm{star}(\zeta(\mathsf{P}_J); \mathrm{Del}(\tilde{\mathsf{P}})).$$

The ρ-relaxed Delaunay complex for P was defined by de Silva [dS08] by the criterion that $\sigma \in \mathrm{Del}^{\rho}(\mathsf{P})$ if and only if there is a ball $B = B_{\mathbb{R}^m}(c; r)$ such that $\sigma \subset \overline{B}$, and $d_{\mathbb{R}^m}(c, q) \geq r - \rho$ for all $q \in \mathsf{P}$. Thus the simplices in $\mathrm{Del}^{\rho}(\mathsf{P})$ all have "almost empty", balls centred on a ρ-centre for σ. We have the following consequence of Theorem 4.14:

Corollary 4.15 (Stability under relaxation) Suppose $\mathsf{P} \subset \mathbb{R}^m$ and $\mathsf{P}_J \subseteq \mathsf{P}$ is such that every m-simplex in $\mathrm{star}(\mathsf{P}_J)$ is secure. If

$$\rho \leq \frac{\Upsilon_0 \mu_0}{18} \delta,$$

then

$$\mathrm{star}(\mathsf{P}_J; \mathrm{Del}^{\rho}(\mathsf{P})) = \mathrm{star}(\mathsf{P}_J; \mathrm{Del}(\mathsf{P})).$$

Proof Suppose that $\sigma \in \mathrm{star}(\mathsf{P}_J; \mathrm{Del}^{\rho}(\mathsf{P}))$. Then there is a ball B enclosing σ such that any point $q \in B$ is within a distance ρ from ∂B. Project all such points radially out to ∂B. Then we have a ρ-perturbation $\zeta : \mathsf{P} \to \tilde{\mathsf{P}}$, and σ has become $\tilde{\sigma} \in \mathrm{star}(\zeta(\mathsf{P}_J); \mathrm{Del}(\tilde{\mathsf{P}}))$. By Theorem 4.14, $\mathrm{star}(\zeta(\mathsf{P}_J); \mathrm{Del}(\tilde{\mathsf{P}})) \overset{\zeta}{\cong} \mathrm{star}(\mathsf{P}_J; \mathrm{Del}(\mathsf{P}))$, and therefore $\sigma \in \mathrm{star}(\mathsf{P}_J; \mathrm{Del}(\mathsf{P}))$. □

4.3.2 Metric perturbation

For a perturbation of the metric, we can exploit the stability results obtained for perturbations of points in the Euclidean metric to ensure that no simplices can appear in $\mathrm{star}(\mathsf{P}_J; \mathrm{Del}_d(\mathsf{P}))$ that do not already exist in $\mathrm{star}(\mathsf{P}_J; \mathrm{Del}(\mathsf{P}))$.

Lemma 4.16 Suppose $\mathrm{conv}(\mathsf{P}) \subseteq U \subset \mathbb{R}^m$ and $d : U \times U \to \mathbb{R}$ is such that $|d(x, y) - d_{\mathbb{R}^m}(x, y)| \leq \rho$ for all $x, y \in U$. Suppose also that $\mathsf{P}_J \subseteq \mathsf{P}$ is such that every m-simplex $\sigma \in \mathrm{star}(\mathsf{P}_J)$ is secure and satisfies $d_{\mathbb{R}^m}(p, \partial U) \geq 2\epsilon$ for every vertex $p \in \sigma$. If

$$\rho \leq \frac{\Upsilon_0 \mu_0}{36} \delta,$$

then

$$\mathrm{star}(\mathsf{P}_J; \mathrm{Del}_d(\mathsf{P})) \subseteq \mathrm{star}(\mathsf{P}_J; \mathrm{Del}(\mathsf{P})).$$

Proof Let $B(c; r)$ be a Delaunay ball for simplex $\sigma \in \mathrm{star}(\mathsf{P}_I; \mathrm{Del}_d(\mathsf{P}))$. Then $d(c, p) \leq d(c, q)$ for all $p \in \sigma$, and $q \in \mathsf{P}$. By the hypothesis on d, this implies that $d_{\mathbb{R}^m}(c, p) \leq d_{\mathbb{R}^m}(c, q) + 2\rho$ for all $p \in \sigma$ and $q \in \mathsf{P}$, and therefore $\sigma \in \mathrm{Del}^{2\rho}(\mathsf{P})$. The result now follows from Corollary 4.15. □

The perturbation bounds required by Lemma 4.16, also satisfy the requirements of Lemma 4.12. This gives us the reverse inclusion, and thus we can quantify the stability under metric perturbation for subcomplexes of secure simplices in Delaunay triangulations:

Theorem 4.17 (Stability under metric perturbation) Suppose $\mathrm{conv}(\mathsf{P}) \subseteq U \subset \mathbb{R}^m$ and $d : U \times U \to \mathbb{R}$ is such that $|d(x,y) - d_{\mathbb{R}^m}(x,y)| \leq \rho$ for all $x, y \in U$. Suppose also that $\mathsf{P}_J \subseteq \mathsf{P}$ is such that every m-simplex $\sigma \in \mathrm{star}(\mathsf{P}_J)$ is secure and satisfies $d_{\mathbb{R}^m}(p, \partial U) \geq 2\epsilon$ for every vertex $p \in \sigma$. If

$$\rho \leq \frac{\Upsilon_0 \mu_0}{36}\delta,$$

then

$$\mathrm{star}(\mathsf{P}_J; \mathrm{Del}_d(\mathsf{P})) = \mathrm{star}(\mathsf{P}_J; \mathrm{Del}(\mathsf{P})).$$

Using Lemma 4.10 and recognizing that the deep interior simplices also satisfy the distance from boundary requirement of Theorem 4.17, we can restate this metric perturbation stability result for Delaunay triangulations on δ-generic point sets:

Corollary 4.18 (Stability under metric perturbation) Suppose P is δ-generic for P_J, with sampling radius ϵ and $\delta = \nu_0 \epsilon$. Suppose also that $\mathrm{conv}(\mathsf{P}) \subseteq U$, and $d : U \times U \to \mathbb{R}$ is such that $|d(x,y) - d_{\mathbb{R}^m}(x,y)| \leq \rho$ for all $x, y \in U$. If

$$\rho \leq \frac{\nu_0^3}{84}\delta = \frac{\nu_0^4}{84}\epsilon,$$

then

$$\mathrm{star}(\mathsf{P}_J; \mathrm{Del}_d(\mathsf{P})) = \mathrm{star}(\mathsf{P}_J; \mathrm{Del}(\mathsf{P})).$$

5 Equating Delaunay structures

We apply the results of the previous sections to the task of triangulating $\mathcal{M}$, a smooth, compact m-manifold, without boundaries embedded in $\mathbb{R}^N$. We wish to build a Delaunay complex on a finite set $\mathcal{P} \subset \mathcal{M}$.

The *restricted Delaunay complex* is the Delaunay complex $\mathrm{Del}_{\mathbb{R}^N|_\mathcal{M}}(\mathcal{P})$ obtained when distances on the manifold are measured with the metric $d_{\mathbb{R}^N|_\mathcal{M}}$. This is the Euclidean metric of the ambient space, restricted to the manifold. In other words, $d_{\mathbb{R}^N|_\mathcal{M}}(x,y) = d_{\mathbb{R}^N}(x,y)$. We use this notation to avoid ambiguities in conjunction with the local Euclidean metrics discussed below. The Delaunay complex $\mathrm{Del}_{\mathbb{R}^N|_\mathcal{M}}(\mathcal{P})$ is a substructure of $\mathrm{Del}_{\mathbb{R}^N}(\mathcal{P})$.

Alternatively, distances on the manifold may be measured with $d_\mathcal{M}$, the *intrinsic metric* of the manifold. This metric defines the distance between x and y as the infimum of the lengths of the paths on $\mathcal{M}$ which connect x and y. This metric is also induced from $d_{\mathbb{R}^N}$. The *intrinsic Delaunay complex* is the Delaunay structure $\mathrm{Del}_\mathcal{M}(\mathcal{P})$ associated with this metric.

Although neither of these metrics are Euclidean, the idea is that locally, in a small neighbourhood of any point, these metrics may be well approximated by $d_{\mathbb{R}^m}$. Then, if the sampling satisfies appropriate δ-generic and ϵ-dense criteria in these local Euclidean metrics, the global Delaunay complex in the metric of the manifold will coincide locally with a Euclidean Delaunay triangulation, and we can thus guarantee a manifold complex.

5.1 Local Euclidean metrics

A *local parameterization* at a point $p \in \mathcal{M}$, is a pair (U, ψ_p), where $U \subset \mathbb{R}^m$ is an open neighbourhood of the origin, and $\psi_p : U \to \psi_p(U) = W \subset \mathcal{M}$ is a homeomorphism onto its image, and maps the origin to p. We will use ψ_p to pull back the metric of the manifold to U, and to simplify the notation we will write $d_{\mathcal{M}}(x, y)$ for $x, y \in U$, where it is to be understood that this means $d_{\mathcal{M}}(\psi_p(x), \psi_p(y))$, and likewise for $d_{\mathbb{R}^N|_{\mathcal{M}}}(x, y)$. Indeed, once W and U have been coupled together by a homeomorphism, we can transfer the metrics between them and the distinction becomes only one of perspective. We refer to the standard metric $d_{\mathbb{R}^m}$ on U as a a *local Euclidean metric* for p on W. Clearly this metric depends upon the choice of ψ_p; there are different ways to impose a Euclidean metric on W. In this work the only local parameterization we consider is defined by the projection map π_p described in Section 5.3.

We wish to generate a sample set $\mathcal{P} \subset \mathcal{M}$ that will allow us to exploit the results of Section 4. We will require that for each $p \in \mathcal{P}$ there be a local parameterization on a neighbourhood W of p that is large enough to ensure that p is a safe interior point in $\mathsf{P} = \mathcal{P} \cap W$ in the local Euclidean metric. This in turn will impose a constraint on the sampling radius ϵ so that W may be small enough that the metric distortion induced by ψ_p does not exceed the bound required for Corollary 4.18.

5.2 Background results for manifolds

The tangent space at $p \in \mathcal{M}$ is denoted $T_p\mathcal{M}$, and we identify it with an m-flat in the ambient space. The normal space, $N_p\mathcal{M}$, is the orthogonal complement of $T_p\mathcal{M}$ in $T_p\mathbb{R}^N$, and we likewise treat it as the affine subspace of dimension $m - k$ orthogonal to $T_p\mathcal{M} \subset \mathbb{R}^N$.

A ball $B = B_{\mathbb{R}^N}(c; r)$ is a *medial ball* at p if $B \cap \mathcal{M} = \emptyset$, it is tangent to $\mathcal{M}$ at p, and it is maximal in the sense that any ball which is centred on the line through p and c, and contains B, either coincides with B or intersects $\mathcal{M}$. The *local reach* at p is the infimum of the radii of the medial balls at p, and the *reach* of $\mathcal{M}$, denoted $\mathrm{rch}(\mathcal{M})$, is the infimum of the local reach over all points of $\mathcal{M}$. In order to approximate the geometry and topology with a simplical complex, manifolds with small reach require a higher sampling density than those with a larger reach. As is typical, an upper bound on our sampling radius will be proportional to $\mathrm{rch}(\mathcal{M})$. Since $\mathcal{M} \subset \mathbb{R}^N$ is a smooth, compact embedded submanifold, it has positive reach.

An estimate of how the tangent space locally deviates from the manifold is given by an observation of Federer [Fed59, Theorem 4.8(7)] (see also Giesen and Wagner [GW04, Lemma 6]):

Lemma 5.1 (Distance to tangent space) If $x, y \in \mathcal{M} \subset \mathbb{R}^N$ and $d_{\mathbb{R}^N}(x, y) \leq r < \mathrm{rch}(\mathcal{M})$, then $d_{\mathbb{R}^N}(y, T_x\mathcal{M}) \leq \frac{r^2}{2\mathrm{rch}(\mathcal{M})}$, and thus $\sin \alpha \leq \frac{r}{2\mathrm{rch}(\mathcal{M})}$, where α is the angle between $[x, y]$ and $T_x\mathcal{M}$.

A complementary result bounds the distance to the manifold from a point on a tangent space [GW04, Lemma 7]:

Lemma 5.2 (Distance to manifold) Suppose $v \in T_x\mathcal{M}$ with $\|v - x\| = r \leq \frac{\mathrm{rch}(\mathcal{M})}{4}$. If $\hat{v} \in \mathcal{M}$ is the closest point in $\mathcal{M}$ to v, then $d_{\mathbb{R}^N}(v, \hat{v}) < \frac{2r^2}{\mathrm{rch}(\mathcal{M})}$.

The previous two lemmas lead to a convenient bound on the angle between nearby tangent spaces. We prove here a variation on previous results [NSW08, Prop. 6.2] [BG11, Lemma 5.5]:

Lemma 5.3 (Tangent space variation) Let $x, y \in \mathcal{M}$ be such that $d_{\mathbb{R}^N}(x, y) = r \leq \frac{\mathrm{rch}(\mathcal{M})}{4}$, and let α be the angle between $T_x\mathcal{M}$ and $T_y\mathcal{M}$. Then, $\sin\alpha < \frac{6r}{\mathrm{rch}(\mathcal{M})}$.

Proof Let $v \in T_y\mathcal{M}$ with $\|v - y\| = r$. We will bound the angle between v and $T_x\mathcal{M}$. We have

$$\begin{aligned} \sin\alpha &\leq \frac{1}{\|v-y\|}\left(d_{\mathbb{R}^N}(y, T_x\mathcal{M}) + d_{\mathbb{R}^N}(v, T_x\mathcal{M})\right) \\ &\leq \frac{1}{\|v-y\|}\left(d_{\mathbb{R}^N}(y, T_x\mathcal{M}) + d_{\mathbb{R}^N}(v, \hat{v}) + d_{\mathbb{R}^N}(\hat{v}, T_x\mathcal{M})\right), \end{aligned} \tag{8}$$

where $\hat{v} \in \mathcal{M}$ is the closest point to v in $\mathcal{M}$.

By Lemma 5.1, we have $d_{\mathbb{R}^N}(y, T_x\mathcal{M}) \leq \frac{r^2}{2\mathrm{rch}(\mathcal{M})}$, and by Lemma 5.2 we get $d_{\mathbb{R}^N}(v, \hat{v}) \leq \frac{2r^2}{\mathrm{rch}(\mathcal{M})}$. For the third term in Equation (8), we find

$$\begin{aligned} d_{\mathbb{R}^N}(x, \hat{v}) &\leq d_{\mathbb{R}^N}(x, y) + \|v-y\| + d_{\mathbb{R}^N}(v, \hat{v}) \\ &\leq 2r + \frac{2r^2}{\mathrm{rch}(\mathcal{M})} \leq \frac{5r}{2} \\ &< \mathrm{rch}(\mathcal{M}), \end{aligned}$$

and so we may apply Lemma 5.1 to obtain $d_{\mathbb{R}^N}(\hat{v}, T_x\mathcal{M}) \leq \frac{25r^2}{8\mathrm{rch}(\mathcal{M})}$.

Putting these observations back into Equation (8) we find

$$\sin\alpha \leq \frac{1}{\|v-y\|}\left(\frac{r^2}{2\mathrm{rch}(\mathcal{M})} + \frac{2r^2}{\mathrm{rch}(\mathcal{M})} + \frac{25r^2}{8\mathrm{rch}(\mathcal{M})}\right) = \frac{45r}{8\mathrm{rch}(\mathcal{M})} < \frac{6r}{\mathrm{rch}(\mathcal{M})}.$$

□

Niyogi et al [NSW08, Prop 6.3] demonstrate a bound on the geodesic distance between nearby points, with respect to the ambient distance. We will use a modified statement of this result:

Lemma 5.4 (Geodesic distance bound) Let $x, y \in \mathcal{M}$ be such that $d_{\mathbb{R}^N}(x, y) \leq \frac{\mathrm{rch}(\mathcal{M})}{2}$. Then

$$d_{\mathcal{M}}(x, y) \leq d_{\mathbb{R}^N}(x, y)\left(1 + \frac{2d_{\mathbb{R}^N}(x, y)}{\mathrm{rch}(\mathcal{M})}\right).$$

Proof The announced result states

$$d_{\mathcal{M}}(x, y) \leq \mathrm{rch}(\mathcal{M})\left(1 - \sqrt{1 - \frac{2d_{\mathbb{R}^N}(x, y)}{\mathrm{rch}(\mathcal{M})}}\right).$$

under the same hypothesis on x and y. Rearranging, we have

$$d_{\mathcal{M}}(x, y) \leq \frac{2d_{\mathbb{R}^N}(x, y)}{1 + \sqrt{1 - \frac{2d_{\mathbb{R}^N}(x,y)}{\mathrm{rch}(\mathcal{M})}}} \leq \frac{d_{\mathbb{R}^N}(x, y)}{1 - \frac{d_{\mathbb{R}^N}(x,y)}{\mathrm{rch}(\mathcal{M})}} \leq d_{\mathbb{R}^N}(x, y)\left(1 + \frac{2d_{\mathbb{R}^N}(x, y)}{\mathrm{rch}(\mathcal{M})}\right),$$

where the second inequality is obtained by squaring away the radical. □

5.3 Sampling criteria for manifolds

A local parameterization at $p \in \mathcal{M}$ will be constructed with the aid of the orthogonal projection

$$\pi_p : \mathbb{R}^N \to T_p\mathcal{M},$$

restricted to $\mathcal{M}$. Niyogi et al. [NSW08, Lemma 5.4] demonstrated that, when $r < \frac{\mathrm{rch}(\mathcal{M})}{2}$, π_p is a diffeomorphism from $W = B_{\mathbb{R}^N|_\mathcal{M}}(p; r)$ onto its image $U \subset T_p\mathcal{M}$. We will identify $T_p\mathcal{M}$ with $\mathbb{R}^m$, and define

$$\psi_p = \pi_p|_W^{-1} : U \xrightarrow{\cong} W. \tag{9}$$

Using ψ_p to pull back the metrics $d_\mathcal{M}$ and $d_{\mathbb{R}^N|_\mathcal{M}}$ to $\mathbb{R}^m$, we can view them as perturbations of $d_{\mathbb{R}^m}$. The magnitude of the perturbation is governed by the radius of the ball used to define W. In the following lemma the requirement $r \leq \frac{\mathrm{rch}(\mathcal{M})}{100}$ is simply a convenient bound that yields a small integer constant in the perturbation bound, and does not constrain subsequent results. The bound could be relaxed to $r \leq \frac{\mathrm{rch}(\mathcal{M})}{4}$ at the expense of a weaker bound on the perturbation.

Lemma 5.5 Using the local parameterization (U, ψ_p), with $\mathbb{R}^m \supset U \overset{\psi_p}{\cong} W \subset \mathcal{M}$, if $W \subseteq B_{\mathbb{R}^N|_\mathcal{M}}(p; r)$, with $r \leq \frac{\mathrm{rch}(\mathcal{M})}{100}$, then for all $x, y \in U$,

$$|d_\mathcal{M}(x, y) - d_{\mathbb{R}^m}(x, y)| \leq \frac{23r^2}{\mathrm{rch}(\mathcal{M})}.$$

Proof Let $u, v \in W \subset B_{\mathbb{R}^N|_\mathcal{M}}(p; r)$, and let θ be the angle between the line segments $[u, v]$ and $[\pi_p(u), \pi_p(v)]$, θ_1 the angle between $[u, v]$ and $T_u\mathcal{M}$, and θ_2 the angle between $T_p\mathcal{M}$ and $T_u\mathcal{M}$. Thus $\theta \leq \theta_1 + \theta_2$, and $d_{\mathbb{R}^m}(\pi_p(u), \pi_p(v)) = d_{\mathbb{R}^N}(u, v) \cos\theta$. Defining $\eta = \frac{r}{\mathrm{rch}(\mathcal{M})}$, Lemma 5.4 yields

$$d_\mathcal{M}(u, v) \leq d_{\mathbb{R}^N}(u, v)\,(1 + 4\eta)\,, \tag{10}$$

and so

$$d_{\mathbb{R}^m}(\pi_p(u), \pi_p(v)) \geq \frac{d_\mathcal{M}(u, v)\,\cos\theta}{1 + 4\eta}.$$

Using Lemma 5.1, we find $\sin\theta_1 \leq \eta$, and Lemma 5.3, yields $\sin\theta_2 \leq 6\eta$. Therefore, since $\sin\theta \leq \sin\theta_1 + \sin\theta_2$, we have $\cos\theta = (1 - \sin^2\theta)^{1/2} \geq 1 - \sin\theta \geq 1 - 7\eta$ and we get

$$\begin{aligned} d_{\mathbb{R}^m}(\pi_p(u), \pi_p(v)) &\geq d_\mathcal{M}(u, v)\left(\frac{1 - 7\eta}{1 + 4\eta}\right) \\ &\geq d_\mathcal{M}(u, v)\,(1 - 7\eta)\,(1 - 4\eta) \\ &\geq d_\mathcal{M}(u, v)(1 - 11\eta). \end{aligned}$$

Using Equation (10) we find $d_\mathcal{M}(u, v) \leq \frac{208r}{100}$, so $d_{\mathbb{R}^m}(\pi_p(u), \pi_p(v)) \geq d_\mathcal{M}(u, v) - 23\frac{r^2}{\mathrm{rch}(\mathcal{M})}$, and the result follows since $d_\mathcal{M}(u, v) \geq d_{\mathbb{R}^m}(\pi_p(u), \pi_p(v))$. $\square$

We need to define the local parameterization on a domain that is large enough to contain an ϵ-sampled ball around p that will ensure that p is a safe interior point. Thus we translate Lemma 5.5 into a sampling density requirement:

Lemma 5.6 If $1 < a \leq 10^4$ and $\epsilon \leq \frac{\mathrm{rch}(\mathcal{M})}{100a}$, and $U = B_{\mathbb{R}^m}(p; (a-1)\epsilon)$, then $\psi_p(U) = W \subseteq B_{\mathbb{R}^N|_\mathcal{M}}(p; a\epsilon)$, and

$$\left|d_{\mathbb{R}^N|_\mathcal{M}}(x, y) - d_{\mathbb{R}^m}(x, y)\right| \leq |d_\mathcal{M}(x, y) - d_{\mathbb{R}^m}(x, y)| \leq \frac{23a^2\epsilon^2}{\mathrm{rch}(\mathcal{M})}.$$

Proof Since $a\epsilon \leq \frac{\mathrm{rch}(\mathcal{M})}{100}$, the perturbation bound follows from Lemma 5.5, and the fact that

$$d_{\mathbb{R}^m}(x,y) \leq d_{\mathbb{R}^N|_{\mathcal{M}}}(x,y) \leq d_{\mathcal{M}}(x,y),$$

for any $x, y \in U = \pi_p(W)$. It remains to show that $B_{\mathbb{R}^m}(p;(a-1)\epsilon) \subset \pi_p(B_{\mathbb{R}^N|_{\mathcal{M}}}(p;a\epsilon))$. Using Lemma 5.1, we have that $B_{\mathbb{R}^m}(p;r) \subseteq \pi_p(B_{\mathbb{R}^N|_{\mathcal{M}}}(p;a\epsilon))$ if

$$\begin{aligned} r^2 \leq a^2\epsilon^2 - \left(\frac{a^2\epsilon^2}{2\mathrm{rch}(\mathcal{M})}\right)^2 &= a^2\epsilon^2\left(1 - \left(\frac{a\epsilon}{2\mathrm{rch}(\mathcal{M})}\right)^2\right) \\ &\leq a^2\epsilon^2\left(1 - \left(\frac{1}{200}\right)^2\right). \end{aligned}$$

Thus we require $r \leq \sqrt{\frac{200^2-1}{200^2}}a\epsilon$, which is satisfied by $r = (a-1)\epsilon$ if $a \leq 79999$. □

With Lemma 5.6 we can establish the sampling requirements that ensure that we locally meet the criteria needed to apply Corollary 4.18. We demand that P be δ-generic for P_J (Definition 3.8) with $\mathsf{P}_J = \{p\}$.

Theorem 5.7 (Equating structures) Suppose $\mathcal{P} \subset \mathcal{M}$ is an ϵ-sample set with respect to the intrinsic metric $d_{\mathcal{M}}$, and that for each $p \in \mathcal{P}$ the set $\mathsf{P} = W \cap \mathcal{P}$ is δ-generic for $\{p\}$ with respect to the local Euclidean metric on $W = B_{\mathbb{R}^N|_{\mathcal{M}}}(p;7\epsilon)$. If $\delta = \nu_0\epsilon$ and

$$\epsilon \leq \frac{\nu_0^4 \mathrm{rch}(\mathcal{M})}{10^5}, \tag{11}$$

then $\mathrm{Del}_{\mathbb{R}^N|_{\mathcal{M}}}(\mathcal{P}) = \mathrm{Del}_{\mathcal{M}}(\mathcal{P})$ and they are manifold complexes.

Proof By Lemma 5.6, $B_{\mathbb{R}^m}(p;6\epsilon) \subseteq W$, and P is an ϵ sample set for $B_{\mathbb{R}^m}(p;5\epsilon)$ with respect to $d_{\mathbb{R}^m}$, since $d_{\mathbb{R}^m} \leq d_{\mathcal{M}}$ on $W \times W$. Thus $d_{\mathbb{R}^m}(p, \partial\mathrm{conv}(\mathsf{P})) \geq 4\epsilon$, and so $p \in \mathsf{P}_I$. Corollary 4.18 guarantees that

$$\mathrm{star}(p;\mathrm{Del}(\mathsf{P})) = \mathrm{star}(p;\mathrm{Del}_{\mathbb{R}^N|_{\mathcal{M}}}(\mathsf{P})) = \mathrm{star}(p;\mathrm{Del}_{\mathcal{M}}(\mathsf{P}))$$

if $\rho \leq \frac{\nu_0^4\epsilon}{84}$, and, by Lemma 5.6, this will be true if $\frac{23a^2\epsilon}{\mathrm{rch}(\mathcal{M})} \leq \frac{\nu_0^4}{84}$ with $a = 7$. We obtain the required bound on ϵ. Thus the star of every vertex in $\mathrm{Del}_{\mathcal{M}}(\mathcal{P})$ is equal to the star of that point in the local Euclidean metric, and likewise for $\mathrm{Del}_{\mathbb{R}^N|_{\mathcal{M}}}(\mathcal{P})$. The claim follows since $\sigma \in \mathrm{Del}_{\mathcal{M}}(\mathcal{P})$ if and only if it is in the local Euclidean Delaunay triangulation of every one of its vertices, and likewise for the simplices in $\mathrm{Del}_{\mathbb{R}^N|_{\mathcal{M}}}(\mathcal{P})$. □

We remark that Theorem 5.7 is equally valid if the sampling radius ϵ is used with respect to $d_{\mathbb{R}^N|_{\mathcal{M}}}$ instead of $d_{\mathcal{M}}$.

Although we can express the sampling density requirement in terms of the intrinsic metric of the manifold, the δ-genericity is expressed in terms of local Euclidean metrics. It would be desirable to remove this awkwardness by expressing the genericity requirements in terms of the intrinsic metric. However, our results in Section 4 depend on the m-simplices of $\mathrm{Del}(\mathsf{P})$ being δ-protected, and exploit the fact that $\mathrm{Del}(\mathsf{P})$ is then necessarily a triangulation. It is not clear that similar assertions can be made if protection is assumed only on the m-simplices of $\mathrm{Del}_d(\mathsf{P})$. Further work is needed to understand genericity in this context.

6 Conclusions

We have quantified the close relationship between the genericity of a point set, the quality of the simplices in the Delaunay complex, and its stability under perturbation.

We have produced extrinsic sampling conditions which will guarantee that the intrinsic Delaunay complex is a manifold and coincides with the restricted Delaunay complex. In a companion paper [BDG12], we present an algorithm to produce point sets which meet these sampling conditions, and exploit existing structural results [BG11] to demonstrate that the intrinsic Delaunay complex is a triangulation of the manifold.

We have relied on an embedding of $\mathcal{M}$ in $\mathbb{R}^N$. In future work we aim to develop intrinsic sampling conditions. Another possible direction is to develop sampling conditions for manifolds that are not smooth, since our stability results impose no such constraints on the metric.

A An obstruction to intrinsic Delaunay triangulations

When meshing Riemannian manifolds of dimension 3 and higher using Delaunay techniques, sliver simplices pose problems which cannot be escaped simply by increasing the sampling density. In particular, developing an example on a 3-manifold presented by Cheng et al. [CDR05], Boissonnat et al. [BGO09, Lemma 3.1] show that the restricted Delaunay triangulation need not be homeomorphic to the original manifold, even with dense well separated sampling.

In this appendix we develop this example from the perspective of the intrinsic metric of the manifold. It can be argued that this is an easier way to visualize the problem, since we confine our viewpoint to a three dimensional space and perturb the metric, without referring to deformations into a fourth ambient dimension. This viewpoint also provides an explicit counterexample to the results announced by Leibon and Letscher [LL00]: In general the nerve of the intrinsic Voronoi diagram is not homeomorphic to the manifold. The density of the sample points alone cannot guarantee the existence of a Delaunay triangulation.

We first give a qualitative argument, in terms of Delaunay balls, that demonstrates the heart of the problem. We then explicitly show how density assumptions based upon the strong convexity radius cannot escape the problem. We work exclusively on a three dimensional domain, and we are not concerned with "boundary conditions"; we are looking at a coordinate patch on a densely sampled compact 3-manifold.

One core ingredient in Delaunay's triangulation result [Del34] is that any triangle τ is the face of exactly two tetrahedra. This follows from the observation that a triangle has a unique circumcircle, and that any circumscribing sphere for τ must include this circle. The affine hull of τ cuts space into two components, and if $\tau \in \mathrm{Del}(\mathsf{P})$, then it will have an empty circumsphere centred at a point c on the line through the circumcentre and orthogonal to $\mathrm{aff}(\tau)$. The point c is contained on an interval on this line which contains all the empty spheres for τ. The endpoints of the interval are the circumcentres of the two tetrahedra that share τ as a face.

The argument hinges on the assumption that the points are in general position, and

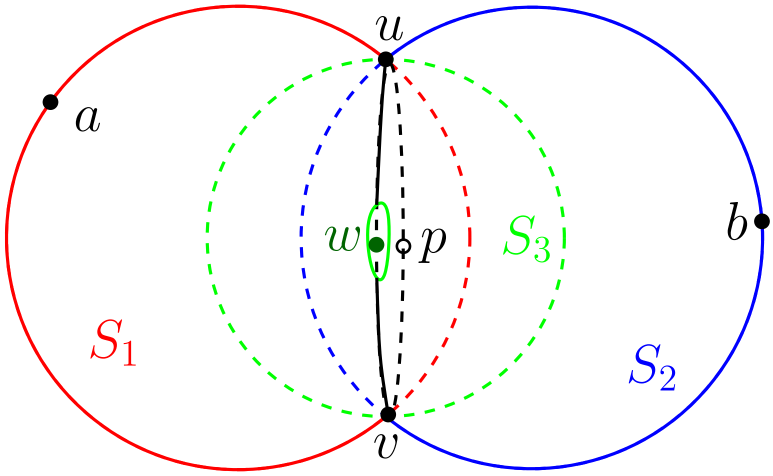

Figure 2: In three dimensions, three closed geodesic balls can all touch three points, u, v, p, on their boundary and yet no one of them is contained in the union of the other two.

the uniqueness of the circumcircle for τ. If there were a fourth vertex lying on that circumcircle, then there would be three tetrahedra that have τ as a face, but this configuration would violate the assumption of general position.

Now if we allow the metric to deviate from the Euclidean one, no matter how slightly, the guarantee of a well defined unique circumcircle for τ is lost. In particular, If three spheres S_1, S_2 and S_3 all circumscribe τ, their pairwise intersections will be different in general. I.e.,

$$S_1 \cap S_3 \neq S_2 \cap S_3.$$

Although these intersections may be topological circles that are "arbitrarily close" assuming the deviation of the metric from the Euclidean one is small enough, "arbitrarily close" is not good enough when the only genericity assumption allows configurations that are arbitrarily bad.

An attempt to illustrate the problem is given in Figure 2, where $\tau = \{u, v, p\}$. Here, the sphere S_3 would be contained inside the spheres S_1 and S_2 if the metric were Euclidean, but any aberration in the metric may leave a part of S_3 exposed to the outside. This means that in principle another sample point w could lie on or just inside S_3, while S_1 and S_2 remain empty. Thus there are three tetrahedra that share τ as a face. We may choose to place w just inside S_3, in which case S_3 is no longer a Delaunay sphere, but two nearby spheres S'_3 and $S"_3$ will take its place. Thus the exact placement of w is not critical: the problem cannot be escaped by an arbitrarily small perturbation.

A.1 Sampling density alone is insufficient

We will now construct a more explicit example to demonstrate that the problem of near-degenerate configurations cannot be escaped with the kind of sampling criteria proposed by Leibon and Letscher [LL00].

Leibon and Letscher [LL00, p. 343] explicitly assume that the points are *generic* which they state as

Definition A.1 The set $\mathcal{P} \subset \mathcal{M}$, is *generic* if $\mathcal{M}$ is an m-manifold and $m+2$ points never lie on the boundary of a round ball.

Here a round ball refers to a geodesic ball. This definition of genericity is natural, and corresponds to Delaunay's original definition [Del34], except Delaunay only imposed the constraint on empty balls. A question that Delaunay addressed explicitly, but which was not addressed by Leibon and Letscher, is whether or not such an assumption is a reasonable one to make. Delaunay showed that any (finite or periodic) point set in Euclidean space can be made generic through an arbitrarily small affine perturbation. That a similar construction of a perturbation can be made for points on a compact Riemannian manifold has not been explicitly demonstrated. However, in light of the construction we now present, it seems that the question is moot when $m > 2$, because an arbitrarily small perturbation from degeneracy will not be sufficient to ensure a triangulation.

Leibon and Letscher proposed adaptive density requirements based upon the *strong convexity radius*. These requirements are somewhat complicated, but they will be satisfied if a simple constant sampling density requirement is satisfied. Exploiting a theorem [Cha06, Thm. IX.6.1], that relates the strong convexity radius to the injectivity radius, $\mathrm{inj}(\mathcal{M})$, and a positive bound on the sectional curvatures, they arrive at the following:

Claim A.2 ([LL00, Lemma 3.3]) Suppose $\mathcal{K}_0$ is a positive upper bound on the sectional curvatures of $\mathcal{M}$, and

$$\eta(\mathcal{M}) = \min\left\{\frac{\mathrm{inj}(\mathcal{M})}{10}, \frac{\pi}{10\sqrt{\mathcal{K}_0}}\right\}. \tag{12}$$

If $\mathcal{P}$ is an $\eta(\mathcal{M})$-sample set for $\mathcal{M}$ with respect to $d_{\mathcal{M}}$, then $|\mathrm{Del}_{\mathcal{M}}(\mathcal{P})| \cong \mathcal{M}$.

In fact, we will show that no sampling conditions based on density alone will be sufficient to guarantee a homeomorphic Delaunay complex in general, even when a sparsity assumption is also demanded. An $\tilde{\epsilon}$*-net* is an $\tilde{\epsilon}$-sparse, $\tilde{\epsilon}$-sample set. We will show:

Theorem A.3 With $\eta(\mathcal{M})$ as defined in Equation (12), for any $\epsilon > 0$, there exists a compact Riemannian manifold $\mathcal{M}$, and and a finite set $\mathcal{P} \subset \mathcal{M}$, such that $\mathcal{P}$ is an $(\epsilon\eta(\mathcal{M}))$-net for $\mathcal{M}$, with respect to the metric $d_{\mathcal{M}}$, but $\mathrm{Del}_{\mathcal{M}}(\mathcal{P})$ is not homeomorphic to $\mathcal{M}$.

A.1.1 A counter-example

We will construct the counter-example by considering a perturbation of a Euclidean metric. This is a local operation, and the global properties of the manifold are only relevant in so far as they affect $\eta(\mathcal{M})$ of Equation (12). We may assume, for example, that the manifold is a 3-dimensional torus $\mathcal{M} \cong \mathbb{S}^1 \times \mathbb{S}^1 \times \mathbb{S}^1$, initially with a flat metric.

Thus assume there is some ϵ_0 such that any compact Riemannian manifold may be triangulated by the intrinsic Delaunay complex when $\mathcal{P}$ is an $\epsilon_0\eta(\mathcal{M})$-net. For convenience, we choose a system of units so that $\epsilon_0\eta(\mathcal{M}) = 1$. We will first construct a point configuration and metric perturbation that leads to a problem, and then we will show that the sampling assumptions are indeed met.

We introduce a number of parameters which we will manipulate to produce the counter-example. We are exploiting the fact that the genericity assumption allows configurations that are arbitrarily close to being degenerate. The assumed ϵ_0 has been fixed.

We will work within a coordinate chart on $\mathcal{M}$, where the metric is Euclidean. We will perturb this metric by constructing a metric tensor $\tilde{g}$, and we will denote by $\tilde{\mathcal{M}}$ the manifold with with this new metric.

Consider points u, v, w, p in the xz-plane arranged with u and v at $\pm a$ on the z axis, and w and p at $\pm(a+\xi)$ on the x axis, with $a = \frac{3}{4}$, and $0 < \xi < r_0\gamma$, where r_0 and γ will be specified below. The Voronoi diagram of these points in the xz-plane is shown in Figure 3. The main point here is that the Voronoi boundary between $\mathcal{V}(u)$ and $\mathcal{V}(v)$ may be arbitrarily small with respect to the distance between the sites, i.e., ξ will be very very small.

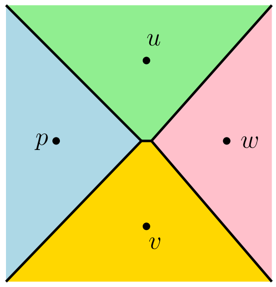

Figure 3: A vertical slice: the xz-plane of the initial Voronoi diagram, seen from the negative y axis.

The three dimensional Voronoi diagram is the extension of this in the horizontal y-direction, so that every cross-section looks the same. Note that since the points are not co-circular, they do not represent a degeneracy by Delaunay's criteria [Del34], but this is irrelevant; we will also argue that the points will not represent a degenerate configuration with respect to the new metric.

We now introduce a small localized metric perturbation so as to change the Voronoi diagram near the origin. For example, we can demand that the matrix of the metric tensor in our coordinate system has the form

$$\tilde{g}(p) = \begin{pmatrix} 1 - f(|p|) & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix},$$

where $|p|$ is the parametric distance from p to the origin. The radial function f is non-negative, and it and its first two derivatives are bounded, e.g.,

$$f(r), |f'(r)|, |f''(r)| \leq \rho.$$

We also demand that there exists a positive $\gamma \leq \rho$ such that $f(r) \geq \gamma$ when $r \leq r_0$, and that $f(r) = 0$ if $r \geq 2r_0$. The parameter r_0, defines the radius of the ball bounding the perturbed region. Now we have $d(w,p) < d(u,v)$ when $\xi < r_0\gamma$.

Since γ may be arbitrarily small compared to ρ, standard arguments supply a function f meeting these conditions. For example, the C^∞ construction described by Munkres [Mun68, p. 6] may be multiplied by a scalar sufficiently small to meet our needs.

The vertical $y = 0$ cross-section of the perturbed Voronoi diagram will look something like Figure 4: $\mathcal{V}(p)$ and $\mathcal{V}(w)$ now meet in the xz-plane, and $\mathcal{V}(u)$ and $\mathcal{V}(v)$ do not. However, since geodesics which do not intersect the ball $B_{\mathbb{R}^3}(0; 2r_0)$ will remain straight lines in the parameter space, the Voronoi diagram is unchanged outside of a neighbourhood of the origin. Thus looking from above at the slice of the Voronoi diagram in the xy-plane, we will see something like Figure 5(a). Figure 5(b) shows the yz-plane.

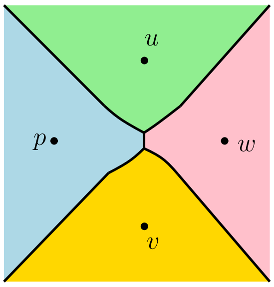

Figure 4: The $y = 0$ slice of the perturbed Voronoi diagram.

Two Voronoi vertices have been introduced, the red and blue points in Figure 5. These are the centres of distinct empty geodesic circumballs for $\{p, u, v, w\}$. Since they cannot lie in the region unaffected by the perturbation, a quick calculation shows that the parametric distance of these Voronoi vertices from the origin is bounded by $4r_0$, when $r_0 \leq \frac{1}{4}$, and it follows from another small calculation that the parametric distance from these Voronoi vertices to any of the four sample points is bounded by $a(1+\frac{3\xi+16r_0^2}{a^2})$. The distances between these Voronoi vertices and the sample points in the new metric will also be subjected to the same bound, since no distances increase. Also, The sparsity condition will not be affected by the perturbation. Thus, since we can make r_0 as small as we please, and ξ is chosen such that $\xi < r_0\gamma$, it follows that the radius of these balls may be made arbitrarily close to $a = \frac{3}{4} = \frac{3}{4}\epsilon_0\eta(\mathcal{M})$. We will argue next that we can make $\left|\eta(\mathcal{M}) - \eta(\tilde{\mathcal{M}})\right|$ as small as desired by reducing the size of ρ. Then other sample points may be placed on the manifold so that the density criteria are met, and no degenerate configuration need be introduced.

This means that the Delaunay complex, defined as the nerve of the Voronoi diagram, will not be a triangulation of the manifold $\tilde{\mathcal{M}}$. As observed by Boissonnat et al. [BGO09], the triangle faces $\{p, w, u\}$ and $\{p, w, v\}$ will be adjacent to only a single tetrahedron, namely $\{p, u, v, w\}$. Thus $\mathrm{Del}_{\mathcal{M}}(\mathcal{P})$ is not a manifold complex as defined in Section 2. This is clearly a problem if the original manifold has no boundary.

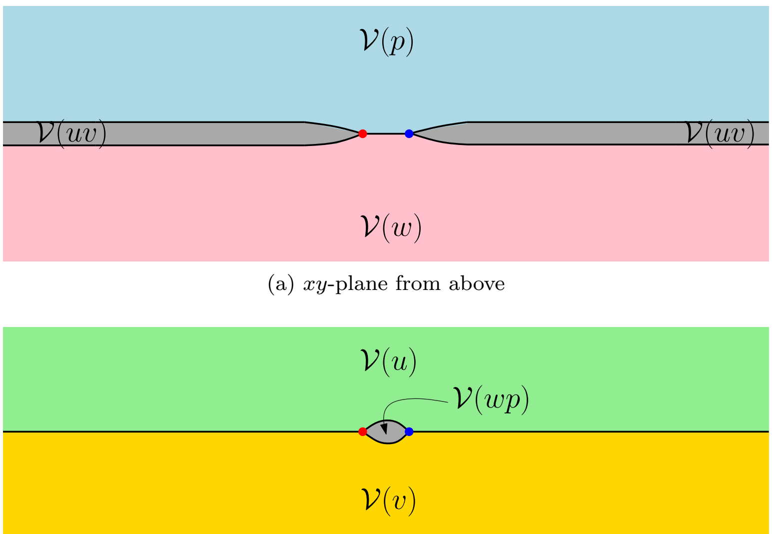

(a) xy-plane from above

(b) yz-plane

Figure 5: Looking at cross-sections; the positive y-direction is to the right. The four points, p, u, v, w, admit two small circumballs with distinct centres (the red and blue points).

A.1.2 The sizing function under perturbation

We need to establish that the metric manipulation that we performed in order to construct the counter-example, does not have a dramatic effect on the sizing function $\eta(\tilde{\mathcal{M}})$. This follows from the fact that we have bounded $g - \tilde{g}$ together with its first and second deriviatives.

Since the sectional curvature may be described as a continuous function of g and it's first and second derivatives [dC92, pp. 56 & 93], the effect of our perturbation on the sectional curvatures can be made arbitrarily small by reducing ρ.

Since we started with a flat metric anyway, the bound $\mathcal{K}_0$ can be made arbitrarily small, and so the second term in Equation (12) will not be the smallest. We need to bound the change in the injectivity radius as well.

This follows from results in the literature [Ehr74, Sak83], which state that for a compact manifold, $\mathrm{inj}(\mathcal{M})$ depends continuously on the metric and its first and second derivatives. Specifically,

Lemma A.4 (Ehrlich) Let $\mathfrak{M}$ be the space of C^3 Riemannian metric structures g on a compact manifold $\mathcal{M}$, and endow $\mathfrak{M}$ with the C^2 topology. The function $g \mapsto \mathrm{inj}_g(\mathcal{M})$ is continuous in this topology.

This means that for any desired bound on $\left|\eta(\mathcal{M}) - \eta(\tilde{\mathcal{M}})\right|$, there will be a ρ that will satisfy the bound.

The construction of the counter-example is complete.

A.2 Discussion

We have shown that for constructing a Delaunay triangulation for an arbitrary Riemannian manifold, a sampling density requirement is not sufficient in general. The solution we propose in the body of this paper, is to constrain the kind of sample sets that we consider. Another approach would be to constrain the kind of metrics that are assumed. However, even with a purely Euclidean metric, allowing configurations to be arbitrarily close to degeneracy means that arbitrarily poorly shaped simplices are to be expected. When the metric is no longer Euclidean, the "shape" of a simplex no longer has an obvious meaning, but the problems associated with point configurations near degeneracy will certainly be present.

Our analysis relied on the ability to make the support of the perturbation small. This is unlikely to be a necessary feature of the construction, but it facilitates our simplistic analysis.

Clarkson [Cla06] remarked that an implication of Leibon and Letscher's claim [LL00] is that for four points close enough together, there is a unique circumsphere with small radius. Our counter-example shows that circumcentres need not be unique under these conditions. In fact the existence of unique circumcentres does not follow from the triangulation result: In our work we do not claim that the m-simplices have a unique circumcentre in the intrinsic metric. However, the argument sketched out by Leibon and Letscher claimed that the intrinsic Voronoi diagram is a cell complex (i.e., it satisfies the *closed ball property* [ES97]), and this does imply unique circumcentres for the top dimensional simplices.

It is worth emphasising that the problems discussed here only arise when the dimension is greater than 2. The same sampling criteria for two dimensional manifolds has been fully validated [Lei99, DZM08], however these works both assume genericity in the sample set, without demonstrating that it is a reasonable assumption.

B Faces inherit protection

In this appendix we will show that if P is δ-generic, then there is a $\tilde{\delta}$ such that the lower dimensional Delaunay simplices are $\tilde{\delta}$-protected.

We defined $\tilde{\delta}$-protection to be a $\tilde{\delta}$-lower bound on the distance from the surface of a Delaunay ball for σ to the nearest point in $\mathsf{P} \setminus \sigma$. However, for computations it is often easier to work with a difference of squared distances. We say that a Delaunay ball $B = B_{\mathbb{R}^m}(c; r)$ for σ is $\tilde{\delta}^2$*-power-protected* if $d_{\mathbb{R}^m}(c, q)^2 - r^2 > \tilde{\delta}^2$ for all $q \in \mathsf{P} \setminus \sigma$. It is often convenient to talk about the protection (or power-protection) of B *from* q for a particular $q \in \mathsf{P} \setminus \sigma$. This means the protection (power-protection) that B would enjoy if q were the only point in $\mathsf{P} \setminus \sigma$.

With the bounds inherent in a δ-generic ϵ-sample set, we can bound the difference between the δ in the definition of δ-protection, and the difference of squares that we often prefer to work with.

Specifically, Let $B = B_{\mathbb{R}^m}(c; r)$ be a Delaunay ball for σ, and let $R = d_{\mathbb{R}^m}(c, \mathsf{P} \setminus \sigma)$. If B is $\tilde{\delta}$-protected, then $R - r > \tilde{\delta}$. Since P is an ϵ-sample set that is δ-sparse, by Lemma 3.9, we have

$$\frac{\delta}{2} \le r \le \epsilon.$$

Also, the same considerations allow us to bound R:

$$\frac{\delta}{2} + \tilde{\delta} \le R \le 3\epsilon.$$

Thus, using that $R - r = \frac{R^2 - r^2}{R + r}$, we have

$$(\delta + \tilde{\delta})(R - r) \le R^2 - r^2, \tag{13}$$

and

$$\frac{1}{4\epsilon}(R^2 - r^2) \le R - r. \tag{14}$$

From Equations (13) and (14) we have that if $\sigma \in \mathrm{Del}(\mathsf{P})$ is δ-protected, then it is $2\delta^2$-power-protected, and if it is $\tilde{\delta}^2$-power protected, then it is $\frac{\tilde{\delta}^2}{4\epsilon}$-protected.

A simplex that is a face of two protected simplices also benifits from some protection:

Lemma B.1 If distinct $(k+1)$-simplices σ^{k+1} and $\tilde{\sigma}^{k+1}$ are $\tilde{\delta}^2$-power-protected and share a k-face, σ^k, then σ^k is $\frac{1}{2}\tilde{\delta}^2$-power-protected.

Proof Let $B_0 = B_{\mathbb{R}^m}(c_0; r_0)$ and $B_1 = B_{\mathbb{R}^m}(c_1; r_1)$ be $\tilde{\delta}^2$-power-protected Delaunay balls for $\sigma^{k+1} = p * \sigma^k$ and $\tilde{\sigma}^{k+1} = q * \sigma^k$, and let $H = \mathrm{aff}(\partial B_0 \cap \partial B_1)$ be the $(m-1)$-flat defined by their intersection. Choose a coordinate system such that c_0 and c_1 lie on the x-axis with the origin where the line generated by $[c_0, c_1]$ intersects H, as shown in Figure 6.

Any circumscribing ball $B = B_{\mathbb{R}^m}(c; r)$ for σ^k, with c in the open segment (c_0, c_1), is a protected Delaunay ball for σ^k, since $B \subset B_0 \cup B_1$. If $w \in \mathsf{P}$ does not belong to σ^{k+1} or $\tilde{\sigma}^{k+1}$, then we must have $d_{\mathbb{R}^m}(w, c)^2 - r^2 > \tilde{\delta}^2$. Indeed, the power

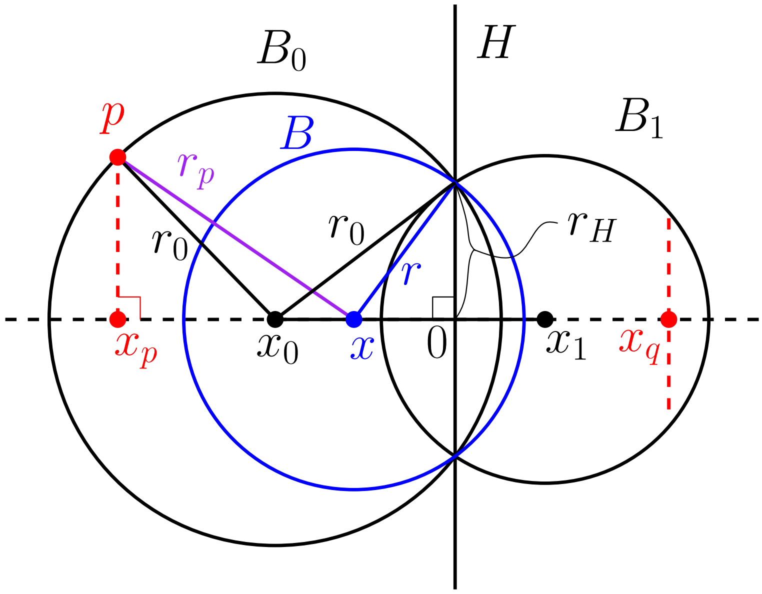

Figure 6: Diagram for Lemma B.1. A planar cross section determined by p, c_0, and c_1. Points on the x-axis are labelled by their x-coordinate. The power-protection of the inner ball (blue) with respect to p is determined by the difference between r_p^2 and r^2; pythagorean expressions for these quantities are extracted from the diagram.

protection of B from w depends linearly on the position of $c \in [c_0, c_1]$, and thus lies in between the power protection of B_0 and B_1 from w. This can be seen from a calculation similar to the one we are about to demonstrate. Thus in order to bound the power-protection of B it suffices to consider only $d_{\mathbb{R}^m}(p,c)^2 - r^2$ and $d_{\mathbb{R}^m}(q,c)^2 - r^2$.

Let x represent the x-coordinate of c, and let x_0 and x_1 be the x-coordinates of c_0 and c_1, and assume that $x_0 < x_1$. Let x_p and x_q be the x-coordinates of p and q. Then $x_p < 0$ and $x_q > 0$. Finally, let r_H be the radius of the circumscribing ball for σ^k whose centre is at the origin in our coordinate system and let $r_p(x) = d_{\mathbb{R}^m}(c,p)$, and $r_q(x) = d_{\mathbb{R}^m}(c,q)$. We wish to bound the quantities $r_p(x)^2 - r(x)^2$ and $r_q(x)^2 - r(x)^2$.

Since $r(x)^2 = x^2 + r_H^2$, and $r_p(x)^2 = (x - x_p)^2 + (r_0^2 - (x_0 - x_p)^2)$, we find

$$\begin{aligned} r_p(x)^2 - r(x)^2 &= ((x - x_p)^2 + (r_0^2 - (x_0 - x_p)^2)) - (x^2 + r_H^2) \\ &= 2x_p(x_0 - x) + (r_0^2 - (x_0^2 + r_H^2)) \\ &= 2x_p(x_0 - x). \end{aligned} \tag{15}$$

Similarly,

$$r_q(x)^2 - r(x)^2 = 2x_q(x_1 - x). \tag{16}$$

Observe that

$$r_p^2(x_1) - r^2(x_1) = r_p^2(x_1) - r_1^2 > \tilde{\delta}^2,$$

and similarly for $r_q(x_0) - r^2(x_0)$. Rather than choosing the optimal point (where $r_p(x) = r_q(x)$), we simply choose $\hat{x} = \frac{1}{2}(x_0 + x_1)$, the mid-point of $[c_0, c_1]$, and find that $r_p^2(\hat{x}) - r^2(\hat{x}) > \frac{1}{2}\tilde{\delta}^2$, and likewise $r_q^2(\hat{x}) - r^2(\hat{x}) > \frac{1}{2}\tilde{\delta}^2$, giving the desired result. □

Lemma B.1 allows us to recursively quantify the protection on simplices of all dimensions in Del(P).

Lemma B.2 If P is δ-generic, then the k-simplices in Del(P) are $\frac{\delta^2}{2^{(m-k)-1}}$-power-protected, and they are $\frac{\delta^2}{2^{(m-k)+1}\epsilon}$-protected. In particular, all $(m-1)$-simplices are $\frac{\delta^2}{4\epsilon}$-protected, and all simplices are at least $\frac{\delta^2}{2^m\epsilon}$-protected.

Proof By Equation (13), the m-simplices are $2\delta^2$-power-protected, and so by inducting with Lemma B.1 we obtain the power protection for the k)-simplices. The claim on the protection then follows from Equation (14). Since the vertices (0-simplices) are δ-protected by Lemma 3.9, the protection of the Delaunay edges (1-simplices), provides a global bound on the protection of all simplices. □

We have found that we did not need to exploit the results presented in this appendix. For our purposes it was sufficient to argue exclusively with the protection given to the top dimensional simplices. The argument of Lemma B.1 may be extended to the case where σ and $\tilde{\sigma}$ are $\tilde{\delta}^2$-power protected and $\sigma^k = \sigma \cap \tilde{\sigma}$, i.e., we need not assume that σ and $\tilde{\sigma}$ are $(k+1)$-simplices. This suggests that we might avoid a recursive argument, and the resulting 2^{-m} constant factor. However, there are k-simplices that cannot be expressed as the intersection of two m-simplcies. This occurs when the Voronoi cell associated to σ^k is itself a simplex. A non-recursive argument for the protection of the lower dimensional simplices would require an analysis of this case.

Acknowledgements

The authors would like to thank David Cohen-Steiner for suggesting the proof technique employed in Lemma 4.4.

References

[AGG+10] P. K. Agarwal, J. Gao, L. Guibas, H. Kaplan, V. Koltun, N. Rubin, and M. Sharir. Kinetic stable Delaunay graphs. In *SoCG*, pages 127–136, 2010. 4

[BDG12] J.-D. Boissonnat, R. Dyer, and A. Ghosh. Delaunay-type structures for manifolds II: Construction. Technical Report CGL-TR-26, Computational Geometric Learning, 2012. http://cgl.uni-jena.de/Publications/WebHome. 27

[BG11] J.-D. Boissonnat and A. Ghosh. Manifold reconstruction using tangential Delaunay complexes. Technical Report N° 7142 v.3, INRIA, 2011. 4, 23, 27

[BGO09] J.-D. Boissonnat, L. J. Guibas, and S. Y. Oudot. Manifold reconstruction in arbitrary dimensions using witness complexes. *Discrete & Comp. Geom.*, 42:37–70, 2009. 27, 31

[BS04] D. Bandyopadhyay and J. Snoeyink. Almost-Delaunay simplices: nearest neighbor relations for imprecise points. In *SODA*, pages 410–419, 2004. 4

[BWY11] J.-D. Boissonnat, C. Wormser, and M. Yvinec. Anisotropic Delaunay mesh generation. Research Report RR-7712, INRIA, 2011. 4

[CDR05] S.-W. Cheng, T. K. Dey, and E. A. Ramos. Manifold reconstruction from point samples. In *SODA*, pages 1018–1027, 2005. 3, 4, 27

[CG12] G. D. Cañas and S. J. Gortler. Duals of orphan-free anisotropic Voronoi diagrams are triangulations. In *SoCG*, 2012. 4

[Cha06] I. Chavel. *Riemannian Geometry, A modern introduction.* Cambridge, 2nd edition, 2006. 29

[Cla06] K. L. Clarkson. Building triangulations using epsilon-nets. In *STOC*. ACM Press, 2006. 32

[Dan00] E. N. Dancer. Degree theory on convex sets and applications to bifurcation. In *Calculus of Variations and Partial Differential Equations*, pages 185–225. Springer-Verlag, 2000. 18

[dC92] M. P. do Carmo. *Riemannian Geometry.* Birkhäuser, 1992. 32

[Del34] B. Delaunay. Sur la sphère vide. *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvennyka Nauk*, 7:793–800, 1934. 3, 11, 20, 27, 29, 30

[dS08] V. de Silva. A weak characterisation of the Delaunay triangulation. *Geometriae Dedicata*, 135:39–64, 2008. 21

[Dug66] J. Dugundji. *Topology.* Allyn and Bacon, Inc., Boston, 1966. 9

[DZM08] R. Dyer, H. Zhang, and T. Möller. Surface sampling and the intrinsic Voronoi diagram. *Computer Graphics Forum (Special Issue of Symp. Geometry Processing)*, 27(5):1393–1402, 2008. 4, 33

[Ehr74] P. Ehrlich. Continuity properties of the injectivity radius function. *Composito Mathematica*, 29:151–178, 1974. 32

[ES97] H. Edelsbrunner and N. R. Shah. Triangulating topological spaces. *Int. J. Comput. Geometry Appl.*, 7(4):365–378, 1997. 3, 4, 32

[Fed59] H. Federer. Curvature measures. *Trans. Amer. Math. Soc.*, 93(3):418–491, 1959. 23

[FKMS05] S. Funke, C. Klein, K. Mehlhorn, and S. Schmitt. Controlled perturbation for Delaunay triangulations. In *SODA*, pages 1047–1056, 2005. 4

[GW04] J. Giesen and U. Wagner. Shape dimension and intrinsic metric from samples of manifolds. *Discrete & Comp. Geom.*, 32:245–267, 2004. 23

[Lei99] G. Leibon. *Random Delaunay triangulations, the Thurston-Andreev theorem, and metric uniformization.* PhD thesis, UCSD, 1999. arXiv:-math/0011016v1. 33

[LL00] G. Leibon and D. Letscher. Delaunay triangulations and Voronoi diagrams for Riemannian manifolds. In *SoCG*, pages 341–349, 2000. 3, 4, 27, 28, 29, 32

[LS03] F. Labelle and J. R. Shewchuk. Anisotropic Voronoi diagrams and guaranteed-quality anisotropic mesh generation. In *SoCG*, pages 191–200, 2003. 3, 4

[Mun68] J. R. Munkres. *Elementary differential topology.* Princton University press, second edition, 1968. 6, 30

[Mun84] J. R. Munkres. *Elements of Algebraic Topology.* Addison-Wesley, 1984. 8

[NSW08] P. Niyogi, S. Smale, and S. Weinberger. Finding the homology of submanifolds with high confidence from random samples. *Discrete & Comp. Geom.*, 39(1-3):419–441, 2008. 23, 24, 25

[Sak83] T. Sakai. On continuity of injectivity radius function. *Mathematical Journal of Okayama University*, 25(1):91–97, 1983. 32

[Whi57] H. Whitney. *Geometric Integration Theory.* Princeton University Press, 1957. 6, 7